MARS ROVER REFERENCE · 2027

URC 2027

Engineering Design Analysis

A first-build reference for a university Mars rover team: the rules, the teams, the subsystems, the numbers and the plan.



Robotics Club @ Mississippi State University

October 2026 · External reference edition

FRONT MATTER

Executive summary

The 2027 University Rover Challenge asks a team to build, from nothing, a rover that finds signs of life, delivers loads across rough terrain, services equipment built for a suited astronaut, and navigates while assisting a human in the field. Four missions plus a written review, 500 points, 36 to 38 teams in the field at the Mars Desert Research Station in June 2027. This page is the argument in one reading: ten findings, ten decisions, eight risks and the next 90 days.

500

POINTS IN THE COMPETITION

Four missions at 100 plus the SAR

2027

S-URC27-RULES

100

SAR POINTS, AND THE FIELD GATE

The top 36 teams are invited

2027

S-URC27-RULES

2.8 %

SHARE OF SCORE VARIANCE FROM THE REVIEW

The other 97 % is the four missions

2022-2026

30.6 %

SHARE FROM AUTONOMY ALONE

The largest single mission share

2022-2026

Key takeaways

- The review decides who competes and the missions decide who wins. The System Acceptance Review is 100 of the 500 points and explains only **2.8 %** of the variation in final totals, while the four field missions carry about **97 %**.
- Autonomy is **20 %** of the score and **95 of its 100 points** need genuine autonomy, so a teleoperated rover caps itself at about **405** of 500 before any other loss.
- The mass and size limits did not move while the work grew: **50 kg** per deployed configuration, **70 kg** fielded, **1.2 m** on every axis.
- Science scores reliably and Autonomy does not. Across 2022 to 2026 Science never produced a zero, and Autonomy averaged a **41 %** zero rate over 162 finals team-years.

The task

The competition is run by The Mars Society at the Mars Desert Research Station, near Hanksville, Utah, from 2 to 5 June 2027 S-URC27-RULES S-URC-LOGISTICS §1.a.i. Registration closes on 30 October 2026, the System Acceptance Review is due 5 March 2027, and the four missions are graded at the field event: Astrobiology, Delivery, Equipment Servicing and Autonomy, at 100 points each, plus the review at 100 S-URC27-RULES §3.a.

About 116 teams register in a typical year. The review admits the top 36, in practice 37 or 38 when scores cluster, and the field is where the ranking is decided S-URC26-SCORES S-CENSUS . A first-time team is therefore solving two different problems in one season: qualification through a written and filmed review, then completion of four missions under a clock.

Ten findings

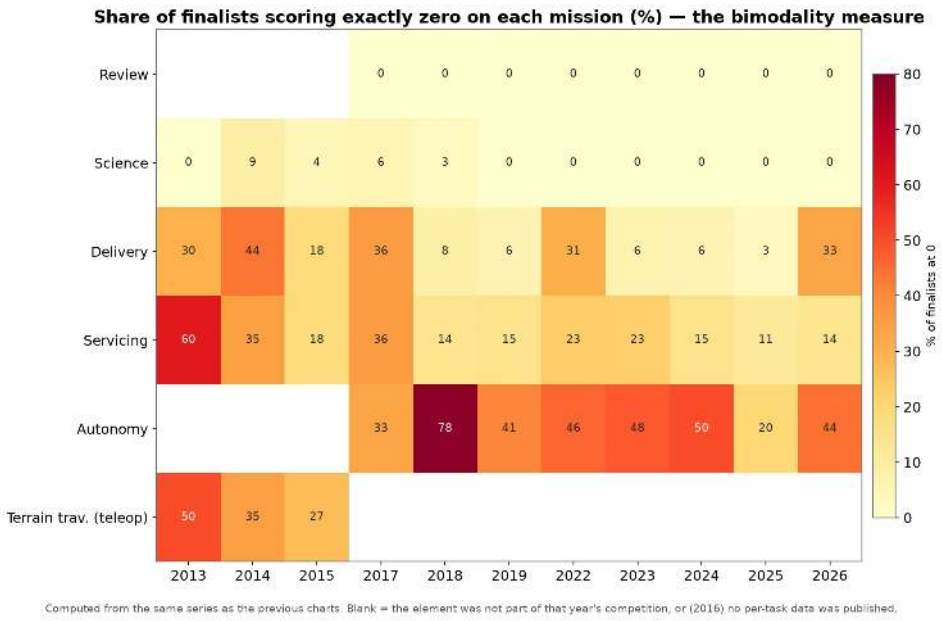
1. 2027 is the largest rule change in five years, and two changes carry the most risk for a first build. The Science Mission became the Astrobiology Mission and now requires two life-detection capabilities, at least one a bioassay that detects a biomolecule such as a protein, a sugar, a nucleic acid, ATP or chlorophyll; a chemical assay alone is not enough S-URC27-RULES S-URC27-QA §1.b.v. The Autonomous Navigation Mission became the Autonomy Mission, whose first half requires the rover to assist a walking astronaut in a suit: follow, stay, fetch a 0.9 kg rock hammer, come and give S-URC27-RULES §1.e. URC ran an astronaut task from 2011 to 2016 and required sign reading from 2017, so the navigation half has precedent; the commanded interaction does not.

2. The limits did not move, so the new capabilities come out of the same 50 kg. The rover must be at or below 50 kg in every deployed configuration and 70 kg across all fielded parts, and must fit inside 1.2 m on every axis at each weigh-in. Overweight costs 5 % of the mission's points per kilogram above 50; a rover that does not fit the box loses 40 % S-URC27-RULES §2.a.ii §2.a.iii. A team that adds a bioassay payload without removing mass elsewhere either fails weigh-in or arrives heavier, slower and hotter.

3. Autonomy is 20 % of the score, and almost all of it needs real autonomy. A teleoperated rover earns about 5 of the 100 Autonomy points, which caps the total near 405 of 500 before other losses S-URC27-RULES §1.e.iv §1.e.x. The follow command is graded by interaction modality: a device command earns 5 points, a visual sign 5, a spoken phrase 10 and a beckoning arm, leg or torso gesture the full 15 S-URC27-RULES §1.e.vi. That is a deliberate, testable choice worth 10 points over the easiest option.

4. Communications is the most common reason teams fail. This is the judges' own headline advice rather than a deduction: they write that "communication failures are far and away the most common problem that stops a team from being a top contender" S-URC-ADVICE . The 2027 environment makes it harder: a 1 km range, no guaranteed line of sight for more than half the Delivery course, a second autonomy target that is deliberately out of sight, a metal command trailer, a 3 m antenna limit, and 900 MHz operation confined to three 8 MHz sub-bands that must be selectable per mission S-URC27-RULES S-URC27-RULES §2.c §3.d.ii.

5. Science is the reliable points; Delivery and Autonomy are where the field breaks. Across 2022 to 2026 Science never produced a zero and averaged between 72 and 81, with the narrowest spread. Autonomy produced 7 to 17 zeros a year and Delivery between 1 and 12, with Delivery's worst season 2026 at 12 zeros S-URC22-SCORES S-URC26-SCORES . A rover that simply completes those two missions gains 20 to 30 places.



The zeros are not spread evenly across the record from 2013 to 2026. Science contributes almost none, while Autonomy and Delivery produce most of the failed runs, which is why completion beats sophistication in the first season.

Robotics Club @ Mississippi State University, 2026. Original analysis.

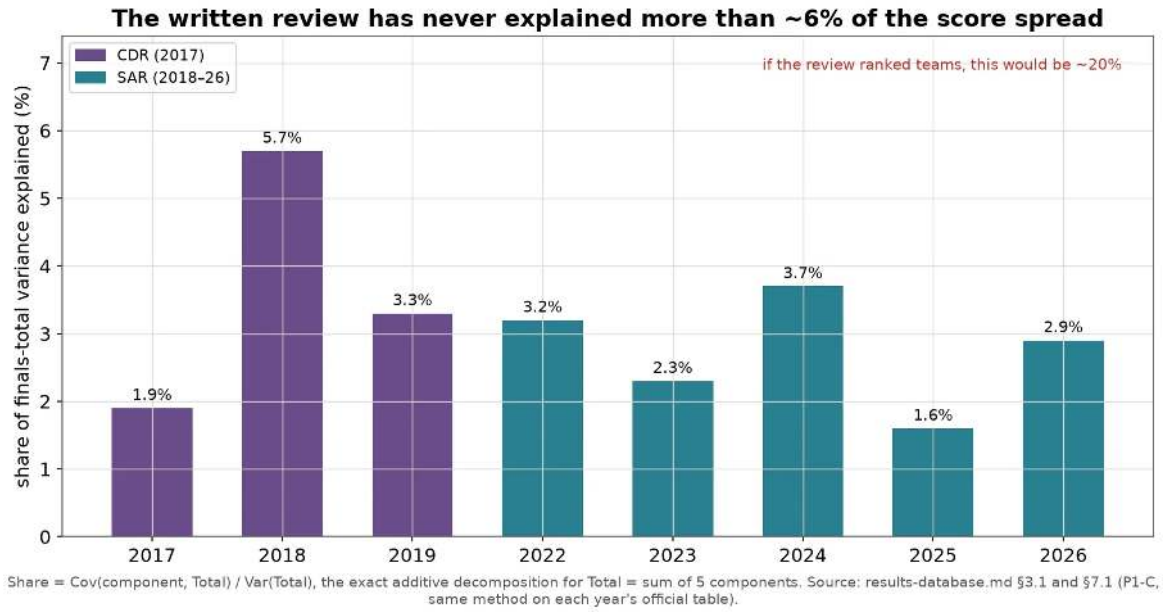
S-URC13-SCORES S-URC19-SCORES Full resolution

6. Mass and power are dominated by peaks, not averages. A 50 kg rover draws about 249 W average on a URC mission and about 3.7 kW at peak, roughly 15 times as much, and the peak sizes the fuse, the contactor and the wiring. A pack sized only for energy sags about 15 % on a hill with the arm loaded, and a battery management system set to trip at a safe current will shut the rover down during a legal arm hold. The binding mission is the 60-minute Delivery clock, so the pack is sized at 400 to 500 Wh rather than at the average.

7. The E-stop rule tightened, and older team designs are now a trap. Since 2026 a single red push button must stop movement and cease all power draw from every battery, switching off all systems S-URC27-RULES §2.a.v. Designs published in 2024 and 2025 that only kill the drive bus are non-compliant, and copying one is a plausible way to fail scrutineering or to start a fire.

8. The 2027 task hardware is named by hyperlink, and its physical properties set real requirements. The typing keyboard is a Redragon K552 with a 19.05 mm key pitch, roughly 2 mm actuation and about 0.45 N of force; the Fetch object is a 0.91 kg all-steel rock pick of 300 mm; the hose coupling is a 1.5 in cam-lock whose gasket changed from removed in 2026 to cut to half thickness in 2027, which makes it stiffer to push onto S-URC27-RULES S-URC27-RULES §1.d.ii §1.e.viii. Key and lock markers are 1 × 1 cm and keyboard markers 2 × 2 cm, which sets camera resolution and arm repeatability.

9. The review is a form, not an essay. In 2026 the System Acceptance Review split into 48 points for a five-minute video, 40 for a seven-page report and 12 for the judges' assessment, and material beyond the limits was not read S-URC26-SAR . The lowest admitted score was 85.10 in 2023, 84.04 in 2024, 86.90 in 2025 and 75.86 in 2026, so the 2026 drop tracks a rebuilt rubric rather than a weaker field S-URC23-SCORES S-URC24-SCORES S-URC25-SCORES S-URC26-SCORES . Cross-year review benchmarks are not comparable.

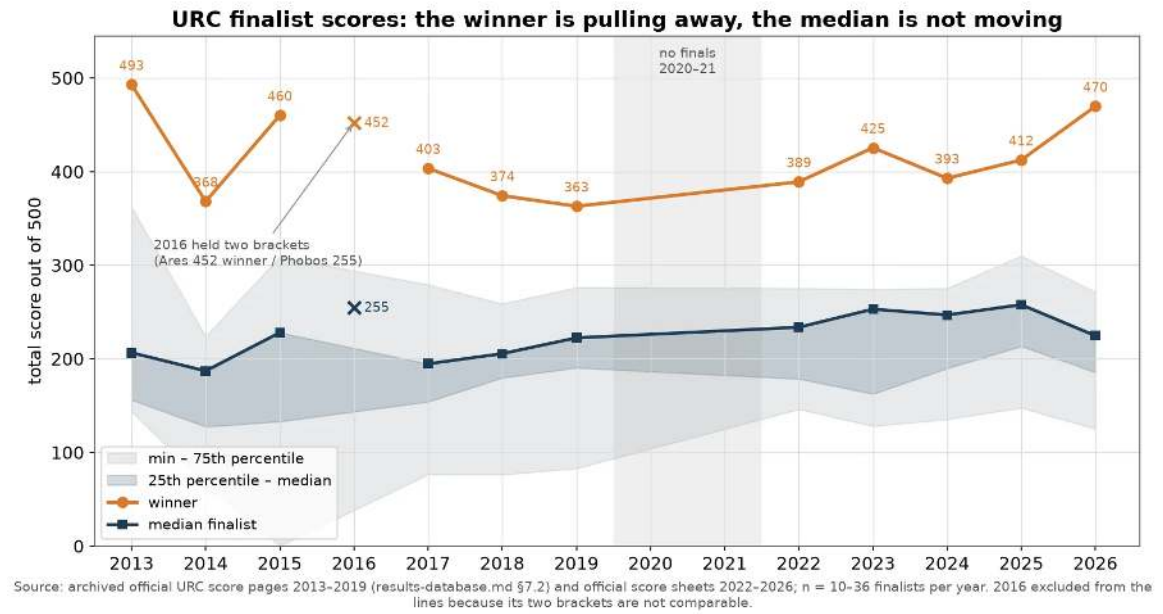


The share of finals-total variance explained by the written review element, by year. The 2.8 % quoted above is the 2022 to 2026 value: the review is worth 20 % of the points and explains almost none of the spread in final totals. Its real leverage is the field slot.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

10. A first-year team's realistic target is the middle of the field, and that is a good result. The median team that recorded mission scores finished between 224 and 258 points of 500 in each of the last five seasons, and 54 teams that appear on two or more official team lists never reached a finals score sheet at all S-URC22-SCORES S-URC26-SCORES S-CENSUS . Qualifying, finishing all four missions at 50 to 60 % of the available points, and returning for a second season is the strategy the historical data supports.



The chart covers 2013 to 2026, and in the five seasons to 2026 the winner is separated from the median finalist by 146 to 245 points. The gap is not built in one subsystem, which is why the plan sequences measurement, integration and review evidence rather than component purchases.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

The ten decisions to take first

These are ordered by how much downstream design each one unlocks. The mission backfit turns each into a trade study, and the problem register turns it into tasks.

#	DECISION	WHY IT COMES FIRST	WHERE IT IS WORKED
D1	Mass and configuration strategy: one rover, or a common chassis with swappable arm and science modules	Sets the 50 kg and 70 kg budgets, the packing problem and the 10-minute Autonomy to Equipment transition	Systems engineering
D2	Mobility architecture: six-wheel rocker-bogie, four-wheel independent, or articulated	Drives wheel size, torque, mass, packaging and terrain capability	Mobility
D3	Arm architecture: degrees of freedom, reach, stow and end-effector strategy	The hardest integration constraint, 1.5 m of reach inside a 1.2 m cube, and the most task-specific	Arm and end-effectors
D4	Science payload concept: the two life-detection methods and the bioassay	A new 2027 requirement that locks mass, power, fluids, time and contamination control	Science
D5	Communications architecture: band, radio, antenna and relay strategy	The dominant failure mode, and it determines mast, power and operator workflow	Communications
D6	Autonomy ambition: teleoperation only, GNSS waypoints, or perception with person following	Determines compute, sensors, software effort and the 20 % of the score that is autonomy	Software and autonomy
D7	Drone: yes or no	Optional and useful in Delivery and Autonomy, never required, and it adds a second vehicle and a pilot	Delivery and drones
D8	E-stop and electrical safety architecture	Pass or fail at scrutineering, and late changes are expensive	Power and safety
D9	Operator station and staffing model	Six seats, no view of the rover, missions of 30 to 40 minutes: the human system is part of the design	Operations
D10	Review strategy: what evidence the video and report will show, and when it must be true	The gate to the field, and it dates the build backwards from March 2027	Review, judging and awards

The eight main risks, in order

1. **Communications failure at the field.** The judges' most-cited cause of a motionless rover.
2. **Mass and packing overrun.** The 50 kg and 1.2 m constraints surface late and force redesign.
3. **Autonomy under-delivery.** A fifth of the score needs real autonomy, and the astronaut sub-mission has no precedent to copy.
4. **Science payload infeasibility.** A bioassay means chemicals, fluids, containment, time and a safety plan for both; a payload deferred until after the review will not work in June.
5. **Thermal and dust.** A sealed box in Utah sun, where electronics that pass on a bench fail at the station.
6. **Schedule.** Preliminary Design Review on 2 December 2026, the System Acceptance Review on 5 March 2027, finals on 2 June 2027: the rover must be demonstrable in March.
7. **Safety and compliance.** E-stop, battery fire, spectrum and weigh-in failures cost the whole event rather than points.
8. **Team capacity.** 24 students, six of them in the command station, learning twelve subsystems at once `S-URC-LOGISTICS`.

The next 90 days

The window runs from 26 October 2026 to 23 January 2027, with the Preliminary Design Review inside it on 2 December. The full sequence is on the first 90 days page; the six items below are the ones that change later decisions.

1. **Registration closes on 30 October 2026** `S-URC27-RULES` `§3.a.i`. It also starts the compliance clock: radio licensing, drone registration and Remote ID if a drone is planned, and the communications documentation due on 30 April 2027 `S-URC27-RULES` `S-URC27-RULES` `§2.b` `§2.c.ii`.
2. **The task hardware must be measured.** Judge-like mock-ups of the keyboard, tubular lock, cam-lock hose, cache handle, drawer, latch and rock pick make every unpublished force, torque and clearance measurable. That single campaign closes most

of the unknowns that block the arm.

3. **The two highest-uncertainty subsystems need prototypes first:** a radio link that reaches 1 km without line of sight, and a drive and suspension mock-up that climbs a loose 30° slope with the intended payload.
4. **The mass and configuration strategy has to be settled before detailed design,** because it constrains wheel size, arm class, packaging and the operator plan.
5. **The review evidence starts now.** The video carried 48 of 100 review points in 2026 and has to show a working rover, so the build schedule is set by March rather than June `S-URC26-SAR`.
6. **The Q&A is binding and it moves.** It changed twice within a month of publication, in October 2026, so the current rulings are re-read before each milestone `S-URC27-QA`.

How the book continues

PART	WHAT IT ADDS TO THIS SUMMARY
The competition	The four missions decoded, the hard limits, the 155 requirements and the Q&A rulings
The field	The 68-team census, five seasons of results and 51 dossiers
Subsystems	Twelve studies with comparison tables, failure modes and a shortlist for a first build
Cross-cutting analysis	Mass, power, the 66 interfaces and what separates the top from the rest
Mission backfit	Each mission and task against what the strongest teams did and what 2027 invalidates
The build plan	110 problems in five waves, 70 risks and 77 open questions
Statistics	Every headline number with its sample, period and chart

title: The 2027 competition description: What the 2027 University Rover Challenge asks for: four field missions, a 100-point acceptance review and the hard limits behind them. part: The competition chapter: "2" toc_depth: 3 tags: - urc-2027 - missions - calendar

The 2027 competition

The University Rover Challenge is a student rover competition with a published rulebook, a field in the Utah desert and a single set of hard limits that every design must fit. In 2027 it scores 500 points: four field missions at 100 points each and a 100-point System Acceptance Review (SAR) that also decides which teams reach the field at all. This page sets the shape of the season: the calendar, the five scored components, the limits, and the changes that make 2027 harder than 2026. The three pages that follow take each part apart, from the rulebook clauses to the review that gates entry.

500

TOTAL POINTS

Four missions at 100 plus the SAR at 100

2027

S-URC27-RULES

36

TEAMS INVITED TO THE FIELD

From about 116 registrants in 2026

2022–2026

S-URC27-RULES

50 kg

MASS PER DEPLOYED CONFIGURATION

70 kg fielded in total; 1.2 m cube at weigh-in

2027

S-URC27-RULES

155

TRACED ENGINEERING REQUIREMENTS

79 mandatory, 61 scoring-driven, 15 advisory

2027

S-URC27-RULES

Key takeaways

- The 2027 mission set changed shape for the first time in five years. Science became Astrobiology, with **two** life-detection capabilities and at least one bioassay. Autonomous Navigation became an Autonomy Mission that opens with a suited astronaut.
- The System Acceptance Review is 20 % of the score and the only route to the other 80 %. It is judged against other teams, and only the 36 highest scores are invited to the field.
- Mass and size did not move: 50 kg per deployed configuration, 70 kg fielded in total, a 1.2 m cube at every weigh-in. The science payload grew inside an unchanged budget, which is the central trade of the cycle.
- A teleoperation-only rover can score at most 405 of 500, because 95 of the 100 Autonomy points require genuine autonomous operation.

What the competition is

The University Rover Challenge is run by The Mars Society at the Mars Desert Research Station (MDRS), a simulated Mars habitat on Utah state land near Hanksville. Student teams design and build a rover, then operate it from a command and control (C2) station with no view of the course S-URC27-RULES §3.d.i. Everything the operators know about the site arrives through the rover: cameras, telemetry and a wireless link that is not guaranteed line of sight S-URC27-RULES §2.c.i.

The station sits at 1,309 m (4,295 ft) in high desert, with full daylight, frequent dust and summer heat up to 37.8 °C S-URC27-RULES S-URC-LOGISTICS §3.d.iii. None of it can be rehearsed on site: URC prohibits testing at MDRS before, during or after the event S-URC27-RULES §3.d.iv. Terrain and reliability therefore have to be proven at home, against an obstacle set the rules do not fully specify.



The ground the 2027 Delivery and Autonomy missions use: loose slopes, rock and boulder fields, and a hill that blocks the radio horizon. A rover that cannot climb and cannot hold a link loses the same points twice.

Mars Society, 2017. CC BY-SA 4.0. *Original source* *Full resolution*

The season at a glance

Every date below is fixed in the 2027 rulebook `S-URC27-RULES`. The milestone that decides the season is the System Acceptance Review on 2027-03-05: it is scored and it controls entry.

MILESTONE	DATE	WHAT IT IS	WHAT IT CONTROLS
Registration closes	2026-10-30	Declaring intent and team details	Eligibility (§3.a.i)
Preliminary Design Review (PDR)	2026-12-02	Project and design readiness, five pages	Readiness assessment, not scored (§3.a.ii)
System Acceptance Review	2027-03-05	Written and video package	100 points, top 36 invited (§3.a.iii)
Communications documentation	2027-04-30	Devices, bands and operator licences	Compliance (§2.c.ii)
Science Plan	2027-05-21	Written plan for the Astrobiology Mission	Feeds the Astrobiology score (§1.b.i)
Rover Expense Report	2027-05-26	Expenses in work-breakdown format; update allowed 2027-06-02	10 % of total points per day late (§3.c.ii)
Finals	2027-06-02 to 2027-06-05	Four missions at MDRS	4 × 100 points (§3.a.iv)

The field schedule itself arrives about a week before the first field day, and the order of tasks is drawn per team, so two teams at the same event do not run the missions in the same order `S-URC27-QA` CQ4. Setup from C2 access is capped at 15 minutes and teardown at 10, and weighing and measuring happen inside the setup window `S-URC27-RULES` §1.a.ii §2.a.ii.

The four missions and the review

Each mission is worth 100 points and is scored on its own. The clock and the autonomy requirement differ sharply, which is what makes the four missions different engineering problems rather than one problem four times.

COMPONENT	CLOCK	POINTS	AUTONOMY REQUIRED	DRONE
Astrobiology	20–30 min plus 5 min cache removal	100	No	Not allowed
Delivery	30–60 min, two stages	100	No	Allowed, optional
Equipment Servicing	30 min	100	Autonomous typing only	Not allowed
Autonomy	40 min, two sub-missions	100	95 of the 100 points	Reconnaissance only, and it must land before route finding
System Acceptance Review	Desk submission due 2027-03-05	100	No	No

Mission scores are independent and cannot fall below zero, and penalties are additive fractions of the points earned rather than of the 100 available S-URC27-RULES §1.a.i. The formula, the worked penalty cases and the exact point value of each autonomy behaviour are on Missions and scoring.

The hard limits

These are the constraints that decide whether a rover competes at all. Each one is stated in the rulebook, and most carry a specific penalty S-URC27-RULES .

LIMIT	2027 VALUE	CONSEQUENCE
Mass per deployed configuration	50 kg	5 % of points earned per kilogram over, in each mission (§2.a.iii)
Mass of all fielded parts	70 kg	Strict; parts are removed until the total is under (§2.a.iii)
Stowed size	1.2 × 1.2 × 1.2 m, no disassembly	40 % penalty at weigh-in (§2.a.ii)
Fielded equipment	$26,000 cap	Covers rover, modules, drone, EVA suit, base station and C2 (§3.c.iii)
E-stop	One red push button	Must stop movement and remove all battery power (§2.a.v)
Drone	5 kg including dummy mass	Above it the drone cannot be used (§2.b.vii)
900 MHz links	8 MHz within three sub-bands	Applies to C2, rover, drone and astronaut links (§2.c.v)
Interventions	20 % each	Additive, and the mission clock keeps running (§3.e.viii)
Setup and teardown	15 min and 10 min	Weigh-in and measurement happen inside setup (§1.a.ii; §2.a.ii)

The mass limit is measured, not declared

The rover is weighed and measured before every mission, and it may fold or articulate to fit but may not be disassembled S-URC27-RULES §2.a.ii. Weight is rounded down to the whole kilogram, so 50.9 kg records as 50 kg, but designing to that edge leaves no margin for a bracket or a cable added after the last weigh-in.

What changed for 2027

The comparison below is the working list for anyone converting 2026 knowledge into 2027 requirements. The 2026 column is reference only; the 2027 column is the rule S-URC27-RULES S-URC26-RULES .

CHANGE	2026	2027	ENGINEERING CONSEQUENCE
Science becomes Astrobiology	One life detection plus a second capability of choice	Two life detections, at least one a bioassay	Payload mass, power, fluids and time-to-result all rise
Stratigraphic profile	Required photograph and discussion	Removed	Less camera reach and geology, more life detection
Autonomy	Seven static targets, 30 min	Astronaut assistance plus route finding, 40 min	Human detection and following, gesture recognition, rough-terrain route planning
Object recognition	Mallet, hammer and bottle detection	Deleted	Perception effort shifts from object detection to pose and gesture
Equipment Servicing	USB-C cable task	USB task removed, key-in-lock added	Finer manipulation, rotational compliance
Hose connector	Gasket removed	Gasket cut to half thickness	Push-on force is higher
Drone launch	From the rover or the start area	Designated field landing pad	The rover no longer carries the drone
Drone dummy mass	Waived above 24 km/h	Never waived	The 5 kg limit holds in all conditions
Mission adjacency	Equipment Servicing, then Autonomy	Autonomy, then Equipment Servicing, 10 minutes apart	Two configurations are weighed back to back
Fielded equipment cap	$24,000	$26,000	The EVA suit sits explicitly inside the cap
Rule numbering	Competition in §2, rover in §3	Equipment in §2, competition in §3	Every clause citation moves

Two consequences are structural rather than incremental. The Autonomy Mission now rewards commanded human-robot interaction, and the command set (Follow, Stay, Fetch, Come, Give) with gesture recognition has no precedent in earlier URC rulebooks S-URC27-RULES S-URC27-QA AUQ11. At the same time the science payload doubled while the 50 kg limit and the 1.2 m cube stayed put, so the extra instrument mass has to come out of mobility, structure or modularity.

From clauses to requirements

The rulebook is written as competition clauses, not as specifications. The requirement set on this site translates those clauses into 155 testable statements, each traced to a rule clause or a Q&A ruling, and marks each one mandatory, scoring-driven or advisory S-URC27-RULES S-URC27-QA .

DOMAIN	ROWS	WHAT IT COVERS
SYS	14	Mass, size, integration and configuration
MOB	11	Terrain, climb, stopping accuracy and payload
ARM	21	Reach, grasp, typing, key transfer and hose connection
SCI	20	Sample acquisition, life detection, cache and documentation
PWR	9	Endurance, E-stop behaviour and thermal limits
COM	15	Range, obstruction, spectrum and antenna geometry
CMP	15	Marker detection, astronaut detection and gesture
SFT	19	Localisation, route planning, following and abort behaviour
OPS	12	C2 station, setup time, roles and schedules
SAF	8	E-stop architecture and deployable safety
DRN	11	Drone mass, flight rules and airspace

Mandatory clauses are the gates: mass, size, E-stop, spectrum, communications documentation and the review milestones. Most mission clauses are scoring-driven, where partial credit exists and the strategic question is where to spend effort. The full set, with trace and test method per row, is on 155 engineering requirements.

How to read this part

- Rules, limits and the field works the cross-mission envelope clause by clause, with the mass, size, spectrum, terrain and penalty rules in one place.
- Missions and scoring takes the four missions and the review apart, with the exact autonomy points and the penalty arithmetic.
- Review, judging and awards covers the PDR, the SAR rubric, the science and chemical safety plans, the expense report and the awards.
- The requirements, the Q&A record, the rule history, how the missions evolved and site and field constraints sit alongside those three.

> **How the 2027 material is dated**
>
> The rulebook was published in September 2026 and the Q&A used throughout this site is current to 2026-10-11. The 2027 PDR, SAR and Science Plan pages are not yet published, and each one says its requirements will follow. Where a 2026 value is used for scale, it is labelled as 2026 and never presented as a 2027 rule `S-URC27-QA` `S-URC27-SAR-PAGE` `S-URC27-PDR-PAGE`.

title: Rules, limits and the field description: The 2027 rulebook clause by clause: mass, size, spectrum, deadlines, terrain and penalties, with every hard limit in one place. part: The competition chapter: "2.1" toc_depth: 3 tags: - rules - limits - penalties - field

Rules, limits and the field

The 2027 rulebook is written as competition clauses, and most of them are read during a design meeting rather than on the field. This page keeps the clauses in order and states what each one fixes: the scoring frame, the envelope of all four missions, the mass, size, power and spectrum limits, the milestone calendar, the site and the penalty arithmetic. Each section points to the page that takes the detail further. Where a number the design needs is missing from the rulebook, the gap is named rather than filled.

50 kg

MASS PER DEPLOYED CONFIGURATION

5 % of points per kilogram over, each mission

2027

S-URC27-RULES

1.2 m

CUBE AT WEIGH-IN

40 % penalty if the rover does not fit

2027

S-URC27-RULES

$26,000

FIELDED EQUIPMENT CAP

Rover, drone, EVA suit, base station and C2

2027

S-URC27-RULES

8 MHz

900 MHZ BANDWIDTH LIMIT

Three selectable sub-bands across 902–928 MHz

2027

S-URC27-RULES

Key takeaways

- The 2027 rulebook renumbers the clauses: §1 holds the missions, §2 holds rover equipment, §3 holds the competition. Every citation written against the 2026 numbering has moved.
- Mass is measured before every mission and rounds down, so 50.9 kg records as 50 kg. The design target is not that edge, because the penalty is 5 % of points earned for each whole kilogram over.
- The oversize penalty is 40 %, and both the mass and size penalties apply to points earned rather than to the 100 available. A bad mission makes a penalty cheap; a good mission makes it expensive.
- Several numbers a design depends on are simply absent from the rulebook: slope and step limits, task weights, object dimensions and insertion forces. The last section lists them.

How the 2027 rulebook is organised

The 2027 rulebook has three parts, and the renumbering matters because the 2026 cycle put competition rules in §2 and rover rules in §3 S-URC27-RULES S-URC26-RULES .

PART	CLAUSES	WHAT IT GOVERNS
Part 1: missions	1.a to 1.e	Scoring frame, schedule, and the four mission envelopes
Part 2: rover equipment	2.a to 2.c	Mass, size, power, safety, drones and communications
Part 3: competition	3.a to 3.e	Milestones, the field, the expense report, the environment, penalties

Each clause group below carries an anchor, so a reference such as §2.a.iii anywhere on this site lands on the section that explains it.

Part 1: the missions

Part 1 sets the scoring frame, the calendar and the four mission envelopes. The detail behind each mission is on Missions and scoring.

Clause 1.a: points, penalties and the clock

The total is 500 points: four missions at 100 points each plus the System Acceptance Review at 100 S-URC27-RULES §1.a.i. Mission scores are independent and cannot fall below zero, and penalties are additive fractions of the points earned, not of the maximum:

```
mission score = max(0, points earned × (1 − sum of penalty fractions))
total score   = sum of the four mission scores + the SAR score
```

Two additive penalties of 10 % and 20 % therefore leave 70 % of the points earned, not 72 % S-URC27-RULES §1.a.i. The practical consequence runs opposite to intuition: a penalty costs least on a mission that scores badly, and most on a mission that scores well.

Time is fixed at both ends of every mission. The window before the rover moves and the window after the clock stops are as binding as the mission itself, and every limit below is a rule limit S-URC27-RULES .

WINDOW	LIMIT	NOTE
C2 setup to ready	15 min	Includes communications, weighing and measuring (§1.a.ii; §2.a.ii)
Teardown	10 min	Radios off immediately (§1.a.ii)
Post-mission retrieval	10 min	Rover, drone and deployed items (§1.a.iii)
Autonomy to Equipment Servicing	10 min	Same C2 station, adjacent course (§1.a.iv; §1.d.i)
Astrobiology	20–30 min	Plus 5 min for cache removal (§1.b.ii; §1.b.vii)
Delivery	30–60 min	Fixed stage times; total includes both stages (§1.c.i)
Equipment Servicing	30 min	Includes up to 100 m of travel to the lander (§1.d.i)
Autonomy	40 min	Both sub-missions, in any order (§1.e.i)

A team that competes in Astrobiology or Delivery competes in only one mission that day (§1.a.iv). The field schedule is issued about a week before the first field day and the order of tasks is drawn per team, so the sequence differs between teams at the same event S-URC27-QA CQ4.

Clauses 1.b to 1.e each set a task list and a numeric envelope. What follows is the operating envelope a design has to fit, not the point values.

Clause 1.b: Astrobiology

Every requirement below is a rule requirement, and the clause links carry the exact text S-URC27-RULES .

- The rover investigates at least two sites within 0.5 km of the C2 station in 20–30 minutes (§1.b.ii).
- Analysis happens on the rover. Samples may not be returned to the team for laboratory work (§1.b.iv).
- The cache sample comes from one site, at least 10 cm deep and at least 5 g, and the portion below 10 cm determines the mass (§1.b.vi).
- Two life-detection capabilities are required, at least one of them a bioassay that detects a biomolecule S-URC27-RULES S-URC27-QA §1.b.v AQ9.
- The cache must not spill when laid on its side or inverted, must be removable within 5 minutes after roving, and must be openable by the judges (§1.b.vii).
- Every site investigated needs a panorama with cardinal directions and scale, a close-up with scale, and GNSS coordinates with elevation and accuracy range (§1.b.iii).
- Chemicals, water and reaction products stay on the rover. Hazardous chemicals need a pre-approved safety plan S-URC27-RULES S-URC27-QA §1.b.ix AQ4.
- A debrief follows the cache handover even if the mission failed (§1.b.x). The Science Plan is due 2027-05-21 and counts toward the score (§1.b.i).

Clause 1.c: Delivery

The envelope below is set by four clauses and the cross-mission sight rule S-URC27-RULES .

- The mission has two stages. A minimum score in Stage 1 is required to proceed, and remaining Stage-1 time is added to Stage 2 (§1.c.i).
- Total time on course is 30–60 minutes. The 2027 stage times and the Stage-1 minimum are not published (§1.c.i).
- Objects are up to 5 kg, smaller than 40 × 40 × 40 cm, with a grasp feature no larger than 7.5 cm (§1.c.iii).
- The course extends up to 1 km from the start gate, over soft sand, gravel, stony ground, boulder fields, vertical drops and steep loose slopes (§1.c.ii).
- Line of sight is not guaranteed for more than 50 % of the course, and later tasks are deliberately beyond it (§3.d.ii).
- Working GNSS is essential, but coordinates for pickup and delivery are only approximate (§1.c.iv).
- A drone may assist. No other mobile deployable, including a mini-rover, may be used (§1.c.i).

Clause 1.d: Equipment Servicing

The rover travels up to 100 m across relatively flat ground and works on a mock lander, with interactive equipment between ground level and 1.5 m S-URC27-RULES §1.d.i. Nine task families are listed, and they may be performed in any order (§1.d.ii).

TASK FAMILY	WHAT THE ROVER MUST DO
Sample tube	Pick up a test-tube-sized tube, insert it into a cache, close the lid, turn the lock
Cache transport	Carry a cache with a handle at least 10 cm long and no more than 5 cm in diameter, under 5 kg
Drawer	Open a lander drawer, insert the cache into a tight space, close the drawer
Latch and panel	Undo a latch on a hinged panel and open it
Autonomous typing	Type a 3–6 letter launch key on a vertically mounted keyboard with no hands on it
Key transfer	Remove a key from a panel-mounted tubular pin tumbler lock, insert it into another lock and turn it
Hose	Push a 1.5 inch cam-lock fitting on, with the sealing gasket cut to half thickness
Valve	Turn a quarter-turn lever handle
Controls	Push buttons, flip switches, turn knobs

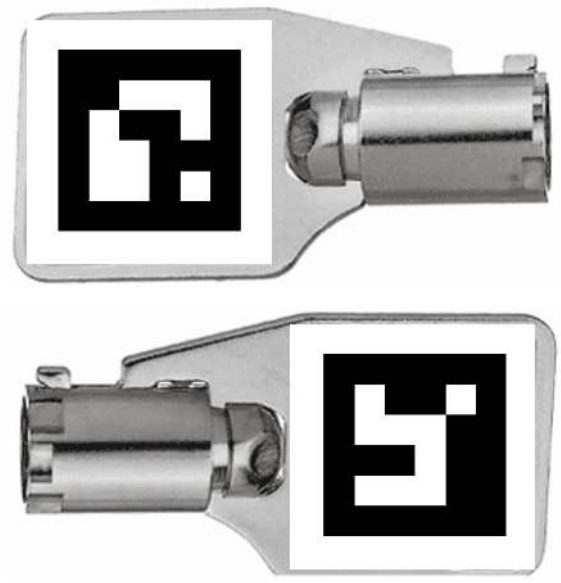
The keyboard carries 2 × 2 cm ArUco markers at its corners, and the key and the keyhole face carry 1 × 1 cm markers (§1.d.ii). Typing allows multiple attempts and carries no penalty for aborting, but the operators must be hands-off while the rover types. The rulebook deliberately leaves the task objects unspecified beyond these interfaces; the stated design intent is that every task should be possible for an astronaut with a single gloved hand S-URC27-QA EQ1.



The specified keyboard is a compact 87-key board with a 19.05 mm key pitch, around 2 mm of travel and about 0.45 N of actuation force. Those three numbers, not the marker size, set the arm's required repeatability and compliance.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*



The 2027 key task: extract a key from one lock, carry it, align it to a second lock and turn it against pin resistance. The marker sits on the key and around the keyhole because the key itself is too small to track from a distance.

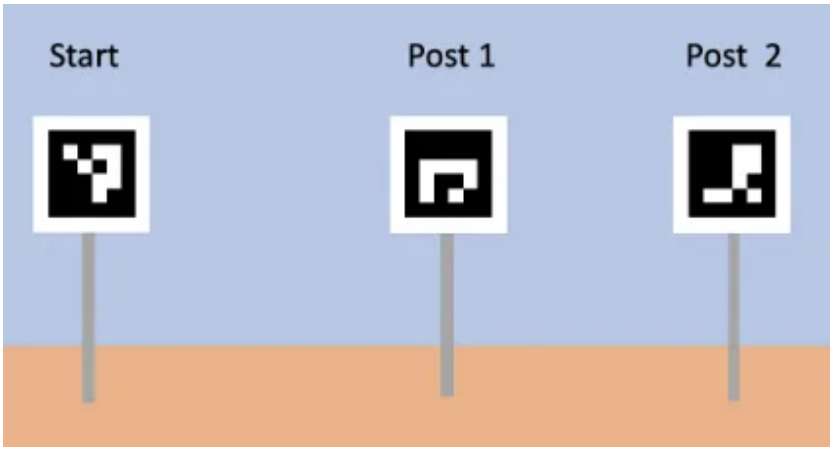
The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

Equipment Servicing may start as soon as 10 minutes after the Autonomy Mission, from the same C2 station (§1.a.iv). A fresh weigh-in and measurement are required inside that window unless the rover is unchanged and already compliant, or the team accepts the same penalty S-URC27-QA EQ4.

Clause 1.e: Autonomy

- The mission runs 40 minutes and holds two 50-point sub-missions, which may be attempted in any order (§1.e.i).
- Astronaut assistance runs the commands Follow, Stay, Fetch, Come and Give, in that order S-URC27-QA AUQ13. The Follow command is graded by modality: a device carried by the astronaut and a visual sign are worth 5 points, spoken audio 10, and a beckoning gesture 15 (§1.e.vi).
- Route finding sends the rover to two targets, 25 points each, reached by autonomously stopping within 1 m and signalling arrival (§1.e.xvii). One target sits on a hill where not all approaches are navigable, and the second is behind the hill, out of radio line of sight (§1.e.xv; §1.e.xvi).
- Targets are marked with three-sided posts carrying 20 × 20 cm ArUco faces in the 4x4_50 library, with 2.5 cm cells and a one-cell white border, set 0.5–1.5 m above the ground (§1.e.xii).
- A rear LED must show red for autonomous, blue for teleoperation and flashing green on arrival (§1.e.ii). There is no photometric specification, so daylight visibility has to be demonstrated S-URC27-QA AUQ3.
- Programming is allowed while the rover is stopped, including waypoints and stay-out zones, but not while it drives (§1.e.iii).
- Aborting to autonomous stop or return costs 20 % of that task, teleoperating back costs 50 %, and the autonomy-exit penalty is capped at 50 % (§1.e.xi; §1.e.xviii).
- The drone may reconnoitre during astronaut assistance but must land at the designated pad before route finding begins (§1.e.xiv).



The route targets are 20 cm faces at knee-to-chest height, not large markers on a wall. Detection range in full sun, at a post that may be 1.5 m above or below the camera axis, is the perception budget the 25 points rest on.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

Part 2: rover equipment

Part 2 holds the limits that decide whether a rover competes at all, plus the rules for drones and radio equipment.

Clause 2.a: mass and size

Clause 2.a carries the limits that disqualify a rover regardless of how well it drives `S-URC27-RULES`.

CONSTRAINT	2027 REQUIREMENT	CONSEQUENCE
Mission mass	50 kg or less when deployed for any mission	5 % of points earned per kilogram over, in each mission affected (§2.a.iii)
Total fielded mass	70 kg of all fielded parts	Strict; parts are removed until the total is under (§2.a.iii)
Rounding	Weight rounds down to the whole kilogram	50.9 kg records as 50 kg; 51.0 kg records as 51 kg (§2.a.iii)
Stowed size	Inside a 1.2 m cube in any orientation	40 % penalty at weigh-in (§2.a.ii)
Folding	Folding and articulating allowed, disassembly not	Wheels, mast, antenna and every protrusion count (§2.a.ii)
After weigh-in	Expand to any size	Only the stowed measurement is limited (§2.a.ii)
Weigh-in frequency	Before every mission, during setup	A re-weigh may fall inside the 10-minute Autonomy-to-Equipment window (§2.a.ii)
Configurations	May differ between missions	Each configuration at or below 50 kg, total at or below 70 kg (§2.a.iii)
Deployed items	Sub-rovers, cameras and relays count toward mass	Spares used to prepare the rover do not (§2.a.iii)
Tethers	No tethered power or communications to C2	Tethers between rover and deployables are allowed (§2.a.i)
Start	One rover leaves the gate carrying its deployables, except the drone	Deployables may not be pre-placed (§2.a.i)
Air-breathing systems	Prohibited for power or propulsion	Cooling fans and gas from a tank are allowed (§2.a.iv)

> **One push button, all power off**
>
> The E-stop is a red push button on the exterior, visible and reachable, not a rotating switch. It must stop movement and cease all power draw from every battery in one action, and every rover system must be off afterwards S-URC27-RULES §2.a.v. Any mobile deployable needs its own E-stop; a fuse is enough for a device that cannot move on its own, and the drone is exempt (§2.a.vi).

The size limit is a packaging requirement in practice. By simple geometry, at 250 mm wheels, a 0.9 m wheelbase and a 1.0 m track, a six-wheel rover is already about 1.15 m long and 1.20 m wide before the arm, mast or antenna, which leaves the stowed height and every protrusion to be managed by design rather than by folding.

Clause 2.b: drones

Drone rules apply across missions, and most of them are gates rather than scoring conditions S-URC27-RULES .

- A drone is optional, and full mission points are possible without one (§2.b.i). It may assist Delivery and Autonomy only S-URC27-QA RQ9.
- It must be rotary-wing and able to hover. Lighter-than-air and fixed-wing craft are excluded (§2.b.iii).
- It starts from the designated landing pad and may take off and land again as the judges allow (§2.b.iv).
- It carries an inert dummy mass equal to its battery mass at all times, and the 2027 rules do not waive that in high wind (§2.b.v; §2.b.vi).
- Mass is 5 kg or less including the dummy. A heavier drone may not be used (§2.b.vii).
- The drone is excluded from the rover's 50 kg and 1.2 m limits (§2.b.vii).
- Wind capability of 48 km/h is strongly suggested, which calls for a top speed of at least 64 km/h (§2.b.vi).
- Flight follows FAA rules: registration above 250 g, Remote ID, a TRUST certificate, below 120 m (400 ft) above ground, class G airspace, no flight in a national park and a visual spotter (§2.b.vii to §2.b.xii).

Clause 2.c: communications and the spectrum

The radio rules apply to every link the team operates S-URC27-RULES .

- There is no wireless internet at the field and satellite internet is not allowed (§2.c.i). Teams set up their own GNSS base station because no correction service is reachable S-URC27-QA RQ11.
- A local repeater or a drone relay is allowed and encouraged S-URC27-QA RQ5.
- 900 MHz operation is limited to 8 MHz of bandwidth inside three sub-bands: 902–910, 911–919 and 920–928 MHz, with the sub-band assigned at the event (§2.c.v).
- The restriction covers every link, including the C2 network, the rover, the drone and the astronaut (§2.c.v).
- The base antenna may be at most 3 m high, its base must be within 5 m of the C2 station, and teams bring at least 25 m of cable (§2.c.iv).
- 2.4 GHz has no URC-specific restriction, and normal radio regulation applies (§2.c.vi).
- Communications documentation and operator licences go to the URC Director by 2027-04-30, and later changes must be reported (§2.c.ii).
- Communications equipment is powered down at event sites when the team is not competing (§2.c.i).

Part 3: the competition

Part 3 holds the calendar, the field, the expense report, the site conditions and the penalty rules.

Clause 3.a: milestones and review gates

The calendar is fixed in the rulebook, and two of its dates are gates rather than deadlines in the ordinary sense S-URC27-RULES .

MILESTONE	DATE	WHAT IT REQUIRES	WHAT IT CONTROLS
Registration	2026-10-30	Intent and team details	Eligibility (§3.a.i)
Preliminary Design Review (PDR)	2026-12-02	Team, resources, project management and design state	Readiness assessment, not scored (§3.a.ii)
System Acceptance Review (SAR)	2027-03-05	Written and video package	100 points; the top 36 are invited (§3.a.iii)
Communications documentation	2027-04-30	Devices, bands and licences	Compliance (§2.c.ii)
Science Plan	2027-05-21	Written plan expanding the review version	Feeds the Astrobiology score (§1.b.i)
Rover Expense Report	2027-05-26	Expenses in work-breakdown format; update allowed 2027-06-02	10 % of total points per day late (§3.c.ii)
Finals	2027-06-02 to 2027-06-05	Four missions at MDRS	4 × 100 points (§3.a.iv)

System Acceptance Review due **128 days** 2027-03-05

The review is not a formality. It is 20 % of the score, and failing to reach the field costs the other 80 % as well. The detail is on Review, judging and awards.

Clause 3.b: the field and the C2 station

The field rules fix what the operators can see, where they sit and when the rover may be touched `S-URC27-RULES`.

- C2 stations are metal trailer units or MDRS structures with the course view blocked, supplied with 120 V at 60 Hz, tables and chairs (§3.d.i).
- Operators stay in C2, and following the rover for feedback is prohibited (§3.b.iv).
- Attendance at the finals is 24 students and 2 advisors, 14 people on site per day, and no more than 6 students in C2 at once `S-URC-LOGISTICS`.
- Nothing may be tested at MDRS before, during or after the event. Testing happens in town or in designated off-road areas (§3.d.iv).
- The field schedule is published about a week before the first field day, and task order is drawn per team `S-URC27-QA` CQ4.

Clause 3.c: the expense report

The financial side of the competition is a submission requirement, and it is enforced with the largest single penalty in the rules.

RULE	2027 VALUE
Project Budget Table	Required in both the PDR and the SAR (§3.c.i)
Final Rover Expense Report	Due 2027-05-26; an updated record may be submitted on 2027-06-02 (§3.c.ii)
Late penalty	10 % of total points per day (§3.c.ii)
Missing report	Disqualified if not submitted by the end of the event on 2027-06-05 (§3.c.ii)
Cap on fielded equipment	$26,000 of cash spent (§3.c.iii)
Inside the cap	Rover, modules, power sources, communications equipment, drone, astronaut EVA suit, base station, antennae, transceivers and all C2 equipment (§3.c.iii)
Outside the cap	Travel, accommodation, equipment transport, general hardware, non-fielded prototypes and one-for-one spare parts (§3.c.v)
Valuation	Used equipment at as-bought price; reused equipment at original cost or the current cost of an equivalent; rented equipment at purchase cost; donations at full value; shipping and tax included (§3.c.vi to §3.c.ix)
Non-US teams	The equivalent of $26,000 at a documented conversion rate between 2026-08-01 and 2027-06-02 (§3.c.x)
Receipts	May be demanded (§3.c.xi)

The report must follow a work-breakdown structure, and URC publishes a template for it `S-URC27-REQPAGE` `S-URC-EXPENSE-TEMPLATE`. The template's top-level nodes run from 0.0 total rover cost through 1.0 frame, suspension and chassis, 2.0 drive and energy, 3.0 electronics and processing, 4.0 command, control and communications, 5.0 arm, 6.0 science payload, and 7.0 miscellaneous parts. Each line carries a manufacturer, a vendor, a unit cost, a quantity and a cost rule that records how the value was obtained.

Two categories the 2027 cap explicitly covers have no node in the published template: the drone and the astronaut EVA suit. A team reporting either has to add a top-level line rather than fold them into miscellaneous parts.

Clause 3.d: site, environment and coordinates

The conditions below hold every competition day `S-URC27-RULES`.

CONDITION	2027 VALUE	DESIGN IMPLICATION
Air temperature	Up to 37.8 °C (100 °F)	Thermal derating, sun load, electronics cooling (§3.d.iii)
Dust	Winds frequently lift dust	Sealing, optical cleanliness, filter maintenance (§3.d.iii)
Rain	Light rain must be tolerated	Basic ingress protection; no heavy-rain requirement (§3.d.iii)
Light	All events in full daylight	Camera dynamic range and a daylight-visible LED (§3.d.i)
Altitude	1,309 m (4,295 ft) at MDRS	Air cooling and fan performance derated `S-URC-LOGISTICS`
Terrain	Soft sand, gravel, stony ground, boulder fields, vertical drops and loose slopes up to 1 km	Traction, tip resistance, ground clearance and suspension travel (§1.c.ii)
Coordinates	WGS 84, in decimal degrees, degrees decimal minutes or degrees, minutes, seconds	Site records also need elevation and accuracy range (§3.d.v; §1.b.iii)
Testing	Prohibited at MDRS	Terrain and environment must be checked before travel (§3.d.iv)



The rock garden is the reference obstacle the rules describe but do not dimension. The rulebook gives no step height, no boulder size and no slope limit, so an obstacle set has to be chosen and measured locally.

Mars Society, 2017. CC BY-SA 4.0. *Original source* *Full resolution*

Clause 3.e: penalties and interventions

Penalties are fractions of the points earned. They are added together, and the result floors at zero (§1.a.i) `S-URC27-RULES`.

EVENT	PENALTY	APPLIES TO
Overweight rover	5 % per kilogram over 50 kg	Each mission affected (§2.a.iii)
Rover outside the 1.2 m cube	40 %	That mission (§2.a.ii)
Team touches the rover or a deployable	20 % per intervention, additive	That mission, with the clock running (§3.e.viii)
Spectator touches the rover unrequested	Counted as an emergency stop	That mission, with points preserved (§3.e.ii)
Operator stops an autonomy task	20 % on that task	That task (§1.e.xi)
Teleoperating back to a visited location	50 % on that task	That task, capped at 50 % (§1.e.xviii)
Physical intervention in Autonomy	20 % each, not capped by the autonomy limit	That mission (§1.e.xi; §3.e.viii)
Late expense report	10 % of total points per day	Total score (§3.c.ii)

The autonomy cap is easy to misread. The 50 % ceiling applies to penalties for exiting autonomous mode on a task. A physical intervention is a separate penalty on the mission and is not capped, so one abort and one intervention cost 40 % rather than 20 %.

Some consequences have no stated points value: exceeding 70 kg, a 900 MHz violation, late communications documentation, a chemical spill and contact with the astronaut or an AR post `S-URC27-QA` `S-URC27-QA` RQ2 AUQ1. Those are compliance and safety questions that judges handle case by case, and the absence of a formula is not permission to plan for them.

What the rulebook does not fix

A design needs numbers the 2027 rulebook does not publish. Each item below is a real gap in the rules.

- **Task weights for three missions.** Astrobiology, Delivery and Equipment Servicing each carry 100 points, but only Autonomy publishes task-level values. The four Astrobiology scoring components and the nine Equipment Servicing tasks have no stated weight (§1.b.xi; §1.d.i).
- **Delivery stage times and the Stage-1 minimum.** The 2027 course description is promised before the competition, and until it appears the stage clock and the gate score are unknown (§1.c.i).
- **Slope, step and drop limits.** The rules describe the terrain and decline to quantify it. When asked, URC said that multiple routes exist at different inclinations and the values are not provided S-URC27-QA AUQ17.
- **Task object dimensions and forces.** The sample tube, cache, key, hose fitting, valve, switches and hammer have no published size, mass or force. Judge-like mockups are the only source (§1.d.ii).
- **The Equipment Servicing cache.** The rules do not say whether the cache the arm carries to the lander is the team's own Astrobiology cache or a judge-provided one, so the arm has to handle both (§1.d.ii).
- **Whether one instrument can satisfy two life detections.** The Q&A defines a bioassay and allows a biochemical assay correlated with biomolecule detection, but does not settle whether a single combined instrument counts as both required capabilities S-URC27-QA AQ9.
- **The penalty above 70 kg.** The rule gives a percentage only for the 50 kg mission limit; the 70 kg total is described as strict with no formula S-URC27-RULES S-URC27-QA §2.a.iii RQ2.
- **The 2027 review documents.** The PDR, review and Science Plan pages are still pending, and the 2027 rubrics will follow them S-URC27-PDR-PAGE S-URC27-SAR-PAGE S-URC27-SCIPLAN-PAGE.

Further reading

- Missions and scoring turns these clauses into the point values and the penalty arithmetic.
- Review, judging and awards covers the milestones and every document that feeds them in full.
- The 2027 competition places the calendar and the limits in the wider season.
- 155 engineering requirements is the clause-by-clause translation into testable statements.
- The Q&A record collects the rulings that change a design choice.
- Site and field constraints covers travel, altitude and the event site.

THE COMPETITION · CHAPTER 2.2

Missions and scoring

Every point in the 2027 competition sits in one of five places: four field missions at 100 points each and a System Acceptance Review at 100. The missions are deliberately different in shape. Astrobiology presses the clock hardest, Delivery gates itself behind a first stage, Equipment Servicing is a dense run of small manipulation tasks, and Autonomy is the only mission with published task values and the highest risk per point. This page gives the point values, the constraints that produce them, and the two arithmetic rules that decide what a score actually is.

500

TOTAL POINTS

Four missions at 100 plus the review at 100

2027

S-URC27-RULES

3.33–5.00

POINTS PER MINUTE, ASTROBIOLOGY

The densest clock of the four missions

2027

S-URC27-RULES

50 + 50

AUTONOMY SUB-MISSIONS

The only mission with published task values

2027

S-URC27-RULES

405

CEILING FOR A TELEOPERATION-ONLY ROVER

95 Autonomy points forfeited before any penalty

2027

S-URC27-RULES

Key takeaways

- The total is fixed at 500, and penalties are additive fractions of the points earned. A 10 % penalty and a 20 % penalty together leave 70 % of a mission score, not 72 %.
- Autonomy is 20 % of the total on paper, but 95 of its 100 points require genuine autonomous operation. A teleoperation-only rover is capped at 405 of 500 before it loses a single Equipment Servicing or review point.
- Three of the four missions publish no task weights, so 300 of the 400 field points cannot be priced at task level. Autonomy is the exception and its values are exact.
- The four missions separate the field by different margins. Across 2022–2026, mission scores explain variance in the order Autonomy 30.6 %, Equipment 27.7 %, Delivery 24.3 %, Science 14.4 %, with the review at 2.8 %.

The 500-point split

Each mission is scored independently and cannot go below zero S-URC27-RULES §1.a.i. The clock below is the on-course time, excluding setup, teardown and retrieval.

MISSION	CLOCK	POINTS	POINTS PER MINUTE AVAILABLE	AUTONOMY REQUIRED	DRONE
Astrobiology	20–30 min, plus 5 min cache removal	100	3.33–5.00	No	Not allowed
Delivery	30–60 min, two stages	100	1.67–3.33	No	Allowed, optional
Equipment Servicing	30 min, including up to 100 m of travel	100	3.33	Autonomous typing only	Not allowed
Autonomy	40 min, two sub-missions	100	2.50	95 of the 100 points	Reconnaissance only, and it must land before route finding
System Acceptance Review	Desk submission	100	Not meaningful per minute	No	No

Point density measures time pressure, not difficulty. Autonomy has the lowest density and the worst risk per point, because its points are concentrated in two 25-point targets and a 15-point gesture command. Astrobiology has the highest density and a broader spread of partial credit.

Astrobiology: 100 points in 20–30 minutes

The rover selects sites itself, analyses samples on board and caches one sample for the judges `S-URC27-RULES`. The mission is the scientific core of the competition and the place where the 2027 payload requirement changed most.

CONSTRAINT	2027 VALUE	NOTE
Site radius	0.5 km from C2	Straight-line radius; path length is not specified (§1.b.ii)
Roving time	20–30 min	Design to the 20-minute end (§1.b.ii)
Sites	At least 2 investigated, 1 sampled	Only the overall site is marked; sample points are not `S-URC27-QA` AQ5
Sample	At least 5 g from at least 10 cm depth	The portion below 10 cm sets the mass (§1.b.vi)
Life detection	Two capabilities, at least one a bioassay	The bioassay must detect a biomolecule `S-URC27-QA` AQ9
Cache	Sealed against inversion, removable in 5 min	Must be openable by the judges (§1.b.vii)
Documentation	Panorama with cardinal directions and scale, close-up with scale, GNSS with elevation and accuracy	Every site investigated (§1.b.iii)
Chemicals	No spill; hazardous chemicals pre-approved	The judges approve a safety plan, not a chemical `S-URC27-QA` AQ4
Science Plan	2027-05-21	Expands the review version and counts toward the score (§1.b.i)

The scoring rubric names four components: thoroughness of the site investigation, quality and applicability of the on-board analysis, quality of the returned sample, and scientific knowledge at the debrief (§1.b.xi). No weights are published for any of them, and the Science Plan is scored separately.

What the 2027 rule changed is the analysis payload. Two life-detection capabilities with at least one bioassay means assay chemistry, fluid handling, containment and a time-to-result that fits inside a 20-minute roving window. The sample requirement did not move: 5 g from below 10 cm, with a documented depth. The cheapest points in the mission are the ones a working camera and GNSS buy: two documented sites and a coherent account of what the data means. The expensive points are the bioassay and a clean, sealed, uncontaminated sample.

Delivery: 100 points across two stages

Delivery is the longest mission and the only one with an explicit stage gate. The rover moves objects for astronauts across terrain built to get harder with distance.

CONSTRAINT	2027 VALUE	NOTE
Structure	Two stages, with a minimum Stage-1 score to proceed	Remaining Stage-1 time is added to Stage 2 (§1.c.i)
Total time	30–60 min on course	Stage times are not published (§1.c.i)
Objects	Up to 5 kg, smaller than 40 × 40 × 40 cm	Grasp feature of 7.5 cm or less (§1.c.iii)
Course	Up to 1 km from the start gate	Soft sand, gravel, stony ground, boulder fields, drops and loose slopes (§1.c.ii)
Line of sight	Not guaranteed for more than 50 % of the course	Later tasks are deliberately out of sight (§3.d.ii)
Navigation	Working GNSS essential	Pickup and delivery coordinates are approximate (§1.c.iv)
Mobile deployables	Mini-rover prohibited	A drone may assist (§1.c.i)

The stage gate is what makes Delivery strategically different. Failing the Stage-1 minimum ends the mission there, so the first stage's points are worth more than their face value: they also unlock the second stage. The 2027 gate score, the stage durations and the task weights are all unpublished.

The most recent complete published course is the 2026 one, and it is a reference point rather than a 2027 rule: 35 minutes for Stage 1 and 20 for Stage 2, a gate at 40 points, a 20 cm placement tolerance and a 100-point task split `S-URC26-DM`. Those numbers describe 2026 and must not be imported into a 2027 plan, but they show the shape a published Delivery rubric takes.

Equipment Servicing: 100 points in 30 minutes

The mission is a dense series of manipulation tasks on a mock lander, reached by up to 100 m of relatively flat driving (§1.d.i). Every task completed earns points, the tasks may be attempted in any order, and no per-task value is published (§1.d.ii).

TASK FAMILY	WHAT IT DEMANDS
Sample tube into cache	Fine grasp, insertion, lid closure and a lock turn
Cache to lander	A carry of a sub-5 kg container by a handle at least 10 cm long
Drawer	Insertion into a tight space and closure
Latch and panel	A pull against a latch, then a hinged swing
Autonomous typing	A 3–6 letter key on a vertical keyboard, hands-off, with 2 × 2 cm markers at the corners
Key transfer	Extract a key, carry it, align to a second lock and turn it against pin resistance
Hose	Push on a 1.5 inch cam-lock fitting whose gasket is cut to half thickness
Valve	A quarter turn of a lever handle
Controls	Buttons, switches and knobs

Two features dominate the design. The first is the vertical keyboard: the rover has to locate a small board, find 2 cm markers in full sun, and strike keys with about 2 mm of travel and roughly 0.45 N of actuation force without knocking the board. The second is the key transfer, which adds a grasp-and-rotate task on an object small enough that the 1 cm markers exist because the key cannot be tracked otherwise.



The target face for the 2027 key task. The four 1 cm markers give the arm a pose reference, and the 7-pin keyway sets the turning torque the wrist has to supply without losing alignment.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

The schedule couples this mission to Autonomy. Equipment Servicing may begin 10 minutes after the Autonomy Mission from the same C2 station, and a fresh weigh-in and measurement may have to happen inside that window `S-URC27-QA` EQ4.

Autonomy: 100 points in 40 minutes

Autonomy holds two 50-point sub-missions that may be attempted in any order (§1.e.i) `S-URC27-RULES`. Unlike the other three missions, every value below is exact.

Astronaut assistance, 50 points

TASK	POINTS	REQUIREMENT
EVA suit with camera and microphone	5	Team-provided suit, monitored from C2 (§1.e.iv)
Drive to the astronaut	5	Autonomous stop within 3 m of the given location (§1.e.v)
Follow	15	Autonomous following, stopping within 3 m when the astronaut stops (§1.e.vi)
Stay	5	No movement while the astronaut moves more than 20 m away (§1.e.vii)
Fetch	10	Autonomous pickup of a rock pick from the ground (§1.e.viii)
Come	5	Autonomous drive to the astronaut, stop within 3 m (§1.e.ix)
Give	5	Place or drop the hammer on command (§1.e.x)

Follow is graded by command modality, and the values are alternatives rather than additions: a device carried by the astronaut is worth 5 points, a visual sign 5, a spoken word or phrase 10, and a beckoning gesture the full 15 (§1.e.vi). The gesture route is worth 10 points more than the device route, which makes gesture recognition the single largest optional capability in the mission.

Gestures may use arms, legs or torso only, with no fingers and no facial movement S-URC27-QA AUQ11. The command order is fixed: Follow, Stay, Fetch, Come, Give S-URC27-QA AUQ13.

Route finding, 50 points

ELEMENT	VALUE
Targets	2, worth 25 points each, in any order (§1.e.xii)
Stop tolerance	1 m, with the arrival indicated (§1.e.xvii)
Markers	Three-sided posts, 20 × 20 cm ArUco 4x4_50 faces, 2.5 cm cells, 0.5–1.5 m above ground (§1.e.xii)
Terrain	One target on a hill where not all approaches are navigable (§1.e.xv)
Link	The second target is behind the hill, out of radio line of sight (§1.e.xvi)
Start	Flat, easy ground, reachable by teleoperation (§1.e.xii)
Abort	20 % for an operator stop or return; 50 % to teleoperate back; capped at 50 % (§1.e.xviii)

The route sub-mission is where the 2027 rule is sharpest. Two 25-point targets, a 1 m tolerance and a link that fails behind a hill mean a single communications decision can cost a quarter of the mission. The Q&A makes the arrival indication easier, since the rear LED is sufficient and a C2 indication is not required for 2027 S-URC27-QA AUQ18.

The System Acceptance Review: 100 points and the gate

The review is a written and video package judged against other teams, and the 36 highest scores are invited to the field S-URC27-RULES §3.a.iii. It is 20 % of the maximum score and the only route to the other 80 %. The 2026 rubric split its 100 points as 48 for the video, 40 for the written report and 12 for the judges' assessment S-URC26-SAR ; the 2027 rubric is not yet published. The mechanics and the 2027 document status are on Review, judging and awards.

The penalty arithmetic

Penalties are fractions of the points earned, added together, with a floor at zero (§1.a.i). The rule is easy to misread, because it makes a penalty cheap on a bad mission and expensive on a good one.

PENALTY	TRIGGER	RATE
Overweight	Above 50 kg in a mission configuration	5 % of points earned per kilogram over, per mission (§2.a.iii)
Oversize	Outside the 1.2 m cube at weigh-in	40 % (§2.a.ii)
Intervention	Any direct team touch of the rover or a deployable	20 % each, additive, clock running (§3.e.viii)
Autonomy operator stop	C2 stops the rover or returns it to the astronaut	20 % on that task (§1.e.xi)
Autonomy teleoperation	Driving back to a visited location	50 % on that task, capped at 50 % (§1.e.xviii)
Late expense report	Report after 2027-05-26	10 % of total points per day (§3.c.ii)

The weight table shows why a mass margin is worth more than it looks. At 100 points earned, one kilogram over costs 5 points; five kilograms over costs 25.

OVERWEIGHT	PENALTY	LOSS AT 100 EARNED	AT 75	AT 50	AT 25
1 kg	5 %	5.00	3.75	2.50	1.25
2 kg	10 %	10.00	7.50	5.00	2.50
3 kg	15 %	15.00	11.25	7.50	3.75
4 kg	20 %	20.00	15.00	10.00	5.00
5 kg	25 %	25.00	18.75	12.50	6.25

Because the fractions add, combinations reach the floor quickly: 5 kg over plus two interventions plus an oversize is 105 %, and the mission scores zero whatever it earned. A single 2 kg overage and one intervention leave 70 % of the points earned, not 72 %.

What a teleoperation-only rover can score

The question that shapes a first season is how much the rover has to do on its own. The table answers it exactly, assuming every other mission and the review score full marks, which isolates the value of autonomy.

CAPABILITY SET	AUTONOMY POINTS	CEILING OUT OF 500	FORFEITED
Teleoperation only	5	405	95
Astronaut assistance only	50	450	50
Route finding only, plus a suited astronaut	55	455	45
Full autonomy, Follow by device or sign	90	490	10
Full autonomy, Follow by audio	95	495	5
Full autonomy	100	500	0

A drone adds no guaranteed points: full marks are possible without one (§2.b.i). It buys time and reduces risk in Delivery and Autonomy, and in Autonomy it must land before route finding starts (§1.e.xiv).

Where the points are at risk

Point values say what a task is worth; the score record says how often teams fail to collect. Across 162 finals team-years from 2022 to 2026, mission scores account for variance in the order Autonomy 30.6 %, Equipment Servicing 27.7 %, Delivery 24.3 % and Science 14.4 %, with the review at 2.8 % S-URC22-SCORES S-URC23-SCORES S-URC24-SCORES S-URC25-SCORES S-URC26-SCORES . The Autonomy Mission produced a zero score in 41 % of team-years, 67 of 162, more than any other mission.

Those are historical figures from the 2022–2026 rules, not a forecast for 2027, and the mission definitions changed. They do show where a first build is most likely to lose points through a single failure rather than through a slow accumulation of small ones. A rover that scores nothing in Autonomy loses the mission; a rover that loses a small task in Equipment Servicing loses a fraction of one.

What changed from 2026

DIMENSION	2026	2027
Astrobiology life detection	One life detection plus a second capability of choice	Two life detections, at least one a bioassay
Astrobiology documentation	Stratigraphic profile required	Removed from documentation and debrief
Delivery tasks	Marked-path task among the options	Not listed for 2027
Delivery scoring	Full rubric published for the course	Stage times, gate and weights unpublished
Equipment Servicing	USB-C cable task	USB task removed, key transfer added
Equipment Servicing hose	Gasket removed	Gasket cut to half thickness
Equipment Servicing hardware	Keyboard linked only	Keyboard and e-paper display both linked
Autonomy	30 min, two GNSS points, two posts, three objects	40 min, astronaut assistance plus two 25-point targets at 1 m
Autonomy penalties	20 % for teleoperation, free autonomous return	20 % for a stop or return, 50 % for teleoperation, capped at 50 %
Drone scope	Delivery only	Delivery and Autonomy, with a designated landing pad
Mass and size	50 kg, 70 kg, 1.2 m cube	Same values, with measurement explicit at every weigh-in

What the 2027 numbers do not yet say

- No task weights for Astrobiology, Delivery or Equipment Servicing. Only Autonomy publishes exact values.
- No Delivery stage durations and no Stage-1 minimum. The 2027 course description is promised before the competition.
- No partial-credit rates for route finding or for Delivery tasks.
- No cumulative-distance cap in the Autonomy Mission, unlike the 2 km cap in 2026.
- No published rubric for the 2027 System Acceptance Review, and no Science Plan specification beyond the Q&A guidance `S-URC27-SAR-PAGE` `S-URC27-SCIPLAN-PAGE`.
- No published dimension, mass or force for any task object, from the sample tube to the hammer.

Where a 2026 number is used on this page, it is labelled as 2026 and treated as a reference point. The 2027 rulebook and the Q&A current to 2026-10-11 are the only authorities for a 2027 value `S-URC27-RULES` `S-URC27-QA`.

Further reading

- Rules, limits and the field states the clauses these points come from.
- Review, judging and awards covers the 100-point review in full.
- The 2027 competition places the season around the five scored components.
- Autonomy: astronaut assistance and Autonomy: route finding work each sub-mission in detail.
- Astrobiology and Equipment servicing carry the mission-level engineering choices.
- What separates the top from the rest explains the variance figures behind the risk section.

THE COMPETITION · CHAPTER 2.3

Q&A rulings

The 2027 Questions and Answers page is part of the rules, not commentary on them: the rulebook states that all matters addressed in the Q&A are applicable to the requirements and guidelines S-URC27-RULES. Its 65 answers settle questions the rulebook leaves open, and a dozen of them change a subsystem choice rather than a detail. This page organises every answer by mission, states the design consequence, and marks the places where the Q&A overrides the rulebook text.

65

ANSWERS PUBLISHED

Six categories, Competition to Autonomy

2027

S-URC27-QA

2026-10-11

LATEST REVISION

Third update after 2026-09-15

2027

S-URC27-QA

12

RULINGS THAT CHANGE HARDWARE

Satellite ban, own GNSS base, bioassay, no-touch keyboard

2027

S-URC27-QA

8 MHz

900 MHZ BANDWIDTH

Applies to the links to rover, drone and astronaut

2027

S-URC27-RULES

Key takeaways

- **Satellite internet is banned and a relay is encouraged.** A local repeater or a drone relay is the endorsed answer to the obstructed links in Delivery and Autonomy, which makes a relay an architecture element rather than an optimisation S-URC27-QA RQ5.
- **The team brings its own GNSS base station.** There is no internet in the field and no correction service, so differential GNSS depends on base hardware the team owns, powers and erects inside the 15-minute setup S-URC27-QA RQ11.
- **The bioassay definition is the payload gate.** A molecule formed only by biology is required; an amino-acid assay alone does not qualify, and a chemical method alone does not count as the bioassay S-URC27-QA AQ9.
- **The keyboard is not a teach-by-touch task.** Touching any part of the keyboard in manual control is penalised, so alignment comes from the four corner tags and the arm servos on them S-URC27-QA EQ2.

How the Q&A binds a design

The page was first published on 2026-09-15 and revised on 2026-10-03 and 2026-10-11; it can change again before the System Acceptance Review, and it changed mid-cycle in earlier years S-URC27-QA. A team that treats it as a fixed document can be designing to an answer that no longer stands. The working rule is that the Q&A is the later ruling: where the two disagree, design to the answer and keep the clause as the trace.

Two structural facts follow. First, every ruling is binding, so a requirement that rests on an answer should cite the answer as its trace rather than a nearby clause. Second, an answer that grants an exception has to be requested early. The judges need lead time to test a Mars-feasibility argument, and the answer lands on the page for every team at once S-URC27-QA RQ12.

The twelve rulings that shape the vehicle

#	RULING	REFERENCE	WHAT IT CHANGES
1	Satellite internet is banned; local repeaters and a drone relay are explicitly encouraged	RQ5	Removes the backhaul shortcut and makes a relay a required architectural option
2	No correction service and no internet: the team runs its own GNSS base station	RQ11	Adds a receiver, antenna, mount, correction link and 120 V power to the setup
3	A magnetometer or digital compass is allowed	RQ3	A legal heading source for dead reckoning and for the hill route
4	Microphones on the rover are allowed	RQ4	A diagnostic audio channel during Equipment Servicing
5	An off-the-shelf drone is allowed; no drone in Astrobiology or Equipment Servicing	RQ9	Scopes the airframe to two missions and moves effort to integration
6	1 km applies to Delivery and Autonomy; the other two courses are much closer	RQ13	Two link budgets, not one
7	Tank-fed pneumatics are legal; ambient-air systems need an exception	RQ12	Bottled gas is available for sample handling, latches and deployment
8	A human drone battery swap is an intervention	RQ14	A 20 % penalty on that mission unless the swap is robotic
9	A bioassay means a molecule formed only by biology	AQ9	Two independent life detections, one meeting the strict test
10	The sample evaluated for life comes from at least 10 cm depth	AQ10	A depth-indexed sampling record, not a scoop
11	25 m of cable is required for the rover and for the drone; operators sit inside the command post with no view	DQ2	Command-post layout, cable runs and pilot ergonomics
12	Touching the keyboard in manual mode is penalised; the display is a 4.3 in e-paper panel	EQ2	Vision-servoed alignment and a compliant end effector

Competition and conduct

CQ1 · The page is the record. Questions may be sent by email and answers are added to the page, so the page itself is the durable reference, and a question answered by email is worth requesting for the page S-URC27-QA CQ1.

CQ2 · No video of the 2027 tasks. The judges did not plan to publish video of the updated tasks or their infrastructure; earlier competition videos are described as a good source on earlier tasks S-URC27-QA CQ2. Task geometry therefore comes from the rulebook text plus older footage, and 2026 hardware cannot be assumed where the 2027 text differs.

CQ3 · Visa support. URC provides a letter of support on request, issued after the team is registered, and advises applying for as many members as possible S-URC27-QA CQ3. Registration gates the letters, and the registration deadline is 2026-10-30.

CQ4 · The schedule arrives about a week ahead, and the task order is drawn by lot. The order differs per team S-URC27-QA CQ4. That is a configuration requirement: the rover has to be mission-agnostic in its transport state, and any module swap has to be a rehearsed operation inside the 15-minute setup.

CQ5 · Regional pricing is ordinary pricing. The judges treat regional cost differences as normal pricing factors and report that across prior years the correlation between declared spend and score is extremely weak, with a wide spread S-URC27-QA CQ5. The cap is a documentation exercise rather than a performance lever.

CQ6 · Student discounts count. A student-only national discount may be used for valuation when any student of that country could obtain it, not only students of one university S-URC27-QA CQ6.

Rover and drone

RQ1 · The rover may change, but not mid-mission. Changes are allowed at any point, including during the field event, subject to the mass and cost limits. An intervention exists to repair a broken rover, not to swap a module S-URC27-QA RQ1. Modularity is legal between missions; a hot swap during one is not a free action.

RQ2 · The 70 kg limit is strict. Above it, parts are discarded or drilled until the rover complies; no percentage penalty schedule exists S-URC27-QA RQ2. The 50 kg figure per configuration keeps the 5 % of points per kilogram sliding penalty, so the aggregate limit is a gate and the configuration limit is a price.

RQ3 · A magnetometer is allowed. The judges said they could not verify that teams avoid a digital compass, and permitted it even though it is not Mars-analogous S-URC27-QA RQ3. It is the cheapest way to bound yaw drift when the differential fix is obstructed.

RQ4 · Microphones are permitted. Audio from the rover is a legitimate diagnostic channel for a stalling arm or a slipping wheel S-URC27-QA RQ4. The same answer states that spectators may not coach operators, which is a conduct rule rather than a hardware rule.

RQ5 · No satellite backhaul, but a relay is encouraged. A satellite relay is considered outside the spirit of the competition, while a rover-deployed repeater or a drone used as a relay is described as strongly encouraged S-URC27-QA RQ5. The relay carries mass and cost from the first design review, not from the week before the event.

RQ6 · Licensed high power on 2.4 GHz is permitted but discouraged. The judges would rather licensed teams use amateur bands away from the crowded 2.4 GHz and 900 MHz allocations S-URC27-QA RQ6.

RQ7 · Compliance is a deadline item. All applicable FCC regulations apply, particularly 47 CFR Part 15 Subpart C, and devices and operator licences go to the URC Director by 2027-04-30 S-URC27-QA RQ7. A change after that date has to be notified immediately.

RQ8 · A licence takes weeks, so it is a schedule item. The judges point at amateur licensing material and note that a licence is easy to obtain in a few weeks, and that foreign nationals may hold a US licence under reciprocal arrangements S-URC27-QA RQ8. The task needs an owner in September, not April.

RQ9 · The drone flies in two missions only. No drone may be used in Astrobiology or Equipment Servicing, and an off-the-shelf airframe is acceptable because nothing in the rules rewards self-build S-URC27-QA RQ9. The FAA requirements still apply in full.

RQ10 · Prior maps are sanctioned. Public LIDAR, photographs and maps may inform the algorithms, and URC intends to post at least one LIDAR dataset S-URC27-QA RQ10. Planning the hill route before the rover moves is legal, and it is the cheapest way to survive the route-finding sub-mission.

RQ11 · The team owns the GNSS base station. With no internet in the field and no correction service allowed, differential GNSS requires a base receiver, an antenna and a link to the rover S-URC27-QA RQ11. The base has to be standing and converged inside the 15-minute setup, and it draws on the 120 V supply at the command post.

RQ12 · Tank gas is legal, ambient air is not automatic. A system whose gas comes from a tank is acceptable, cooling fans need not be custom, and any other ambient-air system needs a case-by-case Mars-feasibility exception S-URC27-QA RQ12. A compressor or vacuum pump that moves ambient air sits squarely in the exception category, so the argument is made early.

RQ13 · 1 km is two missions, not four. Astrobiology and Equipment Servicing courses will be significantly closer, while Delivery and Autonomy may send the rover up to 1 km from the command post with the total traverse exceeding 1 km S-URC27-QA RQ13.

RQ14 · A human battery swap is an intervention. The rover may carry and deploy a spare battery, and anything carried counts in mass and cost, but the drone or the rover has to perform the swap S-URC27-QA RQ14. A hand swap costs 20 % of that mission's points, and a battery left on the ground has to be recovered inside the 10-minute teardown.

RQ15 · The drone may deliver anything except a mini-rover. A drone-dropped repeater on a ridge is legal, at the cost of payload mass inside the 5 kg limit and a recovery inside the teardown S-URC27-QA RQ15. In the route-finding sub-mission a rover of any kind must start from the designated start location.

Astrobiology

AQ1 · Two science plans, and the May one is scored. The review version is about one page in the report with some time in the video; the May version is about five pages with background science and counts toward the mission score S-URC27-QA AQ1. The plan is a scoring artefact due 2027-05-21.

AQ2 · Life detection means bio-molecules. The capability has to search for biological components; it does not need to be expensive or sophisticated S-URC27-QA AQ2. Instrument selection is judged on the molecule it detects.

AQ3 · Any instrument class is acceptable. A purpose-built instrument from about one hundred US dollars to a few thousand is a legitimate route, and price is not the discriminator S-URC27-QA AQ3. The assay's biological specificity is.

AQ4 · URC approves plans, not chemicals. Each chemical needs a safety plan covering transport, storage, use and disposal, anything above US DOT limits is not allowed, and anything below still needs permission from the science judges; most corrections concerned cradle-to-grave tracking rather than concentration S-URC27-QA AQ4. Chemistry selection happens before the review so the plan and the report agree.

AQ5 · The site is marked, the sample points are not. The team chooses where to sample, the soil is in situ and there are no trays, so site triage happens live from rover imagery S-URC27-QA AQ5. That needs a written stop-or-drive rule and the documentation package the mission requires.

AQ6 · Expect hard ground. Soil and sub-surface samples vary over short distances from reasonably easy to very hard, and there is no caliche surface at the science site S-URC27-QA AQ6. The sampling mechanism is designed for a worst-case penetration force and a jam recovery path.

AQ7 · Contamination control is argued, not maximised. The team has to justify its level of cleanliness rather than meet a stated standard S-URC27-QA AQ7. A cleanable scoop, a wipe-down procedure and sample-path segregation are the evidence.

AQ8 · Two people must stay clean. Only members who stayed inside the command post during the science mission and have not seen the site with their own eyes may take part in the science briefing S-URC27-QA AQ8. Anyone who walked the course as a runner or photographer is excluded.

AQ9 · The bioassay test. A bioassay must detect a molecule currently known to be formed only by biology or in the presence of biology: proteins, sugars, chlorophyll, nucleic acids, ATP or carotenoids. An assay for a molecule that could form abiotically, such as amino acids, is not sufficient on its own, and a chemical method alone does not satisfy the requirement S-URC27-QA AQ9. The payload needs two independent capabilities, at least one that passes this test, inside a 20–30 minute roving window.

AQ10 · Depth is provenance, not a lockout. Multiple depths may be compared, but the sample evaluated for life must come from at least 10 cm, it should be a specific sample rather than bulk soil, and the location including depth is part of the mission S-URC27-QA AQ10. An auger with depth control is acceptable; a record without depth is not.

Delivery

DQ1 · No drop, slope or step numbers. The course runs from flat ground to steep loose slopes and even vertical faces, with no values published S-URC27-QA DQ1. The obstacle envelope is a team decision that has to be defended with measurements on an analog site.

DQ2 · 25 m of cable, and the operators cannot see. Both the drone and the rover operators sit inside the command post with no view of the course, not next to the door or the window, and the controller and video wiring runs to antennas outside S-URC27-QA DQ2. Two independent cable runs are needed if the two systems use separate radios.

Equipment servicing

EQ1 · The design criterion is one gloved hand. The equipment details are intentionally withheld to reward flexible design, and every task should be reasonably doable by an astronaut using a single gloved hand S-URC27-QA EQ1. Design for a class of objects, not a fixed geometry list.

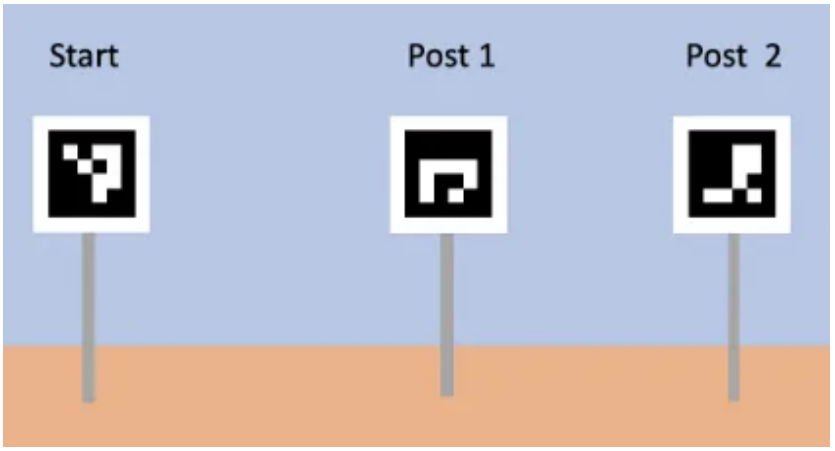
EQ2 · No teaching by touch. The tags exist to aid autonomous alignment; touching any part of the keyboard in manual control is penalised, the display must not be touched, and the surrounding metal of the lander may be touched S-URC27-QA EQ2. Visual servoing on the corner tags is the only legal alignment method.

EQ3 · E-paper, no key repeat, judge reset. The display remains a 4.3 in e-paper panel, keys do not auto-repeat while held, and a judge resets the display to restart an attempt S-URC27-QA EQ3. The read-back delay and the low contrast in sunlight belong in the typing loop, and an abort costs time rather than points.

EQ4 · The 10-minute transition may include a weigh-in. A new weigh-in and measurement is required unless the configuration is unchanged and previously compliant, or the team accepts the same penalty it already carries S-URC27-QA EQ4. Either freeze the configuration or rehearse the swap inside the window.

EQ5 · Confirmation may be requested, scoring may not. For tasks with poor visual feedback a team may ask for confirmation that a task was completed, not for a score, and URC decides which tasks allow it just before the competition S-URC27-QA EQ5. Design for camera-visible evidence of completion instead.

The published page numbers the last two Equipment answers without a matching answer number. The content is unambiguous, and the order here follows the page S-URC27-QA .



The two Autonomy targets are three-sided posts with 20 × 20 cm faces and ArUco 4x4_50 tags, 0.5–1.5 m above the ground. Their coordinates are exact (AUQ16), which is what makes the 1 m full-credit radius reachable with a differential fix and a final visual approach.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

Autonomy

AUQ1 · Nothing is penalised for contact, but contact ends runs. Hitting the astronaut or a post carries no points penalty, yet a wheel caught on a post is common and an astronaut may stop a rover that comes too close `S-URC27-QA` AUQ1. Treat clearance as a functional requirement.

AUQ2 · Off-board compute is legal, human help is not. A base-station computer may do autonomy computation, but no person may assist while the rover drives autonomously, and one target is deliberately out of communications `S-URC27-QA` AUQ2. The far target has to be solvable onboard.

AUQ3 · No photometric specification for the LED. Teams verify daylight visibility themselves, a sunshade is suggested, and any light that clearly shows state is acceptable `S-URC27-QA` AUQ3. This is a cheap item that fails inspection if it is not tested at noon.

AUQ4 · A first abort returns to the command post. If the rover never reached the previous point and the attempt is aborted, it returns to the start area in front of the command post `S-URC27-QA` AUQ4. Return to start is a separate behaviour from return the way the rover came.

AUQ5 · The astronaut is a communications node. The suit microphone and camera may relay information to the rover, and the command post may talk to the astronaut if the link is two-way `S-URC27-QA` AUQ5. A body-worn relay is legal because nothing is deployed.

AUQ6 · Modes may be switched; only Follow! is graded on modality. The 5 / 5 / 10 / 15 split applies to Follow! alone, and every other command pays full points whatever method is used `S-URC27-QA` AUQ6. A simple reliable method is enough for the four binary commands.

AUQ7 · 20 m, in the field of view, from the astronaut. Fetch! and Come! place the astronaut at least 20 m away and directly visible, and all commands must originate from the astronaut `S-URC27-QA` AUQ7. That sets camera field of view, resolution and microphone signal-to-noise.

AUQ8 · The hammer is on the ground nearby. It lies in the immediate area, and the rover should not need to move to see it unless it parked on top of it or missed the 3 m stop `S-URC27-QA` AUQ8. The arm reaches it from the parked pose with a downward-looking camera.

AUQ9 · The LED reports mode, not completion. The field judge reads it to tell autonomous driving from teleoperation, flashing green is not required for Fetch!, Stay! or Give!, and the light should follow the rover's actual state rather than the operator's intent `S-URC27-QA` AUQ9.

AUQ10 · The astronaut sets the distance. The astronaut may walk up until the gesture fills the frame or move as far away as desired `S-URC27-QA` AUQ10. Detection has to work at both ends of that range.

AUQ11 · Gross-body gestures only. Finger and facial gestures are prohibited, justified by suit restriction; arm, leg and torso motion are allowed without further limit `S-URC27-QA` AUQ11. A thumbs-up or a point is non-compliant.

AUQ12 · The suit is not in the mass budget. Coverage is neck, wrists and ankles with closed-toe shoes, colours are free, and the suit itself is excluded while overalls and wearable electronics count `S-URC27-QA` AUQ12. The classification belongs in the expense report.

AUQ13 · The command order is fixed: Follow!, Stay!, Fetch!, Come!, Give! `S-URC27-QA` AUQ13. An operator interface that mirrors the five states reduces both error and time.

AUQ14 · The astronaut holds the geometry. Judges show the astronaut the approximate locations to move between, so the astronaut, not the operators, receives the route information `S-URC27-QA` AUQ14.

AUQ15 · Silence is required for a gesture Follow!, and nowhere else. Astronaut to command-post communication is encouraged at all other times `S-URC27-QA` AUQ15. The audio channel and the gesture channel impose opposite disciplines.

AUQ16 · The tag coordinates are exact. The tags help the judges and spectators as much as the rover reach the 1 m full-credit radius `S-URC27-QA` AUQ16. This is the strongest argument for a team-owned base station.

AUQ17 · The hill has several routes and no published inclination. The team chooses the route on natural terrain and no values are provided `S-URC27-QA` AUQ17. Route planning from elevation data and a traversability estimate is the intended solution; a fixed path cannot work.

AUQ18 · The LED is enough for the blind target. Arrival at the out-of-line-of-sight target needs no command-post indication, and for the visible target a display indication is recommended but not required `S-URC27-QA` AUQ18. A stored arrival event plus the LED is the minimum viable design.

AUQ19 · Repeaters count. They are allowed and encouraged in Autonomy and must be carried in both the mass and the cost budgets `S-URC27-QA` AUQ19. A repeater is fielded equipment, with a recovery inside the 10-minute teardown.

AUQ20 · The elevation model is a bonus. The format is likely GeoTIFF and release approval and timing are not guaranteed `S-URC27-QA` AUQ20. Public LIDAR and imagery stay on the critical path; a URC model does not.

AUQ21 · Hands off while the rover moves. Because the judges cannot tell which devices are connected, operators must not touch the controls while the rover drives autonomously, and programming while stopped is allowed `S-URC27-QA` AUQ21. The rule is enforced by observable behaviour, so the command post needs one visible interface and a rehearsed freeze.

AUQ22 · The spotter is a safety and legal role. The spotter keeps visual line of sight and may take control, and permission is requested before take-off and landing `S-URC27-QA` AUQ22. That is a second field member with no other mission role.

AUQ23 · A mini-rover may score in Autonomy. It is legal there and prohibited in Delivery, it needs its own emergency stop, and its mass counts inside the configuration `S-URC27-QA` AUQ23.

AUQ24 · Any camera may recognise the gesture. The suit camera is as acceptable as the rover's own `S-URC27-QA` AUQ24. Pointing the sensor at the astronaut relaxes the rover's own field-of-view requirement.

AUQ25 · The pilot flies from the command post. The astronaut may not fly the drone, and the pilot is the only person who may touch controls while the rover is autonomous `S-URC27-QA` AUQ25. Two physically separate control positions are needed.

AUQ26 · The astronaut wears but does not deploy. Nothing may be dropped in the field, while any electronics may be built into the suit `S-URC27-QA` AUQ26. Worn relay electronics are legal; a repeater placed on the ground by the astronaut is not.

AUQ27 · Voice commands may arrive either way. Spoken commands may be radioed to the rover or heard directly `S-URC27-QA` AUQ27. The choice has to be deliberate, because the two paths need different hardware.

Where the rulebook and the Q&A diverge

TOPIC	RULEBOOK TEXT	Q&A RULING	WHAT TO DESIGN TO
Drone use	Advantageous for Delivery and Autonomy, not required	Banned in Astrobiology and Equipment Servicing	No drone in those two missions; a drone is a planned option for the other two
Field communications	No normal wireless internet; satellite internet not allowed	Local repeater or drone relay strongly encouraged	A relay is an architecture element
GNSS augmentation	Silent on corrections	No internet, so the team runs a base station	Base station is required hardware
Air systems	No air-breathing power or propulsion	Tank-fed gas and fans are fine; ambient-air systems need an exception	Bottled gas is legal; exceptions are filed early
1 km range	Rovers not expected to travel as far as 1 km	Applies to Delivery and Autonomy; the other courses are closer; traverse may exceed 1 km	Two link budgets
Intervention	20 % and 50 % penalty structure	The mechanism is for repair, not for swapping modules	No mid-mission module swaps
70 kg limit	70 kg total, 50 kg per configuration, 5 % per kilogram over 50	70 kg is strict; parts are discarded to comply	Mass control is a gate, not a penalty line
Weigh-in at each mission	Weighed and measured at every mission setup	Not repeated inside the 10-minute transition under two conditions	Freeze the configuration or rehearse the swap
Keyboard contact	Tags aid alignment	Touching the keyboard in manual mode is penalised	Vision-only alignment on the tags

What the Q&A does not settle

Four questions remain open, and each is worth a written question rather than an assumption.

- **The antenna limits and the GNSS base.** The 3 m height and 5 m from the command post limits are written for the communications antenna; whether they apply to a separate base antenna is not addressed. Planning to 3 m is the safe reading.
- **The repeater and the emergency-stop rule.** A repeater cannot move on its own, so a fuse should suffice, but no answer says so. The reading comes from the wording of the stop rule, not from a ruling.
- **The 2027 suit specification.** The requirement is inherited from the rulebook plus one answer, and no separate 2027 suit document was found.
- **The Science Plan and Chemical Safety Plan formats.** The 2027 science plan page states that details will be posted in early 2027, so planning to the 2026 structure is the only available route.

Further reading

- Rules, limits and the field carries the clauses these answers interpret.
- Missions and scoring places each ruling inside its mission.
- 155 engineering requirements cites the Q&A ruling behind 51 of its rows.
- Deadlines turns the licensing, plan and documentation answers into dates.
- Open questions collects the questions still owed to URC.

THE COMPETITION · CHAPTER 2.4

155 engineering requirements

The 2027 rulebook is written to be judged, not to be designed from. This page turns it into 155 shall-statements, each traced to a rule clause or a Q&A ruling, each given a mission, a type and a primary check. Half of them are gates: a rover can be legal, score well, and still fail on mass, size, power cut-off or a document deadline. The full table is filterable by domain, mission and type, and it is the yardstick every subsystem chapter in this book is measured against.

155
REQUIREMENTS
Eleven domains, each row a testable statement
2027
S-URC27-RULES

79
MANDATORY
Rule breach, penalty or ineligibility if missed
2027
S-URC27-RULES

61
SCORING-DRIVEN
Performance needed for points, not a breach
2027
S-URC27-RULES

15
ADVISORY
Encouraged practice with no direct points
2027
S-URC27-RULES

Key takeaways

- **Half the table is a gate.** Mandatory rows are 79 of 155, and most of them are checked at a weigh-in, a timed setup or a document review rather than on the course.
- **51 rows cite a Q&A ruling**, 40 of those alongside a clause. The operational detail that decides a mission usually lives in the answer, not in the rule.
- **23 rows are derived.** They do not transcribe a sentence; they state how a rule becomes a design or test action, and dropping them leaves the rule unimplemented.
- **The weakest traces are the scoring-driven rows**, because the 2027 task weights for Astrobiology, Delivery and Equipment Servicing are not published. Those rows can be built to and cannot be priced.

What a traced requirement is

Each row pairs a shall-statement with its trace, its missions, a type, a primary check, the affected subsystems and a note. The identifier is stable and is never renumbered when a row is inserted, so a cross-reference from a subsystem chapter stays valid. The trace names the rule clause, the Q&A question, or both, and says `derived` when the requirement follows from a rule without reproducing its wording. The mission field names the events the row applies to. The type says whether the row is a gate or a target. The check names the primary method. The affected subsystems field is the first pass at an interface list. The notes carry the derivation, the penalty, the 2026 contrast or the missing value.

A derived row is not optional. It is the difference between a rule and an activity: the rule says the sample comes from at least 10 cm below the surface, and the derived rows say what the sampling mechanism, the depth record and the operator interface have to do about it.

The rulebook traces are written as clause references, so the table can be read against Rules, limits and the field clause by clause `S-URC27-RULES`. A Q&A trace resolves on Q&A rulings, and every derived row is a statement about this rover rather than a quotation `S-URC27-QA`.

The eleven domains

DOMAIN	ROWS	MANDATORY	SCORING-DRIVEN	ADVISORY	WHAT IT GOVERNS
SYS	14	13	0	1	Mass, the 1.2 m cube, integration, fielded-equipment totals
MOB	11	2	9	0	Terrain, range, slope, stop distance and payload carriage
ARM	21	2	17	2	Reach, grasp, torque, insertion, typing and fine alignment
SCI	20	16	4	0	Life detection, sampling depth, cache, chemical safety, plan
PWR	9	5	2	2	Battery-only operation, power cut-off, thermal limits, 120 V
COM	15	11	4	0	Links, bands, antennas, cable, documentation, relay support
CMP	15	2	10	3	GNSS, tags, cameras, gesture and audio detection, compute split
SFT	19	3	14	2	Autonomy, aborts, teleoperation limits, mode signalling
OPS	12	11	0	1	Attendance, runners, the science briefing, schedule, deliverables
SAF	8	6	0	2	Emergency stop, deployable stops, suit, drone safety
DRN	11	8	1	2	Airframe, mass, dummy mass, launch pad, pilot, airspace
Total	**155**	**79**	**61**	**15**	

Two domains carry the gate risk: SYS, where 13 of 14 rows are mandatory, and OPS, where 11 of 12 are. Two carry the performance risk: ARM and SFT, where 31 of 40 rows are scoring-driven and the task weights are not published `S-URC27-RULES`.

The full table

The table below is the complete set. Each column sorts; domain, mission and type filter; the search box matches the requirement text and the notes. Use the identifier links to reach a single row, and the download link under the table for a spreadsheet.

ID	DOMAIN	REQUIREMENT	TRACES TO	MISSIONS	TYPE
REQ-SYS-001	System, mass, size and integration	The rover, in its stowed transport configuration, shall fit entirely within a 1.200 m × 1.200 m × 1.200 m box in any orientation, including wheels, antenna and every protrusion, without disassembly.	§2.a.ii	All	mandator
REQ-SYS-002	System, mass, size and integration	The rover's judge-recorded mass, after rounding down to the nearest whole kilogram, shall be ≤50 kg in each deployed mission configuration.	§2.a.iii	All	mandator
REQ-SYS-003	System, mass, size and integration	The total mass of all fielded rover parts across all events, after rounding down, shall be ≤70 kg.	§2.a.iii; RQ2	All	mandator
REQ-SYS-004	System, mass, size and integration	The rover shall be a single stand-alone, off-grid mobile platform that carries all deployable items except the drone out of the designated start gate; no tethered power or communications shall run to the C2 station.	§2.a.i	All	mandator
REQ-SYS-005	System, mass, size and integration	The rover shall be weighed and measured by the judges during the setup time of every mission, and the team shall present it for measurement in a configuration that fits the 1.2 m box (folding/articulation allowed, no disassembly).	§2.a.ii	All	mandator
REQ-SYS-006	System, mass, size and integration	Modular payloads (e.g. arm, science module) shall be swappable between missions such that each mission configuration is ≤50 kg and the complete fielded set is ≤70 kg.	§1.a.iv, §2.a.iii derived	All	mandator
REQ-SYS-007	System, mass, size and integration	The rover shall not require disassembly to fit the 1.2 m box; only articulation, folding or bending is permitted.	§2.a.ii	All	mandator
REQ-SYS-008	System, mass, size and integration	The team shall set up all rover systems, including all communications, and be ready to compete within 15 minutes of receiving access to the C2 station.	§1.a.ii	All	mandator
REQ-SYS-009	System, mass, size and integration	The team shall remove all equipment from the C2 station and course within 10 minutes after the mission, switching off radio equipment immediately.	§1.a.ii	All	mandator
REQ-SYS-010	System, mass, size and integration	The team shall return or collect the rover, drone and all deployed items within 10 minutes after the mission time expires.	§1.a.iii	All	mandator
REQ-SYS-011	System, mass, size and integration	The rover shall be ready to start Equipment Servicing within 10 minutes after Autonomy, from the same C2 station and on an adjacent course, including any required re-weigh/measurement.	§1.a.iv, §1.d.i; EQ4	Auto, Equip	mandator
REQ-SYS-012	System, mass, size and integration	No rover power or propulsion system shall ingest ambient air for combustion or any other chemical reaction that yields energy.	§2.a.iv; RQ12	All	mandator
REQ-SYS-013	System, mass, size and integration	The rover shall not use a mini-rover or other mobile deployable in the Delivery Mission; mobile deployables elsewhere shall meet their own E-stop requirement.	§1.c.i, §2.a.vi	Del (prohibition); All (E-stop)	mandator
REQ-SYS-014	System, mass, size and integration	The rover shall be designed for field access and modification throughout the competition, subject to re-weigh/measurement after changes.	§1.a.v derived	All	advisory
REQ-MOB-003	Mobility	The rover shall operate within a 0.5 km radius of the C2 station during the Astrobiology Mission.	§1.b.ii	Astro	mandator
REQ-MOB-011	Mobility	The rover shall operate in full daylight, dust, light rain and ambient temperatures up to 37.8 °C without loss of mobility.	§3.d.iii	All	mandator
REQ-MOB-001	Mobility	The rover shall traverse the Delivery course over soft sand, gravel, rough stony areas, rock and boulder fields, vertical drops and steep loosely consolidated slopes, up to 1 km from the C2 station.	§1.c.ii, §3.d.ii; RQ13	Del	scoring-driven
REQ-MOB-002	Mobility	The rover shall traverse up to 100 m of relatively flat terrain to the lander and complete the Equipment Servicing tasks within the 30-minute mission.	§1.d.i	Equip	scoring-driven
REQ-MOB-004	Mobility	The rover shall autonomously follow a walking astronaut and stop within 3.0 m of the astronaut when the astronaut stops.	§1.e.vi	Auto	scoring-driven
REQ-MOB-005	Mobility	The rover shall autonomously drive to a provided GNSS location and stop within 3.0 m of it.	§1.e.v	Auto	scoring-driven
REQ-MOB-006	Mobility	The rover shall autonomously navigate to a target in hilly natural terrain and stop within 1.0 m of the target location, including an approach up a hill where not all directions are navigable.	§1.e.xv, §1.e.xvii	Auto	scoring-driven

REQ-MOB-007	Mobility	The rover shall climb, descend and traverse slopes and obstacles representative of MDRS natural terrain without loss of control or tip-over.	§1.c.ii; DQ1; AUQ17	Del, Auto	scoring-driven
REQ-MOB-008	Mobility	The rover shall maintain traction and steering control in soft sand and loosely consolidated soil.	§1.c.ii; AQ6	Del, Astro	scoring-driven
REQ-MOB-009	Mobility	The rover shall carry objects up to 5 kg and the arm/cache payload over rough terrain without dropping them or exceeding the 50 kg configuration mass.	§1.c.iii, §1.d.ii derived	Del, Equip	scoring-driven
REQ-MOB-010	Mobility	The rover shall be teleoperable from the C2 station to the Autonomy mission start location on flat, easy terrain.	§1.e.xii	Auto	scoring-driven
REQ-ARM-012	Arm and end effector	During manual alignment, the rover shall not touch any part of the keyboard or the display; it may touch only the surrounding metal areas of the lander.	§1.d.ii; EQ2	Equip	mandator
REQ-ARM-018	Arm and end effector	The arm and its stowed payload shall fit within the 1.2 m cube and the 50 kg mission configuration.	§2.a.ii–iii	All	mandator
REQ-ARM-001	Arm and end effector	The arm/end effector shall reach and manipulate equipment from ground level to 1.5 m above ground at the lander.	§1.d.i	Equip	scoring-driven
REQ-ARM-002	Arm and end effector	The arm shall lift, carry and place objects up to 5 kg with a graspable feature up to 7.5 cm in diameter, including objects up to 40 × 40 × 40 cm.	§1.c.iii	Del, Equip	scoring-driven
REQ-ARM-003	Arm and end effector	The end effector shall pick up a sample tube (test-tube sized), insert it into a larger cache container, close the cache lid and turn the lock.	§1.d.ii	Equip	scoring-driven
REQ-ARM-004	Arm and end effector	The arm shall pick up a cache container with a handle at least 10 cm long and no more than 5 cm in diameter, mass <5 kg, and transport it to the lander.	§1.d.ii	Equip	scoring-driven
REQ-ARM-005	Arm and end effector	The arm shall open a lander drawer, insert the cache into a tight-fitting space in the drawer, and close the drawer.	§1.d.ii	Equip	scoring-driven
REQ-ARM-006	Arm and end effector	The arm shall undo a latch on a hinged panel and open the panel.	§1.d.ii	Equip	scoring-driven
REQ-ARM-007	Arm and end effector	The arm shall remove a key from a panel-mounted tubular pin tumbler lock, insert it into a different lock and turn the key.	§1.d.ii	Equip	scoring-driven
REQ-ARM-008	Arm and end effector	The end effector shall push a 1.5-inch-diameter cam-lock hose fitting onto the lander fitting, with the internal sealing gasket cut to half its original thickness, and achieve the connected state.	§1.d.ii	Equip	scoring-driven
REQ-ARM-009	Arm and end effector	The arm shall turn a ¼-turn lever valve handle through its full 90° travel.	§1.d.ii	Equip	scoring-driven
REQ-ARM-010	Arm and end effector	The arm shall push buttons, flip switches and turn knobs on the lander.	§1.d.ii; EQ1	Equip	scoring-driven
REQ-ARM-011	Arm and end effector	The arm shall autonomously type a 3–6 letter launch key on a vertically mounted Redragon K552 keyboard, using the 2 × 2 cm ArUco 4x4_50 corner tags, with no human contact during autonomous mode.	§1.d.ii; EQ3	Equip	scoring-driven
REQ-ARM-013	Arm and end effector	The arm shall autonomously pick up the specified rock pick hammer from the ground and hold it for transport.	§1.e.viii; AUQ8	Auto	scoring-driven
REQ-ARM-014	Arm and end effector	The arm shall, on command, autonomously place the hammer on the ground or drop it without damaging it.	§1.e.x	Auto	scoring-driven
REQ-ARM-015	Arm and end effector	The sample acquisition system shall collect at least 5 g of material from a depth of 10 cm or deeper and transfer it into the cache without contamination.	§1.b.vi derived	Astro	scoring-driven
REQ-ARM-016	Arm and end effector	The sample acquisition system shall minimise cross-contamination between sites and provide a justified cleaning or replacement method.	AQ7 derived	Astro	scoring-driven
REQ-ARM-017	Arm and end effector	The arm shall handle both the team's sealed sample cache and any judge-provided lander cache with the specified handle.	§1.b.vii, §1.d.ii derived	Astro, Equip	scoring-driven
REQ-ARM-019	Arm and end effector	The perception/arm system shall visually servo on 2 × 2 cm and 1 × 1 cm ArUco 4x4_50 tags on the keyboard, key and keyhole.	§1.d.ii derived	Equip	scoring-driven
REQ-ARM-020	Arm and end effector	The arm shall use force/compliance control or a mechanical compliance device to avoid damaging the delicate screen, keyboard, connectors or lander fixtures.	EQ2, EQ1 derived	Equip	advisory

REQ-ARM-021	Arm and end effector	The arm/rover shall avoid contacting the astronaut and the AR tag posts during Autonomy.	AUQ1 derived	Auto	advisory
REQ-SCI-001	Astrobiology payload and sample	The rover shall carry at least two life-detection capabilities, at least one of which is a bioassay that detects a biomolecule known to be formed only by biological processes or in the presence of biology.	§1.b.v; AQ9	Astro	mandator
REQ-SCI-003	Astrobiology payload and sample	All sample analysis shall be performed onboard and on-site; no sample shall be returned to the team for laboratory analysis.	§1.b.iv	Astro	mandator
REQ-SCI-004	Astrobiology payload and sample	The sample acquisition system shall collect a sub-surface sample from ≥10 cm depth and obtain ≥5 g of material from below 10 cm.	§1.b.vi	Astro	mandator
REQ-SCI-005	Astrobiology payload and sample	The rover shall record and report the depth of each sample and distinguish the portion below 10 cm.	§1.b.vi; AQ10	Astro	mandator
REQ-SCI-006	Astrobiology payload and sample	The rover shall evaluate at least two sites and acquire the cache sample from one of them.	§1.b.ii, §1.b.vi	Astro	mandator
REQ-SCI-007	Astrobiology payload and sample	The rover shall document each investigated site with a wide-angle panorama showing full context, cardinal directions and a scale reference.	§1.b.iii	Astro	mandator
REQ-SCI-008	Astrobiology payload and sample	The rover shall capture a close-up, well-focused, high-resolution image with a scale reference at each sampling site.	§1.b.iii	Astro	mandator
REQ-SCI-009	Astrobiology payload and sample	The rover shall record GNSS coordinates of each site, including elevation and accuracy range.	§1.b.iii; §3.d.v	Astro	mandator
REQ-SCI-010	Astrobiology payload and sample	The cache shall seal the sample so it does not spill when laid on its side or upside down, and shall be reasonably openable by the judges.	§1.b.vii	Astro	mandator
REQ-SCI-011	Astrobiology payload and sample	The cache shall be removable from the rover and handed to the judges within 5 minutes after roving time ends.	§1.b.vii	Astro	mandator
REQ-SCI-012	Astrobiology payload and sample	All chemicals, water and reaction products shall be contained on the rover with no spill to the ground.	§1.b.ix	Astro	mandator
REQ-SCI-013	Astrobiology payload and sample	Any hazardous chemical use shall be pre-approved with a plan covering transportation, usage, disposal, safety precautions and accident response.	§1.b.ix; AQ4	Astro	mandator
REQ-SCI-014	Astrobiology payload and sample	The rover shall complete the roving/analysis phase within the 20–30 minute window and be designed for the 20-minute worst case.	§1.b.ii	Astro	mandator
REQ-SCI-015	Astrobiology payload and sample	The team shall submit a Science Plan by 2027-05-21 that expands the SAR science plan and explains the background science.	§1.b.i; AQ1	Astro (deliverable)	mandator
REQ-SCI-016	Astrobiology payload and sample	The rover shall collect samples from in-situ natural soil/rock with no pre-placed trays; sample sites are not individually marked.	AQ5	Astro	mandator
REQ-SCI-020	Astrobiology payload and sample	The science payload shall fit within the 50 kg mission configuration, the 1.2 m stowed envelope and the rover power/thermal budget.	§2.a.ii–iii, §3.d.iii derived	Astro	mandator
REQ-SCI-002	Astrobiology payload and sample	The two life-detection capabilities shall be sufficiently independent that the failure or contamination of one does not disable the other.	§1.b.v derived	Astro	scoring-driven
REQ-SCI-017	Astrobiology payload and sample	The sampling system shall handle soil/rock ranging from reasonably easy to very hard and difficult to excavate, with no caliche surface.	AQ6	Astro	scoring-driven
REQ-SCI-018	Astrobiology payload and sample	The sampling system shall minimise cross-contamination between sites and document/justify its cleaning or replacement method.	AQ7	Astro	scoring-driven
REQ-SCI-019	Astrobiology payload and sample	The science data pipeline shall produce results that support the post-mission debrief: test results, depth/contamination method, site reasoning, life/habitability interpretation and Mars applicability.	§1.b.x derived	Astro	scoring-driven

REQ-PWR-001	Power and thermal	The rover shall supply all power onboard from batteries; no tethered power shall run to the C2 station.	§2.a.i	All	mandator
REQ-PWR-003	Power and thermal	A battery swap performed by human hand in the field shall be treated as an intervention; a swap performed by the rover/drone itself shall not.	§3.e.vii; RQ14	All	mandator
REQ-PWR-004	Power and thermal	The rover's E-stop shall remove all power draw from all batteries, including compute, communications and payloads, in a single action.	§2.a.v	All	mandator
REQ-PWR-005	Power and thermal	The rover shall operate in ambient temperatures up to 37.8 °C, dust and light rain, and in full daylight, without loss of function.	§3.d.iii	All	mandator
REQ-PWR-006	Power and thermal	The C2 equipment shall operate from the provided 120 V, 60 Hz power, with appropriate power supplies/adapters.	§3.d.i	All	mandator
REQ-PWR-002	Power and thermal	The rover shall complete the longest mission (Delivery, up to 60 minutes) plus setup and any required margin on one battery configuration, or use a permitted battery swap.	§1.c.i, §1.a.ii derived	Del (worst case), All	scoring-driven
REQ-PWR-007	Power and thermal	The rover power system shall allocate and monitor power for mobility, arm, science, compute, communications, LED and any deployables within the 50 kg configuration.	§2.a.iii, §1.e.ii derived	All	scoring-driven
REQ-PWR-008	Power and thermal	The rover shall protect batteries from over-discharge and thermal runaway and provide a documented response to a battery fire.	§2.a.v derived	All	advisory
REQ-PWR-009	Power and thermal	Thermal design shall account for MDRS altitude (~1,309 m) and reduced air density for fan cooling.	S-URC-LOGISTICS derived	All	advisory
REQ-COM-001	Communications	The rover and drone shall be operated by wireless communications with no time delay; no tethered communications shall run to C2.	§2.c.i, §2.a.i	All	mandator
REQ-COM-002	Communications	The system shall not depend on normal internet, cellular service or satellite internet; satellite relay is prohibited.	§2.c.i; RQ5	All	mandator
REQ-COM-003	Communications	The system shall provide control, telemetry and video sufficient to operate without direct visual contact with the rover or site.	§2.c.i, §3.d.i	All	mandator
REQ-COM-004	Communications	The system shall maintain a link for Astrobiology and Equipment with line of sight to C2, and shall operate on the Delivery course where line of sight is not guaranteed for more than 50 % of the course and later tasks are intentionally out of LOS.	§3.d.ii, §1.c.ii	Astro, Equip, Del	mandator
REQ-COM-007	Communications	The team shall comply with FCC rules, submit communication device and licence details to the URC Director by 2027-04-30, and notify changes immediately.	§2.c.ii	All	mandator
REQ-COM-008	Communications	The system shall support omnidirectional and/or directional antennas; any steering shall use electronic control from C2, with no camera on the antenna and no reliance on watching/tracking the rover.	§2.c.iii	All	mandator
REQ-COM-009	Communications	The base-station antenna shall be ≤3 m high, its base within 5 m of the C2 station, guy anchors within 10 m; at MDRS structures the judge may allow up to 20 m; the team shall bring ≥25 m of communications cable.	§2.c.iv; DQ2	All	mandator
REQ-COM-010	Communications	The 900 MHz system shall not use more than 8 MHz bandwidth and shall be operable exclusively in each of 902–910, 911–919 and 920–928 MHz, shifting to the sub-band assigned for each mission.	§2.c.v	All	mandator
REQ-COM-011	Communications	All 900 MHz links, C2-to-rover, C2-to-drone, C2-to-astronaut and any LAN between devices at any location/vehicle, shall comply with the same sub-band and bandwidth restriction.	§2.c.v	All, Auto, Del	mandator
REQ-COM-012	Communications	The 2.4 GHz system shall comply with FCC rules and tolerate interference from other teams; frequency hopping, automatic channel selection/switching are encouraged.	§2.c.vi; RQ6–RQ8	All	mandator
REQ-COM-014	Communications	The team shall power down communications equipment at the event site when not competing.	§2.c.i	All	mandator
REQ-COM-005	Communications	The system shall support Autonomy route finding when the second target is behind a hill and out of radio line of sight with C2.	§1.e.xvi	Auto	scoring-driven
REQ-COM-006	Communications	The system shall operate at ranges up to 1 km from C2 for Delivery and Autonomy, with total path possibly exceeding 1 km; Astro/Equipment are significantly closer.	RQ13; §1.c.ii	Del, Auto	scoring-driven
REQ-COM-013	Communications	The system shall support local repeaters and drone relays; any repeater carried or deployed counts toward the mass budget.	RQ5, AUQ19, RQ15	Del, Auto	scoring-driven

REQ-COM-015	Communications	If an astronaut link is used, it shall provide two-way audio and camera/microphone monitoring in C2; the astronaut may carry electronics but may not deploy repeaters.	§1.e.iv; AUQ5, AUQ26	Auto	scoring-driven
REQ-CMP-001	Compute and perception	The rover shall determine position on WGS 84 and accept/produce latitude/longitude in decimal degrees, degrees decimal minutes, or degrees/minutes/seconds.	§3.d.v	All	mandator
REQ-CMP-013	Compute and perception	The rover shall provide a rear-mounted LED indicator, visible in bright daylight, with red = autonomous, blue = teleoperation and flashing green = successful arrival at a target.	§1.e.ii; AUQ3	Auto	mandator
REQ-CMP-002	Compute and perception	The rover shall use differential/RTK GNSS with its own base station (no CORS or internet) with accuracy sufficient to stop within 1.0 m of a route target and 3.0 m of an astronaut.	RQ11; §1.e.xvii, §1.e.v	Auto, Del	scoring-driven
REQ-CMP-003	Compute and perception	The perception system shall detect 2 × 2 cm ArUco 4x4_50 tags at the keyboard corners and 1 × 1 cm tags on the key and by the keyhole.	§1.d.ii	Equip	scoring-driven
REQ-CMP-004	Compute and perception	The perception system shall detect 20 × 20 cm AR post faces with 2.5 cm cells, 0.5–1.5 m above ground, visible from any of 3 sides, at a range sufficient to navigate to within 1 m.	§1.e.xii, §1.e.xvii	Auto	scoring-driven
REQ-CMP-005	Compute and perception	The perception system shall detect the rock pick hammer on the ground in the rover's field of view while the astronaut may be at least 20 m away.	§1.e.viii; AUQ7–AUQ8	Auto	scoring-driven
REQ-CMP-006	Compute and perception	The perception system shall recognize a sign carried by the astronaut or placed in the field, whose content may be any team-chosen image, AR tag, words or picture.	§1.e.vi, §1.c.iii	Auto, Del	scoring-driven
REQ-CMP-007	Compute and perception	The perception system shall recognize arm/leg/torso beckoning gestures; finger and facial gestures are prohibited.	§1.e.vi; AUQ11	Auto	scoring-driven
REQ-CMP-008	Compute and perception	The audio system shall recognize a spoken word or phrase such as "follow" from the astronaut, either directly or via radio.	§1.e.vi; AUQ27	Auto	scoring-driven
REQ-CMP-009	Compute and perception	The camera system shall capture a wide-angle panorama with a cardinal-direction reference and a scale reference, and a high-resolution close-up with a scale reference.	§1.b.iii	Astro	scoring-driven
REQ-CMP-010	Compute and perception	The camera system shall provide imagery sufficient to select sample sites, read signs and identify objects without direct operator view of the site.	§1.b.ii, §1.c.iii–iv	Astro, Del	scoring-driven
REQ-CMP-011	Compute and perception	Autonomy computation shall run onboard or on a C2 base-station computer, with no human assistance while the rover is driving autonomously.	AUQ2; §1.e.iii	Auto	scoring-driven
REQ-CMP-012	Compute and perception	The system shall ingest satellite imagery and, if provided, a ≤1 m resolution DEM (likely GeoTIFF) for route planning, without depending on the DEM being available.	§1.e.xiii; AUQ20	Auto	advisory
REQ-CMP-014	Compute and perception	The C2 display shall show rover state, camera views, telemetry, maps and autonomy status, and may indicate target arrival.	§1.e.ii; AUQ18 derived	Auto	advisory
REQ-CMP-015	Compute and perception	The perception/software system shall detect obstacles (rocks, posts, astronaut) and avoid collisions during autonomous and teleoperated driving.	AUQ1; §1.c.ii derived	All	advisory
REQ-SFT-008	Software and autonomy	The rover shall clearly indicate and announce entry to and exit from autonomous mode to the C2 judge.	§1.d.ii, §1.e.ii	Auto, Equip	mandator
REQ-SFT-009	Software and autonomy	The rover shall be hands-off while driving autonomously; operators may program only while stopped or teleoperating, and may not drive while stopped.	§1.e.iii; AUQ21	Auto	mandator
REQ-SFT-013	Software and autonomy	During route finding, teleoperation shall not enter areas not already autonomously explored by the rover.	§1.e.xviii	Auto	mandator
REQ-SFT-001	Software and autonomy	The rover shall autonomously navigate from the start location to two target locations in hilly terrain, in any order, stop within 1.0 m of each and indicate arrival.	§1.e.xii, §1.e.xvii	Auto	scoring-driven
REQ-SFT-002	Software and autonomy	The rover shall plan a traversable route up a hill where not all approach directions are navigable, using onboard perception and/or map data.	§1.e.xv; AUQ17	Auto	scoring-driven
REQ-SFT-003	Software and autonomy	The rover shall continue autonomous navigation and decision-making when out of radio line of sight, and indicate arrival without a live link.	§1.e.xvi; AUQ18	Auto	scoring-driven
REQ-SFT-004	Software and autonomy	The rover shall autonomously follow a walking astronaut and stop within 3.0 m when the astronaut stops.	§1.e.vi	Auto	scoring-driven

REQ-SFT-005	Software and autonomy	The rover shall hold position on Stay! while the astronaut moves >20 m away, until a new command is given.	§1.e.vii; AUQ7	Auto	scoring-driven
REQ-SFT-006	Software and autonomy	The rover shall autonomously detect, grasp and fetch the rock pick hammer on command, and later place or drop it on command.	§1.e.viii–x	Auto	scoring-driven
REQ-SFT-007	Software and autonomy	The rover shall accept Follow! commands from at least one of: astronaut-carried device, visual sign, audio word/phrase, or beckoning gesture, and support a fallback mode.	§1.e.vi; AUQ6–AUQ7	Auto	scoring-driven
REQ-SFT-010	Software and autonomy	The rover shall autonomously stop or return to the astronaut on its own for no penalty, and re-accept the command.	§1.e.xi	Auto	scoring-driven
REQ-SFT-011	Software and autonomy	Operator-triggered stop/return-to-astronaut shall cost 20 % on that task; teleoperation back to a visited location shall cost 50 %; further exits are capped at 50 % per task.	§1.e.xi, §1.e.xviii	Auto	scoring-driven
REQ-SFT-012	Software and autonomy	After an abort, the rover shall autonomously return the way it came, or to the mapping start point, or (first attempt) to the start area in front of C2, at the prescribed penalty.	§1.e.xviii; AUQ4	Auto	scoring-driven
REQ-SFT-014	Software and autonomy	The rover shall autonomously type a 3–6 letter launch key on the vertical keyboard, using backspace/delete to correct errors, with multiple attempts and no abort penalty.	§1.d.ii	Equip	scoring-driven
REQ-SFT-015	Software and autonomy	The typing routine shall handle no key auto-repeat and a judge-reset display without assuming prior typed characters were cleared by the rover.	EQ3	Equip	scoring-driven
REQ-SFT-016	Software and autonomy	The C2 software shall allow entry of GNSS points, waypoints and stay-out zones while the rover is stopped, and changes to controls/algorithms.	§1.e.iii, §1.e.xii	Auto	scoring-driven
REQ-SFT-019	Software and autonomy	Teleoperation shall be available for Astrobiology, Delivery and Equipment Servicing, and for driving to the Autonomy start; it shall not be used where the rules require autonomy.	§1.b–d, §1.e.xii, §1.e.iv–xviii	All	scoring-driven
REQ-SFT-017	Software and autonomy	The rover shall implement a failsafe that stops or safely holds the rover on loss of link, low battery, or critical sensor fault.	§2.a.v, §3.e.ii derived	All	advisory
REQ-SFT-018	Software and autonomy	The rover shall log autonomy mode, task starts/completions, aborts and interventions for judge verification and penalty accounting.	§1.e.xi, §1.d.ii derived	All	advisory
REQ-OPS-001	Operations, C2 and schedule	The team shall register by 2026-10-30, submit the PDR by 2026-12-02, the SAR by 2027-03-05, communications documentation by 2027-04-30, the Science Plan by 2027-05-21 and the Rover Expense Report by 2027-05-26.	§3.a.i–iii, §2.c.ii, §1.b.i, §3.c.ii	SAR, all	mandator
REQ-OPS-002	Operations, C2 and schedule	The SAR shall consist of written and video components demonstrating the rover's capability and customisation for each mission, overall system design and progress to date, and the team shall score in the top 36 to be invited to the field.	§3.a.iii	SAR	mandator
REQ-OPS-003	Operations, C2 and schedule	The team shall operate from the designated C2 station with no more than 6 students at any time; advisors are not permitted in C2; all rover operators remain in C2 and nobody follows the rover for feedback.	§3.b.ii, §3.b.iv; S-URC-LOGISTICS	All	mandator
REQ-OPS-004	Operations, C2 and schedule	The team shall comply with Finals attendance limits: 24 registered students, 2 advisors, 14 individuals per day.	S-URC-LOGISTICS	All	mandator
REQ-OPS-005	Operations, C2 and schedule	The team shall designate runners for interventions; a request for intervention may come only from the operators, not spectating team members; a team member who leaves C2 to become a runner cannot return to operate or analyse data for that mission; operators may become runners and spectators may act as runners and carry tools.	§3.e.i, §3.e.iii	All	mandator
REQ-OPS-006	Operations, C2 and schedule	All radio communication during an intervention shall be in English; the team should bring its own radios and use them only during an intervention.	§3.e.vi	All	mandator
REQ-OPS-007	Operations, C2 and schedule	The team shall not test at MDRS before, during or after URC; testing shall be at approved off-road sites (e.g. Factory Butte) or in town, not on undesignated BLM land.	§3.d.iv	All	mandator
REQ-OPS-008	Operations, C2 and schedule	Science briefing participants shall be limited to team members who stayed in C2 during the science mission and have not seen the site directly.	AQ8	Astro	mandator
REQ-OPS-009	Operations, C2 and schedule	The team shall remove the cache and hand it to the judges within 5 minutes after roving, and shall conduct a debrief even if the rover failed the field tasks.	§1.b.vii, §1.b.x	Astro	mandator
REQ-OPS-010	Operations, C2 and schedule	If a drone is used, a pilot shall operate it from C2 and a designated spotter shall maintain visual line of sight in the field; the astronaut may not fly the drone.	§2.b.ii, §2.b.xii; AUQ25	Del, Auto	mandator
REQ-OPS-012	Operations, C2 and schedule	The team shall manage mass/configuration across the four missions so that each configuration is ≤50 kg, the total fielded set is ≤70 kg, and the	§1.a.iv, §2.a.iii; EQ4 derived	All	mandator

		Autonomy → Equipment transition fits 10 minutes including any re-weigh.			
REQ-OPS-011	Operations, C2 and schedule	The team shall plan for a schedule released about one week before the field event and a per-team task order that differs between teams.	CQ4	All	advisory
REQ-SAF-001	Safety and E-stop	The rover shall have a red push-button emergency stop, readily visible and accessible on the exterior.	§2.a.v	All	mandator
REQ-SAF-002	Safety and E-stop	A single E-stop action shall immediately stop the rover's movement and cease all power draw from every battery, switching off all rover systems.	§2.a.v	All	mandator
REQ-SAF-003	Safety and E-stop	The E-stop shall not be a knob or rotating switch.	§2.a.v	All	mandator
REQ-SAF-004	Safety and E-stop	Any mobile deployable (e.g. mini-rover) shall have its own E-stop; a fuse suffices for devices that cannot move on their own; the drone is exempt.	§2.a.vi	All	mandator
REQ-SAF-005	Safety and E-stop	The EVA suit shall cover to the neck, wrists and ankles with close-toed shoes; if a helmet is included it shall be easy for the astronaut to remove.	§1.e.iv; AUQ12	Auto	mandator
REQ-SAF-006	Safety and E-stop	The team shall stop drone operations when a judge directs and shall allow the spotter to take control in an unsafe situation.	§2.b.xii	Del, Auto	mandator
REQ-SAF-007	Safety and E-stop	The rover shall provide a documented battery-fire response and thermal protection consistent with the E-stop's stated purpose.	§2.a.v derived	All	advisory
REQ-SAF-008	Safety and E-stop	The rover shall avoid contacting the astronaut and AR posts, and shall stop safely if a person enters its path.	AUQ1 derived	Auto	advisory
REQ-DRN-001	Drone	The drone shall be a rotary-wing vehicle able to hover in a fixed location; no lighter-than-air or fixed-wing vehicle.	§2.b.iii	Del, Auto	mandator
REQ-DRN-002	Drone	The drone shall weigh ≤5 kg including the dummy mass; a drone over 5 kg shall not be used.	§2.b.vii	Del, Auto	mandator
REQ-DRN-003	Drone	The drone shall carry an inert dummy mass equal to the battery mass at all times, not a spare battery, easily removable for weighing; high wind does not waive it.	§2.b.v	Del, Auto	mandator
REQ-DRN-004	Drone	The drone shall start at the designated landing pad in the field; after initial launch it may take off/land as many times as judges allow.	§2.b.iv	Del, Auto	mandator
REQ-DRN-005	Drone	The drone shall be piloted from C2 with a field spotter maintaining visual line of sight; the pilot will not have direct view.	§2.b.ii, §2.b.xii	Del, Auto	mandator
REQ-DRN-006	Drone	The drone shall not be used in Astrobiology or Equipment Servicing.	RQ9; §1.a intro	Astro, Equip	mandator
REQ-DRN-008	Drone	In Autonomy, the drone shall land at the designated landing pad before the rover proceeds from the mission start to the target locations.	§1.e.xiv	Auto	mandator
REQ-DRN-009	Drone	The drone shall meet FAA requirements: Remote ID, registration if >250 g and marked registration, TRUST proof carried, <400 ft / 120 m AGL, class G, no National Park flight, give way to manned aircraft.	§2.b.vii–xi	Del, Auto	mandator
REQ-DRN-007	Drone	The drone may deploy anything except a mini-rover and may act as a radio relay or carry objects within its limits.	RQ15; §2.b.i	Del, Auto	scoring-driven
REQ-DRN-010	Drone	The drone should be able to operate in 48 km/h (30 mph) wind and have a top speed of at least 64 km/h (40 mph).	§2.b.vi	Del, Auto	advisory
REQ-DRN-011	Drone	The mission plan shall not depend on the drone: full points are possible without one.	§2.b.i	Del, Auto	advisory

Three types, two failure modes

TYPE	ROWS	DEFINITION	WHAT HAPPENS IF IT IS MISSED
Mandatory	79	An explicit rule requirement whose violation is a rule breach, a penalty or an eligibility failure	Points, a mission or the entry is lost, whatever else the rover can do
Scoring-driven	61	A performance level needed to earn points, not itself a breach	Points are lost in proportion, and the mission may still be completed
Advisory	15	Encouraged practice, or a good design habit, with no direct points or penalty	Nothing immediate, and a foreseeable failure later

The distinction decides where the margin goes. A mandatory row with no margin is a risk of losing an event; a scoring-driven row with no margin is a risk of losing points inside an event `S-URC27-RULES`. The two are budgeted differently.

What each check method means

PRIMARY CHECK	ROWS	WHAT IT DEMONSTRATES
Test	78	A measured result on the rover, a mock-up or an analog site
Inspection	36	A physical or documentary attribute, checked at weigh-in or at a review
Demonstration	28	A behaviour performed once, in sequence, under time
Analysis	13	A calculation or a model, for requirements that cannot be measured cheaply

The count is by the first method named, so a row that reads test and inspection is counted as a test. Many mandatory rows are inspections at a weigh-in or a document review, which is why a team can pass every functional test and still fail the table `S-URC27-RULES`.

Where the coverage is thin

Three limits are worth stating before the table is used as a checklist.

- **102 rows trace to a clause and not to a Q&A ruling.** They are only as good as the clause reading, and the 2027 rulebook has clauses whose wording changed in the rewrite, so a requirement inherited from an earlier year has to be re-read rather than carried over.
- **The Q&A can move.** The 2027 page was revised twice after publication, and 51 rows depend on it `S-URC27-QA`. A ruling that changes rewrites the trace of every row that cites it.
- **The scoring rows carry an unpriced gap.** The 2027 task weights for three missions are not published, so a scoring-driven row says what capability is needed and cannot say what it is worth `S-URC27-RULES`.

Rows a first build is most likely to miss

REQUIREMENT	THE TRAP	WHAT IT COSTS
REQ-SYS-001	The cube includes wheels, antenna and every protrusion, in any orientation, with no disassembly S-URC27-RULES §2.a.ii	40 % of the points earned
REQ-SYS-003	The 70 kg aggregate is strict, and deployed items count inside it	Parts removed at the field, or an entry that cannot weigh in
REQ-SYS-008	Fifteen minutes covers all communications, the base station and the weigh-in	The mission starts late or not at all
REQ-PWR-004	One action must stop movement and remove power from every battery, including compute and radios	The E-stop check fails, and a runaway or a battery fire is unmanaged
REQ-COM-009	The mast is under 3 m, the base within 5 m of the command post, and the cable run is at least 25 m	No link through a metal wall, with no time to improvise
REQ-ARM-012	Touching any part of the keyboard in manual control is penalised	The typing task is not attempted legally
REQ-CMP-003	The tags on the key and keyhole are 1 cm squares at task range	The fine-alignment tasks become blind
REQ-SCI-001	Two life detections, at least one a bioassay, are mandatory S-URC27-RULES §1.b.v	The Astrobiology payload is non-compliant before it is tested



The typing task is specified down to the keycap. Key pitch is 19.05 mm, actuation travel about 2 mm and actuation force about 0.45 N, and the four 2 × 2 cm corner tags are the only legal alignment reference. Those numbers set the arm's repeatability and the compliance of the end effector.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*



A 1 cm ArUco tag at task range is the hardest detection target in the competition, and it is the one that has to survive sunlight on a metal panel. The lock task chains detection, grasp, keyway alignment and torque in a single sequence.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

What the rules do not publish

Every row below is a value a design or a test would use, that no 2027 document gives. The right-hand column is the action that produces a usable number `S-URC27-RULES` `S-URC27-QA`.

MISSING VALUE	WHY IT BLOCKS A CHECK	HOW TO SETTLE IT
Astrobiology, Delivery and Equipment task weights	A scoring-driven row cannot be prioritised against another	Watch the science plan and course pages, and ask at the field meeting
Delivery stage durations and Stage-1 minimum	The time budget and the gate risk are unknown	Same; 2026 values are a labelled proxy only
Slope, step and drop limits for the course	The mobility envelope cannot be tied to a published worst case	Measure an analog obstacle set and document it
Sample tube, key, lock and cache dimensions	The interfaces cannot be checked against the real hardware	Build judge-like mock-ups and measure the fits
Forces and travels for the hose, valve, latch, buttons and knobs	The end-effector force and compliance cannot be sized	Measure on the specified fittings, with margin
Hammer mass and handle geometry	The Fetch! grasp cannot be designed	Measure the linked item and build an equivalent
Image resolution and sign-reading distance	The camera and lens cannot be chosen	Test with judge-like signs at plausible distances
Penalty above 70 kg, or for a 900 MHz breach	The risk model is incomplete	Ask in writing; treat both as compliance risks
Whether 3 m and 5 m antenna limits cover a GNSS base antenna	The base geometry may be constrained more than expected	Ask in writing; design to 3 m meanwhile
Whether one instrument can satisfy both life-detection capabilities	The payload architecture changes if it cannot	Ask in writing before the payload is fixed
Whether the Equipment cache is the team's own cache	Two interfaces have to be designed instead of one	Ask in writing; design for both

The check plan in twelve steps

1. **Run the weigh-in as a test.** Build a 1.2 m gauge, weigh every mission configuration and the 70 kg aggregate, and time the procedure inside a 15-minute setup.
2. **Time the clock.** Fifteen minutes of setup, ten of teardown, ten of retrieval and the ten-minute transition between Autonomy and Equipment, with the real command post and the real cable.
3. **Test the stop as one action.** With compute and radios running, confirm that no subsystem keeps power from any battery.
4. **Test the link at 1 km and inside a radio shadow**, in each of the three 900 MHz sub-bands, and confirm that the rover keeps driving and signals arrival by LED.
5. **Measure the stop distances.** Run the 1 m route stops and the 3 m astronaut stops on natural terrain and report the distribution, not the best pass.
6. **Build judge-like mock-ups** for every servicing task and record the forces and travels the rules omit.
7. **Test depth and mass.** Take at least 5 g from below 10 cm in soil that ranges from easy to very hard, and verify the cache seal when inverted.
8. **Run the life detection on blind samples and controls.** A bench result is not a field capability.
9. **Test typing** with 3, 4, 5 and 6 letter keys, with backspace corrections, no key repeat and a judge reset.
10. **Test the suit** for heat, camera and microphone link, two-way audio and helmet removal, with the electronics that count in the mass and cost budgets.
11. **Keep every field test away from MDRS.** Use a permitted off-road site and plan the campaign before travel.
12. **Collect the evidence as the tests run.** The review is a mission worth 100 points and a gate on the other 400.

Further reading

- Rules, limits and the field holds the clauses these rows trace to.
- Q&A rulings holds the answers behind 51 rows.
- Missions and scoring explains what each mission rewards.
- 110 engineering problems turns the largest subset of these rows into build tasks with acceptance criteria.
- Deadlines carries the document rows in the operations domain.

title: Review, judging and awards description: The 2027 review chain: the design review, the 100-point acceptance review, the science and chemical safety plans and the awards. part: The competition chapter: "2.5" toc_depth: 3 tags: - review - judging - awards - documentation

Review, judging and awards

The 2027 season can be lost at a desk. The System Acceptance Review carries 100 of the 500 points and decides which 36 teams are invited to the field at all, and the Preliminary Design Review before it shapes the judges' view of whether a team will be ready. Between them sit the science plan, the chemical safety plan and the Rover Expense Report, each with its own deadline and its own consequences. This page sets out what each milestone asks for, what the published 2026 rubrics value, and how the cut has actually fallen.

100

REVIEW POINTS

20 % of the total and the gate to the field

2027

`S-URC27-RULES`

48 / 40 / 12

VIDEO, REPORT, JUDGES

The 2026 split; the 2027 rubric is unpublished

2026

`S-URC26-SAR`

36

TEAMS INVITED BY RULE

37 to 38 in practice, 2023 to 2026

2023–2026

`S-URC27-RULES`

6.25 s

OF VIDEO PER POINT ON OFFER

48 points inside a 5:00 limit

2026

`S-URC26-SAR`

Key takeaways

- The review is 20 % of the maximum score and the only route to the 400 field points. A team outside the top 36 never deploys a rover, however good the rover is.
- About one team in three registered is admitted. From 2023 to 2026 the field was 37, 38, 38 and 38 places from 104, 102, 114 and 116 registrants.
- The 2026 rubric put 48 of the 100 review points in a five-minute video, 40 in a seven-page report and 12 in the judges' assessment. The 2027 rubric is not published, so the 2026 document is a floor rather than a ceiling.
- The judges score what is shown. A capability that is claimed in a voice-over but not demonstrated on camera does not earn points.

The milestone chain

Every date below is fixed in the 2027 rulebook `S-URC27-RULES`. The chemical safety plan has no separate date and is documented alongside the Science Plan.

MILESTONE	DATE	FORM	WHAT IT CONTROLS
Registration	2026-10-30	Team details through the URC site	Eligibility (§3.a.i)
Preliminary Design Review	2026-12-02	Five pages, non-competitive	Readiness assessment (§3.a.ii)
System Acceptance Review	2027-03-05	Report and video, competitive	100 points and the top-36 gate (§3.a.iii)
Communications documentation	2027-04-30	Devices, bands and operator licences	Compliance (§2.c.ii)
Science Plan	2027-05-21	Written plan, about five pages	Astrobiology score (§1.b.i)
Chemical Safety Plan	With the Science Plan	Per-chemical safety record; the 2027 format is still pending	Any wet assay or hazardous chemical (§1.b.ix)
Rover Expense Report	2027-05-26	Work-breakdown format; update allowed 2027-06-02	10 % of total points per day late (§3.c.ii)
Finals	2027-06-02 to 2027-06-05	Four missions at MDRS	4 × 100 points (§3.a.iv)



The field the review admits a team to. The 36 places are allocated months earlier, on the strength of a written package and a five-minute film, before any of the terrain below is driven.

Mars Society, 2017. CC BY-SA 4.0. Original source Full resolution

The Preliminary Design Review

The Preliminary Design Review (PDR) is a readiness review, not a scored mission. The rulebook describes it as an assessment of each team's overall likelihood of being ready in time, judged on its own merits rather than against other teams S-URC27-RULES §3.a.ii. The 2027 requirements page is still pending, so the 2026 document is the only published format S-URC27-PDR-PAGE S-URC26-PDR.

The 2026 format was five pages: pages 1 and 2 for text and graphics, pages 3 to 5 for graphics and captions only, letter paper with margins of at least 1 inch, Calibri 11 pt or similar, single spaced, submitted as a PDF through a form, with a header block naming the team, university, city, country and student lead S-URC26-PDR. Reports were reviewed on a rolling basis before the deadline, and format violations or missing citations carried penalties.

SECTION	WHAT THE 2026 DOCUMENT EXPECTS
Team structure	Leadership, membership, departments represented, relevant experience, recruiting and training, and a separate outreach plan
Team resources	A project budget presented as a table, in US dollars, with income and expense categories and a plan for any shortfall
Project management	A single Gantt chart with named owners, tasks and dates covering September to June, plus documentation, systems engineering and test plans
Preliminary design	System architecture with each subsystem marked as new, modified or reused, and preliminary designs for mobility, communications, arm, C2 and science

The judges' own comments are the most useful part of the 2026 document. Most historical PDR feedback concerned teams that did not follow the instructions. Blank pages, a Gantt chart of milestones rather than tasks, a budget drawn as a pie chart, outreach merged into recruiting, and a plan that only states that a plan exists were all named in the guidance as things to avoid. Past success is explicitly not a substitute for addressing the required elements.

The System Acceptance Review

The rule fixes the character of the review rather than its rubric. It focuses on demonstration of the capabilities and customisation of the rover for each mission, it is judged against other teams, and the 36 highest scores are invited to the field `S-URC27-RULES §3.a.iii`. The 2027 requirements page still says they will be posted in early 2027 `S-URC27-SAR-PAGE`.

The 2026 rubric is the best available model of what the 2027 one will value `S-URC26-SAR`. It splits 100 points across three components.

COMPONENT	POINTS	SUB-CRITERIA
Report	40	System overview and architecture 8 · integration and testing 6 · project management 11 · science plan 15
Video	48	Mobility and terrain 8 · arm and end effector 8 · autonomous navigation 6 · science payload 6 · overall integration 12 · communications, C2 and operations 8
Judges' assessment	12	Readiness to compete 4 · ruggedness for MDRS 4 · novelty and scope of work 4

The report was seven pages with a fixed page map: pages 1 to 4 free, page 5 for the Gantt chart only, page 6 for the Project Budget Table only, and page 7 for the science plan only. Content after page 7 was not considered, and text in figures had to be legible at 250 % zoom `S-URC26-SAR`. The project-management weight is worth noting: the Gantt chart and the budget table together carry 8 of the 40 report points, as much as the system architecture.

The video was a single film of five minutes or less, publicly accessible and searchable. Audio quality was described as imperative even at the expense of image quality, music was strongly discouraged, and every clip had to be captioned with the month and year it was recorded. Footage from before September of the current cycle was treated as weaker evidence, and a clip of a superseded architecture was assumed not to apply to the current rover. Accelerated segments had to display their speed for the whole segment, and the judges warned explicitly against zooming in on an end effector without showing the whole arm integrated on the rover.

Because the weights are published, the five minutes can be budgeted `S-URC26-SAR`. At 48 points over 300 seconds, each point is worth about 6.25 seconds of film.

SECTION	POINTS	TIME GUIDE	WHAT HAS TO BE VISIBLE
Mobility and terrain	8	~50 s	The same rover on at least three different terrain types, with wheels and suspension in frame
Arm and end effector	8	~50 s	The whole arm mounted on the rover while lifting 5 kg, picking a small item, opening a latched panel, pushing switches, connecting a hose and typing
Autonomous navigation	6	~37 s	Rover footage beside the sensor view and the operator display, with a waypoint drive and a marker detection
Science payload	6	~37 s	Sample collection, the sample entering the cache, and the analysis running with data on screen, all on the rover
Overall integration	12	~75 s	One continuous sequence in which arm, science, autonomy and drive cooperate
Communications, C2 and operations	8	~50 s	Operators in the C2, the antenna outside, a link or telemetry display and a timed field run
Captions	Unscored	Remainder	Month and year per clip, and speed labels for any accelerated segment

The time guide is this reference's reading of the rubric, not a URC instruction. The weights are URC's.

The science plan and the chemical safety plan

Two documents carry the science case, and they are scored in different places.

	REVIEW SCIENCE PLAN	MAY SCIENCE PLAN
Length	1 page, within the report	About 5 pages
Due	With the review, 2027-03-05	2027-05-21
Scored in	The review, 15 of its 100 points in 2026	The Astrobiology Mission
Content	What the science is and why it detects life	Background science, Mars knowledge, payload justification and one essay question

The 2026 May plan structure is the only detailed one published: a cover page naming the science sub-team and the status of the chemical safety plan, at most a page of basic Mars knowledge, at most a page of background research, up to two pages on the payload and its justification, and at most a page answering one question from a fixed list `S-URC26-SCIPLAN`. The judges warned that they will ask one of those questions, or a similar one, at the Astrobiology debrief after the mission.

The Chemical Safety Plan is a per-chemical record, not a declaration. The 2026 document required a named chemical safety officer and up to two alternates travelling to the finals, a training acknowledgement, and for each chemical: its name and concentration, the amount needed per reaction, storage and secondary containment, protective equipment, a justification for its use, a disposal route naming a specific licensed facility, spill prevention and incident response, and safety data sheets both linked and carried on site. Mercury-containing chemicals were prohibited outright. The 2027 rulebook keeps the no-spill policy and requires hazardous chemicals to be pre-approved with a plan covering transport, use, disposal and accidents `S-URC27-RULES` `S-URC27-QA` `§1.b.ix` AQ4.

The practical reading is that the chemistry choice is a project-level decision with a paperwork tail measured in weeks. It has to be settled early enough to write both plans and to name a disposal facility that will actually accept the waste.

The expense report and the fielded-equipment cap

Two financial documents are required, and they are not the same document. The Project Budget Table in the PDR and the review covers the whole year's project. The final Rover Expense Report covers only the fielded rover and must show it inside the $26,000 cap `S-URC27-RULES` `§3.c.i` `§3.c.iii`.

RULE	2027 VALUE
Project Budget Table	Required in both the PDR and the review (§3.c.i)
Final Rover Expense Report	Due 2027-05-26, with an updated record allowed on 2027-06-02 (§3.c.ii)
Late penalty	10 % of total points per day (§3.c.ii)
Missing report	Disqualification if not submitted by the end of the event (§3.c.ii)
Cap	$26,000 of cash spent on fielded equipment (§3.c.iii)
Covered	Rover, modules, power sources, communications, drone, astronaut EVA suit, base station and C2 equipment (§3.c.iii)
Not covered	Travel, accommodation, transport, general hardware, non-fielded prototypes, one-for-one spares (§3.c.v)
Valuation	As-bought for used parts, original or equivalent cost for reused parts, purchase cost for rented equipment, full value for donations, shipping and tax included (§3.c.vi to §3.c.ix)

The report has to follow a work-breakdown structure, and URC publishes a template S-URC27-REQPAGE S-URC-EXPENSE-TEMPLATE. The template runs from node 0.0 total rover cost through frame, drive and energy, electronics, command and control, arm, science payload and miscellaneous parts, with a cost rule on every line that records how the value was obtained. Two categories the 2027 cap covers have no node in the template: the drone and the astronaut EVA suit. A team reporting either has to add a top-level line.

The judges' guidance adds two valuation points that catch teams out. Services donated by a supplier must be valued at a quote for the full value of the work, and sponsorship is reported split into financial and in-kind contributions S-URC-BUDGETGANTT S-URC26-SAR. Equipment reused from a previous rover is valued at its original cost or at the price of a new equivalent, not at zero.

Awards

Two awards sit outside the 500-point score, and neither can rescue a weak review.

The INCOSE Foundation Excellence in Systems Engineering Award was created in 2026, is explicitly separate from competition scoring, and judges systems engineering maturity on its own S-URC-INCOSE. In its first year, every finalist report was pre-screened against a 100-point rubric of eight categories: systems thinking and stakeholder focus 15, requirements engineering 15, system architecture and design 15, V&V strategy 12, technical management and planning 10, risk and configuration management 10, systems modelling and analysis 8, and documentation quality 15. Thirteen teams scoring at least 80 became finalists. The on-site round used a second 100-point rubric of ten categories, with 78 points from a 30-minute interview and 22 from competition observation, and the combined score weighted the pre-screen and the final round equally with a threshold of 75. The 2026 winner was Michigan Mars Rover of the University of Michigan. The thirteen 2026 finalists were A.S.T.R.A. (UAH), BYU Mars Rover, Husky Robotics (Washington), KNR Rover Team (Warsaw), Missouri S&T, Michigan Mars Rover, Monash Nova Rover, MR2 (KAIST), Northeastern, Queen's Space Engineering Team, Robotics for Space Exploration (Toronto), Team Mountaineers (WVU) and Team RoSE (Hawaii).

The John Barainca Memorial Award for Outstanding Scientist commemorates a founding member of The Mars Society who helped select the MDRS site and served as Hab manager S-MARS-INMEMORIAM. No criteria document or rubric for the award is published. The 2026 winner was Grant Quartuccio of A.S.T.R.A. at the University of Alabama in Huntsville S-URC26-SCORES-PAGE S-URC26-NEWS. The award is won in the Astrobiology debrief, which is scored on astrobiology knowledge and on what the data mean for the presence or absence of life, and which takes place even if the rover failed its field tasks (§1.b.x).

How the cut has fallen

The official score sheets carry a review score for every finalist, so the lowest admitted score is visible even though rejected teams are not published. That gives an upper bound on the cut line.

YEAR	REGISTRANTS	ADMITTED	LOWEST ADMITTED	MEDIAN ADMITTED	HIGHEST
2023	104	37	85.10	89.31	94.54
2024	102	38	84.04	89.87	97.80
2025	114	38	86.90	90.31	97.27
2026	116	38	75.86	86.77	96.75

The 2023 to 2025 cut sat between 84 and 87, which means a first-time team had to land within a few points of the field's median to be safe S-URC23-SCORES S-URC24-SCORES S-URC25-SCORES . The 2026 low of 75.86 is not comparable: it is the first year of the rebuilt rubric, and a fall of more than ten points from the 2025 low reads more plausibly as a re-based rubric than as a weaker field S-URC26-SAR S-URC26-SCORES .

Two mechanical facts about the cut matter as much as its level. The rule says 36 teams, and the judges expanded the field to 37 in 2023 and 38 in each year from 2024 to 2026 when scores clustered; there is no standby list for withdrawals S-URC23-NEWS S-URC24-NEWS S-URC25-NEWS S-URC26-NEWS . A team should therefore plan against a cut at 36 and treat any expansion as unearned.

What the 2027 documents do not yet say

- The 2027 PDR requirements page says only that the document will be posted in the near future S-URC27-PDR-PAGE .
- The 2027 review requirements page says the requirements will be posted in early 2027, so the 40/48/12 split and the sub-criteria are unconfirmed S-URC27-SAR-PAGE .
- The 2027 Science Plan specification is not published beyond the Q&A note that the May plan expands the review version and runs to about five pages S-URC27-SCIPLAN-PAGE S-URC27-QA AQ1.
- Whether the review report keeps a single-page science plan and a seven-page limit is unknown.
- Whether the review page six must carry a work-breakdown expense table, which was optional in 2026, is unknown.
- Whether a team can be removed at the PDR stage is unknown. The rules describe the PDR as non-competitive and no published page describes a team being cut there.
- The 2027 INCOSE criteria are promised earlier than in 2026 but were not published as of 2026-10-25 S-URC-INCOSE .

What the judges reward

The published guidance is unusually consistent across both desk milestones.

1. **Demonstrated, current, integrated capability.** A capability that is not shown on camera cannot be scored, recent footage outranks old footage, and a superseded design does not support a current one S-URC26-SAR .
2. **Evidence of readiness.** The PDR asks whether a team will be ready, and the review spends 4 of its 12 judges' points on readiness to compete S-URC26-PDR S-URC26-SAR .
3. **Ruggedness for MDRS.** Four of the 12 judges' points, and the advice page returns to dust, heat, shock and field repair repeatedly S-URC-ADVICE .
4. **Novelty and scope of work.** Another 4 of the 12 judges' points S-URC26-SAR .
5. **Systems engineering discipline.** Traceability, architecture, V&V, risk and configuration management sit across most of the INCOSE rubric S-URC-INCOSE .
6. **Following the instructions exactly.** The most-cited PDR feedback cause, and a stated penalty trigger in both reviews S-URC26-PDR S-URC26-SAR .

The failure modes are the mirror image: a marketing film with no evidence, ignoring the page or time limits, a plan that only states a plan exists, missing or unfindable citations, an engineering page written about science or the reverse, a budget table drawn as a chart, and a per-subteam Gantt chart that reads as absent integration.

Further reading

- Rules, limits and the field states the clauses behind each deadline and penalty.
- Missions and scoring shows where the review's 100 points sit against the four missions.
- The 2027 competition places the milestones in the season.
- Deadlines and milestones collects every date in one place.
- Systems engineering and integration covers the process evidence the review looks for.
- The Q&A record carries the rulings on the science plan, the review and the expense report.

THE COMPETITION · CHAPTER 2.6

Rule history, 2007-2027

2027 is the largest single-year rewrite since the four-mission format settled in 2018, and it is also the year the physical problem stopped being new. The mass limit, the no-tether rule, the blocked view from the command post and the desert heat are the same in 2027 as in 2007. This page reads the record by era, separates the changes that still shape a design from the clauses that moved without changing anything, and dates every claim to the rulebook that carries it.

50 kg

MASS LIMIT SINCE 2007

The same figure in every rulebook read

2007–2027

S-URC27-RULES

2019

1.2 M CUBE INTRODUCED

With a 40 % penalty for not fitting

2019

S-URC19-RULES

2017

AUTONOMY BECAME A SCORED MISSION

A fifth of the 500-point total

2017

S-URC17-RULES

$26,000

FIELDED-EQUIPMENT CAP

From $10,000 in 2007 and $24,000 in 2026

2027

S-URC27-RULES

Key takeaways

- **The envelope never moved.** 50 kg, no tethers, off-grid operation, about 100 °F of desert heat, no direct view of the rover and a real-time link are 2007 rules still in force in 2027 S-URC07-RULES-L1 S-URC27-RULES §2.a.
- **2027 changes capability, not tuning.** The science payload doubles and autonomy is rebuilt around a walking astronaut; nothing was relaxed elsewhere to pay for it.
- **Three quality-judged tasks were deleted in 2027:** the stratigraphic profile, the object-recognition highlight and the precise marked-path drive. Each had been added, tuned and then removed.
- **The size box is eight years old and the most under-priced constraint in the rulebook.** It applies at every weigh-in, costs 40 % of the points if missed, and allows folding but not disassembly S-URC19-RULES S-URC27-RULES §2.a.ii §2.a.ii.

The twenty-year shape

ERA	EVENTS	POINTS	WHAT APPEARED
2007	Engineering and Scientific tasks	Not published	The 50 kg limit, the $10,000 cap, a radio link with no time delay, no tethers
2008	Geology, Soil Characterization, Emergency Navigation, Construction	400	Four scored events, no design review, a delivery race scored on speed
2009–2010	Tasks reported as similar to 2008	Not published	The 2009 and 2010 rulebooks are not available, so nothing is claimed
2011–2012	Site Survey, Sample Return, Astronaut Assistance, Equipment Servicing, Presentation	500	Four tasks plus a scored presentation, the architecture that still exists
2013–2016	Terrain Traversing replaces Site Survey	500	The CDR as a gate (2015), the 70 kg fielded total (2015), the only two-division year (2016)
2017	Science Cache, Extreme Retrieval and Delivery, Equipment Servicing, Autonomous Traversal	500	Autonomy as a scored mission, the CDR scored at 100, a hardware kill switch
2018	The same four missions	500	PDR added, CDR becomes the SAR, still the 100-point review
2019–2021	Science, Delivery, Equipment Servicing, Autonomous Navigation	500	The 1.2 m cube, the top-36 cut, a single E-stop that kills all battery power; the 2020 and 2021 finals were not held
2022–2026	Science, Delivery, Equipment Servicing, Autonomy	500	Drone demo (2022), drone scored (2023), drone confined to Delivery (2024–2026), 2.4 GHz deregulated (2025)
2027	Astrobiology, Delivery, Equipment Servicing, Autonomy in two sub-missions	500	A bioassay requirement, astronaut assistance, a pad-only drone, a widened 900 MHz rule

The envelope, twenty years apart

ITEM	2007	2027
Mass per configuration	50 kg	50 kg, with 5 % of points per kilogram above it
Fielded total	Not stated	70 kg, strict, since 2015
Size at weigh-in	No dimension rule	1.2 m cube, 40 % penalty since 2019
Power	Off-grid, no tethers to the command post	Unchanged
Operator view of the rover	Not permitted	Unchanged, and deliberately blocked
Link	Real-time, no time delay	Unchanged, with no satellite internet
Environment	About 100 °F, dust, light rain	37.8 °C, dust, light rain
Fielded-equipment cap	$10,000	$26,000
Design review	None	SAR, 100 points, plus the top-36 gate



The hill, the dust and the heat are the part of the competition that has never needed a rule change. The 2027 course description repeats the 2007 list almost word for word: soft sand, gravel, boulder fields, vertical drops and steep loose slopes.

Mars Society, 2017. CC BY-SA 4.0. *Original source* *Full resolution*

The 2027 changes that still shape a design

CHANGE	2026	2027	WHAT IT FORCES
Science payload	One life-detection method plus a second capability of the team's choice	Two life detections, at least one a bioassay	Fluid handling, containment and a no-spill plan inside a 20–30 minute window
Autonomy	Seven static targets, 3 m and 2 m tolerances	Astronaut assistance plus two route targets at 1 m	Speech or gesture perception, a 1 m stop, and a 40-minute clock
Judged-quality tasks	Stratigraphic profile, object recognition, marked-path drive	All three removed	Effort moves from presentation tasks to measurable capability
Equipment servicing	USB-C insertion, gasket removed, keyboard keyword only	Key and lock task, half-thickness gasket, specified display, free retries	Fine alignment on 1 cm tags and a force-controlled end effector
Drone	Delivery only, launch from the rover, wind waiver	Delivery and Autonomy, pad launch only, no waiver	A field pad, a spotter, and a 5 kg airframe carrying its dummy mass always
900 MHz scope	Command post and rover, including LAN	All links between command post, rover, drone and astronaut	The astronaut link and the suit video share the sub-band
Line of sight	Autonomy expected line of sight	Autonomy likely without line of sight, and one target deliberately blind	Onboard autonomy as the primary design, not a fallback
Cost coverage	Fielded rover	Fielded equipment, explicitly including drone and suit	Wearable electronics count; the garment does not

Operations got harder while the clock got shorter

YEARS	SETUP	TEARDOWN	INTERVENTION PENALTY
2011–2013	25 min	15 min	10 points flat per intervention
2015	20 min	10 min	20 % of that task's points
2017–2027	15 min	10 min	20 %, additive, clock runs

Setup time halved between 2011 and 2017 while the equipment count went up: a 900 MHz link, a 25 m cable run, a GNSS base station, a drone and, from 2027, a suited astronaut with video and audio S-URC27-RULES §1.a.ii. The smaller number is the harder constraint.

What the removals say

The 2027 rewrite deleted the stratigraphic profile from the science documentation and the debrief, removed the three-object recognition task from autonomy, and removed the precise marked-path drive from Delivery. Earlier removals include the 2008 spectral-analysis requirement, the 2011–2012 site survey and the 2022 USB-A insertion S-URC19-RULES S-URC22-RULES S-URC26-RULES S-URC27-RULES §1.b.iii §1.b.iii.

The pattern is consistent, and the reading of it is an inference rather than a published policy: tasks judged on the quality of a human performance tend to disappear, while tasks judged on completion, time or a measured tolerance survive. No URC document states that rationale. What follows for a design is concrete. A payload bay and an end effector should be modular enough to accept next year's object list, because the list changes, and a team should not hard-code a 2027 sub-task into a mechanism that cannot be adapted.

Dates that moved

MILESTONE	2026	2027
Registration	2025-10-29	2026-10-30
Preliminary Design Review	2025-12-03	2026-12-02
System Acceptance Review	2026-02-27	2027-03-05
Communications documentation	2026-04-24	2027-04-30
Science plan	2026-05-15	2027-05-21
Rover Expense Report	2026-05-20	2027-05-26
Finals	2026-05-27 to 2026-05-30	2027-06-02 to 2027-06-05
Review cut	Top 36	Top 36

The dates shift by about a week, and the top-36 cut does not move at all. Two deadlines are new in substance rather than date: the communications documentation now covers the drone and the astronaut link as well as the rover S-URC27-RULES §2.c.ii, and the expense report now covers the suit's wearable electronics S-URC27-RULES §3.c.iii.

Why the history is worth reading

- **Design for the stable core first.** The 50 kg limit, the off-grid power rule, the blocked view and the real-time link have survived every year read. A chassis, a thermal margin and a reliable link are never wasted effort S-URC07-RULES-L1 S-URC27-RULES §2.a.

- **Read the removals as a warning about hard-coding.** The task list is not permanent, and the end effector and payload bay are where adaptability pays.
- **Trace every inherited requirement to the current year.** The 2027 rulebook renumbered its sections: what was a rover rule in §3 is now an equipment rule in §2, and a citation copied from a 2026 document points at a different clause `S-URC26-RULES` `S-URC27-RULES`.

Further reading

- How the missions evolved follows the five mission families task by task.
- Rules, limits and the field states the 2027 clauses themselves.
- Missions and scoring covers what each mission rewards now.
- 155 engineering requirements carries the traceable version of the current rules.
- Where the competition is going reads the same record for the 2028 outlook.

THE COMPETITION · CHAPTER 2.7

How the missions evolved

The 2027 competition is the visible end of five separate task lineages, four of them older than the autonomy mission. Read as task designs rather than as rules, the history explains why each 2027 mission looks the way it does: what the rover has to do, on which objects, measured against which tolerance, and what was deliberately taken out. Two facts surprise most first-time readers. Autonomy became a scored mission only in 2017, and the drone arrived in 2022 as an explicitly unscored trial.

2017

AUTONOMY'S FIRST SCORED YEAR

Nine seasons before 2027

2017

S-URC17-RULES

2022

THE DRONE'S FIRST APPEARANCE

An unscored demonstration

2022

S-URC22-RULES

100 m

EQUIPMENT SERVICING APPROACH

Down from 0.25 km in 2025

2026–2027

S-URC27-RULES

2

LIFE DETECTIONS REQUIRED

Up from one method in 2023–2026

2027

S-URC27-RULES

Key takeaways

- **Autonomy is the youngest mission and the most changed one.** It began in 2017 as a waypoint traverse, became object recognition by 2026, and in 2027 becomes human assistance plus terrain route finding with a 1 m tolerance.
- **Science has tightened into a payload specification.** The 2023 mission accepted any one method and a 1–10 g surface sample; 2027 requires two life detections, one of them a bioassay, from at least 10 cm depth.
- **Equipment servicing lost a fifth of its drive and gained dexterity.** The approach fell from 0.5 km to 100 m while the manipulation set grew, and the recurring core is torque plus insertion at 1.5 m with a small visual marker.
- **The drone became a field asset.** It entered as a consultation, was briefly allowed to replace the rover, and is now a pad-launched Delivery and Autonomy assistant with no wind waiver.

The five mission families

FAMILY	FIRST APPEARS	NAME IN 2027	LINEAGE
Science	2007, as the Scientific Task	Astrobiology Mission	Scientific Task 2007, Geology and Soil Characterization 2008, Site Survey and Sample Return 2011–2012, Sample Return 2013–2015, Science Cache 2016–2018, Science 2019–2026, Astrobiology 2027
Delivery	2008, as Emergency Navigation	Delivery Mission	Emergency Navigation 2008, Astronaut Assistance 2011–2016, Extreme Retrieval and Delivery 2017–2022, Extreme Delivery 2023, Delivery 2024–2027
Equipment servicing	2008, as Construction	Equipment Servicing Mission	Construction 2008, Equipment Servicing Task 2011–2016, Equipment Servicing Mission 2017–2027
Autonomy	2017	Autonomy Mission	Autonomous Traversal 2017–2019, Autonomous Navigation 2020–2026, Autonomy in two sub-missions 2027
Aerial	2022, as an unscored demo	A regulated equipment class	Demo 2022, rover substitute 2023, Delivery assistant 2024–2026, Delivery and Autonomy assistant 2027



The 2009 BYU rover on the same site: six wheels, a small arm, an instrument deck and a mast. Two decades of rule changes have not changed the shape of the machine, which is why a chassis and a thermal margin designed to the stable core still transfer to 2027.

Peter Hyatt, 2009. CC BY 3.0. *Original source* *Full resolution*

Science: the measuring stick

YEARS	ROVING TIME	LIFE DETECTION MINIMUM	SAMPLE	CACHE	STRATIGRAPHIC PROFILE
2007	1–3 h	Investigate geology and palaeontology; digging allowed, no explosives, on-site analysis only	Not specified	None	No
2008	30–60 min plus equal preparation	Two separate events, one using visible and infrared spectral analysis	Soil and subsurface probes	None	No
2011–2012	30–60 min	Site Survey and Sample Return as separate tasks	Not specified	None	No
2016–2018	Not stated	A science cache with a judged hand-off	Sample in a cache	Container opened by judges	No
2023	20–35 min	One method or assay of the team's choosing, onboard	1–10 g, surface, no depth rule	Closed container	No
2024–2026	20–30 min	One capability plus a second capability of the team's choosing	At least 10 cm deep, at least 5 g	Sealed, spill-proof when inverted, removable in 5 min	Required
2027	20–30 min	Two life detections, at least one a bioassay	At least 10 cm deep, at least 5 g, unchanged	Sealed, unchanged	Removed

Three breaks matter more than the trend line. In 2008 science was two scored events with two different instrument classes, so the year rewarded breadth in a way no later year did. In 2013 the survey event disappeared and science became one task. In 2024 the mission gained the 10 cm depth rule, the 5 g floor, the sealed cache, the two-site comparison and the stratigraphic profile all at once, and only the profile came out in 2027 `S-URC08-RULES` `S-URC13-RULES` `S-URC24-RULES` `S-URC27-RULES` §1.b §1.b.

The design consequence is blunt. An auger producing 1–10 g of surface material satisfied 2023 and fails every rulebook since. Sub-surface acquisition to at least 10 cm with a documented depth has been mandatory for four consecutive seasons, and 2027 adds the bioassay on top of it S-URC24-RULES S-URC27-RULES §1.b.vi §1.b.v.

Delivery: from a speed race to a staged mission

The family starts as Emergency Navigation in 2008: deliver a supply container to a simulated astronaut as fast as possible, with the container capped at 30 × 30 × 15 cm and 6 kg, orientation-independent, and covered in hook-and-loop so the rover had to strap it on rather than grip it. Delivery was within 1 m of the suit, the astronaut was up to 1.5 km away, and line of sight was guaranteed for only 80 % of the course S-URC08-RULES . From 2011 the task became Astronaut Assistance, then a staged retrieval and delivery mission in 2017, and Delivery from 2024.

By 2027 the menu is a set of capabilities rather than a course: open boxes with hinged lids, pick up and deliver objects, read signs carried by astronauts or placed in the field, find equipment in a large search area, and find objects where radio reception is poor. The 2027 rewrite removed one item, the precise marked-path drive, and emphasised that objects appear in a variety of sizes, shapes and weights S-URC26-RULES S-URC27-RULES §1.c.iii §1.c.iii.

OBJECT ENVELOPE	VALUE	YEARS IN FORCE
Mass	Up to 5 kg	2026–2027, with 6 kg for the 2008 container
Bounding size	Under 40 × 40 × 40 cm	2018 onward, unchanged
Grasp feature	7.5 cm diameter	2026–2027, widened from 5 cm
Object identity	Disclosed at the delivery location, not before	2026–2027
Aerial help	Drone allowed, no mobile deployable	2026–2027

The 7.5 cm grasp feature is the sharpest mechanical constraint in the competition. Rocks with natural features are legal targets, so the end effector cannot be designed around a known handle. A 5 kg object with an unknown centre of gravity and a 40 cm envelope is the design case, not an object with a moulded grip.



The Delivery terrain description has not changed since 2008: soft sand, gravel, rough stony ground, boulder fields, vertical drops and steep loose slopes. What changed is that later Delivery tasks are now deliberately out of radio line of sight.

Mars Society, 2017. CC BY-SA 4.0. Original source Full resolution

Equipment servicing: the drive collapsed, the dexterity grew

The mission began in 2008 as a construction task: secure half-inch hex bolts on a mock panel between 15 cm and 75 cm above the ground, at up to 45° of inclination, on a drive of up to 0.5 km. In 2011 it became a read-the-instructions task with a printed sequence and 3-prong plugs S-URC08-RULES S-URC11-RULES §3.e.

YEARS	APPROACH DISTANCE
2011–2016	0.5 km
2017–2025	0.25 km
2026–2027	100 m

A fivefold reduction in the drive over sixteen years, with the manipulation content increasing. That is the clearest statement in the rulebook that this mission is an arm mission, not a mobility mission.

TASK	YEARS PRESENT	WHAT IT DEMANDS
Push buttons, flip switches, turn knobs	2011–2027	Force and travel against an unpublished value
Read a printed sequence on the panel	2011	Text legibility on a vertical surface
Connect a three-prong plug	2011	Insertion with orientation
Sample tube into cache, close lid, turn lock	2026–2027	Insertion plus torque in one sequence
Carry the cache, open the drawer, insert and close	2019–2027	Grasp, alignment and a tight fit
Tighten a captive screw	2019, 2022	Torque with an Allen key
Autonomous typing on a vertical keyboard	2019–2027	Precision with a 2 cm fiducial at the corners
Insert a USB connector	2022, 2026	Fine connector alignment
Turn a quarter-turn lever valve	2026–2027	Torque through 90°
Push a 1.5 in cam-lock hose onto a fitting	2026–2027	Push force, travel and retention
Remove a key, insert it into a second lock, turn it	2027	Grasp, keyway alignment and torque on a 1 cm tag

The churn is constant and the handedness is stable. Every year since 2019 has had at least one torque task and at least one insertion task, and the recurring core is reach at 1.5 m to the ground, grasp of a 5–10 cm feature, torque application, and a small visual marker for feedback S-URC19-RULES S-URC27-RULES §3.d §1.d.ii. A first build can treat that core as the durable requirement and the object list as the annual variable.

The 2027 geometry also reverses the back-to-back order: 2026 ran servicing into autonomous navigation, while 2027 runs Autonomy into Equipment Servicing, possibly only ten minutes apart, from the same command post on an adjacent course S-URC26-RULES S-URC27-RULES §1.a.iv §1.a.iv. That makes the payload and sensor configuration part of the schedule, not only part of the design.

Autonomy: from waypoints to a walking astronaut

Before 2017 no mission required autonomy. The rulebooks said the rover was not required to be autonomous, and the 2017 rules still confined the requirement to one task. It then went from optional to a fifth of the score in a single season S-URC07-RULES-L2 S-URC16-RULES S-URC17-RULES §2.c.

PERIOD	SHAPE	TOLERANCE	ABORT COST	HUMAN ROLE
2017	Reach an object and waypoint set, partial credit within 5 m	5 m for partial credit	Not stated	None
2019–2022	Staged legs between physical markers: tennis balls, then gates, then AR tags	Marker coordinates approximate, search required	Not stated	None
2026	Two coordinate points, two tagged posts and three objects to highlight on the display	3 m and 2 m	Autonomous return free, teleoperation 20 %	None
2027	Astronaut assistance worth 50 and route finding worth 50	1 m at targets, 3 m at the astronaut	20 % on a stop or return, 50 % for teleoperation, capped at 50 %	A suited astronaut who walks, waits, receives and hands over

The 2027 Follow! task is the one place where the method itself is scored, and the spread is large enough to be a design decision rather than a preference S-URC27-RULES §1.e.vi.

FOLLOW! METHOD	POINTS	SHARE OF THE 15 AVAILABLE
Astronaut-carried device	5	One third
Visual sign	5	One third
Spoken word or phrase	10	Two thirds
Beckoning gesture	15	Full

The gesture path is worth three times the device path, and it is a perception problem rather than a radio problem. Everything else in the sub-mission is binary, which is why a reliable simple method for Stay!, Fetch!, Come! and Give! is enough.

The drone: five years from demo to field asset

YEAR	ROLE	LAUNCH AND MASS
2022	Optional demonstration, not scored, with an open invitation for feedback on making it a mission	5 kg cap, 250 g registration threshold, dummy mass guidance
2023	Allowed to fly instead of the rover in Autonomous Navigation	5 kg, dummy mass separate; required dummy mass waived above 24 km/h
2024	Delivery assistant only	Drone plus payloads plus dummy mass at 5 kg; waiver above 24 km/h
2025	Delivery assistant	5 kg including dummy mass; waiver above 24 km/h
2026	Delivery assistant	Launch from the rover or a start location near the command post; waiver above 24 km/h
2027	Delivery and Autonomy assistant	Designated landing pad only; never waived; over 5 kg not permitted to fly

The permanent constraints arrived with the demonstration and never left: rotary wing only and able to hover, no lighter-than-air vehicle, no fixed wing, FAA Remote ID, a TRUST certificate carried by every pilot, below 120 m above ground in Class G airspace, not in the national park, give way to manned aircraft near Hanksville, and a spotter with continuous visual line of sight S-URC22-RULES S-URC27-RULES §4 §2.b. The 2027 change is not a new capability but a narrowing: the airframe is a field asset with a pad, a pilot in the command post and a spotter on the ground.

The review never got easier

YEARS	REVIEW ELEMENT	SCORED	CUT
2007–2014	None; a video presentation from 2011	Presentation only	None
2015–2016	Critical Design Review	No, a down-selection gate	Pass or fail
2017	Critical Design Review	Yes, 100 of 500	Pass or fail
2018	Preliminary Design Review added, CDR becomes the System Acceptance Review	Yes, 100 of 500	Pass or fail
2019–2027	PDR and SAR	Yes, 100 of 500	Top 36

The review has been a fifth mission since 2017: 100 points and a gate on the other 400. For a first-time team it is the highest-leverage document of the season, and it rewards collected test evidence rather than prose.

Three readings for a first build

1. **The physical problem is twenty years old.** 50 kg, off-grid, about 100 °F, no direct view and a real-time link have survived every rulebook read. Time spent on a robust chassis, a thermal margin and a reliable link is never wasted S-URC07-RULES-L1 S-URC27-RULES §2.a.
2. **The score has migrated from manipulation and driving to perception and science.** The servicing drive fell fivefold, autonomy went from nothing to 100 points and now includes human interaction, and the science payload went from any instrument to two life detections with a bioassay S-URC11-RULES S-URC27-RULES §1.b.v.
3. **Treat every sub-task as temporary.** The 2008 spectral analysis, the 2011–2012 site survey, the 2019 board swap, the 2022 USB insertion, the 2026 object recognition and the 2026 stratigraphic profile all disappeared. Modularity at the end effector and the payload bay outlasts any specific object list. The reading of a judging-reliability motive behind the removals is an inference, not a published policy.

Further reading

- Rule history, 2007-2027 dates every change to the rulebook that carries it.
- Missions and scoring states the 2027 tasks and their point values.
- Mission-by-mission backfit turns each 2027 task into a design choice.
- Astronaut interaction and the EVA suit covers the newest mission in depth.
- Where the competition is going extends the same record into 2028.

title: "Site, travel and field constraints" description: The desert site as an engineering constraint: the command post, the clock, the spectrum, the road and the heat. part: The competition chapter: "2.8" toc_depth: 3 tags: - logistics - operations - field

Site, travel and field constraints

At URC the logistics are engineering constraints. The command post is a metal box with no view, so every antenna is outside and every display is read blind. The setup window is fifteen minutes, so the ground station is a deployable system rather than a packing problem. The radio spectrum is monitored and split into three sub-bands that can be assigned per mission. The site is at 1,309 m, reaches 37.8 °C, has no cell service and permits no testing at the competition location. Each of these becomes a mass item, a time budget or a test-plan constraint.

15 min

SETUP FROM COMMAND-POST ACCESS

Covers all communications and the weigh-in

2027

S-URC27-RULES

25 m

COMMUNICATIONS CABLE

One run for the rover and one for the drone

2027

S-URC27-QA

1,309 m

SITE ALTITUDE

Thinner air, so less fan cooling

MDRS

S-URC-LOGISTICS

37.8 °C

AIR TEMPERATURE LIMIT

Plus dust, light rain and full daylight

2027

S-URC27-RULES

Key takeaways

- **Six students, one box, no view.** The command post holds at most six students, excludes advisors, blocks the course from sight and blocks radio through its walls S-URC-LOGISTICS S-URC27-RULES §3.d.i.
- **Fifteen minutes in and ten minutes out** is the whole ground-station design budget, and a drone link on its own antenna adds a second 25 m run S-URC27-RULES S-URC27-QA §1.a.ii DQ2.
- **The spectrum is a referee.** The 900 MHz link must work in each of three 8 MHz sub-bands, and from 2027 the rule covers the drone and the astronaut link as well as the rover S-URC27-RULES §2.c.v.
- **There is no testing at the site and no parts nearby.** The nearest electronics source is a six-hour round trip and the nearest trauma hospital is at least two hours away S-URC-ADVICE S-URC-DESERT-SAFETY .

Who may be on site

CONSTRAINT	VALUE
Registered students	24
Registered advisors	2
Individuals on site per competition day	14
Students inside the command post during a mission	6 or fewer
Advisors inside the command post	Not permitted
Minimum age	18 on the first competition day, unless the Director grants prior approval
Team size during the year	Unlimited, with students enrolled at least half-time

The 2027 rulebook defers the attendance figures to the URC website, so the published numbers above are the working assumption S-URC27-RULES S-URC-LOGISTICS §3.b.iii.

The six-person limit is a console count, not a suggestion. Driver, science operator, arm operator, drone pilot, autonomy operator and mission commander fill it, and there is no seventh seat and no advisor to take one S-URC-LOGISTICS S-URC27-RULES §3.b.ii. Two further roster rules follow from it. A team member who leaves the command post to act as a runner cannot return to operate or analyse data for that mission, so the field crew of spotter, runners and astronaut has to be rostered separately. And the Astrobiology briefing is open only to members who stayed inside and never saw the site, so two people have to stay clean all week S-URC27-QA AQ8.



Fourteen people on site, six in the command post and a rover in the field is the shape of a competition day. The site is a desert research station with no cell service, so every roster decision is also a communications decision.

Mikołaj Stasiak, 2026. CC BY 3.0 pl. *Original source* *Full resolution*

The command post

> These stations will be metal trailer units (such as the back of a small moving truck provided by URC) or structures at the Mars Desert Research Station. Visibility of the course to the operators in the C2 station will be blocked. Basic power (120 V, 60

> Hz), tables and chairs will be provided. S-URC27-RULES §3.d.i

The judges' advice page is blunter about the station at the research station: "This is a metal structure and radio signals, including WiFi, will NOT go through the walls" S-URC-ADVICE . Both statements point the same way. The antenna is outside, the cable is long, and the operator sees only what the rover sends back.

PROVIDED	NOT PROVIDED
120 V, 60 Hz power	A stated current limit or a guarantee of an uninterruptible supply
Tables and chairs	Network, internet or cell service
A structure with a blocked view	Any allowance for a slow setup

That turns the ground station into a deployable system with a stopwatch on it: base station, GNSS base, antenna mast, two radio links, six operator positions, a drone control position and a power distribution board, carried in cases, erected by hand in the sun, and reversed in ten minutes S-URC27-RULES §1.a.ii.

Setup, teardown and the mission clock

CONSTRAINT	VALUE
Setup from command-post access to ready, including all communications	15 minutes or less
Teardown and removal of all equipment, then radios off	10 minutes or less
Recovery of equipment left in the field	Inside the same 10 minutes, after mission time ends
Autonomy to Equipment Servicing turnaround	As little as 10 minutes, same command post, adjacent course
Mission order	Drawn by lot, with the schedule known about a week ahead

The schedule arriving about a week before the field event and the order being drawn by lot are configuration requirements rather than inconveniences. A rover has to be mission-agnostic in its transport state, and any module swap has to finish inside the setup window S-URC27-QA CQ4.

Cable, mast and antenna

CONSTRAINT	VALUE
Communications cable the team brings	At least 25 m
Cable for drone operations	A second 25 m run when the drone link has its own antenna
Base-station antenna height	3 m or less
Antenna base position	Within 5 m of the command post
Guy anchors	Within 10 m of the command post
Judge-designated positions at site structures	Up to 20 m, to avoid buried pipes and blocking structures
Steered antennas	Allowed, with no camera on the antenna and no reliance on watching the rover

A 25 m run costs real decibels at 2.4 GHz and at 5 GHz, so the cable loss budget belongs in the link budget rather than after it. The 20 m exception exists to keep the team away from buried services at the structures, and it changes the cable length and the guy geometry rather than the mast S-URC27-RULES §2.c.iii.

The field environment

CONSTRAINT	VALUE	DESIGN CONSEQUENCE
Air temperature	Up to 37.8 °C	Surface temperatures in sun are higher, so the enclosure heat load is set by sunlight
Wind and dust	Dust raised frequently; rover must withstand it	Sealed connectors, filtered intakes, a cleaning routine
Rain	Light rain must be tolerated; heavy rain and thunderstorms are not expected	Splash protection rather than a waterproof claim
Light	All events in full daylight	Displays need a hood or high brightness; the LED needs a daytime check
Altitude	1,309 m	Thinner air reduces air-cooled heat-exchanger performance
Terrain	Soft sand, gravel, rough stone, boulder fields, vertical drops, steep loose slopes	The mobility envelope in one line
Transit	A washboard dirt road from the highway to the site	Connectors and harnesses have to survive vibration

Drone wind is separate and advisory: the rules strongly suggest operating in 48 km/h wind with a top speed of at least 64 km/h, and the dummy mass requirement is never waived S-URC27-RULES §2.b.vi.

Spectrum: the invisible referee

CONSTRAINT	VALUE
900 MHz bandwidth	8 MHz maximum per use
900 MHz sub-bands	902–910, 911–919 and 920–928 MHz, usable exclusively
Sub-band choice	Assigned per mission by the competition schedule, so the radio must shift
Scope	All links between command post, rover, drone and astronaut, including any local network between devices
Monitoring	On site, to keep teams inside their assigned channels
2.4 GHz	No URC restriction, but FCC rules apply
Other bands	Amateur bands are encouraged for licensed teams
Documentation	Devices and licences to the URC Director by 2027-04-30
Non-competing time	Radios powered down at the event site

The scope line is the 2027 change that catches designs: a link that was a rover link in 2026 is a rover, drone and astronaut link now, and an internal network that happens to run in the 900 MHz band is inside the same limit S-URC27-RULES §2.c.v. A team on the site is also within signal range of other teams on other sub-bands, so interference is expected rather than exceptional.

Where testing is allowed and what is not

- **No testing at the site.** The rules prohibit it before, during and after the competition S-URC27-RULES §3.d.iv.
- **Permitted areas only.** Testing is limited to places where off-road vehicles are permitted, such as the Factory Butte open area, or to town land. Undesignated public land is not allowed.
- **No cell service.** There is no telephone and no cell coverage at the site, so documentation, licence copies, spares lists and vendor contacts have to be carried offline S-URC-ADVICE S-URC-DESERT-SAFETY.
- **No parts nearby.** The nearest electronics source is a six-hour round trip, which makes a declared spares kit a requirement rather than a preference.
- **Drone airspace.** The site is Class G, below 120 m above ground, about five miles from an uncontrolled airport, and drones may not enter the national park S-URC27-RULES §2.b.xi.
- **Drone paperwork.** Remote ID, registration above 250 g, a TRUST certificate carried by every pilot and a spotter with continuous visual line of sight.

Travel, vehicles and lodging

CONSTRAINT	VALUE
Public transport to the site	None; the team must have its own vehicle
Permitted vehicles	Pickups, cargo vans, minivans, sport utility vehicles and 15-passenger vans
Not permitted	Box trucks of 14 ft or more, and any stakebed truck
Airports	Grand Junction 159 mi, about 2.5 h; Salt Lake City 236 mi, about 4 h; Las Vegas 405 mi, about 6.5 h
Fuel spacing	Over 100 mi between stations is common, so do not pass a station with under 150 mi of range
Hanksville to the site	20–25 min over dirt road
Lodging	Hanksville, with overflow in Caineville, Torrey and Green River
Lodging rules	No rover work or hand-tool use indoors
Fees	$197 application, $397 event registration after review selection

Two details cost teams time at the wrong moment. Lodging blocks are released after the review results, so a cancellable fallback is worth holding. And components, including lithium batteries, have been seized in customs, so a US purchase backup belongs in the travel plan for any international team `S-URC-LOGISTICS`.

Health and safety

CONSTRAINT	VALUE
Nearest trauma hospital	At least 2 h by road, or 90 min by air ambulance
Nearest medical clinic	At least 1 h by road
Emergency communications	No telephone and no cell service at the site
Footwear	Fully enclosing footwear is required at all events
Water	Drinking-water re-supply is provided at the mission sites
Hazards	Venomous snakes, wasps and hornets, scorpions, spiders and cougars
Chemical response	No chemical safety response equipment on site

A written medical plan, a designated first-aid lead, an evacuation route and medication for the whole trip plus a margin are all part of the field kit. The absence of chemical response also sets the science posture: containment and a no-spill regime, because there is nothing to clean a spill with `S-URC-DESERT-SAFETY` `S-URC26-SCIPLAN`.

The season to the clock

WHEN	WHAT	WHAT IT GATES
2026-10-30	Registration closes	Support letters, and the competition itself
2026-12-02	Preliminary Design Review	A five-page report in the 2026 format
2027-03-05	System Acceptance Review	A seven-page report and a 5:00 video; the top 36 go to the field
2027-04-30	Communications devices and licences	The radio hardware freeze for every link
2027-05-21	Science plan	About five pages, scored in Astrobiology
2027-05-26	Rover Expense Report	A work-breakdown level accounting, with a deadline-day penalty
One week before finals	Mission schedule and 900 MHz sub-band per mission	Radio agility and mission-order rehearsal
2027-06-02 to 2027-06-05	Finals	Everything above, in order

What the site does not publish

- **Command-post current capacity.** The rules give 120 V at 60 Hz and no current limit, no circuit count and no guaranteed backup. A design load of about 10 A continuous with a battery buffer is a defensive reading, not a published figure.
- **Whether the antenna limits cover the GNSS base.** The 3 m and 5 m figures are written for the communications antenna. Planning to 3 m is the safe interpretation.
- **The 2027 command-post details.** Whether the station is a trailer or the station structure, and the power arrangement at each site, are not published for 2027; the 2026 text is the available description.
- **The vehicle list revision date.** The permitted-vehicle document is linked from the logistics page and carries no revision date.

- **Lodging block timing.** The release of room blocks after review results is 2026 practice, not a stated 2027 rule.
- **Spectrum enforcement detail.** The rules state that monitoring exists and do not say what follows a breach.
- **Medical support at the site.** Distance to care is published; whether a medic is present during missions is not.

Further reading

- Rules, limits and the field carries the clauses quoted here.
- Q&A rulings holds the antenna, cable and command-post answers in full.
- Deadlines turns the season table into dated milestones.
- Operator station and operations designs the ground station against these limits.
- Environment, dust, thermal and repair turns the desert into a test plan.

THE FIELD · CHAPTER 3

The URC landscape

Sixty-eight teams appear in the official URC score sheets from 2022 to 2026, and roughly one in three registered teams reaches the field. Five seasons of results describe a field that is stable in the middle and fragile at the edges: the median finalist scored between 224.6 and 257.8 of 500 in every season, while a handful of missions produced almost all of the zeros. The pages that follow carry the census, the results explorer, the team browser and the study of sister competitions; this page connects them.

68

TEAMS IN THE CENSUS

187 team-years, 2022–2026

2022–2026

S-CENSUS

116

REGISTRANTS IN 2026

36 to 38 reach the field

2026

S-URC26-TEAMS

36–38

FINALS FIELD

36 invited after the SAR

2027

S-URC27-RULES

30.6 %

AUTONOMY SHARE OF TOTAL VARIANCE

Largest of the five components

2022–2026

S-URC26-SCORES

Key takeaways

- About one in three applicants reaches the field. Registrations ran from 99 in 2022 to 116 in 2026, while the finals field stayed at 36 to 38 teams S-URC22-TEAMS S-URC26-TEAMS .
- Completion beats optimisation. Science produced no zero scores in five seasons; Autonomy produced 67 zeros in 162 team-years, a 41 % base rate S-URC22-SCORES S-URC26-SCORES .
- The written review gates the field and does not rank it. The SAR explains 2.8 % of the variance in final totals; Autonomy explains 30.6 % S-URC22-SCORES S-URC26-SCORES .
- The middle of the field has not moved. The median finalist scored between 224.6 and 257.8 of 500 in every season, while the winner rose from 388.97 to 469.57 S-URC22-SCORES S-URC26-SCORES .

§1.c.ii." aria-label="Open full-size figure">



The terrain above the Mars Desert Research Station at Hanksville, Utah, where every edition since 2007 has been held except the two pandemic years. The 2027 course runs on the same terrain `S-URC27-RULES §1.c.ii.`

Mars Society, 2017. CC BY-SA 4.0. *Original source* *Full resolution*

The field in one table

Five seasons, five official score sheets. The registration counts come from the URC team lists of each year; the field counts come from the score sheets themselves.

SEASON	TEAMS REGISTERED	ROWS IN THE SHEET	ROWS WITH A MISSION SCORE	COUNTRIES REPRESENTED	WINNER AND TOTAL
2022	99	36	26	15	Michigan MRover, 388.97
2023	104	37	31	17	WVU Team Mountaineers, 425.35
2024	102	38	34	18	AGH Space Systems, 392.76
2025	114	38	35	18	Missouri S&T, 412.27
2026	116	38	36	18	Missouri S&T, 469.57

Registration counts come from the season team lists `S-URC22-TEAMS` `S-URC23-TEAMS` `S-URC24-TEAMS` `S-URC25-TEAMS` `S-URC26-TEAMS` *; field counts and winners from the official score sheets* `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES` *. The 2026 country count includes two multi-national entries.*

One in three reaches the field

The score sheet is not the registration list. Registered teams clear the System Acceptance Review to be invited to the 36-team field, and the 2023 sheet has exactly 37 rows, the announced finalist count for that year, which supports reading a score-sheet row as finalist status `S-URC23-NEWS` . For the other four seasons the published sheet is the only record of who competed.

Two facts follow from that reading. Roughly two-thirds of registrants never reach Utah, so the SAR is the competition's first cut rather than a formality S-URC27-RULES §3.a.iii. A missing season is also not evidence of a worse rover: AGH Space Systems won in 2024 and recorded a SAR score with no finals result in 2025 S-URC25-SCORES.

Twenty-five of the 187 rows across the five seasons carry a SAR score and no mission score at all. Those teams cleared at least part of the desk and left no finals record; the census page lists them with their SAR scores.

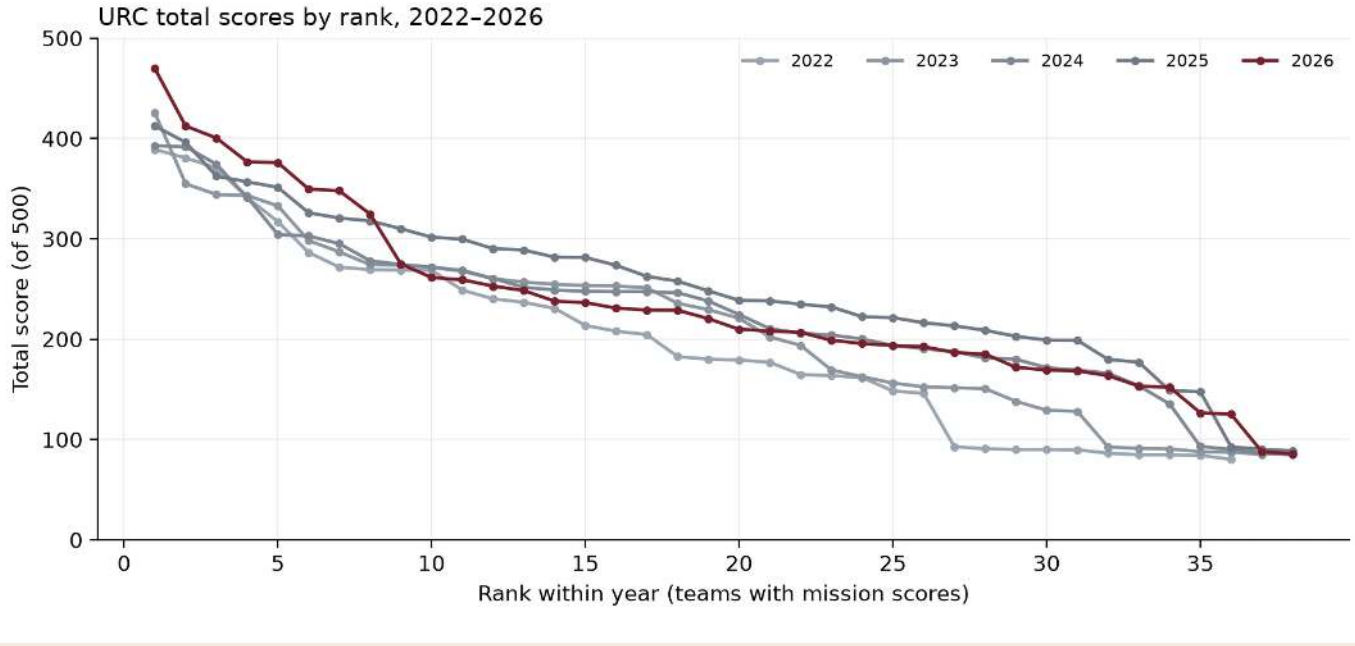
Where the teams come from

The American block dominates by headcount. Of the 116 teams registered for 2026, 60 came from the United States, followed by India (11), Bangladesh (9), Canada (9), Türkiye (6), Poland (5) and Japan (3), with two each from Colombia, México and Korea, and single entries from Australia, Czechia, Italy, Perú, Puerto Rico, Singapore and one Japan/US joint team S-URC26-TEAMS.

Depth and quality are not the same thing. The United States supplies 21 of the 68 census teams and 70 of the 162 finals team-years, yet its mean total is close to Poland's and Poland's mean final rank is better S-CENSUS S-URC26-SCORES. Australia is a single team, Monash Nova Rover, and that team carries the highest mean total in the dataset at 343.1.

Fifteen country labels appear across the five score sheets and 14 of them in a finals field; Egypt's Pharaohs Team carries SAR-only rows and never converted one into a finals result S-CENSUS.

How performance is distributed



Every team-year total, sorted within each season. The shape is consistent: a small leading group, a long middle, and a tail below 150. The 2026 season is the most stretched of the five, with the winner at 469.57 and the last finalist at 125.14 S-URC26-SCORES.

Robotics Club @ Mississippi State University, 2022. S-URC22-SCORES

S-URC26-SCORES *Full resolution*

The two halves of the distribution behaved differently. The median finalist scored between 224.6 and 257.8 points for five straight seasons; a clean, reliable rover that finishes all four missions lands near that band. The winner's total rose 21 % over the same period, from 388.97 in 2022 to 469.57 in 2026, and 2026 was the first season in which any team passed 450 S-URC22-SCORES S-URC26-SCORES.

SEASON	SCIENCE MEAN (ZEROS)	DELIVERY MEAN (ZEROS)	EQUIPMENT MEAN (ZEROS)	AUTONOMY MEAN (ZEROS)
2022	81.3 (0)	25.1 (8)	21.8 (6)	23.9 (12)
2023	72.3 (0)	15.9 (2)	35.3 (7)	26.4 (15)
2024	73.5 (0)	22.9 (2)	32.0 (5)	23.9 (17)
2025	79.6 (0)	30.9 (1)	31.5 (4)	31.5 (7)
2026	78.5 (0)	22.9 (12)	29.1 (5)	26.7 (16)

Mission means and zero counts for teams that recorded a mission score `S-URC22-SCORES` `S-URC26-SCORES`. Denominators vary by season because not every finalist records every mission.

Three patterns hold in every season:

1. **Science is the reliable points.** No team scored zero on Science in five years, the mean is the highest of the four, and the spread is the narrowest. A competent, well-documented science effort returns 70 to 90 points almost regardless of how the rover drives.
2. **Autonomy and Delivery are where the field fractures.** Across the 162 finals team-years, Autonomy recorded 67 zeros and Delivery 25, against none in Science `S-URC22-SCORES` `S-URC26-SCORES`.
3. **Equipment Servicing rewards a complete manipulator.** Its mean stays in the low 30s with a standard deviation almost as large, and 4 to 7 zeros a season.

The variance decomposition is the sharpest statement of the same result. Because the total is the sum of five components, the share of total variance each component explains can be computed directly over the 162 finals team-years:

COMPONENT	SHARE OF FINALS-TOTAL VARIANCE
Autonomy	30.6 %
Equipment Servicing	27.7 %
Delivery	24.3 %
Science	14.4 %
SAR	2.8 %

Computed as the covariance of each component with the final total, divided by the variance of the total, pooled over 2022–2026 `S-URC22-SCORES` `S-URC26-SCORES`. The decomposition covers finalists only, every one of whom has already passed the SAR.

The written review has never explained more than about 6 % of the spread of final scores in any era: the CDR explained 1.9 % in 2017 and the SAR 5.7 % in 2018, 3.3 % in 2019 and 2.8 % pooled across 2022–2026 `S-URC17-SCORES` `S-URC18-SCORES` `S-URC19-SCORES` `S-URC22-SCORES`. The rank correlation tells the same story. The pooled Spearman correlation between SAR score and final rank is 0.48, falling to 0.34 and 0.36 in 2025 and 2026, and the top ten by SAR overlaps the top ten by total by only five or six teams each year. The desk gates the field; the field ranks it.

The teams to study first

Tier A holds the eight programmes that reached a podium or combined a top-eight finish with full public documentation. The mean total covers finals seasons only.

TEAM	COUNTRY	BEST RANK	MEAN TOTAL 2022–26	WHAT SETS IT APART
Missouri S&T	United States	1	378.5	Champion in 2025 and 2026; the most complete recent documentation
WVU Team Mountaineers	United States	1	371.7	2023 champion; top four in four of five seasons
Michigan MRover	United States	1	336.4	2022 champion; INCOSE systems-engineering award in 2026
AGH Space Systems	Poland	1	336.0	2024 champion; the strongest single-season peaks
Monash Nova Rover	Australia	2	343.1	2026 runner-up; open-source software stack
BYU Mars Rover	United States	3	318.4	Third or fourth in 2023–2025, then thirteenth in 2026
UIU Mars Rover	Bangladesh	3	308.2	2026 third; the strongest sustained climb in the field
Northeastern	United States	3	289.7	2025 third; deep autonomy work

Best rank and mean total from the census, finals seasons only `S-CENSUS`*. Tier labels are defined on the census page.*

What the leading teams share

Five patterns recur across the Tier-A programmes. They are inferences from results structure, tested against the dossiers and the subsystem chapters rather than measured directly.

- **They finish missions.** Their score profiles carry few zeros, and the advantage comes from breadth across three or four missions rather than one spectacular run.
- **They have working autonomy.** Even a modest autonomy score earns 40 to 90 points, which the field's zero-heavy distribution makes disproportionately valuable.
- **Their science is systematic rather than clever.** High Science scores follow documentation discipline and a repeatable analysis workflow, not an exotic instrument.
- **They invest in communications as a subsystem.** Poor comms is the most common reason a rover sits motionless, and it is the judges' own headline advice.
- **They come back.** Seven of the eight recorded a finals score in all five seasons; AGH raced four editions and reached three finals `S-CENSUS`.

Strategic consequences for a first build

1. **Survival first.** Two-thirds of registrants never reach the field, so the first goal is a SAR that ranks inside the top 36, which rewards demonstrated capability rather than CAD.
2. **Completion beats optimisation.** A rover that finishes all four missions at half of the available points beats a rover that excels at one and fails two.
3. **Science is the safest 70 points; autonomy is the biggest differentiator.** Engineering effort follows that split, and both rest on a rover that reliably drives, sees and communicates.
4. **The horizon is three years.** The leading programmes are repeat competitors; a first-year team that qualifies and scores in the middle of the field has done well.
5. **Failures repay study as closely as wins.** Most of the failure evidence in this reference comes from teams outside the top ten.

Where to go next

- The 68-team census carries the full team list, the tiering rule and the teams without a profile.
- Results explorer, 2022–2026 reads the score sheets row by row, including the difference between a zero and a non-entry.
- Team browser filters all 70 records by tier and country and links to each profile.
- Team dossiers holds the 51 team studies, grouped by tier.
- Sister competitions covers five other rover events and two NASA student events, and where their evidence misleads.

- What separates the top from the rest works the variance and zero analysis at length.
- The field in statistics collects the census and registration figures on one page.

THE FIELD · CHAPTER 3.1

The 68-team census

The census is the team list behind every statistic in this part of the reference: 68 teams and 187 team-years drawn from the five official URC score sheets of 2022 to 2026. Each team carries a tier derived only from its results and from whether it publishes design material, never from reputation. The page sets out the rule, the table and the nineteen teams that publish too little to support a profile of their own.

68

TEAMS IN THE CENSUS

Tier A 8, Tier B 24, Tier C 36

2022–2026

`S-CENSUS`

187

TEAM-YEAR ROWS

162 with mission scores, 25 SAR-only

2022–2026

`S-CENSUS`

36

TEAMS IN A FINALS FIELD

37 and 38 in practice

2023–2026

`S-URC26-NEWS`

51

TEAMS WITH A PROFILE

19 census teams have none

2022–2026

`S-CENSUS`

Key takeaways

- The census holds 68 teams over five seasons, and thirteen of them appear in every edition `S-CENSUS`.
- Sixty-two of the 68 recorded at least one finals result. Of the six that did not, five hold only a SAR-only row and one, UBC Rover, still carries a profile.
- Nineteen census teams have no profile: their public record is a score-sheet line and little more. Eighteen hold at most one finals appearance, and thirteen hold exactly one.
- Tiering depends on results first. Documentation can move a team one band, and only when its results are already close to the band above.
- A zero in the score sheet is not a diagnosis, and a missing season is not a decline. Both readings are explained below.

What the census counts

The unit is a team, not a club name. The census keeps one row per team, with the seasons it appears in a score sheet, its finals results, its best rank, its mean final total and the website that URC listed for it `S-CENSUS`.

Two counts run through every page of this part:

- **187 team-years.** Every row of the five score sheets, 2022 to 2026, after normalising team names.
- **68 teams.** The number of distinct programmes those rows resolve to.

Of the 187 rows, 162 record a mission score and count as finals results. The other 25 carry a SAR score and no mission score, so they appear in the census but not in any ranking `S-CENSUS`.

How the table was assembled

Three properties of the source material matter for anyone using the numbers.

Ranks are derived, not published. No score sheet has a rank column. Each sheet is printed with teams in descending total, followed by teams holding only a SAR score, and the rank in the census comes from that row order among teams with mission scores. This is an inference from the official ordering, not a published field `S-CENSUS`.

Team names are not stable. The same programme appears under two or three names across seasons: Northeastern University Mars Rover Team, Northeastern Mars Rover Team and Northeastern Mars Rover are one club. Every merge is recorded with the evidence that justifies it, and pairs that look similar but are different teams were deliberately kept apart, such as A.S.T.R.A. of Alabama-Huntsville and Astra Robotics of RV College of Engineering `S-CENSUS`.

A zero is not a diagnosis. A mission score of zero can mean attempted and failed, did not attempt, or withdrew. The sheets cannot separate those cases, so any causal account of a specific failure comes from a report, a video or a news item, not from the score `S-CENSUS`.

The tiering rule

The rule is applied mechanically to the census. It uses four inputs per team: best finals rank, number of finals appearances, mean finals rank, and a public-documentation flag.

TIER	CONDITION, APPLIED IN ORDER	TEAMS
A	Any podium in 2022–2026; or a mean finals rank of 6.0 or better over two or more finals; or a best rank of 6 or better with a full public report set	8
B	Not A, and a mean finals rank of 20.0 or better over two or more finals, or a best rank of 12 or better, or a public documentation page with a finals record	24
C	Everything else, including every team that never recorded a mission score	36

Tier counts from the census `S-CENSUS`.

Three features of the rule are deliberate. A team that never reached the finals cannot rise above C, whatever its SAR score, because qualifying is the first test. Documentation never buys a tier on its own: it can move a team one band, and only if that team is already close to the band above. And the inputs are all public, so any team's label can be recomputed from the score sheets.

The census

The table carries 70 records: the 68 census teams, plus two registered teams that hold no finals result and are profiled because they publish design material. It filters by tier and country, and searches team, university and country.

TEAM ↕	UNIVERSITY ↕	COUNTRY ↕	TIE
Mars Rover Design Team (Missouri S&T)	Missouri University of Science and Technology	United States	A
Team Mountaineers (WVU)	West Virginia University	United States	A
Michigan Mars Rover Team (MRover)	University of Michigan	United States	A
AGH Space Systems	AGH University of Krakow	Poland	A
Monash Nova Rover	Monash University	Australia	A
BYU Mars Rover	Brigham Young University	United States	A
UIU Mars Rover	United International University	Bangladesh	A
Northeastern University Mars Rover Team	Northeastern University	United States	A
Legendary Rover (PRz)	Rzeszow University of Technology	Poland	B
Yonder Dynamics (UC San Diego)	University of California, San Diego	United States	B
IMPULS	Kielce University of Technology	Poland	B
A.S.T.R.A. (Alabama-Huntsville)	University of Alabama in Huntsville	United States	B
BRACU Mongol-Tori	BRAC University	Bangladesh	B
Queen's Space Engineering Team (QSET)	Queen's University	Canada	B
Sooner Rover Team	University of Oklahoma	United States	B
ARES Project (Tohoku/Keio)	Multi-institution (Tohoku, Keio,…)	International (multi-country)	B
KNR Rover Team	Warsaw University of Technology	Poland	B
Quantum Robotics	Instituto Tecnológico y de Estudios Superiores de Monterrey	Mexico	B
MIST Mongol Barota	Military Institute of Science and Technology	Bangladesh	B
Toronto MetRobotics (TMR)	Toronto Metropolitan University	Canada	B
PCz Rover Team	Czestochowa University of Technology	Poland	B
MAVRIC (Iowa State)	Iowa State University	United States	B
Wisconsin Robotics	University of Wisconsin-Madison	United States	B
SKA Robotics	Warsaw University of Technology	Poland	B
Project Scorpio	Wroclaw University of Science and Technology	Poland	B
New Haven Robotics	University of New Haven	United States	B
DAM Robotics (Oregon State)	Oregon State University	United States	B
Robotics for Space Exploration (RSX)	University of Toronto	Canada	B
SC Robotics	Saddleback College	United States	B
Mars Rover Manipal	Manipal Institute of Technology	India	B
IITB Mars Rover Team	Indian Institute of Technology Bombay	India	B
Husky Robotics	University of Washington	United States	B
Mars Rover UdeG Space	Universidad de Guadalajara	Mexico	C
ITU Rover Team	Istanbul Technical University	Turkey	C
YILDIZ ROVER	Yıldız Technical University	Turkey	C
Zeus (Université de Sherbrooke)	University of Sherbrooke	Canada	C
Project Altair (IUT)	Islamic University of Technology	Bangladesh	C
RoverOva (VSB-TUO)	VSB - Technical University of Ostrava	Czechia	C
Binghamton University Rover Team	Binghamton University (SUNY)	United States	C
Project Kratos	BITS Pilani KK Birla Goa Campus	India	C
Ceres (Universidad de los Andes)	Universidad de los Andes	Colombia	C
Carleton Planetary Robotics Team (CPRT)	Carleton University	Canada	C
MR2 (KAIST)	KAIST	South Korea	C
Robocol (Universidad de los Andes)	Universidad de los Andes	Colombia	C

BILL-EE (Cal Poly Pomona)	California State Polytechnic University, Pomona	United States	C
Team Interplanetar (BUET)	Bangladesh University of Engineering and Technology (BUET)	Bangladesh	C
Team RoSE (Robotic Space Exploration)	University of Hawaii at Manoa	United States	C
Space Concordia (Concordia University)	Concordia University	Canada	C
Cornell Mars Rover	Cornell University	United States	C
Zenith Space (SeoulTech)	Seoul National University of Science and Technology	South Korea	C
RoboJackets (Georgia Tech)	Georgia Institute of Technology	United States	C
OzU Rover	Özyeğin University	Turkey	C
Inferno DTU	Delhi Technological University	India	C
Society of Martian Engineers - Team Sofia (USU)	Utah State University	United States	C
KARURA (international)	Multi-institution (Texas A&M + Japanese universities)	International (multi-country)	C
Team Anveshak (IIT Madras)	Indian Institute of Technology Madras	India	C
GTU Rover Team	Gebze Technical University	Turkey	C
UMD Loop	University of Maryland, College Park	United States	C
Ad Astra (Sairam)	Sairam Institutions	India	C
AAUB Rover71	Aviation and Aerospace University, Bangladesh	Bangladesh	C
NAFT Mars Rover	Multi-institution (Nagoya, Tokyo, Meijo, Chiba Tech)	Japan	C
Team DIANA (Politecnico di Torino)	Politecnico di Torino	Italy	C
Team RoverX	Vellore Institute of Technology	India	C
Team RUDRA - SRM Mars Rover	SRM Institute of Science and Technology	India	C
Pharaohs Team	Military Technical College	Egypt	C
AIUB Robotic Crew (ARC)	American International University-Bangladesh	Bangladesh	C
METU ROVER	Middle East Technical University	Turkey	C
UBC Rover	University of British Columbia	Canada	C
Space Exploration Alberta Robotics (SPEAR)	University of Alberta	Canada	C
UW Robotics / UWRT	University of Waterloo	Canada	C

Fields: team, university, country, tier, seasons in a score sheet, finals appearances, best finals rank and mean final total. The two registered teams carry a registration-only season record and no rank. Tier labels follow the rule above `S-CENSUS`.

The nineteen teams without a profile

Nineteen census teams carry no dossier. In almost every case the public record is one score-sheet line and a team website, with no design report, paper or code archive behind it. Eighteen of the nineteen hold at most one finals appearance; eight have at least one SAR-only season; five never recorded a mission score at all.

TEAM	COUNTRY	SEASONS IN THE SHEET	WHAT IS PUBLIC
Zeus (Université de Sherbrooke)	Canada	2022, 17th	Team site
RoverOva (VSB-TUO)	Czechia	2023, 18th	Team site
Project Kratos	India	2022, 20th; SAR-only 2023 and 2024	Team site
MR2 (KAIST)	South Korea	2026, 22nd	Documentation site
BILL-EE (Cal Poly Pomona)	United States	2024, 23rd	Team site
Zenith Space (SeoulTech)	South Korea	2026, 26th	Team site
RoboJackets (Georgia Tech)	United States	2024, 26th; 2025, 34th	Club site
OzU Rover	Türkiye	2023, 26th; SAR-only 2022	Documentation page
Inferno DTU	India	2022, 26th	Team site
Society of Martian Engineers, Team Sofia (USU)	United States	2025, 27th	No site listed
GTU Rover Team	Türkiye	2023, 31st; SAR-only 2022	Team site
AAUB Rover71	Bangladesh	2026, 34th	Social page
NAFT Mars Rover	Japan	2026, 35th	Project site
Team DIANA (Politecnico di Torino)	Italy	2026, 36th	Team site
Team RoverX	India	SAR-only 2023 and 2024	Team site
Team RUDRA (SRM Mars Rover)	India	SAR-only 2026	Team site
Pharaohs Team	Egypt	SAR-only 2022 and 2023	No site listed
AIUB Robotic Crew (ARC)	Bangladesh	SAR-only 2022	Club site
METU ROVER	Türkiye	SAR-only 2022	Team site

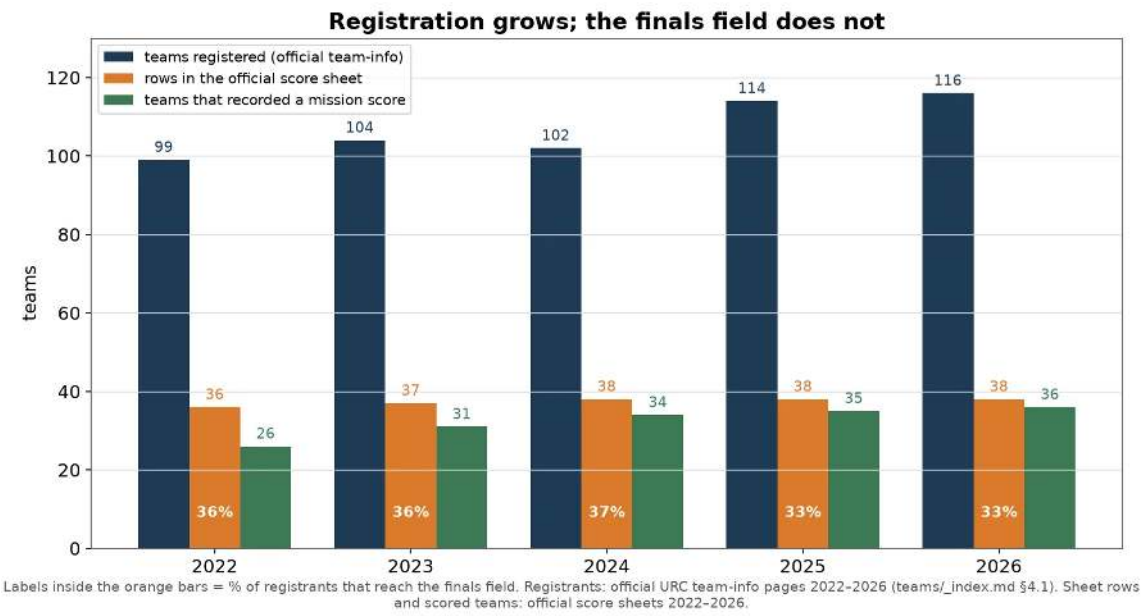
Seasons, ranks and the public record of each team from the census and the official season lists `S-CENSUS`. MR2 and OzU Rover each maintain a documentation page, but neither publishes a design report or a paper.

The two teams with a documentation page but no dossier, MR2 and OzU Rover, are the closest to a future profile. The seventeen others are thinner still: at most two seasons on a sheet, usually one, and a website that describes the club rather than the rover.

Teams that never reached the field

The census counts score-sheet rows, so it excludes teams that registered and never appeared in a sheet. That group is large and worth naming, because it carries the strongest evidence about where a season is actually lost.

Fifty-four teams appear on two or more official team lists and in no score sheet from 2022 to 2026 `S-CENSUS`. Twenty of them registered in three or more editions. SuRover and Titan Rover appear in all five; Sun Devil Robotics, Team Deimos, TrickFire Robotics and UNAM Space appear in four; Berkeley's BEAR, the Boston University Mars Rover Club, GU Robotics, Maverick Rover Team, Panther Robotics, PROJECT AV, Vicharaka, Boiler Robotics, CU Rover Team, Lehigh University Space Initiative, PolyOrbite, RIT Space Exploration, SJSU Robotics and uORover appear in three each. Several are large, long-running university clubs.



Registrations, score-sheet rows and rows with a mission score, by season. The gap between the first bar and the third is the desk: teams cut at the System Acceptance Review or lost before the field never appear in a score sheet `S-CENSUS`.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC13-SCORES` `S-URC19-SCORES` *Full resolution*

Two readings are consistent with the data. The score sheet may list the invited field, in which case the teams absent from it were cut or withdrew before travelling. Or the sheet may list only teams that recorded a finals result, in which case the SAR-only rows are finalists that withdrew. The 2023 season, where 37 sheet rows match 37 announced finalists, favours the first reading for that year; the other seasons are inferred `S-URC23-NEWS`.

Why a team is absent cannot be determined from the public record. The possible causes include a SAR that was not submitted, a SAR below the cutting line, a qualified team that could not travel, a withdrawal, or a name that matched nothing in the merge table. The list is therefore not a list of weak teams. It is the clearest available evidence that the binding constraint for a first-time entrant is the desk rather than the desert.

What the census cannot settle

- **The rank derivation.** A published rank column, or a finalist list for 2022 and 2024–2026, would replace the inferred ordering.
- **The Robocol and Ceres question.** Both are Universidad de los Andes teams and both appear in the sheets; no document states a rename, so they are counted as two teams `S-CENSUS`.
- **The documentation flag.** It records what was reachable on a given day, not a judgement of a team's engineering. A team can move one band on a second look.
- **The missing seasons.** URC held no finals in 2020 or 2021, and complete per-mission scores before 2022 were not recovered, so the census covers five seasons rather than the full history `S-URC-HISTORY`.

Where to go next

- The URC landscape places these 68 teams in the wider field.
- Results explorer, 2022–2026 works the 187 rows directly.
- Team browser filters the same records into cards.
- Team dossiers lists the 51 profiles by tier.

THE FIELD · CHAPTER 3.4

Sister competitions

Seven other competitions, five independent rover events and two run by NASA, publish engineering documents that are often deeper than a URC SAR submission. Their ranges are larger, their mass limits are looser and their radio links are shorter, so their rovers transfer badly. What transfers is method: how a requirements table is built, how a review is scored, how a mission is decomposed into scored behaviours.

50 kg

URC MASS LIMIT

Tightest of the rover events covered here

2027

S-URC27-RULES

75 kg

ERC MASS LIMIT

25 kg looser, with no stowed cube

2026

S-ERC26-RULES

1 km

URC RADIO RANGE

ERC's rover-to-mast link needs under 100 m

2027

S-URC27-RULES

400

ERC DOCUMENTATION POINTS

Of a 3,000-point total

2026

S-ERC26-RULES

Supplementary evidence only

Nothing on this page establishes or changes a URC 2027 requirement. The requirements live on the 155 engineering requirements page and trace to the rulebook and the Q&A. Sister-competition evidence is labelled by competition and used here for method and structure.

The competitions at a glance

COMPETITION	ORGANISER AND VENUE	LATEST FIELD	MISSION SET	POINTS	MAIN DIFFERENCE FROM URC 2027
URC 2027, for comparison	The Mars Society, Mars Desert Research Station, Utah	36 teams invited, 2–5 Jun 2027	Astrobiology, Delivery, Equipment Servicing, Autonomy, plus the SAR	500	The reference set for this whole site
ERC	European Space Foundation, AGH Kraków campus, Poland	25 finalists from more than 120 entries, Sep 2026	Science, navigation, maintenance, probing, presentation, documentation	3,000	Six times the points, 400 of them documentation, and rover weight is itself a scored category
CIRC	CSTAG, Drumheller, Alberta	27 teams, Aug 2026	Five short site tasks, one of them at night	Per-task rubrics, no fixed total	Task names and content change every year, and there is no science mission
Australian Rover Challenge	Adelaide University, Roseworthy, South Australia	17 teams, Mar 2026; 18–21 Mar 2027	Post-landing, space resources, excavation and construction, mapping and autonomy	600	The closest mass and size regime to URC, but lunar framing and drones prohibited
IRC	SPROS, Udupi, Karnataka, India	26 finalist places, Jan–Feb 2026	Astrobiology expedition, autonomous delivery, instrument maintenance	Per-mission rubrics, total not published	The closest mission-for-mission analogue, run as 30-minute missions
Anatolian Rover Challenge	Space Exploration Society, Türkiye	About 20 teams, Jul–Aug 2026	Science, night, collaboration, autonomous	500	A night mission, a multi-team mission and a mandated autonomy module
NASA HERC	NASA Marshall, Huntsville, Alabama	Two divisions, Apr 2026	Half-mile obstacle course with three sampling tasks	100	Half the points are earned in three design reviews before the course
NASA Lunabotics	NASA Kennedy, Florida	47 qualifying teams, 10 finalists, May 2026	Excavate regolith and construct a berm	Four weighted criteria	A construction robot judged on systems-engineering documents, not a rover

Competition facts from the official rules, handbooks and results pages of each event `S-ERC26-RULES` `S-CIRC25-RULES` `S-ARCH27-RULES` `S-IRC26-RULES` `S-ARC26-MANUAL` `S-HERC26-HANDBOOK` `S-LUNA-SITE` `S-LUNA26-NEWS` `S-URC27-RULES`.

The envelope, which is why a sister rover is not a URC rover

COMPETITION	MASS LIMIT	STOWED SIZE	RADIO RANGE	DOCUMENTATION GATE
URC 2027	50 kg per configuration, 70 kg fielded	1.2 m cube, any orientation, no disassembly	Up to 1 km, with non-line-of-sight sections	SAR at 100 points, top-36 gate
ERC 2026	75 kg including payload	No stowed-volume limit in the rules or the requirements sheet	Under 100 m from rover to mast	400 points of reports, video and forms
CIRC 2025	50 kg target, 1 % per kilogram bonus below, 70 kg hard stop	Fit through a 0.8 m door, in pieces if needed	Site scale, 5.8 GHz	Safety report due about seven weeks before the event
Australian Rover Challenge 2027	60 kg per task, 90 kg fielded, 10 % per kilogram over	1.6 × 1.6 m lander, 1.6 m tall	Team base station	CDR 30 points plus SAR 70 points
IRC 2026	65 kg per mission, 85 kg fielded	1.5 × 1.2 m, no height limit	Base station isolated from view	Written report and video at qualification
Anatolian Rover Challenge 2026	60 kg per mission, 80 kg fielded	Designed for a 1.5 m cube, airlock entry	Team base station	Design report and video presentation
NASA HERC 2026, remote-controlled division	27.2 kg	0.76 m cube	Commercial radio control	Three reviews, 50 of 100 points

Envelope values from the rulebooks, requirement sheets and handbooks of each event `S-ERC26-REQ` `S-CIRC25-RULES` `S-ARCH27-RULES` `S-IRC26-RULES` `S-ARC26-MANUAL` `S-HERC26-HANDBOOK` `S-URC27-RULES`.

URC holds the tightest combination of mass and stowed volume in this set, and it asks for an order of magnitude more radio range than ERC or CIRC. A rover designed to win ERC is likely to weigh 25 kg more than URC allows and to carry a link sized for 100 m. The HERC remote-controlled class is the only sister class lighter than a URC rover, and at 27.2 kg inside a 0.76 m cube it is a different robot entirely.

What transfers

DESIGN PROBLEM	WHAT SISTER TEAMS DID	STRENGTH	WHERE IT STOPS
Mass accounting to a hard per-configuration cap	ARCh weighs at every task and applies a 10 % per kilogram penalty; CIRC pays a 1 % per kilogram bonus below 50 kg	Direct, as a method	ARCh, IRC, the Anatolian challenge and ERC set their per-configuration limits 10 to 25 kg above URC's 50 kg, and only CIRC matches it, so only the accounting transfers
Stowage without disassembly	ARCh fits a 1.6 × 1.6 m lander; IRC allows 1.5 × 1.2 m with no height limit	Partial	URC's cube is smaller and symmetric on all three axes
Weigh-in inside the task clock	ARCh measures mass and power before every task and applies an efficiency multiplier	Direct	ARCh rewards low mass; URC only penalises mass above 50 kg, so chasing ARCh's curve can cost traction
Rocker-bogie and double-bogie suspension	Monash ran a four-wheel rocker with pivot steering; Team DIANA's FORTUNA uses double-bogie mobility; Scorpio uses a rocker-bogie-inspired layout	Direct, as an architecture	URC publishes no slope or step figure, so travel numbers and gear ratios do not transfer
Printed or elastomer wheels	Carleton's DUG uses six TPU wheels that deform to 50 % of diameter; Scorpio prints its own tyres	Partial	Utah rock is sharper than Drumheller soil, so tread that survives CIRC may chunk at MDRS
Modular payloads swapped between missions	Monash keeps one platform for two years and swaps arm and science modules; Scorpio Infinity is built for module access	Direct, as a concept	URC requires a swap inside 10 minutes, including a possible re-weigh
Six-axis arm with inverse kinematics	Carleton's SPIKE carries a 6-DoF arm with full IK; Team DIANA offers interchangeable end effectors	Direct, as an architecture	None has typed on a vertical keyboard or turned a cam-lock hose fitting, both URC 2027 tasks
Panel manipulation	ERC's maintenance panel has lever and rotary switches, an electromagnetic lock and a mains plug	Partial	URC's objects differ completely: a key transfer between two locks, a cam-lock fitting, a vertical keyboard
Two life-detection capabilities, one a bioassay	ERC's Astro-Bio sub-task is 300 points and teams run assays such as ninhydrin and luciferase	Partial	URC's two-capability rule with at least one bioassay is stricter, and it is new for 2027
A radio link sized for the course	ERC needs under 100 m from rover to mast; CIRC is site-scale	Negative	URC needs 1 km with terrain masking, so an ERC-grade link is a tenth of the requirement
Autonomous traverse with degraded positioning	ERC pays full marks only with no video feedback and penalises GNSS-only autonomy	Partial, as a skill superset	URC permits RTK GNSS with a team base station and only requires the second target to be out of radio line of sight
Hot, dusty, wet field conditions	ERC designs for 0 to 30 °C and low light; CIRC runs one task at night in Alberta weather	Partial	MDRS reaches 37.8 °C in full daylight, so a thermal design validated in Europe is under-tested

What each team did is taken from the competition documents and the teams' own pages `S-ERC26-RULES` `S-CIRC25-RULES` `S-ARCH27-RULES` `S-IRC26-RULES` `S-ARC26-MANUAL` `S-CPRT-ROVERS` `S-MONASH-ROVERS` `S-SCORPIO-PROJ` `S-DIANA-SITE`. *URC's silence on slope and step height comes from the Q&A* `S-URC27-QA`. *The strength judgement and the limits column are inferences from those sources.*

The five places sister evidence misleads

1. **Mass.** ERC allows 75 kg, IRC 65 kg per mission and 85 kg fielded, ARCh 60 and 90, the Anatolian challenge 60 and 80, and CIRC 50 kg as a target with a 70 kg hard stop. URC's 50 kg per configuration inside a 1.2 m cube is the tightest combination in the set. A rover copied from any of them is likely to fail the weigh-in or the cube.
2. **Radio range.** ERC's 100 m and CIRC's site links are an order of magnitude short of URC's 1 km requirement with non-line-of-sight sections. A link that works at ERC is not evidence of anything at URC.
3. **Autonomy framing.** ERC rewards harder autonomy than URC asks for: no video feedback, GNSS penalised. Building to ERC's constraint is safe but expensive, and building to URC's is sufficient.

4. **Environment.** ERC's 0 to 30 °C and night operation, CIRC's Alberta weather and the Anatolian night mission all describe a cooler, darker duty cycle than the Utah desert in June.
5. **Indicator colours.** URC uses red for autonomous, blue for teleoperation and flashing green for arrival. The Anatolian challenge uses green for teleoperation, yellow for autonomous and red to disarm. One lamp is enough, but its logic must be switchable, or judges will read it backwards.

Teams that compete in both

TEAM	URC RECORD	SISTER RECORD	WHAT THE DUAL CALENDAR APPEARS TO BUY	CAVEAT
Monash Nova Rover	2026 2nd, 412.41	Australian Rover Challenge 1st 2021–2023, 2nd 2024, 3rd 2025 and 2026	Two field campaigns a season and a two-year platform life with a mid-life overhaul	Its 2024 Autonomy score was 0, so dual experience does not prevent a subsystem collapse
AGH Space Systems	2024 champion, 2023 8th	ERC 1st 2022 and 2023, 2nd 2024 and 2026, 3rd 2025; Anatolian 2026 1st against 19 opponents	One rover programme serving four rule sets in a single season	AGH recorded no URC 2025 finals result
Carleton (CPRT)	2026 21st, 208.06	CIRC 1st 2026, 382.19; 3rd 2025	A CIRC-winning core platform and a tuned operations team	The same team is mid-field at URC, because CIRC's short tasks do not rehearse a 1 km delivery mission
Northeastern	2026 6th, 349.56	CIRC 2nd 2026, 256.21	A second field event with different rubrics, run in parallel	URC fell from 3rd to 6th while CIRC rose from 4th to 2nd
Michigan MRover	2026 5th, 375.89	CIRC 3rd 2026, 233.08	A compact, task-diverse rover and a strong software core	Michigan's CIRC points came from tasks that are not autonomy
DAM Robotics, Oregon State	2026 19th, 220.39	CIRC 1st 2025, 403.46; 7th 2026	Task speed and arm dexterity practice	The URC placing fell over the same period
KNR Rover Team	2026 9th, 274.91	ERC 9th 2025, 14th 2026	A sharp URC improvement in the same season as a solid ERC campaign	One season is not a trend; URC 2025 was 28th
Legendary Rover (PRz)	2026 10th, 261.39	Australian Rover Challenge 2nd 2026; CIRC 4th 2026	Three sister events in one season, the widest rule-set exposure in the field	The URC placing slipped from 5th to 10th in the same season
Mars Rover Manipal	2024 18th	IRC 2018 winner; ERC 10th in 2025 and 2026	The clearest IRC, then ERC, then URC pipeline	The URC results lag the sister results by several seasons

Sister results from the official results pages and team announcements `S-ERC25-RESULTS` `S-ERC26-SCORES` `S-CIRC25-RESULTS` `S-CIRC26-RESULTS` `S-ARCH25-REPORT` `S-ARCH26-NEWS` `S-ARC26-AGH` `S-MONASH-ARCHPAGE` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. *The "what it buys" column is an inference from the two records, not a claim by any team.*

The pattern in that table is consistent and sobering. The sister results and the URC results move independently: Carleton won CIRC and finished 21st at URC, DAM Robotics won CIRC in 2025 and finished 15th at URC, and Northeastern improved at CIRC while dropping at URC. Legendary Rover and DAM both slipped at URC in a season of sister success, and KNR is the counter-example, improving sharply at URC while its ERC placing fell. The skills that cross over are integration repetitions, operator practice and documentation discipline. The mission that does not cross over is Delivery, which has no equivalent in CIRC or the Australian challenge: 1 km, up to 60 minutes, with the second target out of radio line of sight.

Documents worth reading

These are the free, public documents with the most transferable structure.

DOCUMENT	WHAT IT CONTAINS	HOW IT HELPS A URC SAR
ERC 2026 Appendix 3, Requirements	71 numbered requirements, each with assumptions, a method of proof and a compliance column	The best free model for turning prose rules into a traceable requirements table
ERC 2026 Appendix 2, Scoring	Per-item rubrics for the reports, the video and every task	Shows how a mission decomposes into scored behaviours
ERC 2026 Rules, Revision 3	Task definitions, review schedule, autonomy point table, RF rules	A mature review cycle in one document
Australian Rover Challenge 2027 rules	Four tasks, 600 points, the 60 kg and 90 kg regime, weigh-in procedure	The closest structural analogue to URC's mass and penalty architecture
CIRC 2025 Rover Safety Requirements	Fusing, battery handling, E-stop, inspection and report deadlines	The most concrete safety-argument structure in the set
SPROS IRC 2026 Rulebook	Astrobiology, delivery and maintenance missions with base-station rules	The closest mission-for-mission analogue to URC
NASA HERC 2026 Handbook	Design review and readiness review criteria, obstacles, points	A complete review and readiness scoring system to plan a first SAR against
CMU Moon Miners design review	A student review package with real dates and feedback	What a review package looks like in practice

Availability of each document was checked in October 2026 `S-ERC26-REQ` `S-ERC26-SCORING` `S-ERC26-RULES` `S-ARCH27-RULES` `S-CIRC-SAFETY` `S-IRC26-RULES` `S-HERC26-HANDBOOK` `S-CMU-CDR`.

What is not settled

- **CIRC's edition count.** The site calls both 2025 and 2026 the ninth competition, so no ordinal is quoted here `S-CIRC25-RULES`.
- **IRC's total score.** The rulebook's scoring table is an image, so per-mission behaviour is usable and the total is not `S-IRC26-RULES`.
- **The Anatolian challenge's full scoreboard.** It sits behind a download button, so only AGH's first place against 19 opponents is confirmed `S-ARC26-AGH`.
- **The HERC human-powered champion.** The 2026 human-powered winner was not located; the confirmed results are UAH's third place in that division and first place in the remote-controlled division `S-HERC26-UAH`.
- **A URC requirement from any of this.** There is none. A design decision that rests only on a sister-competition source is an option awaiting a check against the URC requirements.

Where to go next

- Rules, limits and the field carries the URC envelope these competitions are measured against.
- Where the competition is going reads the sister events as evidence about task direction.
- Monash, AGH, Carleton, KNR and Michigan carry the dual-competition records in detail.
- The field places all of this back in the URC results.

SUBSYSTEMS · CHAPTER 4

Subsystem studies: scope and evidence

The twelve pages in this part answer one question each: for a given subsystem, what are the options and which one fits a 50 kg rover. A team dossier records what one team built and what happened. A subsystem study aggregates every team that could be documented and adds the component-level evidence that no single team publishes. Each study states the requirement, maps the design space, compares the options on benefits, drawbacks, relative cost and integration risk, and closes with a shortlist and the tests that remain open.

12

SUBSYSTEM STUDIES

From mobility to astronaut interaction

2027

68

TEAMS IN THE CENSUS

Tier A 8, Tier B 24, Tier C 36

2022–2026

`S-CENSUS`

51

TEAM DOSSIERS

Eight Tier A teams in depth

2022–2026

`S-CENSUS`

155

TRACED REQUIREMENTS

Each row tied to a rule or Q&A clause

2027

`S-URC27-REQ`

Key takeaways

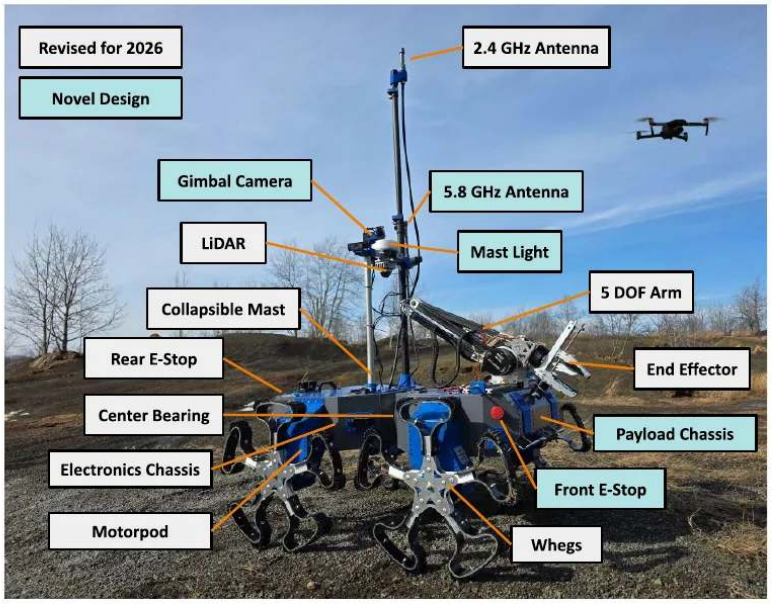
- A dossier answers what a team built. A study answers which option fits, and that question cannot be answered from a sample of one.
- Every comparison row carries benefits, drawbacks, relative cost, complexity and integration risk, plus the evidence behind the row.
- Relative cost runs from `$` to `$$$$$` inside one table only. A `$$$` antenna and a `$$$` gearbox are both average for their category, not the same price.
- The evidence is uneven. Where nothing is published, the row says so, and the shortlist states what evidence would change the recommendation.

Why subsystem studies, after team dossiers

A team dossier answers what a team built and what happened. A subsystem study answers what the options are and which one fits. The second question cannot be answered by describing one team, because a single team is a sample of one. The twelve studies therefore aggregate every team that could be documented: 68 in the census, with the eight Tier A teams studied in depth, plus sister-competition and industry analogues where URC evidence is thin `S-CENSUS`. They then add component-level evidence that no single team publishes, from vendor datasheets to the soil mechanics that decide whether a wheel grips.

The structure is fixed, which is what makes the pages comparable. Twelve studies that each invented their own headings would be twelve essays. The same eight questions asked of every subsystem produce a set of tables that can be read against each other, and they are the reason the cross-subsystem analysis can be written at all.

PART OF A STUDY	WHAT IT ANSWERS
Purpose and requirements	The subsystem's job in each 2027 mission, traced to requirement identifiers
Design space	The architectures that exist, how they differ, and the variables that separate them
Comparison tables	Every option, with benefits, drawbacks, relative cost, integration risk and evidence
Cross-team patterns	What the top teams converge on, where they diverge, and what correlates with a score
Failure modes	The failures observed, their causes, and the design change that prevents a repeat
Interfaces	The inputs the subsystem needs and the outputs it owes the rest of the rover
Technical shortlist	The strongest option for each set of conditions, with no option removed for cost
Open problems	The tests that must settle each question, and the decision each one unblocks



One rover, twelve subsystems. Team Mountaineers labelled Deimos in its 2026 SAR: wheels, motor pods, electronics bay, payload chassis, mast, LiDAR, E-stops, both radio bands, the 5-DoF arm and the end effector, with the drone Phobos in the background. The labels are almost exactly the list of studies in this part.

West Virginia University Team Mountaineers, 2026. `S-WVU-SAR2026`

Full resolution

The twelve studies

The order is the reading order, and it runs from the subsystem every mission depends on to the one that exists only because the 2027 rules added it.

CHAPTER	STUDY	CORE QUESTION IT ANSWERS
4.1	Mobility	How does the rover cross a kilometre of sand, rock and steep loose slope without getting stuck or tipping?
4.2	Arm and end-effectors	How does the arm reach from the ground to 1.5 m and work keys, locks, latches, hoses and a 5 kg cache?
4.3	Science and astrobiology payload	How does the rover find and prove, or exclude, life with two methods, one a bioassay, inside 30 minutes?
4.4	Delivery, payload handling and drones	How does the rover find, grasp, carry and deliver a 5 kg object over rough ground, and is a drone worth it?
4.5	Communications and networking	How does the rover keep a real-time video and control link at 1 km, out of line of sight, inside a metal trailer, on a shared band?
4.6	Power, electrical architecture and safety	How does the rover power a 40-minute mission with peaks, and guarantee that one button removes all battery power?
4.7	Compute, sensors and perception	What is needed to see ArUco tags at 1 cm, a walking astronaut, terrain and 5 kg objects?
4.8	Software stack and autonomy	How does the rover know where it is, plan a route, follow a person and recover safely when something goes wrong?
4.9	Operator station and operations	How do six people with no view of the rover run a 30–40 minute mission without mistakes?
4.10	Systems engineering and integration	How does a 50 kg rover fit inside a 1.2 m cube and pass the System Acceptance Review?
4.11	Environment, dust, thermal and repair	How does the rover survive the Mars Desert Research Station and the journey there, and how is it repaired in the field?
4.12	Astronaut interaction and the EVA suit	What does the 2027 astronaut-assistance sub-mission demand, and how are the suit and the interaction built?

Delivery and drones is the exception. No comparable published account of the 2027 Delivery mission or its drones exists, so that page is a synthesis drawn from the rules, the Q&A, the team dossiers and the arm and astronaut studies. It says so in its opening paragraph, and its evidence is thinner than its siblings.

What a comparison row carries

The comparison table is the centre of each study, and it uses the same columns throughout. The column set is deliberately wider than a shopping list, because the option that wins on paper often loses on integration.

COLUMN	WHAT IT HOLDS
Option and key specifications	What the thing is, with the numbers that separate it from the next row
Teams and years using it	Who built it, in which years, with the census tier beside the name
Benefits	The capability it buys, in physical terms rather than adjectives
Drawbacks and costs	What it costs in mass, volume, time, calibration or maintenance
Relative cost	`$` to `$$$$$` within this table only, or a list price with currency and date
Complexity and integration risk	What can go wrong in integration, how many subsystems it drags in, whether a first-year team can bring it up in one season
Evidence and links	A source identifier, and where useful a figure, a frame or a video timestamp

> **Cost is one attribute among several**
>
> Relative cost is measured inside a single table and nowhere else. Nothing is removed from a comparison because it is expensive, and no row reasons about a team's budget or funding. Where a list price was published it carries a currency and an access date; where it was not, the row says that a quote is required or that no price is published.

How strong the evidence is

URC teams publish unevenly, and a study that hid the difference would be less useful than one that shows it. Every row carries one of four grades.

GRADE	WHAT IT RESTS ON
Strong	A team's own report, repository or video, plus competition results, plus a manufacturer datasheet
Moderate	A video observation plus a team statement, or a single primary document
Weak	One photograph, one frame, or an analogue from another competition
Absent	Nothing published; the row records where it was looked for

Two consequences follow. First, where the only evidence is a video frame, the study records what was seen and treats a part identification as an inference rather than a fact: an anodised can of about 40 mm with a planetary front cover is consistent with a Maxon GP42, which is not the same statement as knowing that it is one. Second, the shortlist states its confidence and what evidence would overturn it. A recommendation that cannot be tested is a preference.

Seventeen of the 68 census teams publish too little to support a dossier, and some subsystems are documented by only a handful of teams. The arm is the clearest case: only six of the roughly 36 teams that scored in 2026 publish enough to reconstruct an arm at all `S-CENSUS`.

Three reading paths

The twelve studies are not independent, and the most important decisions cut across them. Three paths are worth following in order.

1. **The mass-limited path.** Power and energy mass, then mobility mass, then arm mass, then packaging in the 1.2 m cube, then mast and cable mass. This path establishes what 50 kg actually buys, and it ends at the cube rather than at a component.
2. **The task-limited path.** Arm dexterity, then the science payload, then delivery, then astronaut interaction. This path establishes what the rover can do with its hands, and it is the one that decides whether the manipulator is a light chain of links or a 14 kg machine.
3. **The reliability path.** Communications, then power and the E-stop, then perception and software failure handling, then operations, then environment and field repair. This path addresses why teams lose missions, and it is the path that converts a working rover into a scoring one.

Mass and power works the first path with numbers. Interfaces and couplings shows where the paths collide.

What the studies leave open

Three limits run across the whole part. The 2027 per-task weights are not published for Astrobiology, Delivery or Equipment Servicing, so every statement about the value of one task against another is a prediction S-URC27-RULES §1.b §1.c §1.d. Delivery has no dedicated study, for the reason given above. And most teams publish no rover mass or centre of gravity, which is why the mass figures in these pages carry ranges rather than single values.

The studies are also dated. Every claim carries the year it belongs to, because a 2024 design decision can be made wrong by a 2027 rule change: the two life-detection capabilities in the science payload, the key and lock transfer in the arm task set, and the removal of the marked-path task from Delivery are all changes that invalidate a design copied from an earlier year S-URC27-RULES §1.b.v §1.d.ii §1.c.iii.

Further reading

- The 2027 competition states the four missions and the scoring, which is the context every study assumes.
- 155 engineering requirements is the traced list the studies work from.
- The 68-team census gives the tiers and the results behind every “teams and years” column.
- Team dossiers hold the per-team detail the studies aggregate.
- Mission backfit turns the subsystem shortlists into a decision for each 2027 mission.

SUBSYSTEMS · CHAPTER 4.1

Mobility

Mobility is the subsystem every mission depends on. It carries the rover and its payload across up to a kilometre of soft sand, gravel, boulders and loose slopes, stops within a metre of an AR target, and does all of it inside a 1.2 m cube and 50 kg at every weigh-in. The engineering surprise is that torque is rarely the binding constraint. Traction, packaging and the clock are.

1 km

COURSE RANGE

Delivery and Autonomy, from the C2 station

2027

S-URC27-RULES

7.72 N·m

TORQUE PER WHEEL AT 30°

45 kg rover, 250 mm wheel, six wheels

Worked example

S-UMD-TERRAMECHANICS

14-29 kg

MOBILITY MASS

28-58 % of the 50 kg limit

First-build budget

S-MST-SITE-ATHENA

1.2 m

CUBE AT WEIGH-IN

Includes wheels, mast and every protrusion

2027

S-URC27-RULES

Key takeaways

- The 30° climb needs **7.72 N·m per wheel** and about **371 W** electrical, but it also needs a traction coefficient of **0.727**, just above the **0.70** limit of a smooth wheel on cohesionless soil. Grousers and ground pressure decide the climb, not motor size.
- A skid-steer spin turn demands **7.6 N·m per wheel**, about as much as the climb, and dissipates roughly **1 kJ**. Three-point turns preserve both tread and energy.
- At 250 mm wheels, a 0.9 m wheelbase and a 1.0 m track, the rover is already **1.15 m long and 1.20 m wide** before the arm, mast or antenna. The cube is a packaging requirement, not a packing problem.
- Speed is a mission requirement because the clock is fixed. The practical target is **1.0-1.5 m/s on flat gravel** and **0.4-0.6 m/s on a 25-30° slope**.

What the subsystem must do

The 2027 rules set the mobility envelope in five places. The rover must traverse "soft sandy areas, gravel, rough stony areas, rock and boulder fields, vertical drops and steep loosely consolidated slopes" up to **1 km** from the start gate S-URC27-RULES §1.c.ii. It must find a navigable route up a hill where "not all approach directions will be navigable" S-URC27-RULES §1.e.xv. It must stop within **1 m** of an AR target S-URC27-RULES §1.e.xvii. It must fit inside **1.2 m** on every axis and stay at or below **50 kg** in each deployed configuration S-URC27-RULES §2.a.ii. And the E-stop must stop movement and remove battery power, which puts a brake or a self-locking ratio in the drivetrain rather than relying on a zero command S-URC27-RULES §2.a.v.

The rules do not publish a step height or a slope limit for the course. URC declined to give a number when asked S-URC27-QA D1. The design therefore has to choose its own obstacle set and measure against it, rather than design to a published worst case.

The numbers that size the design

The design-speed, climb-torque, skid-scrub, spin-energy and fabrication rows are engineering estimates from the stated parameters. The course range and traction rows rest on the rules and the cited soil mechanics. The packaging and mass rows are budget arithmetic.

QUANTITY	VALUE	BASIS
Course range	up to 1 km from C2; total path may exceed 1 km	Delivery and Autonomy S-URC27-QA RQ13
Design speed	1.0-1.5 m/s on flat gravel; 0.4-0.6 m/s on 25-30°	Estimate: a 2 km Delivery round trip is 33 min at 1.0 m/s, 67 min at 0.5 m/s
30° climb torque	7.72 N·m per wheel at the gearbox output; about 371 W total	Estimate from 45 kg, 250 mm wheel, efficiency 0.75, rolling resistance 0.15
Traction at 30°	required coefficient 0.727 against a smooth-wheel limit near 0.70	tan 35° = 0.70 for cohesionless soil S-UMD-TERRAMECHANICS
Skid-spin scrub	164 N·m yaw moment; 7.6 N·m per wheel	Estimate from 0.9 m track, 0.85 m wheelbase, friction 0.6
Spin energy	about 1.0 kJ per 360° spin	Estimate; equal to climbing about 2.3 m
Packaging at 250 mm wheels	length about 1.15 m, width about 1.20 m before arm or mast	0.9 m wheelbase, 1.0 m track, 80 mm tyre, 60 mm hub overhang
Mobility mass	14-29 kg, or 28-58 % of the 50 kg limit	Frame 5-11 kg; wheels and drive 7-14 kg; controllers 2-4 kg
Fabrication effort	250-400 person-hours for a six-wheel rover	Estimate; 150-250 hours for a four-wheel rocker

The inversion that matters: the AK80-9 (9:1) on a 250 mm wheel already gives **5.1 m/s** at rated speed S-CUBEMARS-AK80-9. Speed is not what the motor lacks. Torque margin and traction are.

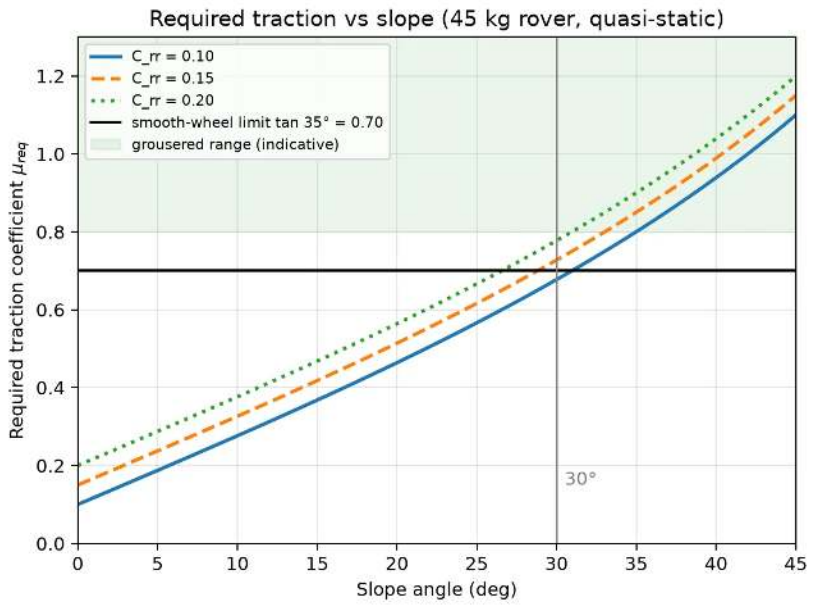


Athena, the 2026 championship rover, on the boulder terrain the 2027 Delivery and Autonomy missions use. Six wheels, printed TPU treads and a compact suspension that no longer folds to fit the cube.

Missouri S&T Mars Rover Design Team, 2026. ***Full resolution***

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Wheel count and layout	4 · 6 · wheel-legs	6 for redundancy and boulder bridging; 4 when mass and the cube bind, as on Monash and UIU in 2026
Suspension	6-wheel rocker-bogie · 4-wheel rocker · double-lambda · articulated body · independent double wishbone · rigid	6-wheel rocker-bogie with a differential bar: the largest top-team evidence base and no control law
Steering	skid · pivot · swerve	Skid steer: zero steering hardware; pivot or swerve only when fine positioning is the bottleneck
Wheel	printed TPU · moulded polyurethane · composite · pneumatic	250-300 mm printed TPU with 10-15 mm V-chevron grousers and an exchangeable tread
Drive	integrated BLDC · BLDC with planetary or cycloidal reduction · hub motor · brushed	48 V integrated actuators if machining is weak; in-house cycloidal or planetary if it is strong
Frame	laser-cut 6061 · welded torsion box · T-slot · carbon fibre	Laser-cut 6061 sheet with bolted or waterjet plate nodes; carbon fibre only with composites experience
Folding	fixed compact · folding	Fixed if it fits, as on Missouri S&T 2026; folding adds joints and failure modes, as on BYU and WVU



The traction coefficient a 45 kg rover needs at constant speed, against the smooth-wheel limit. At 30° the requirement is 0.73, above the 0.70 limit, so grousers or a larger contact patch are part of the climb design.

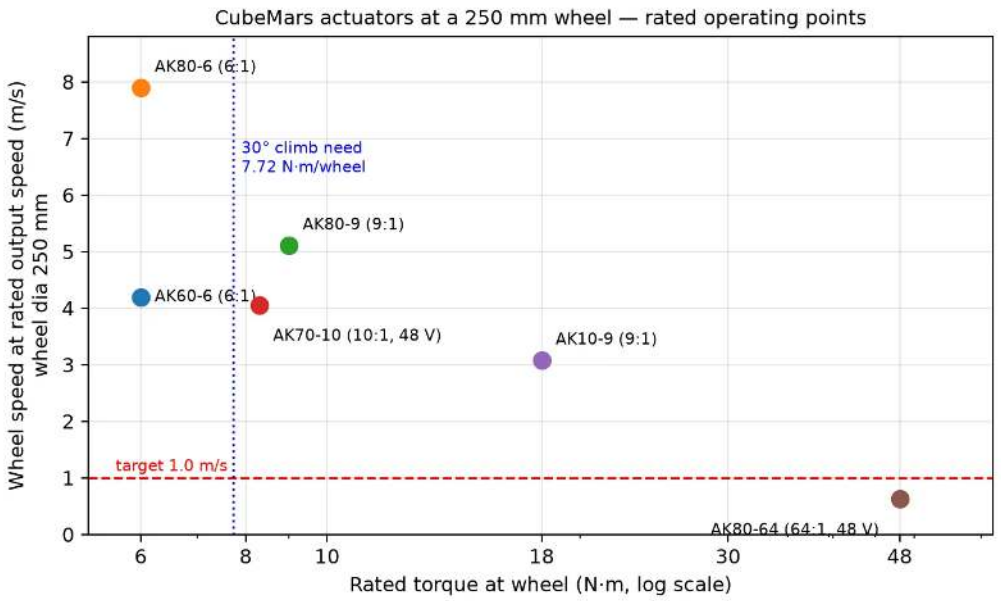
Robotics Club @ Mississippi State University, 2026. Original analysis.

S-UMD-TERRAMECHANICS *Full resolution*

Traction and scrub, not torque, size the drivetrain

The worked 45 kg, six-wheel, 250 mm example is the calculation a first-time team should reproduce before buying motors. A 30° loose slope needs **7.72 N·m per wheel** and about **371 W** electrical, but the required traction coefficient is **0.727**, just above the **0.70** smooth-wheel limit of a cohesionless soil with a 35° friction angle `S-UMD-TERRAMECHANICS`. The motor has torque to spare and the wheels still spin. That is why grousers, a wider contact patch and lower ground pressure matter, and why a larger motor does not fix a loose 30° slope. At 45° the requirement is 1.15, which traction alone cannot deliver. The 2027 answer is route choice, not a gearbox.

Skid steering is the second surprise. A spin turn about the centre produces a **164 N·m** yaw moment and demands **7.6 N·m per wheel**, more than the 30° climb, and dissipates about **1.0 kJ**, the energy of climbing 2.3 m. Skid steer is still right for a first build because the hardware is simple and the wear is gradual. Operators should three-point turn rather than spin, and the tread should have a few millimetres to lose.



Rated torque and speed for six integrated actuators at a 250 mm wheel. Speed is ample for every option. The binding question is rated torque margin at the 7.72 N·m climb requirement.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-CUBEMARS-AK80-9 *Full resolution*

The 1.2 m box and the speed arithmetic

At 250 mm wheels, a 0.9 m wheelbase, a 1.0 m track and 80 mm tyres, the rover is already **1.15 m long and 1.20 m wide** before the arm, mast or antenna. That is why Missouri S&T shrank its suspension until folding was unnecessary, and why Monash packages at about 1.2 × 1.1 × 0.8 m excluding the mast. BYU and WVU fold. The box allows any orientation but not disassembly `S-URC27-RULES §2.a.ii`, and a diagonal does not rescue a bounding box that exceeds 1.2 m on an axis.

Speed is a mission requirement because the clock is fixed. Delivery is 30-60 minutes over a 1 km course. Astrobiology is 20-30 minutes within 0.5 km, and two sites are easily 1 km of driving. At 0.5 m/s the 2 km Delivery drive alone is 67 minutes, longer than the mission. The practical target is **1.0-1.5 m/s on flat gravel and 0.4-0.6 m/s on a 25-30° slope**, not as fast as possible. Monash quotes 3.6 m/s with 32 N·m continuous at the wheel, and Michigan designs for 30° `S-MONASH-2026-SAR` `S-UMICH-2026-SAR`.

Video available in the web edition.

Monash Nova Rover, URC SAR 2026

Monash Nova Rover

Watch on YouTube

Interfaces that bite

Mobility is coupled to almost every other subsystem, and four couplings decide the architecture:

- **Arm mass and reach against the wheel diameter.** A longer arm needs a stiffer deck; a larger wheel raises the deck and changes the arm's working height. The arm mount and the suspension cannot be designed independently.
- **Mast height against the cube.** The mast sets the radio horizon and the camera height, and it is the item most likely to break the 1.2 m envelope when stowed.
- **Wheel diameter against torque and packaging.** A larger wheel raises ground clearance and speed but increases the wheelbase, the packaging length and the motor torque demand.
- **Battery mass against the mass budget.** The pack is the largest single mass that can move without changing the structure, and it sits low, which helps the tip angle.

The 2027 rules also require the rover to stop within 1 m of an AR target. That is a mobility accuracy requirement at the end of a drive, not only a perception requirement.

Failure modes

The failure record is consistent across teams. Suspension joints and mounts fail under shock, printed treads wear through or delaminate, and a single motor or controller loss turns a six-wheel rover into a five-wheel rover with a dragging corner. Wheels that are too stiff transmit shock into the frame; wheels that are too soft roll off the bead or tear at the hub. The cheapest mitigations are a mechanical fuse in the wheel mount, a tread that can be replaced in the field, and a drivetrain that can isolate one corner.

The 2027 environment adds two more: dust that enters sealed bearings through the shaft, and transport shock into a crate that is also the weigh-in box. Both are testable before the season starts.

Shortlist for a first build

ITEM	CHOICE	WHY
Layout	Six wheels, rocker-bogie, differential bar	Redundancy and boulder bridging with no control law
Steering	Skid	No steering hardware, no extra actuators, simple failure modes
Wheel	250-300 mm printed TPU, 10-15 mm V-chevron grousers, exchangeable tread	Traction on loose soil; replaceable wear item
Drive	48 V integrated BLDC, 9:1 class, rated torque at least 9 N·m at the wheel	The 30° climb needs 7.72 N·m; 9 N·m leaves a small margin
Frame	Laser-cut 6061 sheet, bolted nodes	Predictable stiffness, repairable, no weld distortion
Suspension travel	60-80 mm at the wheel	Enough for rocks without a complex linkage
Mass target	16-20 kg for the mobility system	Leaves room for arm, science, power and compute inside 45 kg

Open problems

- **What is the real step height?** The rules do not publish one. A measured obstacle set from the 2026 course video and a local rock pile is the only defensible basis.
- **What is the soil strength at MDRS?** The traction calculation uses a 35° friction angle and a 0.15 rolling resistance. Both need a field measurement.
- **How much does the tread wear in 2 km?** A 20 km test on gravel and sand gives the answer; a 200 m test does not.
- **What is the tip angle with the arm extended?** The static tip angle with the arm stowed is not the worst case.
- **Can one corner fail and still drive?** A five-wheel configuration with one dragging wheel needs a test, not an assumption.

Further reading

- Missouri S&T, Monash, WVU and Michigan carry the strongest mobility evidence.
- Mass and power places the mobility mass in the full 50 kg budget.
- Interfaces and couplings works the wheel-diameter and arm-mass chains with numbers.

- Mission backfit maps the mobility requirements to each 2027 mission.

SUBSYSTEMS · CHAPTER 4.2

Arm and end-effectors

The arm is the only subsystem that touches the lander. In 2027 it must perform nine physically different actions, from pressing a 19.05 mm key to lifting a 5 kg cache and picking a 0.91 kg hammer off the ground, between 1.5 m height and the ground. It must do that while the whole rover folds inside a 1.2 m cube. Two numbers decide the architecture more than any other: 78.3 N·m at the shoulder, and 4.4–5.2 mm of backlash at the tip.

1.5 m

WORKSPACE, GROUND TO PANEL

Stowed inside a 1.2 m cube

2027

S-URC27-RULES

78.3 N·m

SHOULDER TORQUE, STATIC

5 kg cache at 1.05 m, plus a 6.5 kg arm

Worked example

4.4–5.2 mm

TIP ERROR FROM BACKLASH

15–18 arcmin at one metre of reach

Vendor data

S-CUBEMARS-AK

29.1

MEAN EQUIPMENT SERVICING SCORE

Median 25, five zeros, best 100

2026

S-URC26-RESULTS

Key takeaways

- **Nine actions, one arm** S-URC27-RULES §1.d.ii. The 2027 task set added the key and lock transfer and deleted the 2026 USB-C insertion, so a design frozen to 2026 tasking is already wrong.
- **Size the shoulder for 78.3 N·m static**, 117 N·m dynamic and 176 N·m with shock. No integrated actuator reaches that at one metre, which is why every high-scoring arm has a high-ratio reducer, a brake or a counterbalance at the shoulder.
- **Backlash, not encoder resolution, decides typing.** A 14-bit encoder is worth 0.38 mm per metre. A 15–18 arcmin planetary joint loses 4.4–5.2 mm per metre, the whole typing tolerance, before any vision error S-CUBEMARS-AK.
- **The camera belongs on the wrist.** A 720p sensor with a 70° lens resolves the 2 cm keyboard tags at 0.91 m and the 1 cm key tags only at 0.46 m.

What the subsystem must do

In 2027 the arm covers nine physically different actions: press a 19.05 mm key, turn a key in a tubular lock, pull a drawer, undo a latch, push a 1.5-inch cam-lock hose on, turn a quarter-turn valve, work buttons, switches and knobs, lift a cache under 5 kg by a handle of at least 10 cm, and, in Autonomy, pick a 0.91 kg rock-pick hammer off the ground S-URC27-RULES §1.d.ii §1.e.viii–x. The workspace runs from **1.5 m to the ground** S-URC27-RULES §1.d.i, while the whole rover must still fold inside the **1.2 m cube** S-URC27-RULES §2.a.ii.

Those two facts conflict. An arm with 1.5 m of reach cannot stow straight out, and the ground reach has to clear the rover's own bumper, so the shoulder sits near the front face of the chassis, on a turret that swings outboard, or on a deck high enough to point down over the nose. The three traced requirements behind that geometry are REQ-ARM-001 for the reach envelope, REQ-ARM-011 for autonomous typing and REQ-ARM-018 for the stow box. The tolerances below are inferences from the hardware geometry except where a source is given; they matter because the field failure is a one to five millimetre alignment error, not a lack of strength.

TASK	WHAT THE ARM MUST PHYSICALLY DO	TOLERANCE OR FORCE
Sample tube into cache, close the lid, turn the lock	Pinch, axial insert, two-finger twist	±3–5 mm, insertion under 5 N
Lift the cache	Power grasp on a handle of at least 10 cm	±20 mm, load under 5 kg
Drawer	Hook or pinch pull, place, push closed	±10 mm; pull force not published
Latch and hinged panel	Lever the latch, then a broad push on the panel	±5 mm
Autonomous typing	Single-point press, retreat, repeat	±3 mm design, 0.45–5 N S-URC27-TASKHW
Key in a tubular lock	Pinch the bow, extract, align the keyway, insert, turn	±1–2 mm at the keyway
Cam-lock hose	Power grasp on the fitting, axial push of 15–25 mm	±5 mm; push force not published
Quarter-turn valve, buttons, switches, knobs	Grasp and rotate 90°, or push and toggle	±5 mm, 1–10 N
Rock-pick hammer	Ground pick of a 0.91 kg, ≈300 × 170 × 25 mm steel hammer, retained while driving	±20 mm

The numbers that size the design

The torque, backlash and camera figures are worked estimates from the stated parameters. The task geometry and the mass range rest on the rulebook, the vendor datasheets and the two published arm masses.

QUANTITY	VALUE	BASIS
Reach	ground to at least 1.5 m, at a 0.5–0.8 m standoff	Panel height is in the rules; the standoff is a practical estimate
Cache load	under 5 kg, by a handle of at least 10 cm and no more than 5 cm diameter	S-URC27-RULES §1.d.ii
Shoulder torque, static	78.3 N·m	5 kg at 1.05 m plus 6.5 kg at 0.42 m
Shoulder with factors	117 N·m dynamic, 176 N·m shock	×1.5 twice
Typing geometry	19.05 mm pitch, about 0.45 N actuation, about 4 mm travel	Redragon K552, ANSI standard S-URC27-TASKHW
Typing tolerance	±3 mm design, ±5 mm geometric limit	13–14 mm cap top with a 3 mm tip
Joint backlash	15–18 arcmin = 0.25–0.30° = 4.4–5.2 mm per metre	CubeMars AK and Robotis XM540 quotations S-CUBEMARS-AK S-ROBOTIS-XM540
Camera standoff	2 cm tag at 0.91 m, 1 cm tag at 0.46 m, 0.30 m at 30 px	Pinhole model at 720p and 70°, by inference
Arm mass	4–14 kg, or 8–28 % of the 50 kg limit	RSX 4.42 kg to ARES 14 kg S-RSX-ARM-URDF S-ARES-PAPER-2025

The inversion that matters is the same one that governs the drivetrain. An AK80-9 turns at 570 rpm and supplies 22 N·m at peak S-CUBEMARS-AK , far more joint speed than any of the nine tasks needs. What the arm lacks is torque margin at the shoulder and lost motion at the wrist.

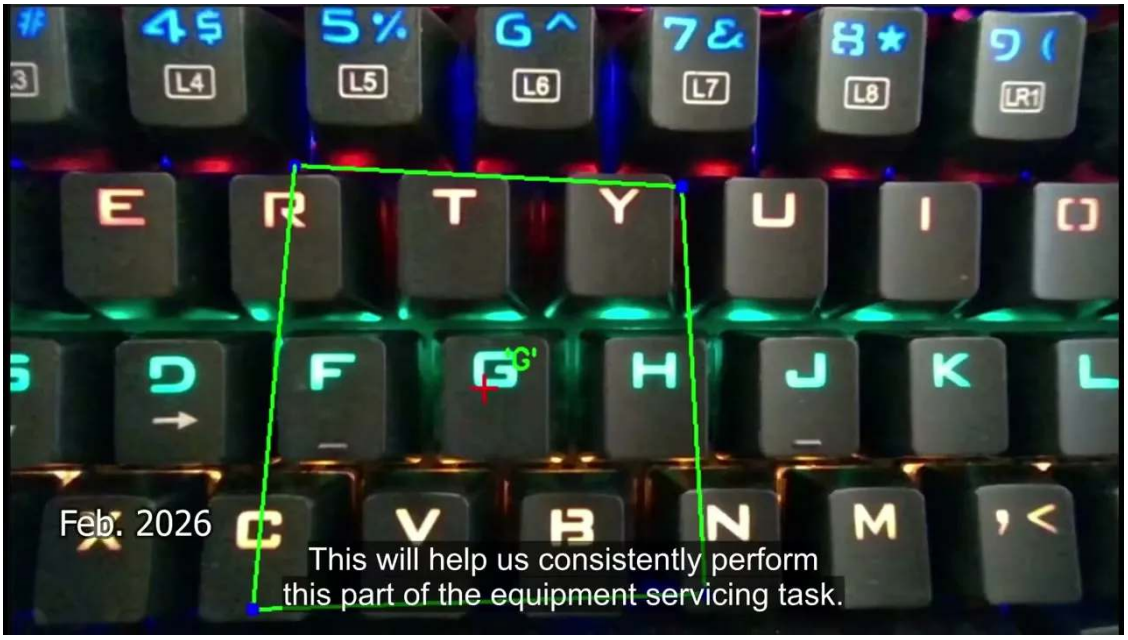
Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Kinematics	6R spherical wrist · 5R plus prismatic X · 5R · 4R · dual arm · arm on a lift	6R of 0.9–1.1 m with a wrist that faces the panel, or 5R plus one stiff prismatic axis parallel to the lander
Prismatic axis	none · 400 mm ball-screw stage	Worth adding when typing and key insertion are the priority: the best documented Equipment Servicing run used this layout
Joint actuation	integrated BLDC · harmonic or cycloid · custom BLDC · brushed with high ratio	Integrated actuators at elbow and wrist, with a harmonic, cycloid or 2:1 stage at the shoulder; a planetary reduction at the wrist blows the typing budget
Gravity holding	brake · spring counterbalance · gravity-compensating current	A fail-safe brake or a counterbalance at the shoulder; current-only holding heats the joint
End effector	two-finger parallel · dual head · swappable heads · adaptive or suction	Two-finger parallel, 100–110 mm opening, one motor, current-limited, with swappable tips
Sensing	incremental plus limit switches · absolute magnetic · multi-turn absolute · force and torque	An absolute encoder on every joint, current-based torque limits, and a wrist camera with ArUco detection

Backlash decides whether typing works

The encoder is fine. At 14 bits, one count is 0.022°, which is 0.38 mm per metre at the tip. The gearbox is not. A 15–18 arcmin planetary joint loses 4.4–5.2 mm per metre, which is the entire ±5 mm typing limit consumed by lost motion before any vision error `S-CUBEMARS-AK` `S-ROBOTIS-XM540`. A harmonic or cycloid reducer is roughly an order of magnitude better and costs three to ten times the mass and price per joint.

That is why the two teams that publish a typing method both close the loop with a camera rather than trusting joint counts. ARES uses a keyboard detector plus a single-point tip and the instruction “extend the tip and press the key”; BYU uses ViSP visual servoing `S-ARES-SAR26` `S-BYU-SAR26`. ARES placed 8th overall at URC 2026 with 324.59 points and 5th in Equipment Servicing with 62.11 `S-URC26-RESULTS`. The practical recipe is a two-stage loop: inverse kinematics to a pose 20–30 mm in front of the key, visual servoing over the last 30 mm from the 2 cm corner tags, then a short compliant press.



The arm’s eye view of the typing target, from BYU’s 2026 SAR: a full-height mechanical keyboard with a detected quad around E-R-T/Y and a crosshair on G. This is why the keyboard is a vision problem. A flat mini keyboard assumed from a diagram would put the tip in the wrong place.

BYU Mars Rover, 2026. *Original source* *S-BYU-SAR26* *Full resolution*

Video available in the web edition.

ARES Project, URC 2026 SAR: the two-stage typing method

ARES Project

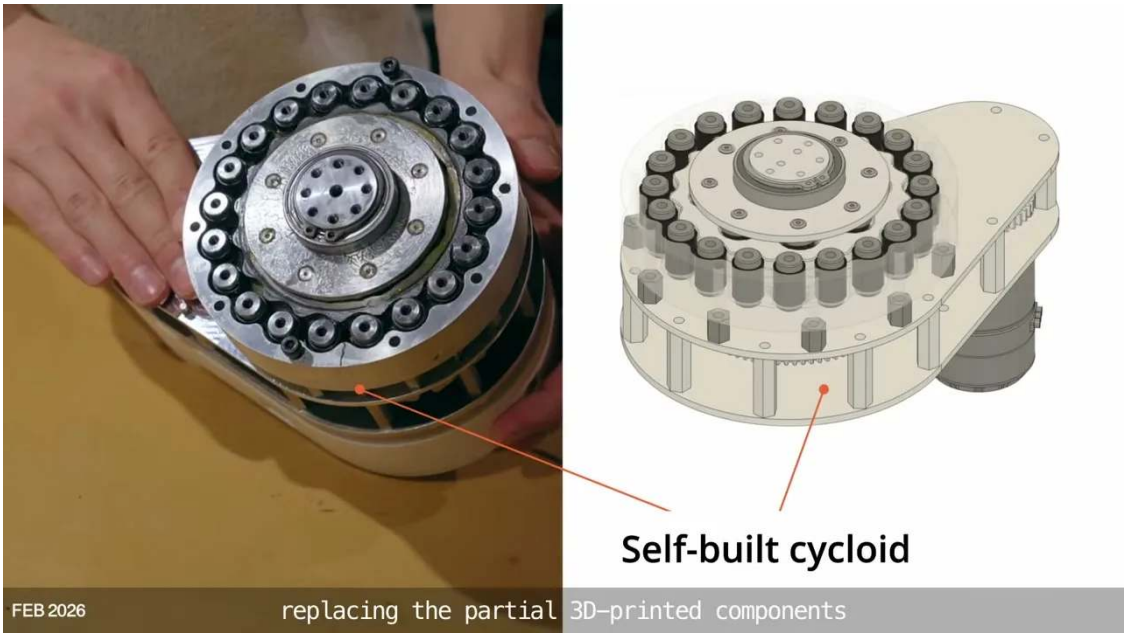
Watch on YouTube

A 720p camera with a 70° lens resolves the 2 cm keyboard tags at 20 px from 0.91 m, and the 1 cm key and keyhole tags only from 0.46 m, or 0.30 m for a comfortable 30 px. The key task therefore forces the camera within about 0.3–0.4 m of the keyway, which rules out a mast or chassis camera and means the wrist must come that close without the end effector occluding the tag. A 4K sensor buys a factor of about three in standoff and costs field of view, so the usual answer is a wide camera for the approach and a narrow one for the fine task.

The shoulder is a reducer problem

At the 5 kg cache limit and 1.05 m of reach the shoulder needs 78.3 N·m static, and the dynamic and shock factors take it to 117 N·m and 176 N·m. No single integrated actuator in the comparison class reaches that: the AK10-9 peaks at 53 N·m and needs a 2.2:1 stage, the AK70-9 at 29.2 N·m needs 4.0:1, and the AK80-9 at 22 N·m needs 5.3:1. ARES solves it with a self-built 1:38 cycloid rated 98 N·m in a 14 kg arm; a 100:1 harmonic is the bought alternative `S-ARES-PAPER-2025` `S-CUBEMARS-AK`.

Holding is the second half of the problem. A non-counterbalanced arm holds about 78 N·m continuously for tens of seconds during the drawer and cache tasks. An AK10-9 rated at 18 N·m cannot do that without a 4.3:1 reduction, and the joint will heat. The options are a high-ratio reducer, a normally engaged brake, or a spring counterbalance, and it is the brake or the counterbalance that removes the heat rather than the torque mode. The cache is held at only 0.4–0.6 m when the arm is folded, where 20–29 N·m is manageable.



A 1:38 cycloid built in-house, shown as a sectional drawing in ARES's 2026 SAR and rated at 98 N·m at the joint. This is how a 14 kg arm meets the 176 N·m shock case without buying a harmonic drive, and it is the clearest published answer to the shoulder problem.

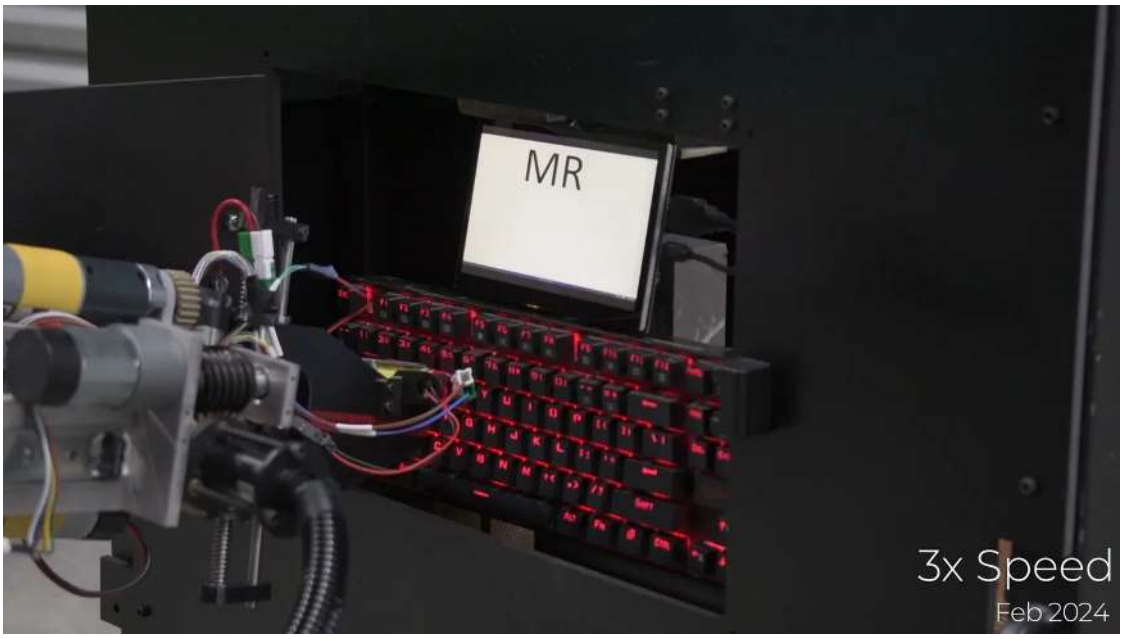
ARES Project, 2026. *Original source* `S-ARES-SAR26` `S-ARES-PAPER-2025` *Full resolution*

The two architectures that win

Six-revolute arms dominate the record, and the arithmetic is why: a shoulder, an elbow and a three-axis wrist can point a tip normal to any panel face and still fold the elbow inside the box. RSX publishes a 6-revolute chain with a 384 mm upper arm, a 305 mm forearm and 4.42 kg of total arm mass including the gripper and camera `S-RSX-ARM-URDF`. Monash's 0.50 m and 0.45 m links with a 0.10 m wrist come to 6.54 kg `S-NOVA-BULLETIN-2025`.

The competing layout keeps the rotary joints and adds one prismatic axis parallel to the lander face. Missouri S&T (MRDT) has run a hybrid Cartesian and rotary arm since 2024, with the linear axis taking the fine motion across the panel, and scored 94, 93 and 100 in Equipment Servicing across 2024–2026 `S-MRDT-SAR26`. ARES independently built a 400 mm ball-screw stage for the same reason, describing it as designed for tasks such as keyboard input, in a 1170 mm arm of 14 kg `S-ARES-PAPER-2025`.

The mechanisms differ and the capability does not: both buy backlash-free fine motion in the plane where typing, key insertion and drawer pulls happen. The strongest architectural signal in the record is convergence on that capability, not on a joint count.



The linear X rail, the rotary wrist and the small end effector of Missouri S&T's 2024 arm at the equipment-servicing panel, with the screen showing MRD. The rail supplies the precise in-plane motion that the rotary joints alone could not hold inside the backlash budget.

Missouri S&T Mars Rover Design Team (MRDT), 2024. *Full resolution*

What the field scores

Equipment Servicing is the arm's report card, and it is the weakest mission in the field. The mean has stayed between 21.8 and 35.3 for five years while single teams reach the high nineties, which says the difficulty is the manipulation itself rather than the task list.

YEAR	TEAMS SCORED	MEAN	MEDIAN	EXACT ZERO	BEST
2022	26	21.8	19.5	6	58
2023	31	35.3	35.0	7	93
2024	34	32.0	26.3	5	95
2025	35	31.5	27.0	4	93
2026	36	29.1	25.0	5	100

An arm team should expect a median near 25 and a realistic ceiling of 60 to 100 with a long tail of zeros S-URC22-RESULTS S-URC23-RESULTS S-URC24-RESULTS S-URC25-RESULTS S-URC26-RESULTS. The mission is scored task by task, and one early failure consumes the thirty minutes.

The team record supports an arm reconstruction for only six of the roughly 36 teams that scored in 2026, which is itself a finding: this is the least documented subsystem in the public record S-CENSUS.

Interfaces that bite

- **Mobility.** Reaction torque at the shoulder reaches 176 N·m, so the arm mounts on a structural node, not a suspension link. Its 4–14 kg at 0.4–0.6 m raises the centre of gravity and unloads the front wheels on a slope.
- **Power.** Peak current is 28 A per AK10-9 joint and 10–12 A per AK70-class joint. A stalled arm during a task is a brownout risk, so it needs its own fused rail with soft start and current telemetry.
- **Timing and compute.** Fine alignment wants under 150 ms of latency, and visual servoing over a 300 ms link is unstable, so the control loop belongs on the rover. The hands-off rule forbids human input, not offboard computation S-URC27-QA AUQ2.
- **Science.** A compliant wrist makes a 30 cm drill wander, so the drill mounts on the body rather than the wrist.
- **Safety.** A 5 kg cache held by a current-limited gripper falls when the E-stop removes all power. That has to be an accepted decision before the mission, not a surprise beside a judge.

Failure modes

Most of the failure record is mechanical rather than computational. A prismatic or vertical axis designed for static load fails under driving shock, a high-ratio reducer breaks after a collision with the lander, and a lost encoder zero ends the mission because re-homing is impossible hands-off.

MECHANISM	OBSERVED FAILURE	DESIGN CHANGE THAT PREVENTS IT
Prismatic axis and structure	Missouri S&T's 2024 Delivery run lost the vertical axes of the arm under driving loads	Design the vertical axis for shock, add tensioning and travel limits, and test the arm while driving
Alignment and compliance	Monash's 2022 panel slipped and closed twice, costing about nine minutes	A hook-shaped latch tip, passive compliance, and an approach along the panel edge rather than the face
Gravity and heat	A shoulder holding 78 N·m continuously heats and sags	A high-ratio reducer, a brake or a counterbalance, plus temperature monitoring
Lost zero	Incremental encoders with limit switches lose position at power-off	An absolute encoder on every joint, or a battery-backed multi-turn counter
Gripper slip	A friction grasp drops the cache, the hose or the hammer	A closed jaw with a mechanical stop, grip-force monitoring, and a secondary retention for the cache

Shortlist for a first build

ITEM	CHOICE	WHY
Architecture	6R serial, 0.9–1.1 m links, wrist that faces the panel	The field default, with closed-form inverse kinematics and the largest evidence base
Alternative	5 rotary joints plus a stiff prismatic axis parallel to the lander	The layout that produced 94, 93 and 100 in Equipment Servicing; rail backlash must stay under 0.5 mm
Shoulder	Integrated actuator with a harmonic, cycloid or 2:1 stage, or a brake	78 N·m static and 176 N·m shock are past any single integrated unit
Elbow and wrist	Integrated BLDC of the AK70 or AK80 class at the wrist, with a 2:1 stage or an AK10-9 at the elbow	The elbow needs 35.6 N·m static and 53 N·m with the dynamic factor; the wrist needs 7.4 N·m and 11 N·m
End effector and tips	Two-finger parallel, 100–110 mm opening, one motor, current-limited, plus a compliant typing tip, a chamfered key funnel and a hook	Covers the cache handle, the hammer shaft, the valve lever and the hose fitting; a hook cannot slip off a latch like a friction grasp
Sensing and mass	Absolute encoder on every joint, wrist camera, 6–8 kg for the arm and tips	Re-homing is impossible hands-off, and the mass leaves room inside 50 kg

Open problems

- **What force does the cam-lock hose need?** The 2027 gasket is cut to half thickness and no figure is published. A mock with the exact fitting and a load cell settles it.
- **What does the K552 need in force, travel and repeatability?** The 0.45 N figure is a standard switch number, not a measurement on the task keyboard, and a hundred automated presses at ±3 mm gives the real distribution. The chosen joint also has to be measured under load and at 38 °C, because a datasheet figure taken cold and unloaded is not the number that decides typing.
- **What is the arm's real mass and centre of gravity on the rover?** No 2026 team publishes both, and the slope analysis needs them.
- **What happens after one joint fails?** No team in the 2022–2026 record documents a mid-mission recovery, and injected encoder, limit-switch and CAN faults on the bench give the answer.

Further reading

- Missouri S&T, Monash, BYU, ARES, RSX and ROBOCOL carry the reconstructed arms.
- Equipment servicing sets the task sequence the arm has to survive.
- Mass and power places the arm in the full 50 kg budget.
- Interfaces and couplings works the arm-mass and backlash chains with numbers.

SUBSYSTEMS · CHAPTER 4.3

Science and astrobiology payload

The 2027 payload rule is the hard change of this rulebook. Onboard equipment must now include two life-detection capabilities, at least one of them a bioassay of a molecule that only biology makes. A camera, a spectrometer or a Raman probe no longer counts on its own, which puts at least one wet biochemical reaction, with reagent storage, fluid handling and waste containment, on the rover. All of it has to happen at two sites inside a 20–30 minute roving window.

2

LIFE-DETECTION CAPABILITIES

At least one a bioassay of a named biomolecule

2027

S-URC27-RULES

20–30 min

ROVING AND ANALYSIS WINDOW

Two sites within 0.5 km of C2

2027

S-URC27-RULES

≥5 g

SAMPLE FROM BELOW 10 CM

A 25–30 mm core, or a repeated scoop

2027

S-URC27-RULES

85

MEDIAN 2026 SCIENCE SCORE

13 teams at 90 or above

2026

S-URC26-SCORES

Key takeaways

- **The clause is not an extra instrument.** Q&A AQ9 accepts an assay of a molecule "currently known to be formed only by biological processes", and names proteins, sugars, chlorophyll, nucleic acids, ATP and carotenoids. An amino-acid assay alone is explicitly not enough S-URC27-QA AQ9.
- **A payload that scored 94 in 2026 can fail the 2027 clause unchanged.** WVU called ninhydrin its primary life-detection experiment, and under AQ9 that channel is insufficient on its own S-WVU-2026-SCIPLAN S-URC27-QA AQ9.
- **The clock binds, not the chemistry.** Sample acquisition and documentation take 10.5 of 25 minutes, so the slowest assay starts first, at site 1, and runs while the rover drives.
- **Depth is a record, not a claim.** AQ10 makes the depth of any life detection part of the mission, so the sample needs a graduated liner and a logged stroke, not a verbal assurance.

What the subsystem must do

Inside a **20–30 minute** window within **0.5 km** of C2, and using rover-relayed data only, the payload selects sites, documents **at least two**, analyses samples onboard, collects a sub-surface sample from **at least 10 cm** with **at least 5 g from below 10 cm**, seals a cache that does not spill when inverted, hands it over within **5 minutes**, and supports a post-mission **debrief** even if the rover failed S-URC27-RULES §1.b.ii–xi.

The three hardest constraints are traced as REQ-SCI-001 for the two capabilities, REQ-SCI-004 for depth and mass, and REQ-SCI-010 for the cache. Five clauses shape the design more than the mission flow does:

- **Analysis is onboard.** Samples may not be returned to the team for laboratory work; the cache handed to the judges is a sample, not an analysis S-URC27-RULES §1.b.iv.
- **The second capability is defined narrowly.** Two life-detection capabilities are required, at least one a bioassay S-URC27-RULES §1.b.v. Whether one combined assay can serve as both is not published.
- **Chemicals stay on the rover.** Water and every reaction product follow a no-spill policy in "a remote desert location with very limited water supplies", and hazardous chemicals need transport, use, disposal and accident plans approved in advance S-URC27-RULES §1.b.ix. Judges approve plans rather than chemicals, and anything above US DOT limits is not allowed S-URC27-QA AQ4.
- **Cleanliness is argued, not measured.** There is no swab test and no bioburden number, so the scored artefact is a written justification of the level claimed S-URC27-QA AQ7.
- **The Science Plan counts.** It is due **2027-05-21**, expands the System Acceptance Review version to about five pages, must explain the background science, and is factored into the score S-URC27-RULES S-URC27-QA §1.b.i AQ1.

What the rule actually changed

The change is not one more channel. It reclassifies most of the recent field.

DIMENSION	2023	2024–2026	2027
Mission name	Science	Science	Astrobiology
Minimum life-detection methods	one	one, plus an optional second capability	**two, at least one a bioassay**
Is a biomolecule required	no	no	**yes**
Stratigraphic profile	not required	required	**deleted**
Sub-surface sample	1–10 g, no depth	at least 10 cm, at least 5 g	at least 10 cm, at least 5 g
Cache	closed container	sealed, invertible	sealed, invertible, removable in 5 minutes
Sites evaluated	multiple	at least two	at least two

Three consequences follow from AQ9, which is binding and narrow. An amino-acid assay such as ninhydrin can support a result but cannot be the bioassay. A habitability measurement, whether pH, moisture, mineralogy or carbonates, is not life detection at all: "for URC life detection capability is required, not habitability" S-URC27-QA AQ9. And a physical detection such as chlorophyll autofluorescence names a listed biomolecule without clearly being an "assay", so it is best paired with a wet channel.

That is why Missouri S&T's 2022 payload, with three wet biomolecule tests, was closer to the 2027 rule than the Raman-only 2024 payload that replaced it. Their science score went 96, 79.17, 82, 96 and 93 across the transition, which is not evidence that spectroscopy helped S-MST-ITC-2022 S-MST-HYPERION-2024 S-MST-TALOS-2025. WVU's 2026 plan is the clearer warning: it scored 94 with ninhydrin as the primary life-detection experiment and resazurin for metabolic activity, and metabolism is not a molecule on the AQ9 list S-WVU-2026-SCIPLAN.

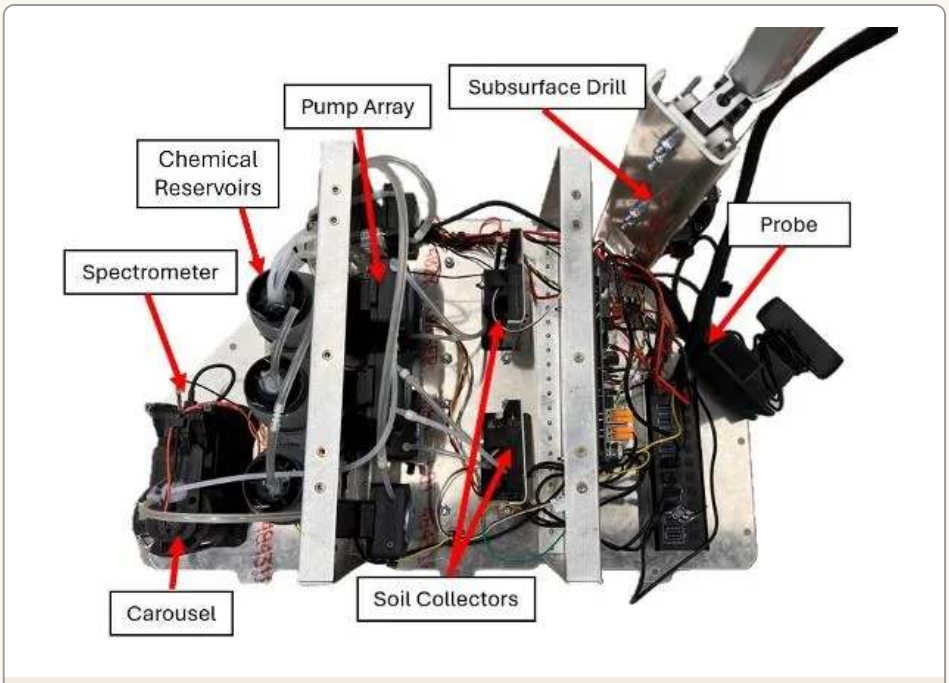
The numbers that size the design

The mass, power and fluid figures are class estimates for a first-order budget. The mission numbers are rules.

QUANTITY	VALUE	BASIS
Life detection	two capabilities, at least one a bioassay	S-URC27-RULES S-URC27-QA §1.b.v AQ9
Roving window	20–30 minutes, plus 5 minutes for cache removal	S-URC27-RULES §1.b.ii §1.b.vii
Site radius	0.5 km from C2, at least two sites evaluated	S-URC27-RULES §1.b.ii
Sample	at least 10 cm deep, at least 5 g from below 10 cm	S-URC27-RULES §1.b.vi
Sample volume	5–8 cm^3 for 5 g; a 20 mm core to 10 cm is only about 3 cm^3	Geometry
Payload mass	3.5–4.5 kg in the recommended configuration; 3–6 kg and 3–8 L overall	Class estimate, 6–12 % of the 50 kg limit
Peak power	60 W drilling, 15 W assay, about 8 Wh for the window	Estimate from pump, heater and drill classes
Fluid inventory	about 12 mL total, about 10 mL waste	2 sites × 4 reactions × 250 µL, plus extraction and flush
2026 Science score	median 85, maximum 100, 13 teams at 90 or above	Official scores S-URC26-SCORES

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Payload architecture	A: ATP and BCA, one detector · B: ATP and autofluorescence · C: LAMP and lateral flow	**A**: two unambiguous bioassays, one carousel, one detector
Bioassay channel	ATP luminometry · BCA or Bradford protein · lateral flow · LAMP · DNA stain	ATP first, because it reads in 10–60 seconds, then a rapid BCA
Sample acquisition	body-mounted auger · arm-held auger · core tube · scoop and trench	A body-mounted auger on a linear axis with a clear graduated liner, plus a scoop as backup
Fluid path	peristaltic pumps and a carousel · single-use path per site · paper microfluidics	Peristaltic pumps and an 8-cuvette carousel; single-use if the contamination argument must be strongest
Cache	screw-cap tube with an O-ring · dual topsoil and subsoil cache · bayonet · clamshell	A 50 mL screw-cap tube with an O-ring, or a dual cache that separates above and below 10 cm
Documentation	panorama with a placed scale bar, GNSS with elevation and accuracy	8–12 frames, a north arrow, a scale bar, and a GNSS record with elevation and accuracy



The architecture in one photograph, annotated by Team Mountaineers in its 2025 System Acceptance Review: visible spectrometer with halogen backlight, three reagent reservoirs, six peristaltic pumps, an 8-slot cuvette carousel, soil collectors, a 30 cm-stroke drill and a probe for soil moisture and temperature. This is what "two capabilities, one detector" looks like as hardware.

Team Mountaineers, West Virginia University, 2025. *Original source*

Full resolution

The 25-minute clock, not the chemistry

Two sites cannot be run serially. The worked sequence puts sample acquisition and documentation in the first 12 minutes, starts assay 1 at that point and lets it run while the rover drives to site 2. The second assay then has about 3.5 minutes of loading and 2 minutes of read.

STEP	TIME	WHAT RUNS ALONGSIDE
Site 1 documented and sampled	0:00–12:00	Nothing; this is the critical path
Assay 1 loaded and started	12:00–13:00	The slowest assay starts first
Drive to site 2	13:00–15:00	Assay 1 reads and the result is captured
Site 2 documented and sampled, assay 2 loaded	15:00–21:30	None; the clock is the constraint
Assay 2 read, negative control run	21:30–23:30	The operator writes the result summary
Cache sealed and stowed	23:30–24:30	None
Parked for handover	24:30–25:00	The team then has 5 minutes to remove the cache

Three methods drop out of that timetable. Colorimetric LAMP needs 30 minutes at 65 °C and its master mix ships at −20 °C, which a desert rover cannot store, and acidic soil can turn its dye yellow for a false positive `S-NEB-LAMP`. PCR is slower still. A standard BCA run is 30 minutes at 37 °C, but the rapid variant is 5 minutes at room temperature, which is why it survives as the second channel `S-THERMO-PROTEIN`.

Two design rules come out of the arithmetic. Start the slowest assay first, at site 1. And treat a failed drill as the whole mission, because acquisition and documentation, not chemistry, hold the critical path: 10.5 of the 25 minutes.

Depth is a record, not a claim

AQ10 makes depth provenance part of the mission. The sample evaluated for life must come from at least 10 cm, must be a specific sample rather than a bulk mixture, and the location and the depth of any life detection must be identified `S-URC27-QA` AQ10. Five grams of soil is only about 5–8 cm^3, and a 20 mm core to 10 cm holds about 3 cm^3, so a 25–30 mm core or a repeated scoop is needed. That figure is geometry, not a published requirement.

The design rule is to make the depth visible. Instrument the stroke with a linear encoder or limit switches and log it, keep the sample in a transparent liner with 1 cm bands so the portion below 10 cm can be photographed in place, and state the depth in the metadata and the debrief. The record costs about 50 g and one camera frame.

The cache then faces a real test rather than a statement: fill it with 10 g of analogue soil and 5 mL of water, shake it for 10 minutes at rover-representative levels, then invert it 20 times over 60 seconds above a pre-weighed absorbent pad. A mass gain above 0.1 g fails, and gloved removal should take no more than 60 seconds against a 5-minute rule.



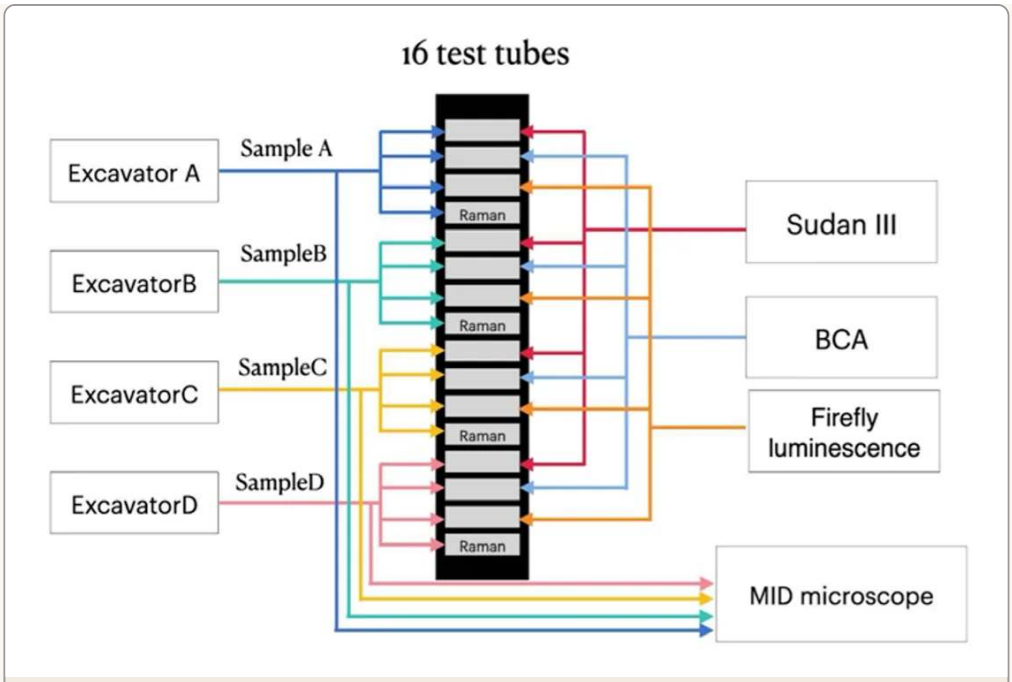
Missouri S&T's 2024 Hyperion drilling on a vertical axis, with the sample zone visible around the hole. A body-mounted drill like this keeps the reaction torque off the wrist, and the stroke becomes the depth measurement when the liner is graduated.

Missouri S&T Mars Rover Design Team (MRDT), 2024. Full resolution

Contamination is argued, not measured

AQ7 asks for a justification rather than a standard, so the deliverable is a short table: the contamination source, the control, the residual risk and the evidence. Five sources carry almost all of it: the sampler between sites, the wetted path, the operator's gloves, the sample liner and the detector optics. Each needs a single-use component or a documented purge with a blank read, plus one sentence stating the level claimed.

A reagent blank in every mission is the cheapest credibility upgrade in the payload: one cuvette and a few hundred microlitres turn an assertion about cleanliness into a number on the results slide.



KARURA's 2024 answer to contamination: four independent excavators, each feeding its own set of tubes for Raman, lipid, protein, ATP and a one-chip microscope, with no shared wetted surface between samples. It is the only published URC diagram that shows a complete cross-contamination strategy rather than asserting one.

*KARURA (Hirai et al.), *Planetary Science Society of Japan* journal, Vol. 33 No. 3, 2024, 2024.*

Original source *Full resolution*

What the field scores

Science is the most forgiving of the four missions and the one where the new rule bites hardest, because a payload can score well under the old rules and fail the new clause.

YEAR	TEAMS SCORED	MEDIAN	MAXIMUM	TEAMS AT 90 OR ABOVE	TEAMS BELOW 70
2022	26	82.0	100	11	6
2023	31	75.7	100	6	14
2024	34	73.0	100	1	10
2025	35	82.0	100	11	7
2026	36	85.0	100	13	11

At least one team reached 100 every year, and the median has never fallen below 73 across 162 team-years `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. The 2024 dip is the year the deeper sample and the sealed cache arrived, which suggests the field lost points on execution rather than on ideas. Scores in the forties and fifties look like a documented site with a failed sample, and the 2027 weights are not published, so that reading is a calibration hypothesis.

The documentary record is thin where it matters. Four of the six payloads that scored 100 have no public description, so they cannot be assumed to have done anything exotic. The one that is documented, UIU's 2025 AXIOS, was a straightforward wet system: a deep soil auger, multi-chamber reagent testing read spectrophotometrically, and a Bradford protein assay `S-UIU-AXIOS-2025`. KARURA earned 74 in 2024 after its sampler was destroyed by a rock, which is the empirical case that observation and a good debrief carry real points `S-KARURA-2024-PAPER`.

Interfaces that bite

- **Arm.** The attachment interface needs at least 10 N of downforce at full reach and wrist rotation for cap torque, but a compliant wrist makes a 30 cm drill wander. The drill belongs on the body, not the wrist.
- **Mobility.** The rover must hold position while the drill applies downforce, and drilling on a slope loads the suspension. A braced stance or a ground-contact leg is worth about 0.5 kg.
- **Power and heat.** Drilling at 30–100 W and heating at 10–20 W can coincide with drive and arm peaks; the energy is small at about 8 Wh, but the peak is the problem. Reagents must also survive 38 °C, so a sealed reagent bay with desiccant fights the dust sealing that the rest of the payload needs.
- **Communications.** An 8–12 frame panorama at about 3 MB per frame is 24–36 MB, which no rover link carries live, so the panorama is stored and forwarded or compressed hard.
- **Operations.** Only team members who stayed in C2 and never saw the site may take part in the debrief `S-URC27-QA` AQ8. The person who understands the assay has to be the C2 operator, so the shift needs two people, not one.

Failure modes

MECHANISM	WHAT HAPPENED	DESIGN CHANGE THAT PREVENTS IT
Sampler struck an immovable rock	KARURA 2024 damaged its sampler during the mission, abandoned sampling and analysis, and still scored 74 through observation and discussion	A current-limited sampler with a slip clutch, force sensing, a second method such as a scoop, and rock rejection from the camera
Instrument too weak for the claim	KARURA's Raman could not reach a level where organic compounds could be identified	Do not count Raman as life detection; prove sensitivity on an analogue soil before the System Acceptance Review
Assay time exceeded the window	KARURA's BCA needed 30 minutes at 37 °C, so the extract was pre-heated before the mission	Start the slowest assay first, and choose the 5-minute room-temperature variant where one exists
Reagent cold chain	LAMP master mix ships and stores at −20 °C	Select room-temperature-stable reagents and test them at 38 °C for five days, or drop the method
False positive from soil pH	Acidic samples turn the colorimetric LAMP dye yellow	Buffer the extract, run a matrix blank, or choose a method that does not read pH
Debrief given by someone who saw the site	AQ8 restricts participation to C2-only members	Name the speaker in the operations plan, and keep that person in C2

Shortlist for a first build

CONFIGURATION	CAPABILITY 1	CAPABILITY 2	VERDICT
A: two wet assays, one detector	ATP luminometry	Rapid BCA protein	Recommended. Two clearly sufficient bioassays, one carousel, one detector, 8–12 minutes to two results
B: one wet, one optical	ATP luminometry	Autofluorescence microscopy	Lower fluid risk, but the second capability is ambiguous under AQ9; choose it only with a written answer from the URC Director
C: amplification	LAMP	Lateral flow	Not viable in 20–30 minutes with a −20 °C cold chain and soil DNA extraction

Configuration A itemises to about 3.5–4.5 kg and 60 W peak: a body-mounted auger with a clear liner and a scoop as backup, two or three peristaltic pumps with an 8-cuvette carousel, an ATP channel with a dark cup, a BCA channel read by camera or photodiode, a panorama camera with a 20 cm scale bar, a GNSS receiver with the team's own base, and a 50 mL screw-cap cache. Every mass is a class estimate and has to be replaced by a weighed part during the build.

Open problems

- **Does one combined assay count as two capabilities?** The rule is silent and AQ9 does not address it. A written answer from the URC Director changes the payload mass by about a kilogram, and the question should be asked well before the 2027-03-05 System Acceptance Review.
- **How much ATP does the extraction recover from soil of the kind at the Mars Desert Research Station?** A method that extracts a small fraction of the ATP reads negative on a living sample regardless of the biology.
- **How large are the false positives?** Extracellular ATP, clay-stabilised ATP and reagent contamination all mimic life, and a false positive is worse than a false negative. Oven-sterilised soil and a reagent blank give the answer.
- **Do humic substances mimic protein?** BCA depends on copper reduction, which soil organics also perform, and a standard-addition recovery of 80–120 % is the pass mark. A timed dry run with a mock sample is the only evidence that two results fit inside 25 minutes.
- **Do the reagents survive five days at 38 °C?** A denatured enzyme is a dead payload, and the chamber test takes five days, so it has to start early.

Further reading

- KARURA, WVU, Missouri S&T, UIU and Monash carry the documented payloads.
- Astrobiology maps the mission tasks to the payload.
- Mass and power places the payload in the 50 kg budget.
- Interfaces and couplings works the drill-torque and data-volume chains.

Delivery, payload handling and drones

Delivery is a staged mission: the rover finds, picks up, carries and delivers objects over terrain of increasing difficulty, and it must reach a minimum score in Stage 1 before Stage 2 opens. Objects reach 5 kg, the course reaches a kilometre, and line of sight is not guaranteed for more than half of it. This page is a synthesis rather than a dedicated study. No comparable published account of the 2027 Delivery mission or its drones exists, so the evidence is drawn from the 2027 rules and Q&A, the team dossiers, and the arm and astronaut studies. It is thinner than the pages around it, and it says where.

5 kg
OBJECT LIMIT
Under 40 × 40 × 40 cm, grasp feature under 7.5 cm
2027
S-URC27-RULES

30–60 min
TOTAL MISSION TIME
Stage 1 gate, leftover time carried into Stage 2
2027
S-URC27-RULES

22.9
MEAN 2026 DELIVERY SCORE
Median 20, twelve zeros of 36
2026
S-URC26-SCORES

5 kg
DRONE LIMIT
Includes an inert dummy mass equal to the battery mass
2027
S-URC27-RULES

Key takeaways

- **The Stage-1 gate is the mission.** Failing the minimum ends the run, and in 2026 the gate was 40 of 100. The cheap tasks have to be banked before the expensive ones are attempted.
- **The object envelope is a gripper specification.** A wide-opening, current-limited parallel gripper with a hook covers every documented task, and a mini-rover is explicitly prohibited.
- **The drone is optional.** It is capped at 5 kg with an inert dummy mass, allowed in Delivery and Autonomy only, and full points are possible without one, so it never belongs on the critical path.
- **Delivery is a link problem before it is an autonomy problem.** The rover may be teleoperated throughout, but the link has to survive a kilometre of terrain that blocks line of sight.

What the subsystem must do

Each stage has a fixed time and several tasks. The team must achieve a specified minimum score on Stage 1 to proceed, and remaining Stage-1 time is added to Stage 2, for a total on-course time of **30–60 minutes** S-URC27-RULES §1.c.i. Objects are **up to 5 kg**, smaller than **40 × 40 × 40 cm**, with a graspable feature **no larger than 7.5 cm**, and they may be small hand-held items or instruments, rocks, or supply containers, grasped by a handle or by a natural feature S-URC27-RULES §1.c.iii.

The course runs up to **1 km** from the start gate, across soft sand, gravel, rough stone, boulder fields, vertical drops and steep loose slopes, and the total path may exceed the kilometre S-URC27-RULES S-URC27-QA §1.c.ii RQ13. Line of sight is "not guaranteed for more than 50 % of the course", and later tasks are deliberately placed out of it S-URC27-RULES §3.d.ii. Working GNSS is described as essential, but the coordinates are only approximate, so the rover has to search at the waypoint and identify the right object S-URC27-RULES §1.c.iv.

The rulebook lists the task families a team can expect, and the order is not known in advance: open a box with a hinged lid, pick up and deliver an object, read a sign carried by an astronaut or placed in the field, find equipment in a large search area, find objects in an area with poor radio reception, and use approximate GNSS coordinates for pickup and delivery S-URC27-RULES §1.c.iii–iv. A drone may help. No mini-rover or other mobile deployable may be used S-URC27-RULES §1.c.i.

The 2027 change is small and easy to miss: the 2026 rule listed following a marked path as a possible task, and the 2027 rule does not S-URC26-RULES S-URC27-RULES §1.c.iii §1.c.iii. Everything else about the object sizes, the range and the scoring basis is carried forward.

The numbers that size the design

The mission numbers are rules. The gripper strengths and the time arithmetic are inferences from the arm analysis and from the 2026 course, and they are labelled as such.

QUANTITY	VALUE	BASIS
Object envelope	up to 5 kg, under 40 × 40 × 40 cm, grasp feature under 7.5 cm	S-URC27-RULES §1.c.iii
Course range	up to 1 km; total path may exceed 1 km	S-URC27-QA RQ13
Mission time	30–60 minutes, with fixed time per stage	S-URC27-RULES §1.c.i
Stage-1 gate	a specified minimum; 2026 used 40 of 100	2027 value not published; 2026 proxy S-URC26-DM
Placement tolerance	2026 required the object within 20 cm, with no rover part touching	2027 value not published; a useful design target S-URC26-DM
Drone mass	5 kg or less, including the inert dummy mass	S-URC27-RULES §2.b.v–vii
Point density	1.67–3.33 points per minute across a 30–60 minute mission	100 points over the stated window; the 2026 course was 1.82 S-URC26-DM
2026 Delivery score	mean 22.9, median 20, twelve zeros of 36	Official scores S-URC26-SCORES

Design space, in one table

DECISION	OPTIONS	WHAT THE EVIDENCE RECOMMENDS
Payload carrier	arm only · cargo box or crane · towed cart · mini-rover · drone only	Arm only. Every documented top team carries with its arm, and a mini-rover is prohibited
End effector	two-finger parallel · hook or pin · dual head · swappable heads · suction	Two-finger parallel with 100–110 mm opening, plus a hook for latches and drawers
Search	mast camera with map coverage · drone · ground pattern · learned detector	Mast camera with map-driven coverage and an object detector; a drone is optional
Drone	none · commercial quad · custom quad	None for a first build, or a commercial quad under 5 kg; never on the critical path
Autonomy	teleoperation · GNSS waypoints with object detection · full autonomy	Teleoperation with GNSS waypoints and an object detector. Delivery does not require autonomy

The Stage-1 gate is the mission

The stage structure turns the first minutes into the highest-value minutes of the run. A rover that fails the gate scores whatever it banked in Stage 1 and stops, so a slow start is not a slow mission but a short one. In 2023 the field median was **15 of 100** S-URC23-SCORES.

The 2026 course is the only published proxy for the shape of the scoring. Stage 1 ran 35 minutes for 70 points split 10, 10, 30 and 20, with a **40-point gate**; Stage 2 added 10 and 20 points, a 20 cm placement tolerance, a 20 m wreckage radius and cones about 100 m apart S-URC26-DM. None of those values is a 2027 rule, and the 2027 stage durations, the gate and the task weights are all unpublished.

The team record shows what a gate failure looks like over several seasons. Missouri S&T scored **80** in 2022, **15** in 2023 and **15** again in 2024, then **70** in 2025 and **100** in 2026 S-URC22-SCORES S-URC23-SCORES S-URC24-SCORES S-URC25-SCORES S-URC26-SCORES. The 2023 account is explicit: the rover ran out of time in Stage 1 without enough points to proceed, and the team blamed its inability to locate a specific rock through its cameras S-MST-NEWS2023. The 2024 failure was different and equally instructive: mechanical faults in the rover-side signals and in the vertical axes of the arm S-MST-NEWS2024. One season was a perception failure, the next a mechanism failure under driving loads, and both ended the mission before the expensive tasks.

The same spread appears across the field. Monash went 70, 15, 80, 71 and 76; WVU won in 2023 with 90 and in 2025 with 80 but fell to 52 and 61 in the other years; UIU went from 18.75 to 62 S-MONASH-IEEE-RAM S-WVU-SAR2024 S-UIU-2026-SAR. Between one and twelve teams score exactly zero each year, and the distribution is not graded: a mediocre run is not a safe plan. Missouri S&T reached 100 in 2026, and it is the same team that scored 15 twice, so the difference between a zero and a win is mostly whether the first three tasks worked.

The object envelope is a gripper specification

The rules give a mass, a box and a grasp feature, not a shape. That is what makes a single-purpose hook a bad design and a wide gripper the starting point. The arm analysis puts numbers on the classes: a jaw of at least 60 mm and a grip of at least 40 N for a 5 kg cache, at least 70 mm and 60 N for a 0.91 kg hammer, and a 100–110 mm opening to cover the cache handle, the hammer shaft, a valve lever and a hose fitting. Those strengths and openings are inferences from the arm study, not published figures.

Two details decide whether the grasp survives the drive. The first is that objects can be grasped by a natural feature, so the grip has to work on a rock as well as on a handle, which argues for a closed jaw with a mechanical stop rather than a friction pinch. The second is the 2026 placement tolerance of 20 cm, which means the mechanism has to place the object rather than throw or drop it. No rover part may touch the object once placed S-URC26-DM.

The end effector is therefore the same one the Equipment Servicing mission needs, with one addition: a hook or pin for latches and drawer pulls. That sharing is a systems decision rather than a Delivery decision, and it is worth taking, because a second end effector costs mass at the wrist and a second calibration.



The payload class in one frame: Monash's 2026 arm lifting the cache container out of a mock lander drawer, with the inset bench scale reading 5.085 kg. The same gripper, jaw and wrist serve the Delivery carry and the lander tasks, which is why one end effector is the right systems answer.

Monash Nova Rover, 2026. *Full resolution*

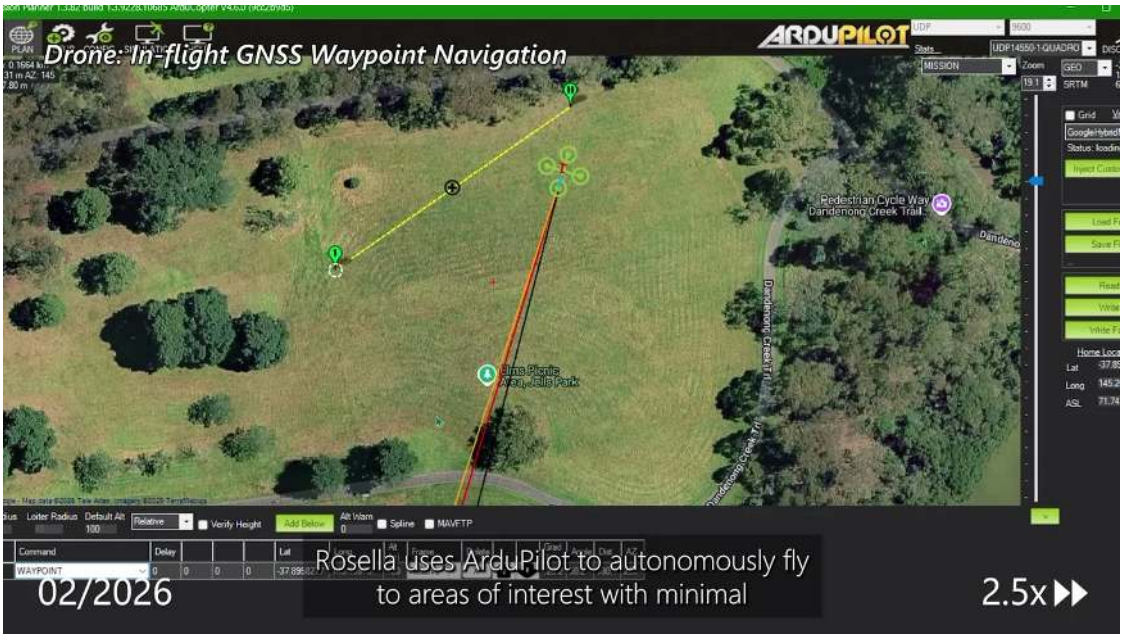
The drone, capped at 5 kg

A drone is optional and full points are possible without one S-URC27-RULES §2.b.i. If one is flown, it must be rotary-wing and able to hover, it must start from the designated landing pad, and it must weigh **5 kg or less including an inert dummy mass equal to its battery mass**, carried at all times; a spare battery is not a dummy and high wind does not waive the requirement S-URC27-RULES §2.b.iii–vii. It may assist **Delivery and Autonomy only**, never Astrobiology or Equipment Servicing, and in Autonomy it must land before the rover proceeds S-URC27-QA S-URC27-RULES §1.e.xiv RQ9. The constraints are traced as REQ-DRN-002 for the mass, REQ-DRN-006 for the missions in which it may fly, and REQ-DRN-011 for the requirement that the plan must not depend on it.

Three operational costs sit behind that envelope. A human battery swap is an **intervention**, and interventions are additive at 20 % each S-URC27-QA S-URC27-RULES §3.e RQ14. The pilot sits in the C2 station with no view, so a field spotter is the only line of sight, and the base needs at least 25 m of cable S-URC27-QA DQ2. And the airspace rules apply in full: registration above 250 g, Remote ID, a TRUST certificate, below 400 ft (120 m) above ground level, class G airspace, and no flight in a National Park

S-URC27-RULES §2.b.vii–xii. The rules strongly suggest a drone that can operate in 48 km/h wind, which implies a top speed of at least 64 km/h, but that is a suggestion rather than a requirement.

The team record is thin and mostly negative. Missouri S&T's Zephyr never flew at competition in 2024, while Noctua in 2026 scouts and acts as a signal repeater; WVU uses a commercial DJI Mavic 2 Pro; UIU's drone identifies rocks and reads signs but does not deliver S-MST-SITE-ATHENA S-WVU-SAR2024 S-UIU-2026-SAR. No team is documented flying a relay that changed the outcome. The honest reading is that a drone buys time and reduces risk on search and sign-reading tasks, and that it is not a scoring line of its own.



A commercial-class quad flown by Monash Nova Rover in 2026, with a Mission Planner waypoint mission over Victoria on the screen. This is the realistic role for a first drone: an overhead view of a search area, not a carrier. The 5 kg cap including the dummy mass leaves almost no margin for a payload.

Monash Nova Rover, 2026. Full resolution

Interfaces that bite

- **Arm and end effector.** One arm carries 5 kg, opens a hinged box, reads a sign and, in Autonomy, fetches a hammer. A 100–110 mm gripper with swappable tips is the common denominator, and the same wrist has to pass through a 1.2 m cube when stowed.
- **Communications.** Up to 1 km, with line of sight broken for more than half the course, makes Delivery a link problem even though it is teleoperated. Both operators sit inside a metal trailer with no view, and every intervention costs 20 %.
- **Mobility and power.** Soft sand, gravel, boulder fields, vertical drops and steep loose slopes up to 1 km, with a 5 kg object raising the centre of gravity. At up to 60 minutes this is the longest mission.
- **GNSS.** Approximate coordinates plus a local search means the navigation solution has to survive a waypoint that is wrong by tens of metres, and the team must run its own base station because there is no correction service on site S-URC27-QA RQ11.

What fails

The failure record is short, because most of it is buried in post-season notes rather than published. Three mechanisms account for most of the zeros.

- **Object finding.** A rover that cannot pick the required rock out of a boulder field cannot start. Missouri S&T's 2023 account is the clearest statement of it S-MST-NEWS2023.
- **Mechanism failure under driving load.** The 2024 failure in the arm's vertical axes happened in Delivery rather than in Equipment Servicing, which is the point: a mechanism that works while stationary can fail on a slope S-MST-NEWS2024.
- **Time.** The gate converts a slow start into a zero, and the 2023 field median of 15 shows how common that is S-URC23-SCORES.

Shortlist for a first build

ITEM	CHOICE	WHY
Carrier	The arm, with no cargo mechanism	Every documented top team does it, and a cargo box adds mass and failure modes for no rule requirement
End effector	Two-finger parallel, 100–110 mm opening, plus a hook	Covers a handle, a rock, a hammer shaft and a latch with one wrist
Grasp	Closed jaw with a mechanical stop, grip force monitored	A natural-feature grasp cannot rely on friction alone over rough ground
Placement	Designed to 20 cm, placed rather than dropped	The 2026 tolerance is the only published figure and a sensible target
Sequence	Bank the cheap Stage-1 tasks first, with a timed rehearsal	The gate makes the first minutes worth more than their face value
Navigation	GNSS waypoints plus an object detector, teleoperated	Delivery does not require autonomy, and the operator can resolve what the detector cannot
Drone	None, or a commercial quad under 5 kg with a pilot and a spotter	Full points are possible without one, and the cap leaves little margin

Open problems

- **What are the 2027 stage durations, the Stage-1 gate and the task weights?** All are unpublished and promised before the competition, and all three decide the task order.
- **What do the objects actually weigh and measure?** The rules give limits, not objects. A mock set built to 5 kg, 40 cm and a 7.5 cm feature is the only way to test a grasp.
- **How much force does a hinged lid need?** No figure is published for the box lid, the latch or the drawer.
- **What is the real link margin out of line of sight?** The course is the test, and it is not available before the event.
- **Does a drone earn back its mass, its pilot and its intervention risk?** No team documents a case where it did, and no team documents a case where its absence cost a mission.
- **What is the placement tolerance for 2027?** The 20 cm figure is a 2026 value and may change.

Further reading

- Missouri S&T, Monash, WVU and UIU carry the delivery record and the drone evidence.
- Delivery maps the mission tasks to the design.
- Arm and end-effectors sizes the gripper and the carry.
- Interfaces and couplings works the payload, link and mass chains.

Communications and networking

Communications decides how much of the rover the operators can actually use. Three of the four 2027 missions are teleoperated, and a link that fails does not lose a share of the points, it loses the mission. The rules also make the hard case deliberate: line of sight is not guaranteed, one route target sits behind a hill, and the 900 MHz system has to move between assigned sub-bands. This chapter fixes those limits, compares the candidate bands on computed margins, and recommends a three-plane architecture.

48.2 dB
900 MHZ MARGIN AT 1 KM
Still +28.2 dB with a +20 dB obstruction
Link model
S-VENDOR-RFDESIGN-RFD900X

8 MHz
900 MHZ SUB-BAND CAP
One of three windows, assigned per mission
2027
S-URC27-RULES

4.6×
LINE-OF-SIGHT TO OBSTRUCTED COLLAPSE
750 m becomes 164 m behind terrain
WVU 2026
S-WVU-SAR2026

10.7 dB
25 M LMR-400 LOSS AT 5.8 GHZ
3.2 dB at 900 MHz, 5.5 dB at 2.4 GHz
Datasheet
S-VENDOR-TIMES-LMR400

Key takeaways

- **At 1 km every candidate band closes** with margin at compliant EIRP. What kills a link is the close-in case, where a high-gain base puts **+2.6 dBm** into a receiver specified for **−10 dBm**, and terrain, where **+20 dB** of obstruction drives 2.4 GHz to **−7.6 dB**.
- **900 MHz is the plane that survives the intentional non-line-of-sight**, at **+28.2 dB** against the same 20 dB obstruction. WVU measured 750 m of line of sight collapsing to **164 m** in 2026.
- **The 8 MHz sub-band cap is a documentation problem before it is a radio problem.** 250 kHz channels give **32 hops**, which complies; 500 kHz gives **16**, which does not. The RFD900x ships with sub-bands that match neither 2027 window.
- **Three planes on different bands**, each with its own antenna and power feed, is the strongest convergence in the published record. No documented team runs a single radio.

What the subsystem must do

The 2027 rules state the radio problem in one clause: the network must carry control, telemetry and video with **no internet, cellular or satellite relay** S-URC27-RULES §2.c.i. The course then removes the easy geometry. Line of sight is "not guaranteed for more than 50 %" of the distance S-URC27-RULES §1.c.ii, the second route target is deliberately out of sight S-URC27-RULES §1.e.xvi, and the operators sit inside a metal C2 trailer with no view of the course S-URC27-RULES §3.d.i. A rover that cannot be seen has to be seen through the link.

Four more clauses fix the hardware. The 900 MHz system may occupy at most **8 MHz** and must work inside one of three assigned sub-bands, shifting between them per mission S-URC27-RULES §2.c.v. The base antenna is limited to **3 m** of height, its base within **5 m** of the C2 station, with at least **25 m** of cable S-URC27-RULES S-URC27-QA §2.c.iv DQ2. No camera may be used to aim an antenna, which rules out the tracking systems that would otherwise make a dish practical S-URC27-RULES §2.c.iii. Every radio must also be documented as compliant by **2027-04-30**, and the judges are explicit that being able to buy equipment is not the same as being allowed to operate it S-URC27-RULES S-URC27-QA §2.c.ii RQ7.

The whole station has to be standing inside the **15-minute** set-up window, which makes pre-assembled masts and colour-coded harnesses part of the communications design rather than an operations detail S-URC27-RULES §1.a.ii.

The numbers that size the design

The margins assume compliant EIRP and a stated interference allowance. They are the arithmetic to reproduce before choosing a band; the obstruction, latency and video rows are the ones that decide a mission.

QUANTITY	VALUE	BASIS
Distance the link must cover	up to 1 km from the C2 station, and the total path may exceed it	Delivery and Autonomy S-URC27-QA RQ13
Free-space loss at 1 km	91.7 dB at 915 MHz, 100.2 dB at 2.45 GHz, 107.7 dB at 5.8 GHz	Computed from frequency and distance
900 MHz margin at 1 km	48.2 dB line of sight; **+28.2 dB** with a +20 dB obstruction	1 W, 3 dBi rover, 16 dBi base, 3 dB noise rise
2.4 GHz margin at 1 km	12.4 dB line of sight; **−7.6 dB** with a +20 dB obstruction	28 dBm, 5 dBi rover, 18 dBi base, 10 dB noise rise
5.8 GHz margin at 1 km	15.1 dB with 25 m of coax; 25.6 dB with the radio at the mast	25 dBm, 13 dBi panel, 29 dBi base
23 cm margin at 1 km	52.0 dB line of sight; +32.0 dB with obstruction	Licensed amateur, 2 W, 12 dBi rover, 15 dBi base
25 m LMR-400 loss	3.2 dB at 900 MHz, 5.5 at 2.4 GHz, 10.7 at 5.8 GHz	Datasheet S-VENDOR-TIMES-LMR400
900 MHz channel plan	32 hops at 250 kHz spacing; 16 at 500 kHz	FCC needs at least 25 hops S-FCC-15-247
EIRP ceiling	36 dBm, so a 29 dBi dish may be fed about 7 dBm	The fixed point-to-point exemption does not cover a moving rover S-FCC-15-247
Measured obstruction collapse	850 m to 300 m, and 750 m to 164 m	WVU 2023 and 2026 S-WVU-SAR2023 S-WVU-SAR2026
Latency budget	70 to 230 ms glass to glass; target ≤250 ms	Camera, encode, air, decode, command; estimated
Video demand	six 720p30 streams at 3 Mbit/s is 18 Mbit/s	Impossible on 900 MHz, marginal on a contested 2.4 GHz link

Design space, in one table

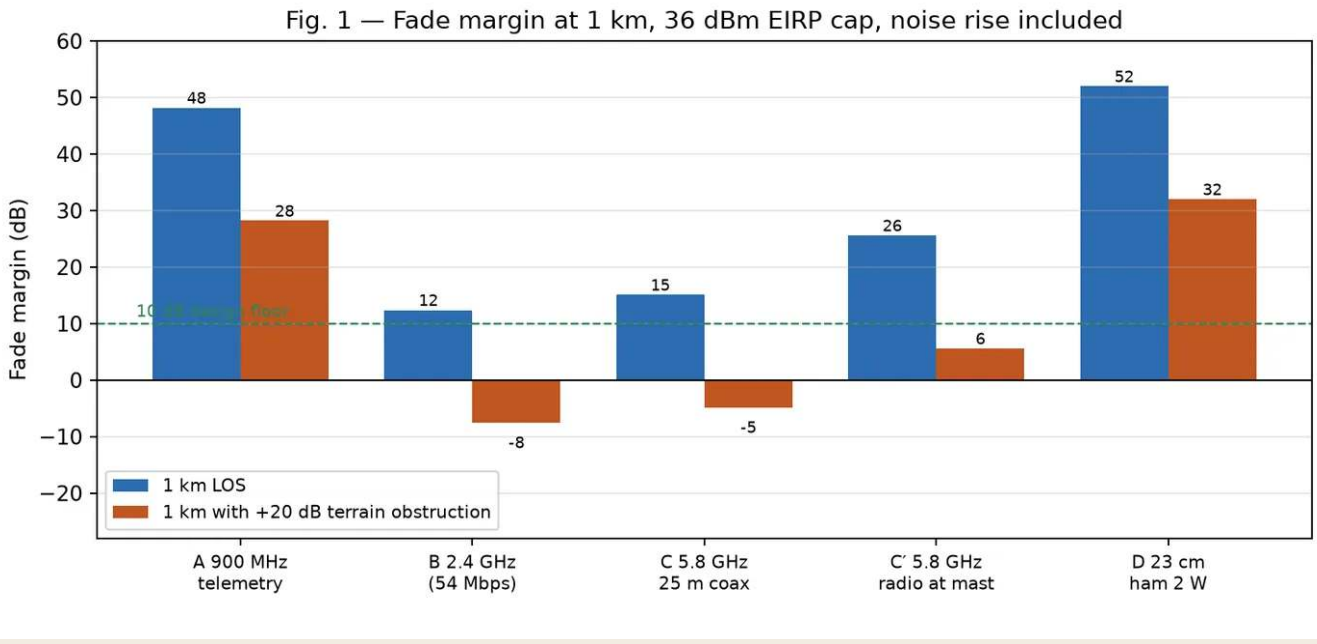
DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Band	900 MHz · 2.4 GHz · 5.8 GHz · 23 cm · 70 cm	2.4 GHz primary, 900 MHz recovery, 5.8 GHz as an optional third plane
Architecture	single band · two planes · three planes	Three independent planes, each with its own antenna and power feed
Base antenna	omni · 90–120° sector · 29 dBi dish	16–18 dBi 90–120° sector: gain with the coverage a moving rover needs
Rover antenna	3–5 dBi omni · 13 dBi panel	5–10 dBi omni on a 2–3 m mast, or a 13 dBi panel if the rover can be aimed
Feedline	25 m coax · radio at the mast on PoE	Radio at the mast: 10.5 dB recovered at 5.8 GHz, and one fewer connector chain
Video transport	MJPEG · RTSP H.264 · WebRTC · analog	Hardware H.264 at 720p with an operator-selectable 6 Mbit/s fallback
Rate control	fixed rate · operator-selected rate ladder	Operator-selectable: falling from MCS7 to 6 Mbit/s buys about 9× the range at 2.4 GHz

At 1 km every option closes, and two effects still kill the link

The headline budgets are comfortable: 900 MHz closes 1 km with 48 dB of margin, 2.4 GHz with 12 dB, 5.8 GHz with 15 dB, all at compliant EIRP. A team that fails at URC did not fail a free-space calculation. Two effects outside that calculation decide the outcome.

Close-in overload. At 5 m the 900 MHz architecture with a 16 dBi yagi puts **+2.6 dBm** into the receiver, as does 5.8 GHz with the radio at the mast. Most radios specify a **−10 dBm** maximum input. Every mission begins with a close-in leg, so the design needs transmit power control or a fixed attenuator, tested at 5 m rather than at 1 km.

Terrain. With 20 dB of obstruction the two Wi-Fi bands go negative, at −7.6 dB for 2.4 GHz and −4.9 dB for 5.8 GHz on 25 m of coax, while 900 MHz stays at **+28.2 dB**. WVU measured the collapse in its own tests: 850 m of line of sight became 300 m obstructed in 2023, and 750 m became 164 m in 2026 `S-WVU-SAR2023` `S-WVU-SAR2026`. The judges' advice to design for both near and far operation is a direct statement of these two columns `S-URC-ADVICE`.



Computed fade margin at 1 km under the 36 dBm EIRP cap. The chart rounds the table: 900 MHz holds 28 dB when obstructed while both Wi-Fi bands cross zero. The 10 dB line is the usual design floor.

Robotics Club @ Mississippi State University, 2026. Original analysis.

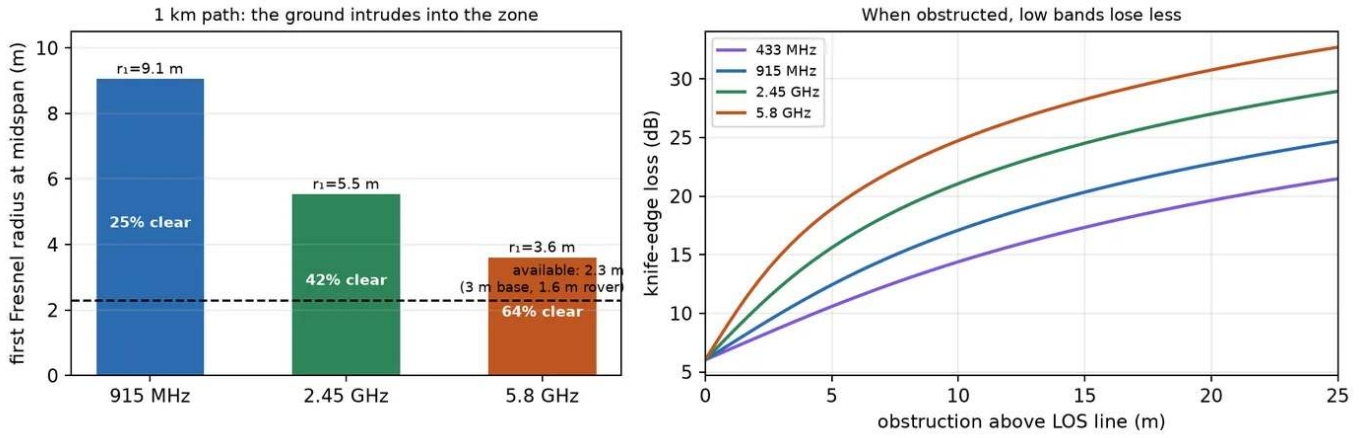
Full resolution

Height is worth more than band choice in the open

The first Fresnel radius at the middle of a 1 km path is **9.1 m** at 915 MHz, 5.5 m at 2.45 GHz and 3.6 m at 5.8 GHz. With a base antenna at its 3 m limit and a rover antenna at 1.6 m, the line between them is only 2.3 m above the ground at midspan, so the 60 % clearance criterion is met at 5.8 GHz and nowhere else. That is the opposite of the obstruction result, and the two effects pull against each other: low bands need more clearance but lose less when they do not get it.

The practical consequence is that **mast height is the cheapest gain on the link.** Raising the rover antenna from 0.6 m to 3.0 m recovers about **14 dB** at 915 MHz and 14.7 dB at 2.4 GHz on flat ground, which is the same order as switching bands. WVU grew its rover mast from 2 m in 2023 to 3 m in 2025 and reduced antenna mass to pay for it `S-WVU-SAR2023` `S-WVU-SAR2025`. The rover mast is therefore a communications component that the structures and mobility pages also have to carry, and it is the item most likely to break the 1.2 m cube when stowed.

Fig. 3 — Two effects that decide a 1 km URC link: clearance (favours high bands) vs obstruction loss (favours low bands)



The two competing effects at 1 km. Clearance favours the high bands, where the Fresnel zone is small enough for a 2.3 m line to clear; obstruction loss favours the low bands by 3 to 8 dB for the same ridge.

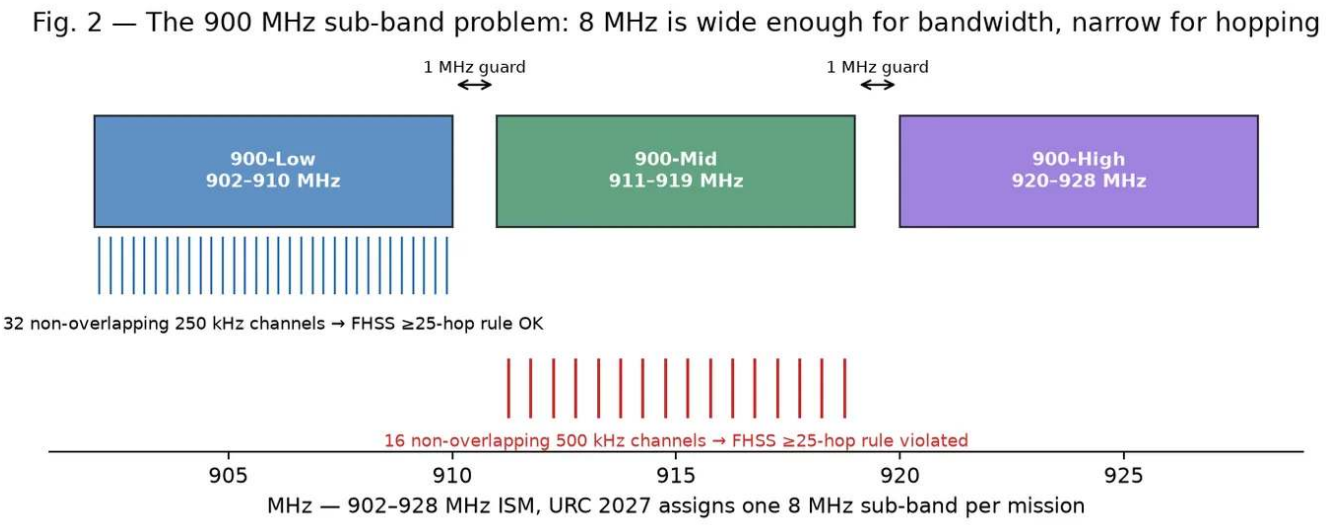
Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

The 900 MHz sub-band is a rules problem before it is a radio problem

The 2027 rule is unusually specific. Occupied bandwidth is capped at 8 MHz, operation must sit exclusively inside **902–910**, **911–919** or **920–928 MHz**, the schedule assigns the sub-band per mission, and the restriction covers every link between the C2 station, the rover, the drone and the astronaut, including local-area networks between devices at any location, and URC monitors the spectrum on site to police it `S-URC27-RULES` §2.c.v.

The FCC hop-count rule then depends on the occupied bandwidth, not the channel spacing `S-FCC-15-247` . Below 250 kHz the requirement is at least 50 hopping frequencies, which an 8 MHz window cannot hold. Between 250 and about 320 kHz the requirement is at least 25, and 250 kHz spacing gives 32, which complies. At 500 kHz the window holds 16, which fails. A team that buys a radio on its data rate and its price has not answered this question.



The sub-band plan that turns a radio specification into a compliance test. At 250 kHz spacing the window holds 32 non-overlapping channels and the hop rule is met; at 500 kHz it holds 16 and is not. The vendor datasheet has to state the 20 dB occupied bandwidth.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

The concrete trap is in the shipped configuration. The RFD900x-US is a 1 W modem with a published 224 kbit/s ceiling, and it ships with two selectable sub-bands, 902–915 and 915–928, neither of which is one of the three 2027 windows `S-VENDOR-RFDESIGN-RFD900X`. The workable answer is three saved profiles with the frequency limits set to the assigned window, each verified in set-up. That is a configuration task on every run, and the failure it prevents is a rules breach rather than a performance loss.

Three planes, not one band

WVU is the closest thing to a controlled experiment in the published record: first in 2023, second in 2024 and 2025, fourth in 2026, with a communications paragraph in every System Acceptance Review.

YEAR	PRIMARY LINK	RECOVERY	ROVER ANTENNA AND MAST	RANGE TESTED	SOURCE
2023	900 MHz, Cambium PTP 450	none	two 5 dBi whips on a 2 m telescoping mast	850 m line of sight, 300 m obstructed	`S-WVU-SAR2023`
2024	2.4 GHz, 90° sector at the base	900 MHz, 120° sector	omni	250 m on 2.4 GHz, 1 km on 900 MHz	`S-WVU-SAR2024`
2025	2.4 GHz, 90° sector	70 cm packet radio and 433.25 MHz analog television	collinear omni on a 3 m mast	750 m, with a seamless transition to the recovery plane	`S-WVU-SAR2025`
2026	2.4 GHz multi-video plus 5.8 GHz	70 cm	collapsible mast with a gimbal camera	750 m line of sight, 164 m obstructed in noise	`S-WVU-SAR2026`

The sequence shows a team moving **up in frequency for bandwidth and down in frequency for survival**, and ending with three planes rather than a better band choice. The E-stop makes the pattern necessary rather than elegant: the single kill removes all battery power, so the rover radio dies with everything else, and the C2 station sees silence that is expected rather than a fault `S-URC27-RULES` §2.a.v. A recovery plane that carries a heartbeat and drive commands at low rate is what keeps a bad minute recoverable.

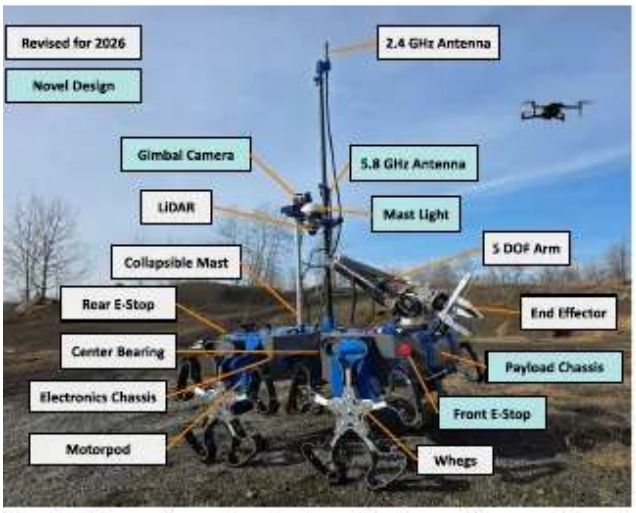


Fig. 1: Deimos' mechanical architecture diagram for all

The 2026 layout: a 2.4 GHz and a 5.8 GHz antenna on one collapsible mast, next to the gimbal camera. Two bands with two different jobs on a single mast is the architecture this chapter recommends.

Team Mountaineers, West Virginia University, 2026. *Original source*

Full resolution

The System Acceptance Review that describes this arrangement is itself the primary source for the 2026 row of the table, and it shows the mast, the antennas and the gimbal camera in service.

Video available in the web edition.

WVU Team Mountaineers, URC SAR 2026

Team Mountaineers, West Virginia University

Watch on YouTube

Antenna gain belongs at the receiver. Teams routinely put a 1 W radio on a 20 to 29 dBi antenna. That is legal for a fixed point-to-point link in 5.8 GHz and illegal for a moving rover anywhere, because EIRP is capped at 36 dBm `S-FCC-15-247`. The cap costs nothing in practice: EIRP is what the far end receives, the receiving gain is unregulated, and the correct move is to back the transmitter down and take the gain at the base station. A 29 dBi dish thereby becomes a receive-gain device, which is why the lower-gain sector is the better first choice: the judges warn that the rover will not always be in front of the antenna `S-URC-ADVICE`.

Interfaces that bite

- **Power.** The kill removes battery power, so the radio dies with the rover. Link loss after a kill is expected, not a fault, and the field judge has to read the rover's own status indicator instead. The radio also needs a quiet 12 V rail, because a radio sharing a rail with a motor controller drops packets when the drive moves `S-URC27-RULES` §2.a.v.
- **Sensors and perception.** Six 720p30 streams at 3 Mbit/s is 18 Mbit/s, which 900 MHz cannot carry and a contested 2.4 GHz link carries badly. The camera count is a link-budget decision, taken before the cameras are bought.
- **Autonomy.** The link-loss behaviour is a scoring feature. Autonomy allows a stop-and-return at a 20 % penalty or teleoperation back at 50 %, and Delivery does not allow either `S-URC27-RULES` §1.e.xviii. The state machine that holds, stops and returns has to work over the recovery plane.
- **Operator station.** Signal strength, noise floor and latency have to be continuously visible, with a one-action drop to 6 Mbit/s. Without that control the operator cannot trade rate for range.
- **Delivery and astronaut.** A drone relay is allowed and explicitly encouraged but unproven in this record, and a battery swap counts as an intervention `S-URC27-QA` `S-URC27-QA` RQ5 RQ14. A suited astronaut adds a third node on the same bands,

and may carry electronics but may not deploy a repeater S-URC27-QA AUQ26.

Failure modes

The failure record is consistent, and it is dominated by physics rather than by component choice.

MECHANISM	WHAT HAPPENS	PREVENTION
Terrain and non-line-of-sight	The link works on the flat and dies behind the hill, at 164 m instead of 750 m	A lower-band recovery plane with its own antenna; route planning that keeps the rover inside the base sector
Feedline loss	25 m of LMR-400 costs 10.7 dB at 5.8 GHz, which is the whole gain advantage of a panel over an omni	Radio at the mast on PoE, or LMR-600, and count the feedline in the budget
Receiver overload	The close-in leg saturates the front end and the link appears to fail at short range	Transmit power control, automatic gain control, or a fixed 10–20 dB attenuator, tested at 5 m
Interference from the field	About 36 teams arrive with 2.4 GHz and often 900 MHz gear; the quiet parking lot is not the venue S-URC27-TEAMINFO	Sub-band agility, a channel plan written before arrival, and a scan during set-up
Regulatory non-compliance	Non-compliant EIRP or missing documentation fails the April deadline	One compliance line per radio: band, conducted power, antenna gain, EIRP, rule
Vibration and transport	Wiring and connectors work in the parking lot and fail at the site S-URC-ADVICE	Locking connectors, strain relief, a spare antenna cable, and a link test after every transport leg
Single plane of failure	One radio is one failure domain	Two planes on different bands, independently powered

The gap in that record is worth stating: no public URC post-mortem reports measured link data such as throughput, signal strength or packet loss, so the record is judges' statements plus each team's own tested ranges. A team that instruments its link and publishes the numbers would be ahead of the field.

Shortlist for a first build

ITEM	CHOICE	WHY
Primary plane	2.4 GHz, 802.11ac, radio at the mast on PoE	Cheapest established ecosystem, video-native, and 12 dB of margin at 1 km with a 10 dB interference allowance
Recovery plane	900 MHz, 1 W modem, three saved sub-band profiles	+28 dB when obstructed, and the band that survives the intentional non-line-of-sight
Third plane	5.8 GHz, radio at the mast, optional	Quietest unlicensed band on site, 25.6 dB of margin at 1 km with the radio at the antenna
Base antenna	16–18 dBi, 90–120° sector on a 3 m mast	Gain with the coverage a moving rover needs, and no aiming problem
Rover antenna	5–10 dBi omni on a 2–3 m mast	No steering hardware; mast height buys about 14 dB
Video	Hardware H.264 at 720p, operator-selectable 6 Mbit/s fallback	One driving stream plus on-demand detail streams fits the link
Documentation	One compliance line per radio, plus the licence copies	The April deadline is a documentation task, not a radio task

Open problems

- **What is the noise floor at the competition site?** There is no measured spectrum data at MDRS. The 3, 6 and 10 dB interference allowances are estimates, and they are the difference between a 12 dB and a 22 dB margin at 1 km. A 60-second scan at a

comparable multi-team event settles it.

- **What does the 2027 terrain profile actually look like?** Base siting and the value of a relay depend on geometry that is not published. A digital elevation model of the area plus the candidate base positions gives Fresnel clearance to the far corners.
- **Does the candidate radio shift sub-bands on command?** The vendor documents two windows, the rules require three. A bench check with a spectrum display, counting hops inside 8 MHz, is the test.
- **What is the receiver's maximum input?** Most radios do not publish a compression point, and the close-in case exceeds the typical limit. A bench test that feeds a known level into each radio finds it.
- **What is the real glass-to-glass latency?** The 250 ms target is an engineering estimate. A camera pointed at a millisecond clock measures it at 1 km, at 100 m and at three video rates.
- **Does a relay earn its place?** Two hops turn a failing obstructed link into two comfortable ones, but no team in this record documents a deployment. A two-hop test with a mast at 100 m measures throughput and latency on both legs.

Further reading

- WVU carries four years of published communications detail and the only measured obstruction numbers in the record. Waterloo published its own radio research and is the clearest example of a team documenting the trade rather than the outcome.
- Power, electrical architecture and safety owns the other half of the E-stop problem: the kill that removes battery power also removes the radio.
- Compute, sensors and perception sets the camera count that the link has to carry.
- Interfaces and couplings traces the communications, power and autonomy couplings with numbers.
- The Q&A rulings carries the answers on relays, sub-bands and the base station.

SUBSYSTEMS · CHAPTER 4.6

Power, electrical architecture and safety

The electrical architecture has to satisfy a pass/fail safety rule before it meets any performance target. One red push button must stop the rover and remove all battery power, and that requirement tightened in 2026, so a design copied from an earlier season may no longer comply. This chapter fixes the numbers that size the pack, ranks the kill-chain options, and sets out the enclosure thermal problem that decides whether the electronics survive a June afternoon in the Utah desert.

249 W
AVERAGE MISSION DRAW
50 kg rover, 35-minute mission
Bottom-up load table
S-URC27-RULES

3.7 kW
PEAK AT ABOUT 145 A
Sizes the fuse, contactor and wire, not the pack
Peak case
S-VENDOR-MOLICEL-P42A-DS

273–289 Wh
NAMEPLATE FOR ONE MISSION
400–500 Wh is the recommended pack
2027 mission set
S-BU-808

88 °C
INSIDE A SEALED BLACK BOX
From solar gain alone at 38 °C ambient
First-order model
S-URC27-RULES

Key takeaways

- **Average power sizes the battery; peak power sizes the protection.** The 2027 mission set needs **273 to 289 Wh** of nameplate capacity, and a **400 to 500 Wh** pack is the sensible buy. The same rover draws about **145 A at 3.7 kW** for a few seconds, which is what the fuse, contactor, connector and wire have to survive.
- **The E-stop is the highest-severity item in the subsystem.** The 2026 and 2027 wording requires one button to switch off **all** power from **any and all batteries**. A drive-only kill fails the rule.
- **The kill removes compute power too.** Filesystems and logs have to be designed for an unannounced power cut, not a shutdown sequence.
- **A sealed black box reaches about 88 °C** internally at 38 °C ambient. A sunshade and a filtered fan bring it to roughly 42 °C, and no pack survives a season at the first number.

What the subsystem must do

The rules state the safety requirement without ambiguity, and this page quotes it because every word drives a part choice:

> “All rovers shall have a red push button emergency stop that is readily visible and accessible on the exterior of the rover. It must be a push button, not a knob or rotating switch. This push button shall immediately stop the rover’s movement and cease all power draw from any and all batteries in the event of an emergency such as a battery fire. All systems on the rover must have the power switched off by this single E-stop.” S-URC27-RULES §2.a.v

Six constraints follow. The button is red and push-actuated. It is exterior and findable. It stops movement **and** removes battery power, which a wheel disable does not do. No second switch or software menu may stand between the button and a dead rover. Its action is immediate, so a shutdown sequence does not qualify. And it must work when the pack’s own electronics are the thing that failed, which rules out any path through a microcontroller.

The surrounding requirements set the operating envelope. Mobile deployables such as a mini-rover need their own E-stop, while a fuse is enough for anything that cannot move S-URC27-RULES §2.a.vi. No tether to the C2 station may carry power S-URC27-RULES §2.a.i. No air-breathing power source is allowed, but cooling fans are explicitly permitted and need not be custom S-URC27-RULES S-URC27-QA §2.a.iv RQ12. The mission clocks run from 20 minutes for Astrobiology to 60 minutes for Delivery, with no recharge in between S-URC27-RULES S-URC27-RULES §1.b.ii §1.c.i. Set-up is 15 minutes S-URC27-RULES §1.a.ii. Site power is 120 V at 60 Hz S-URC27-RULES §3.d.i. June ambient is 37.8 °C with dust and light rain

S-URC27-RULES §3.d.iii. And any intervention to swap a pack or reset a tripped protection device costs 20 % of that mission's points S-URC27-RULES §3.e.

The numbers that size the design

QUANTITY	VALUE	BASIS
Average mission power	≈249 W for a 50 kg rover on a 35-minute mission	Ten-row load table, by inference
Mission energy	≈146 Wh for that 35-minute mission	Same load table
Nameplate for one mission	273 Wh for the 60-minute Delivery; 289 Wh for Autonomy followed by Equipment Servicing	Mission energy, 25 % margin, 0.662 usable fraction
Recommended pack	400–500 Wh nameplate at 8S, 2.4–2.9 kg with BMS and enclosure	Covers the back-to-back case with 38–73 % headroom
Peak current	≈145 A, ≈3.7 kW at the sagged 25.6 V bus	Climb, arm hold and radio together, for a few seconds
Pack sag at peak	4.6 V on an 8S4P pack, 15.6 % of nominal	32 mΩ pack resistance S-VENDOR-MOLICEL-P42A-DS
Cell energy and mass	15.5 Wh and 70 g per 21700 cell; 496 Wh and 2.24 kg for 32 cells	Datasheet S-VENDOR-MOLICEL-P42A-DS
Precharge energy	2,000 µF at 33.6 V is 1.13 J; about 646 A without a resistor	Computed loop resistance; 100 Ω precharge gives 0.34 A and a 0.2 s time constant
Sealed-box temperature	about 88 °C internal at 38 °C ambient	Solar gain on a 0.5 × 0.5 × 0.25 m box
Cooling requirement	about 10.6 CFM for 60 W and a 10 K rise	Enclosure cooling rule of thumb, verified against a fan curve

Two numbers govern the whole design. The first is the average, which sets the pack. The second is the peak, which sets everything downstream of the pack: fuse, contactor, connector, wire gauge and the battery management system's trip curve. A team that sizes protection for the average melts, trips or welds on the first loose slope.

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Chemistry	Li-ion 21700 · LiPo pouch · LiFePO4	8S4P Li-ion 21700 at about 496 Wh and 2.3 kg of cells; a bought 8S LiPo if the first mission is under six months away; LiFePO4 only to avoid lithium fire risk, at 2.3× the mass
Voltage	6S · 8S · 12S · 24S	8S, at 29.6 V nominal and 33.6 V full. A 6S pack forces 1.33× the current through the same harness; 12S only if every converter and contactor is 60 V rated
Kill device	DC contactor · MOSFET e-fuse · controller disable · wireless	A hardwired contactor driven by the button, with precharge. A MOSFET is distribution protection, not the primary kill
Protection settings	fixed BMS trip · programmable trip with delay	A programmable BMS whose current limit sits above a legitimate two-joint stall
Distribution	central board with per-output sensing · fuse block and bought converters	Fused branches with per-output sensing if an electronics subteam exists; otherwise one 24 V bus and three bought converters
Thermal	passive · forced air · sunshade	Sunshade, white exterior and filtered forced air

Average power sizes the battery, peak power sizes the protection

The load table below is a defensible first-order model for a six-wheel skid-steer rover with a six-joint arm, one Jetson-class computer, a 2.4 and 900 MHz radio pair and a small science package. The structure is the deliverable; every number should be

replaced with a measurement in the first month of driving.

LOAD	QUANTITY	RAIL	AVERAGE EACH	DUTY	AVERAGE TOTAL
Drive motors and controllers	6	24 V	20 W	0.55	120 W
Arm joints, hold and move	6	24 V	8 W	0.5	48 W
Perception computer	1	12 to 5 V	20 W	1.0	20 W
Auxiliary computer and MCU	2	5 V	4 W	1.0	8 W
Radio and switch	1	12 V	9 W	1.0	9 W
Cameras, GNSS, IMU	6	5 and 12 V	2 W	1.0	12 W
Science package	1	isolated 12 V	15 W	0.8	12 W
Status LEDs and lights	1	12 V	8 W	0.6	5 W
Fans	4	12 V	1.5 W	1.0	6 W
Converter losses	Not applicable	Not applicable	8 % of the loads above	Not applicable	9 W
Total					**≈249 W**

The peak is a different calculation, and it comes from three sources. A 50 kg rover climbing a 25° loose slope at 0.5 m/s needs about 190 W electrical, which is nothing. Wheel slip and obstacle climbing at the controller's 20 A limit are the real drive case: six motors at 20 A is **120 A at 24 V, about 2.9 kW**. One shoulder joint holding 8 kg at 0.9 m needs about 71 N·m at the output, which through a 50:1 reducer and a torque constant of 0.06 N·m/A is about **23 A and 560 W per joint**, and two joints together are about 1.1 kW. Climbing while the arm holds a sample while the radio transmits is **≈145 A and 3.7 kW** for a few seconds.

That peak has four consequences. It is about **15 times** the average draw, so it sets the fuse, contactor, connector and wire while the average sets the pack. On an 8S4P pack it is an 8.6C discharge, inside the pack's 180 A rating but hot, and far beyond a single 30 A protection device or a 45 A connector. It sags the bus by 4.6 V, so converters need a wide input range and the low-voltage cut-off has to sit below the sagged voltage or the rover dies on the hill by design. And it is usually the **arm**, not the drivetrain, that causes a trip, because a stalled joint draws that current continuously while a person decides what to do. Missouri S&T's published protection trips at 80 A for longer than 10 ms and restores after 3 s, which would trip on a legitimate two-joint stall `S-MST-2023-POWER`.

Two fixes are cheaper than a bigger pack: interlock the firmware so the drive current limit drops while the arm moves, and give the arm brakes so a hold costs no current at all.

The E-stop is a pass/fail clause, not a lost point

The wording has tightened twice, and the history explains why this subsystem needs a procedure rather than a component:

SEASON	REQUIREMENT TEXT	WHAT CHANGED
2023	a kill switch, readily visible and accessible, that stops movement and ceases power draw	no colour, no actuator, no dimension `S-URC23-RULES` §3.a.v
2024	a red push button at least 3 cm across, mounted away from moving parts, plus a strongly encouraged heartbeat monitor	first dimension, first mention of a comms-loss failsafe `S-URC24-RULES` §3.a.v
2025	a red push button emergency stop, dimension language dropped	colour retained, detail removed `S-URC25-RULES` §3.a.v
2026 and 2027	adds "not a knob or rotating switch", "any and all batteries", and all systems off from this single button	the strongest version `S-URC27-RULES` §2.a.v

Four ways to remove power exist, ranked here by what a field judge can verify:

OPTION	STRENGTH	WEAKNESS
DC contactor in the battery lead, coil broken by the button	Physical gap a judge can see; fails open; survives a dead BMS	Coil current, mass, and it welds without precharge
High-side MOSFET or e-fuse	Microseconds, no wear, can report its own state	No physical gap; a shorted device keeps the rover live
Motor-controller disable	Free, already wired	Does not comply: converters, radios and computers keep drawing
Wireless kill	Useful as a supplement	Cannot be an exterior button on the rover, and depends on a working link

The recommended chain is short and entirely hardware. Pack, then a 125 A class-T fuse within 150 mm of the positive terminal, then the normally-open contactor with a precharge bypass across its contacts, then distribution. The latching red mushroom sits in series with the contactor coil through two normally-closed contacts, one of which also feeds a green loop-healthy indicator. A second, smaller enable contact from the protection board sits in series, never in place of the manual kill. Everything the rover owns is downstream of the contactor, which is what makes “all systems switched off” true.

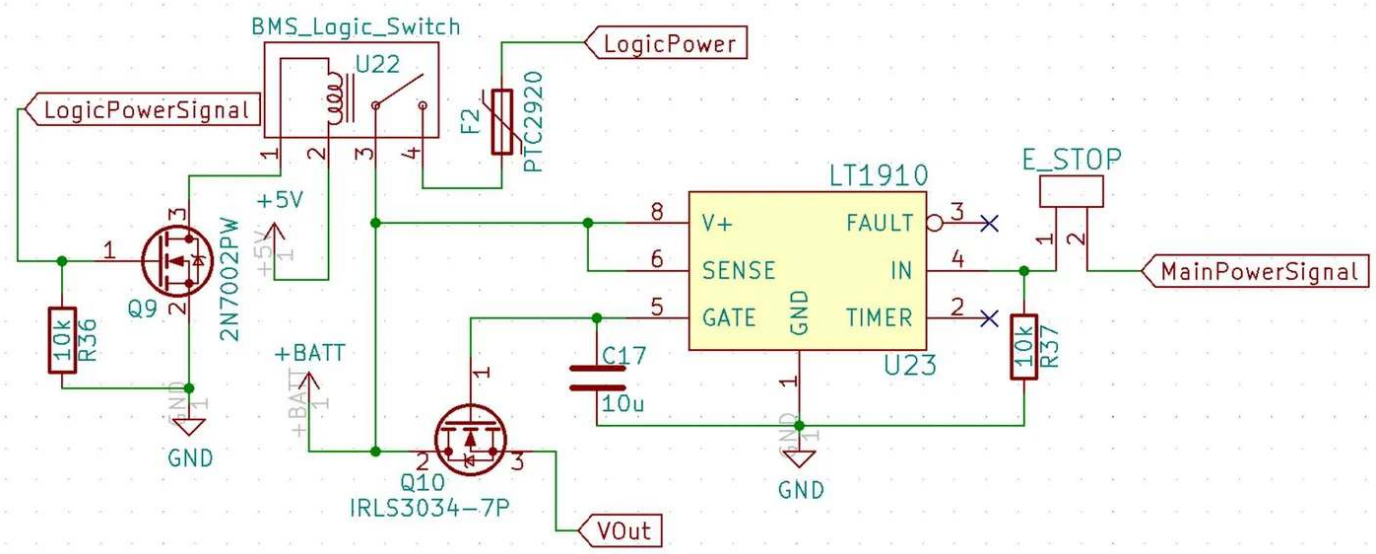
The timing targets are measurable: contacts open within **50 ms** of the press, bus voltage reaches zero within **100 ms**, and the wheels stop within **0.5 s** through brakes or friction.

§3.b.iv." aria-label="Open full-size figure">



The arrangement the rule asks for: a red mushroom on the rear electronics enclosure and a second on the front sensor box, both reachable by a person walking up to a stopped rover, which is the runner case in §3.b.iv.

Missouri S&T Mars Rover Design Team (MRDT), 2023. *Full resolution*



A real team's kill path: an LT1910 high-side driver whose IN pin is driven by the E_STOP net, with a P-channel MOSFET as the switch. Fast and telemetry-rich, but the stop is a logic signal into a driver, which is why a parallel hardware contactor is the safer primary.

Missouri S&T Mars Rover Design Team, 2023. *Full resolution*

A second design finding matters more than the schematic: mature teams end up with **more** than one stop path. Missouri S&T built a dedicated arm stop board and a wireless stop board on top of the mandated button, and a separate "rover insurance" board `S-MST-2023-POWER`. Extra paths are good practice as supplements and a compliance risk if used instead of the button.

Because the kill is literal, the compute dies with it. Logs therefore belong in a continuously flushed file or on a microcontroller, the root filesystem should be read-mostly, and the record of the kill belongs on the C2 side, where the moment telemetry stops is the timestamp. A supercapacitor that flushes a card for 200 ms is not power draw from a battery; a battery-backed uninterruptible supply is, and it does not comply.

The kill is measured, not assumed

- Continuity across the coil circuit, with the button pressed and released.
- Contactor open time on a scope, and every rail below 0.5 V within 200 ms of a press.
- Stopping distance under 0.5 m from 1 m/s.
- The rover stays dead after the button is released, until it is deliberately re-armed.
- Precharge completes before the main contactor closes.
- The whole set repeated with the pack at full charge and the arm moving.

Precharge, and why contactors weld

A motor bus is a capacitor bank, and closing a contactor onto an uncharged bank is a short circuit for milliseconds. For 2,000 µF at 33.6 V the stored energy is **1.13 J**, small in absolute terms, but with 52 mΩ of loop resistance the theoretical peak is **646 A**, limited in practice to a few hundred amps for a millisecond by the arc and the fuse. A welded contactor is the worst failure in the system: the button works, the coil clicks, and the rover stays live.

A 100 Ω, 25 W wirewound resistor across the contacts fixes it. Peak current becomes 0.34 A, the time constant is 0.2 s, and the bus is at 95 % in 0.6 s, so the main contactor closes with about 1.7 V across it and no arc. The symptom of skipping it is a click and a dead rover; the symptom of a precharge resistor that has gone open is a brownout on every power-up while the main contactor arcs.

Solar gain decides the enclosure

At 1,000 W/m^2 and an absorptivity of 0.9 for black anodised aluminium, the top face of a 0.5 × 0.5 × 0.25 m enclosure absorbs **225 W**. Natural convection and radiation off 0.75 m^2 give about 4.5 W/K, so the steady-state rise is **50 K**, or **88 °C internal at 38 °C ambient**. That single number explains why electronics work on the bench and fail at the competition site.

MEASURE	INTERNAL TEMPERATURE, SAME MODEL	EFFORT
Nothing, black sealed box	88 °C	none
Reflective sunshade and white lid	55 °C	a few grams of sheet and standoffs
White or polished exterior	55 °C	paint
Sunshade plus filtered forced air	42 °C	a fan, two vents, about 1 to 3 W
Separate battery compartment	pack cooler than the electronics	volume and a second thermal path

Two cautions attach to this table. A fan is a silent failure, so monitor tachometer output or temperature and never let forced air be the only thing between the computer and thermal shutdown. And the dust is fine bentonitic clay that will clog a filter in a day of driving, so filter cleaning belongs in the daily routine rather than the pre-competition checklist.

The pack has its own thermal rules. Dwell above 30 °C is elevated temperature for ageing purposes, a cell held at 40 °C and full charge keeps about 65 % of its capacity after a year `S-BU-808`, and charging is only allowed between 0 and 45 °C `S-VENDOR-MOLICEL-P42A-DS`. Self-heating at peak is real: 145 A through 32 mΩ is 0.67 kW of heat inside the pack for the duration of the peak, so peaks have to be short or the pack has to be larger.

Interfaces that bite

- **Mobility.** A 6S pack forces 33 % more current through the same harness for the same power, and a regen-capable controller can over-voltage a full pack on a descent. The drive branch fuse, connector and precharge all follow from the controller's input capacitance and current limit.
- **Arm.** Long holds without brakes are a continuous current the pack and the protection must tolerate, and they heat a motor the size of a fist with hundreds of watts. Brakes remove the load, the heat and the trip.
- **Communications.** The kill removes the radio, so C2 sees silence after a press and must not read it as a fault. The radio also wants a quiet 12 V rail, because sharing with a motor controller costs packets every time the drive moves.
- **Compute.** Ten to forty watts of compute with peaks near 60 W need a wide-input converter, bulk capacitance and a low-voltage cut-off below the sagged voltage. Motor stall current browning out the cameras is the classic coupling failure.
- **Delivery and the drone.** A spare drone battery carried on the rover adds mass and cost, and the phrase "any and all batteries" may be read to cover it. A written ruling is worth having before the System Acceptance Review `S-URC27-QA` RQ14.
- **Systems and environment.** A 500 Wh pack with BMS and enclosure is about 3 kg and is routinely underestimated in a mass budget. Filter access, pack mounting and drip loops belong in the enclosure drawing, not in the field routine.

Failure modes

MECHANISM	WHAT HAPPENS	PREVENTION
Protection trip during an arm stall	Sustained 80 to 150 A against a stuck object trips the pack and the rover goes dark mid-task	Trip above legitimate stall current, or current-limit each joint, and add brakes
Brownout on a climb	Pack sag drops the bus below converter drop-out and the computer resets	Oversize the pack, use wide-input converters, and set the cut-off below the sag
Welded contactor	No precharge, so a later close welds the contacts and the button cannot remove power	Precharge resistor, sequenced closing, a DC-rated contactor and an anti-spark connector
Kill loop opened by vibration	An inline connector or crimp fails on rough ground, so the rover dies or the button stops working	One harness, at most one connector in the loop, strain relief, and a continuity check every field day
Hot connector	An undersized connector on a 100 A branch opens the circuit or starts a small fire	Size connectors to at least 1.5× the steady current, and check temperature after a full-load test
Dust ingress	Fine clay bridges contacts and produces intermittent faults after the first dusty run	Sealed connectors outside the box, drip loops, conformal coating and a cleaned filter
Over-discharge during storage	Quiescent drain takes one cell group below its limit and the pack is unusable on the morning of the event	A physical disconnect and a stored-charge schedule
Charging failure at the site	A 230 V-only charger on a 120 V supply, or a wrong connector	Dual-voltage charger, a spare, and a bench test on a 120 V circuit
Filesystem corruption after a kill	The hard cut damages the filesystem exactly when the log matters	Read-only root, a journalled data partition, or logging to a microcontroller

The honest limit of that table is that no published URC post-mortem names an electrical failure with a root cause. It is a mechanism-first list, confirmed from the rules and from one detailed team paper rather than from a field incident.

Shortlist for a first build

ITEM	CHOICE	WHY
Pack	8S4P Li-ion 21700, about 496 Wh, 2.3 kg of cells, bought programmable BMS	Enough energy for the mission set, 180 A of continuous rating for the peak, and no assembly beyond spot-welding
Voltage	8S, 29.6 V nominal, 33.6 V full	The centre of the converter and contactor market; 6S costs 33 % more current
Kill	Latching red mushroom into a DC contactor coil, with precharge	Inspectable, fails open, and no firmware in the path
Distribution	One 24 V bus, a fuse block, three bought converters for 12 V, 5 V and isolated 12 V	A custom board pays off only if an electronics subteam can freeze it three months out
Connectors	XT90-S at the pack, Powerpole for 12 V field power, M12 or Deutsch DT for exterior sensors	One family per voltage, and the correct crimper owned for each
Charging	120 V-native 600 W charger and two packs, charged at 1C between missions	A second pack removes the recharge from the critical path
Thermal	Reflective sunshade, white lid, filtered forced air, pack in its own compartment	The cheapest 46 K of the whole design

Open problems

- **What is the real peak current?** Published team numbers are averages. A logging shunt and three full mock missions answer it, and every protection setting depends on the answer.

- **Do the protection thresholds clear a legitimate peak?** Inject a controlled 150 A pulse and confirm no trip, then a 200 A pulse and confirm a trip.
- **How hot does the pack run in the real enclosure?** The solar-gain model is one-dimensional. Thermocouples on the pack, the computer heatsink and the enclosure air on a 38 °C day replace it with data.
- **Must a carried spare drone battery be killed by the single button?** The rule says any and all batteries, and a stowed pack is not connected to anything. A written ruling before the review settles it.
- **Does the contactor or the MOSFET survive dust and vibration better?** Both designs exist in the record and no comparative reliability study has been published. A vibration and dust test with contact-resistance measurement before and after is the honest comparison.
- **Does a dust filter survive a day of driving?** Measure airflow drop across a filter in a dust chamber or on a dusty field day.

Further reading

- Missouri S&T is the only team in the record with a detailed published electrical paper, covering the pack, the protection logic and a 19-output distribution board.
- Mobility sets the drive current limit and the regen behaviour that this subsystem has to absorb.
- Communications and networking carries the other half of the E-stop consequence: the radio dies with the rover.
- Mass and power places the pack in the 50 kg budget and works the pack-sizing arithmetic in full.
- Interfaces and couplings follows the power, arm and thermal couplings with numbers.

SUBSYSTEMS · CHAPTER 4.7

Compute, sensors and perception

Perception answers four questions fast enough to score: where the rover is, what the object in front of it is, what the astronaut just commanded, and whether the view is good enough to drive. Each question has a number attached, and the numbers are smaller than they look. The most useful finding on this page is that the 2027 post tag is a last-five-metres sensor rather than a twenty-metre one, and that the mission design already accounts for it.

5 px

PER MARKER MODULE, MOVING ROVER

30 px across a six-module tag

Measured

S-P407-ARUCO-TEST

8.3 m

POST TAG AT 1080P AND 60°

4.6 m at 720p and 69°

20 cm post

S-URC27-RULES

±53 mm

STEREO DEPTH AT 3 M

±594 mm at 10 m; the model is optimistic by 1.6×

D435i class

S-RS-D435

67 TOPS

PERCEPTION COMPUTE

7 to 25 W, USD 249, no carrier board

Orin Nano Super

S-NV-JETSON-ORIN

Key takeaways

- **Pixels, not compute, bound perception.** A moving rover needs **5 pixels per module** on an ArUco tag, measured across five degradations. That gives **8.3 m** on a 20 cm post tag through a 1080p camera at 60°, and **4.6 m** through the common 720p wide camera.
- **The tag does not have to work at 20 m.** Target coordinates are exact, so GNSS carries the rover to the post and the tag closes the final metre. A 4K driving camera bought to see the post from far away solves a problem the rules have removed.
- **Depth is a close-range sensor.** Stereo error grows with the square of range: ±53 mm at 3 m becomes ±594 mm at 10 m, and a rolling shutter turns motion into **1.3 m** of depth error at 3 m.
- **The compute is adequate and the cooling is not.** 67 TOPS runs every model the missions need at 25 W; a 40 W module in a sealed box at 38 °C sits at about **65 °C** before solar gain is counted.

What the subsystem must do

The 2027 rules ask for four capabilities with numbers attached. The rover must stop within **1 m** of a route target and within **3 m** of the astronaut S-URC27-RULES §1.e.xvii. It must detect a **20 cm** post tag, a **2 cm** keyboard tag and a **1 cm** key tag S-URC27-RULES §1.e.xii. It must recognise five astronaut commands, one of which scores 15 points only as a gesture S-URC27-RULES §1.e.vi. And it must give three of the four missions a driveable view, because those missions are teleoperated S-URC27-RULES §1.d.ii.

Two rules then set the architecture. The team must bring and run **its own RTK base**, with no correction service and no internet in the field S-URC27-QA RQ11. And the second route target sits **behind a hill, out of radio line of sight**, so anything score-critical has to run on the rover rather than at the C2 station S-URC27-RULES §1.e.xvi. A base-station computer may help, but it cannot be in the critical path S-URC27-QA AUQ2.

The numbers that size the design

QUANTITY	VALUE	BASIS
Pixel budget, moving rover	**5 px per module**, 30 px across a six-module tag; 12 px is the identification-only floor	Measured across five degradations S-P407-ARUCO-TEST
Post tag, 1080p at 60°	**8.3 m** reliable, 20.8 m for identification only	1,663 px focal length, 15 cm tag extent
Post tag, 720p at 69°	**4.6 m** reliable, 11.6 m for identification only	931 px focal length, the common student camera
Post tag, 4K at 60°	16.6 m reliable	3,325 px focal length
Key tag, any camera	**0.31 to 0.78 m**	1.67 mm per module; sub-metre at 5 px per module
Stereo depth, D435i class	±53 mm at 3 m, ±594 mm at 10 m; ±1.3 m at 3 m while moving	Error model with 0.2 px disparity error, optimistic by 1.6× S-RS-D435
RTK relative and cold fix	better than 1 cm, 35 s	Receiver figures S-ARDUSIMPLE-F9P
Base coordinate error	maps one to one into the 1 m stop tolerance	By inference from RTK geometry
Perception compute	67 INT8 TOPS at 7 to 25 W	Orin Nano Super 8 GB S-NV-JETSON-ORIN
Uncompressed video	1080p30 is 186 MB/s; six 720p30 streams are 498 MB/s	More than one USB 3 controller can carry
Astronaut size in frame	84 px at 20 m through 720p at 69°; 150 px through 1080p at 60°	A suited adult at 1.8 m

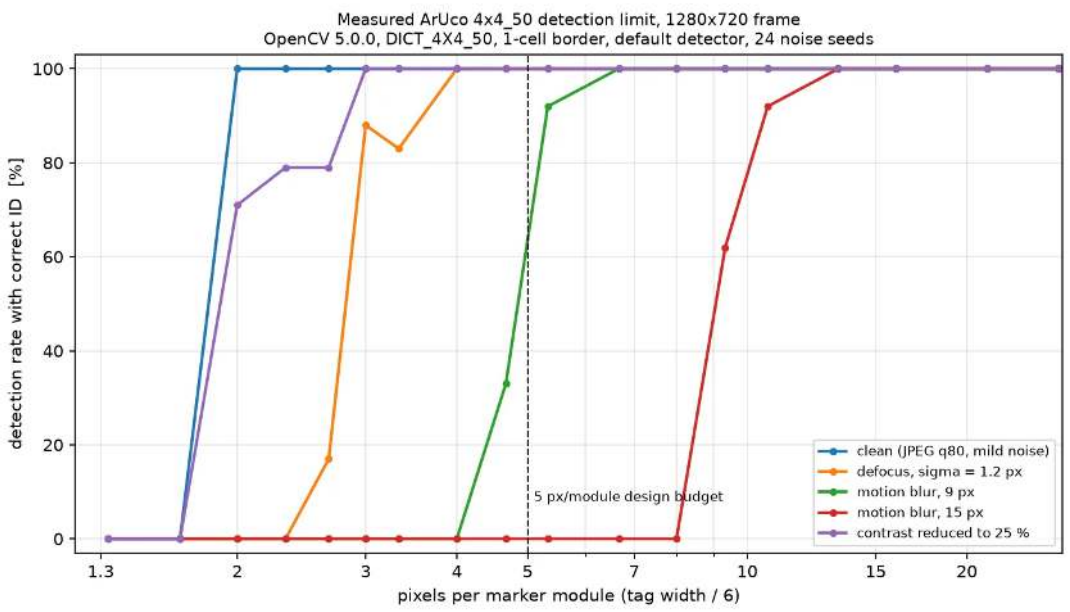
Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Perception computer	Orin Nano Super 8 GB · Nano 4 GB · Orin NX · AGX · Pi 5 with an accelerator · x86	Orin Nano Super 8 GB: enough for detection, pose, ArUco and one stereo stream at 25 W, and no carrier board
Camera split	one wide · wide and narrow · stereo · wrist · zoom · thermal	Two or three 1080p global-shutter driving cameras at 60 to 90°, plus a wrist camera fixed at 0.2 to 0.5 m
Depth	D435i · D455 or D456 · OAK-D · ZED 2i · GMSL	OAK-D or D455 class, with the two cameras at least 75 mm apart, global shutter
LiDAR	Mid-360 · Avia · 2D lidar · none	None in year one; add a 360° unit once GNSS and stereo already reach the 1 m stop
GNSS	two F9P and an own base · one F9P and a borrowed base · single band	Two ZED-F9P receivers with multiband antennas; single band cannot meet 1 m
Heading	two receivers on a fixed span · calibrated magnetometer	Magnetometer at 2 to 3°, which the rules allow S-URC27-QA RQ3; dual antenna only if the panorama cardinal direction must be exact
Mast	carbon tube · glass fibre · aluminium · two masts · isolators	30 mm carbon tube, 0.8 to 1.2 m, folding, with the GNSS ground plane on top and the camera plate isolated
Compute split	one computer · single-board computer and accelerator · single-board computer and MCU	Jetson plus an MCU on CAN; all three documented teams split the load and none runs one computer for everything

The pixel budget is the subsystem

An ArUco `4x4_50` tag with a one-cell border is six modules across. With the cell fixed at 2.5 cm, the black extent on a 2027 post is **15 cm inside a 20 cm face**, inferred from the face and cell dimensions in the rules. If the tag fills the whole face instead, every range below improves by a factor of 4/3, which makes one photograph of the real post worth taking.

Detection was measured rather than assumed. OpenCV 5.0.0 on a synthetic sunlit-sand background, with 24 noise seeds per point, decodes a clean tag at 12 px, needs 24 px under defocus, **40 px with 9 px of motion blur**, and manages 18 to 24 px at 25 % contrast `S-P407-ARUCO-TEST`. That table converts directly into a design budget of 5 px per module.



Measured detection rate with the correct identity across five degradations, from a clean tag to 15 px of motion blur. The dashed line is the 5 px per module budget; motion blur costs more range than defocus or low contrast.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

Two consequences follow, and both are counter-intuitive.

Full sun is not the blur problem. At the roughly 10^5 lux of a June afternoon the exposure is short: 1 m/s at 1/1000 s through a 1080p camera at 60° blurs **0.17 px** at 10 m, which is invisible. The sun problems are **contrast**, where a shaded tag sits against bright ground, and **flare**, where the sun is near the field of view and no exposure setting helps. The answer is a hood on every camera and an exposure set so that the bright background is not clipped.

The keyboard is tractable and the key is not. The 2 cm keyboard tags are twice the 1 cm numbers, which is about 1 m through a 1080p wrist camera against the 0.2 to 0.5 m the key needs. That is why teams put a camera on the arm rather than relying on the mast, and it is why the wrist camera is an arm requirement created by a perception number.

The same arithmetic sizes the human tasks. A suited astronaut is 84 px tall at 20 m through 720p and 150 px through 1080p, which is enough for driving to the astronaut and for “Come!”, since both need only a bearing and a range. Pose and gesture models want 150 to 250 px, so the gesture belongs at 3 to 8 m. That is a team decision rather than a rule `S-URC27-QA` AUQ10, and voice is an explicitly encouraged fallback `S-URC27-QA` AUQ6. The hammer is a 0.91 kg object about 300 mm long, which makes it a 5 m object for detection and a 3 m object for grasp planning `S-URC27-HARDWARE`.

Depth is close range, and the base coordinate owns the whole metre

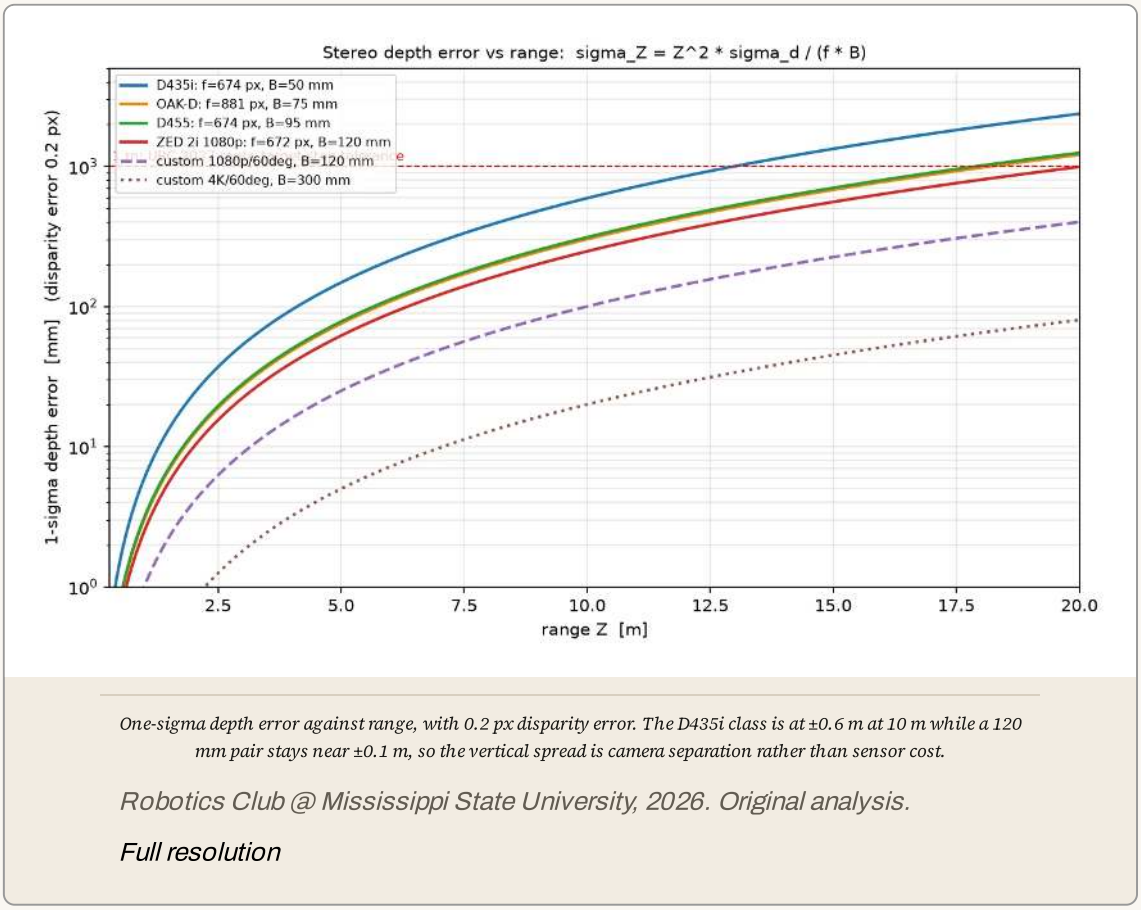
Stereo depth error grows as the square of range divided by focal length and by the distance between the two cameras. With the cameras 50 mm apart that is ±53 mm at 3 m and ±594 mm at 10 m, and the model assumes 0.2 px of disparity error while real hardware needs about 0.32 px, so the real error is roughly **1.6× the table** `S-P407-STEREO-CALC` `S-RS-D435`. Three design rules follow.

Depth must not be the primary localiser at range. A D455 at 10 m is within ±0.31 m on one measurement, which is a third of the stop tolerance, and the worst case is worse than the tolerance.

Camera separation beats pixels. Moving the two cameras from 50 mm apart to 120 mm apart improves depth at 10 m by 2.4×. Moving from 720p to 1080p at the same field of view improves it by about 1.0×, because the field of view sets the focal length in pixels. A narrow lens and a wide camera spacing are the only route to centimetre depth at 10 m.

Rolling shutter is unusable on a moving stereo pair. At 1.5 m/s with a 15 ms readout the scene moves 22 mm during a frame, which is about 4.9 px of disparity error, which is **1.3 m of depth error at 3 m.** Global shutter and hardware sync are not a luxury on

a driving pair; they are the difference between depth and noise.



One-sigma depth error against range, with 0.2 px disparity error. The D435i class is at ±0.6 m at 10 m while a 120 mm pair stays near ±0.1 m, so the vertical spread is camera separation rather than sensor cost.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

The 1 cm key task never uses stereo. A mono camera with ArUco pose estimation gives range directly from the known tag size, at about **0.3 % of range**, or 1 mm at 0.3 m, which is enough for a 2 to 3 mm insertion.

RTK is only as good as the base station's own coordinate, and a 1 m base error moves every waypoint by 1 m, which is the entire route-target tolerance. The mitigations in order of cost are a **30 to 60 minute** survey-in, because 35 s is the time to first fix rather than to a decimetre absolute solution; reusing one base position file all week so the map and the rover share a single error; and letting the visual tag arbitrate at the end, which is what the task design already does.

One coupling matters most. The correction link is line-of-sight limited, so behind the hill the rover loses RTK corrections in the same shadow that blocks the C2 link. The metre at the second route target is bought with dead reckoning, a pre-loaded route and the tag, not with a live fix.

Compute is adequate; heat is the constraint

An Orin Nano Super 8 GB runs 67 INT8 TOPS at 7 to 25 W and needs no carrier board `S-NV-JETSON-ORIN`. In practice that is enough for person detection at 60 to 120 fps, pose at 30 to 60 fps, ArUco detection at 60 fps or better on the CPU, one stereo stream and a SLAM back-end. It is not enough to add a 4K stream and a second large model at the same time, and 8 GB is the ceiling rather than a margin.



The reference perception computer: 67 TOPS at 7 to 25 W in a developer kit that needs no carrier board. The heat has to leave the enclosure, and that is the part of the choice that fails in the field.

NVIDIA, 2026. *Original source* *Full resolution*

The thermal arithmetic is the failure that goes unnoticed. A 40 W module in a sealed 300 × 250 × 150 mm box, with natural convection at about 5 W/m²K over 0.3 m², sits **27 K above ambient**, or about **65 °C at 38 °C**. That is before solar gain on the enclosure itself, which the electrical design also has to remove. Nothing crashes; the model simply runs slower, and the team reads the lost frames as a software problem `S-URC27-RULES §3.d.iii`.

The fixes are the ones the power page already asks for: size the enclosure for the **sustained** power mode rather than the peak, add a fan with a filtered inlet, which the rules permit `S-URC27-QA` RQ12, and soak-test at 40 °C for 40 minutes rather than 5.

The mast carries the consequences. It holds the cameras, the GNSS ground plane and often the radio antennas, it must fold inside the 1.2 m cube, and the terrain excites it at 2 to 15 Hz, so its first mode has to sit above 20 Hz to keep the images still. Every joint in a folding mast is a stiffness question and a cable-fatigue question at the same time.



A folding mast in service: the slender tube rises behind the arm and carries the antenna and camera hardware, with the cable run left outside the tube. It is the tallest, most fragile item on the rover when stowed.

Monash Nova Rover (Banksia, 2026), 2026. *Full resolution*

Interfaces that bite

- **Autonomy.** Perception owes a tag identity, range, bearing, covariance and timestamp at 10 to 30 Hz, never a six-degree-of-freedom pose computed from a 30 px tag. It owes an obstacle map at 5 to 10 Hz with 0.1 to 0.2 m resolution over a 10 m radius.
- **Arm.** The 1 cm key tag creates the wrist camera, and the wrist camera lands on the arm's mass and cable life. A cable through every joint is a fatigue item, so the camera's driver board belongs in the wrist and the service loop has to be designed rather than discovered.
- **Power.** Compute draws 10 to 40 W with peaks near 60 W, and the kill has to cut it or the safety rule is not met. The classic coupling is motor stall current browning out the shared USB rail and taking the cameras with it.
- **Communications.** Three to six concurrent 720p30 streams at 100 to 300 ms is the video load, plus RTCM3 corrections at 0.5 to 2 kB/s. The second route target needs zero link, which is why the map and route must be on the rover.
- **Structures and environment.** The mast, the camera plate and the GNSS ground plane are one assembly, and the camera count is an operator-load decision as much as a perception decision: one documented 2026 rover carries 11 cameras `S-WVU-SAR2026`, which costs bandwidth and attention as well as mass.

Failure modes

MECHANISM	WHAT HAPPENS	PREVENTION
Sun glare and backlight	The target is in shadow against bright ground, or the lens flares, and detection stops	A hood on every camera, exposure set for the bright background, and the exposure value logged
Dust on the lens	A thin film kills contrast and a single grain becomes a permanent false blob	Recessed lens with a replaceable flat window, a cleaning step in the pre-run check, and a flat-field image each morning
Calibration drift	A knock to the camera plate changes an extrinsic by half a degree, which looks like a slow depth drift	Machined mounts with dowel pins, torque marks, and a checkerboard check at the start of each field day
USB dropout	A camera disappears mid-run, the node dies, and the stack blocks waiting for a frame	Locking connectors, a powered hub, one camera per controller, and a watchdog with a degraded mode
Thermal throttling	Frame rate falls after 10 to 20 minutes in a sealed box and nothing reports an error	Sustained-power enclosure sizing, filtered forced air, and a 40-minute soak at 40 °C
GNSS shadow	Position jumps by a metre near a hill face or an AR post and RTK drops from fixed to float	Antenna on a ground plane at the mast top, fix type logged and surfaced, and dead reckoning through the shadow
Base coordinate error	Every waypoint is a metre out and the rover stops 1.1 m from the target	A proper survey-in, one base file for the week, and the visual tag as the final arbiter
Tag not found at the target	The rover arrives on the exact coordinate and the tag is behind it or edge-on	A slow yaw search at the waypoint; the tags repeat on all three post faces
Collision with the post or the astronaut	A wheel catches the post, or the astronaut stops a rover that came too close	A 1.5 m standoff, a low-speed approach mode, and the astronaut in the map as a 1 m obstacle

Shortlist for a first build

ITEM	CHOICE	WHY
Perception computer	Orin Nano Super 8 GB	67 TOPS at 25 W, USD 249, and it runs every model the missions need
Real-time computer	STM32 or Teensy class MCU on CAN	A deterministic loop for drive, kill state and the status indicator
Driving cameras	Two or three 1080p global shutter at 60 to 90°	8.3 m of post-tag range and a usable input for depth and odometry
Depth	One stereo pair with the cameras at least 75 mm apart and global shutter	±76 mm at 5 m is enough for obstacle avoidance
Wrist camera	1080p at 70°, fixed focus at 0.2 to 0.5 m	The only way to read the 1 cm key tag
GNSS	Two ZED-F9P receivers, two multiband antennas, a 900 MHz correction link	The own base is mandatory, and relative accuracy is better than 1 cm
Heading	Calibrated magnetometer	Allowed at 2 to 3°, and adequate for a panorama label
LiDAR	Optional 360° unit, added later	It buys a terrain map but adds extrinsics, time sync and tuning that a first build cannot absorb early
Mast	30 mm carbon tube, folding, isolated camera plate, ground plane on top	Above 20 Hz keeps the images still and keeps the assembly inside the cube

Open problems

- **Is the post tag 15 cm or 20 cm across?** The rule fixes the face and the cell but not the margin. One photograph of a real post at the first practice task settles a 4/3 factor on every range on this page.

- **What is the real detection range in Utah sun?** The measured numbers are synthetic, with uniform degradations and no glare. A printed 20 cm tag at 5, 10, 15 and 20 m, logged against the time of day and the angle to the sun, replaces them.
- **How accurate is the base coordinate?** Survey-in for 30, 60 and 120 minutes at a known point, or run two bases and compare. The error maps one to one into the stop tolerance.
- **Does the computer throttle in the real enclosure?** The thermal figure is a one-dimensional estimate. A 40-minute soak at 40 °C with the real enclosure and the real workload is the test.
- **Does gesture recognition work on a baggy suit at 3 to 8 m?** No competition data exists. Film 200 gestures from a real suit at four ranges in sun and count the false positives per minute.
- **Does a 360° lidar pay for itself?** The teams using one also win, which is selection bias rather than proof. The same rover with and without it, on the hill route, measured by the share of runs reaching the 1 m stop, is the comparison.
- **How often does a camera drop out?** No team publishes a rate. Log every dropout across 20 hours of driving with the arm stalling.

Further reading

- WVU publishes the most detailed perception stack in the record, including a learned marker detector tested to 12 m `S-WVU-SAR2023`, and Monash shows the wrist-camera and mast arrangement in service. Missouri S&T documents the split between a central computer and motor controllers on CAN.
- Software stack and autonomy consumes every interface this page produces.
- Arm and end-effectors owns the wrist camera mount and the cable route that this page's 1 cm requirement creates.
- Power, electrical architecture and safety owns the thermal design of the enclosure that the computer sits in.
- Communications and networking sets the video budget that the camera count has to fit.
- Interfaces and couplings follows the mast, camera and compute couplings with numbers.

SUBSYSTEMS · CHAPTER 4.8

Software stack and autonomy

Autonomy decides 100 points of the 2027 score, and it is the only subsystem that can lose a task for a reason invisible in the hardware. Two route targets are exact coordinates inside 1 m, the hill has no guaranteed approach, and one target sits out of radio line of sight by design. The stack that wins is unremarkable in framework terms and precise in three places: the error budget at the stop, the terrain model on the hill, and the state machine around the astronaut.

5 of 100

SCORE WITHOUT AN AUTONOMY STACK

The suit camera and microphone are the whole score

2027

S-URC27-RULES

1 m

ROUTE-TARGET STOP TOLERANCE

Two 25-point targets, taken in either order

2027

S-URC27-RULES

0.25–0.55 m

POSITION ERROR AT THE STOP

Root-sum-square of six sources, estimated rather than measured

2027

S-URC27-QA

15

FOLLOW! POINTS FROM A GESTURE

Device 5, sign 5, audio 10, gesture 15

2027

S-URC27-RULES

Key takeaways

- A perfect chassis with no autonomy scores **5 of 100**, the EVA suit camera and microphone. Autonomy is not a bonus task.
- The **1 m** stop is an error budget, not a feature. Six sources root-sum-square to **0.25–0.55 m**, and heading error multiplied by distance dominates as soon as the rover is more than 10 m from its last fix.
- The framework is not the differentiator. **9 of 11** teams with public code use ROS 2, and UMD Loop publishes the most textbook-correct Nav2 configuration in the field and finished **32nd overall in 2026** S-URC26-SCORES . Missouri S&T uses **no ROS at all** and won 2025 and 2026.
- Follow! is a ladder, not a rung. Mode switching inside the mission is allowed [AUQ6], so a design that attempts gesture, falls back to audio and ends at the button adds probability at each step and never subtracts points.

What the subsystem must do

The Autonomy Mission is 40 minutes and 100 points, and software touches all of it S-URC27-RULES §1.e.i. Two route targets carry 25 points each, must be reached within **1 m** of exact coordinates in either order, and require an arrival indication the judge can see S-URC27-RULES S-URC27-QA S-URC27-QA §1.e.xvii AUQ16 AUQ18. The astronaut sub-mission carries the other 50: a drive to a suited team-mate, five commands in binding order, a 15-point Follow!, a 5-point Stay!, a 10-point Fetch!, a 5-point Come! and a 5-point Give! S-URC27-RULES S-URC27-RULES S-URC27-QA §1.e.iv §1.e.x AUQ13. One target sits behind the hill and out of radio line of sight, so the stack must keep deciding with no link at all S-URC27-RULES §1.e.xvi.

Leaving autonomy is priced. An operator stop costs **20 %** of that task, teleoperation back to a previously visited point costs **50 %** capped at 50 % per task, and a physical intervention costs **20 % of the mission** with the clock still running S-URC27-RULES S-URC27-RULES S-URC27-RULES §1.e.xi §1.e.xviii §3.e.viii. The rover's own autonomous stop is free and the command may be re-given S-URC27-RULES §1.e.xi, which turns self-recovery from an error path into a feature worth engineering.

One requirement is not written as one. Judges score hands-off behaviour by observation rather than by reading code S-URC27-QA AUQ21, so the mode topic, the rear LED and the operator posture belong to the software contract S-URC27-RULES §1.e.ii.

The numbers that size the design

The budget, ladder and latency rows are engineering estimates from stated parameters. The rule, framework and configuration rows are sourced.

QUANTITY	VALUE	BASIS
1 m route stop, error allocation	root-sum-square **0.25–0.55 m**, leaving 0.45–0.75 m margin	Six-source budget, engineering judgement
Stop distance at 1.0 m/s, 0.3 m/s²	**≈1.7 m**, twice the tolerance	Kinematics: the stop is a controlled terminal approach
Follow! ladder	gesture **15** · audio **10** · sign **5** · device **5**, maximum 15	The method, not the capability, sets the score `S-URC27-RULES` §1.e.vi
Expected Follow! value	gesture **8.25** · audio **7.00** · device **4.75** points	Points multiplied by an assumed success probability
Teams on ROS 2	**9 of 11** with public code	Read from the public code, last commit checked 2026-10-25
ROS 2 Humble end of life	**May 2027**, one month before the 2–5 June finals	`S-ROS-REP2000`
ArUco at 10 m, 1080p/60°	**25 px** on the 15 cm tag, 4.2 px per module	A terminal aid only, inside about 5 m
Glass-to-glass latency	**80–230 ms** one way; published guidance below 500 ms	Capture and encode 40–90 ms, radio 20–80, decode 20–60 `S-TELEOP-LATENCY-IEEE`
UMD Loop published configuration	Smac Dubin, 1.5 m turn radius · pure pursuit 1.0 m/s, 2.5 m lookahead · goal tolerance 1.0 m · DEM slope cap 15°	A real 2027-shaped Nav2 setup `S-GH-UMD-ATHENA`

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
OS and middleware	Jazzy 2024–2029 · Humble, end of life May 2027 · Rolling · ROS 1 · a custom RoveComm-class protocol	Ubuntu 24.04 with ROS 2 Jazzy when two or more software people continue; Humble only if the JetPack image forces it, with migration planned for spring 2027
Localisation	dual-antenna RTK with a dual EKF · custom Kalman filter · wheel odometry · visual or lidar odometry	Dual-antenna RTK, robot_localization dual EKF, wheel odometry and a magnetometer as a third yaw check
Map source	GNSS waypoints only · DEM slope costmap · onboard SLAM · hybrid	A 1 m DEM into a slope costmap plus a live obstacle layer; waypoints alone fail the hill by construction
Global planner	NavFn · Smac 2D with Dubin · Route Server · custom A · *RRT*	SmacPlanner2D over the DEM costmap, with the Route Server where graph-level semantics matter
Local planner	Regulated Pure Pursuit · MPPI · DWB · lattice · custom model-predictive control	Pure pursuit at or below 1.0 m/s with a 2.5 m lookahead; MPPI when someone enjoys tuning critics
Behaviour layer	Nav2 behaviour tree · hand-written state machine · mission sequence	Nav2 behaviour tree with three to five custom nodes: LED state, task start and end, person visible, pick visible
Human perception	YOLO · MediaPipe Pose · depth clustering · lidar clustering	YOLO plus depth for range, MediaPipe Pose for gesture, designed for the 6–8 m band
Simulation	Gazebo · Unity · Unreal 5 · a viewer built from the flight code	Whatever runs the flight code, because a general simulator that has to be fought into shape for a semester loses
Release discipline	tagged competition build with CI · branch per field test · no tags	Monorepo, pull-request review, a tagged competition release and a field-test branch

The 1 m stop is an error budget

The tolerance is a **total** error, so every contribution has to be named before it can be reduced. The allocation below is engineering judgement rather than measurement, and it is the calculation a first build should reproduce and then improve.

ERROR SOURCE	BUDGET	WHY
Base-station position, RTK fixed	0.05–0.15 m	An own base station is required, because no correction service is available S-URC27-QA RQ11
Rover position, RTK fixed	0.02–0.10 m	Fix noise and multipath near posts and rocks
Antenna-to-body lever arm	0.02–0.05 m	Missouri S&T subtracts a measured 0.20 m northing and 0.85 m altitude offset S-MST-DOCS-NAVBOARD
Heading error multiplied by distance	0.05–0.30 m	1° at 3 m is 0.05 m; at 20 m it is 0.35 m
Controller overshoot and stop detection	0.15–0.35 m	A 1.0 m/s approach at 0.3 m/s² needs about 1.7 m to stop
Target coordinate interpretation	0	The coordinates are exact S-URC27-QA AUQ16
Root-sum-square	**0.25–0.55 m**	0.45–0.75 m of margin against the rule

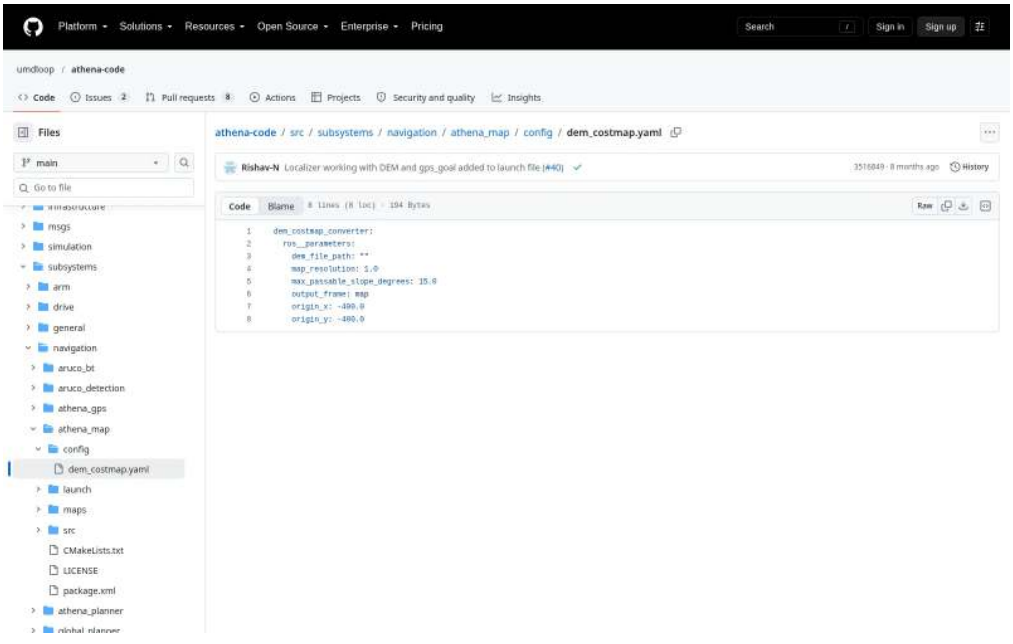
Three consequences follow, and each one changes a design decision.

- **ArUco is a terminal aid, not the navigation source.** The 15 cm black tag is about 25 px wide at 10 m through a 1080p camera with a 60° field of view, which is 4.2 px per module. GNSS closes the route; the tag confirms the last few metres.
- **Heading dominates once the rover is 10 m or more from the last fix.** One degree of yaw error is 0.05 m at 3 m and 0.35 m at 20 m, which is why dual-antenna RTK yaw with a magnetometer as a third check is worth its mass.
- **The last metre is a control problem.** Inside 3 m the terminal approach is made at or below 0.3 m/s, the stop is commanded at 0.6 of the tolerance, and a 3 s hold gives the judge time to see the green LED. UMD's 1.0 m goal tolerance only works because the controller closes that last metre.

The hill turns a straight line into a contour

A slope limit on a 1 m elevation model is what makes route finding a real task. At the 15° passable-slope cap that UMD Loop uses, the direct 200 m line to a hilltop target becomes a 400–600 m contour-following route, so the planner has to be cost-aware and the time budget has to absorb the detour. Nothing in the rules guarantees a shorter path: not all approach directions will be navigable S-URC27-RULES §1.e.xv.

The dependency most teams take on is the wrong one. URC's own elevation model is described as an attempt with an unconfirmed timeline S-URC27-QA AUQ20, while USGS 3DEP serves 1 m coverage that has been verified at the site. Building against the public model costs an afternoon of raster work and removes a schedule risk. The only public implementation in the field is eight lines long, and those eight lines are worth copying.



Eight lines are the whole hill capability: a 1 m map resolution, a 15° passable-slope cap, and an origin offset 400 m in each direction. The 15° is a choice, not a measured limit for any particular rover, so it belongs in the test plan rather than in a comment.

UMD Loop (University of Maryland), athena-code, 2026. *Original source*

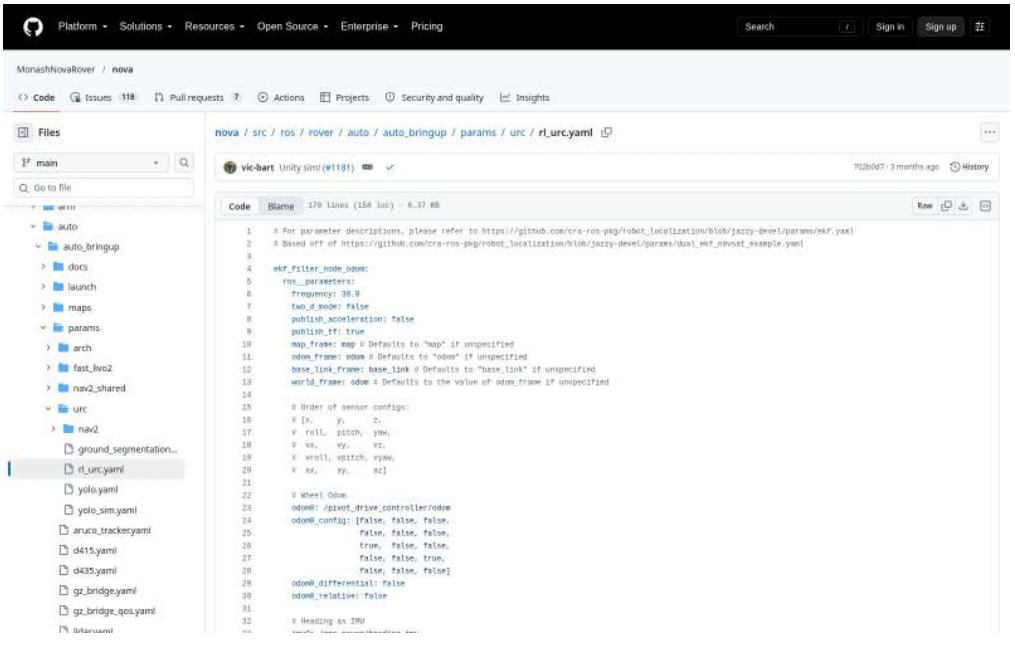
Full resolution

Four planning details separate a costmap that works from one that refuses the hill. The **unknown space** must be traversable, or a mostly unmapped start area makes the only route look lethal. **Inflation** must stay below about 0.4 of the narrowest gap the rover intends to drive, because an inflated costmap seals a gap long before the rover is too wide for it; UMD runs 1.5 with a cost-scaling factor of 3.0 and Monash uses 1.28 at twice the robot radius. A **progress checker** converts a six-minute oscillation into a bounded recovery, and UMD sets 0.5 m in 10 s. A **stay-out zone** drawn while the rover is stopped is legal `S-URC27-RULES §1.e.iii` and is the only mechanism that turns “do not go there” into a rule the planner obeys.

Localisation lives in one file of constants

Most localisation failures are not filter failures. They are a coordinate frame, a lever arm or a yaw source that nobody wrote down. Missouri S&T converts incoming WGS-84 coordinates to UTM and subtracts a measured antenna offset `S-MST-DOCS-NAVBOARD`; one file of constants and a ten-minute field calibration are the whole defence.

The recommended structure is a dual extended Kalman filter, one for the smooth odometry frame and one for the jump-tolerant map frame, fed by wheel odometry, IMU, GNSS and a magnetometer. Tuning is the work: days of log replay are normal, and a monitor on the innovation and covariance makes divergence visible rather than discovered at the target. Heading noise from a pair of antennas scales as the position noise divided by their separation, so 1 cm over 0.5 m is about 1.15° and 1 cm over 1 m is about 0.57°.



The per-source mask is the part that matters: each of the fifteen state entries is enabled or disabled per sensor, so wheel odometry contributes position and zero-rate yaw while GNSS heading contributes yaw alone. Missouri S&T won 2025 and 2026 without this stack, which is a reminder that the configuration is not the achievement.

Monash Nova Rover, 2026. *Original source* *Full resolution*

The astronaut state machine

The astronaut sub-mission has a fixed command order `S-URC27-QA` AUQ13 and an unfixed number of retries, because the rover may self-abort for free and re-accept the command `S-URC27-RULES` §1.e.xi. That combination is exactly what a behaviour tree or a disciplined state machine expresses and a chain of conditionals does not. Five properties are worth copying outright.

1. **Every free recovery is a state.** A self-stop and a safe hold exist so that a fault costs time rather than points. A team that turns a recoverable fault into an operator abort donates 20 % of that task.
2. **Link loss is two policies, not one.** During the astronaut sub-mission the rover holds, because a human is nearby. During route finding it keeps driving and queues telemetry, because the second target is out of radio line of sight by design. A stop-on-link-loss handler is a 25-point defect `S-URC27-RULES` §1.e.xvi.
3. **The armed state is a listening state.** The rover waits visibly parked with the LED red while it scans, so the operators must already be hands-off before the astronaut commands.
4. **Fetch is a camera and arm sweep from where Follow ended.** The pick lies in the immediate area `S-URC27-QA` AUQ8, so driving with an arm extended is an avoidable tip-over and self-collision risk.
5. **Come! closes the last 20 m on the person detector.** The rules guarantee the astronaut is in the field of view at that moment `S-URC27-QA` AUQ7, which makes the geometry a free safety margin.

Interfaces that bite

- **Compute and sensors.** Camera streams at 720p30 or better with valid camera intrinsics, RTK fixed position with dual-antenna heading, and an IMU at 100 Hz or more. Uncalibrated intrinsics or unsynchronised odometry halve the value of everything above them.
- **Communications.** A round-trip budget of 80 ms and 4 Mbit/s for one 1080p stream, and route finding that survives a total outage. The camera count is frozen with the communications design before the operator interface is written, because video bandwidth is the binding constraint.
- **Power.** A brownout during an arm stall looks exactly like a software crash. The E-stop also removes power from the log storage, so mode transitions must be flushed to non-volatile storage and logged at the station as well `S-URC27-RULES` §2.a.v.
- **Arm.** A 100–200 Hz joint loop and a Cartesian mode at or below 0.1 m/s for the grasp, with inverse kinematics and joint limits shared between the arm and the autonomy code.
- **Operator station.** Hands-off is a procedure as much as a setting: a physical autonomy key kept out of reach, one E-stop owner, and only the drone pilot holding a live controller `S-URC27-QA` AUQ25.

Failure modes

FAILURE	MECHANISM	PREVENTION
Rover stops dead 40 m behind the hill with 25 points pending	One link-ok flag gates the drive, and the watchdog calls the E-stop path	Two fault policies; an autonomous return is a planned behaviour S-URC27-RULES §1.e.xvi
Yaw drifts 20° over two minutes of driving	Wheel slip plus magnetometer disturbance and no heading observation	Dual-antenna heading as the primary yaw source; re-initialise when the fix quality drops
Rover drives 1.4 m past the target and scores zero	The tolerance is measured at rest, so the stop is not treated as a control problem	Terminal approach at or below 0.3 m/s and a 3 s hold for the LED
Planner returns no path because inflation sealed the gap	Inflation radius at or above half the gap width	Keep inflation below 0.4 of the intended gap and check it on the elevation model
Zero autonomy credit because the judge did not see red	The LED washes out in sun with no shade	A mechanical shade, a daylight test at 20 m, and a second LED copy visible from the side S-URC27-QA AUQ3
Penalty because an operator touched a gamepad during a near miss	Reflex during autonomous driving	The autonomy key out of reach, a rehearsed freeze, and a visible autonomous banner S-URC27-QA AUQ21

Shortlist for a first build

ITEM	CHOICE	WHY
OS and middleware	Ubuntu 24.04 with ROS 2 Jazzy	Supported to 2029, and it pairs with the current Gazebo release S-ROS-REP2000
Localisation	Dual-antenna RTK, dual EKF, wheel odometry, magnetometer	The cheapest way to buy the dominant error term
Mapping	A 1 m USGS 3DEP elevation model into a slope costmap, plus a live obstacle layer	Removes a dependency on an unreleased model and answers the hill directly S-USGS-3DEP-EPQS
Global planning	SmacPlanner2D with a Dubin model	Kinematically feasible for a skid-steer rover S-NAV2-REPO
Local planning	Regulated Pure Pursuit at or below 1.0 m/s	Stable on rough ground, few parameters, terminal accuracy from the goal checker
Behaviour	Behaviour tree with custom nodes for LED and task state	Re-entrant recovery comes free, which the retry rule rewards
Perception	YOLO with depth, MediaPipe Pose for gesture	Two open models beat one custom pipeline in a first season
Simulation	Whatever runs the flight code	A passing simulation test only means something when it exercises the same code path
Follow!	Audio with the device as the floor, gesture only after a measured test	10 points for three and a half weeks is the best return on the ladder

Open problems

- **What slope can the rover actually climb?** URC publishes no inclinations S-URC27-QA AUQ17, and 15° is a choice with no cited basis. An instrumented wheel-slip test at 10, 15, 20 and 25° sets the costmap cap from measurement.
- **Does dual-antenna heading hold under the rover's own multipath?** A 30-minute static test and a 200 m run against a surveyed bearing give the heading error; the target is 1° at one standard deviation with no drift after a fix loss.
- **Is the 1 m stop repeatable?** Thirty runs around a surveyed point give the hit rate. The pass criterion is 90 % inside 1.0 m and 100 % inside 1.5 m, because 25 points turn on it.

- **Does the audio channel survive motor noise and wind?** Thirty utterances at three distances with the drive motors at mission speed settle a false-accept rate of one per ten minutes or better.
- **How long does an autonomous return actually take?** The 20 % penalty is cheap only if a return from 300 m plus a re-attempt fits inside the 40 minutes with five minutes to spare.
- **Does the hands-off procedure survive a near-collision?** Ten injected near misses with an instructed operator measure whether reflex beats procedure. Zero control touches is the only acceptable result.

Further reading

- UMD Loop publishes the reference Nav2 configuration, the elevation-model costmap and a route-finding result worth studying: the best software in the corpus and 32nd overall in 2026.
- Missouri S&T and Monash carry the two winning patterns: a bespoke protocol and a full ROS 2 stack, both with a simulator built from the flight code.
- Michigan MRover shows the custom-navigation route inside ROS 2, without a navigation framework.
- Compute, sensors and perception owns the camera, GNSS and compute choices this page consumes.
- Interfaces and couplings works the video-budget and mast-height chains with numbers.
- Autonomy: astronaut assistance and Autonomy: route finding map these requirements onto the two sub-missions.

SUBSYSTEMS · CHAPTER 4.9

Operator station and operations

The operator station turns a rover into points, through six people in a metal box with no view of the course. It is not a desk: it is a deployable system with a stopwatch on it, and the stopwatch is scored twice, once in the System Acceptance Review video and once in the mission. Everything here is a 2027 constraint, not an operating preference.

15 / 10 min
SETUP AND TEARDOWN
All communications inside the setup, radios off at the end
2027
S-URC27-RULES

6
STUDENTS INSIDE THE STATION
No advisors, and no view of the course
2027
S-URC-LOGISTICS

40 min
AUTONOMY MISSION CLOCK
Two sub-missions, one target behind the hill
2027
S-URC27-RULES

20 %
COST OF ONE INTERVENTION
The clock keeps running and penalties add
2027
S-URC27-RULES

Key takeaways

- The **15 minutes** covers all necessary systems including all communications S-URC27-RULES §1.a.ii. Mast, 25 m cable runs, GNSS base, switch, six positions and power all have to be live inside it, in full daylight, with a judge timing.
- Two hats cannot be worn at once: **the commander does not drive and does not touch the arm**, and **the science operator is not the arm operator**. Both rules exist to protect the clock and the evidence, not to spread workload.
- Every mission hands the crew a different clock, from **20–30 minutes** for Astrobiology to **40 minutes** for the two Autonomy sub-missions and as little as **10 minutes** between Autonomy and Equipment.
- The abort ladder must be written before the clock starts. In the moment, crews default to “keep trying”, which is how a 5-point Follow! consumes a 50-point route.

What the subsystem must do

The station is a metal trailer unit or an MDRS structure, with the course deliberately invisible from inside, 120 V 60 Hz power, and tables and chairs provided S-URC27-RULES §3.d.i. Six students may be inside and advisors may not; advisors spectate from the field S-URC27-RULES S-URC-LOGISTICS §3.b.ii. Twenty-four students and two advisors register, and **14 people may be on site per competition day**, which makes the crew a logistics plan as much as a roster S-URC-LOGISTICS .

Three clocks are fixed by rule and one is fixed by geometry. Setup runs from C2 access to ready to compete in ≤**15 minutes**, teardown and removal of all equipment in ≤**10 minutes**, and radio equipment goes off immediately after S-URC27-RULES §1.a.ii. Items deployed in the field, such as repeaters or cameras, may stay until the mission ends but must be recovered inside the same ten minutes S-URC27-RULES §1.a.iii. The antenna may be **3 m** tall at most, its base **5 m** from the station, guys anchored **10 m** out, with **25 m** of cable or two runs if the drone carries its own radio S-URC27-RULES S-URC27-RULES S-URC27-QA §2.c.iii §2.c.iv DQ2. Operators must be inside with no view, not standing by the door S-URC27-QA DQ2.

The station is scored twice. The System Acceptance Review awards **8 of 100** points for communications, command and control and operations, judged on the station being used during demonstrations and on the state of operational testing, and the written report awards a further **1 %** for approach to team training S-URC26-SAR . A crew that treats the station as furniture loses both quietly.

The numbers that size the design

The rule rows are sourced. The latency and staffing rows are design targets, because no team publishes a measurement.

QUANTITY	VALUE	BASIS
Setup and teardown	**≤15 min** and **≤10 min**, then radios off	Includes antenna, cable, GNSS base, network and six stations S-URC27-RULES §1.a.ii
People	**≤6** inside; 24 students and 2 advisors registered; **14 per day**	S-URC-LOGISTICS
Antenna and cable	mast **≤3 m**, base **≤5 m** from the station, guys **≤10 m**, **≥25 m** of cable	Two runs if the drone has its own radio S-URC27-QA DQ2
Mission clocks	Astrobiology **20–30** · Delivery **30–60** staged · Equipment **30** · Autonomy **40** · Autonomy to Equipment **10**	Two clocks inside one Autonomy mission
Operator-facing latency targets	glass-to-glass **250–300 ms**, **15 fps** drive stream, command round trip **100 ms**, jitter **50 ms**	No team publishes a measurement, so these are targets
Sub-band change	**≤2 min** by one person, procedure on the card	The schedule assigns the sub-band per mission S-URC27-RULES §2.c.v
Scoring exposure	**8 of 100** review points for station and operations; **1 %** for training	The station is scored twice S-URC26-SAR
Power allowance	assume **≤10 A continuous** on one circuit	The rules state 120 V 60 Hz and say nothing about current

Design space, in one table

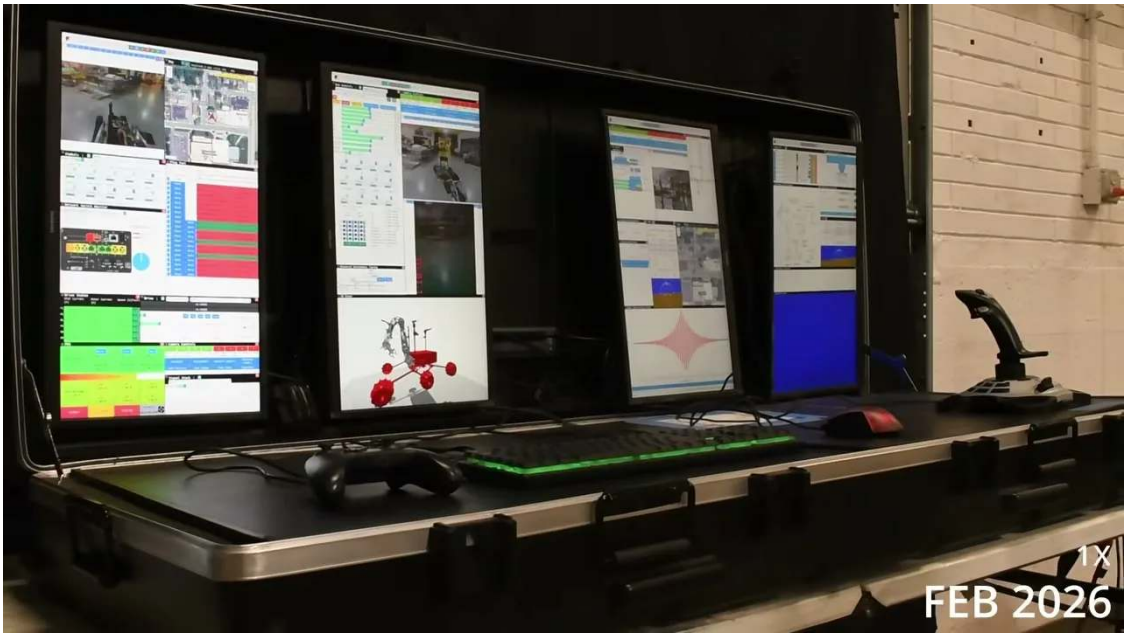
DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Physical estate	URC trailer only · console in a moulded case · a laptop per operator · one large monitor · a video wall	Console in a moulded case: one case, four screens, a known-good loom, seconds to deploy
Interface architecture	monolithic dashboard · per-role pages · multi-client web app · mixed third-party viewers	Per-role pages in one app, camera-dominant, with a map inset and a state banner, served to several machines
Input devices	gamepad · joystick · keyboard and mouse · a labelled button box	Gamepad for driving and a joystick for the arm, plus labelled buttons for autonomy, radio off and the hands-off display
Station compute and network	two business switches trunked over the radio links · browser clients on a local network · recording at both ends · a battery buffer	A managed switch with a saved known-good configuration, a battery buffer for the switch and one laptop, and recording at both ends
Staffing	six fixed seats plus field crew · six plus a station-clean science pair · four seats · five	Six fixed seats with named alternates, plus two station-clean people for the science debrief
Procedure	laminated card per mission · digital checklist · memory	A laminated card carrying setup order, roles, clock plan, abort thresholds and the named owner of radio-off

The 15 minutes is a system-integration test

The clock covers all necessary systems, including all communications S-URC27-RULES §1.a.ii. That makes it a test of the whole rover programme, not of packing speed, and the review judges score exactly this in the video S-URC26-SAR . A worked sequence that fits: cases and tables by 2:00, both cable runs with mast and guys by 6:00, the GNSS base logging by 8:00, every page loading and the mission's sub-band set by 10:00, weigh-in and cube check by 12:00, the voice check with spotter, runners and astronaut by 14:00, and the commander's opening call at 15:00.

Three margins make that repeatable rather than lucky. A **3-minute setup and 2-minute teardown margin** turns 15 and 10 into targets of 12 and 8. A **"last two minutes" pair** whose only job is cable and mast removes the end-of-setup scramble. **Radio-off as the last checklist line with a named owner** prevents the small rule violation that happens after everyone relaxes S-URC27-RULES §2.c.i.

The environment is part of the test. Events run in full daylight, up to 37.8 °C, in a metal box S-URC27-RULES S-URC27-RULES §3.d.i §3.d.iii. Glossy laptop panels are the default failure at these ambient levels; no URC statement covers screen brightness, so matte panels, hoods, water and a rotation plan are design items rather than comfort items.



Missouri S&T's 2026 station: four screens, a gamepad, a joystick and a keyboard inside the same case that carries them. The case is the interface definition, because the loom inside it is the same loom every mission, which is what makes a 12-minute setup realistic.

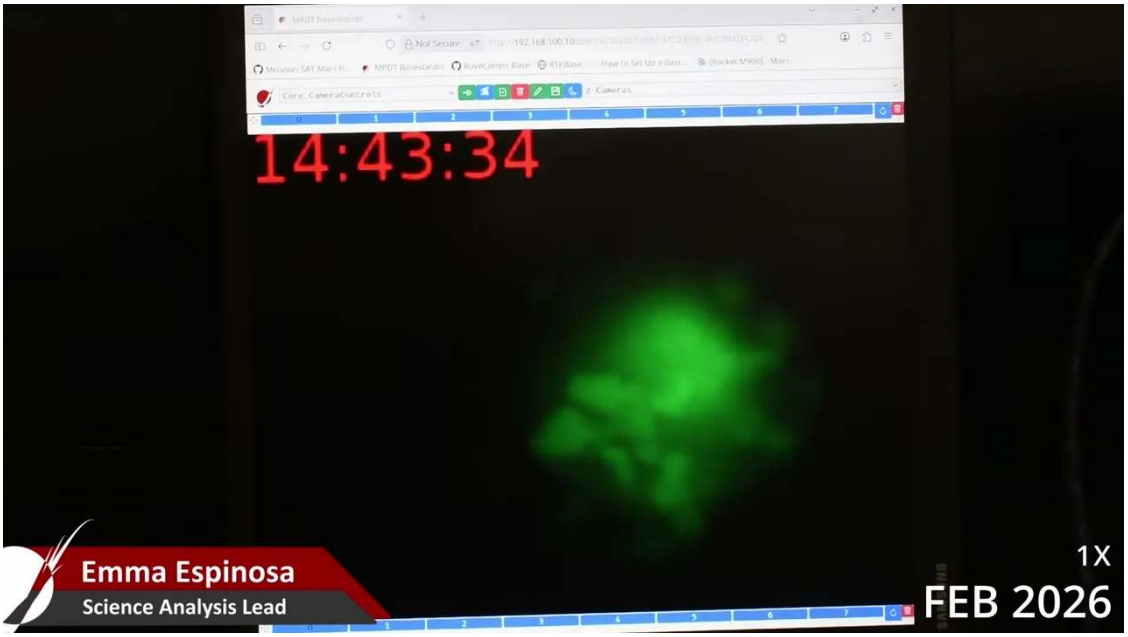
Missouri S&T Mars Rover Design Team (MRDT), 2026. *Full resolution*

The clock is a role, not a gadget

Every mission hands the crew a different clock, and the operator's biggest lever is how it is divided. A planning model for a 30-minute Astrobiology run shows where the decisions sit; the numbers are a construction from the rule limits and the task structure, not a measured trace of any team.

STAGE	CLOCK	OWNER	DECISION AT THE END
Opening call, final checks, start	0:00–1:00	Commander	The clock start is confirmed to the crew aloud
Drive to site 1	1:00–6:00	Driver, with science calling candidates	Site 1 is committed, or the drive continues
Site 1 survey	6:00–10:00	Science and arm	The site is accepted for a sample or rejected
Sample to 10 cm depth	10:00–16:00	Arm, with the driver repositioning	Has 10 cm been reached, and how many more attempts are justified?
First instrument run	16:00–21:00	Science	Is the signal interpretable?
Second site or second capability	21:00–25:00	Commander decides	Site 2 is abandoned at 25:00 regardless
Cache seal and hand-off	25:00–29:00	Arm and commander	The cache must be removable within 5 min of the roving end S-URC27-RULES §1.b.vii
Stop and debrief	29:00–30:00	Commander and the station-clean pair	The debrief starts even if the rover failed S-URC27-RULES §1.b.x

Three rules make a timekeeper real. **One clock**, read aloud at stage boundaries and embedded in the page the operator is watching. **Pre-declared decision points**, because without them abandoning a task feels like failure and crews do not do it. And a **hard last useful minute**, so the second Autonomy target is never started so late that the rover is still driving at 40:00.



The clock is not a wristwatch beside the screen; it is in the screen, over the instrument feed, in the largest type on the page. The science operator reads it without looking away from the data, and the commander reads the same number on a different display.

Missouri S&T Mars Rover Design Team (MRDT), 2026. *Full resolution*

The abort ladder is written down before the mission

Every 2027 mission has a legal way to give up ground and keep the rest, and the ladder only works if it is decided while the crew is calm.

SITUATION	CORRECT 2027 ACTION	COST
Autonomy rover not progressing	Operators send stop and return to the astronaut	**20 %** of that task
Rover must be driven out of a bad position	Teleoperate back to a previously visited point	**50 %**, capped per task
Rover stuck or broken	Request an intervention, runners go out	**20 % of the mission**, clock running
Spectator touches the rover unasked	Treated as an emergency stop	Mission over, points to that moment kept
Typing attempt going wrong	Abort and restart; judges reset the display	No penalty, clock only

The ladder is written because the alternative is a habit. The intervention ladder in particular punishes the second bad decision, not the first: two interventions in one mission leave 60 % of the points earned S-URC27-RULES §3.e.viii. A team member who leaves the station as a runner cannot operate again that mission S-URC27-RULES §3.e.iii, so the runners are named before the mission and the forfeit is stated out loud.

Six seats, and the two rules that make them work

URC mandates a maximum of six students and a set of operator conduct rules, not a roster. The seat plan below is a design recommendation that fits both.

SEAT	OWNS	MUST SEE ON THEIR OWN SCREEN	FAILURE IF THE SEAT IS EMPTY
Commander	The clock, go or no-go, judge liaison, the abort decision	Mission clock, rover state, all task states, link status, battery, a copy of the driver's camera	The mission drifts past the last useful minute and nobody dares abandon a failing task
Driver	Teleoperation: the only hands on drive controls	Three or four cameras, heading and attitude, speed, wheel current, position, map	The rover is slow, gets stuck, or drives over the sample site
Arm and payload	Arm joints, gripper, end-effector changes, task order at the lander	Arm and wrist cameras, joint angles, a kinematic twin, gripper state, force or current	Equipment Servicing collapses, and the arm takes its own camera with it `S-MONASH-IEEE-RAM`
Science	Site triage, instrument runs, sample provenance, the report	Instrument data, sample depth and identity log, panoramas with cardinal directions, position with accuracy	Sites are chosen on impression, and the 10 cm provenance is lost at the debrief
Autonomy	Targets, route, state machine, the behind-the-hill task	Map with planned path and stay-out zones, the autonomous banner, target list, progress against the clock	The rover sits still while the crew debugs, or someone touches a control and voids the run
Drone pilot	Drone flight, video, relay deployment	Drone camera, battery, altitude and attitude, spotter voice channel	The one asset that sees beyond the hill is never used, or is used unsafely

Two structural rules carry the roster. **The commander does not drive and does not touch the arm:** the moment the person with the clock is inside the task, the clock stops being anyone's job. **Science is not the arm operator:** sampling is a manipulator task and a provenance task, and combining them forces a choice between moving soil and recording why.

Two compliance surfaces follow. Hands-off is enforced by observation rather than by code `S-URC27-QA` AUQ21, so the procedure is visible: announce entering autonomous mode, confirm the rear LED is red, put both hands on the desk edge with palms visible, and hang the gamepad on a hook. Drone control is a separate chair: the pilot sits inside with no view, the spotter holds visual line of sight and may take control, and only the pilot touches controls while the rover drives autonomously `S-URC27-RULES` `S-URC27-QA` §2.b.xii AUQ25.



The drone pilot's own view: a head-up overlay of battery, current, altitude and speed over the camera, with a map inset below. It is a second control station inside the box, which is exactly why the pilot must be obviously flying only the drone while the rover moves autonomously.

Monash Nova Rover, 2026. Full resolution

Interfaces that bite

- **Communications.** The 15 minutes is shared: if the mast takes six, the stations get nine. The station needs signal strength per link, a sub-band change one person can run in two minutes, and antenna-aim feedback that does not need a camera.
- **Software.** Autonomy state must be published as data rather than inferred from behaviour, or the hands-off check becomes an argument S-URC27-QA AUQ21. Telemetry needs timestamps and a replayable log.
- **Arm.** One USB hub feeding every end-effector camera is a single point of failure: Monash nearly lost all of its arm cameras to one snagged cable in 2023 S-MONASH-IEEE-RAM.
- **Astronaut.** Quiet is required only for a gesture Follow!, and talking is encouraged at every other moment S-URC27-QA AUQ15. The station owns that switch deliberately.
- **Systems engineering.** The schedule must show integrated operations tests and operator training, because the review scores both. The station kit is not rover mass, but it is luggage against the 14-per-day limit.

Failure modes

- **The link that worked in the motel.** The judges name this as the single most common way a team fails: the rover works in the parking lot and not at the site, because of a loose wire or connector S-URC-ADVICE. The prevention is a labelled loom, one spare of every connector, and a link test as a checklist item on arrival.
- **Setup overrun.** The mast and guys are the slowest items and everything wants the same three minutes at the end. Pre-attached guys, a cable reel with the connector already fitted, and a radio-off master switch reachable without crawling under a table fix most of it.
- **The rule violation that looks like bad luck.** An operator leaves the station to help and comes back to keep operating S-URC27-RULES §3.e.iii; a spectator touches the rover S-URC27-RULES §3.e.ii; a keyboard is touched during autonomous typing S-URC27-QA EQ2. Each is a procedure failure with a written prevention.
- **The task that eats the clock.** A staged Delivery mission makes Stage 1 a gate: Missouri S&T scored 15 in 2023 because Stage 1 was not cleared S-URC23-SCORES. Stage 1 is not a warm-up, it is the mission.
- **The search that never converges.** A search area plus a low-contrast target plus no declared abandonment threshold is how a 30-minute mission ends with 15 points S-MONASH-IEEE-RAM.
- **Heat and glare.** Operators who cannot read their screens make different and worse decisions. Matte or hooded panels, water, and a rotation plan are part of the station, not comforts.
- **The intervention spiral.** A crew that intervenes once early tends to intervene again. The commander's job is to decide whether the mission is still worth points before the runner leaves the station.

Shortlist for a first build

ITEM	QUANTITY	WHY
Moulded hard cases used as console tables	1–2	Doubles as furniture and protection, and carries the loom
Matte 15–17 inch USB-C monitors	3–4	One per information role; matte beats glossy in sun
Laptops, one per operator or one plus thin clients	4–6	A dead machine must be swappable in minutes
Managed gigabit switch with labelled leads	1 plus 8	A known-good configuration saved to a stick
Gamepad and joystick	2 plus 1–2	Driving needs two analogue axes; the arm needs fine proportional control
Antenna mast of 3 m or less with pre-attached guys	1	The slowest single setup item
25 m low-loss cable runs, connectors fitted, on reels	2	Rover and drone; the rules require at least 25 m
GNSS base with tripod and correction link	1	An own base station is required, since no correction service is available S-URC27-QA RQ11
Power strips, 10 A fused	2	120 V 60 Hz only, and one circuit until measured
Battery buffer for the switch and one laptop	1	No battery buffer is provided by the organisers
Laminated mission cards and grease pencil	5	Survives dust, heat and a power loss

Open problems

- **What is the station's power capacity?** No current limit, circuit count or battery buffer is published. Asking the competition organisers and measuring at the site meeting settles it; until then, design to 10 A continuous.
- **What is the table geometry inside the trailer?** The six-seat layout assumes a 20 ft box. Photographing and measuring at the first site visit settles it.
- **What is the real glass-to-glass latency?** No team publishes a measurement, so 250–300 ms is a target. A measurement at 25 m and at 1 km turns it into a specification.
- **How many camera streams can one operator page carry?** A load test on the actual laptops and the actual radio settles it, and the figure belongs in the interface specification.
- **May the debrief pair watch field video during the mission?** The rule excludes members who have seen the site with their own eyes, and video is not direct sight. Asking the science judges in writing before the review settles it.
- **What are the 2027 Delivery stage boundaries and time splits?** They set the Stage 1 gate and the carry-forward decision, and they are promised before the competition.

Further reading

- Missouri S&T shows the console-in-a-case pattern and a station built by a team that wins with it.
- Monash shows the per-role web interface, six station crew and seven field crew per mission `S-MONASH-NEWS2025`.
- WVU Team Mountaineers publishes the clearest fault-injection practice in the field, including a documented unplugging test `S-WVU-SAR2026`.
- Software stack and autonomy owns the state machine and the hands-off behaviour this page depends on.
- Communications and networking owns the link, the antenna and the sub-band plan.
- Review, judging and awards sets out how the station is scored in the System Acceptance Review.

SUBSYSTEMS · CHAPTER 4.10

Systems engineering and integration

A rover is fielded as one system, and four of its constraints are paid for by the same score: mass, volume, time at the site, and the evidence a report and a video present. Each one crosses subsystem boundaries, so none of them can be owned locally. This page sizes those couplings from published team practice, reduces them to the artefacts a first build can sustain through a seven-page review, and names the interfaces where one subsystem quietly breaks another.

50 kg

MASS PER CONFIGURATION

5 % of an event's points per kilogram above it

2027

S-URC27-RULES

1.2 m

CUBE AT EVERY WEIGH-IN

No disassembly, and a 40 % penalty if it misses

2027

S-URC27-RULES

100 of 500

REVIEW POINTS

2026 split: 40 report, 48 video, 12 judges

2026

S-URC26-SAR

1.14 m

STOWED STACK, 6 WHEELS AND A 2 M MAST

0.06 m of margin inside the 1.2 m cube

2026

S-WVU-SAR2026

Key takeaways

- Four constraints are paid from one score. Mass costs **5 % of an event's points per kilogram** over 50 kg, volume costs **40 %** of a mission if the cube is missed, and the review is **100 of 500 points** for showing that the first three were managed.
- The 1.2 m box is a **measurement**, not a packing problem. The vertical stack of a 6-wheel rover with a 2 m mast lands at **≈1.14 m**, which leaves 0.06 m, and three axles on a 1.05 m wheelbase are **1.55 m long**, 0.35 m over the limit before any arm.
- Mass follows architecture. WVU took **20 kg** out of its rover in one year with a fabrication decision, and the lightest championship rover in the record is **33 kg**, so 50 kg is a ceiling at 1.5 times a winning design.
- The review is a compressed exam, not a soft landing. The admitted set is decided inside a few points in a normal year, and documentation explains about **18 %** of rank variance, so it is a gate and a bank of points, not a podium predictor.

What the subsystem must do

Systems engineering at URC exists to hold the couplings that no subsystem owns. **Mass:** 5 % of an event's points per kilogram over **50 kg per configuration**, with a **70 kg** ceiling on total fielded parts S-URC27-RULES §2.a.iii. **Volume:** a 1.2 m cube measured at every weigh-in, with no disassembly permitted and a **40 % penalty** if the rover does not fit S-URC27-RULES §2.a.ii. **Time:** 15 minutes of setup, a 10-minute teardown, and as little as 10 minutes between two missions S-URC27-RULES S-URC27-RULES §1.a.ii §1.a.iv. **Evidence:** 100 of 500 points for a page-limited report and a five-minute video showing that the first three were under control.

The syllabus is published almost verbatim, as a list of what the milestone reviews look for: project management, systems engineering, systems development life cycle, work breakdown structure, schedules and Gantt charts, organisational chart, and test plan S-URC-RECTOPICS . None of that is optional decoration; it is the scored description of the process, and the team that already owns the artefacts scores them without writing anything in February.

The numbers that size the design

The rule rows are sourced. The alignment, allocation and correlation rows are computations from published team figures and official score sheets.

QUANTITY	VALUE	BASIS
Mass limits	**50 kg** per configuration, **70 kg** total fielded parts, **5 % per kilogram per event**	Judges round down, so 50.9 kg records as 50 kg S-URC27-RULES §2.a.iii
Stow envelope	**1.2 m** cube, any orientation, no disassembly, **40 % penalty**	Measured at every mission inside the 15-minute setup S-URC27-RULES §2.a.ii
Vertical stack, 6 wheels and a 2 m mast	**≈1.14 m**, leaving a 0.06 m margin	Wheels and axle 0.51 m, chassis 0.16, arm 0.12, mast 0.30, antennas 0.05, by inference from published geometry
Plan-view length	**1.55 m** on a 1.05 m wheelbase with 0.25 m overhangs	0.35 m over the limit; 1.2 times the square root of two is 1.70 m of diagonal envelope
Arm as a line item	**9.4–10.4 kg**, or **19–24 %** of a 50 kg rover	WVU 5-DoF arm S-WVU-SAR2025
Worked allocation	37.5 kg of estimates plus 15 % growth is **43.1 kg**, so the design target is **≤48.5 kg**	A template rather than a measurement, with 1.5 kg kept for scale error
Lightest championship rover	**33 kg**, WVU 2023, **425.35 points** and first place	The rover carried a 5-DoF arm with 8 kg at full extension S-WVU-SAR2023
Review weight	**100 of 500 points**; 2026 split **40 report, 48 video, 12 judges**, seven pages and five minutes	The 2027 split is unpublished, so 2026 is the model S-URC26-SAR
Documentation against rank	Pearson **r = −0.419** over 162 team-years, $r^2 \approx 0.18$	Documentation explains about 18 % of rank variance

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Process	a full life-cycle model · internal gates with dated exit artefacts · a one-year winter cycle · informal	An internal gate chain with dated exit artefacts, labelled in the same language the review panel uses
Requirements management	spreadsheet matrix · machine-readable manifest · the robot model as truth · formal systems modelling · narrative	One spreadsheet plus one printable traceability figure, because nothing larger fits in seven pages
Interface control	a document per discipline · a register map · an envelope drawing · a bench electrical replica · none	Interface control document at the preliminary review, frozen at the critical review, plus a bus map
Mass method	CAD roll-up · catalogue plus scale · design to a cap · weigh at the end	Design to a cap at concept, CAD at the preliminary review, and a scale at integration
Modularity	front-half chassis swap · payload insert · one integrated rover · end-effector swap only	One integrated rover plus a two-bolt end-effector swap; treat the 70 kg allowance as reserve, not architecture
Stow strategy	geometry first · fold the wheels · hinge the chassis · telescope the mast · a diagonal fit	Geometry first, with the mast as the only thing that folds, decided while the wheel diameter is still free
Test programme	a ladder of tests · finite-element analysis · simulation first · a bench replica · a mock lander · a road shake	A ladder plus a bench replica plus a mock lander, with operator hours logged
Review production	write to the published rubric · storyboard the video · write the report last · pre-screen the systems award separately	A package written to the rubric page by page, a video storyboarded to its weights, and the systems award targeted on its own

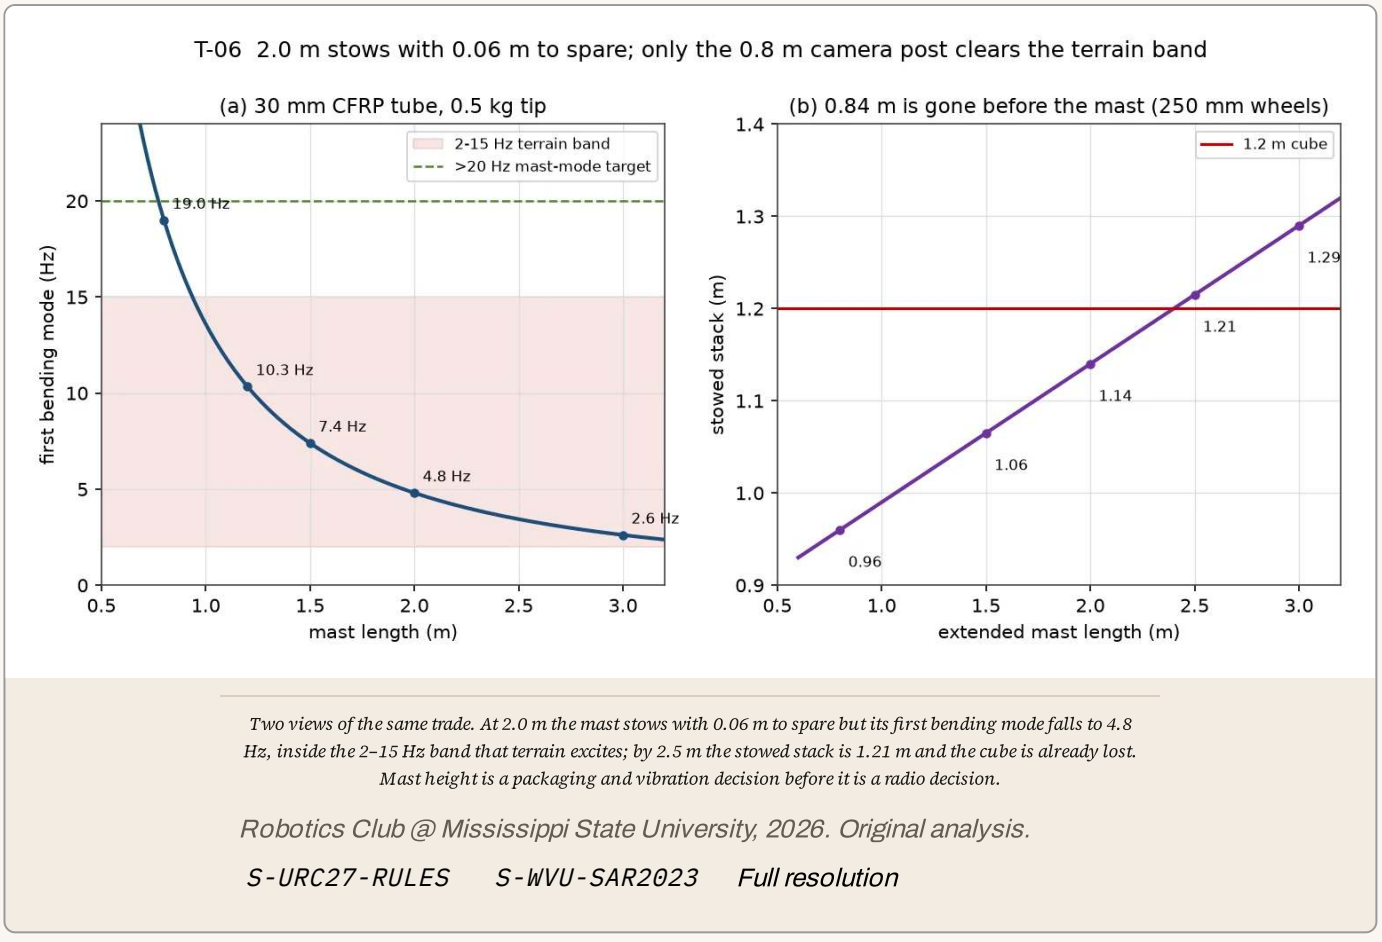
The 1.2 m box is a measurement, not a packing problem

The rule is unusually strict. The rover may be placed in any orientation and may be folded or articulated, but it may not be disassembled to fit, and that includes wheels, antenna and any other protruding system S-URC27-RULES §2.a.ii. A mast that

unbolts and rides in the crate is therefore illegal at the weigh-in, which is the single most common packaging misconception in the field.

BAND IN THE VERTICAL STACK	HEIGHT
Wheels and axle height, at a 0.51 m wheel or wheg	0.51 m
Chassis deck and electronics box above the wheel top	≈0.16 m
Arm parked along the deck	≈0.12 m
Mast, collapsed	≈0.30 m
Antennas and gimbal above the mast	≈0.05 m
Total	**≈1.14 m**

The 0.06 m that is left is the whole problem. **Wheels and axle height spend 43 % of the cube and the mast spends 25 %**, so a design that wants a 3 m mast has to add segments, tilt the mast, or give up height somewhere else. Length is worse than height: three axles on a 1.05 m wheelbase with a 0.25 m overhang at each end is **1.55 m**, 0.35 m too long.



Two views of the same trade. At 2.0 m the mast stows with 0.06 m to spare but its first bending mode falls to 4.8 Hz, inside the 2–15 Hz band that terrain excites; by 2.5 m the stowed stack is 1.21 m and the cube is already lost. Mast height is a packaging and vibration decision before it is a radio decision.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *S-WVU-SAR2023* *Full resolution*

Three escapes exist. **Shorten the wheelbase**, which is what a 20 inch wheg does. **Hinge the chassis** at a central bearing, which costs a bearing, a harness across a moving joint and a lock. Or **prove a diagonal fit**: the rule permits any orientation, and 1.2 times the square root of two is **1.70 m** of envelope in plan and 2.08 m in space, which is legal and almost never documented.

The sharpest trend in the record is that the two best-resourced teams independently removed folding from their stow strategy in 2026. WVU went from 24 inch folding whegs to 20 inch rigid ones, citing fewer points of failure `S-WVU-SAR2026`, and Missouri S&T moved to a compacted suspension that no longer needs to fold `S-MST-SITE-ATHENA`. A fold is a hinge, a lock, a tolerance stack and a failure mode that appears at the weigh-in, in front of judges, on a clock. Geometry designed to fit is cheaper than a mechanism designed to shrink, and the fold that remains should be the mast's: single, over-centre and needing no separate hardware.

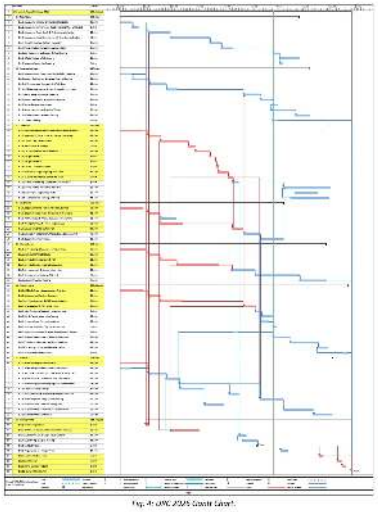
Mass follows architecture, and the schedule is the real budget

A **33 kg** rover won the 2023 championship with a 5-DoF arm carrying 8 kg at full extension `S-WVU-SAR2023`. Those 20 kg came off in one year from a fabrication decision, not a diet: semi-monocoque sheet metal and composite wheels replaced tube-and-printed construction. Mass is therefore recorded rather than negotiated.

No URC team publishes a mass budget or a centre-of-gravity table, so the discipline rests on a handful of figures: WVU's 33 kg, its predecessor at about 53 kg, Northeastern at about 46 kg, and a 24.4 kg rolling chassis that can be rebuilt from Monash's public robot description `S-WVU-SAR2023` `S-NEU-NEWS-2026` `S-MONASH-GH-NOVA`. The arm is the largest single line anyone publishes, at 9.4–10.4 kg for a 5-DoF unit `S-WVU-SAR2025`.

A workable allocation for a first rover, given as structure rather than as numbers, is 37.5 kg of component estimates plus a **15 %** growth allowance, which lands at **43.1 kg** and is held to a ≤**48.5 kg** gate with 1.5 kg kept for scale error. The weigh-in plan is the part teams forget: weigh the bare chassis, then each subsystem, then the integrated rover, at three points, and record the scale, the date and the configuration. The number that matters is the heaviest legal configuration, because the judges weigh at every mission and round down.

The schedule evidence is a ratio rather than a date list. WVU's 2026 plan released procurement in **December**, assembled in **January**, filmed in **February**, and spent **March to May** testing, which is roughly 3:1 test-to-build in the second half of the year `S-WVU-INCOSE2026`. A team that builds until May has inverted the ratio and will arrive with an untested rover and an unwritten review. Backfit to 2027, integration is complete by mid-February, the software freeze falls on 2027-04-10 and the mechanical freeze on 04-24.



The Gantt gets its own page, its work-package numbers are the same ones used in the budget, and every line has an owner. It is a review artefact rather than a planning aid: the panel reads it as evidence that the rest of the package is real.

Team Mountaineers, West Virginia University, 2026. Publicly released by the team; reproduced for analysis with attribution.

`S-WVU-SAR2026` `S-URC-BUDGETGANTT` *Full resolution*

The review is a compressed exam

The System Acceptance Review decides admission, and in a normal year it decides it inside a few points. The distribution is narrow enough that a mid-field package is a risk even with a good rover.

YEAR	FINALISTS	LOWEST ADMITTED	MEDIAN	HIGHEST	STANDARD DEVIATION
2022	26	80.60	88.42	94.71	3.71
2023	31	85.61	89.31	94.54	2.71
2024	34	84.04	89.87	**97.80**	3.57
2025	35	86.90	90.33	97.27	2.54
2026	36	**75.86**	86.77	96.75	**6.39**

The population is teams that recorded mission scores, so the lowest column is the lowest admitted score rather than a cut line; scores of teams that were not admitted are not published `S-URC22-SCORES` `S-URC26-SCORES`. Two readings matter for planning. In 2023 to 2025 the standard deviation was **2.5–3.7 points**, so a team five points below the field median was outside the admitted set on that evidence. And 2026 is not comparable with the years before it: the lowest admitted score fell by about nine points while the highest held near 97, which is consistent with a substantially updated rubric rather than with a weaker field.

A high review score is real and moderate. Across **162 team-years** the correlation with final rank is **r = −0.419**, where a negative value means a higher review score goes with a better finish, and $r^2 \approx 0.18$ means documentation explains about a fifth of the variance while mission execution carries the rest. The team with the highest review score won the event once in five years, in 2025. The review is a gate and a bank of 100 points.

The 2026 rubric is the model used here, at **40 report, 48 video and 12 judges**, seven pages and five minutes, because the 2027 split is unpublished. The finalist set is counted two ways, with 38 packages pre-screened and 36 finalists carrying mission scores, the difference probably being packages submitted without a mission run. Content past the page or time limit is not read, and reused material from a previous year is a disqualification risk, which makes the format a design constraint rather than a presentation preference.

The systems award is a separate instrument. The INCOSE Foundation award ran for the first time in 2026, independent of competition scoring: the 38 finalist reports were pre-screened on a 100-point, eight-category rubric, **13 teams scored 80 or more**, and each finalist then did a 30-minute interview worth 78 points plus 22 from on-site observation, combined 50/50 with a threshold of 75 `S-URC-INCOSE`. It is a second reason to hold the artefacts, and a separate narrative.

What top packages do, that a first build can copy

Five practices are visible across the four public WVU packages, which scored 94.35, 97.80, 93.17 and 96.75, and across the Missouri S&T and Michigan public record. These are documented practices, not proven causes; the sample is far too small to claim that.

- **Every architecture diagram is colour-coded new against reused**, with a legend that tags each changed component `S-WVU-SAR2026`.
- **The Gantt occupies its own page**, work-package numbered, owner-labelled and dated to the week.
- **The budget is a table to work-package level 2**, with cost categories that match the work breakdown.
- **Testing is a ladder with named facilities and measured numbers**, not adjectives: a rock garden, a replica lander, 750 m line of sight and 164 m without it, 30° slopes, a 10 kg end-effector lift.
- **Each mission gets the same four-part structure**: what the mission demands, what the design does about it, what has been tested, and what remains.



One page answers the two questions a review asks: what is new, and how does it integrate. The mechanical architecture diagram carries a two-entry legend, revised against new, so the change set is visible in one screen, with changed components separated from inherited ones, and the key-innovations paragraph names the decisions behind them.

Team Mountaineers, West Virginia University (Team Lead Logan Gold; advisor Dr. Yu Gu), 2026. Publicly released by the team on its open-source documentation page; reproduced for analysis with attribution.

`S-WVU-SAR2026` *Full resolution*

Interfaces that bite

- **Every subsystem.** Systems engineering owes an allocation and a requirement identifier, and needs measured mass, current, geometry and interface deltas in return. Without both, subsystems optimise locally and the rover arrives at 55 kg.
- **Arm.** A 1.5 m reach requirement fights a 0.12–0.15 m parked height, so the reach trade and the stow trade have to be made together, with the parked pose drawn at concept.
- **Mobility.** Wheel diameter is a four-way coupling of climbing, centre-of-gravity height, stow height and mass. WVU's change from 24 inch to 20 inch moved all four at once, so someone has to own the number.
- **Communications.** The link requirement comes from mission geometry, so the range test at 750–850 m with line of sight and 164–300 m without it has to precede the review video.
- **Power and operations.** One red button must stop all power draw across power, compute, communications and arm `S-URC27-RULES` §2.a.v, which is a system function rather than a feature of a power board. Systems engineering owns the setup-clock budget; operations owns whether the clock is achievable.

Failure modes

Six of the eleven documented integration failures in the record are process failures rather than hardware failures, which makes them the cheaper half to prevent.

FAILURE	WHAT HAPPENED	PREVENTION
Transport shock	Parts were found damaged on arrival in Utah in 2026 S-UMICH-AERO-2026	Shake the packaged rover before departure, then repeat the power-on and link checks
Connector or link failure	Rovers that work in the motel parking lot and fail at the site because of a loose wire S-URC-ADVICE	Strain relief, latched connectors, and a post-transport check as a procedure
Mechanism on the signal path	Missouri S&T lost Delivery in 2024 to a rover-side signal-stack mechanism S-MST-NEWS2024	Lock the mechanism and test the stack on a shaker
Single arm axis	Delivery was lost again in 2024 to an arm vertical axis S-MST-NEWS2024	One-fault tolerance on the vertical axis, and an axis check in the setup procedure
Staged-mission time-out	A 15-point Delivery score in 2023 and 2024 because Stage 1 was not cleared S-MST-NEWS2023	Rehearse the clock as a requirement; the stage gate turns a slow start into a 15
Missing budget or envelope	A 5 % per kilogram penalty, and 40 % for the cube S-URC27-RULES S-URC27-RULES §2.a.ii §2.a.iii	The allocation, the weigh-in plan and a physical 1.2 m mock box

Shortlist for a first build

ITEM	CHOICE	WHY
Requirements	One spreadsheet and one printable traceability figure	Nothing larger fits in seven pages
Interfaces	Interface control document at the preliminary review, frozen at the critical review, plus a bus map	The best artefact in the field is a machine-readable protocol manifest S-MST-GH-ROVECOMM
Mass	Design to a cap, CAD at the preliminary review, a scale at integration, **≥1.0 kg** in reserve	Judges round down, so the reserve is real headroom
Packaging	Geometry first, a 1.2 m mock box before metal is cut, and the mast as the only fold	Two top teams removed their folds in 2026 S-WVU-SAR2026 S-MST-SITE-ATHENA
Modularity	One integrated rover and a two-bolt end-effector swap	The 70 kg allowance is reserve, not a second rover
Schedule	Integration complete by mid-February, film in February, freezes in April	A rover integrated in March cannot produce a February video

Open problems

- **What is the real mass of a 2027-class rover?** No team publishes a mass budget or a centre of gravity, so four rover masses, two of them inferred, are the entire evidence base. A short survey of ten teams at the event would close it.
- **Is a tilted or diagonal stow accepted in practice?** The rule permits any orientation and no team documents using it. A measured tilted fit in a mock box, plus a written question to the organisers, settles it.
- **What rubric applies in 2027?** Both the review rubric and the systems-award rubric are unpublished. Re-reading them monthly from December 2026 is the only defence.
- **Is the 15-minute setup achievable with a fold and a weigh-in?** No team publishes a measured setup time. Timing ten setups produces the distribution, which is the design input.
- **What thermal qualification practice exists?** No team document describes an environmental test at 37.8 °C in full sun. Logging enclosure temperature against ambient on the hottest field day, and publishing it, would create the first.

Further reading

- WVU Team Mountaineers publishes four years of review packages, which is the richest evidence for this page.
- Missouri S&T shows the compacted-suspension route out of the packaging problem.
- Northeastern is the closest published mass to the limit, at about 46 kg in 2026.
- Mass and power builds the full 50 kg budget that this page only allocates.
- Interfaces and couplings works the arm-reach, mast-height and wheel-diameter chains with numbers.
- System Acceptance Review sets out what the milestones demand, and Review, judging and awards covers the published weights.

SUBSYSTEMS · CHAPTER 4.11

Environment, dust, thermal and repair

A rover that passes its bench test on campus can still fail at the Mars Desert Research Station, and the gap between the two is thermal, mineral and logistical. The 2027 rules state the environment in 27 words: 37.8 °C, dust, light rain, no heavy rain S-URC27-RULES §3.d.iii. This page sizes the thermal problem from a steady-state box model, treats the dust as the clay and salt mixture it actually is rather than as sand, and prices field repair against the clock the rules run. Where a number does not exist, the page says so.

37.8 °C

RULEBOOK AMBIENT

100 °F; measured June maxima at the site run 25.3–36.7 °C

2027

S-URC27-RULES

77 °C

SEALED BLACK BOX AT 38 °C

40 W inside, no shade, steady state at noon

Model

S-URC27-RULES

1,309 m

SITE ALTITUDE

85 % of sea-level air density, so 15 % less fan mass flow

2027

S-URC-LOGISTICS

20 %

COST OF ONE MID-MISSION REPAIR

Per intervention, with the mission clock running

2027

S-URC27-RULES

Key takeaways

- The rulebook number is not the design number. A sealed **black** box at 38 °C ambient with 40 W inside reaches 77 °C internal, and the same box painted **white** reaches 55 °C, with a sunshade 50 °C and with a filtered fan 42 °C.
- **Paint is the biggest single lever at 22 K**, a shade adds 5 K and forced air another 8 K. A fan without a speed alarm is not a thermal design, and a top-only shade stops working when the sun drops.
- The dust is **montmorillonite and sulfate salt at pH 8–9**, finer than the dust an ingress rating is tested with, swelling when damp and conductive when wet. Connectors and crimps are the leading mission stopper, and they are the cheapest thing to fix.
- Repair **between** missions is free and repair **during** one costs **20 %** with the clock running S-URC27-RULES §3.e.viii. The strategy is a fast module swap, a cloned boot medium and a field kit traceable to a failure list.

What the subsystem must do

The environmental requirement is short and leaves four things out. It does not say that the ground and an enclosure run far hotter than the air, that the finals window is not one temperature, that the site sits at **1,309 m** where air density is about 85 % of sea level, or that **testing at MDRS is forbidden** before, during and after the competition S-URC27-RULES §3.d.iv. It also does not mention that there is no mobile coverage, that the nearest parts are about six hours away, and that the last 3.2 miles are a dirt road with a cattle grate and a short steep incline S-URC-ADVICE S-URC-LOGISTICS.

Three consequences follow for a first build. The rover must be its own test site, because the real one is unavailable. Every mission-critical function must be repairable from what the team carried in, because there is no ordering of parts. And the thermal design must be sized for a parked rover in the sun, because that is the worst case, not a moving rover on a cool morning.

The numbers that size the design

Ambient rows are sourced from the rules and from a reanalysis of the finals window. Internal temperatures are predictions from a steady-state box model, not measurements.

QUANTITY	VALUE	BASIS
Ambient, daily maximum, 2–5 June	**25.3–36.7 °C** for 2020–2025; **32.6–35.9 °C** in 2026	Size the enclosure for 40 °C air, not 37.8 °C S-ERA5-MDRS-JUNE
Ambient, daily minimum, same window	**12.3–21.4 °C**	Cool nights; a pack at 12 °C has higher internal resistance
Hanksville normals	June high **32.8 °C**, July high **35.6 °C**	The finals are about 3 °C cooler than the site's worst month S-WEATHERSPARK-HANKSVILLE
Solar noon elevation, 21 June	**75.0°**	At noon a flat shade works; by 18:00 the west face is the problem
Solar resource	**8.5 kWh/m²/day** in June; peak about **1,000 W/m²**	A black lid is a 200–400 W heater
Wind, 10 m maximum, 2–5 June	**14–36 km/h**, gusts to **66.6 km/h**; drones must be designed for **48 km/h**	Dust load, sandblast, antenna and crate tie-down loads S-ERA5-MDRS-JUNE S-URC27-RULES §2.b.vi
Dust carried into the habitat	**more than 50 g** over 12 simulated EVAs	Terrain, not wind alone, loads dust onto hardware S-BOS-SCOTT-2004-DUST
Soil mineralogy	montmorillonite, nontronite and illite, sulfate-dominated, **pH 8–9**	The clay in bentonite: swells when wet, clogs filters, conducts when damp S-STOKER-2011-MDRS-SOILS
Highest material limit that bites	PLA softens at **55 °C**	Inside the box band and inside a closed vehicle S-PRUSAMENT-PLA
Sealed black box, 38 °C ambient, 40 W	**77 °C** internal, 39 K above ambient	Steady state, still air, noon
Site surface temperature	not published	An infrared thermometer on the lid is the measurement that closes it

Design space, in one table

DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Enclosure	fully sealed and gasketed · filtered fan-pressurised · open with conformal coating · hybrid · heat exchanger	A sealed compute box with splash-shielded drive electronics, putting the sealing effort where the fragile and hot parts are
Thermal	white skin · sunshade · filtered forced air · heat spreader · a paraffin heat store · thermoelectric cooler	White box, a shade with at least 50 mm of overhang, one filtered fan with a speed alarm, and a heat spreader to the skin
Connectors	sealed automotive connectors · circular industrial connectors · hobby connectors · sealed bearings · conformal coating	Latched sealed connectors, strain relief within 50 mm, dielectric grease, and a coating that stops at the contact
Sensors that see dust	resistive potentiometers · magnetic encoders · sealed switches	Contactless magnetic encoders and Hall end-stops, because a contaminated potentiometer track fails intermittently
Fasteners	threadlocker · nylon-insert nuts · plain nuts with witness marks	Threadlocker on metal joints, nylon-insert nuts on plastic, and paint-pen marks on every critical fastener
Transport	strapped in a vehicle · a plywood crate with a bolted pallet and foam · a flight case	A crate with 25–50 mm closed-cell foam and a logged road test, never a van for the finals trip

The thermal ladder

A 0.40 × 0.30 × 0.20 m box has **0.52 m²** of skin. Still-air conductance is **6.0 W/K**, and internal forced air at 1–2 m/s raises it to **16.4 W/K**. At 38 °C ambient, 1,000 W/m² of beam, a solar elevation of 75° and 40 W of electronics, the same box four ways:

CASE	SOLAR ABSORBED	RISE ABOVE AMBIENT	INTERNAL
Black, sealed, no shade	195 W	39 K	**77 °C**
White, sealed, no shade	62 W	17 K	**55 °C**
White, shaded	30 W	12 K	**50 °C**
White, shaded, filtered fan	30 W	4.4 K	**42 °C**
Reference: black, shaded, fan	96 W	8.4 K	**46 °C**

Two readings matter. Paint is worth **22 K** and it is free at the design stage, which makes a black lid the most expensive aesthetic decision on the rover. And a sealed, white, shaded box is still **12 K above ambient** at 40 W, so "sealed and white" alone is not a thermal design; it is the starting point for one.

The shade geometry is usually wrong. A shade at height h displaces its shadow by h divided by the tangent of the solar elevation: at noon a 25 mm standoff moves the shadow only **6.7 mm**, but at 29° elevation the offset is **45 mm** and the west face becomes the absorber. A top-only shade is therefore a noon device. The shape that works is an open-bottomed cap with **at least 50 mm of overhang on all four sides**, about 0.35 kg in 0.5 mm aluminium for a box this size.

The fan needs two corrections that a bench test hides. At 1,309 m and 38 °C the air density is **0.96 kg/m³** against 1.13 at sea level, which is 85 %, so mass flow and heat transfer fall by **15 %** for the same fan speed. And the filter adds back-pressure. Two independent sizings, 10.3 CFM by rule of thumb and 12.8 CFM from the physics, agree within 25 %, which gives **13 CFM at the site**; adding 30–50 % for the filter points to a 40–60 mm fan rated at **20 CFM or more free delivery**, with the filter mounted outside the box. Filter face area, not filter grade, sets service life: a fan's own 25 × 25 mm intake has 6.25 cm², while a remote 80 × 80 mm filter has 64 cm² and ten times the dust capacity for the same flow.



Monash's Banksia on loose sand in 2026. This is the dust source: the wheels lift the material that then settles on the lid, the lens and the connector shells, and the same surface is what the tyres have to climb. Loose sand is a thermal problem too, because a wheel spinning in it heats the motor housing that sits closest to the ground.

Monash Nova Rover (Tier A), 2026. Original source Full resolution

Two more thermal numbers are worth carrying. A stalled arm joint dissipates about **140 W** in a winding with a thermal mass near 500 J/K, which is **0.28 K per second** at first and an equilibrium **70 K above ambient**, so a hold that is comfortable for a minute passes a derating threshold in four or five. And a compute module that throttles at minute twelve of a 40-minute run looks exactly like a software defect, which is why the temperature channel and a "too hot" state belong on the operator display.

The dust is clay and salt, which is worse than sand

The soils around the station are rich in montmorillonite, nontronite and illite, sulfate-dominated and basic at pH 8–9 `S-STOKER-2011-MDRS-SOILS`. Three properties follow, and none of them is captured by a rating printed on a component.

1. **It is finer than the dust an ingress rating uses.** Clay minerals are defined below about 2 µm, while the standard ingress test uses a dry, free-flowing talc-based dust. Passing the test on talc is not a promise about a hygroscopic clay.
2. **It swells and cakes when damp.** With morning dew or one shower, a dry powder becomes a paste that bridges connector contacts and packs into gear teeth.
3. **It is salty.** Sulfate salts are soluble, so a damp film on a circuit board is an electrolyte. Dust at this site is a corrosion problem as well as an abrasion problem.

No team publishes per-component failure counts, so the ranked failure list is a mechanism-based prior rather than a census: **1 connectors and crimps, 2 fans, 3 lenses, 4 USB and HDMI ports, 5 potentiometers, 6 limit switches, 7 bearings and gearboxes, 8 heat-sink fins.** The leading two are cheap to fix and expensive to ignore. Connectors want a latch, a pull-tested crimp, a service loop and dielectric grease; fans want a speed alarm and an oversized remote filter.

Cleaning is a scheduled activity, not a reaction. A twelve-minute routine of brush, filter, lens, connector check and a look for dust trails fits the ten-minute turnaround only when it is split across two missions, which makes it a between-mission task rather than a between-task one. And **PLA has no place anywhere it can get hot**: at a 55 °C softening point it sits inside the box band and inside a closed June vehicle, and the failure is silent creep that gets misdiagnosed as software.



The Hab at the Mars Desert Research Station, where the event runs. Pale ground, bright sky and no natural shade: the visual environment the rover is designed into, and the reason a bare aluminium or black anodised lid is the wrong finish.

McKay Salisbury (own work), 2004-06-16. *Original source* *Full resolution*

A repair costs 20 % of a mission, so build for the swap

A runner may fix the rover in place, take it to the station or return it to the start `S-URC27-RULES` `S-URC27-RULES` §3.e.i §3.e.iv, but a team member who leaves the station cannot operate again that mission `S-URC27-RULES` §3.e.iii. Every intervention costs **20 % of the mission's points** and the penalties are additive, so two interventions leave 60 % of the points earned `S-URC27-RULES` §3.e.viii.

The design rules that follow are not heroic. Repair between missions is free, so the kit's job is to make that repair fast. The swap must be possible without a laptop, a manual or a phone call, which means a keyed connector and a printed removal procedure for every mission-critical module. And some repairs are not worth doing: if the remaining points are worth less than 20 % of what is already banked, the mission is abandoned and the hardware is saved for the next one.

FAILURE	FIELD REPAIR	TARGET TIME	SPARE NEEDED
Connector unlatched or crimp failed	Re-seat, replace a pre-made harness	5 min	Pre-crimped harness per branch
Blown fuse	Replace, then find the cause before re-powering	2 min	Fuse assortment, two of each rating
Motor controller dead	Swap the controller and re-flash parameters	15 min	One spare per class
Motor or gearbox failure	Swap the actuator assembly	30–60 min	One spare actuator or a rebuild kit
Tyre puncture or delamination	Swap the wheel	10–15 min	One spare wheel and nuts
Seized bearing	Press or pull the bearing and replace	30–45 min	Bearing set and puller
Camera or USB failure	Swap the camera and re-plug the cable	10 min	One spare camera and cable
Compute will not boot	Boot from a known-good image, or swap the board	10–20 min	Cloned boot medium, spare board
Firmware regression	Roll back to the tagged known-good build	15 min	Tag plus an offline copy of the code
Battery pack failure	Swap the pack and isolate the bad one	10 min	A second pack

Spares classify by one question: if this fails, can the rover still score? **Class A** is mission-stopping and cheap, so it travels at two to three times: fuses, connectors, crimps, wire, heat-shrink, tape, threadlocker, grease, cleaning fluid, boot media, fasteners. **Class B** is mission-stopping and expensive, so it travels once and is proven before departure: one spare controller per class, one actuator, one camera, one radio, one compute path, one battery. **Class C** is recovery only and can live in the vehicle: chassis plate, rivets, spare wheel, spare mast. Spares and equipment used to maintain the rover are excluded from the 50 kg limit `S-URC27-RULES` §2.a.iii, which is why the kit can be generous.

Transport is the strategy's first test. A 100 kg crate in a 0.8 g stop needs **785 N** of restraint; a 100 mm drop onto 25 mm foam is about **4 g**, and onto 10 mm foam about **10 g**. A 50 kg rover at 10 g puts **4.9 kN** into four anchors, which an M5 bolt survives and a T-nut in 12 mm plywood does not. The qualification is one logged afternoon: an accelerometer at 100 Hz or better on the crate floor, 60 minutes of driving including 10 of washboard, the same checklist before and after, and every fastener marked with a paint pen.

Interfaces that bite

- **Mobility and arm.** Sealed hubs add drag and vented hubs ingest clay, so the decision is made before the wheel is machined. A long arm hold is a continuous thermal load that a brake would remove, at the worst place to add mass: the arm's far end.
- **Communications.** A 3 m antenna limit, a 25 m cable run and a 66.6 km/h gust are in tension, because a larger antenna buys link margin and costs mast stiffness.
- **Power.** A 355 Wh 6S pack or a 592 Wh 8S pack both exceed the 160 Wh aircraft limit, so the pack is freight, a purchase in the United States, or a drive `S-FAA-LITHIUM`.
- **Software and operations.** Throttling at minute twelve looks like a software bug, so software owes a temperature channel and a visible "too hot" state. The crate is the weigh-in box, the shock absorber and the field bench at once.
- **Systems engineering.** The crate and the kit belong in the transport plan and the mass model, even though spares are excluded from the rover mass.

Failure modes

- **The connector that let go in transit.** The judges name loose wires and connectors as the most common way a rover works in the parking lot and fails at the site `S-URC-ADVICE`. Latched connectors, strain relief and witness marks are the prevention.
- **The fan that blinded its own filter.** Clay loads the blades, the bearing ingrates and the filter blinds, and the only symptom is a temperature that climbs. An oversized remote filter and a speed alarm convert a silent failure into an alarm.
- **The potentiometer with a dead spot.** A contaminated track jumps, calibration drifts, and a joint position becomes untrustworthy. Contactless magnetic encoders and Hall end-stops remove the mechanism entirely.
- **The PLA bracket that crept.** Softening at 55 °C produces sag and misalignment in a sun-loaded box. Anything hot is metal, ASA or PC.
- **The fastener that seized.** A sulfate film plus morning dew between stainless and aluminium seizes a fastener in a week. Anodising, nylon washers and dielectric grease are the answer.

- **The firmware that would not boot.** A hard kill can corrupt a filesystem and there is no network to recover from. A read-only root, cloned boot media and tagged builds make the recovery a swap.
- **The missing two-dollar part.** The nearest supplier is six hours away, so a missing connector ends an event after the diagnosis.

Shortlist for a first build

1. **Sealed white compute box, a shade with 50 mm of overhang, one filtered fan with a speed alarm, and a heat spreader to the skin**, when the compute dissipates 15 W or more. A fanless sealed box works below about 10 W with 0.5 m^2 of skin.
2. **Latched sealed connectors, service loops, dielectric grease and a conformal coat that stops at the contact.** This is the cheapest fix for the most common mission stopper.
3. **A plywood crate with at least 1.25 m internal, a bolted pallet, closed-cell foam, ratchet straps and a logged 60-minute road test.** The crate is also the weigh-in box.
4. **A field kit built from the failure list**, with the cheap mission-stopping class at two or three times, and module swaps rather than field soldering.
5. **A colour scheme and a finish chosen for thermal reasons**: white or bright on every sun-facing surface, no black lids, no printed parts in the hot path.

Open problems

- **What is the dust size distribution at rover height?** The published MDRS work measures habitat intrusion, and the particle counts are behind a paywall. An optical particle counter at wheel height for 30 minutes produces the number.
- **What is the real surface temperature of a sun-loaded lid?** The model says 50–77 °C and no site measurement exists. An infrared thermometer on the lid, the battery box and the wheel hubs at 13:00 and 17:00, over five days, settles it.
- **Does an ingress rating mean anything for clay?** The standard test dust is talc. A team-defined clay-and-vibration test, published with its method, is the only honest answer.
- **What is the filter service life in clay dust?** Running the fan through a pre-weighed filter for an hour in the dustiest local conditions gives a flow drop and a service interval.
- **What does the road do to the rover's mounts, rather than to the crate?** Logging an accelerometer on the chassis instead of the crate produces the load case the design actually sees.
- **Can a first-season team hit a 20-minute repair?** The times above are targets. A "dark rover" drill, in which a team-mate disables one item at random and the crew diagnoses it outdoors against the clock, turns them into data.

Further reading

- Missouri S&T and Monash both document sealed or fan-cooled electronics and the failures that followed.
- Power, electrical architecture and safety owns the pack, the fusing and the thermal limits this page assumes.
- Compute, sensors and perception owns the module that has to survive the box.
- How rovers fail places these mechanisms in the wider failure record.
- Site, travel and field constraints covers the road, the altitude and the rules of the site.

title: Astronaut interaction and the EVA suit description: Fifty points built on a human: the Follow! ladder, the suit heat limit, and a safety case where contact costs nothing but an intervention ends the run. part: Subsystems chapter: "4.12" toc_depth: 3 tags: - astronaut - eva-suit - human-robot - autonomy

Astronaut interaction and the EVA suit

The Autonomy Mission asks a rover to find a suited person, obey them, follow them and hand them a pick, which makes a human both the commander and the target of the same system. This is new for 2027: nine years of earlier tasking cover finding a person and none of it covers obeying one. Three things decide the 50 points, and none of them is the rover's hardware: the Follow! ladder, the heat inside the suit, and a safety case that inverts the usual incentive.

50

ASTRONAUT SUB-MISSION POINTS

Suit 5 · drive 5 · Follow! 15 · Stay! 5 · Fetch! 10 · Come! 5 · Give! 5

2027

S-URC27-RULES

5 / 5 / 10 / 15

FOLLOW! BY COMMAND METHOD

Device, sign, audio, gesture, and they do not stack

2027

S-URC27-RULES

3 m

STOP DISTANCE AND STAND-OFF

The scoring distance is also the safety distance

2027

S-URC27-RULES

0.50 m

FULL STOP FROM 1.0 M/S

2.5 m/s² with 0.3 s latency, leaving 2.5 m of margin

Model

S-URC27-RULES

Key takeaways

- Follow! is priced by **method**, not by capability: **15** points for a gesture, **10** for a spoken word, **5** for a sign and **5** for a device, and the four do not stack S-URC27-RULES §1.e.vi. The entire modality choice is therefore a 10-point swing out of 50.
- A **ladder**, not a rung, is the design. Switching methods inside the mission is allowed and encouraged S-URC27-QA AUQ6, so a design that attempts gesture, falls back to audio and ends at a button adds probability at each step and can never subtract points.
- **Contact carries no points penalty, and that is the trap.** An unrequested intervention counts as an emergency stop and terminates the mission S-URC27-RULES §3.e.ii, so the real cost of driving too close is every remaining point in both sub-missions.
- The suit is an **integration problem, not a costume.** Heat is the failure mode, and a consumer digital video unit shuts itself down after **five minutes at 35 °C** S-DJI-03 , inside a suit during a 14-minute task.

What the subsystem must do

The sub-mission delivers **50 of the Autonomy Mission's 100 points** inside a shared 40-minute clock S-URC27-RULES §1.e.i. The rover drives to a suited team-mate, is commanded by them through five ordered commands, follows them, waits while they walk more than 20 m away, fetches a rock pick from the ground, returns and gives it back S-URC27-RULES S-URC27-RULES §1.e.iv §1.e.x. Four of the five commands are worth 5 points each and Follow! is where the margin sits.

ELEMENT	POINTS	EARNED BY	FAILS IF
Suited astronaut with camera and microphone	5	A suited figure plus a live feed on a station display	The feed exists but is not monitored in the station
Drive to the astronaut	5	An autonomous stop within 3 m of the given point	Teleoperated, or a stop beyond 3 m
Follow!	5–15	The method used, then a stop within 3 m of the stopping astronaut	Command not recognised, target lost, final stand-off too large
Stay!	5	No motion while the astronaut walks more than 20 m away	Drift, creep, or a re-plan that moves the rover
Fetch!	10	An autonomous grasp and lift of the pick from the ground	A teleoperated grasp scores zero; a dropped pick fails
Come!	5	An autonomous drive to the astronaut, stopping within 3 m	Obstacles, overshoot, or a stop beyond 3 m
Give!	5	The pick placed on the ground or dropped	The pick is left in the gripper when time expires

The command order is fixed S-URC27-QA AUQ13, and **every command except the first drive originates from the astronaut** S-URC27-QA AUQ7. The first drive is the only point in the mission where the station may choose a destination, and the walk route may never be known to it at all, because the judges show the approximate locations to the astronaut during setup S-URC27-QA AUQ14.

The numbers that size the design

The heat and wind rows come from an open weather archive. The ladder, stand-off and expectation rows are engineering estimates from stated parameters and are marked as such where they appear.

QUANTITY	VALUE	BASIS
Follow! ladder	gesture **15** · audio **10** · sign **5** · device **5**, maximum 15	The method sets the value and the methods do not stack S-URC27-RULES §1.e.vi
Expected value per person-week	sign **3.5** · device **3.2** · audio **2.0** · gesture **1.0**	Points multiplied by an assumed success probability, over an effort estimate
Stand-off and stopping distance	3.0 m target, never below 2.0 m; **0.50 m** to stop at 1.0 m/s, **3.33 m** at 2.5 m/s on loose sand	The margin disappears quickly with speed and latency
Heat at the site, June	mean daily maximum **33.3 °C**, 95th percentile **38.9 °C**, maximum **40.7 °C**; **10 %** of June days reach 37.8 °C	2015–2025, 330 June days S-OPENMETEO-ARCHIVE
Wind	mean daily maximum **7.3 m/s**, 95th percentile **12.3**, gust maximum **25.6**	An exposed microphone at 7 m/s is a roar
Walking pace to design for	astronaut **0.6–1.0 m/s**; rover follow at **1.2 m/s** or less, and **1.0 m/s** with a person in front	A suited student on sand, watching footing, by inference
Digital suit video	one consumer unit shuts down after **5 min at 35 °C** standby	Disqualifying inside a suit S-DJI-O3
Recommended suit video	about **USD 163** for an analogue 5.8 GHz pod	Camera 36.99, transmitter 77.99, receiver 48.49 S-GETFPV-FOXEER S-GETFPV-RUSHTANK S-GETFPV-SKYDROID
Fetch! object	**920 g**, 300 mm, alloy steel with a rubber grip, **USD 14.99**	The pick the rules link §1.e.viii\ : KEILEOHO 32 oz Rock Pick Hammer, alloy steel head and handle, 11.8 × 6.7 × 0.98 in, 920 g, anti-skid rubber grip, USD 14.99 (1-pack flat head) · Amazon / KEILEOHO · n.d.">S-AMAZON-ROCKPICK

Design space, in one table

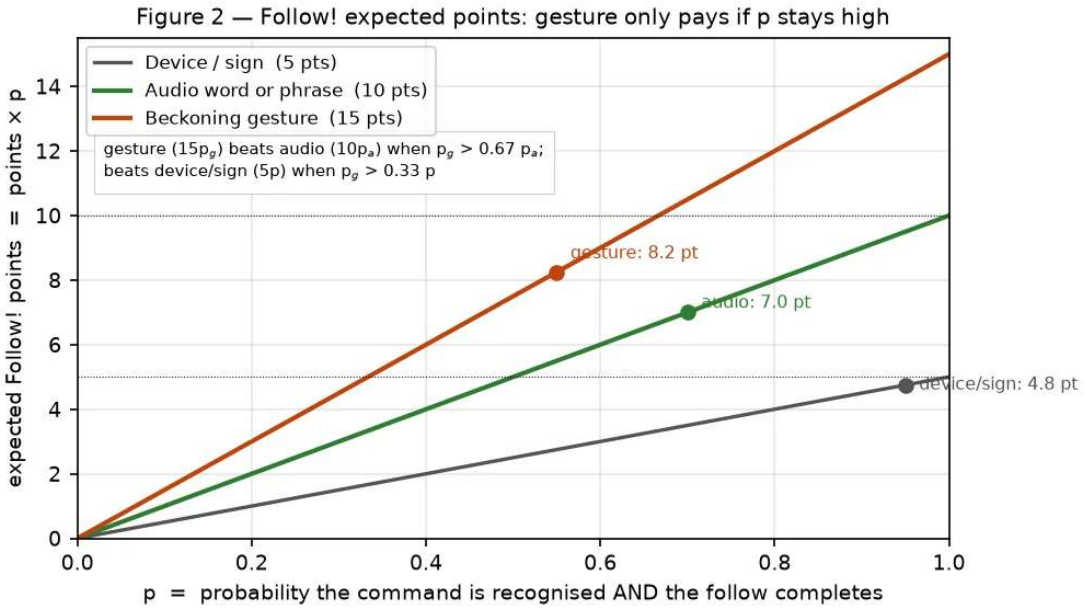
DECISION	OPTIONS	RECOMMENDED FOR A FIRST BUILD
Garment and heat	coverall · coverall plus an evaporative vest · a frozen pack vest · helmet or no helmet	A light compliant coverall, an evaporative cooling vest and a wide hat; skip the helmet or use a one-hand release, and test removal with gloves on
Suit camera and link	analogue 5.8 GHz · digital high definition · a small computer with an IP camera · a phone · an action camera	An analogue 5.8 GHz pod, because its failure modes are analogue and the station sees a plain video device
Command device	LoRa button · direct radio peer link · a phone application · tone decoding · a tether	A LoRa button independent of the video band, sized for a glove, decoded entirely by machine
Sign	an ArUco board · a printed word · a pictogram · an electronic display	A printed 20 cm 4x4_50 tag on rigid matte board, reusing the detector route finding already needs
Audio	rover microphone · suit microphone to the station · a wake word	Suit microphone into station recognition, so one microphone earns both the 5-point feed and the 10-point command
Gesture	pose landmarks plus a wrist rule · an arm fiducial · a trained sequence model · a sleeve inertial sensor	Pose at 6–8 m with an arm fiducial, promoted only after zero false positives on the real suit
Identity of the follow target	any person · a fiducial panel · colour · a beacon	A chest and back fiducial panel, because the field judge walks with the astronaut and “nearest person” is a defect

The Follow! ladder is an expected-value argument

Follow! is worth 5, 5, 10 or 15 points by the method the astronaut uses, and the methods do not stack S-URC27-RULES §1.e.vi. Pricing each rung by the effort a first build needs, with deliberately generous success probabilities, gives an ordering that is not the obvious one.

METHOD	POINTS	SUCCESS ASSUMED	EXPECTED POINTS	EFFORT	EXPECTED POINTS PER PERSON-WEEK
Device button	5	0.95	**4.75**	1.5 weeks	**3.2**
Visual sign	5	0.70	**3.50**	1.0 week	**3.5**
Spoken word	10	0.70	**7.00**	3.5 weeks	**2.0**
Beckoning gesture	15	0.55	**8.25**	8.0 weeks	**1.0**

The conclusion is not to pick a rung but to build a ladder. The design attempts gesture, speaks a command if the gesture is not recognised inside a timeout, then presses the button. Because switching methods mid-mission is permitted S-URC27-QA AUQ6, each rung is strictly additive in probability and cannot subtract points; the worst case is the 5-point device.



The three lines are the three point values. Gesture only beats audio while its success probability stays above about two thirds of the audio probability, and beats the device above a third of it. A gesture pipeline that works two times in three, at four times the effort of the button, is worth less than the boring option used as the top rung of a ladder.

Robotics Club @ Mississippi State University, 2026. Original analysis.

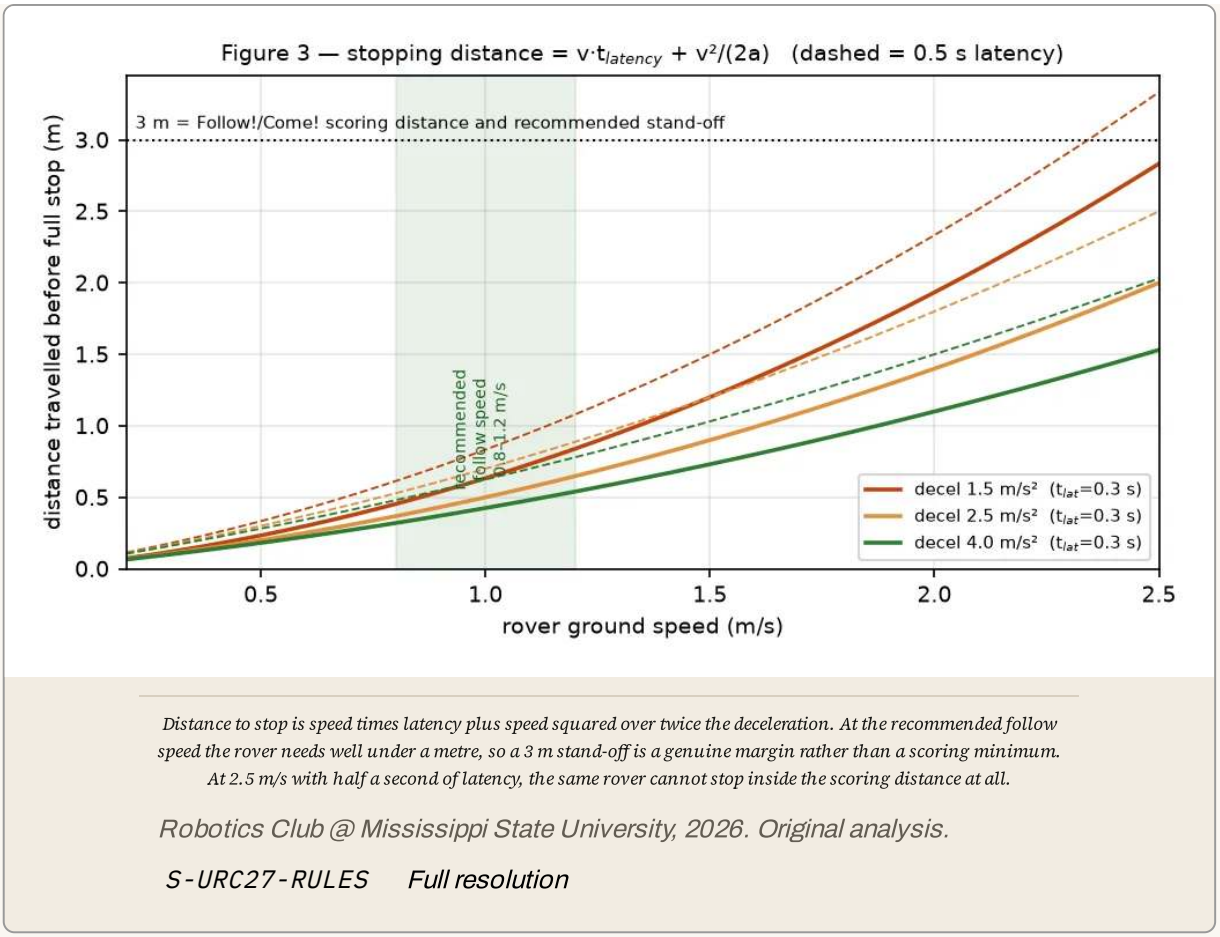
S-URC27-RULES *Full resolution*

There is a trap in the cheap end of the ladder. **Every device channel must be decoded by machine, end to end.** The moment a station operator hears a beep and presses a button, the command no longer originates from the astronaut, and the team is teleoperating while the rover drives autonomously, which the rules forbid S-URC27-RULES S-URC27-QA S-URC27-QA §1.e.iii AUQ2 AUQ21. A 5-point button is a one to two person-week firmware project, not a weekend. The gesture rung carries the opposite risk: it is the only rung that can drive the rover at a person on a false positive, so it stays conditional on a measured 90 % true-positive rate at 6–8 m on the real suit.

Safety: contact has no points penalty, and that is the trap

Three rules interact, and teams misread the middle one. **Contact carries no points penalty**; avoidance is highly recommended, and the competition organisers would sympathise with an astronaut who stops a rover that is getting too close S-URC27-QA AUQ1. **The E-stop is on the rover**, reachable from beside it, and it removes all battery power S-URC27-RULES §2.a.v. And **an unrequested intervention counts as an emergency stop**, terminating the mission with points preserved to that moment S-URC27-RULES §3.e.ii.

So the real cost of getting close is **every remaining point in both sub-missions**, and the arithmetic inverts the usual incentive. At 1.0 m/s with 2.5 m/s^2 of deceleration and 0.3 s of latency the rover stops in **0.50 m**, leaving 2.5 m of margin inside a 3 m stand-off. At 2.5 m/s with 0.5 s of latency and 1.5 m/s^2 on loose sand it needs **3.33 m**, which is more than the scoring distance. The operating envelope follows: follow at **1.2 m/s** or less, stand off at **3.0 m**, never close below **2.0 m**, and treat “a person in front” as a hard limit of **1.0 m/s**.



Distance to stop is speed times latency plus speed squared over twice the deceleration. At the recommended follow speed the rover needs well under a metre, so a 3 m stand-off is a genuine margin rather than a scoring minimum. At 2.5 m/s with half a second of latency, the same rover cannot stop inside the scoring distance at all.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *Full resolution*

The safety protocol fits on a card that both the astronaut and the station carry. Walk the intended path together before the clock, agree which side the rover stays on, and agree an abort word. **Voice is the first safety action**: the astronaut says stop, the station gives the free autonomous stop, and the E-stop is reserved for imminent contact. The astronaut never walks toward the rover; if the rover is in the wrong place, the astronaut stops and gives the command again from a standstill. Bystanders are told that a red LED means autonomous and where the E-stop is, and the drone pilot is told not to fly low over the suit. Hydration, shade and a planned 12-minute exposure per attempt are part of the safety case, not part of comfort.

The suit is an integration problem, not a costume

The rules make the suit easy and then hard. No pressurisation, no backpack and no oxygen are required, and colours and patterns are free, but it must cover the neck, wrists and ankles with closed-toe shoes S-URC27-QA AUQ12, carry a camera and microphone that are **monitored in the station** for 5 points S-URC27-RULES §1.e.iv, include a helmet only if the astronaut can remove it alone, and be practical at 37.8 °C S-URC27-RULES §3.d.iii.



Analogue suits from a desert research station in 2008: light coverall-style garments, neck rings, gloves, boots, chest harness and a backpack. The 2027 rules require none of the pressure hardware, which is why the realistic target is a light coverall with a cooling layer and a camera, not a replica.

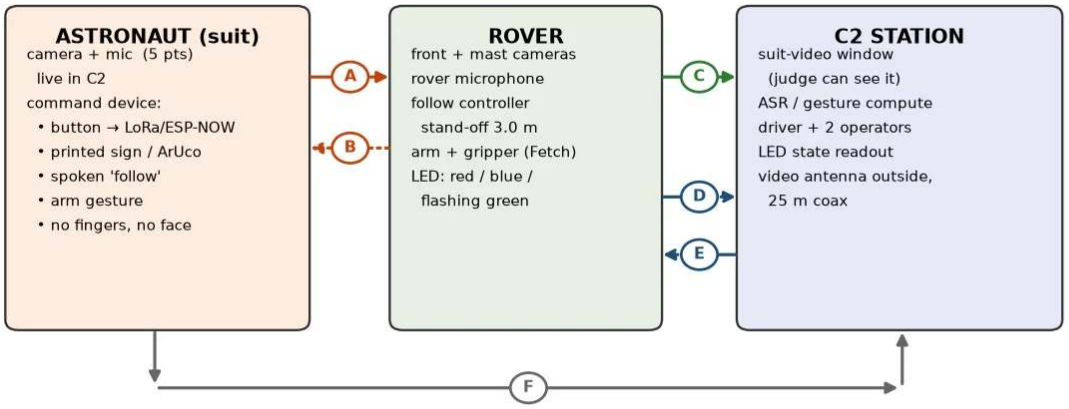
Bandgirl807 / Wikimedia Commons, 2008. CC BY-SA 3.0.

`S-WIKIMEDIA-MDRS-SUITS` `S-URC27-RULES` *Full resolution*

Four consequences follow, and each one is a decision.

1. **The 5 points live on the path from the astronaut to the station, not on the camera.** A recorded feed earns nothing; a feed that is always visible on a station display earns the points. The topology matters too: a direct suit-to-station link survives the rover driving away, while a relay through the rover dies exactly when the astronaut is more than 20 to 30 m off, which is the Stay! and Fetch! geometry by design.
2. **Heat is the failure mode, and digital video loses to it.** A consumer digital unit shuts down after five minutes at 35 °C standby `S-DJI-O3` , and a phone sits at or beyond its own 0–35 °C operating range `S-APPLE-TEMPERATURE` . An analogue pod draws one or two watts, needs no operating system and degrades to a fuzzy picture, and the suit battery is sized for twice the planned exposure.
3. **Identity is a design problem.** The field judge and spectators walk with the astronaut `S-URC27-QA` AUQ22, so person following has to be identity following. A chest and back fiducial gives identity, range and bearing, and the back panel is what the rover sees during the walk-away geometry. Any colour or pattern is legal `S-URC27-QA` AUQ12, so this costs a printout.
4. **The paperwork is part of the subsystem.** The suit transmitter, the command device and any repeater are communications equipment and must be declared in the communications documentation by **2027-04-30** `S-URC27-RULES` §2.c.ii. The suit itself is excluded from rover mass `S-URC27-QA` AUQ12, while the rules list it inside the fielded-equipment accounting; itemising each suit line and marking its treatment is the defensible procedure.

Figure 4 — three-node command and video architecture for the Astronaut Assistance sub-mission



A suit camera/mic frame relayed through the rover (2.4 GHz Wi-Fi or 5.8 GHz, <30 m) B rover acknowledgement / command echo back to the suit device
C suit camera/mic direct to the C2 (5.8 GHz analog or 433/915 MHz) — stands alone if the rover link dies D primary rover link, 2.4 / 5.8 GHz
E telemetry, stop and return-to-astronaut commands F two-way voice (QA AUQ5): the astronaut can be warned and can report hazards
Each path must survive 10–20 dB body shadowing when the astronaut turns away, and the rover driving beyond the suit transmitter range.

Three nodes and six links. The suit camera and microphone travel either directly to the station or through the rover, and only the direct path survives the astronaut walking out of the rover's radio range. The same diagram carries the command channel back to the suit, which the rules encourage and which a two-way link makes free.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *S-URC27-QA* *Full resolution*

Nine years of precedent cover finding, not obeying

The competition ran a task called Astronaut Assistance for six consecutive years, and the geometry it enforced is the geometry of the 2027 drive and return tasks. In 2011 to 2016 the position fix was good to only about **25 m**, so a visual search was mandatory, and the astronaut was identified by a simulated space suit rather than by a coordinate. The delivery radius was **1 m**, tighter than the **3 m** the 2027 tasks allow `S-URC11-RULES` §3.d.

What transfers is precision at the end of a position-led approach, plus the habit of treating the suit as the visual marker. What does not transfer is any experience of a human giving commands: the sign channel has appeared before as an object identity held up in Delivery, but it has never been a command.

Two constraints on the practice regime follow. The station has **no view of the course** `S-URC27-RULES` §3.d.i, so every rehearsal runs blind, with operators in a closed room and the astronaut outdoors. And the published bar from a top team is **100 hours** of operator practice `S-WVU-SAR2026` ; a first build needs about **38 hours** of dedicated sessions, from a bench session on the suit link through a command bench, open-ground following, a full blind run, fault injection and a dress rehearsal on unfamiliar ground.

Interfaces that bite

- **Communications.** The suit pod competes for the same 5.8 GHz band and the same station antennas as rover video, and its two to five megabits per second belong in the link budget.
- **Arm.** Give! must lower to 10 cm or less before release, with the astronaut at least 2 m away, because a dropped 0.92 kg steel pick near a person is a safety event rather than a lost point.
- **Mobility.** Following needs a controllable speed of 1.2 m/s or less and a measured deceleration of 2.5 m/s^2 on loose soil. Both are software limits, and neither may fight the Delivery mission's 1.0–1.3 m/s average.
- **Software and operations.** A five-state machine with latched transitions, duplicate suppression, and a link-lost stop that needs no station action; plus one always-visible suit window and a transcript display.
- **Environment.** A dust-tolerant lens and a camera that does not thermally shut down. This is a thermal problem before it is a video problem.

Failure modes

FAILURE	MECHANISM	PREVENTION
The rover follows the wrong person	The field judge or a spectator walks near the astronaut	Require the suit fiducial before closing distance; ignore any person without it
The target is lost mid-walk	The astronaut turns away, is occluded, or a light suit blows out in sun	Predict from the last bearing, coast for a moment, then stop and wait rather than searching
A bystander's arm is read as a command	Natural arm swing or waving triggers the gesture pipeline	Two-of-three confirmation, a rest pose before accepting a trigger, and the fiducial as a gate
The command is recognised twice	A repeated trigger re-arms the controller mid-task	Command identifiers with a latch, and duplicate suppression inside 2 s
Stay! creeps	A re-plan toward a stale goal, or obstacle recovery nudging the rover	An explicit hold state that disables the planner, verified on the wheel encoders
Fetch! closes on air	The grasp point is the image centroid of the mask, and the handle is under the head	Bias the grasp toward the head and confirm with depth before closing
The suit feed dies mid-task	Body shadowing, battery sag, or thermal shutdown	A shoulder-mounted antenna, an analogue link, and twice the battery capacity needed

Shortlist for a first build

ITEM	CHOICE	WHY
Suit	Light compliant coverall, evaporative vest, wide hat, no helmet or a one-hand release	Meets coverage, survives heat, and leaves the removal rule trivially satisfied
Suit video	Analogue 5.8 GHz pod with the antenna on the shoulder	Analogue failure modes, a plain video device at the station, about USD 163
Command ladder	Printed ArUco sign, LoRa button, then audio; gesture only after a measured test	The three cheapest rungs cover 10 points reliably
Identity	Chest and back fiducial panel	Removes the wrong-person failure for the cost of a printout
Follow behaviour	1.2 m/s or less, 3.0 m stand-off, voice-first stop protocol	A 3 m stand-off is a real margin at the recommended speed
Practice	About 38 hours of dedicated sessions, blind station, with a fault-injection session	The station cannot see the field, so blind practice is the only rehearsal that counts

Open problems

- **Is a sleeve inertial gesture legal?** The rule says visual recognition and the ruling names cameras `S-URC27-QA` AUQ24. Asking the organisers in writing settles it; until then the assumption is no.
- **How far can a 20 cm tag be read in full sun?** A range sweep at 5, 10, 15, 20 and 25 m at two sun angles gives a detect rate, which is what the sign rung depends on.
- **Does pose-based beckoning work at 6–8 m on a light suit against desert?** A hundred gestures at three distances with three subjects gives the true and false positive rates.
- **What is the measured emergency deceleration on MDRS-like soil?** Every number in the safety case depends on it. An instrumented stop test at four speeds on three surfaces closes it.
- **How does speech recognition behave in 7 m/s wind with the motors running?** Thirty utterances per condition give a false-accept rate, which is the gate on the audio rung.

- **Does the astronaut's body shadowing kill the suit link when they turn?** A 360° walk at 10, 20 and 30 m with a signal-strength log settles the antenna placement.
- **What happens if the astronaut presses the rover's E-stop?** The rules imply the mission is terminated and the rulings sympathise with the action. A written answer from the organisers, and a protocol that means it never happens, are both required.

Further reading

- WVU Team Mountaineers sets the published operator-practice bar and validates ArUco-based approach stopping in its 2026 review `S-WVU-SAR2026`.
- Missouri S&T trains a detector class for the rock pick and publishes its confidence thresholds `S-MST-DOCS-CONSTANTS`.
- Monash documents the wrong-place failure that this sub-mission punishes most: a failed detection in the field after a laboratory success.
- Autonomy: astronaut assistance maps every element of the task onto the rules.
- Software stack and autonomy owns the state machine this page assumes.
- Communications and networking owns the band plan the suit shares.

CROSS-CUTTING ANALYSIS · CHAPTER 5

Cross-cutting analysis

The twelve subsystem chapters answer what the options are for one part of the rover. Five questions cross all of them and decide whether a first build is coherent: what the whole rover masses and draws, what fights what, how rovers fail, what separates the top teams, and where the competition is going. The subsystem chapters remain the authority on any single subsystem. What crosses between them is the subject of this part.

50 kg

PER-CONFIGURATION MASS LIMIT

5 % of that event's points per kilogram over

2027

S-URC27-RULES

66

INTERFACES BETWEEN TWELVE SUBSYSTEMS

37 high criticality, 25 medium, 4 low

2027 design

54

SOURCED FAILURE ENTRIES

48 classifiable, 6 outcome-only

2015–2026

162

FINALS TEAM-YEARS BEHIND THE ZERO RATES

2022–2026 official score sheets

2022–2026

S-URC-RESULTS

Key takeaways

- **The envelope, not the 2027 task list, sets the budget.** The 50 kg limit has been in the rules since 2007, the 70 kg fielded union since 2015 and the 1.2 m cube since 2019, with no exception in the record S-URC07-RULES-L1 S-URC15-RULES S-URC19-RULES S-URC27-RULES §2.a.ii–iii.
- **A first build should target 45–46 kg with 4 kg of reserve, 400–500 Wh of pack and a 3.7 kW peak capability.** The measured field sits at 33–53 kg, and the pack is sized by peak current as often as by energy.
- **Sixty-six interfaces exist between twelve subsystems and 37 of them are high criticality.** The arm and perception each touch nine, so deferring either is a choice to re-open the most work.
- **The largest point loss in the field is an autonomy capability gap, not comms.** Sixty-seven of 162 finals entries scored exactly zero on autonomy, 41 % of the sample, worth about 610 of the roughly 945 mission points the field forgoes each year.

The five questions

QUESTION	THE ANSWER IN ONE LINE	PAGE
What does the whole rover mass and draw?	45–46 kg with 4 kg of reserve; 249 W average and 3.7 kW at peak; 400–500 Wh of pack	Mass and power
What fights what?	66 interfaces, 37 of them high; the arm and compute each touch nine, and a wheel diameter change re-opens seven	Interfaces and couplings
How do rovers fail?	54 sourced entries in thirteen classes; the expensive class is autonomy, the published class is the arm	How rovers fail
What separates the top from the rest?	The three hard missions and the absence of zeros, not the review; one zero is worth about 11 places and 85 points	What separates the top from the rest
Where is the competition going?	The physical envelope is frozen and the capability content is not	Where the competition is going
What does it add up to?	Ten system-level conclusions for a first-time team	Ten system-level conclusions

What the field publishes

The evidence base is thin and the thinness is part of the finding. Fourteen of 187 team-years (7.5 %) in the 2022–2026 record carry a published total mass, exactly one mass was read off a judge's scale, one team publishes a subsystem breakdown, and no team publishes a measured watt-hour figure for a mission S-CENSUS S-MIST-VURC2021 S-MONASH-GH-NOVA.

QUANTITY	PUBLISHED IN THE FIELD	CONSEQUENCE FOR THE ANALYSIS
Total mass of a competition rover	14 of 187 team-years (7.5 %)	Every distribution rests on teams that chose to publish
Judge-weighed mass	1 (MIST Phoenix, virtual URC 2021, 42.1 kg)	The only number that is a measurement S-MIST-VURC2021
Subsystem mass breakdown	1 (Monash Banksia 2026, rolling chassis and arm only)	Mass fractions are one case plus two allocations S-MONASH-GH-NOVA
Measured energy per mission	0	Every energy figure on these pages is a model
Centre-of-gravity table	0	Every tip-angle figure is an estimate with its basis stated

Each quantitative page states the sample behind its numbers, and where a chart plots fewer configurations than the table beside it. The full dataset can be filtered on the mass and power page.

The numbers worth carrying into a first build

NUMBER	VALUE	WHY IT MATTERS
Design mass	45–46 kg with 4 kg of reserve	The cap is 50 kg, the last 3 kg of a first build arrive in the final six weeks, and the penalty is 5 % of that event's points per kilogram over S-URC27-RULES §2.a.iii
Pack	400–500 Wh at 8S, two packs rather than one	Above the 289 Wh the back-to-back mission pair needs, and the current rating comes free with 4P
Interface count	66, of which 37 high	A late arm change re-opens nine interfaces; a wheel-diameter change re-opens seven
Zero rate	41 % on autonomy	One zero costs about 11 places and 85 points S-URC-RESULTS

How to read this part

A team sizing its first budget starts with mass and power, because the mass budget and the pack are the two constraints that every other decision spends. The interface page turns that budget into a freeze order: which numbers must be fixed at the preliminary design review and which can wait. The failure page decides where the test effort goes.

The last two pages are the field view. Top versus rest shows which missions produce the score gap. Trends shows which rules have moved and which have not. The ten conclusions collect the decision-relevant ones in one place.

Open problems

- **No measured energy figure exists.** A single logged 30-minute discharge with a power meter would turn the whole power model into a calibration, and no team in the field publishes one.
- **No team publishes a centre-of-gravity table or a mass budget**, so every stability and fraction number rests on stated bases rather than measurement.
- **The record is a publication census.** Teams write up the failures that drove a redesign and leave the rest; the judges' statement that comms is the most common problem and the database's tie between comms and the arm are both true statements about different things.

- **Two rule questions remain open** and change a design: whether a drone counts toward the 70 kg fielded total, and whether the Equipment Servicing cache is the team's own astrobiology cache `S-URC27-RULES §2.b.vii`.

Further reading

- Mass and power is the budget chapter, with the 17 mass configurations and the pack arithmetic.
- Interfaces and couplings carries the 12 × 12 matrix, the forty-one-row register and the three chains.
- How rovers fail carries the taxonomy, the zero record and the FMEA tripwires.
- Subsystem studies are the authority on each subsystem; Mission backfit maps the 2027 tasks to them.

CROSS-CUTTING ANALYSIS · CHAPTER 5.1

Mass and power

A first 2027 build should mass **45–46 kg with 4 kg of reserve**, draw **95–160 W average and 3–4 kW at peak**, and carry **400–500 Wh**, about 2.4–2.9 kg of pack. The rest of this page gives the evidence and the arithmetic behind those numbers. The evidence is thin on purpose: fourteen of 187 team-years publish a total mass, exactly one mass was read off a judge's scale, one team publishes a subsystem breakdown, and no team publishes a measured energy figure for a mission.

50 kg
PER MISSION CONFIGURATION
5 % of that event's points per kilogram over
2027
S-URC27-RULES

249 W
MEAN DRAW OF A 50 KG ROVER
Against 3.7 kW at peak, about 15 times

273–289 Wh
NAMEPLATE REQUIREMENT
60-minute Delivery, or the autonomy pair

400–500 Wh
RECOMMENDED PACK
Two 8S packs, 4P rated 180 A continuous
S-VENDOR-MOLICEL-P42A-DS

Key takeaways

- **The 50 kg limit is a constraint, not a target.** Seventeen configurations from eight teams have a mean mass of 47.9 kg and a median of 49.1 kg, with several pinned just under the cap and three at exactly 50.0 kg, which is not an engineering optimum.
- **The 1.2 m cube is the higher-stakes limit.** The oversize penalty is 40 % of one event's points, the equivalent of 8 kg of overweight at a score of 80, against 5 % per kilogram for mass.
- **The hotel load is nearly mass-independent and drive is strictly proportional.** The fixed load grows from 46 W at 35 kg to 63 W at 50 kg, while drive power grows with mass, so cutting 10 kg saves about 25 W of drive and almost nothing off the hotel.
- **Peak current, not energy, often sets the pack.** The credible simultaneous peak is 120–160 A at 24 V, three to four kilowatts, on a 95–160 W average, and a 4P cell stack rated 180 A lands at 496 Wh, inside the recommended band.

Three mass budgets, and what one kilogram costs

The 2027 rules create three mass budgets, not one S-URC27-RULES §2.a.ii–iii. The first is **50 kg per mission configuration**, with **5 % of that event's points per kilogram over**. The second is **70 kg across all fielded parts**, the union of rover, arm, science module and deployables, which is not a second free rover. The third is a **1.2 m cube** at every weigh-in, wheels and antenna included, foldable but not disassemblable, penalised at **40 % of that event's points**.

BUDGET	LIMIT	CONSEQUENCE OF EXCEEDING
Per-mission configuration	50 kg	5 % of the points scored in that event, per kilogram over
All fielded parts	70 kg	The union is over budget even when each configuration is legal
Stowed envelope	1.2 × 1.2 × 1.2 m	40 % of that event's points

Three consequences follow. **Mass is the mild penalty and size is the severe one:** scoring 80 and being 1 kg over costs 4 points, while a 40 % oversize penalty costs 32 points, the equivalent of 8 kg of overweight, per event. **Weight rounds down:** 50.9 kg is recorded as 50 kg and the cliff is at 51.0 kg S-URC27-RULES §2.a.iii. The weigh-in target should still be 49.5 kg or less on the team's own scale, with the rounding treated as unearned. **Penalties multiply the points scored** and add across events, so 2 kg heavy costs 10 % of the score, and the better the missions go the more mass costs.

The enforcement record is nearly empty. Across 187 team-year rows in the 2022–2026 sheets the penalty column is non-empty exactly once: **Mars Rover UdeG Space, 2025, 16.39 points**, with the cause not stated S-URC25-SCORES . USU measured 51.8 kg in

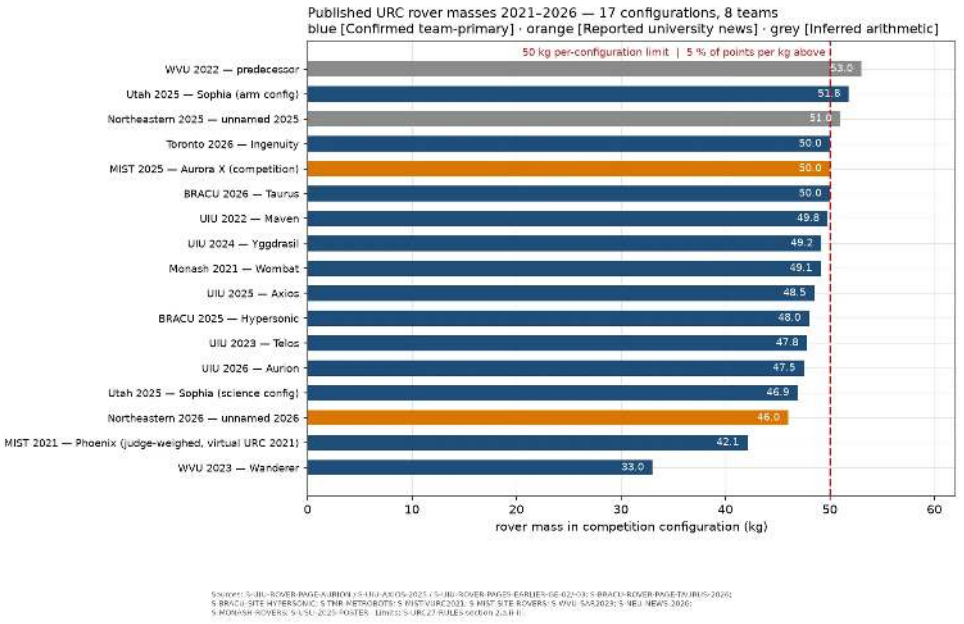
arm configuration in 2025 and carries no penalty in the sheet `S-USU-2025-POSTER`. The practical implication is that a first team cannot learn the real cost of being overweight from past results, only from the rule arithmetic.

The mass database: seventeen configurations from eight teams

Seventeen configurations across eight teams have a mass figure of some kind. Each row gives its basis: a team specification, a judge's scale, a university report or arithmetic from a stated change.

TEAM	YEAR	ROVER	CONFIGURATION	KG	EVIDENCE BASIS
UIU	2022	Maven	competition	49.8	stated by the team `S-UIU-ROVER-PAGES-EARLIER-G`
UIU	2023	Telos	competition	47.8	stated by the team `S-UIU-ROVER-PAGES-EARLIER-G`
UIU	2024	Yggdrasil	competition	49.2	stated by the team `S-UIU-ROVER-PAGES-EARLIER-G`
UIU	2025	Axios	competition	48.5	stated by the team `S-UIU-AXIOS-2025`
UIU	2026	Aurion	competition	47.5	stated by the team `S-UIU-ROVER-PAGE-AURION`
BRACU Mongol-Tori	2025	Hypersonic	competition	48.0	stated by the team `S-BRACU-SITE-HYPERSONIC`
BRACU Mongol-Tori	2026	Taurus	competition	50.0	stated by the team `S-BRACU-ROVER-PAGE-TAURUS-2`
WVU Team Mountaineers	2022	predecessor	competition	≈53.0	from a stated change `S-WVU-SAR2023`
WVU Team Mountaineers	2023	Wanderer	competition incl. 5-DoF arm	33.0	stated by the team `S-WVU-SAR2023`
Northeastern	2025	unnamed	competition	≈51.0	from a stated change `S-NEU-NEWS-2026`
Northeastern	2026	unnamed	competition	≈46.0	university news `S-NEU-NEWS-2026`
Monash Nova Rover	2021	Wombat	ARC competition	49.1	stated by the team `S-MONASH-ROVERS`
Utah State	2025	Sophia	science configuration	46.9	stated by the team `S-USU-2025-POSTER`
Utah State	2025	Sophia	arm configuration	51.8	stated by the team `S-USU-2025-POSTER`
Toronto MetRobotics	2026	Ingenuity	competition	50.0	stated by the team `S-TMR-METROBOTS`
MIST Mongol Barota	2021	Phoenix	virtual URC, judge-weighed	42.1	judged at weigh-in `S-MIST-VURC2021`
MIST Mongol Barota	2025	Aurora X	competition	50.0	stated by the team `S-MIST-SITE-ROVERS`

Mean 47.9 kg, median 49.1 kg, range 33.0–53.0 kg; fifteen of the seventeen sit between 42.1 and 51.8 kg.



The thirteen published masses against the 50 kg per-configuration limit, whose penalty is 5 % of points per kilogram over; the table above carries every configuration with its basis, including the inferred and university-reported ones. The pile-up just under the cap is the clearest pattern in the mass record.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-UIU-ROVER-PAGE-AURION` `S-UIU-AXIOS-2025` *Full resolution*

- **The median is 49.1 kg because the cap is what gets measured.** Several configurations sit at 49.8 kg or more, three of them at exactly 50.0 kg, and the pile-up under the line is not explained by any engineering argument.
- **WVU's 33 kg is the most important datum here.** Wanderer won **425.35 points** at 33 kg with a 5-DoF arm lifting 8 kg at full extension, and the 20 kg reduction came from fabrication, semi-monocoque sheet metal and composite wheels, not from a diet `S-WVU-SAR2023` `S-URC23-SCORES`. Fifty kilograms is 1.5 times a championship rover.
- **UIU is the best natural experiment in the field.** Five generations sit at 47.5–49.8 kg while the score rises from **236.67 to 400.44** and the battery grows from **432 to 840 Wh** `S-UIU-ROVER-PAGE-AURION` `S-UIU-AXIOS-2025`. Mass is the stable variable and capability is the moving one.
- **The only judge-weighed mass disagrees with the team's card.** The virtual-2021 stream read **42.1 kg** where the team reports 45 kg, and the judge's number wins `S-MIST-VURC2021` `S-MIST-SITE-ROVERS`.
- **Four of the five championship titles from 2022 to 2026 were won without a published mass:** Michigan 2022, AGH 2024, and Missouri S&T in both 2025 and 2026 `S-URC22-SCORES` `S-MST-GH-ORG`. Any mass quoted for Missouri S&T is invented until sourced.

The full dataset behind the numbers on this page, including the per-row source for every mass, is the table below. It switches between its seven sections and searches every row.

#	TEAM	YEAR	ROVER	CONFIGURATION	MASS (KG)	EVIDENCE BASIS	SOURCE ID	BASIS / N
1	UIU Mars Rover	2022	Maven	competition	49.8	team-published	S-UIU-ROVER-PAGES-EARLIER-GE-03	team rov page spe table
2	UIU Mars Rover	2023	Telos	competition	47.8	team-published	S-UIU-ROVER-PAGES-EARLIER-GE-02	team rov page spe table
3	UIU Mars Rover	2024	Yggdrasil	competition	49.2	team-published	S-UIU-ROVER-PAGES-EARLIER-GE	team rov page spe table
4	UIU Mars Rover	2025	Axios	competition	48.5	team-published	S-UIU-AXIOS-2025	team rov page spe table
5	UIU Mars Rover	2026	Aurion	competition	47.5	team-published	S-UIU-ROVER-PAGE-AURION	team rov page spe table; pa also stat 15 kg pay capacity
6	BRACU Mongol-Tori	2026	Taurus	competition	50.0	team-published	S-BRACU-ROVER-PAGE-TAURUS-2026	team spe table; ex at the 50 limit
7	WVU Team Mountaineers	2023	Wanderer	competition incl. 5-DoF arm	33.0	team-published	S-WVU-SAR2023	SAR p.1 verbatim “With this arm equipped Wanderer weighs 3 resulting 20 kg reduction from its predeces
8	WVU Team Mountaineers	2022	predecessor of Wanderer	competition	≈53.0	estimated from the team’s stated reduction	S-WVU-SAR2023	33 kg + stated 20 reduction a measu statemer
9	Northeastern University	2026	unnamed 2026 rover	competition	≈46.0	reported by the university	S-NEU-NEWS-2026	universit news: 10 “10 % lig than 202
10	Northeastern University	2025	unnamed 2025 rover	competition	≈51.0	estimated from the later figure and the team’s stated reduction	S-NEU-NEWS-2026	46 kg / 0.
11	Monash Nova Rover	2021	Wombat	ARC competition (no URC in 2021)	49.1	team-published	S-MONASH-ROVERS	team rov page
12	Utah State University	2025	Sophia	science configuration	46.9	team-published	S-USU-2025-POSTER	final-repo poster, measure
13	Utah State University	2025	Sophia	arm configuration	51.8	team-published	S-USU-2025-POSTER	final-repo poster, measure 1.8 kg ov the limit
14	BRACU Mongol-Tori	2025	Hypersonic	competition	48.0	team-published	S-BRACU-SITE-HYPERSONIC	team rov page
15	Toronto MetRobotics	2026	Ingenuity	competition	50.0	team-published	S-TMR-METROBOTS	team spe page; no URC 202 score (di compete
16	MIST Mongol Barota	2021	Phoenix	competition, virtual URC 2021	42.1	judge-weighed at the event	S-MIST-VURC2021	the only judge-weighed mass in t

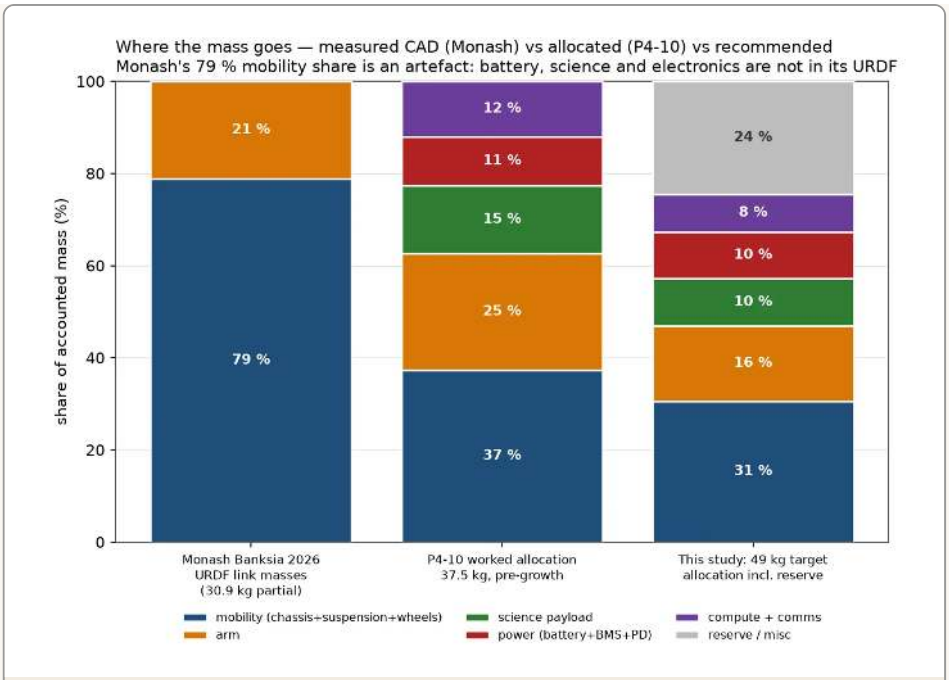
#	TEAM	YEAR	ROVER	CONFIGURATION	MASS (KG)	EVIDENCE BASIS	SOURCE ID	BASIS / N
								record: official U stream, 1 the team own card says 45 k
17	MIST Mongol Barota	2025	Aurora X	competition	50.0	reported by the team	S-MIST-SITE-ROVERS	team rov card

Where the mass goes

The only measured breakdown in the public record is Monash's 2026 robot description, covering the rolling chassis and arm only: no battery, electronics, science or drone S-MONASH-GH-NOVA . Of its **30.88 kg**, the wheels are the largest single line.

ELEMENT	MASS (KG)	SHARE OF THE 30.9 KG ACCOUNTED
Wheels ×4	7.40	24.0 %
Suspension legs	6.72	21.8 %
Arm incl. gripper	6.54	21.2 %
Chassis	5.46	17.7 %
Ankles and pivots ×4	4.32	14.0 %
Differential bar	0.42	1.3 %
Accounted total	**30.88**	100 %

Three things follow. **Wheels outweigh the arm**, so a 0.6 kg printed wheel instead of a 1.85 kg composite-spoke wheel saves 5 kg on a four-wheel rover. **Suspension and pivots are 36 %** of the accounted mass, which is where a six-wheel rover pays its bill against a four-wheel one. And **arm, wheels and chassis are comparable**, so a 50 kg rover is five or six decisions of 5–8 kg, any one of which going 50 % over eats the reserve.



Subsystem shares in three cases: Monash's measured chassis and arm, a 37.5 kg allocation before growth, and the 49 kg target used here. The Monash bar is read for its relative composition, not its mobility share, because the battery, electronics and science are missing from it.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-MONASH-GH-NOVA *Full resolution*

At a 49 kg target the allocation is explicit, and **the reserve is the payload line:**

LINE	KG	SHARE OF 50 KG
Mobility, six wheels and six motors	15.0	30 %
Arm, 5–6 DoF, about 1 m, gripper	8.0	16 %
Power, 500 Wh pack with BMS and distribution	5.0	10 %
Compute, sensors, comms, antennas	4.0	8 %
Science payload	5.0	10 %
Structure, fasteners, cable, mast hardware	2.0	4 %
Subtotal	**39.0**	78 %
Reserve	**10.0**	**20 %**
Total	**49.0**	98 %

The science payload itself is 3.5–5 kg, so most of the 10 kg reserve is not science. A team that plans a 9 kg science module against the 50 kg limit and then adds a mast, a drone battery and two kilograms of brackets discovers that the module was never the last line in the budget.

The arm is the only subsystem with a real spread. Cornell and Toronto RSX both fly six-degree-of-freedom arms at **3.53 kg** and **4.42 kg**, against **9.4 kg** for WVU's five-degree-of-freedom arm and **14.0 kg** for ARES with a 400 mm linear stage, a spread of **4.0 times** `S-RSX-ARM-URDF` `S-WVU-SAR2025` `S-ARES-PAPER-2025`. The reducer and the linear stage carry the mass, not the link length.

What the rover draws

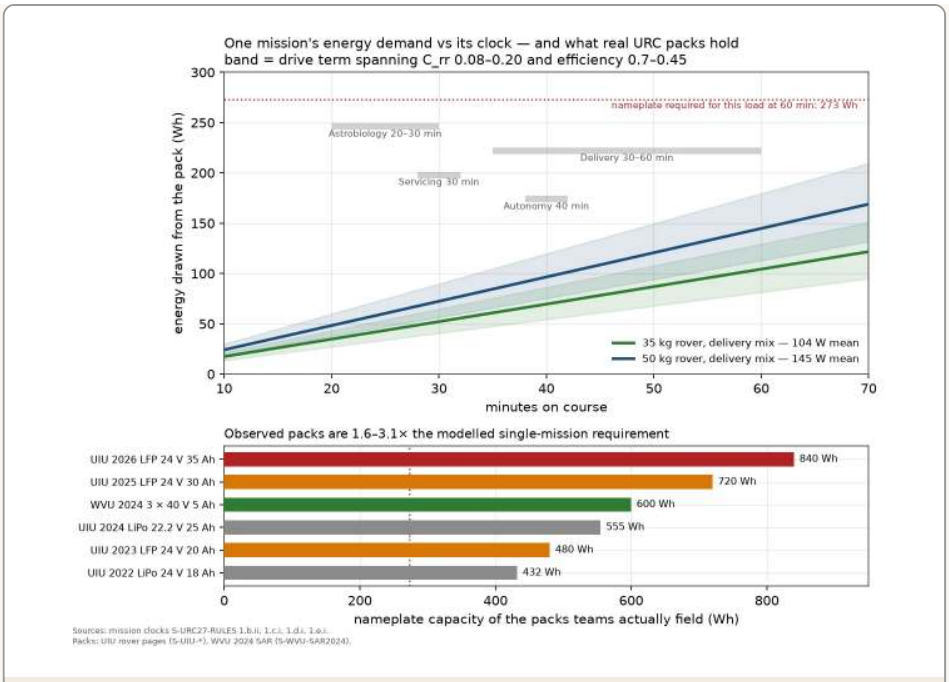
The power model is a first-order estimate, not a measurement; its purpose is to show which assumptions dominate. Loads are separated into the hotel, the drive, the arm and the science payload, and every mission is priced at 35 kg and at 50 kg.

LOAD	35 KG DESIGN	50 KG DESIGN	BASIS
Hotel at the battery	46 W	63 W	Compute, radios, cameras, fans and LEDs, plus 8 % conversion loss
Drive, while driving	86 W	123 W	2.46 W/kg at a rolling resistance of 0.12, a 3° mean grade and 0.8 m/s
Arm, while arm-active	20 W	29 W	3.0 W/kg of arm; the arm is 19 % of rover mass
Science, while science-active	25 W	25 W	Mass-independent

Cutting 10 kg saves about 25 W of drive power and almost nothing off the hotel. A light rover is hotel-dominated and a heavy rover is drive-dominated, and that is the whole mass-power relationship in one sentence. The power study's independent 50 kg table lands in the same place: **249 W mean and 146 Wh** over a 35-minute mission, against **123 W while driving and 74 W averaged over a 60-minute Delivery** in this model, and 3.7 kW at the simultaneous peak. The mission clocks are the rule clocks: 20–30 minutes for astrobiology, 30–60 for Delivery, 30 for servicing and 40 for autonomy S-URC27-RULES §1.b.ii–§1.e.i.

MISSION AND CLOCK	35 KG	50 KG
Astrobiology, 30 min	46 Wh	62 Wh
Delivery, 30 min	52 Wh	72 Wh
Delivery, 60 min	104 Wh	**145 Wh**
Equipment Servicing, 30 min	35 Wh	47 Wh
Autonomy, 40 min	76 Wh	106 Wh
Autonomy then Servicing, 70 min	not modelled	**153 Wh**
Nameplate, one 60-minute Delivery	197 Wh	**273 Wh**
Nameplate, the back-to-back pair	not modelled	**289 Wh**

The worst case is the clock, not the mission. Every mission except the 60-minute Delivery fits inside about 106 Wh at any mass, and the rules allow servicing to begin ten minutes after autonomy, so the back-to-back pair totals **153 Wh before margin** S-URC27-RULES §1.d.i. Sizing a pack on a generic 30-to-40-minute mission sizes for about 70 % of the worst case.



Energy drawn against minutes on course for a delivery-type mission at 35 and 50 kg, with the rolling-resistance and efficiency bands, the four 2027 mission windows and the six packs teams actually field. The gap between the 273 Wh requirement and the observed 432–840 Wh is current rating, chemistry and hotel growth.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC27-RULES` `S-UIU-ROVER-PAGES-EARLIER-GE-02`

Full resolution

Peak current is the number that sizes the fuse and often the pack. At 25.6 V, a 30° climb at 50 kg draws 11 A, a single stalled shoulder joint 24 A, and six motors at their controller limits 120 A. The credible simultaneous peak, climb plus one joint plus hotel, is **156 A or about 4.0 kW**; the power study's estimate of everything at once is **145 A and 3.7 kW**.

CASE	W	A
Hill climb, 50 kg, 30°, 0.5 m/s	269	11
One shoulder joint at stall	624	24
Wheel slip at the controller limits	3 072	120
Credible simultaneous peak	3 994	156
Independent estimate, everything at once	3 712	145

The peak is 20 to 30 times the power a rover draws while rolling, and about 15 times the 249 W mission average. The ratio drives four decisions: the parallel count of the pack, a time delay on the battery-management trip so a legitimate peak does not shut the rover down, a software interlock between drive and arm, and a fuse rated above the peak rather than above the average. The protection chain is ordered **controller limits, then the battery-management trip, then the fuse curve, then wire ampacity, then connector rating.** Missouri S&T's 80 A trip at 10 ms would open on a legitimate 156 A peak, which is why their published design restores automatically after 3 s `S-MST-2023-POWER` .

Sizing the pack: 273 to 289 Wh, then choose 400 to 500 Wh

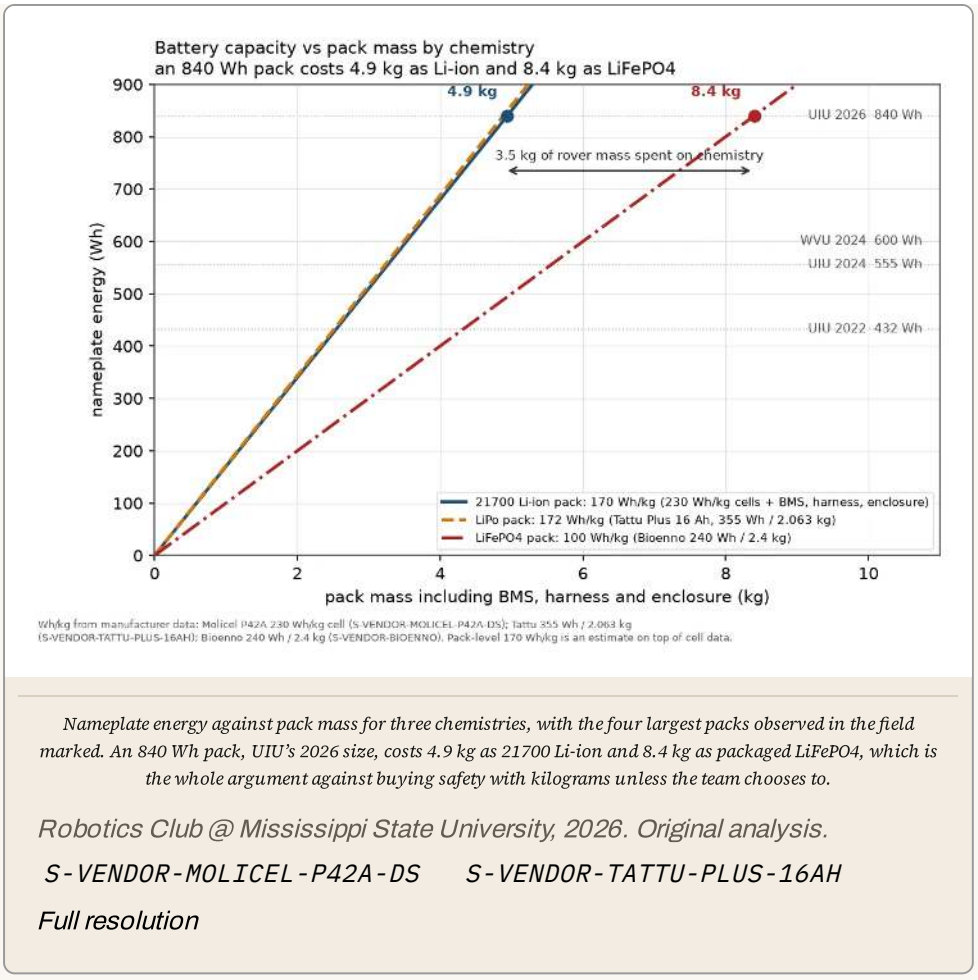
From mission energy to nameplate, the derating chain multiplies four factors, and the result for the 60-minute Delivery at 50 kg is **273 Wh:**

STEP	VALUE	BASIS
Mission energy	145 Wh	60-minute Delivery at 50 kg
Plus 25 % mission margin	181 Wh	Unmodelled hills, radio at maximum, extra arm moves
÷ depth of discharge 0.80	227 Wh	80 % discharge is about 400 cycles against 600 at 60 % S-BU-808
÷ temperature and ageing 0.92	246 Wh	38 °C ambient plus end-of-life capacity loss
÷ path efficiency 0.90	**273 Wh**	Sized at the bus, not at the cells
Combined usable fraction	0.662	0.80 × 0.92 × 0.90

Chemistry changes the derating and the mass in opposite directions. Packaged LiFePO4 accepts a deeper discharge, 0.90 instead of 0.80, so it needs 243 Wh of nameplate, but at 100 Wh/kg instead of 170 Wh/kg it masses 2.4 kg against 1.6 kg for the same mission.

CHEMISTRY	USABLE FRACTION	NAMEPLATE FOR THE SAME MISSION	WH/KG AT PACK LEVEL	PACK MASS
21700 Li-ion (P42A class)	0.662	273 Wh	170	1.6 kg
LiPo pouch	0.662	273 Wh	172	1.6 kg
Packaged LiFePO4	0.745	243 Wh	100	2.4 kg

The recommendation is 400–500 Wh of nameplate at 8S, Li-ion 21700, 2.4–2.9 kg including battery management and enclosure, and two packs rather than one big one. Two 450 Wh packs mass 5.3 kg fielded, one flies while the other charges, and a pack carried as a spare but not fitted does not count against the 50 kg S-URC27-RULES §2.a.iii. The band covers the 289 Wh back-to-back case with 38–73 % headroom and keeps the current rating comfortable, because 4P of that cell is rated 180 A continuous and 8S4P is 496 Wh S-VENDOR-MOLICEL-P42A-DS .



Nameplate energy against pack mass for three chemistries, with the four largest packs observed in the field marked. An 840 Wh pack, UIU's 2026 size, costs 4.9 kg as 21700 Li-ion and 8.4 kg as packaged LiFePO4, which is the whole argument against buying safety with kilograms unless the team chooses to.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-VENDOR-MOLICEL-P42A-DS` `S-VENDOR-TATTU-PLUS-16AH`

Full resolution

Why the packs teams field are 1.6 to 3.1 times the model. Current-driven sizing explains 400–500 Wh without any energy argument. Two missions on one charge adds about 100 Wh. Conservative derating adds 10 %. Unmodelled hotel growth is real: WVU ran eleven cameras, a three-dimensional lidar, three radios and two computers in 2026, and UIU grew from 432 to 840 Wh while mass stayed flat `S-WVU-SAR2026` `S-UIU-ROVER-PAGE-AURION`. LiFePO4 costs 50 to 100 % more watt-hours for the same mass. And nobody can calibrate, because no team publishes a measured figure.

What the model cannot do. Varying the five main assumptions moves the nameplate requirement across **221 to 410 Wh**, and the two that matter are mission duration and mean speed. The energy budget is a driving-distance budget in disguise, so a rover that is 20 % more efficient but drives 30 % further uses more energy.

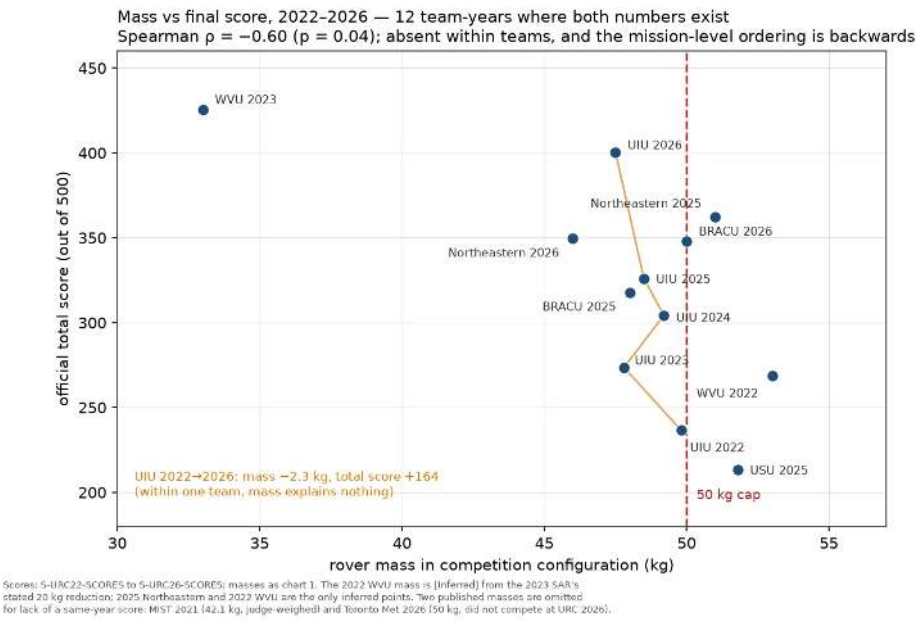
Mass against score: a marker, not a lever

The mass analysis asks whether mass predicts score. The answer is a correlation that fails both causal tests.

TEST	VALUE	READING
Spearman correlation, mass against total, n = 12	**−0.60, p = 0.04**	Statistically real in this sample
Score normalised within year	−0.52, p = 0.08	Partly a year effect in disguise
The two inferred masses dropped, n = 10	−0.72, p = 0.02	Strengthens, so interpolation is not driving it
Leave-one-out range	−0.48 to −0.75	No single point dominates

The mission decomposition is backwards from the physics. The most mass-sensitive mission, Delivery, correlates weakest at **−0.46 (p = 0.13)**, while Equipment Servicing, whose workload barely depends on mass, correlates at **−0.61 (p = 0.03)**; autonomy is **−0.67 (p = 0.02)** and science **−0.12 (p = 0.70)**. That is what a confounder looks like: heavy teams are heavy because the design did not converge, and the same failure costs points everywhere.

Within a team the correlation disappears entirely. UIU moved **−2.3 kg for +164 points**, WVU **−20 kg for +157**, Northeastern **−5 kg for −12** and BRACU **+2 kg for +30** `S-UIU-ROVER-PAGE-AURION` `S-WVU-SAR2023` `S-NEU-NEWS-2026` `S-BRACU-ROVER-PAGE-TAURUS-2026`. Heavier and better, lighter and worse.



Mass against final score for the twelve configurations where both exist. The association is negative and reasonably robust, and it is not causal: the mission ordering is the reverse of the physics and within-team changes show no relationship at all. Mass is a marker of design maturity.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

There is no interior optimum in the observed range. The 33 kg winner and the 42–50 kg cluster both score between 200 and 425. The question for a first build is not how heavy but **how much reserve**: 45 kg with 4 kg unspent beats 49.5 kg with none, because the last 3 kg of a first build arrive in the final six weeks.

Margin policy and the weigh-in plan

STAGE	ALLOWANCE	WHAT IT COVERS
Preliminary design	+20 % on estimated structure and mechanism mass	Parts not yet designed
Critical design, CAD complete	+10 % on unbuilt parts, 0 % on measured	Machining tolerance, fasteners, brackets
Integration, hardware exists	+5 % on the whole rover	Cable, tape, the second E-stop
Weigh-in target	49.5 kg or less on the team's own scale	Scale variance between events, dust, humidity, rounding
Science module reserve	6 kg or more of the 50 kg	A module added after the first weigh-in

The weigh-in plan matters as much as the budget. USU measured **51.8 kg in arm configuration and 46.9 kg in science configuration from the same programme**, a 4.9 kg spread between two legal-looking builds, one of them over `S-USU-2025-POSTER`. That is the failure mode: not a rover that is obviously too heavy, but one whose heaviest configuration is discovered at the field. The defence is to weigh the bare chassis, each subsystem as it is completed, the integrated rover at three points, and the heaviest mission configuration rather than the science one.

Open problems

- **What is Missouri S&T's competition mass?** No public number exists for the 2025 and 2026 champion, after a documented search of the team site, 224 public code projects, four SAR videos and twelve papers `S-MST-GH-ORG` . It is the highest-value missing datum in the field.
- **What is the real rolling resistance at MDRS?** The whole power model rests on an assumed 0.12 for a six-wheel skid-steer on sand, gravel and rock. A tow test at 40 kg would close it in an afternoon.
- **How many watt-hours does a mission actually take?** No team publishes a measured figure, so every energy number here is a model with a stated basis.
- **Was the 51.8 kg Utah State configuration penalised?** The official sheet shows no penalty, and whether the penalty column is complete is not published.
- **Does a drone count toward the 70 kg fielded total?** The rules exclude it from the rover's own mass and size limits and are silent on the union `S-URC27-RULES` `§2.b.vii`.

Further reading

- Interfaces and couplings works the arm-mass and wheel-diameter couplings that decide this budget.
- How rovers fail explains why the pack's sag and the arm's zero cost more than the mass penalty.
- Mass and power data carries the full appendix behind the table on this page.
- Systems engineering and power hold the subsystem-level budgets.
- Missouri S&T and UIU are the two dossiers with the most mass evidence.

Interfaces and couplings

Sixty-six interfaces exist between the twelve subsystems, and every pair has at least one. Thirty-seven of the 66 are high criticality, 25 medium and 4 low. That saturation is itself the finding: at this level of integration there is no isolated subsystem change, and the order in which a team freezes its numbers matters more than the order in which it designs its parts.

66
INTERFACES BETWEEN TWELVE SUBSYSTEMS
37 high, 25 medium, 4 low
2027 design

37
HIGH-CRITICALITY INTERFACES
Each needs an owner, a line in the interface document and a test
2027 design

9
HIGH INTERFACES ON THE ARM
The most coupled mechanical subsystem
2027 design

9
HIGH INTERFACES ON PERCEPTION
The most coupled data subsystem
2027 design

Key takeaways

- **The arm and perception each touch nine high-criticality interfaces**, the most of the twelve. A team that plans the arm last and the cameras as the design proceeds has chosen the two worst places to defer decisions.
- **A change to the arm's length, mass or stow pose re-opens nine interfaces; a wheel-diameter change re-opens seven** and is irreversible once the wheels and gear ratios are made.
- **Delivery is a mission, not an interface hub.** It has one high interface, the arm's 5 kg grasp, because it reuses everyone else's interfaces and its risk sits in the mission sequence.
- **The arm's mass silently chooses the drivetrain.** A 4.4, 9.5 or 14 kg arm puts the rover at 30.6, 35.7 or 40.3 kg, and the 9.0 N·m actuator class runs out at 40.6 kg, so one more bracket can force a 2.7 kg actuator step.

Twelve nodes, five kinds of interface

Each subsystem is a node that trades in a small set of numbers. Every interface is one of five types, because the five fail in different ways and are owned by different engineers.

CODE	TYPE	WHAT CROSSES IT	TYPICAL FAILURE WHEN IT IS UNMANAGED
M	mechanical	Mounting patterns, envelopes, mass, centre of gravity, stiffness, retention	Interference found at assembly; a stow that does not fit at weigh-in
E	electrical	Rails, currents, fusing, grounds, inrush, the scope of the E-stop	Brownout, a welded contactor, an E-stop that does not cut everything
D	data	Protocols, rates, schemas, timestamps, calibration files	Un-synchronised odometry; a camera that resets when a motor stalls
T	thermal	Conduction paths, dissipation, ambient, solar gain, derating	The sealed-box oven that works in the lab and fails at the site
t	timing	Clocks, sequences, setup and teardown, mission windows	A correct rover that runs out of mission clock

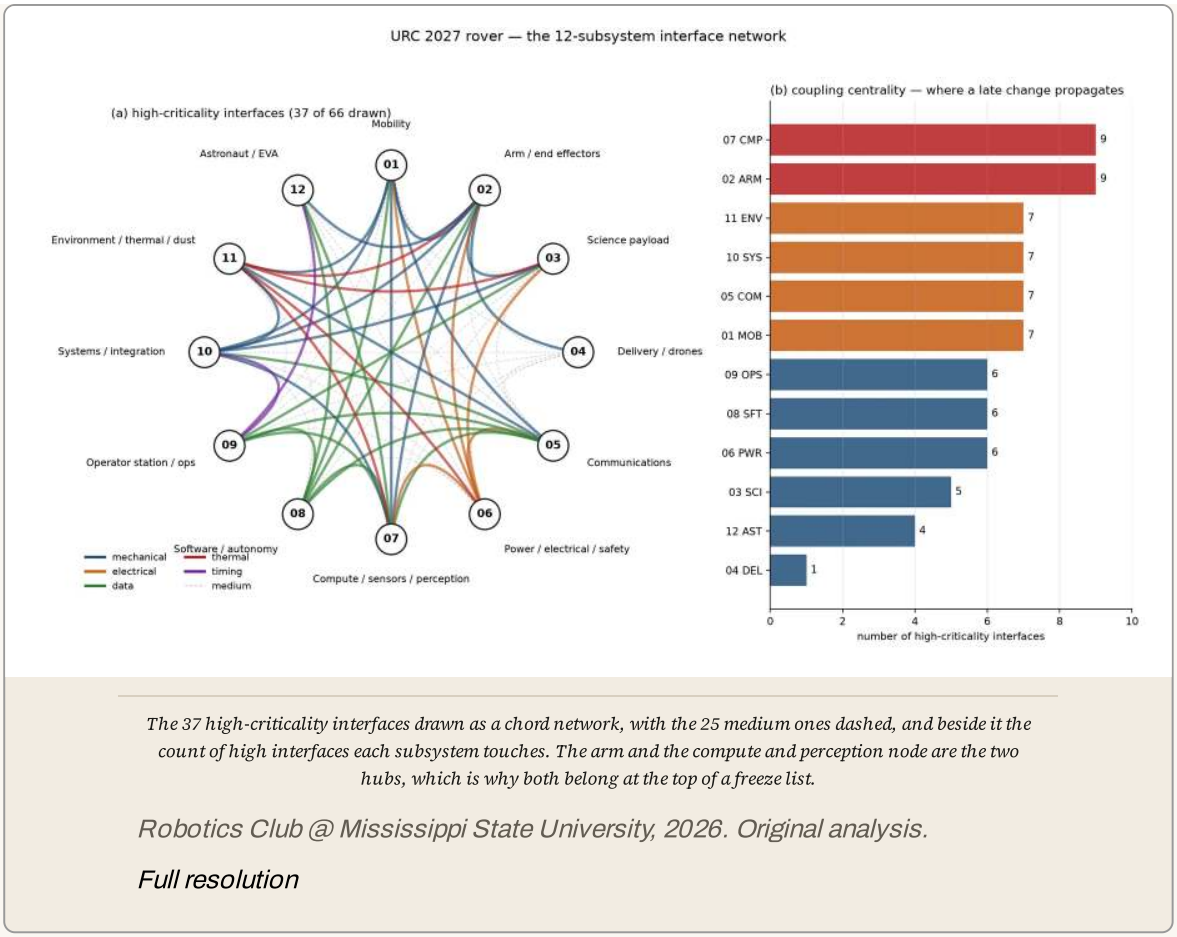
Timing is an interface type rather than a management afterthought because the competition scores it. The astrobiology clock is **20–30 min**, the autonomy-to-servicing transition can be as short as **10 min**, and setup is **15 min including the weigh-in and all communications** S-URC27-RULES S-URC27-QA §1.a.ii–§1.e.i EQ4. A design that is two minutes too slow loses the same points as a design that does not work.

The 12 × 12 matrix

Each cell is the interface as seen from the subsystem in that row. The letters give the interface type, in descending order of dominance, followed by its criticality: **H** high, **M** medium, **L** low. The upper triangle is the authoritative entry and the lower triangle repeats it so the table can be read in either direction. The diagonal is a subsystem's own business and is marked n/a.

ROW, COLUMN	01	02	03	04	05	06	07
01 MOB	n/a	M+E·H	M+t·M	M+D·M	M+E+D·H	E+T·H	M+D
02 ARM	M+E·H	n/a	M+D·H	M·H	D+t·M	E·H	M+D
03 SCI	M+t·M	M+D·H	n/a	t+M·M	D·M	E·H	D·M
04 DEL	M+D·M	M·H	t+M·M	n/a	D+t·M	E·M	D·M
05 COM	M+E+D·H	D+t·M	D·M	D+t·M	n/a	E·H	D·H
06 PWR	E+T·H	E·H	E·H	E·M	E·H	n/a	E·H
07 CMP	M+D·H	M+D·H	D·M	D·M	D·H	E·H	n/a
08 SFT	D+t·H	D+t·H	D+t·M	D·L	D+t·H	E+D·M	D+t·
09 OPS	D+t·M	D+t·M	D+t·H	t·M	D+t·H	E·M	D·H
10 SYS	M·H	M·H	M·H	M+t·M	D·H	M+E·M	M·H
11 ENV	M+T·H	T+M·H	T·H	M·L	M+T·H	T·H	T+M
12 AST	M+t·M	M·H	t·L	t·L	D·M	E·M	D·H

Five readings follow from the matrix. The two most coupled subsystems are the arm and perception at nine high interfaces each, one mechanical and one data-dominated. Seven subsystems sit at six or seven, so there is no clean partition into core and periphery. Delivery has a single high interface, which is evidence that its difficulty is in the mission sequence rather than in a novel coupling. The astronaut node is data and timing dominated, with the arm as its only high mechanical interface. Every power interface carries an electrical term, and the only one without it is the thermal link to the environment node, because the E-stop rule makes one subsystem carry a system function `S-URC27-RULES` §2.a.v.



The 37 high-criticality interfaces drawn as a chord network, with the 25 medium ones dashed, and beside it the count of high interfaces each subsystem touches. The arm and the compute and perception node are the two hubs, which is why both belong at the top of a freeze list.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

RANK	SUBSYSTEM	HIGH INTERFACES	WHAT IT MEANS FOR THE SCHEDULE
1–2	Arm (02), compute and perception (07)	9	Freeze their mounting, power and data contracts first; both are long-lead
3–6	Mobility (01), comms (05), systems (10), environment (11)	7	These host other subsystems' mass, masts and heat, so a change invalidates several interface documents at once
7–9	Power (06), software (08), operations (09)	6	Broad but shallow: each interface is a number or a schema that can change independently
10	Science (03)	5	Contained as a bay, which is exactly why its internal integration of fluids, heat and time is where it fails
11	Astronaut interaction (12)	4	Mostly data; the suit is allowed to be a separate system
12	Delivery (04)	1	A mission, not an interface hub

Eight interfaces with a number that forces a decision

INTERFACE	WHAT CROSSES THE BOUNDARY	RISK IF IT IS IGNORED
Arm to mobility	Shoulder reaction **176 N·m peak**; an arm from 4.4 to 14 kg raises the centre of gravity from **0.222 to 0.282 m** and cuts the static tip angle from **64.9° to 59.3°** on a 0.95 m track	An arm mounted on a suspension link moves relative to the wheels; the deck must be a structural node
Wheel diameter to stow	A 250 mm wheel on a 0.90 m wheelbase gives **1.15 m**; 330 mm gives **1.23 m** and fails; 510 mm gives **1.41 m**	Decides whether a fold mechanism exists at all, with a 40 % penalty if the cube fails
Camera mast to structure	A 30 mm carbon tube, 0.8 m, 0.5 kg tip, has its first mode at **19 Hz**; a 20 mm tube gives **9 Hz**, inside the 2–15 Hz terrain band	A flexible post destroys stereo calibration and visual odometry
Mast loads to mobility	A 65 km/h gust gives **155 Pa**; **23 N·m** for a 0.15 m² antenna at 1 m, **78 N·m** for a 0.5 m² dish at 3 m	A mast load path through the science box cracks the box
Power to environment	40 W in a 0.52 m² box: **195 W** of solar absorbed by a black skin gives **77 °C** internal, white 55 °C, shade 50 °C, a filtered fan **42 °C**	The sealed-box oven is the classic works-in-the-lab failure
Power to arm	**28 A peak per joint**; about 1.1 kW for two joints moving, and hold current unless brakes are fitted	A stalled arm is the usual cause of a battery-management trip, at 561 W per joint continuously
Perception to software	ArUco id, range and bearing at **10–30 Hz**, range error 0.3 % of distance at 60 pixels or more; command events carrying confidence and method	A six-degree-of-freedom pose from a 30-pixel tag is worse than the bearing, and the method field is what scores
Systems to operations	**15 min** of setup including weigh-in and all communications; a six-minute mast leaves nine minutes for everything else	Every subsystem's setup ritual competes for the same clock

The full register is 41 rows across the five types: ten mechanical and stow, seven electrical, ten data, six thermal and eight timing. Every high cell in the matrix is represented by at least one row. The table below searches and filters by interface type.

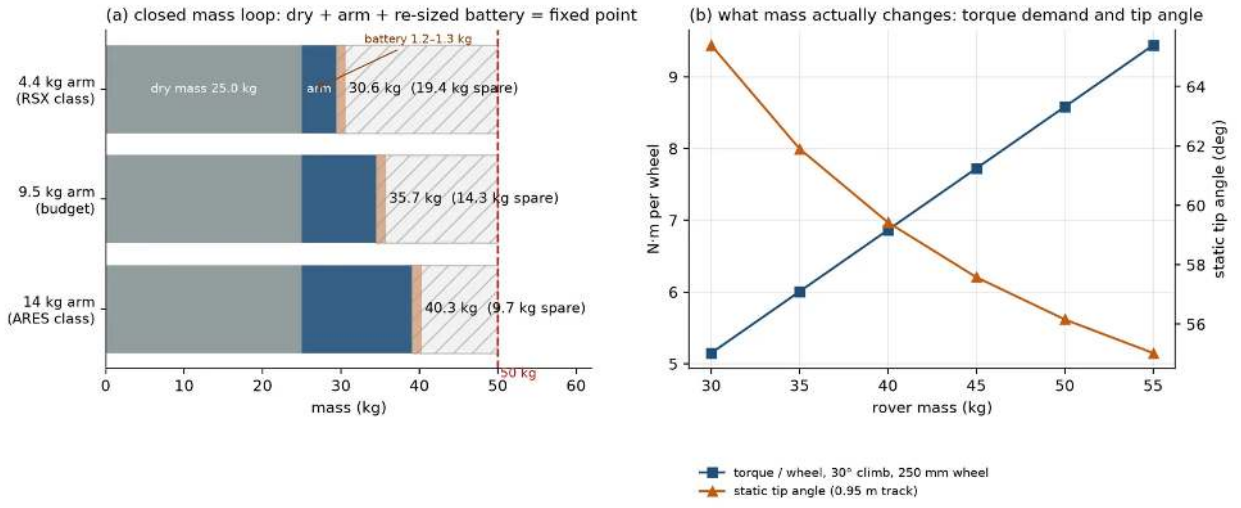
ID	FROM → TO	TYPE	WHAT IS EXCHANGED	OWNER	RISK
IF-001	02 → 01	M	Deck/nose hardpoint with a known transform to the frame; shoulder reaction ≤176 N·m peak, 233 N·m static; arm 4.4–14 kg carried at the deck	02	H, an arm on a suspension link moves relative be a structural node
IF-002	02 → 01	M	CG shift: arm 4.4→14 kg raises h_CG 0.222→0.282 m and cuts static tip angle 64.9→59.3° (0.95 m track)	01	H, unloads the front axle on a climb; §4.1
IF-003	02 → 10	M	Stow envelope: arm parked along the deck within 0.12–0.15 m of vertical space; folded footprint max(a,b) ≈0.55 m; a 200 mm tool must rotate parallel to the forearm	10	H, the 1.2 m cube bites vertically, not horizon fit
IF-004	01 → 10	M	Wheel diameter vs box: 250 mm wheel on a 0.90 m wheelbase gives 1.15 m overall length; 330 mm gives 1.23 m (fails); 510 mm gives 1.41 m	01	H, decides whether a fold mechanism exists
IF-005	05 → 01	M	Mast 0.8–3.0 m with antennas and gimbal on the chassis; 65 km/h gust → 155 Pa → 23 N·m for a 0.15 m² antenna at 1 m, 233 N·m for a 0.5 m² dish at 3 m	01	H, a mast load path through the science box i box
IF-006	07 → 10	M	Camera/GNSS mast or post: 30 mm CF tube, 0.8 m, 0.5 kg tip → 19 Hz first mode; 5 N tip load → 0.70 mm static	07	H, a 20 mm tube at 9 Hz sits inside the 2–15 H stereo calibration
IF-007	03 → 01	M	Drill downforce ≥10 N at full reach reacts into the chassis; rover must hold position on a slope while drilling	03	H, reaction torque path is the classic arm/drill
IF-008	12 → 10	M	Suit and wearable electronics are not rover mass [AUQ12]; they are fielded equipment for the budget only	12	M, accounting treatment, not a physical interf
IF-009	04 → 10	M	Drone ≤5 kg including an inert dummy mass equal to the battery; excluded from the 50 kg / 1.2 m limits; treatment inside the 70 kg total is	10	M, a spare drone battery left on the rover is a
IF-010	10 → 11	M	Crate = the 1.2 m weigh-in box; 10 g transport shock sends 4.9 kN into four anchors (1.2 kN each); crate as a sail at 100 km/h → 516 N	11	H, transport damage is the best-documented
IF-011	06 → 01	E	Drive bus 24 V (48 V emerging); 120 W average rolling, 371 W at 45 kg up 30°, 120 A peak at current limit, regen	06	H, 6S vs 8S changes harness current by 1.33 can over-volt a full pack
IF-012	02 → 06	E	Arm rail 24–48 V; 28 A peak per AK10-9 joint; ~1.1 kW for two joints moving; hold current unless brakes fitted	06	H, arm stall is the usual cause of a BMS trip (5
IF-013	03 → 06	E	Isolated 12 V science rail: 30–100 W drilling, 10–20 W heater, pumps ~200 mA each; ~5–10 Wh per mission	06	H, a pump motor on a shared rail injects noise cameras
IF-014	07 → 06	E	5 V (cameras), 12 V (Jetson, lights), 19 V (x86); 20 W average / 40 W peak compute; ripple <100 mV, brownout <10 ms	06	H, motor stall on a shared bus resets USB ca
IF-015	06 → 09	E	C2 equipment on 120 V 60 Hz; UPS buffer; measured C2 current per configuration	09	M, no URC current limit and no UPS provided
IF-016	06 → 05	E	Radio rail 12 V, 9–15 W (0.8 A at 5 V for an RFD900x class); the E-stop kills the radio and the link	05	H, the judge must rely on the rover LED after
IF-017	06 → 08	E+D	LED status array driver powered and driven independently of the compute, plus an "E-stop latched" state bit	06	M, if the LED driver dies the judge cannot see
IF-018	07 → 08	D	ArUco id/range/bearing/covariance at 10–30 Hz (range error 0.3 % of Z at ≥60 px); costmap 5–10 Hz at 0.1–0.2 m; person track 10–30 Hz; command events with confidence and method	07	H, a 6-DoF pose from a 30 px tag is worse tha is what scores 15 points
IF-019	08 → 05	D	Telemetry ≥2 Hz, ≤2 kB/packet (mode, pose, battery, task, faults); video H.264/H.265 720p30 at 2–4 Mbit/s adaptive	08	H, DDS multicast over a lossy radio collapses
IF-020	05 → 09	D	3–6 concurrent streams at ≥720p30, glass-to-glass 100–300 ms; RSSI/noise/latency continuous; one-click bitrate drop	09	H, without link-quality display the operator ca
IF-021	02 → 07	D	Wrist camera 720p–4K at 30 Hz; MIPI/USB routed through 6 joints; 1 cm tag resolvable at ≤0.45 m (720p/70°) or ≤1.4 m (4K/40°)	07	H, a USB cable through a 6-DoF arm is a fatig points to a damaged USB-C cable)
IF-022	07 → 02	D	Eye-in-hand pose of the 1 cm key tag and 2 cm keyboard tags: 30 Hz, 1–3 mm relative, ≤5 mm absolute, at 0.2–0.5 m standoff	07	H, sets whether typing is possible at all; §4.2

IF-023	01 → 08	D	Wheel odometry ≥50 Hz on the same clock as the IMU; GNSS PPS → PTP/gPTP, ≤5 ms camera/LiDAR, ≤1 ms IMU; tilt/IMU for slope limiting	01	H, un-synchronised odometry halves the valu
IF-024	02 → 08	D	Joint position/velocity/current at 100–200 Hz on CAN 1 Mbit/s; Cartesian mode ≤0.1 m/s for grasp; per-joint limits shared with IK	02	H, IK and limits must be one implementation,
IF-025	08 → 09	D	One autonomy-state enum, last-telemetry age, map, target list, stay-out editor; hands-off lockout visible	08	H, if state is inferred from behaviour the hand argument
IF-026	07 → 05	D	RTCM3 corrections 0.5–2 kB/s; the base station is the team's own because there is no CORS/internet	07	M, a 30–60 min survey-in fights the 15-min se
IF-027	12 → 05	D	Suit camera + microphone into a C2 display (5 points); two-way link permitted; astronaut may carry but not deploy a repeater	12	M, rover-relayed path dies beyond ~20–30 m, geometry
IF-028	06 → 11	T	40 W internal dissipation in a 0.4×0.3 m box (0.52 m²): 195 W solar absorbed black → 77 °C; white → 55 °C; white + shade → 50 °C; + filtered fan → 42 °C	11	H, the sealed-box oven is the most common '
IF-029	02 → 11	T	Stall in a joint: 20 A through 0.35 Ω = 140 W in the winding; 0.28 K/s initial, 70 K equilibrium, ~4–5 min to +60 K	11	H, a stalled hold is survivable for a minute and
IF-030	03 → 11	T	Reagent storage ≤38 °C; a 65 °C heater for LAMP-class assays; venting a heated bay fights sealing it against dust	03	H, sealing against dust and venting a heater a
IF-031	01 → 11	T	~9 W per motor at 371 W electrical over six motors; 50 °C actuator ambient limit; sealed bearings add drag	01	M, thermal risk is stall, not the 30° climb
IF-032	12 → 11	T	37.8 °C air, ~33 °C June daily max, 12-minute planned exposure per attempt, water every 15 min	12	M, heat is a medical risk, not comfort
IF-033	07 → 11	T	15 W Jetson in a 40 W box; throttling at minute 12 of an autonomy run looks like a software bug	07	H, an "autonomy failure" that is really a therm mission
IF-034	10 → 09	t	15 min setup including weigh-in and all comms; a 6-minute mast leaves 9 minutes for everything else; teardown 10 min, retrieval 10 min	09	H, the clock is shared: every subsystem's se
IF-035	10 → 04	t	Autonomy→Equipment transition as short as 10 min, including a possible re-weigh (EQ4); payload swap must fit inside it	10	H, modularity is only legal if the swap fits the '
IF-036	03 → 09	t	Astrobiology 25-min plan inside a 20–30 min window; sample acquisition + documentation is 10.5 min of the 25; cache out within 5 min	09	H, a failed drill costs the whole mission, not o
IF-037	04 → 09	t	Drone pilot + dedicated visual spotter from C2; drone must land before the rover proceeds in Autonomy; battery swap by hand = intervention	09	M, the drone buys time but consumes a seat
IF-038	12 → 09	t	Five ordered commands; quiet required only for gesture Follow!; ≤3 m stand-off, ≥20 m stay test; C2 stop costs 20 % of the task	12	H, a human is inside the rover's motion envel
IF-039	01 → 09	t	Teleop at 0.5–2.0 m/s over up to 1 km; 70–230 ms latency budget; 1 km at 0.5 m/s = 33 min of a 30–60 min mission	09	H, link quality converts directly into stage tim
IF-040	11 → 09	t	Any team touch = 20 % penalty per intervention (additive); repair must fit inside a mission window or the mission is lost	11	M, repair is a scheduled activity, not an accid
IF-041	09 → 10	t	6 students in C2 at any time; 24 registered, 2 advisors, 14 people per day; a runner who leaves C2 cannot operate that mission	09	H, swapping a payload can consume the next

The three chains that carry the most consequence

CHAIN	FIRST STEP	SECOND STEP	THIRD STEP	FOURTH STEP
1. Arm mass sets the drivetrain	Arm 4.4, 9.5 or 14 kg	Rover 30.6, 35.7 or 40.3 kg	Cliff at 40.6 kg against a 9.0 N·m actuator	Next actuator class at +2.7 kg for six, 43.0 kg fixed point
2. Wheel diameter sets the box and the centre of gravity	250 mm: 1.15 m, 61.5°	330 mm: 1.23 m, fails the cube	510 mm: 1.41 m, 48.6°, 25 kg of printed wheels	Torque at a 30° climb grows from 6.9 to 14.0 N·m
3. Camera count sets the link and the clock	6 × 3 Mbit/s = 18 Mbit/s	2.4 GHz: 6 Mbit/s at 2.62 km; 900 MHz: 0.2 Mbit/s at 1 km	0.5 m/s teleoperation = 33 min for 1 km	A 30-minute stage runs out of clock

Chain 1: arm mass to drivetrain to points. The mass loop is stable and converges in two iterations with a loop gain of **0.0137 kg of battery per kilogram of rover**, so forgetting to re-size the pack after adding the arm is wrong by 0.16 kg. The danger is two cliffs. At 250 mm wheels, a 45° loose slope and 45 kg requires **9.97 N·m per wheel**, and an AK80-9 is rated **9.0 N·m** `S-CUBEMARS-AK80-9`. Inverting the formula gives **40.6 kg**. The 14 kg ARES-class arm lands at **40.3 kg, 0.3 kg below the cliff**, so one more bracket forces a 0.94 kg actuator class at **+0.45 kg per wheel or +2.7 kg for six**, re-opening the budget at **43.0 kg** with 7.0 kg of headroom. The arm is not 4.5 kg lighter than the alternative; it is 4.5 kg plus the actuator class it preserves. No published team process owns the margin-to-cliff number.

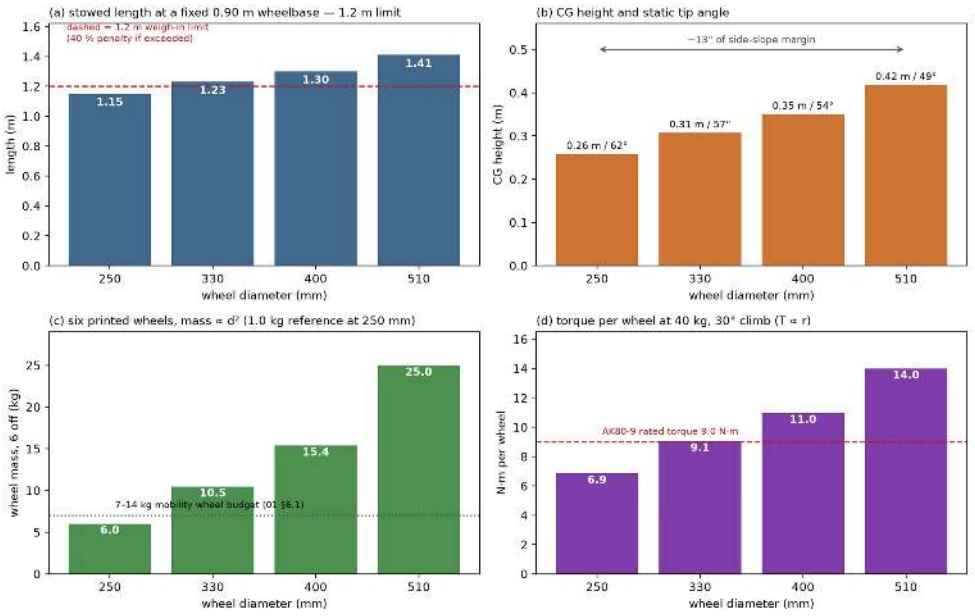


A 25.0 kg dry mass plus the arm plus a re-sized pack, closed for three arm classes, with the headroom to the cap on the left; on the right, what mass actually changes, which is wheel torque on a climb and static tip angle. The 14 kg arm leaves 9.7 kg of headroom against the cap.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

Chain 2: wheel diameter is a four-way, irreversible coupling. It changes climbing, centre-of-gravity height, stow length and torque in conflicting directions. From 250 to 510 mm costs **12.9°** of static tip angle, against 5.6° for an arm change from 4.4 to 14 kg. Six printed wheels grow from **6.0 to 25.0 kg**, torque at 40 kg on a 30° slope grows from **6.9 to 14.0 N·m**, and overall length grows from 1.15 to 1.41 m. Only top speed improves, from **5.1 to 10.4 m/s** at rated motor speed, against a design target of 0.5–2.0 m/s.



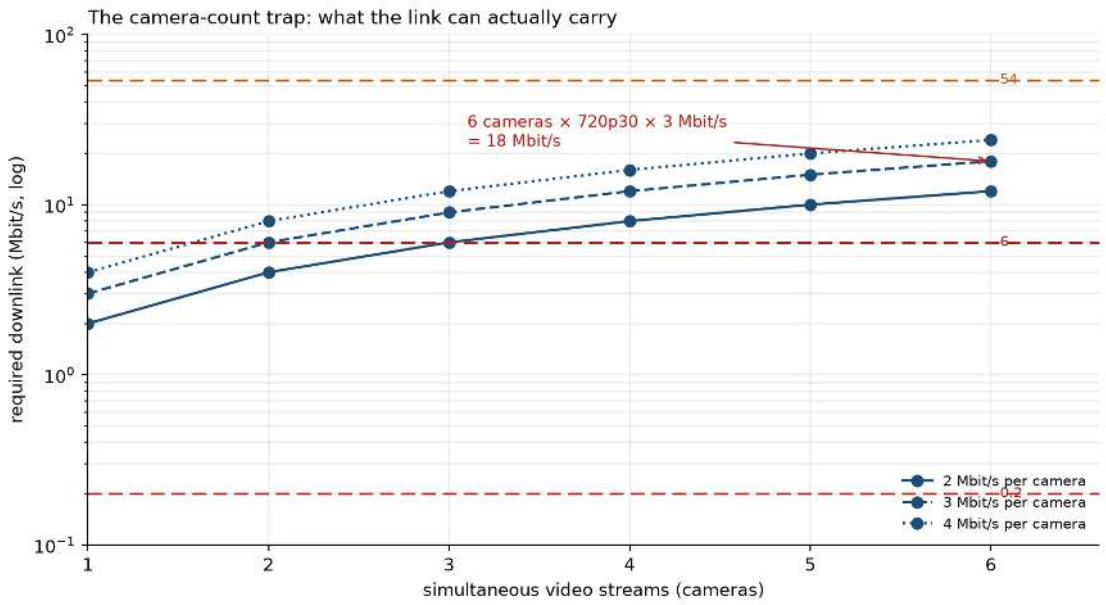
One panel per consequence of wheel diameter at a fixed 0.90 m wheelbase: stowed length against the cube, centre-of-gravity height and tip angle, printed wheel mass, and torque per wheel on a 30° climb. Only top speed improves with diameter, and top speed is the requirement no team is short of.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-CUBEMARS-AK80-9` `S-URC27-RULES` *Full resolution*

The field evidence agrees. WVU ran 20-inch wheels in 2024, 24-inch in 2025 and **back to 20-inch with no folding mechanism in 2026**, giving up 2 inches of climbing radius to fit the cube without a fold `S-WVU-SAR2025` `S-WVU-SAR2026`. Missouri S&T moved to a suspension that no longer needs to fold to fit the 1.2 m box `S-MST-SITE-ATHENA`. Two of the strongest teams independently concluded that geometry beats mechanisms. The defensible point for a first build is **250–300 mm with a 0.85–0.95 m wheelbase and no fold**, which leaves 40–50 mm per side for suspension links and tolerance.

Chain 3: camera count is a link-budget decision that converts into stage time. Six cameras at 720p30 and 3 Mbit/s each demand **18 Mbit/s**. The 900 MHz link carries **200 kbps at 1 km**, which is **1.1 %** of the demand; 2.4 GHz carries 6 Mbit/s at 2.62 km and 54 Mbit/s at 0.42 km; 5.8 GHz with the radio at the mast carries 6 Mbit/s at 21.4 km. Dropping from the fastest modulation to 6 Mbit/s buys about **9 times the range**. At 1 km and 0.5 m/s the Delivery traverse takes **33 minutes**, which does not fit a 30-minute stage, and WVU's 2026 tri-band system was tested to only **750 m line of sight and 164 m obstructed** `S-WVU-SAR2026`. The camera count must be frozen against the link: three or four streams, adaptive bitrate, and a separate stills path.



Required downlink for two, three and four simultaneous streams at 2, 3 and 4 Mbit/s per camera, against the measured capacity of the three bands. At the 6 Mbit/s long-range operating point only two streams fit, while six streams need 18 Mbit/s available only close in.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

Two shorter chains. Arm reach against stow and typing is not a horizontal problem: a 1.05 m arm folds to 0.55 m and fits with 0.65 m to spare, but the parked vertical budget is only **0.12–0.15 m**, and the fold joint that solves it adds backlash at the worst lever. At 4.4–5.2 mrad, the shoulder's 1.05 m lever gives **4.6–5.5 mm** of lost motion at the tip, four joints sum to **9.7–11.4 mm**, and the typing target is **±3 mm** against a 5.0 mm geometric limit. The resolution is a wrist camera and visual servoing, which turns the shoulder term into 1.8 mm at a 0.35 m lever, not a low-backlash reducer at every joint.

Mast height against stability and the cube is the other. The stowed stack already uses **0.84 m** of the 1.2 m envelope before the mast, so a 2.0 m mast leaves **0.06 m** and a 2.5 m mast fails. The same 30 mm tube has its first mode at **19 Hz at 0.8 m**, **4.8 Hz at 2.0 m** and **2.6 Hz at 3.0 m**, inside the 2–15 Hz terrain band. A 20 mm tube is at 9 Hz at 0.8 m. The resolution is **two masts**: a tall guyed radio mast and a short stiff camera post, the only fold a first team should keep.

Conflicts that no subsystem can settle

Fourteen decisions cannot be made inside one subsystem, because each has an owner on both sides. The table gives the two sides and the review by which the answer is needed.

ID	CONFLICT	THE TWO SIDES, WITH NUMBERS	DECIDE BY
CF-01	Arm mass against drivetrain margin	A 9.5 kg arm gives 35.7 kg, 4.9 kg below the 40.6 kg cliff; a 14 kg arm gives 40.3 kg, 0.3 kg below	Preliminary design
CF-02	Reach against stow against typing precision	1.5 m reach, 0.12–0.15 m parked height, ±3 mm typing	Preliminary design, then SAR
CF-03	Wheel diameter against cube, centre of gravity and torque	250 mm fits and is light; 510 mm climbs but costs 12.9° of tip angle and 19 kg of wheel mass	Preliminary design, irreversible
CF-04	Mast height against camera stability and the cube	3 m helps the link; 3 m is 2.6 Hz and 1.29 m stowed	Preliminary design
CF-05	Camera count against link and operator load	Six streams demand 18 Mbit/s against 6 Mbit/s at range, and an operator watches three to six	Preliminary design
CF-06	Compute power against thermal and volume	A 25 W computer in a 77 °C sealed box against a separate battery bay	Preliminary design and SAR
CF-07	Bioassay against arm stiffness and the 25-minute clock	Wet chemistry, a 10.5-minute acquisition block, a mandatory assay	December 2026, before the SAR video
CF-08	Modularity against swap time and staffing	20 kg of second configuration against a 10-minute window and six seats in the control station	Preliminary design
CF-09	E-stop scope against comms and logging	One button must cease all power draw, which also kills the link	Preliminary design, verified at SAR
CF-10	Drone accounting	5 kg including a dummy is excluded from the rover's limits; the treatment in the 70 kg total is not stated	Preliminary design, by a question to the Director
CF-11	Whether the servicing cache is the astrobiology cache	A reusable cache against two designs; the rulebook leaves it open	December 2026
CF-12	Astronaut link topology	A rover-relayed path dies beyond about 20–30 m, which is the stay-test geometry, against a second receiver at the control station	Preliminary design
CF-13	Base-station survey against the setup clock	A 30–60 minute survey against 15 minutes of total setup	SAR
CF-14	900 MHz sub-band agility against a multi-link network	8 MHz in one of three sub-bands for all links, including drone and suit	2027-04-30, the licence date

The freeze order follows the penalty, not the subsystem number

The interfaces whose consequence is a rule penalty are frozen first; the rest are a re-test. A worse but frozen choice beats a better but late one, which is also the judges' own advice on design stability `S-URC-ADVICE`.

INTERFACE	FREEZE BY PRELIMINARY DESIGN	FREEZE BY SAR	CONSEQUENCE OF A LATE CHANGE
Mass allocation and the 50 and 70 kg plan	Allocation, 15 % growth, 1.5 kg reserve	Weighed values replace estimates	5 % of points per kilogram per event, and a re-allocation re-opens every mounting interface
Stow concept: fold or geometry	Concept and vertical stack, with the 0.06 m margin identified	CAD envelope check into a timber mock box	40 % penalty at every weigh-in; a new fold is a new failure mode with no test time
Wheel diameter and wheelbase	Fixed, because tooling and gear ratio depend on it	Wheel mass and bearing choice	Seven interfaces re-open, and the 250 to 330 mm step adds 2.2 N·m per wheel at 40 kg
Arm link lengths, mount and end effector	Reach class and mount location	Link lengths, end-effector change, wrist camera	Chassis hardpoint, stow pose, inverse kinematics and typing strategy all change together
Bus voltage and connector families	The 12, 24 or 48 V decision	Connector part numbers	The E-stop rating, fusing, every harness and every board input
E-stop topology	Single contactor loop with the status driver independent of compute	Test evidence in the SAR video	A pass or fail rule breach, with no partial credit
Comms bands, antennas and licence documents	Bands, antennas, base station	Range and throughput at realistic geometry	The licence date is 2027-04-30, and a band change invalidates the link budget
Camera count and video pipeline	Count, codec, bitrate ladder	Operator display and adaptive-bitrate behaviour	The station layout, link budget, thermal load and power all move
Science assay pair and fluid bay	Assay pair chosen in December	Fluid path, heater, waste containment, cache interface	A post-SAR change is unscored in the review and unproven in the field
Autonomy schema: state, telemetry	Schema version one	Version two with logging and replay proven	Every operator screen and every log depends on it

Open problems

- **No team publishes a mass, power or centre-of-gravity budget**, so every mass-distribution number in the chains is an engineering estimate with its stated basis.
- **The weigh-in duration is not published**, which is one reason not to build a mission plan around a fast re-weigh.
- **Teleoperation ground speed and glass-to-glass latency are not measured by any team**; both are inputs to the camera-chain arithmetic and both need a field test.
- **Whether a drone counts toward the 70 kg total, and whether a spare drone battery falls under the one E-stop rule**, is not stated in the rules and changes a design `S-URC27-RULES §2.b.vii`.
- **Whether the Equipment Servicing cache is the team's own astrobiology cache** is open, and it changes both the cache and the end-effector design.

Further reading

- Mass and power carries the mass and pack numbers these chains consume.
- How rovers fail shows which of these interfaces actually breaks on the field.
- Coupling and trades works the fourteen conflicts as trade studies.
- Systems engineering owns the allocation and the review gates.

CROSS-CUTTING ANALYSIS · CHAPTER 5.3

How rovers fail

The failure record holds **54 sourced entries** from 15 teams plus the judges' field-wide statements, set against 162 finals team-years. Forty-eight entries name a mechanism or a symptom inside a known subsystem and six are zeros whose cause was never published. The record does not say that comms is the problem in the way the judges describe it. It says that the largest and most persistent loss in the field is an autonomy capability gap, and that the most expensive single failures are the ones that take several missions down at once.

54

SOURCED FAILURE ENTRIES

48 classifiable, 6 outcome-only

2015–2026

48

ENTRIES WITH A NAMED MECHANISM

A physical mechanism or a named symptom inside a known subsystem

2015–2026

41 %

AUTONOMY ZERO RATE

67 of 162 finals team-years

2022–2026

`S-URC-RESULTS`

~945

MISSION POINTS FORGONE EACH YEAR

About 610 autonomy, 196 servicing, 140 delivery

2022–2026

Key takeaways

- **Autonomy is the largest points sink in the field.** Sixty-seven of 162 finals entries scored exactly zero on autonomy, 41 % of the sample, worth about **610 of the roughly 945 mission points** the field forgoes each year. It is a capability gap, not a broken rover: 58 % of those zeros were the team's only zero.
- **A large share of the rest is one broken rover.** Equipment and Delivery zeros coincide in **12 of 162 team-years**, 2.9 times the independent expectation, and only 11 % of Equipment zeros were a team's only zero.
- **The counts measure publication, not occurrence.** The arm and comms tie at seven entries each, while the judges call comms "far and away the most common problem" `S-URC-ADVICE` . Both statements are true about different things.
- **The emptiest rows are the most dangerous.** Power has one sourced electrical failure in five years and dust has none, yet one electrical fault can end a whole event. Those rows record unquantified risk, not proven absence of risk.

The database and its three biases

Every count on this page is a count of write-ups. Three biases shape the record and each is worth stating before a number is used.

Survivor bias in team reports. System acceptance videos are a qualification deliverable, so they show successes. Teams explain a failure mainly when it drove a redesign; a 2023 failure that ended a season and was never written up leaves no trace. The competition publishes no did-not-finish cause codes and no judging sheets.

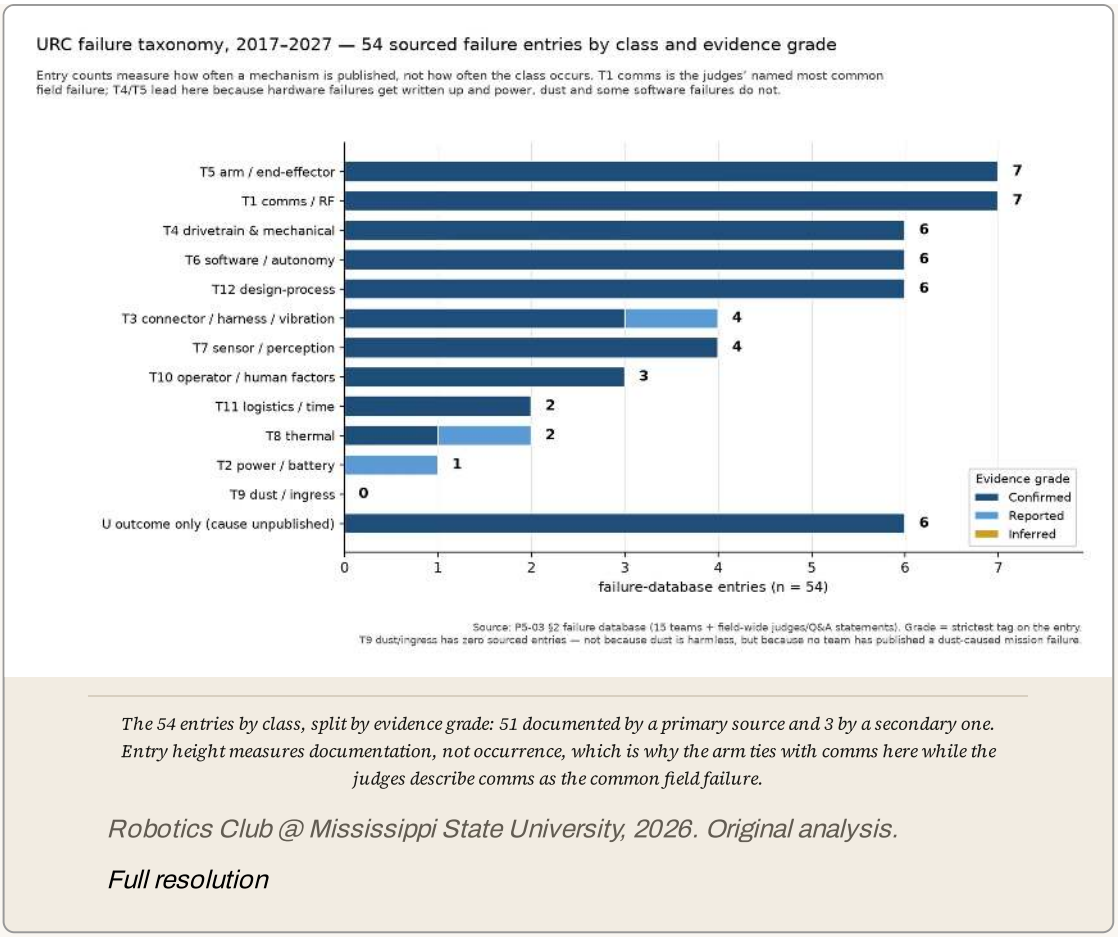
Attribution bias in news releases. A release names the incident but rarely the mechanism. A "communication loss" is a symptom, not a cause, and where the team itself gave the mechanism, as WVU did for a damaged cable in 2026, that is the entry used `S-WVU-NEWS2026` .

A zero is not one thing. A zero can mean attempted and failed, did not attempt, or withdrew, and the score sheets do not distinguish them `S-URC22-SCORES` `S-URC26-SCORES` . Every causal reading on this page is therefore an inference rather than a measured cause.

The taxonomy

Counts are entries, not incidents: one bad mission can produce two, as Missouri S&T's 2024 Delivery did with a failed signal mechanism and a failed arm axis. "Mechanism confirmed" means the source names the physical mechanism rather than only the score.

CLASS	ENTRIES	MECHANISM CONFIRMED	WHERE IT FIRES	TYPICAL COST	SHARE OF THE POINTS LOST TO ZEROS
T1 Comms and RF	7	7	Delivery, servicing, autonomy recovery, setup	40–50 points, or the whole event at the start line	Not separable in the score data
T5 Arm and end effector	7	6	Servicing, Delivery, science	30–65 points	Hides inside servicing and Delivery zeros
T4 Drivetrain and mechanical	6	5	All, because a broken drivetrain zeroes everything	25–100 points, often across three missions	Not separately measurable
T6 Software and autonomy	6	3	Autonomy	0–70 points	**About 65 %**, or 610 of 945
T12 Design process	6	5	Every mission, by construction	A whole season	Indirect, because it causes the others
T3 Connector, harness, transport	4	3	Servicing, arm, setup	45 points to the entire event	Under-counted
T7 Sensor and perception	4	4	Delivery search, servicing tags, waypoint tags	15–65 points	Feeds Delivery zeros
T10 Operator and human factors	3	3	Delivery, autonomy	The whole run	Not visible as zeros
T11 Logistics and time	2	1	Servicing, all	25–70 points or the event	Indirect, through stage time-outs
T8 Thermal	2	1	Compute, drive	Minutes of degraded autonomy; a dead motor	Small and under-instrumented
T2 Power, battery, electrical	1	0	All	Event-ending	Presumed large, evidenced thinly
T9 Dust and ingress	0	0	No sourced entry	Not applicable	**Zero sourced entries**
U Outcome only	6	0	No named cause	A task to the whole event	The residual the taxonomy cannot explain



The 54 entries by class, split by evidence grade: 51 documented by a primary source and 3 by a secondary one. Entry height measures documentation, not occurrence, which is why the arm ties with comms here while the judges describe comms as the common field failure.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

The count measures publication, not occurrence

The class with the most entries is not the class the judges name. The arm class and the comms class tie at seven. That is not evidence that the arm breaks as often as the link, because the judges state that comms is "far and away the most common problem that stops a team from being a top contender" `S-URC-ADVICE`. It is evidence of which failures get written up: a team explains a broken gearbox because it redesigned the gearbox, and explains a loose connector less often because the fix was a cable tie in a car park.

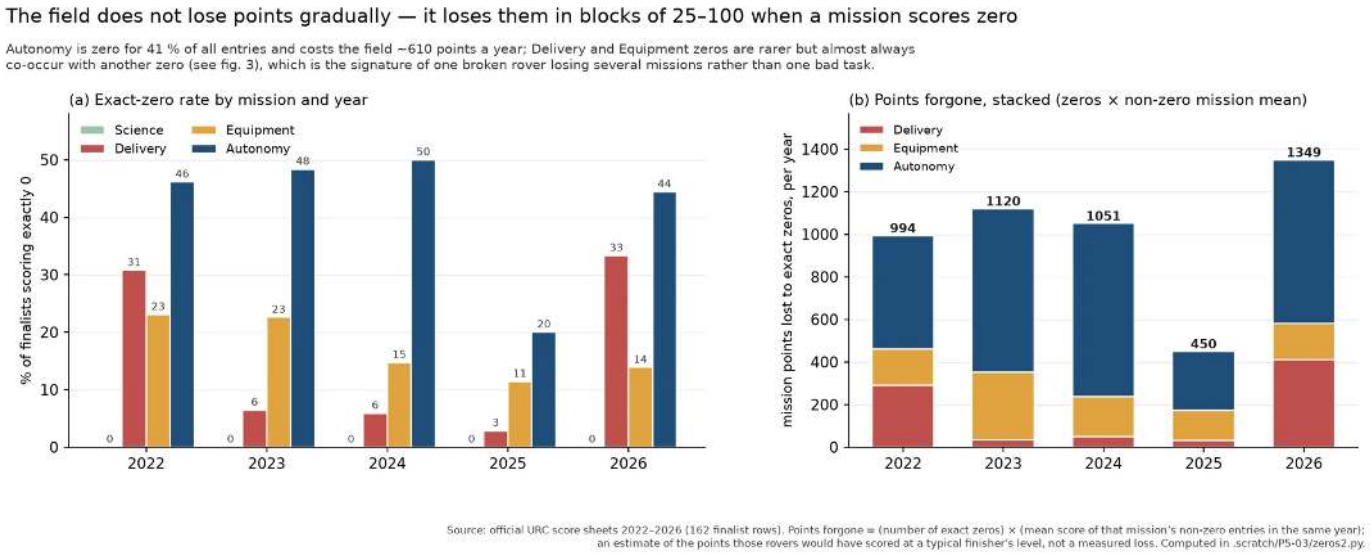
The same caution applies to the thin rows. **Power has one sourced electrical failure in five years** and **dust has none at all**, despite every subsystem study recommending sealed bearings, rated connectors and protected lenses. Two readings are possible and the data cannot separate them: those failures are rare, or they are fatal and under-published. Given that the judges describe rovers dying from a loose wire as routine, the second reading is at least as likely. The design consequence is to build to the mechanism rather than to the count.

Three further caveats belong with every percentage below. The database is a publication census rather than an incident census. Three teams account for 24 of the 54 entries because they are the best documented, and the zero rates are an upper bound on failure frequency and a lower bound on capability gaps.

The zero-score record

YEAR	N	SCIENCE ZEROS	DELIVERY ZEROS	SERVICING ZEROS	AUTONOMY ZEROS
2022	26	0	8 (31 %)	6 (23 %)	12 (46 %)
2023	31	0	2 (6 %)	7 (23 %)	15 (48 %)
2024	34	0	2 (6 %)	5 (15 %)	17 (50 %)
2025	35	0	1 (3 %)	4 (11 %)	7 (20 %)
2026	36	0	12 (33 %)	5 (14 %)	16 (44 %)
Pooled	**162**	**0**	**25 (15.4 %)**	**27 (16.7 %)**	**67 (41.4 %)**

Multiplying each zero rate by the mean non-zero score in that mission, 27.9 for Delivery, 36.3 for servicing and 45.5 for autonomy, gives the field's points forgone: about **945 mission points a year, of which about 610 are autonomy, 196 servicing and 140 Delivery** `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES` . It is an upper bound, because a team that scored zero might have scored below the mean rather than at it.



The exact-zero rate by mission and year, and the points forgone stacked. Autonomy is the persistent loss and Delivery is the volatile one: one zero in 2025 and twelve in 2026, which points at a course change rather than a change in the field's reliability.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC22-SCORES` `S-URC26-SCORES` Full resolution

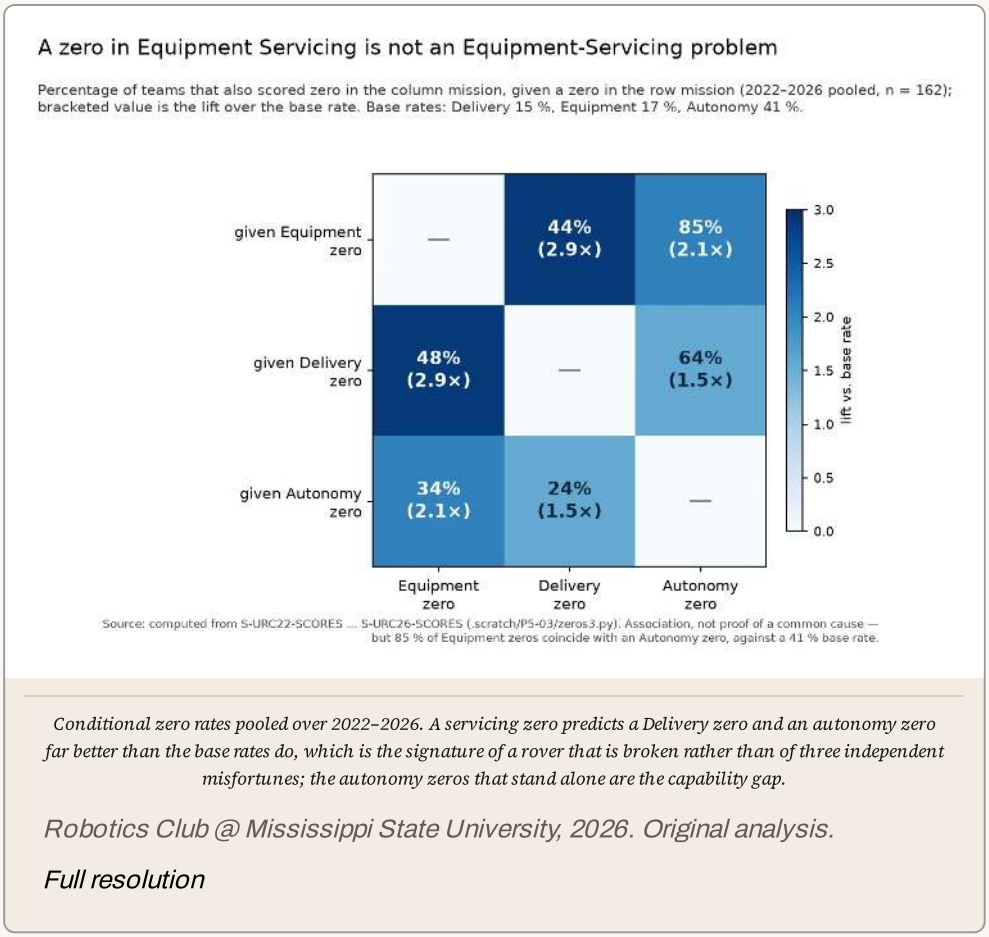
Science has scored no zero since 2022, and the review element has never explained more than 5.7 % of the spread in nine years. The 300 points in Delivery, servicing and autonomy are where the field is decided, and all three are all-or-nothing for a substantial share of teams.

What the zeros are made of

If the three missions failed independently, servicing and Delivery zeros would coincide in about **2.6 %** of team-years. The observed count is **12 of 162, or 7.4 %**, 2.9 times the independent expectation.

CONDITIONAL RATE	OBSERVED	BASE RATE	LIFT
Delivery zero given a servicing zero	44 %	15 %	2.9×
Autonomy zero given a servicing zero	85 %	41 %	2.1×
Autonomy zeros that were the team's only zero	58 % (39 of 67)	No base rate applies	Capability gap
Servicing zeros that were the team's only zero	11 % (3 of 27)	No base rate applies	Rover-level failure

Twenty-nine of 162 team-years had two or more zeros and eleven had three or more. Those are the rover-level failures, and they are where the one broken rover story lives. The isolated autonomy zero is the opposite case: the rover drove, serviced and sampled, then scored nothing when it had to drive itself.



Conditional zero rates pooled over 2022–2026. A servicing zero predicts a Delivery zero and an autonomy zero far better than the base rates do, which is the signature of a rover that is broken rather than of three independent misfortunes; the autonomy zeros that stand alone are the capability gap.

Robotics Club @ Mississippi State University, 2026. Original analysis.

Full resolution

That produces **two failure economies** with different investments. The **blocks of 100** are drivetrain, power, connectors and the arm as a single point: they cost several missions at once and are bought off with redundancy, retention and test. The **recurring 45** is autonomy, a capability 41 % of the field never demonstrates, bought with one simple planner that finishes. A first-season team that fixes the first and ignores the second lands near the median of about 230–260 points `S-URC-RESULTS`.

The decision behind each class, and the cheapest fix

Each row names the design decision that makes the class likely, not the component that broke. Cost is relative within this table only.

CLASS	THE DECISION THAT MAKES IT LIKELY	THE MECHANISM	CHEAPEST DESIGN CHANGE	THE TEST THAT REVEALS IT	COST
T1 comms	A link added late, sized by free-space loss at line of sight, with one radio and no mast	Terrain diffraction collapses range: measured 850 to 300 m, and 750 to 164 m	Radio at the mast, a 900 MHz recovery plane on its own antenna, a rehearsed behaviour on loss of link	Drive the full course out of sight and log signal strength and packet loss	$$$
T2 power	A pack sized to energy and not to sag, with the trip current below the sum of legitimate stall currents	An 8S4P pack sags about 4.6 V at 145 A; a two-joint arm stall looks like a fault	Per-joint current limit, brakes so holding is free, wide-input converters	Log bus voltage at 100 Hz through a simultaneous climb and arm stall	$$
T3 connectors	Hobby-grade connectors on an exterior harness with no strain relief and no post-transport check	Vibration work-hardens crimps; one damaged USB-C cable took out cameras and two motors in 2026	Latched rated connectors, service loops, one connector maximum in any safety path, a post-transport check	Shake the packaged rover for 30 minutes, then power on every peripheral	$
T4 drivetrain	Single-shear joints, one fastener per wheel, no secondary retention	Shock loads at a rock impact are 3–10 times static	Bolted flange and a retention ring on every wheel, double shear on every load path	Drop 0.5 m onto one wheel and drive 20 km of rock field, then check fasteners	$
T5 arm	An arm designed for static, table-top loads with no absolute encoders and no current limit	The arm that scored 94 in servicing failed under driving loads; belts slip; a lost zero needs a re-home	Absolute encoder on every joint, belt tensioner and spare, current clamp, lift before retract	Drive 500 m with the arm in carry pose, then run servicing without re-homing	$$$
T6 autonomy	The planner with the best paper performance, changed in the season it must work, with no progress monitor	Monash went 0, 14 and 86 across 2024 to 2026 by simplifying	One conservative planner, a progress checker, bounded recovery, a keep-out halo around posts	Ten full-length autonomy runs from cold boot; count completed gates	$
T7 perception	One camera, fixed exposure, one tag scale, no yaw search at the waypoint	A tag needs about 18–24 pixels to decode, and sun, dust and edge-on geometry remove it	Sunshade on every camera, exposure bracketing, slow yaw search at the waypoint, a second wide camera	Arrive with the sun behind the tag and the tag edge-on; require detection within 30 s	$
T8 thermal	Sealing the electronics in a black box and sizing cooling on a fan on the lid, with no temperature telemetry	40 W in a sealed box at 38 °C ambient implies about 77 °C internal without forced air	Size for sustained mode, filtered inlet and fan, white skin, log die temperature with an alarm	A 40-minute run at 40 °C in direct sun under full autonomy load	$$
T9 dust	Component ingress ratings treated as a system guarantee, no lens protection, filter behind the fan	Clay dust bridges contacts and blinds encoder windows, and each failure looks like software	Recessed lens with a sacrificial window, contactless encoders, external filter with access, drip loops	A two-hour dust loop, then inspect every connector and re-run calibration	$
T10 operator	The person holding the clock also driving, with no tilt limit and no geofence	Rovers have been driven into gulleys and flipped; a close approach invites an astronaut E-stop	Low centre of gravity, wide track, a tilt limit in firmware, a commander who never touches a control	Ten field trials with a stopped clock and an abort decision called aloud	$
T11 logistics	Building until the week before, an unrehearsed setup order, no freeze date	A slipped arm bought a year late; stage gates convert a slow start into a fixed low score	Order by long-lead date, an eight-minute setup drill, a freeze six weeks out, spares traced to the failure list	A full dress rehearsal with setup, both missions and teardown, timed	$$
T12 design process	Changing the architecture in the season it must win, testing at the event, treating the review as paperwork	Three of the four largest single-year losses follow a whole-system change	One system per season, a dated freeze, a known-good tag	A log showing each mission run end to end before departure	$

The root cause in one line: **rovers fail where a single element is load-bearing and untested under field loads.** The field under-tests exactly those elements, because testing them is boring, slow and happens in a car park rather than a laboratory.

The four load-bearing single points

The lift matrix says a servicing zero drags Delivery and autonomy with it, so reliability effort belongs on the elements that can zero three missions at once.

SINGLE POINT	WHAT IT DOES	THE EVIDENCE	THE MITIGATION
Wheel retention	Loses the drivetrain, and therefore every mission	BYU lost wheel mounts and fixed the symptom trackside with ties around the rims `S-BYU-2025-SAR`	Bolted flange, secondary retention ring, double shear, nothing below the contact plane
The connector in the safety path	Kills the rover or silently disables the E-stop	A damaged USB-C cable cost WVU **45 points** of servicing and cascaded to cameras and two motors `S-WVU-NEWS2026`	One connector maximum in the safety path, lockwire, continuity check every field day
Pack sag under peak current	Browns out the computer mid-run and loses the logs	A 156 A legitimate peak against an 80 A trip at 10 ms, and a published 3 s auto-restore `S-MST-2023-POWER`	Size for sag, a trip with a delay, converters rated below the sagged voltage, bulk capacitance
The arm's joint zero	Makes servicing and typing impossible without a re-home	A power glitch leaves the arm without a reference, and a re-home under time pressure can cost a mission	Absolute encoders or backed-up counters, the last pose in memory, a hands-off homing routine

The FMEA tripwires that follow from these rows are **FMEA-02, the connector, at 560; FMEA-06, pack sag, at 384; FMEA-05, spectrum congestion, at 336; FMEA-16, the arm zero, at 336; and FMEA-20, the one-metre arrival, at 294.** FMEA-31, an untested feature shipped the week before the event, is also at 294 and is the row that creates the others.

Failure modes the 2027 rules add

No team has hit these yet, because the task did not exist in this form. They belong in the failure list before the field discovers them.

2027 CHANGE	THE FAILURE MODE NOBODY HAS HIT	WHY IT IS NEWLY POSSIBLE	THE CHEAPEST COUNTERMEASURE
A bioassay is mandatory, with two life-detection capabilities and at least one bioassay	An assay that runs but cannot answer the life question, a chemistry demonstration rather than a bioassay	Chemical methods alone do not satisfy the requirement, and a dry instrument is no longer enough	A room-temperature assay with a documented positive and negative control, justified before the SAR
Wet chemistry on every science payload	A leak in motion, or a waste path that spills when the rover tilts	The no-spill rule now binds on the scoring path, in a desert with limited water	Closed fluid path, sealed waste, an inversion test with power off, a drip tray
Astronaut following, graded by command modality	A modality that works at 3 m but not at the 20 m the judges use, or a gesture fired by a bystander	Detection range is a hard requirement in daylight, and quiet is required for the gesture path	Audio as the primary, device as backup, gesture only after outdoor testing
One-metre route-finding accuracy	Arriving but stopping 1.1 m out, or a base-coordinate error that puts both targets off by a metre	Two targets at 25 points each, and the tag must also be read to indicate arrival	Own base station surveyed the day before, terminal approach at 0.3 m/s, stop at 0.6 m and hold
The second route target is deliberately out of radio sight	The rover stops dead behind the hill, and the recovery routes eat the score	Loss of line of sight is designed into the mission rather than incidental	Autonomous return as a planned behaviour, rehearsed with the recovery radio off
Servicing may begin ten minutes after another mission	Thermal and state carry-over: a hot arm, a pack at 40 %, the previous mode still latched	Back-to-back scheduling removes the cool-down that teams assume	A between-missions checklist in eight minutes, rehearsed as a pair

What a first build should do first

- **The reliability budget belongs on the failure that zeroes three missions, not on the one that costs a task.** The four single points above are the whole first season's effort, and each is a small change compared with a redesign.
- **The rover should be instrumented before the evidence is needed.** No team publishes measured link, thermal or bus data. Signal strength, bus voltage at 100 Hz and enclosure temperature with an alarm turn every field day into data.

- **Autonomy points are cheaper with simplicity than with sophistication.** The cleanest natural experiment in the dataset is Monash's autonomy score of 0 in 2024, 14 in 2025 and 86 in 2026, with the last step explicitly a simplification `S-MONASH-SAR2025` `S-MONASH-SAR2026`. One planner, a progress checker, bounded recovery and ten cold-boot full runs beat a better algorithm.
- **The transport leg is a mission of its own.** The judges name it, Michigan found damaged parts on arrival and IIT Madras found a burned board `S-UMICH-AERO-2026` `S-IITM-URC2017-REPORT`. A packaged shake test and a post-arrival check are the cheapest reliability measures available.
- **The freeze date is written down before travel.** Design-process entries appear at every level of the field, and three of the four largest single-year losses follow a whole-system change.
- **Recovery is rehearsed as well as success.** WVU's requirement that each operator train at least 30 hours turned a damaged cable into 22.5 points instead of a zero `S-WVU-NEWS2026`.

Open problems

- **The database is a census of write-ups.** Fifty-four entries from 15 teams against 162 finals team-years in five years means absence of an entry is absence of documentation, not absence of failure.
- **The zero-to-cause link is an association.** The lift matrix cannot separate one broken rover from one team that had a bad weekend; a published intervention log would settle it in one table.
- **Three teams supply 24 of the 54 entries**, which is a documentation bias rather than a claim that those programmes fail more.
- **No team publishes measured thermal, link or bus data**, so the thermal and electrical classes are argued from mechanism rather than from field measurement.
- **A new failure mode appears with each rule change**, and the 2027 bioassay and astronaut tasks need wet rehearsal and a suited trial, neither of which can be arranged at short notice.

Further reading

- Mass and power carries the pack sag and the peak-current chain behind the power failures.
- Interfaces and couplings explains why a connector in the safety path is an interface rather than a component.
- Risk register turns these classes into a ranked register.
- Environment, dust and thermal covers the classes the field documents least.
- Missouri S&T, Monash and WVU hold the best-documented loss records.

CROSS-CUTTING ANALYSIS · CHAPTER 5.4

What separates the top from the rest

The top teams and the bottom teams build the same class of machine: four or six wheels, passive articulation, an arm, a camera mast and a laptop at a command station. What separates them is a machine that works on the day, not a better idea. Across 162 finals team-years from 2022 to 2026, a team with no zero mission averages rank 10.2, one zero costs about 11 places, and the written review, the System Acceptance Review or SAR, explains 2.8 % of the spread while three missions explain 82.6 % `S-URC-RESULTS` . This page sets out the evidence for that claim, the factors that survive a causal test, and the factors that do not.

2.8 %

REVIEW SHARE OF SCORE VARIANCE

The SAR gates the field, then stops predicting

2022–2026

`S-URC-RESULTS`

82.6 %

SHARE FROM THREE HARD MISSIONS

Autonomy, equipment servicing and delivery

2022–2026

`S-URC-RESULTS`

41 %

AUTONOMY ZERO BASE RATE

67 of 162 finals team-years

2022–2026

`S-URC-RESULTS`

+175.9

POINTS BETWEEN TOP FIVE AND BOTTOM TEN

Means of finals totals, 44 of 62 teams eligible

2022–2026

`S-URC-RESULTS`

Key takeaways

- The bottom group is not slightly worse at everything. It is a normal science team (68.7, a gap of 20.8) with a broken manipulator (9.8) and no autonomy (3.7).
- One zero is a step, not a slope: mean rank **10.2** with none, **21.1** with one, **29.2** with two. The 2026 top eight scored all **32 of 32** mission-runs.
- The review is the wrong place to look for a ranking lever. It is worth 100 points and decides the field, but it has never explained more than **5.7** % of the spread in nine years.
- Framework fashion, team size, open-source output and one-year rank jumps carry no signal. The copyable behaviour is rehearsal and a no-zero discipline.

Two comparison sets and what they measure

All values come from the five official score sheets for 2022 to 2026 `S-URC-RESULTS` `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES` . A team-year enters a mean only when the team competed and recorded a total, and the 25 rows that carry a review score and no mission score are analysed separately.

Set A compares the top five and the bottom ten by mean finals total over 2022 to 2026, restricted to the 44 of 62 finals teams with at least two appearances so that a single visit cannot define either end. **Set B** compares the 2026 top eight with places 9 to 36. The cut at eight is not arbitrary: it separates the only eight teams above 300 points, and the gap from 8th (324.59) to 9th (274.91) is 49.68 points, the widest in the table apart from first to second.

COMPARISON	SAMPLE	MEAN TOTAL	MEAN SAR
Top five by mean, 2022–2026	5 teams, 23 team-years	354.6	93.03
Bottom ten by mean, 2022–2026	10 teams, 26 team-years	178.7	87.86
2026 top eight	8 teams	382.1	88.91
2026 places 9 to 36	28 teams	203.1	84.85

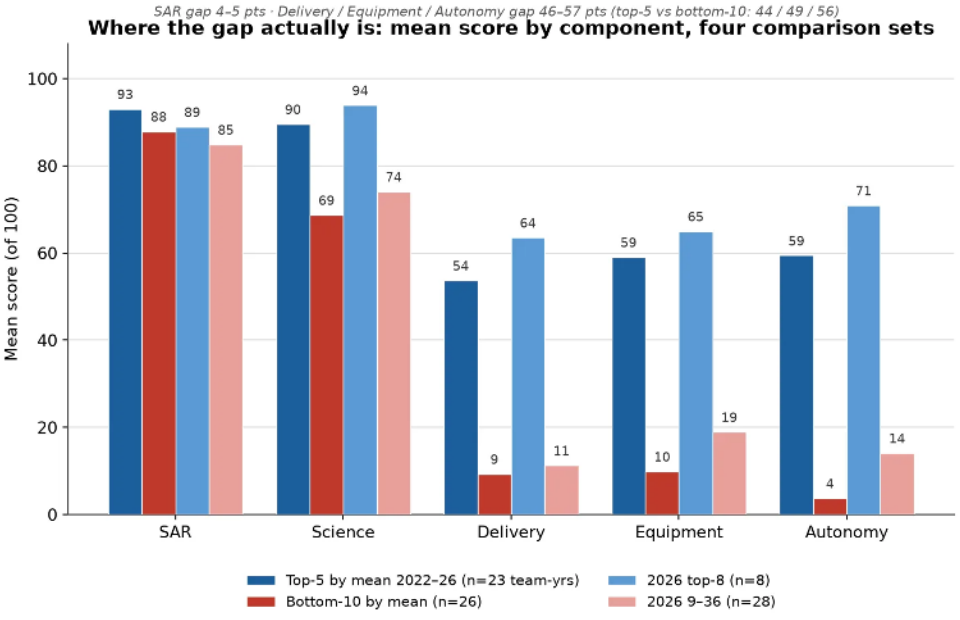
The mean final ranks of the two five-year sets are 4.6 and 27.3. The top five by mean are Missouri S&T (378.5), WVU (371.7), Monash (343.1), Michigan (336.4) and AGH (336.0). The bottom ten run from Binghamton (195.8) to Team Interplanetar (157.6). Both sets are populations of convenience, not samples, and the next section shows why the gap between them is mostly three missions.

Sample

Means are finals-only and cover 162 team-years across five editions. The 25 review-only rows are excluded from every mean on this page and are discussed on their own below. The within-season spread uses the same team-years; the year-to-year spread uses the 28 teams with at least three appearances.

The gap is in three hard missions, not the review

COMPONENT	TOP FIVE (N = 23)	BOTTOM TEN (N = 26)	GAP	2026 TOP EIGHT	2026 PLACES 9–36	2026 GAP
SAR	93.03	87.86	+5.17	88.91	84.85	+4.05
Science and astrobiology	89.53	68.69	+20.83	93.88	74.05	+19.83
Delivery	53.67	9.27	+44.40	63.50	11.24	+52.26
Equipment servicing	59.02	9.84	+49.18	64.94	18.90	+46.04
Autonomous navigation	59.39	3.71	+55.68	70.90	14.07	+56.83
Total	**354.6**	**178.7**	**+175.9**	**382.1**	**203.1**	**+179.0**



The same gaps as a chart. The SAR bars are the only pair that nearly touch, which is why a review score is a poor guide to final position. Means over 23 top-five and 26 bottom-ten team-years, 2022–2026.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC-RESULTS *Full resolution*

The direction matters more than the size. Lifting each hard mission by 20 points gains more than lifting the SAR by 5, because the four mission scores are not independent: all four need a rover that drives, so the gaps are diagnostic rather than additive. The bottom-ten profile is a team that can do science and cannot manipulate or navigate. That is a capability threshold, not a slope.

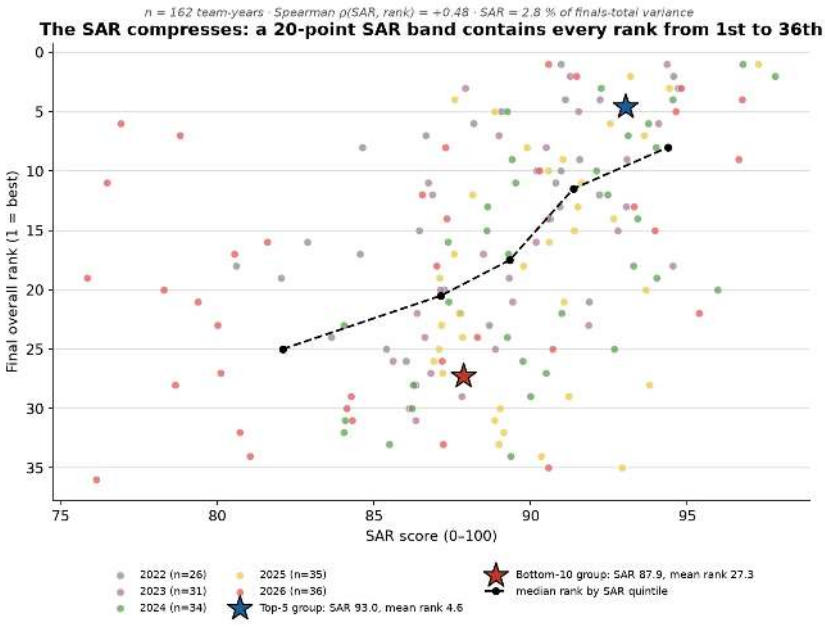
The decomposition and the zero rate for each component belong in one table, because a mission that is often zero is also the mission that moves the total.

COMPONENT	SHARE OF FINALS-TOTAL VARIANCE	ZERO SHARE, 2022–2026
Autonomous navigation	30.6 %	41.4 %
Equipment servicing	27.7 %	16.7 %
Delivery	24.3 %	15.4 %
Science and astrobiology	14.4 %	0 %
Review	2.8 %	0 %

The shares are an upper bound on each mission's causal weight, because a team that zeroes everything carries covariance with three missions at once. The ordering is robust under every tie treatment of the underlying ranks, even when individual correlations move: in 2023 Delivery, 14 of 31 teams scored exactly 15, and the correlation for that mission moves from 0.69 to 0.87 depending on how the tie is broken `S-URC-RESULTS`.

The review gates the field, then stops predicting

The review is worth 100 of the 500 points and it sets the field: about one registered team in three is invited, and a team outside the top 36 never competes `S-URC27-RULES` §3.a.iii. Of its 100 points, 48 are the video, 40 the report and 12 the judges' assessment `S-URC26-SAR`. Once a team is inside the field, the review stops discriminating. Its pooled rank correlation is +0.48, the weakest of the five components, and its Pearson correlation with the final total is −0.42. It has never explained more than 5.7 % of the spread in nine years, and 2.8 % pooled over 2022 to 2026 `S-URC-RESULTS`.



Every rank from 1st to 36th sits inside a 20-point band of review scores, so a strong document does not lift a team and a weak one does not sink it. Stars mark the mean review score and mean rank of the top-five and bottom-ten sets; the dashed line is the median rank of each quintile.

Robotics Club @ Mississippi State University, 2026. Original analysis.

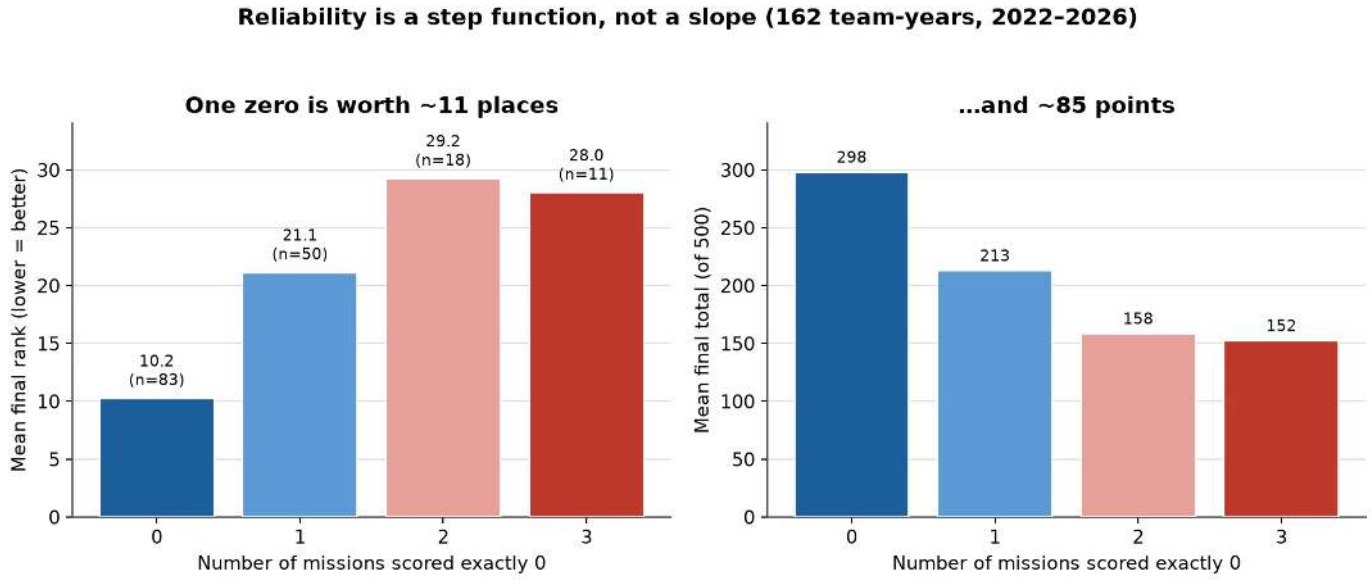
S-URC-RESULTS *Full resolution*

The 2026 field makes the point bluntly. The winner, Missouri S&T, had the 10th-highest review score (90.57). KNR had the second-highest (96.63) and finished 9th. MR2/KAIST had the third-highest (95.38) and finished 22nd. The inverse also happens: BRACU was

30th on the review (78.81) and finished 7th; Northeastern was 33rd (76.93) and finished 6th `S-URC26-SCORES`. A respectable review is also not evidence that a team reached the field. The 25 team-years with a review score and no mission score average 88.28, range 80.17 to 93.07, and 12 of the 25 sit above the median of that year's finals field. Whether those rows are withdrawals or teams that cleared the cut and were not admitted cannot be settled from the sheets.

One zero is a step change, not a slope

MISSIONS SCORED EXACTLY ZERO	TEAM-YEARS	MEAN FINAL RANK	MEAN FINAL TOTAL
0	83	10.2	297.9
1	50	21.1	212.8
2	18	29.2	158.4
3	11	28.0	152.2



Reliability is a step function, not a gradient. The move from no zero to one zero is about 11 places and 85 points, more than any plausible review improvement. Beyond two zeros the total flattens because the floor is reached; n = 83, 50, 18 and 11 team-years.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC-RESULTS *Full resolution*

The effect is not confined to weak teams. The 2026 top eight scored a non-zero on all four missions, 32 of 32 mission-runs. The other 28 teams produced 33 zeros in 112 runs, 12 of them in Delivery and 16 in Autonomy. A zero is worth between 9.7 and 14.0 places depending on the mission: mean rank 15.4 with a Delivery zero against 25.1 without, 14.6 against 28.6 for Equipment, and 11.7 against 24.3 for Autonomy. The Autonomy gap stays between 11.9 and 15.2 places in every year from 2022 to 2026 `S-URC-RESULTS`. Two of the top-five Autonomy zeros belong to Monash in 2024 and AGH in 2023, both otherwise podium-class that year. A subsystem failure is not a weakness of weak teams.

Consistency is the floor, not the ceiling

Two measures of consistency give opposite answers, and the difference is the finding.

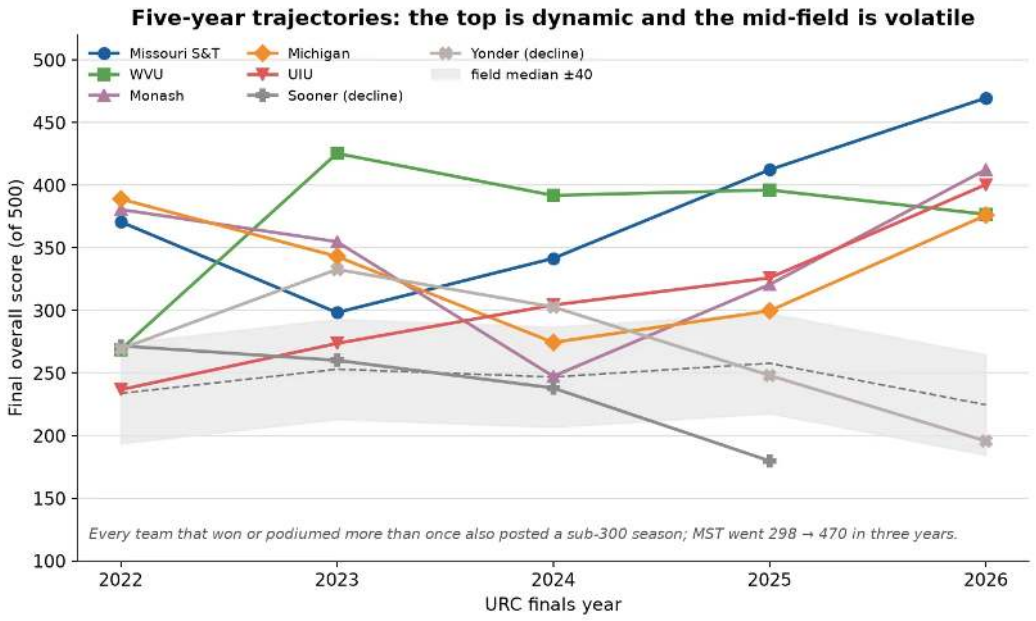
SET	WITHIN-SEASON MISSION SD	YEAR-TO-YEAR TOTAL SD
Top five by mean	22.4	52.3
Bottom ten by mean	27.3	26.4
2026 top eight	16.2	single year
2026 places 9 to 36	27.9	single year

The top five are **more** variable year to year than the bottom ten, because the bottom is pinned near the floor. Across the 28 teams with at least three finals appearances, mean total and year-to-year SD correlate at +0.56: in this dataset, better teams are more variable, not less. The most consistent totals belong to KARURA (SD 5.6, mean 191), TMR (10.2, mean 256) and Project Scorpio (10.7, mean 246), all mid-field or lower; the most volatile include Missouri S&T (58.8), Monash (56.6) and WVU (53.9), which are the winners `S-URC-RESULTS`. What the top teams do not do is zero a mission. A flat four-mission profile in 2026 is a consequence of four working missions, not an independent virtue.

The top is a small, changing club

TEAM	2022	2023	2024	2025	2026
Missouri S&T	371 (3)	298 (6)	342 (4)	412 (1)	470 (1)
WVU	269 (10)	425 (1)	392 (2)	396 (2)	377 (4)
Monash	381 (2)	355 (2)	247 (16)	321 (7)	412 (2)
Michigan	389 (1)	343 (4)	274 (9)	300 (11)	376 (5)
AGH	341 (4)	274 (8)	393 (1)	no finals result	no finals result
UIU	237 (13)	274 (9)	304 (5)	326 (6)	400 (3)
Northeastern	208 (16)	257 (13)	272 (10)	362 (3)	350 (6)
Sooner	272 (7)	260 (12)	238 (19)	180 (32)	no finals result

Finals totals out of 500, with final rank in brackets. Every repeat winner and podium team also had a sub-300 season inside the five years.



Seven trajectories against a shaded field median band. The leaders swing by more than 100 points; the climb by UIU from 13th to 3rd is the largest sustained move in the set, and Sooner's 7th to 32nd is the largest fall.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC-RESULTS *Full resolution*

The interactive version plots the same seven totals, without the median band.

Four distinct teams won the five editions and eight distinct teams filled the fifteen podium places. Thirteen teams appear in every edition, and 11 of those recorded a finals score every year `S-URC-RESULTS`. The top is not a closed group, but it is a small one, and the way in is not a single good year: UIU's rise is a change of +1.91 in within-year z-score, and the team publishes no code. Platform-change years cost a season in the two best-documented cases. Monash fitted a new arm and a new autonomy stack in the same year and scored 18 on Equipment and 0 on Autonomy, falling to 16th. Missouri S&T lost the 2023 Delivery stage gate and then a signal stack and an arm axis in 2024 before winning in 2025 and 2026 `S-MONASH-SAR2026` `S-MST-NEWS2023` `S-MST-NEWS2024`.

Returns also track position, but the gradient is survivorship, not learning. One finals appearance averages 179.6, five appearances average 303.9, and the teams that keep returning are the ones that were good enough to keep travelling. The cleanest natural experiment is a first appearance in the window: in 2026 the six first-time teams finished 22nd, 26th, 32nd, 34th, 35th and 36th, averaged 161.1 against 259.2 for the 30 returning teams, and none finished in the top half `S-URC-RESULTS`.

Mass is a marker of maturity, not a cause of score

Mass correlates with score, and the correlation fails every causal test applied to it. Over the 12 competition configurations with a published mass and a finals total, Spearman ρ is −0.60 at p = 0.04; normalised within year it is −0.52 at p = 0.08; excluding the two inferred masses it is −0.72 at p = 0.02; and leave-one-out values stay between −0.48 and −0.75 `S-URC-RESULTS`. The association is real in this sample. The mechanism is not mass.

TEAM AND YEARS	MASS CHANGE	SCORE CHANGE
UIU, 2022 to 2026	−2.3 kg	+164 points
WVU, 2022 to 2023	−20 kg	+157 points
Northeastern, 2025 to 2026	−5 kg	−12 points
BRACU, 2025 to 2026	+2 kg	+30 points

The mission decomposition is backwards from the physics. The most mass-sensitive mission, Delivery, correlates weakest with mass (−0.46, p = 0.13), and Equipment Servicing, whose workload barely depends on mass, correlates as strongly as the total (−0.61, p = 0.03); Autonomy is −0.67 (p = 0.02) and Science is −0.12 (p = 0.70). Heavy teams are heavy because the design did not converge, and the same failure costs points everywhere. There is also no interior optimum in the observed range: the 33 kg winner of 2023 S-WVU-SAR2023 and the 42 to 50 kg cluster both score in the 200 to 425 band. The limit is a constraint, not a target. The useful question is not how heavy the rover is but how much reserve it carries, because the last 3 kg of a first build arrive in the final six weeks.

Causal candidates, consequences and noise

The test applied here is stated so that it can be attacked. A **causal candidate** needs a mechanism that operates before the outcome, the right timing, and a dose-response or a natural experiment. A **consequence** is plausibly produced by winning. **Noise** behaves like a draw.

FACTOR	STATUS	WHAT THE RECORD SHOWS	WHAT A FIRST BUILD TAKES FROM IT
No-zero reliability	Causal candidate, strongest	Rank 10.2 with no zero, 21.1 with one, 29.2 with two; the 2026 top eight 32 of 32 non-zero	Zero avoidance is the first design filter, ahead of performance
A link that works at the site	Causal candidate	No bottom-eight team reached 60 on a teleoperated mission in 2024, 2025 or 2026 (0 of 8 each year); obstructed range is about one-fifth of line of sight (850 m to 300 m, 750 m to 164 m)	Two bands, a 2 to 3 m mast, a directional base and one deliberate out-of-sight test
Autonomy that completes the course	Causal candidate	67 of 162 team-years zeroed; rank 11.7 against 24.3; a 2026 gap of +56.8	One conservative planner, a progress checker and ten cold-boot full-length runs
Manipulator precision	Causal candidate	Highest pooled rank correlation of any mission (ρ = 0.77); top-five mean 59.0 against bottom-ten 9.8	A wrist camera, a panel-parallel prismatic axis and task-specific end effectors
Rehearsal	Causal candidate, weak measurement	WVU publishes 30, 50+ and 100 operator hours in different documents; UIU reports 80+ hours of drive time	The structure (timed mock missions, a replica lander, degraded modes), not the number
Analogue-terrain testing	Causal candidate	Tucumcari and Hanksville campaigns, a beach test, coal-ash fields	One off-campus campaign before the review video
Mass discipline and weigh-in margin	Causal candidate, thin evidence	ρ = −0.60, but within-team mass changes explain none of the score change	4 kg of margin, not a ranking lever

FACTOR	WHY IT IS DOWNSTREAM	THE COUNTEREXAMPLE
Team size	Winning attracts members and institutional support	Size is known for 4 of 62 teams S-CENSUS ; WVU won with 50 members in 2023 S-WVU-SAR2023 and finished 4th with more than 80 in 2026 S-WVU-SAR2026
Documentation volume	A documentation subteam needs continuity and spare capacity	The three teams with a full public report set average 344.2 S-CENSUS S-URC-RESULTS , but one of them, BRACU, finished 7th in 2026 and 21st in 2024 S-URC-RESULTS
Open-source output	Publishing follows slack and staff, not the reverse	UIU publishes no code and finished 3rd in 2026; BYU publishes none
A drone or a specialised instrument	Discretionary projects appear after the base platform works	Missouri S&T's 2024 drone never flew; the Raman instrument moved its Science score 96 → 79 → 82 → 96 → 93 across 2022 to 2026

Four more candidates are noise or an artefact, and none should carry a theory:

- **One-year rank swings.** ARES went 30 → 8, UdeG 35 → 14 and Sooner 7 → 32. A single jump is not a method.

- **The 2026 Autonomy tie at 86.** Five teams scored exactly 86 and they are exactly ranks 1 to 5. The 2026 Autonomy column holds only six distinct non-zero values across 36 teams, so the podium was decided by the other four components.
- **Penalties.** One finals row in 187 carries a non-empty penalty cell (UdeG 2025, 16.39). Penalties are a rules risk, not a score driver.
- **Planner sophistication.** Monash went 70 → 0 → 14 → 86 while simplifying, and UMD Loop's textbook Nav2 configuration finished 32nd of 36 in 2026 `S-MONASH-SAR2026` .

> The strongest counter-reading
>
> A team that zeroes one mission loses about 85 points, but the sheets cannot distinguish a mission that was attempted and failed, one that was not attempted and one that was withdrawn. The reliability claim is a statistical pattern, not a failure taxonomy. The per-mission failure evidence sits on How rovers fail.

What this comparison cannot show

The sets are populations of convenience: five top teams, ten bottom teams and four documented programmes, all four of them Tier A and all four in the 2026 top eight. There is no bottom-ten dossier, so the qualitative half of this page reports what winners do, not what losers fail to do. Everything is correlational except the arithmetic identities, and two of the factors are measured weakly: the documentation rating is a first-pass check of team websites rather than a formal review `S-CENSUS` , and the operator-hour claims are self-reported.

Two further limits are specific to 2027. The 2026 Autonomy mission was scored in blocks, so a 2026 autonomy score is a threshold, not a continuous capability measure. The 2027 mission replaces that course with astronaut assistance and two 1 m route targets `S-URC27-RULES` §1.e. None of the four documented programmes describes a person-follow, gesture, sign or audio capability, so a 2027 team that copies a 2026 stack scores the route targets and forfeits the astronaut points.

What a first build takes from the record

1. **Zero avoidance comes before performance.** One zero costs about 11 places and 85 points, more than any plausible review improvement, and the 2026 top eight had none.
2. **The link is a subsystem with its own test campaign.** Two bands, a 2 to 3 m mast, a directional base and one deliberate test in the shadow of a hill. Obstructed range should be expected at about one-fifth of the line-of-sight figure.
3. **Autonomy is chosen for completion, not sophistication.** One conservative planner, a progress checker and ten full-length cold-boot runs beat a cleverer stack; the 41 % zero base rate is the number to beat.
4. **A wrist camera and one task-specific end effector carry Equipment Servicing.** The arm's pooled rank correlation of 0.77 is the highest of any mission, and actuator brand and degrees of freedom show no signal.
5. **The review is completed, on time and in the required format, and then treated as insurance.** It gates the field and then explains 2.8 % of the spread, so effort past a complete package buys almost nothing.
6. **Mass margin is kept and one major subsystem changes at a time.** The record's losses cluster in seasons where two large subsystems changed together.

Further reading

- How rovers fail works the failure record behind the zero rates.
- Mass and power carries the mass model and the reserve argument in full.
- Where the competition is going sets these findings against the rule history.
- Top versus rest gives the same statistics with definitions, samples and data links.
- Results explorer, 2022 to 2026 holds the score table every figure on this page is computed from.
- The strongest documented programmes are Missouri S&T, WVU, Monash and UIU.

CROSS-CUTTING ANALYSIS · CHAPTER 5.5

Where the competition is going

The 2027 rover must do more than the 2019 rover inside the same 50 kg, the same 1.2 m cube and the same off-grid field. That is the shape of the competition's history: the physical envelope froze, the scored content inside it kept moving, and the median score has not improved for five years. This page plots the quantitative record from 2013 to 2026, separates what is frozen from what moves one way and what churns every season, and states a labelled outlook to 2030.

469.6

WINNING TOTAL, 2026

Out of 500; the first total above 450

2026

S-URC26-SCORES

224.6

MEDIAN FINALIST, 2026

The median has sat between 224.6 and 257.8 since 2022

2022–2026

S-URC-RESULTS

245.0

GAP FROM WINNER TO MEDIAN

Widest of the modern era

2026

S-URC26-SCORES

116

TEAMS REGISTERED IN 2026

Against a finals field of 36 to 38

2022–2026

S-URC26-TEAMS

Key takeaways

- The physical envelope has not moved since 2019: 50 kg per deployed configuration, a 1.2 m cube and a 70 kg fielded limit. A 40 % oversize penalty is larger than the winner-to-median gap in most years.
- The top is separating and the median is not. The last five medians sit inside a 33-point band; the winning total rose from 389.0 to 469.6.
- Difficulty is concentrated rather than drifting. Delivery, Equipment Servicing and Autonomy carry zero shares of 15.4 %, 16.7 % and 41.4 %, and together they explain 82.6 % of the variance; the review carries a 0 % zero share and 2.8 % of the variance.
- Every 2027 capability change lands on the three missions that are already bimodal, which is why the scope decision outweighs the architecture decision.

The series and what it cannot support

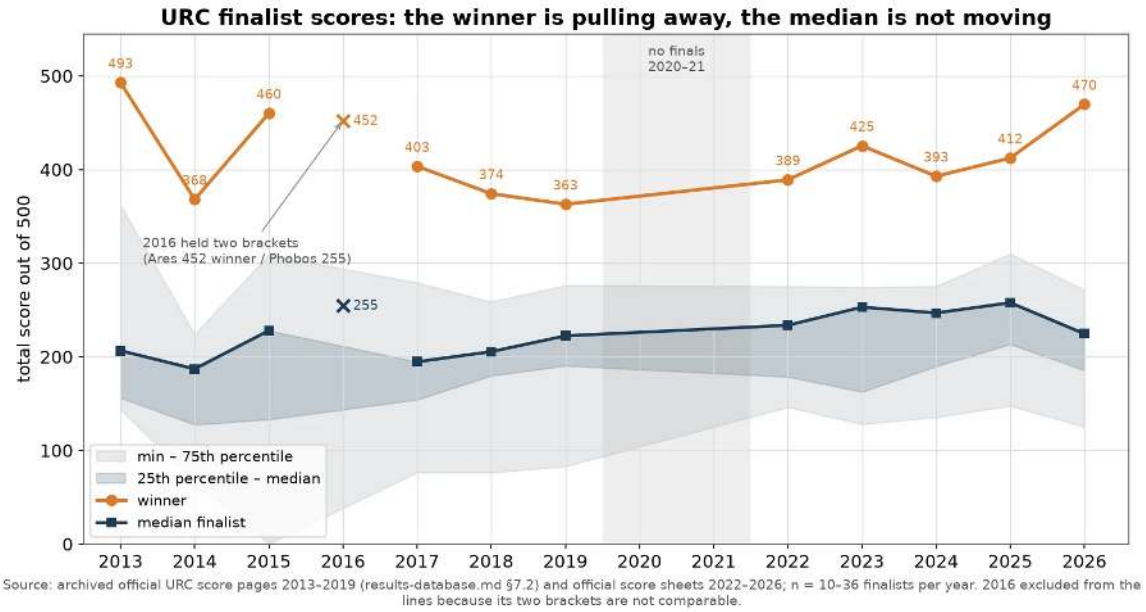
The series covers eleven comparable editions: 2013 to 2019, excluding 2016, taken from the archived official score pages, and 2022 to 2026 from the official score spreadsheets S-URC13-SCORES S-URC14-SCORES S-URC15-SCORES S-URC17-SCORES S-URC18-SCORES S-URC19-SCORES S-URC22-SCORES S-URC23-SCORES S-URC24-SCORES S-URC25-SCORES S-URC26-SCORES . No finals were held in 2020 or 2021 S-URC-HISTORY .

Four limits apply to every number on this page. Cross-era pooling is invalid: the 2013 to 2016 tasks were different objects under different scoring, and those totals include a scored presentation, so each series is read year by year and never as an average since 2013. 2016 ran two brackets, with winners on 452.3 and 255.0, and is excluded from every line S-URC16-SCORES . The score sheet is not the registration list: the 2022 to 2026 sheets hold 36 to 38 rows against 99 to 116 registrants. And a change in difficulty cannot be separated from a change in who was in the field, so no causal claim rests on a single year's move.

The top is separating, the median is not

YEAR	N	WINNER	MEDIAN	SD	WINNER – MEDIAN
2013	10	492.8	206.5	115.3	286.3
2014	23	368.0	187.0	81.1	181.0
2015	23	460.0	228.0	108.4	232.0
2017	36	403.4	194.7	75.0	208.8
2018	36	374.2	205.5	61.5	168.7
2019	36	362.9	222.4	66.2	140.5
2022	26	389.0	233.6	71.6	155.3
2023	31	425.4	253.0	73.9	172.3
2024	34	392.8	246.8	65.1	146.0
2025	35	412.3	257.8	65.6	154.5
2026	36	469.6	224.6	84.6	245.0

Totals out of 500. 2016 is omitted because it ran two non-comparable brackets.



The interquartile band has stayed within 60 points at either edge since 2017, while the winner has pulled away. 2026 is the most stretched modern field: the standard deviation is 84.6 against a previous modern high of 75.0, and the winner sits 245.0 points above the median.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

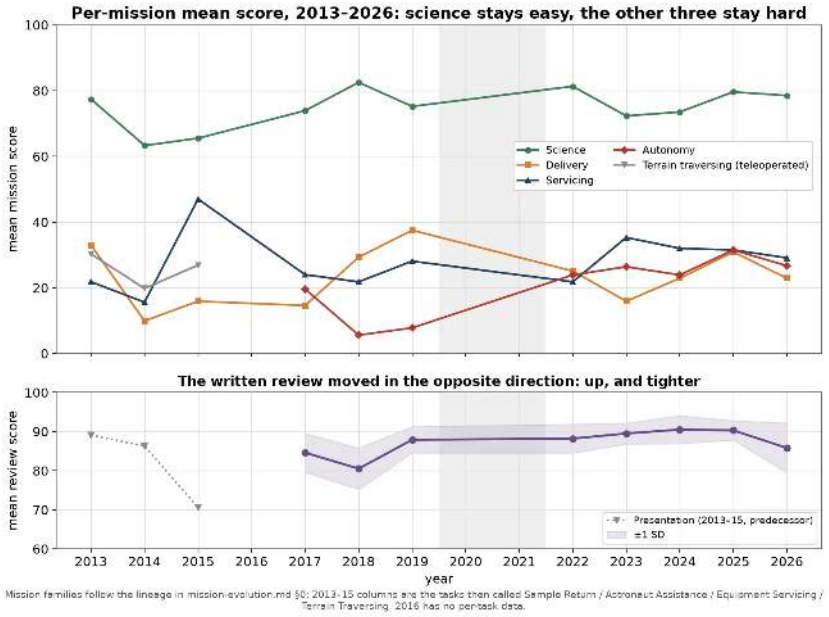
Fitted over 2017 to 2026, the median rises 5.2 points a year ($R^2 = 0.61$) and the winner 6.7 ($R^2 = 0.46$), but the last five medians sit inside a 33-point band with no monotone trend and the 2026 median is 30 points below its own fit line. The defensible statement is that the top is separating, not that the field is improving. The winning margin is noise rather than trend: first-to-second gaps across the modern era run 67.1, 35.2, 33.4, 8.4, 70.5, 1.0, 16.1 and 57.2 points, so a single good or bad run still decides a title.

The last-place total is a reliability statistic, not a capability measure. It runs 125.1 to 147.5 across 2022 to 2026, and the 2015 low of 0.0 is a team that scored only its presentation. The 2013 total of 492.8 is also not a high-water mark: ten teams, five tasks, and one team scoring 100 on four of the five `S-URC13-SCORES`.

Mission means are static while the review rises

MISSION FAMILY	2013–15 (N = 55)	2017–19 (N = 106)	2022–26 (N = 162)	DIRECTION
Science	66.8	77.3	76.9	one step up, then flat for nine years
Delivery	16.5	26.9	23.6	low and stuck
Servicing	29.3	24.6	30.3	flat, very wide (SD 22–34)
Autonomy	not a mission	11.0	26.7	new in 2017, up 143 % since
Review	80.4	84.3	88.9	up, and tighter

Means per family, with renamed tasks mapped onto a common line: Sample Return and Science Cache to Science, Astronaut Assistance and Extreme Retrieval and Delivery to Delivery, CDR and SAR to Review. Autonomy did not exist before 2017, and the 2013 to 2015 presentation is the review's predecessor.



Mission means have held between about 15 and 38 points out of 100 for a decade while the review has climbed. The 400 mission points are hard to score and the 100 review points are easy to score, which is the single most useful asymmetry in the competition for planning purposes.

Robotics Club @ Mississippi State University, 2026. Original analysis.

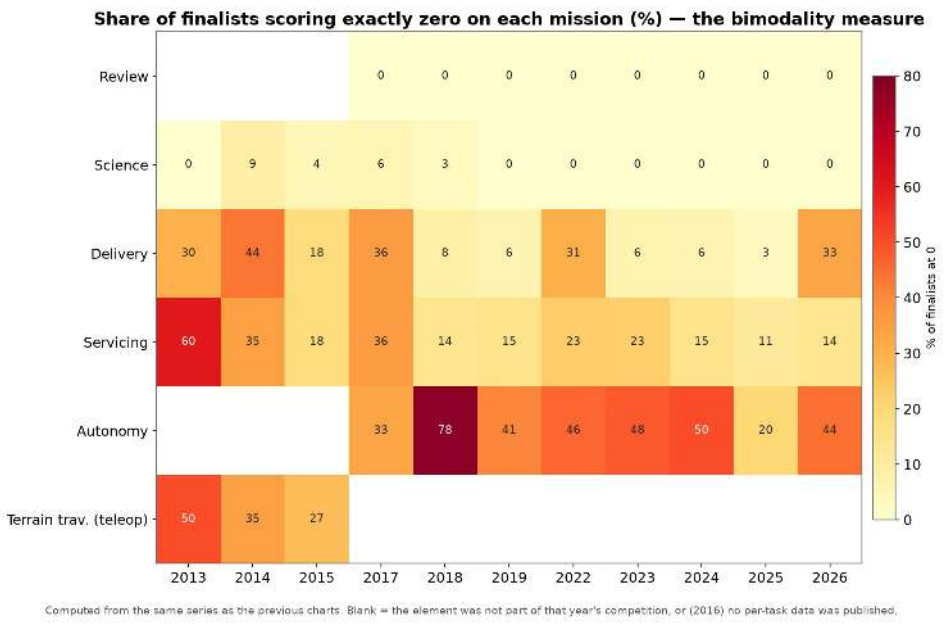
S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

The review has become the most generous and most compressed component. Its 2022 to 2026 mean is 88.9 with a standard deviation of 4.5, 95 % of finalists clear 80 and none score zero; the mission standard deviations are 16 to 28. It has never explained more than 5.7 % of the spread of final scores in nine years, and 2.8 % pooled over 2022 to 2026 S-URC-RESULTS . It gates entry and contributes 100 points, and once a team is in the field it does not predict position. The review is a gate that can lose a season and cannot win one. The 2026 re-base is the one exception to the compression: the mean fell from 90.3 to 85.8, the standard deviation rose from 2.5 to 6.4, and it was the first year since 2018 that fewer than 95 % of finalists cleared 80.

Difficulty is concentrated, not rising

Pooled over 2022 to 2026, the zero shares are 0 % for Science, 15.4 % for Delivery, 16.7 % for Equipment Servicing and 41.4 % for Autonomy. Delivery, Servicing and Autonomy therefore carry 300 points that are all-or-nothing, and they explain 82.6 % of the

variance in final totals; the review's 100 points explain 2.8 % S-URC-RESULTS .



The zero share is the cliff measure of difficulty: a mission either works or it does not. Delivery and Servicing have been bimodal for thirteen years, and neither has ever had more than 27 % of the field above 80. Science has had no zero since 2022.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

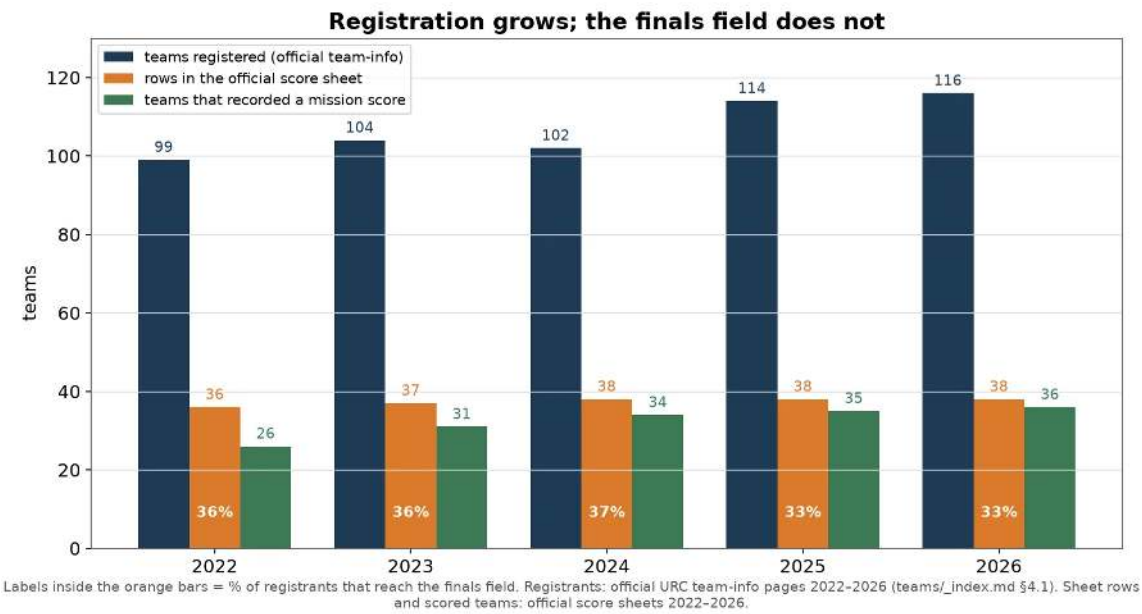
Autonomy is the hardest mission and it improves slowly. Its zero share was 51 % pooled over 2017 to 2019 and 41 % pooled over 2022 to 2026, but 2026 rebounded to 44 % after a 2025 low of 20 %, and only 14 % of the 2026 field cleared 80. A high zero share is a joint product of task difficulty, an all-or-nothing scoring rule and course design, so the measure is not difficulty alone. The 2018 result of 78 % zeros followed the first full year of a scored autonomy mission, and 2027 replaces more of the mission than 2018 did, so 2018 is the closest available comparison rather than a prediction.

The funnel is narrowing

YEAR	REGISTERED	SCORE-SHEET ROWS	RECORDED A MISSION SCORE	ROWS ÷ REGISTERED
2022	99	36	26	36 %
2023	104	37	31	36 %
2024	102	38	34	37 %
2025	114	38	35	33 %
2026	116	38	36	33 %

Registrations from S-URC22-TEAMS S-URC23-TEAMS S-URC24-TEAMS S-URC25-TEAMS S-URC26-TEAMS ; sheet rows and scored teams from the official score spreadsheets.

Registration grew 17 % in four years while the finals field stayed at 36 to 38 rows and the top-36 cut has been unchanged since 2019 S-URC19-RULES S-URC27-RULES §1.a.iii §3.a.iii. The funnel therefore narrowed from about 2.8 to 1 to about 3.1 to 1. Between 33 % and 37 % of registrants appear in the score sheet at all, and the number that recorded a mission score grew from 26 to 36 while the field itself stabilised.



Registration keeps rising and the field does not. The line that moved is the number of teams that actually completed a mission, from 26 to 36, while the score-sheet count held at 36 to 38.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC13-SCORES *S-URC19-SCORES* *Full resolution*

The field is also narrower than the registrant base. It is drawn from 10 country labels in 2022 to 2025 and 12 in 2026, against 15 registered countries in 2022 S-URC22-TEAMS S-URC-RESULTS . The United States share of the field runs 29 % to 45 % and fell in team count from 2024 to 2026, while the United States mean total of 264.4 is barely above Poland's 266.3. Making the field is the hard part; once there, the United States advantage is headcount, not performance. Registrations before 2022 were not recovered, so any claim that the competition used to be easier to enter is unsupported.

The envelope has not moved since 2019

Three rules set the physical design, and none has changed in eight years or more:

- **50 kg per deployed configuration**, with a penalty of 5 % of the points scored per kilogram over, rounded down, since 2011 S-URC07-RULES-L1 S-URC11-RULES S-URC27-RULES §2.b §2.a.iii.
- **70 kg of fielded equipment**, in force from 2015 S-URC15-RULES S-URC27-RULES §2.a.iii.
- **A 1.2 m cube at weigh-in**, in any orientation, including wheels and antenna, with folding allowed and disassembly not, and a **40 % penalty** for exceeding it, in force from 2019 S-URC19-RULES S-URC27-RULES §2.a.ii §2.a.ii.

The cube is the most consequential mechanical rule for a new team, because 40 % of a 400-point score is 160 points, more than the median-to-winner gap in most years of the series (140.5 to 245.0 in the modern era). A rover that is excellent and 1.3 m long scores 60 % of what it earns. The 2027 rover is the 2019 rover's envelope with 2027's perception and sample handling inside it, and that is the strongest single predictive statement about any design decision made now.

What moves inside the envelope

Autonomy: from optional to a human-centred mission

Autonomy was not a scored requirement before 2017: the early rulebooks say the rover need not be autonomous at all, and the 2017 text still limits the requirement to the new Autonomous Traversal task S-URC07-RULES-L1 S-URC15-RULES S-URC17-RULES §2.c. The 100-point mission arrived in 2017, the same year the CDR became scored. It ran on staged legs between physical markers from 2019, with the most explicit per-leg partial-credit scale URC has published in 2021, then on AR-tag waypoints and object recognition at 3 m and 2 m tolerances from 2022 S-URC19-RULES S-URC21-RULES S-URC26-RULES §3.e §3.e §1.f. 2027 replaces the course with two 50-point sub-missions, astronaut assistance and route-finding to two AR posts at 1 m tolerance S-URC27-RULES §1.e.

The technical content of the change is the modality ladder. One 15-point Follow task pays 5 for a carried command device, 5 for a visual sign, 10 for a spoken word and 15 for a silently given beckoning gesture S-URC27-RULES §1.e.vi. That is a perception problem, audio and gesture classification from a rover several metres away, replacing a radio problem, and it is the first time the rules have made the command modality itself worth points.

Science: three rewrites in five years

Science was rewritten from a presentation task to a cache-selection task in 2023, and the mean fell from 81.3 to 72.3 with the share above 80 falling from 50 % to 42 %. The 2024 rewrite changed everything at once: 10 cm depth, 5 g minimum, a sealed cache, a two-site comparison and a mandatory stratigraphic profile. The mean held at 73.5 but only one team reached 90 and the share above 80 fell from 42 % to 24 %, the lowest in the series. The 2027 rules demand two life-detection capabilities with at least one bioassay and delete the stratigraphic profile S-URC27-RULES §1.b. There is no 2027 result to measure; the closest analogue is 2024, when a comparable rewrite depressed the share above 80 by 18 points for one year before teams learned the task.

Communications: harder every year, in three ways

Spectrum scope has widened monotonically. The 900 MHz Low, Mid and High sub-band scheme with an 8 MHz maximum dates from 2014, 2025 removed the 2.4 GHz restrictions in the only relaxation in the record, and 2027 extends the 900 MHz rule from the command station and rover to every link between the station, rover, drone and astronaut, including the suit's video and audio, inside the same 8 MHz S-URC14-RULES S-URC25-RULES S-URC27-RULES §3.b §3.b.vi §2.c.v. Line-of-sight guarantees have shrunk over the same period: 2026 promised line of sight for Science, Equipment Servicing and Autonomous Navigation and "not more than 50 %" for Delivery, while 2027 keeps it only for Astrobiology and Equipment, says autonomy is likely not line of sight and makes later Delivery tasks intentionally obstructed S-URC26-RULES S-URC27-RULES §2.b §3.d.ii. And the operating window has not grown while the equipment count has: command-station setup fell from 25 minutes to 15 by 2017 and stayed there, while the deployed equipment grew to include a 900 MHz link, a 25 m cable run, a drone and, in 2027, a suited astronaut with video and audio S-URC11-RULES S-URC27-RULES S-URC27-RULES §3.g §1.a.ii §3.e.viii.

The drone: consulted, then regulated for five years

The drone entered as an optional, unscored demonstration in 2022, when the rules invited teams to comment on making it a scored event S-URC22-RULES S-URC22-RULES §4.a §4.g. In 2023 it could fly instead of the rover in autonomy S-URC23-RULES . From 2024 to 2026 it was a Delivery assistant, launched from the rover or from a start location near the command station S-URC26-RULES §1.d.iv. In 2027 it moves from the Delivery mission to the equipment regulations, launches only from a designated pad, carries an unwaivable dummy mass within a 5 kg allowance, and an airframe over 5 kg "will not be permitted to be used" S-URC27-RULES §2.b. The five-year direction is tightening, not loosening, and the drone is a strategic option rather than a default: it can help on two of four missions and it costs 5 kg of the 50 kg limit, an airframe, a pilot, a spotter and pad discipline.

Documentation: from a presentation to a hard gate

The review element has been redesigned more often than any mission. A video presentation carried 100 points from 2011 to 2014; the 2015 CDR was pass-fail; the 2017 CDR was scored at 100 of 500; 2018 renamed it the SAR and added the PDR; and the top-36 cut has stood since 2019 S-URC15-RULES S-URC17-RULES S-URC18-RULES S-URC19-RULES §1.h.iii §3.a.ii §1.i §1.a.iii. Three consequences are still true in 2027. Documentation is a fifth mission worth as much as any field mission, and it gates the other 400 points. The documentation load grows as documents rather than as points: 2027 adds a dated Astrobiology Plan due 2027-05-21 and a communications document due 2027-04-30 on top of the PDR, the SAR and the expense report S-URC27-RULES S-URC27-RULES S-URC27-RULES §1.b.i §2.c.iv §3.c.ix. And the rubric rewards specific artefacts, with the science page worth about 15 % of the SAR report and the payload 6 % of its video S-URC26-SAR .

The object list churns every year

Six Equipment Servicing sub-tasks have been introduced and then deleted inside eight years, and three of the 2027 tasks are new in 2026 or 2027. The keyed tubular pin-tumbler lock, in which a key is removed, inserted into a second lock and turned, appears in 2027 only, and the insertion it replaced, a USB-C cable, is gone S-URC19-RULES S-URC22-RULES S-URC26-RULES S-URC27-RULES §3.d §3.d.i §1.e.ii §1.d.ii. What recurs every year is the underlying skill set: reach 1.5 m to ground level, grasp a 5 to 10 cm object, apply torque, insert something and read a 1 to 2 cm fiducial marker. The Delivery grasp feature has also widened from 5 cm in 2018 to 2025 to 7.5 cm in 2026 and 2027, and the rules now say objects will come in a variety of sizes, shapes and weights S-URC26-RULES S-URC27-RULES §1.c.iii §1.c.iii. The competition is telling teams that its objects are becoming less predictable, not more.

What 2027 changes

2027 CHANGE	WHAT THE RULE SAYS	WHAT IT FORCES ON THE DESIGN
Bioassay mandatory	Two life-detection capabilities, at least one a bioassay S-URC27-RULES §1.b.v	Fluid handling with containment, a no-spill regime for every onboard chemical, a pre-approved hazardous-chemicals plan, and an assay inside a 20 to 30 minute window that also covers two site visits and a sample of at least 10 cm S-URC27-RULES §1.b
Astronaut assistance	Suit camera and microphone in the command station (5 points), stop within 3 m (5), Follow graded 5, 5, 10 and 15 by modality, Stay beyond 20 m (5), Fetch a rock-pick hammer with teleoperated pickup scoring zero (10), Come (5), Give (5) S-URC27-RULES §1.e	Audio and gesture perception of a target that is not a fiducial marker, arm precision to a hammer with no teleoperated fallback, and a human-safe stop behaviour
1 m route-finding	Two AR posts worth 25 each, stop within 1 m and indicate arrival, target 2 behind a hill and out of line of sight, a 1 m digital elevation model provided S-URC27-RULES §1.e	Terrain traversability mapping on board, 1 m absolute accuracy from RTK-grade localisation or visual servoing, and autonomy that works with the link down
Key and lock	Remove a key from a tubular pin-tumbler lock, insert it into a second lock and turn it, with a 1 × 1 cm ArUco marker on key and keyhole S-URC27-RULES §1.d.ii	Sub-centimetre visual servoing on a 1 cm marker, a grasp of a small flat key, and torque applied while alignment is held
Autonomy then servicing	Servicing may begin as soon as 10 minutes after Autonomy, at the same command station and on an adjacent course S-URC27-RULES S-URC27-RULES §1.a.iv §1.d.i	A configuration change in 10 minutes including a possible re-weigh, with stowage and quick-change designed in
Comms scope widens	The 900 MHz 8 MHz sub-band covers every link between station, rover, drone and astronaut, with line of sight guaranteed only for Astrobiology and Equipment and later Delivery tasks intentionally obstructed S-URC27-RULES S-URC27-RULES §2.c §3.d.ii	A written frequency plan, rehearsed loss-of-link behaviour and less reliance on video at range
Drone pad-only	A designated field pad only, 5 kg including an unwaivable dummy mass, and an over-5 kg airframe “will not be permitted to be used” S-URC27-RULES §2.b	A separate operating system with an airframe, a pilot, a spotter and pad discipline, not a rover accessory

The penalties for losing the link are rising at the same time as the course moves behind hills. Teleoperating back from route-finding costs **50 %** of that mission’s points, against 20 % before, and the autonomy abort signal costs 20 % S-URC27-RULES §1.e.

The labelled outlook to 2030

Nothing in this section is a statement about a published future rule. URC amends its rules annually and has published no roadmap. Every row is an inference from a published rule or score series, and the confidence is a judgement about how likely the pattern is to repeat.

Three regimes explain why some things can be forecast at all.

REGIME	CONTENTS	BEHAVIOUR, 2013–2027	FORECAST
Frozen	50 kg, the 1.2 m cube, 70 kg fielded, the four-mission architecture, the top-36 cut, the 20 % intervention penalty, no tether, no direct view, dust	no change found over 8 to 20 years	trust through 2030, high confidence
Ratcheted	autonomy content and tolerance, comms scope, science payload requirement, documentation count	one direction, with two exceptions: the 2.4 GHz relaxation in 2025 and the drone scope narrowing in 2024	one more notch in 2028 to 2030, medium-high confidence
Churning	servicing sub-tasks, the delivery menu, autonomy targets	six sub-tasks added then deleted in eight years; three 2027 tasks new in 2026 or 2027	the 2027 object list is not assumed to survive to 2029, high confidence

FORECAST FOR 2028 TO 2030	EXTRAPOLATED FROM	CONFIDENCE
The physical envelope does not change	20 years of the 50 kg limit and 8 years of the cube, with no exception found	high
Attendance limits stay and are enforced	the 2027 clause gives a reason outside the competition, permitting and environmental impact S-URC27-RULES §3.b.iii	high
Human-robot interaction becomes larger and more independent	2027 already grades command modality 5, 5, 10 and 15, and every past autonomy change has added content	medium-high
Science deepens toward quantitative, multi-site life detection with a control	the 2027 bioassay clause, the Q&A ruling that a chemical method on its own is insufficient S-URC27-QA AQ9, and the 2024 precedent of adding depth a year after a simpler rule	medium
The stratigraphic-profile class of task does not return	it was removed and deweighted in 2027 after three years	medium-high
Comms gets stricter, including an explicit behaviour on loss of link	2027 extends the 900 MHz rule to the astronaut and says autonomy is likely not line of sight; the rules have recommended a heartbeat monitor since 2024 S-URC24-RULES §3.a.v	medium
The drone stays a regulated field asset, not its own scored mission	2027 moves drones into the equipment class and removes rover launch and the wind waiver; the 2023 substitution was not repeated	medium
More all-or-nothing scoring blocks and fewer quality-judged elements	2027 scores autonomy in discrete blocks and deletes three judgement-heavy tasks	medium

Each row has an obvious falsifier, and a falsified row should be deleted rather than reinterpreted. A 2028 rulebook that relaxes the 900 MHz scheme, removes the astronaut content, lowers the payload count or restores stratigraphy deletes the corresponding forecast.

What the forecast means for a design

Designing for 2027 and designing for 2030 are mostly the same job, because the frozen regime dominates the mass, volume, power and thermal allocations. The forecast changes three decisions, and each is cheap to keep open:

1. **Payload bay and fluid handling sized for a second assay and a control.** That is the most likely margin to be needed, and it is a volume and containment problem rather than an instrument problem.
2. **Perception compute and audio sized for human interaction, not for AR tags.** Audio and gesture classification is compute-hungry in a way that waypoint navigation is not.
3. **Autonomy tested with the link down, not merely weak.** If loss-of-link behaviour becomes a scored requirement, only teams that have already tested it will score.

One thing the forecast should not change: the 2027 object list is not a design target. The same clause that added the key and the lock deleted a USB-C insertion, and the 2028 list is at least as likely to differ as to repeat.

Three ways a first-time team loses a season

Ranked by probability and damage, with one upstream cause: attempting too much of the new content.

1. **It never reaches the Utah field.** Between 33 % and 37 % of registrants appear in a score sheet at all, and 25 team-years in 2022 to 2026 carry a review score and no finals result, at 84 to 93; MIST Mongol Barota's 92.85 in 2022 would have been third-best among that year's 26 finalists S-URC22-SCORES . In 2027 the science plan must describe two life-detection capabilities including a bioassay and is due 2027-05-21, and a communications document is due 2027-04-30.
2. **It reaches the field with nothing that works reliably.** The 300 points in Delivery, Servicing and Autonomy are bimodal, and each intervention costs a further 20 % of that mission's points with the clock running. The roughly 240-point median is the profile of a team with everything working and an autonomy score of 30 to 50.
3. **It over-reaches on the new content.** 2027 adds four capability families at once on an unchanged envelope. Rule discontinuities show up as one-year collapses: the first full scored autonomy year, 2018, had 78 % zeros, and the first rewritten science year, 2024, dropped the share above 80 from 42 % to 24 % and produced one team at 90 or above against 13 in 2026. A first build that meets a discontinuity in its rookie year has no recovery year.

Further reading

- What separates the top from the rest measures the score gap and the cost of a zero.
- How rovers fail collects the failure record that produces the zero shares.
- Mass and power turns the frozen envelope into a mass and energy allocation.
- Rule history, 2007 to 2027 carries the full year-by-year rule record.
- How the missions evolved explains the family mapping used in the tables above.
- What 2027 invalidates works each 2027 change through to the design decisions it forces.
- 70 risks puts the failure modes of this page into a ranked register.

CROSS-CUTTING ANALYSIS · CHAPTER 5.6

Ten system-level conclusions

These ten conclusions hold across the rules, the scoring record, the mass and power model, the interface matrix and the failure record. Each one is stated with the number that drives it and a link to the page where the working is set out. They are ordered by the decision they settle, from the envelope that must be fixed first to the measurement that closes the plan, and a first build can adopt all ten before any detail design starts.

500

POINTS ON THE TABLE

Four missions at 100 plus the SAR

2027

S-URC27-RULES

50 kg

MASS LIMIT

Unchanged from 2007 to 2027

2027

S-URC27-RULES

1.2 m

CUBE AT WEIGH-IN

40 % of points lost above it

2027

S-URC27-RULES

82.6 %

VARIANCE IN THREE MISSIONS

Autonomy, servicing and delivery

2022–2026

S-URC-RESULTS

The four that carry the most weight

- The envelope is the design target: **50 kg**, a **1.2 m cube** and a **70 kg** fielded limit. The 2027 task list is the part that churns.
- One zero costs about **11 places and 85 points**. An Equipment zero multiplies the chance of a Delivery zero by **2.9×** and of an Autonomy zero by **2.1×**.
- The design point is **45–46 kg with 4 kg in reserve**, and a pack of **289 Wh** of nameplate for the back-to-back pair, with **400–500 Wh** recommended.
- The first competitive target is the median of **224–258 out of 500**. Lifting each hard mission by **20 points** beats any plausible review improvement.

1. The design is set against the frozen envelope, not the 2027 task list

The 50 kg mass limit, the 1.2 m cube and the 70 kg fielded limit have not moved in 8 to 20 years, while the scored content inside them has been rewritten repeatedly. The envelope is therefore the part of the requirement that can be trusted through 2030, and the task list is the part most likely to differ by the time the rover is built S-URC27-RULES S-URC19-RULES §2.a §2.a.ii.

The practical allocation follows: the envelope gets the design margin, and the new content gets the prototype. Mass, volume, power and thermal margin are held against the limits, not against the 2027 objects, and the objects are met with a modular end effector and a modular science bay.

2. The team records the capabilities it will not attempt

2027 is the largest single-year rewrite since the four-mission format settled. It adds a mandatory bioassay, astronaut assistance, 1 m route-finding in hills and a key-and-lock manipulation task in the same rulebook S-URC27-RULES S-URC27-RULES S-URC27-RULES §1.b.v §1.d.ii §1.e. A first build cannot integrate four unproven capability families in the weeks it should be testing one.

A dated list of deferrals, with the scoring cost of each one, is the cheapest risk control in the project. The deferral is a decision recorded in advance, not a capability quietly dropped in March. The scoring cost is taken from the published point values where they exist, and from the mission structure where the 2027 task weights are not yet published: the bioassay and the astronaut sub-mission each sit inside a 100-point mission, and the drone costs 5 kg of the 50 kg limit.

3. The design target is 45–46 kg with 4 kg of reserve

Four of the five champions publish no mass at all, as Mass and power records. WVU won in 2023 at 33 kg S-WVU-SAR2023, and the winning totals in the 200 to 425 band are reached by both the lightest rover and the 42 to 50 kg cluster, so mass does not

predict score within the observed range. What mass does is set the failure modes: a rover at 49.5 kg has no room for the last ideas, and the last 3 kg of a first build arrive in the final six weeks.

The reserve is 4 kg at a target of 45 to 46 kg. The 1.2 m cube is the higher-stakes limit of the two, because 40 % of the points is worth the same as 8 kg of overweight at 5 % per kilogram. Keeping the cube is a packaging plan rather than a hope: a folding or removable mast, a removable science module and a full-size mock box that is checked before every configuration change.

---

4. The pack is sized against the 60-minute Delivery and the back-to-back pair

The duty-cycle arithmetic in Mass and power gives 273 Wh of nameplate for a 60-minute Delivery and 289 Wh for the Autonomy-then-Servicing pair, which is the true single-pack requirement. The recommendation is 400 to 500 Wh of nameplate at 8S Li-ion 21700 in two packs rather than one large one, with 4P good for 180 A comfortably.

The protection chain is ordered, not improvised: controller current limits below the battery management trip, below the fuse curve, below the wire ampacity, below the connector rating. A pack that sagged under peak current would take the arm, the drive and the link down together, which is the failure class the rest of this list is written to avoid.

---

5. The arm's mass class is fixed at concept, because it silently chooses the drivetrain

A 4.4, 9.5 or 14 kg arm puts the rover at 30.6, 35.7 or 40.3 kg with the rest of the design held constant. The AK80-9 class of actuator runs out at 40.6 kg, so a single bracket choice can force a +2.7 kg step up in actuator mass. The arm is a mobility decision made in another room.

That is why the arm's mass class is fixed at concept and carried in the mass allocation from the first week, rather than discovered when the arm is designed. The arm page and the mobility page carry the two sides of the same number.

---

6. The wheel diameter comes from the box first, then the wheelbase from the terrain, then the motor from the torque

A 250 to 300 mm wheel with a 0.85 to 0.95 m wheelbase and no fold leaves 40 to 50 mm per side for suspension links and tolerances inside the 1.2 m box. A 330 mm wheel already fails the 1.2 m length at a 0.90 m wheelbase `S-URC27-RULES §2.a.ii`.

The order matters because each step constrains the next. The box fixes the wheel and wheelbase, the terrain fixes the suspension travel and the tread, and only then does the climb torque fix the motor. A rover designed in the reverse order arrives at a motor with ample torque and a chassis that cannot be stowed.

---

7. The freeze order follows the penalty, not the alphabet

The order of freezes is a risk decision. From most to least expensive to change late: mass and stow, wheel diameter, arm mount, bus voltage, E-stop, spectrum licence, camera count, thermal architecture, mast, autonomy schema, assay pair and configuration plan. The top of that list is a rule breach if it moves after the weigh-in; the bottom is a re-test.

Freezing in that order each time a design review closes keeps the cheap decisions open for as long as possible and the expensive ones closed early. The first 90 days sequence applies the same order to the schedule.

---

8. The design covers the four load-bearing single points, and the transport leg is rehearsed

Four failures zero several missions at once: wheel retention, the connector in the safety path, pack sag under peak current and the arm's joint zero. The dependency between missions is measurable: an Equipment zero multiplies the chance of a Delivery zero by 2.9× and an Autonomy zero by 2.1×.

How rovers fail collects the evidence behind those four. The cheapest reliability test is a 30-minute packaged shake test plus a post-arrival checklist, run on the crate that is also the weigh-in box.

---

9. Comms and autonomy are one reliability problem

Obstructed range should be expected at about one-fifth of line of sight: WVU measured 850 m falling to 300 m in 2023 and 750 m falling to 164 m in 2026 `S-WVU-SAR2023` `S-WVU-SAR2026`. Two bands, a 2 m mast and a directional base antenna are the floor, not a premium option.

Autonomy points are then bought with one simple planner, a progress checker and ten cold-boot full-length runs. The cleanest natural experiment in the record went 0 → 14 → 86 across three seasons while the autonomy stack was simplified, and the reading is that completion rather than sophistication produced the gain `S-MONASH-SAR2026`. The communications page and the software and autonomy page work the link and the planner as one problem.

10. The first target is the median, with one person owning the coupling numbers

The median finalist has scored between 224 and 258 out of 500 for five years, and lifting each hard mission by 20 points beats any plausible review improvement. The median is the first target because it is the profile of a team whose four missions all work; the winner's profile adds capability on top of that floor.

Six numbers decide the couplings, and no team in the published record reports all six: arm mass, wheel diameter, mast height, camera count, pack watt-hours and configuration swap time. They should have one owner and one page. One number is worth measuring above the rest: a 30-minute discharge log, which no team publishes.

What would change this list

Four findings would move a conclusion. A 2027 season of data would test the scoring claims in conclusions 2, 3 and 10. A per-task autonomy scoresheet would convert the zero rates into a failure taxonomy. Published masses for the bottom half of the field would turn the mass association into a testable relationship. A team-published reliability metric, such as missions completed without intervention, would replace the zero count as the primary measure of the thing this list is really about.

Further reading

- What separates the top from the rest carries the scoring evidence for conclusions 2, 8 and 10.
- Where the competition is going carries the rule history behind conclusions 1 and 2.
- Mass and power carries the model behind conclusions 3, 4 and 5.
- Interfaces and couplings works the chains that conclusions 5, 6 and 9 depend on.
- How rovers fail carries the failure record behind conclusion 8.
- The first 90 days turns the freeze order into a schedule.
- Mass and power data holds the 118-row table behind the mass conclusions.

MISSION BACKFIT · CHAPTER 6

Mission-by-mission backfit

URC 2027 is four field missions and one written gate, and each of the five rewards a different kind of work. These pages take the tasks one at a time: what each physically demands, what the strongest teams did on the equivalent task between 2022 and 2026, whether the 2027 change invalidates that approach, and which subsystem options serve it. One warning governs the whole part. Task weights are published only for Autonomy, so every other point figure below is either an exact Autonomy value or an explicit statement that the weight is not published.

500

TOTAL POINTS

Four missions at 100 plus the SAR

2027

S-URC27-RULES

162

FINALS TEAM-YEARS

The sample behind every score shape here

2022–2026

S-URC22-SCORES

41 %

AUTONOMY ZERO RATE

67 exact zeros in 162 team-years

2022–2026

S-URC22-SCORES

30.6 %

AUTONOMY SHARE OF VARIANCE

Against 27.7 % Equipment, 24.3 % Delivery and 14.4 % Science

2022–2026

S-URC22-SCORES

Key takeaways

- **The four missions have four score shapes, and the shape matters more than the mean.** Astrobiology has never produced a zero in 162 team-years. Autonomy produced 67.
- **Variance, not average, decides rank.** Autonomy explains **30.6 %** of finals-total variance, Equipment 27.7 %, Delivery 24.3 %, Science 14.4 % and the SAR 2.8 % S-URC22-SCORES S-URC26-SCORES .
- **The cheapest points in the competition are in Autonomy.** A team suit with a camera and microphone on a C2 display is **5 exact points** for a coverall and a video link S-URC27-RULES §1.e.iv.
- **2027 invalidates more payload than robotics.** The bioassay definition, the key-and-lock and hose hardware, and the commanded astronaut interaction are the three genuinely new requirements. The rest carries forward with a raised bar.

The four missions have four score shapes

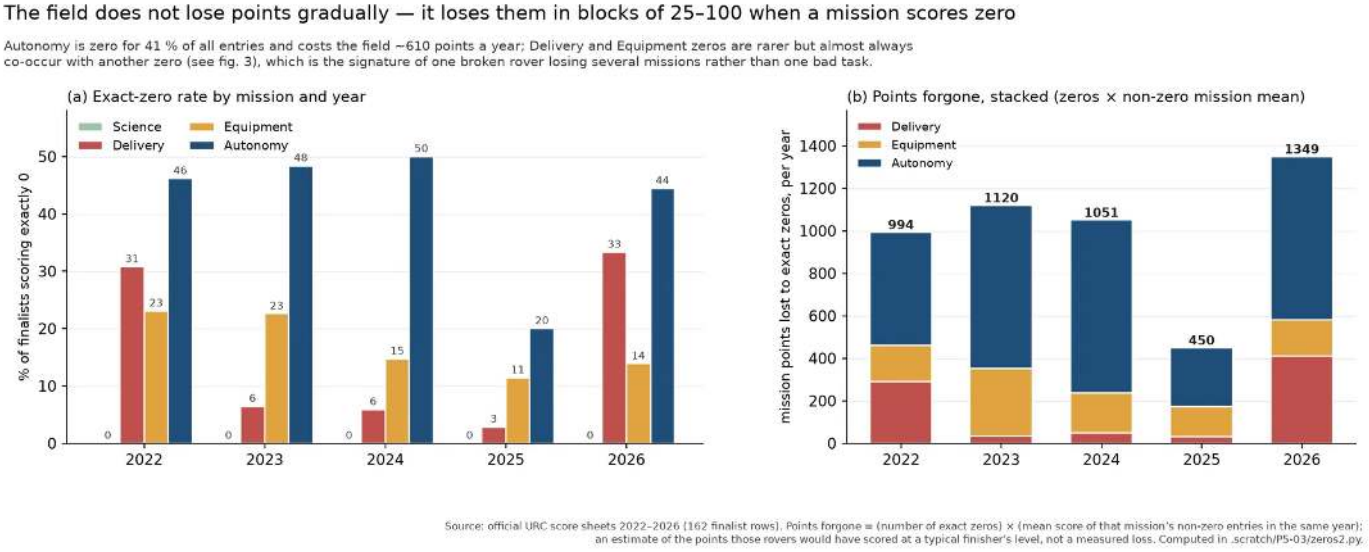
Across 162 finals team-years from 2022 to 2026, the official score sheets describe not which mission is hardest but how each one distributes its points. Astrobiology is a high-floor, low-ceiling mission. Delivery and Equipment Servicing give broad partial credit but a thin top. Autonomy is close to all-or-nothing.

MISSION	MEAN	MEDIAN	EXACT ZEROS	SCORES ≥80	STANDARD DEVIATION	SHAPE
Astrobiology (Science)	76.9	78.5	0	46 %	16.4	High floor, low ceiling
Delivery	23.6	15.0	25 (15 %)	4 %	22.3	Bimodal, gate-structured
Equipment Servicing	30.3	25.4	27 (17 %)	5 %	24.7	Broad partial credit, thin top
Autonomy (Autonomous Navigation)	26.7	22.4	67 (41 %)	7 %	27.8	All-or-nothing
SAR	88.9	89.3	0	86 %	4.5	Compressed, high stakes

Pooled 2022–2026 finalist scores, derived from the five official result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. *A zero is not always a capability zero: it can be a weigh-in penalty, a no-show or a mission the rover never reached, and the score sheet cannot separate those.*

Three readings change how a first build is planned.

1. **Astrobiology is where a team's process pays first.** No team in 162 team-years scored zero, and the median is about 79 %. A rover that reaches a site, documents it and gives an honest debrief earns credit even when sampling fails, because the debrief happens regardless `S-URC27-RULES` §1.b.x. The 2027 bioassay is the one new technical gate in an otherwise process-driven mission.
2. **Delivery and Autonomy are where seasons are lost.** Delivery has a Stage-1 gate that turns a slow start into 15/100; in 2023 **14 of 31 teams scored exactly 15** `S-URC23-SCORES`. Autonomy is worse: 41 % of all scores are zero, because autonomy either works or does not.
3. **The SAR is the least discriminating score and the most important gate.** A standard deviation of 4.5 against 16–28 for the missions means a good package cannot win the competition, but a weak one means no field at all, because only the top 36 are invited `S-URC27-RULES` §3.a.iii.



The zero rate is the clearest single picture of mission risk. Science never touches zero, Delivery and Equipment Servicing lose between 3 % and 33 % of their entries to zeros depending on the year, and Autonomy sits between 20 % and 50 % in every season, with sixteen finals teams on zero in 2026 `S-URC22-SCORES` `S-URC26-SCORES`.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC22-SCORES` `S-URC26-SCORES` *Full resolution*

Caution on 2026 Autonomy

Five teams scored **exactly 86.0** in 2026 and nobody scored higher `S-URC26-SCORES`. That is a rubric ceiling, not 86 % capability, so the 2026 Autonomy column ranks teams but does not measure how much of the task a team could complete.

The master table

Points are exact only in Autonomy; where a 2027 weight is unpublished, the cell says so. Best past evidence cites the highest-scoring team-year observed and, where it matters, the failure that shows the risk is real. The rule or Q&A clause is cited wherever a row turns on a specific clause.

TASK	2027 POINTS	REQUIRED CAPABILITY	BEST PAST EVIDENCE	2027 DESIGN IMPLICATION	MAIN RISK
A1 Science Plan	not published	About 5 pages, with background science [AQ1]	Every high-scoring Astrobiology team files one	One owner; due 2027-05-21	The plan owner is then ineligible for the debrief (AQ8)
A2 Two or more sites documented	not published	Panorama with cardinal directions and scale, close-up, GNSS with elevation and accuracy	West Virginia 2025 built this workflow into its science plan	Automated north arrow and scale overlay; a heading source on board	Missing cardinal or scale reference
A3 Sample of 10 cm and 5 g	not published	A 10 cm stroke, depth evidence, 5 g from below 10 cm, hard soil	West Virginia 30 cm drill; UIU auger; Monash depth-confirmed 2026	Body-mounted drill on a linear stage, with a scoop fallback	One rock ends the sample (KARURA 2024)
A4 Two life detections, one bioassay	not published	Two orthogonal biomolecule channels S-URC27-QA AQ9	UIU 2025 Bradford protein (100); West Virginia 2023 chlorophyll (100); Missouri S&T 2022 DNA and protein (96)	Two assays on one detector, with the AQ9 justification written	A 94-point 2026 payload fails AQ9 unchanged
A5 Sealed cache and 5-minute handover	not published	No spill when inverted, removable by hand, openable by judges	Unchanged since 2024; no documented URC failure	50 mL screw-cap in a holster, inversion-tested	Spill; cache trapped by the arm
A6 Debrief	not published	Results, depth and contamination, site reasoning, life and habitability, Mars	KARURA 2024 salvaged 74 after losing its sampler	Two staff: a C2-only operator and a debrief speaker	The speaker leaves C2 (AQ8)
D1 Stage-1 gate	not published; 2026 proxy 40/100	Two or three fast tasks banked early	2023: 14 of 31 teams at exactly 15	Sequenced by points per minute; the riskiest task does not go first	A 60/100 becomes 15/100
D2 Open a hinged-lid box	not published	Hook or pinch; latch force not measured	Routine where the object comes to the end effector	Non-slip hook; wrist camera on the seam	Lid springs shut; grip slips
D3 Pick, carry and deliver 5 kg	not published	About 25 N·m at the shoulder for 5 kg at 0.5 m, estimated; retention over rough ground	West Virginia 2023 8 kg lift for 90; Missouri S&T 2026 Delivery 100	100–110 mm jaw plus a hook or pin; carry-pose lock	Dropped object; centre of gravity outside the jaw
D4 Read a sign	not published	Character or ArUco recognition at range in sun	Sign reading was a Delivery task 2017–2026	Zoom lens or drone; a team-owned ArUco fallback	Glare, distance, angle
D5 Search a large or poor-radio area	not published	Sweep pattern, detection at range, link survival	West Virginia 2023 found the target with 90 s left (90); Monash 2023 scored 15	Sweep, detection range and abandon rule fixed before the mission	Search eats the clock; the detector fails outdoors
D6 Traverse 1 km with more than half out of sight	not published	1.0–1.5 m/s on flat gravel, 0.4–0.6 m/s on 25–30°, by inference; a link in terrain shadow	West Virginia 2025 at 1.6 m/s for Delivery 80; BRACU comms rebuild to 61	Passive six-wheel with printed wheels; a 900 MHz recovery band	Stuck, rollover, link loss
E1 Tube into cache, lid, lock	not published	Pinch, insert, close, rotate	ARES 2025 banked several Equipment tasks	A cache designed for one approach pose	Tube dropped; lock torque
E2 Cache transport	not published	A 5 kg grasp and a 100 m drive	Missouri S&T 2026 Equipment 100; ARES 10.4 kg lift	Stowed for the drive, or a vibration-proof grasp	Drop during the traverse
E3 Drawer open, insert and close	not published	One-axis pull; sub-centimetre insertion	Monash 2022 opened at 32:08 and closed at 32:52	Compliance or vision in the insertion axis	Friction jam; rover pulled
E4 Latch and hinged panel	not published	Release, then hold the panel open	Monash 2022 lost about 9 minutes to two panel closures	An end effector and pose that prop the panel	Panel closes mid-task
E5 Autonomous typing	not published	Few-millimetre repeatability, about 0.44 N over 2 mm, 2 cm ArUco, no auto-repeat	Missouri S&T 2026 100 requires it; BYU ViSP; ARES YOLO; UIU 72.63	Two-stage visual servo and a compliant tip, with each key confirmed	Backlash of 4.4–5.2 mm at 1 m exceeds the target
E6 Key in a tubular lock	not published	Grasp a flat key, align to a 1–2 mm keyway with a 1 cm tag in view, then turn	No URC precedent	Chamfered funnel end effector; wrist camera; a tethered key	A dropped key ends the task
E7 Cam-lock hose, half gasket	not published	Axial push with compliance and alignment	BYU listed a hose connection in 2026, on hardware without a gasket	An exact half-gasket mock, with force measured against depth	Force and travel unmeasured; jams

TASK	2027 POINTS	REQUIRED CAPABILITY	BEST PAST EVIDENCE	2027 DESIGN IMPLICATION	MAIN RISK
E8 Quarter-turn valve	not published	90° rotation, torque not measured	Present 2022–2026; routine	A cheap early score	Lever slips; reach blocked
E9 Buttons, switches, knobs	not published	Push, flip, rotate; no manual keyboard contact	Monash 2022 completed all of them	Lowest risk, sequenced first	Wrong button; tip slides
F1 Suit with camera and microphone	**5**	Team suit and a live feed on a C2 display	No precedent for a monitored suit feed	Coverall, cooling vest and an analogue 5.8 GHz pod	Feed not on screen means judged absent
F2 Drive to the astronaut within 3 m	**5**	Autonomous GNSS waypoint and stop	2011–2013 required delivery within 1 m from a fix of ±25 m	A stop at 1.5 m internally, with visual acquisition for the last 25 m	Stopping short; base offset
F3 Follow!	**15 max**: device 5, sign 5, audio 10, gesture 15	Person tracking on natural terrain, stopping within 3 m	No precedent for a commanded follow	A fallback ladder, so the worst case is 5 and the best 15	A false positive drives at a person
F4 Stay! beyond 20 m	**5**	A latched hold with no creep	No precedent; nearest is a waypoint hold	Planner disabled in the hold; encoder zero-motion check	Drift; a re-plan moves the rover
F5 Fetch! the hammer	**10**	Detect and grasp a 0.91 kg, 300 mm steel hammer; teleoperation scores zero	Missouri S&T trained a `RockPick` class; West Virginia reached the 80th percentile of mAP across four implement classes	Wide jaw or a hook; lift 100 mm before moving	Low-profile grasp; drag on retract
F6 Come! within 3 m	**5**	Drive to a person at least 20 m away and in the field of view `S-URC27-QA` AUQ7	No precedent, but the geometry is guaranteed	Person detection or a suit-carried GNSS beacon	Person not detected at 20 m
F7 Give!	**5**	Place or drop at ground level	No precedent	Release at 10 cm and confirm on the ground	Still in the gripper at time-out
R1 Target 1, on the hill, within 1 m	**25**	A hill route where not every approach is navigable, a 1 m stop and an arrival signal	Five teams at the 86 ceiling in 2026; Missouri S&T used an elevation map and LiDAR; Monash used navfn	Slope costmap; terminal approach at 0.3 m/s; stop at 0.6 × the tolerance	Overshoot; the planner refuses the gap
R2 Target 2, out of radio sight, within 1 m	**25**	Autonomous continuation with no link; an arrival LED	No documented intentional out-of-sight run in 2022–2026	Keep planning on board; never gate the drive on the radio link	One link-status check costs 25 points
R3 Penalties and arrival LED	−20 %, −50 %, cap 50 %, physical uncapped	A daylight LED on red, blue and green; hands-off discipline	An LED and hands-off operation were already required in 2026	LED on the authoritative mode topic, with a sunshade	A judge misses red; a reflex touch
S1–S3 SAR report, video and gate	100 plus the top-36 gate	Evidence for each mission and about one page of science plan [AQ1]	The cut was decided inside a ±3-point band in 2023–2025	The report maps to the four missions and shows test data	A good rover that never reaches the field

Where the points actually are

The ranking below is an engineering judgement built from the rule numbers and the score distributions. It is not a rule quantity, and where the 2027 weight is unpublished the entry says so.

RANK	CAPABILITY	WORTH	RISK
1	EVA suit with camera and microphone	5 exact	Very low
2	Short autonomous behaviours: drive, Stay, Come, Give	20 exact	Low to medium
3	Astrobiology documentation and debrief	not published; historically the difference between about 40 and about 80	Low
4	Equipment cheap tasks: buttons, switch, knob, valve, latch	not published	Low
5	Follow! by sign or device	5 exact	Low
6	Fetch!	10 exact	Medium
7	Autonomous typing	not published	Medium to high
8	Route targets at 1 m	25 each, exact	High
9	Gesture Follow!	+10 over the device modality	High
10	Key in lock	not published	Very high
11	Cam-lock hose with a half gasket	not published	High

The sequencing consequence is the largest schedule risk in this part. Two of the three genuinely new 2027 requirements, the bioassay payload and the astronaut interaction, must be demonstrated in the SAR by **2027-03-05**, about three months before mission rehearsals can begin. The robotics can improve until June; the assay chemistry, a suit and an operator-training dataset cannot. The schedule must front-load the two least-familiar subsystems and back-load the most familiar ones, which is the opposite of the natural build order.

> **The ordering that follows**
>
> The guaranteed-floor work comes first: documentation, sample, cache and debrief in Astrobiology, and the suit plus the four short autonomous behaviours in Autonomy. That floor is not traded for the high-ceiling work of 1 m route targets, gesture Follow! or the key in the lock.

The mission pages

- Astrobiology mission: the 20-minute clock, the mandatory bioassay and the payload that a 2026 score of 94 would fail.
- Delivery mission: the staged gate, the 5 kg object over 1 km and the task list still to come.
- Equipment servicing: nine task families, millimetre precision and the two hard tasks that change in 2027.
- Autonomy: astronaut assistance: fifty points, five commands and a fallback ladder that spans 5 to 15 points.
- Autonomy: route finding: two targets, one deliberately out of radio sight and a 1 m stop.
- System Acceptance Review: 100 points and the only route to the other 400.
- What 2027 invalidates: the synthesis of what changes, what carries forward and what the schedule has to do about it.

Further reading

- Missions and scoring carries the rule clauses, task lists and deadlines behind every row above.
- What separates the top from the rest derives the variance shares and zero rates used on this page.
- 110 engineering problems turns each design implication in the master table into a task with an acceptance test.
- Mass and power sizes the rover that carries all four missions.

MISSION BACKFIT · CHAPTER 6.1

Astrobiology mission

Astrobiology is the only mission where the science is the deliverable and the rover is the instrument platform. The physical envelope is modest: a sample from at least 10 cm below the surface, of at least 5 g. The evidence envelope is not. Two documented sites, two life-detection capabilities, one of them a bioassay, and an honest debrief have to fit inside a 20-minute worst case.

100

MISSION POINTS

The field has never scored a zero on this mission

2022–2026

`S-URC22-SCORES`

20–30 min

ROVING AND ANALYSIS

Plus 5 minutes to remove the cache after the clock

2027

`S-URC27-RULES`

10 cm, 5 g

MINIMUM SAMPLE DEPTH AND MASS

Both measured from below the surface

2027

`S-URC27-RULES`

2

LIFE DETECTIONS REQUIRED

At least one must be a bioassay

2027

`S-URC27-QA`

Key takeaways

- **Mass is easy; the evidence is the mission.** 5 g of soil is about **3–4 cm^3** at a bulk density of 1.3–1.6 g/cm^3, roughly a fifth of a 16 mm auger filled to 10 cm. Proving depth, avoiding contamination and running two assays inside the clock is the hard part.
- **The clock decides the architecture.** A planning model built from the rule's own sequence leaves **about 2.5 minutes of slack** in the 20-minute case. Assays must therefore run while the rover drives, which makes the payload body-mounted and self-sequencing.
- **A 94-point 2026 payload would fail the 2027 clause without a hardware change**, because amino-acid chemistry alone does not meet the bioassay definition `S-URC27-QA` `S-WVU-2026-SCIPLAN` AQ9. The payload that scored 96 in 2022 was closer to the 2027 rule than the spectrometer that replaced it at Missouri S&T.
- **This is the field's strongest mission.** The median has stayed between 73 and 85 every year, a 100 has been achievable every year, and no team in 162 team-years has scored zero.

What the mission physically demands

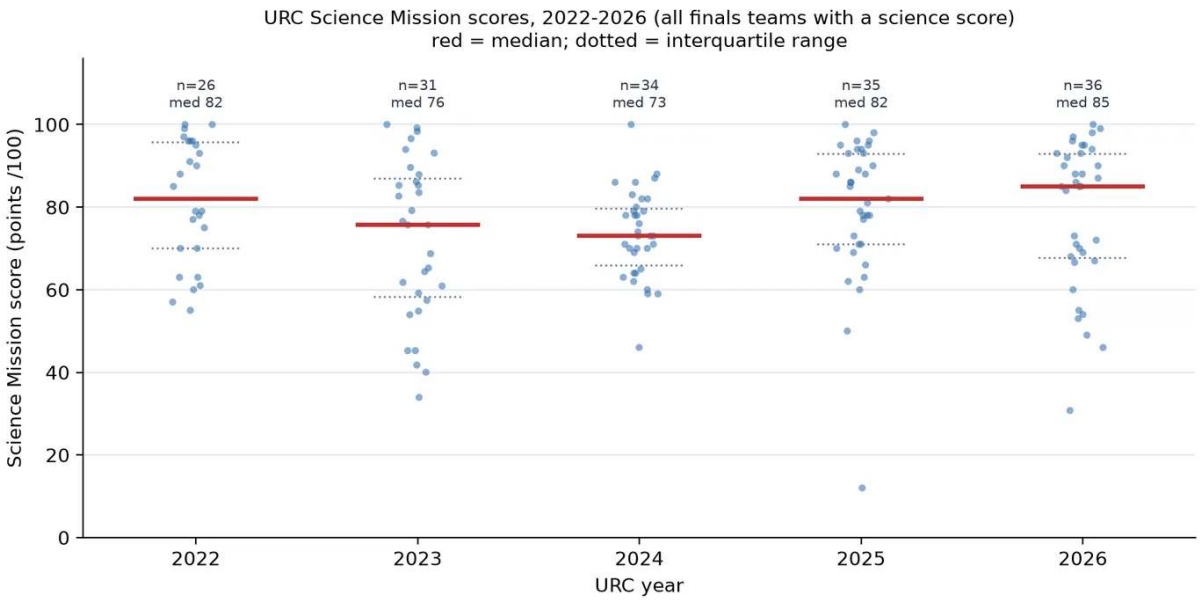
DEMAND	NUMBER	BASIS
Depth	At least 10 cm below the surface, from a single sample; the portion below 10 cm sets the mass	S-URC27-RULES §1.b.vi
Mass	At least 5 g from below 10 cm	S-URC27-RULES §1.b.vi
Volume implied	About 3–4 cm^3 of soil at 1.3–1.6 g/cm^3	Arithmetic, by inference
Soil resistance	From reasonably easy to very hard, with no caliche; no force value is published	S-URC27-QA AQ6
Downforce implied	About 30–80 N of downforce and 1–3 N·m of torque for a 16–20 mm auger in medium clay	Standard soil mechanics, by inference; the chassis carries the reaction, not a compliant wrist
Sample point	Sites are not marked beyond the overall site, and the soil is in situ with no trays	S-URC27-QA AQ5
Life detection	Two capabilities, at least one a bioassay that detects a molecule formed only by biology or in its presence	S-URC27-RULES S-URC27-QA §1.b.v AQ9
Documentation	Panorama with cardinal directions and scale, close-up at high resolution with scale, GNSS with elevation and an accuracy range	S-URC27-RULES §1.b.iii
Time and radius	20–30 minutes inside a 0.5 km radius, plus 5 minutes for cache removal	S-URC27-RULES §1.b.ii §1.b.vii
Containment	All chemicals, water and reaction products contained; hazardous chemicals pre-approved with a transport, use, disposal and accident plan	S-URC27-RULES S-URC27-QA §1.b.ix AQ4
Perception	Site selection from rover data only; the operators cannot see the site	S-URC27-RULES §2.c.i §3.d.i

The 20-minute case

The planning model below is constructed from the rule's sequence. It is a planning aid, not a rule.

STEP	SITE 1	SITE 2	NOTE
Drive and select	3 min	3 min	100–200 m between sites inside a 0.5 km radius
Document	2 min	2 min	Automate the overlay and the metadata capture
Acquire a sample below 10 cm	3 min	none, the cache comes from one site	A 3-minute allowance is optimistic in very hard soil
Load and start assay 1	1.5 min	none	The slowest assay starts first and runs during the traverse
Load and start assay 2	1 min	1 min	Both results must be read before the clock ends
Stow and seal the cache	1 min	none	The 5-minute removal is after the clock
Sub-total	**11.5 min**	**6 min**	**17.5 of 20 minutes, leaving 2.5 minutes**

A design in which the arm must hold a cuvette in front of a camera for three minutes has already lost the mission. The payload must be a body-mounted fluidic system that a pump or carousel can drive without the arm, with a readout the C2 can interpret while the rover moves.



Every finals science score from 2022 to 2026. The median has never fallen below 73 and the lower tail never touches zero, which is why this mission rewards documentation discipline before it rewards instrumentation.

analysis, 2026. `S-URC22-SCORES` `S-URC26-SCORES` *Full resolution*

What the best teams did, and what 2027 does to it

Astrobiology is the field's strongest mission: a 100 has been achieved every year, and the top ten sit 9–23 points above the rest of the field.

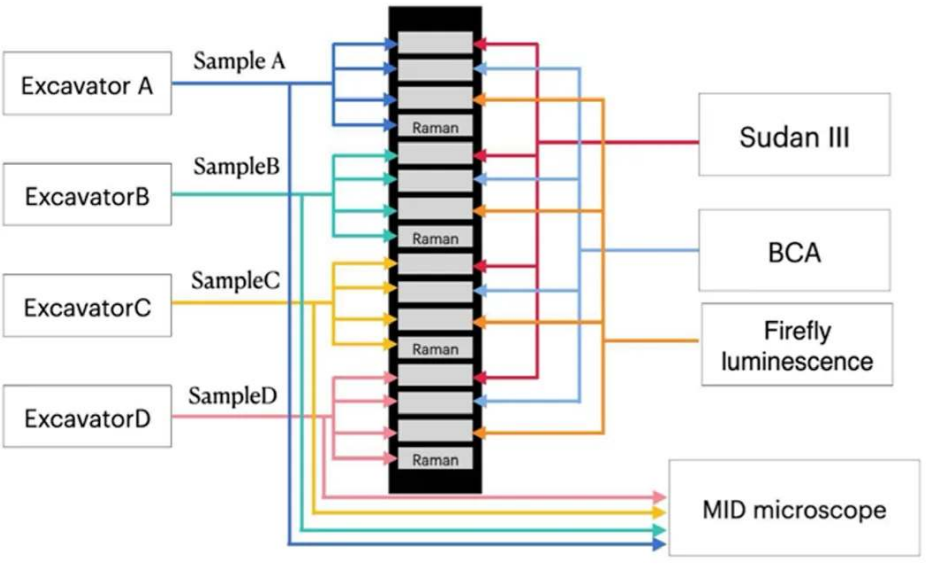
YEAR	TEAMS	MEDIAN	TEAMS ≥90	MAXIMUM	TOP-10 MEAN	REST MEAN
2022	26	82.0	11	100	95.4	72.4
2023	31	75.7	6	100	83.2	67.1
2024	34	73.0	1	100	80.1	70.7
2025	35	82.0	11	100	92.1	74.6
2026	36	85.0	13	100	94.0	72.5

Derived from the five official result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES` *. The 100s went to AGH and Yonder Dynamics in 2022, West Virginia in 2023, BYU in 2024, UIU in 2025 and Monash in 2026.*

The payload record is the more useful table, because it converts a rule change into a list of approaches that survive it.

TEAM, YEAR	DOCUMENTED PAYLOAD	SCORE	UNDER THE 2027 CLAUSE
Missouri S&T 2022	DPA for DNA, propidium iodide for nucleic acids, biuret for protein	**96**	Compliant, and the most 2027-ready payload in the window
West Virginia 2023	An STS spectrometer covering the visible range, plus a TMB peroxidase assay; chlorophyll, carbonate and a “no life” result	**100**	Compliant, and a negative result still scored 100
KARURA 2024	ATP by luciferase, BCA protein, lipid, Raman and a microscope	74	Compliant, and the best public contamination design
Missouri S&T 2024	A Raman spectrometer replacing all chemical tests	82	Not compliant as the sole method
UIU 2025	Deep auger, spectrophotometric tests and a Bradford protein assay	**100**	Compliant, and the best-documented template
Monash 2026	Near-infrared and wet-chemistry UV-Vis, including BCA protein	**100**	Compliant if the protein channel is kept
West Virginia 2026	Ninhydrin as the primary life detection, with resazurin and methylene blue	**94**	Not compliant: AQ9 excludes amino acids alone

The decisive finding. A payload that scored **94/100 in 2026 would fail the 2027 clause without any hardware change**, because amino-acid chemistry alone does not meet the bioassay definition `S-URC27-QA` `S-WVU-2026-SCIPLAN` AQ9. Conversely, the payload that scored **96 in 2022** was closer to the 2027 rule than the Raman instrument that replaced it. Science scores from 2022 to 2026 therefore transfer as evidence of documentation, sampling and debrief skill, and not as evidence of payload compliance.



KARURA routed four independent excavators to separate test-tube sets for Raman, lipid, protein and ATP assays. The separation is the contamination control, and it is the design that best matches what the 2027 clause asks for `S-KARURA-2024-PAPER`.

*KARURA (Hirai et al.), *Planetary Science Society of Japan* journal, Vol. 33 No. 3, 2024, 2024.*

Original source *Full resolution*

What 2027 removes is the stratigraphic profile, from both documentation and debrief `S-URC26-RULES` `S-URC27-RULES` §1.b.iii §1.b.iii. That frees one to three minutes of site time and one debrief bullet. It does not devalue the rest of the documentation, and it raises the bar on the biology, because there is less geology to hide behind.

Which subsystem options serve it

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Payload	**Two wet biomolecule assays sharing one detector**: ATP and BCA protein on one photodiode or spectrometer, about 3.5–4.5 kg by inference	One wet and one optical assay (ATP with autofluorescence), only with a written answer from the judges on the second method's status
Sampler	**A body-mounted auger on a 200 mm linear stage**, the West Virginia and UIU approach	An arm-held scoop, as the backup sample path only
Assay timing	**Room-temperature rapid chemistry**, such as the 5-minute BCA variant	A 37 °C incubation only if a heater is already aboard, because pre-heating is what KARURA had to do
Depth proof	**A graduated clear liner with a measured stroke**	A depth camera, if it can resolve the soil column through dust
Heading	**Dual-antenna GNSS** for exact cardinal annotation	A calibrated magnetometer, allowed, for ±2–3°
Operations	**Two staff**: a C2-only science operator and a debrief speaker	One person, only if that person never leaves C2 (AQ8)

Minimum viable versus stretch

- **The minimum that scores:** reach a site, relay a panorama and a close-up with a scale, record GNSS with elevation and accuracy, and debrief honestly. The historical cluster for this behaviour is a 40–55 band, by inference.
- **The minimum that is also compliant:** two life-detection capabilities, one of them a bioassay. That is a mandatory requirement rather than a scoring option, so a rover with no bioassay is non-compliant however well it documents.
- **The stretch:** two sites analysed, two biomolecule results each, depth-resolved provenance, a written contamination justification, and a quantitative Mars-applicability argument in the debrief.

Risks to retire early

RISK	WHY IT IS REAL	THE TEST THAT RETIRES IT	LEAD TIME
The assay chemistry does not survive the desert	LAMP master mix ships at –20 °C and a rover has no freezer S-NEB-LAMP	Five days at 38 °C, then run the assay; signal at 80 % of day 0 or better	3 months
The assay cannot finish in 20 minutes	KARURA's BCA needed 30 minutes at 37 °C and had to pre-heat S-KARURA-2024-PAPER	A timed dry run of the full sequence; both results logged by T+17:00	6 months
False positives	Extracellular ATP, humic interference with BCA, acidic soil turning a dye yellow S-NEB-LAMP	Sterile-soil blank, a method blank and standard-addition recovery	3 months
Sample too shallow or too light	The 10 cm depth and 5 g mass limits are single-point failures, and hard soil is confirmed S-URC27-QA AQ6	Drill three analog soils and photograph the liner; 5 g below 10 cm in under 4 minutes	4 months
One rock ends the sample	KARURA 2024 lost its sampling mechanism to a rock S-KARURA-2024-PAPER	A slip clutch, a scoop fallback and rock rejection from imagery	4 months
The cache spills or jams	No documented URC incident, but the rule is explicit S-URC27-RULES §1.b.vii	Twenty inversions over a weighed pad, then a timed gloved removal	4 months
Cardinal direction or scale missing	Both are required S-URC27-RULES §1.b.iii	Automate the north arrow and scale bar, then check a blind sample of 20 images	2 months
The debrief speaker is ineligible	AQ8 limits the debrief to members who stayed in C2	Name the speaker in the operations plan and keep that person in C2	1 month

One risk sits above the table: **the payload definition itself**. It is the only item in the mission with a hard pre-SAR deadline, and every other risk in the list is cheaper to retire once it is fixed.

What is not published

- The 2027 weights of the four Astrobiology components, so no statement of the form “life detection is worth X points” is possible.
- Whether one combined assay can satisfy both life-detection capabilities; AQ9 is silent on double counting.
- The Science Plan format, which is promised for early 2027.
- Whether the Equipment Servicing cache is the team’s own cache or a judge-supplied one, so the arm interface must work for both.

Further reading

- Science and astrobiology payload works the payload options and the fluidic architecture with numbers.
- UIU, West Virginia, Monash, Missouri S&T and KARURA carry the payload evidence.
- Mass and power places a 3.5–4.5 kg payload inside the 50 kg limit.
- What 2027 invalidates sets this mission’s payload change beside the other two genuinely new requirements.

MISSION BACKFIT · CHAPTER 6.2

Delivery mission

Delivery is the mission where time, terrain and a gate interact. A 5 kg object, smaller than 40 cm on each side, travels up to a kilometre over sand, gravel, boulder fields, vertical drops and steep loose slopes, with line of sight lost for more than half of the course. The gate, not the terrain, is what breaks seasons.

100

MISSION POINTS

Staged, with a minimum Stage-1 score that gates Stage 2

2027

`S-URC27-RULES`

5 kg

OBJECT MASS

Under 40 × 40 × 40 cm, with a grasp feature of at most 7.5 cm

2027

`S-URC27-RULES`

1 km

COURSE RANGE

Line of sight is not guaranteed beyond half the course

2027

`S-URC27-QA`

15

THE SCORE 14 OF 31 TEAMS TOOK IN 2023

A slow start behind the Stage-1 gate

2023

`S-URC23-SCORES`

Key takeaways

- **The gate is the physics.** A fixed Stage-1 clock plus a minimum Stage-1 score means the *rate* at which the first tasks are banked matters more than total capability. In 2023, **14 of 31 teams scored exactly 15** `S-URC23-SCORES` .
- **The lift is cheap; the retention is not.** A 5 kg object at 0.5 m needs about **25 N·m** at the shoulder, by inference. What fails is the grip over rough ground, not the motor.
- **Speed is a mission requirement.** Two kilometres of driving at 1 m/s is **33 minutes** of a 30–60 minute mission, and at 0.5 m/s the drive alone is longer than the mission. The practical target is 1.0–1.5 m/s on flat gravel and 0.4–0.6 m/s on a 25–30° slope.
- **2027 leaves most of this valid.** The marked-path task is removed, while the object envelope, the 5 kg mass, the 1 km course, the out-of-sight requirement and the drone option are unchanged. What is missing is the task list and the Stage-1 minimum.

What the mission physically demands

DEMAND	NUMBER	BASIS
Object mass	5 kg or less per object	S-URC27-RULES §1.c.iii
Object envelope	Smaller than 40 × 40 × 40 cm	S-URC27-RULES §1.c.iii
Grasp feature	7.5 cm diameter or less, as a handle or a natural feature such as a rock edge or a shaft	S-URC27-RULES §1.c.iii
Lift moment implied	About 25 N·m static at the shoulder for 5 kg at 0.5 m	Arithmetic, by inference
Course range	Up to 1 km from the start gate; the total path may exceed 1 km	S-URC27-RULES S-URC27-QA §1.c.ii RQ13
Line of sight	Not guaranteed for more than half the course, and later tasks are deliberately out of sight	S-URC27-RULES §1.c.ii §3.d.ii
Terrain	Soft sand, gravel, rough stony areas, rock and boulder fields, vertical drops and steep loose slopes; the Q&A adds vertical faces	S-URC27-RULES S-URC27-QA §1.c.ii DQ1
Terrain limits	No numeric slope or step limit is published	S-URC27-QA AUQ17
Speed implied	2 km at 1 m/s is 33 minutes; 0.5 m/s makes the drive alone longer than the mission	Arithmetic, by inference
Placement tolerance	Not published for 2027; the 2026 rule required the object to rest within 20 cm with no rover part touching	S-URC26-DM as a 2026 reference
Stage structure	A fixed time per stage, a specified Stage-1 minimum that gates Stage 2, and leftover Stage-1 time transferred forward	S-URC27-RULES §1.c.i
Autonomy	Not required; teleoperation is legal for the whole mission	S-URC27-RULES §1.c
Drone	Allowed and pad-launched, at 5 kg or less including a dummy mass; full points are possible without one	S-URC27-RULES §2.b.i §2.b.iv–vii
Mini-rover	Prohibited	S-URC27-RULES §1.c.i



The 2027 object limit is the same 5 kg that Monash lifted here, with a bench scale inset as the evidence. A lift claim needs a measured mass, not an estimate from the CAD model S-MONASH-SAR2026 .

Monash Nova Rover, 2026. ***Full resolution***



The terrain list for Delivery is unchanged from 2026, and the rock garden is the part of it that eats wheels, links and time. The 2027 course keeps the same terrain; the stage gate and its timings are still to be published.

Mars Society, 2017. CC BY-SA 4.0. Original source Full resolution

What the best teams did

Delivery is the most volatile mission in the record: mean 23.6, median 15.0, 15 % zeros and only 4 % of scores at 80 or above. The distribution is bimodal rather than graded, because a slow start produces 0–20 and a competent run produces 50–80.

YEAR	TEAMS	MEAN	MEDIAN	EXACT 0	EXACT 15	SCORES ≥80
2022	26	25.1	22.5	8	2	2
2023	31	15.9	15.0	2	14	1
2024	34	22.9	15.0	2	5	1
2025	35	30.9	30.0	1	1	1
2026	36	22.9	18.0	12	0	1

Derived from the official mission column in the five result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. *The 2023 cluster at exactly 15 is the Stage-1 gate: one or two cheap tasks banked, then a time-out.*

TEAM	2022	2023	2024	2025	2026	THE LESSON
Missouri S&T	80	**15**	**15**	70	**100**	Three seasons to close one gap: object search and grasp inside the clock
Monash	70	**15**	80	71	76	2023 was a pure search-and-recognition failure
West Virginia	42.5	**90**	52	**80**	61	90 was double the next-best, built on an 8 kg lift and a large workspace
UIU	18.75	45	30	35	62	The 2026 jump follows a mobility and arm rebuild
Michigan	**80**	**15**	**15**	30	44	Champion in 2022, then two 15s, with no public explanation
BRACU	did not enter	15	15	22	61	15 to 61 in four seasons, through a communications rebuild

Official Delivery column, from the five result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`; BRACU's first recorded Delivery entry is 2023.

The clearest transferable lesson from the best Delivery runs is time management. West Virginia's 2023 search discipline, which covered the search area from an overhead map and kept a detection window open to the last 90 seconds, and Missouri S&T's 2026 run both came from a reliable wrist camera and a rehearsed order of tasks, not from unusual hardware `S-WVU-SAR2023` `S-MST-SITE-ATHENA`.

Does 2027 invalidate it

Mostly no, with one planning caveat.

- **Removed:** the marked-path task, "follow a marked path that requires precise maneuvering over a designated terrain feature", is absent in 2027 `S-URC26-RULES` `S-URC27-RULES` §1.c.iii §1.c.iii. That removes one path-following family. It does not reduce the terrain requirement, because reaching the objects still means crossing the same terrain list.
- **Unchanged:** the object envelope, the 5 kg mass, the 1 km range, the out-of-sight requirement, the approximate GNSS locations, scoring on the correct object at the correct location with partial credit, the drone option and the mini-rover prohibition `S-URC27-RULES` §1.c.
- **Not yet published:** the stage durations, the Stage-1 minimum and the task list, all of which the rulebook defers to a document published before the competition `S-URC27-RULES` §1.c.i. A team training to the 2026 timing of 35 minutes and 20 minutes with a 40-point gate may be training to the wrong clock `S-URC26-DM` as a 2026 reference. The capability transfers; the rehearsal timings must be set when the task list lands.

Verdict: the approach carries forward with a raised planning risk. The design implication is not a different rover but a rehearsed gate, run as a sequence of bankable tasks with an abandon rule, and re-timed once the 2027 task list is published.

Which subsystem options serve it

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Mobility	**Six-wheel passive rocker, printed wheels and skid steer**: the field default and the safest first build	Four wheels when mass binds and the terrain is mostly gravel and sand
Arm	**A two-finger 100–110 mm jaw, current-limited, with swappable tips**, plus a hook or pin	A dedicated cargo hook when the object's centre of gravity is unpredictable
Communications	**2.4 GHz primary and 900 MHz recovery**, the only combination that survives intentional loss of sight	A 5.8 GHz third band for video when most of the course is in sight
Search	A wrist or mast camera with a YOLO-class detector	A **drone** when the search area dominates the clock and a pilot and spotter pair exists
Drone	Flown only after 20 sorties with the dummy mass in wind	**Left at home**: full points are possible without one, and Missouri S&T's 2024 drone never flew
Load handling	The object stowed for the traverse, keeping the centre of gravity low	Carried in the gripper only for short moves, with a carry-pose lock

Minimum viable versus stretch

- **The minimum that scores:** reach the course and complete one Stage-1 task, such as reading a sign, opening a lid, pushing a button or moving a small object. Teleoperation is legal.
- **The minimum that passes a gate:** two or three of those banked in the first third of Stage 1, which needs a known task order. The mitigation is to rehearse a generic fast opening, then resequence when the 2027 list publishes.
- **The stretch:** every task, with the correct object at the correct location, including a 5 kg load over rock, a large-area search and a later task beyond radio line of sight. On the 2026 task set that was worth 100 to one team.

Risks to retire early

RISK	WHY IT IS REAL	THE TEST THAT RETIRES IT	LEAD TIME
Stage-1 time-out	14 of 31 teams at exactly 15 in 2023; Missouri S&T ran out of time in Stage 1 `S-URC23-SCORES` `S-MST-NEWS2023`	A timed mock Stage 1 with an abandon rule; three bankable tasks inside one third of the clock, five runs in a row	4 months
Object not found	Missouri S&T could not locate a specific rock through its cameras in 2023; Monash scored 15 on the same task `S-MST-NEWS2023` `S-MONASH-IEEE-RAM`	Detection-range test with judge-like objects at 5, 10 and 20 m in sun; a sweep that covers the area inside one third of the clock	3 months
Dropped or misplaced 5 kg object	ARES demonstrated a 10.4 kg lift, so the risk is retention rather than strength `S-ARES-SAR26`	Twenty-metre carry over rock with a 5 kg mass; no slip and a stable jaw current in 20 of 20 runs	3 months
Link loss out of sight	West Virginia measured 750 m in sight against 164 m out of sight in 2026 `S-WVU-SAR2026`	A deliberate drive behind a hill on the 900 MHz band alone, with a recovery switch under 10 seconds	4 months
Stuck or rolled rover	Judges describe rovers driven into gulleys and flipped over `S-URC-ADVICE`	Twenty traverses of a local rock field with no stuck event longer than 60 seconds	5 months
Arm fails under driving loads	Missouri S&T lost a vertical axis under driving in 2024; West Virginia had an over-torque cascade in 2023 `S-MST-NEWS2024` `S-WVU-YT-DOORFAIL`	Two hours of rough driving with the arm installed, logging strain and current	4 months
Drone not flyable on the day	Missouri S&T's 2024 drone was kept on the ground by complications `S-MST-SITE-HYPERION`	Twenty sorties with the dummy mass in wind, plus a pad-launch rehearsal	3 months
Placement outside tolerance	The 2026 rule was 20 cm with no rover part touching `S-URC26-DM`	Place-and-withdraw against a 5 cm ring, 18 of 20 runs inside the tighter of the 2026 and 2027 tolerances	3 months

What is not published

- The 2027 stage durations and the Stage-1 minimum, so a gate rehearsal cannot yet be timed.
- The 2027 task list, so the points per minute of each task cannot be ranked.
- The 2027 placement tolerance.
- Any numeric slope or step limit for the course, which means the obstacle set has to be chosen and measured locally.

Further reading

- Mobility sizes the speed and the terrain capability this mission needs.
- Arm and end-effectors covers the jaw, the retention problem and the lift.
- Delivery, payload handling and drones is the synthesis for the drone option, which has no dedicated study behind it.
- Missouri S&T, Monash, West Virginia and BRACU carry the Delivery evidence.
- What 2027 invalidates places the removed marked-path task beside the other changes.

MISSION BACKFIT · CHAPTER 6.3

Equipment servicing

Equipment Servicing is the arm's report card, and the field's weakest mission after Autonomy. Nine task families span ground level to 1.5 m, and the smallest targets in the rulebook are 2 × 2 cm ArUco tags at the corners of a vertical keyboard and 1 × 1 cm tags on a key and a keyhole. The difficulty is manipulation itself, not the task list.

100

MISSION POINTS

Nine task families in 30 minutes, in any order

2027

S-URC27-RULES

1.5 m

WORKSPACE HEIGHT

From ground level, at an unspecified standoff

2027

S-URC27-RULES

19.05 mm

KEYBOARD KEY PITCH

About 0.44 N over 2 mm, with no auto-repeat

2027

S-URC27-TASKHW

29–35

THE FIELD'S MEAN FOR FOUR SEASONS

Flat despite four years of rule changes

2023–2026

S-URC22-SCORES

Key takeaways

- **The flat mean is the finding.** The mission mean has stayed between 29 and 35 for four seasons despite rule changes, because the difficulty is manipulation itself rather than the task list.
- **Open-loop precision does not exist at this scale.** Planetary joint backlash of **15–18 arcmin is 4.4–5.2 mm at 1 m**, wider than a keyway. Vision has to close the loop on typing and on the key.
- **2027 adds two hard tasks and removes one.** The key in the tubular lock has no precedent, the hose is stiffer, and the USB Type-C task is gone. Per-task weights are unpublished in both years, so the five-year mean is the base rate.
- **A mid-cost arm reaches the middle of the field.** UIU's 72.63 came from a 6-DoF arm, an ArUco-aligned typing routine and a rehearsed mock lander S-UIU-2026-SAR .

What the mission physically demands

DEMAND	NUMBER	BASIS
Workspace	Ground level to 1.5 m above ground, from an unspecified standoff; a 100 m traverse; 30 minutes	S-URC27-RULES §1.d.i
Task families	Nine, in any order: sample tube, cache transport, drawer, latch and panel, typing, key and lock, hose, quarter-turn valve, buttons and switches	S-URC27-RULES §1.d.ii
Keyboard	Redragon K552, 87-key tenkeyless, 354 × 123 × 37 mm, 875 g, 19.05 mm key pitch, about 0.44 N actuation over 2 mm of roughly 4 mm travel, no auto-repeat	S-URC27-TASKHW
Keyboard tags	2 × 2 cm ArUco 4x4_50 markers at the four corners	S-URC27-RULES §1.d.ii
Key and lock	A flat key with a 1 × 1 cm tag on the bow, a standard 7-pin tubular keyway with 1 × 1 cm tags at the four corners of the face, and a 1–2 mm keyway to align to	S-URC27-RULES §1.d.ii
Accuracy budget implied	Backlash of 15–18 arcmin is 4.4–5.2 mm at 1 m	Datasheet values, by inference
Time per task implied	30 minutes minus a 100 m traverse and nine families leaves about 2–3 minutes each, including approach and retreat	By inference; Monash 2022 spent about 4 minutes on the door latch alone across two attempts S-MONASH-ES22
Contact rule	Every task must be doable with one gloved hand, and the keyboard must not be touched by hand	S-URC27-QA S-URC27-RULES §1.d.ii EQ1

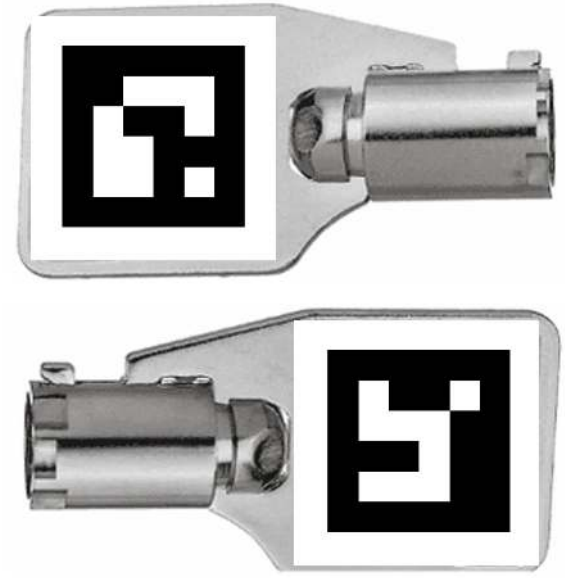


The keyboard is a standard 354 × 123 × 37 mm tenkeyless unit with 19.05 mm pitch. The tags at the corners give the arm a visual reference, but the press still has to land on a key that is 19 mm wide
S-URC27-TASKHW.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

§1.d.ii." aria-label="Open full-size figure">



The key task is the smallest target in the rulebook: a flat key, a 1 × 1 cm tag on the bow and a keyway to align to. Removal, transfer between two locks and a turn are three separate failure points in one task
S-URC27-RULES §1.d.ii.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

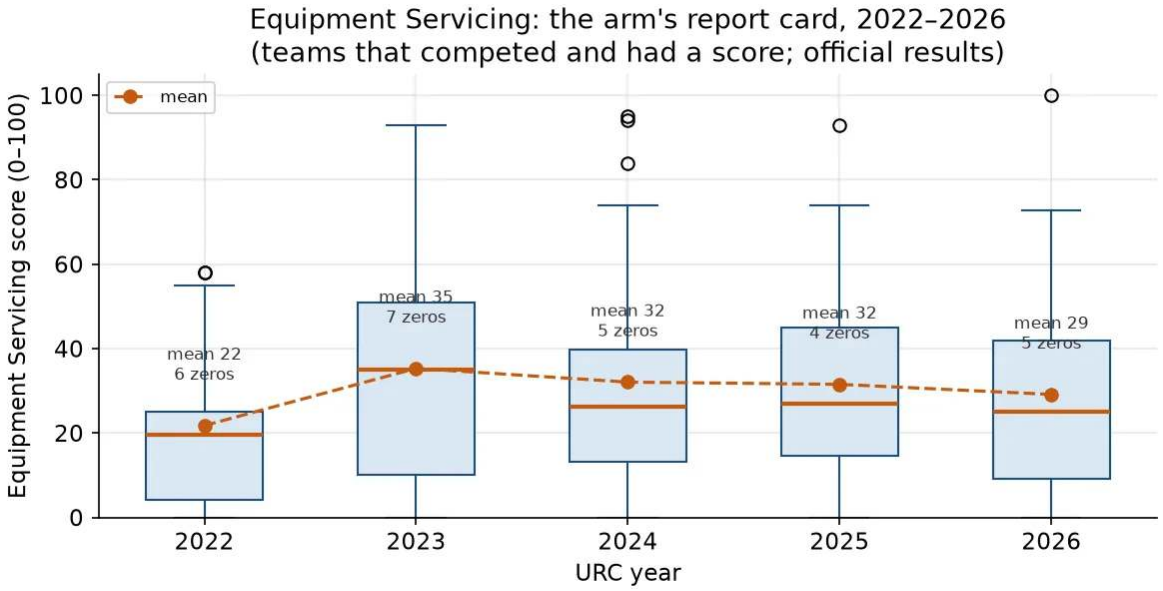
Original source *Full resolution*

What the best teams did

The record is consistent: a long tail of zeros, a handful of 93–100 runs, and a mean that does not move.

YEAR	TEAMS	MEAN	MEDIAN	EXACT 0	SCORES ≥80	MAXIMUM
2022	26	21.8	19.5	6	0	58 (Missouri S&T and IMPULS)
2023	31	35.3	35.0	7	3	93 (AGH)
2024	34	32.0	26.3	5	3	95 (BYU)
2025	35	31.5	27.0	4	1	93 (Missouri S&T)
2026	36	29.1	25.0	5	1	100 (Missouri S&T)

Derived from the official mission column in the five result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`.



Five seasons of Equipment Servicing scores. The top of the distribution moves between teams, but the mean and the zero count do not, which is what a mission limited by manipulation quality looks like.

Source material, 2026. `S-URC22-RESULTS` `S-URC26-RESULTS`

Full resolution

TEAM	2022	2023	2024	2025	2026	ARCHITECTURE AND LESSON
Missouri S&T	58	68	**94**	93	**100**	A hybrid of one prismatic and five rotary axes, mounted on the top deck and chosen for precision on the panel face
BYU	18	58	**95**	59.2	22.11	A wrist sensor head with ViSP visual servoing; bimodal, 95 then 22
AGH	16	**93**	74	did not enter	did not enter	93 in 2023 was the field maximum
West Virginia	25.2	85	84	**43**	**38.95**	An over-torque cascade in 2023 and a damaged USB-C cable in 2026 that took out cameras and two motors
Monash	53	82	**18**	54	58.95	Every arm redesign scored worse than the design it replaced, 2023 excepted
UIU	23	50	72	42.4	**72.63**	A 6-DoF arm, ArUco typing and a rehearsed mock lander: the most replicable result in the record
MIST	2021: **100**	29	30	41	23.68	The only complete judged run on video, from 2021; hobby servos once the tasks hardened

Official Equipment Servicing column `S-URC22-SCORES` `S-URC26-SCORES`*, with MIST's 2021 result from its own competition record* `S-MIST-VURC2021`*.*

The per-task pattern. Tasks where the object comes to the end effector are routine: buttons, switches, the drawer handle and the latch release. Tasks where the end effector must arrive at a precise pose in free space fail: typing, key insertion, the hose push and re-opening a panel. 2027 adds two of the second kind.

What 2027 changes

2027 CHANGE	VERDICT
Key in a tubular pin-tumbler lock added	**No precedent.** 1 cm tags, a 1–2 mm keyway, and insertion and turn torque not yet measured
USB Type-C task removed	**Invalidated.** Fixtures, poses and rehearsal built for USB are obsolete, and one bankable task disappears
Hose: GatorLock with no gasket becomes a 1.5-inch cam lock with a half-thickness gasket	**Invalidated hardware.** 2026's fitting was the easier push, so any 2026 tuning is to the wrong part
Display formally linked, at 4.2 inches on e-paper while the Q&A says 4.3 inches	**Raised bar.** A no-touch zone, the launch key read from the display, and judges who reset it
Buttons, switches, knobs, quarter-turn valve, latch, drawer, tube and cache	**Carry forward**
Per-task weights	**Not published in either year**, so no pricing of tasks is possible

§1.d.ii." aria-label="Open full-size figure">



The 7-pin tubular keyway at the centre of the four alignment tags. The arm has to hold a key on a target this small while applying turning torque, which is why the task carries the highest uncertainty in the competition S-URC27-RULES §1.d.ii.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

Original source *Full resolution*

The net effect is not simply “harder”. Two hard tasks were added, one was removed, and the weights are unpublished. Anyone who assumes the 2026 mean carries over exactly is guessing in both directions.

Which subsystem options serve it

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Arm architecture	**Five rotary axes plus one prismatic axis**, with the rail parallel to the lander face, when Equipment is the priority: the best scores in the field use this layout	**Six rotary axes** for the least-surprising path with in-house machining; five axes plus a 400 mm lateral stage when the arm would otherwise be too long to stow
Rail caution	A prismatic axis **must** be designed for driving shock: Missouri S&T’s own vertical axis failed under driving in 2024 S-MST-NEWS2024	A pure rotary arm if the rail cannot be made shock-proof in one season
Actuators	Integrated brushless modules with CAN and torque mode, to buy past driver development	Harmonic or cycloidal drives at the wrist when typing or the key task must work open-loop
End effector	**A two-finger 100–110 mm jaw with a compliant typing tip, a chamfered key funnel and a hook or pin**	A turret of several end effectors, never on a first build
Perception	**A wrist camera at 1080p, 60–70° field of view, fixed focus at 0.2–0.5 m**, for the 2 cm keyboard tags	A 5 MP camera if the 1 cm key tag needs margin; the mast camera cannot resolve 1 cm at task range
Typing	**Two-stage visual servo**: coarse inverse kinematics, then tag-based fine alignment, with each key confirmed	Open-loop typing only if wrist backlash is 0.5 mm or less
Order	The cheap tasks banked first, because any order is allowed	The hardest task first only if its weight dominates the gate, and that weight is unpublished
Transition	One arm configuration for both Autonomy and Equipment, so no re-weigh is triggered S-URC27-QA EQ4	A tip swap inside the 10-minute window, never a module swap

Minimum viable versus stretch

- **The minimum that scores:** one completed manipulation task. The cheapest are buttons, switches, knobs, the quarter-turn valve and the hinged latch.
- **The mid-field:** four to six tasks banked. UIU's 72.63 in 2026 shows this band is reachable with a mid-cost 6-DoF arm and a rehearsed panel.
- **The stretch:** all nine families, including typing, the key transfer and the hose. One team did it for 100 in 2026, with a prismatic-axis arm and a dedicated typing tip.

Risks to retire early

RISK	WHY IT IS REAL	THE TEST THAT RETIRES IT	LEAD TIME
Key dropped or lost	The task is new and a dropped key ends it, possibly for both locks	A tethered key, a mechanical stop in the jaw and a drop test over a padded tray: 50 of 50 grasps, retained through 20 shakes and a 1 m drop	5 months
Key will not insert or turn	A 1–2 mm keyway against 4.4–5.2 mm of open-loop backlash at 1 m	The same class of 7-pin lock, with insertion force and turn torque measured and the keyhole tags used for visual servoing: 20 of 20 insertions	5 months
Hose will not seat	The gasket is now half thickness, and 2026 tuning was to a fitting with no gasket	An exact 1.5-inch cam-lock mock, with force logged against depth over 50 cycles and a retention pull to confirm the connected state	4 months
Over-torque cascades to other subsystems	West Virginia in 2023 lost its arm cameras through a USB hub, and in 2026 lost cameras and two motors to a damaged cable S-WVU-YT-DOORFAIL S-WVU-NEWS2026	Per-joint current limits and fault-isolated camera paths, so that no single cable takes out more than one camera and no motor	4 months
Typing misalignment or a double press	Backlash, motion blur and no auto-repeat	One hundred automated presses at 3 mm offsets, each key confirmed by camera: 95 % correct on 3–6 letter keys, including a backspace correction	4 months
Panel slips closed	Monash 2022 lost about 9 minutes to two closures S-MONASH-ES22	A prop or hook that holds the panel, rehearsed as its own step: the panel stays open hands-off for 60 seconds	3 months
Drawer jam or the rover pulled	The insertion is tight-fitting and the pull force is unmeasured	A mock drawer across a range of clearances: ten insertions at the tightest clearance with no rover motion	3 months
Time runs out	About 2–3 minutes per task and a 100 m traverse inside 30 minutes	A timed full run on a mock panel in the real order: six tasks banked in 20 minutes	3 months
The 10-minute transition	Any change in mass or configuration forces a re-weigh S-URC27-QA EQ4	Autonomy, weigh-in and Equipment rehearsed inside the window: 8 minutes end to end, five runs in a row	2 months
Wrong button or switch	One mis-press can fail a task	A mock labelled exactly like the rulebook figures: 20 of 20 correct on a randomised sequence	2 months

What is not published

- The per-task weights, which is why tasks can be ranked by risk but not by points per minute.
- Every force, torque and travel for the key, the lock, the hose, the valve, the drawer, the latch and the switches. A mock and a load cell are the only route to a number.
- Whether the cache in this mission is the team's own Astrobiology cache or a judge-supplied one, so the interface must work for both.
- Whether digits can appear in the typing sequence, since the rule says 3–6 letters.
- The 2027 partial-credit split for typing corrections.

Further reading

- Arm and end-effectors works the architecture, backlash budget and end-effector choices behind this page.
- Compute, sensors and perception covers the wrist camera and the tag geometry.
- Missouri S&T, UIU, BYU, Monash and West Virginia carry the Equipment evidence.
- What 2027 invalidates places the key, the USB removal and the hose change beside the other 2027 changes.

MISSION BACKFIT · CHAPTER 6.4

Autonomy: astronaut assistance

The astronaut-assistance sub-mission is half of the Autonomy score and the only part of URC 2027 where a person stands inside the control loop. A team member in a suit walks, stops, walks away and holds out a hammer, and the rover answers with five commanded behaviours. Forty-five of the fifty points require autonomy; only the suit itself is manual.

50

POINTS IN THIS SUB-MISSION

Half of the 100-point Autonomy mission

2027

S-URC27-RULES

45

OF THE 50 REQUIRE AUTONOMY

A teleoperated hammer pickup scores zero

2027

S-URC27-RULES

15

MAXIMUM FOR FOLLOW!

Device 5, sign 5, audio 10, gesture 15

2027

S-URC27-RULES

0.91 kg

THE HAMMER

About 300 mm of one-piece steel, lying on the ground

2027

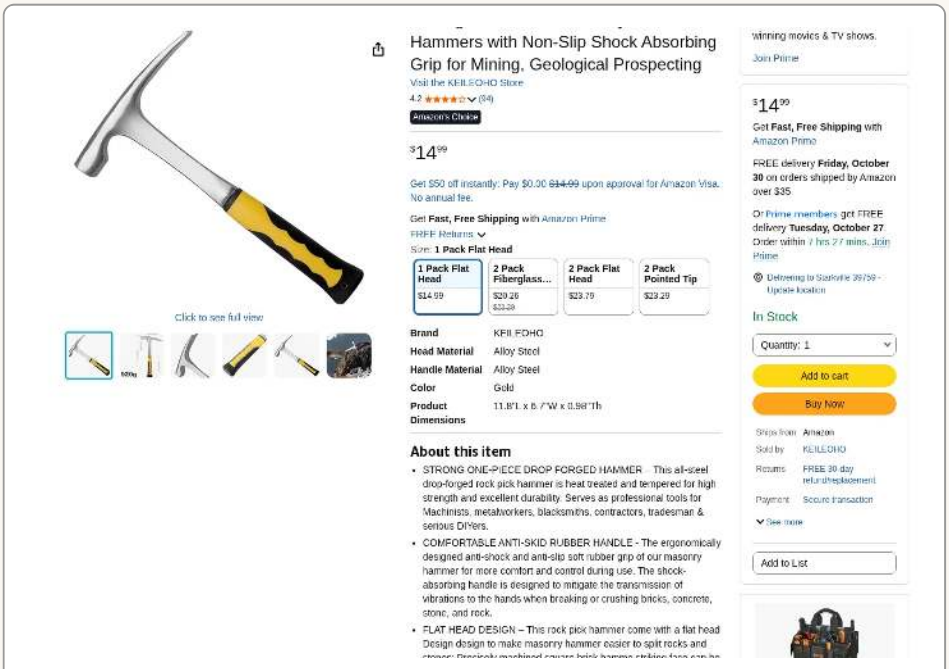
S-URC27-TASKHW

Key takeaways

- **The interaction is new; the geometry is not.** The Mars Society ran suited-astronaut tasks for nine years, and the 2011–2013 version required delivery **within 1 m** of a suited figure from a GPS fix accurate to **±25 m**. Only the 15-point Follow! has no precedent at all.
- **Follow! is a ladder, not a modality.** Gesture, then audio, then a sign or a device button: a rung that fails cannot subtract points, so the floor is 5 and the ceiling 15 S-URC27-QA AUQ6.
- **The rules give geometry away.** The astronaut's GNSS point, the guaranteed 20 m standoff with the astronaut in the field of view, and a two-way link that is encouraged all convert hard perception into easy localisation.
- **The minimum is 30 of 50** without person-following or a hammer grasp; 40 with Fetch!; 50 with gesture Follow!.

What the sub-mission physically demands

DEMAND	NUMBER	BASIS
Suit	Team-provided, covering neck, wrists and ankles, with closed-toe shoes, a camera and a microphone monitored in C2, and a helmet the astronaut can remove alone	S-URC27-RULES S-URC27-QA §1.e.iv AUQ12
Heat	The rules cite 37.8 °C ambient; the June daily maximum at the site averages 33.3 °C and reaches 40.7 °C	S-URC27-RULES S-OPENMETEO-ARCHIVE §3.d.iii
Drive to the astronaut	Autonomous, stopping within 3 m of the GNSS point	S-URC27-RULES §1.e.v
Follow!	Follow a walking person on natural terrain and stop within 3 m when the person stops; one modality, 15 points maximum	S-URC27-RULES §1.e.vi
Stay!	No motion while the astronaut moves more than 20 m away	S-URC27-RULES §1.e.vii
Fetch!	Detect, grasp and lift a 0.91 kg, roughly 300 mm one-piece steel hammer lying on the ground; a teleoperated pickup scores zero	S-URC27-RULES S-URC27-TASKHW §1.e.viii
Come!	Autonomous drive to the astronaut and a stop within 3 m, with the astronaut at least 20 m away and in the field of view	S-URC27-RULES S-URC27-QA §1.e.ix AUQ7
Give!	Place the object on the ground, or drop it	S-URC27-RULES §1.e.x
Command order	Follow, Stay, Fetch, Come, Give	S-URC27-QA AUQ13
Time	Inside a 40-minute mission shared with route finding; no separate clock for this sub-mission	S-URC27-RULES §1.e.i
Safety	No points penalty for contact, but an astronaut E-stop ends the mission; avoidance is highly recommended	S-URC27-QA AUQ1
Human-robot safety frame	ISO 13482 protective stop and speed-and-force limits are good practice, although URC does not require compliance	S-ISO-13482



The specified Fetch! object is a 32 oz rock pick of about 0.91 kg and 300 mm, and the grip is the only place the jaw can hold it. The task is a ground pickup of a low-profile steel object, not a lift S-URC27-TASKHW.

Amazon product listing, 2026. Product listing; reproduced for reference.

S-URC27-RULES S-URC27-QA Full resolution

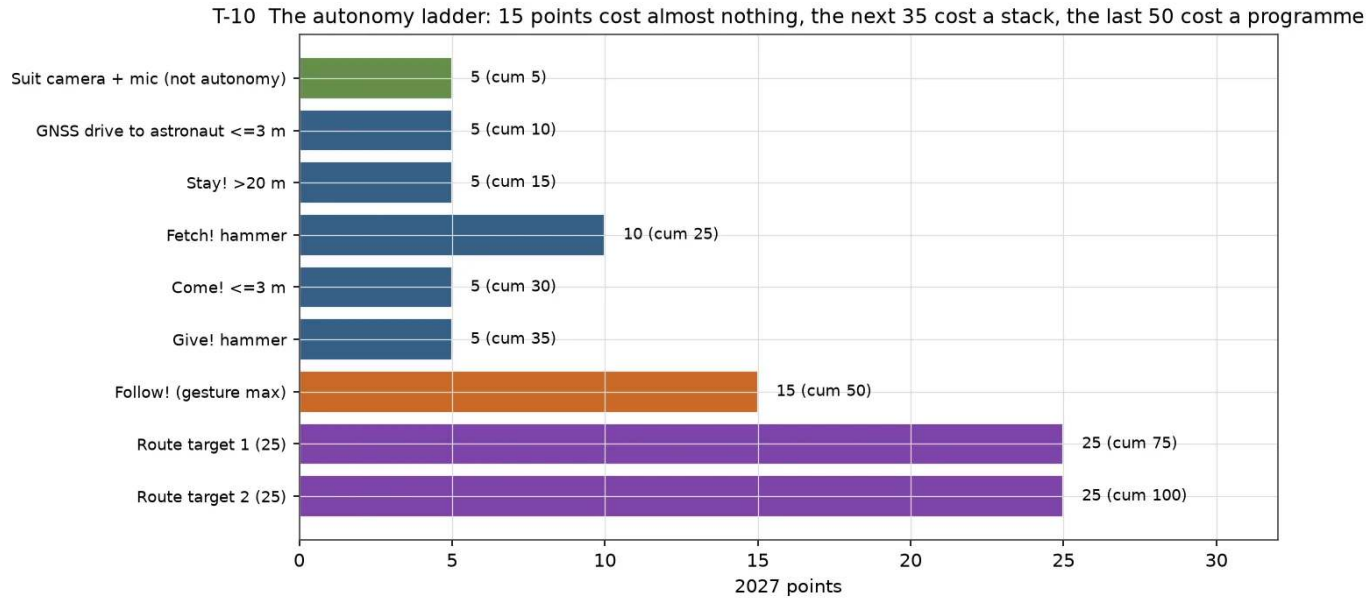
The nine-year precedent

There is no precedent for the commanded interaction, but there is a long one for the geometry.

PRECEDENT	WHAT IT REQUIRED	YEAR
Emergency Navigation	Deliver a container **within 1 m** of a simulated distressed astronaut, identified by the presence of a space suit; the astronaut may have strayed from the last known position	2008
Astronaut Assistance	Deliver within **1 m** of two to five suited astronauts, with GPS accurate to only **±25 m**, so visual acquisition was mandatory, and up to half of them out of line of sight	2011–2013
Astronaut Assistance, evolved	The same, plus collecting equipment and returning samples, including small hand implements such as a screwdriver, a hammer or a wrench	2014–2016
Delivery sign reading	Read signs carried by astronauts or placed in the field	2017–2026
Commanded interaction	Follow, Stay, Fetch, Come and Give	2027, new

Reading the precedent correctly. Thirty-five of the fifty points rest on capabilities the field exercised before 2017: approaching a suited figure, stopping close, and handling a hand implement. The precedent is about geometry and perception, not about autonomy level. The 2008–2016 tasks were largely teleoperated, whereas 2027 requires autonomous operation for 45 of the 50 points, so the history de-risks the problem rather than the implementation.

§1.e.iv §1.e.xviii." aria-label="Open full-size figure">



Every 2027 Autonomy scoring element in the order a first build can earn it. The suit camera and microphone (green) is manual work; the astronaut ladder pays 15 points at the bottom and 35 with Fetch, Come and Give; Follow! pays by method; the two route targets are 25 each S-URC27-RULES §1.e.iv §1.e.xviii.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES Full resolution

What 2022–2026 adds

There was no astronaut task in this window, but its parts were trained.

PRACTICE	EVIDENCE	USE IN 2027
ArUco detection with visual servoing and marker-proximity stops	West Virginia validated the combination in 2026 S-WVU-SAR2026	The sign channel and the route posts become nearly free
Object detection with an explicit class	Missouri S&T trained a rock-pick class at 0.60 confidence S-MST-DOCS-CONSTANTS	Fetch! is a class the field already trains
Ground-object detection benchmarked by mean average precision	West Virginia reported the 80th percentile across mallet, keyboard, water bottle and rock hammer S-WVU-SAR2026	Sets a realistic bar for the hammer
A reached-goal radius tighter than the rule	Missouri S&T uses 1.5 m with 10 seconds of stationary GNSS averaging S-MST-DOCS-CONSTANTS	A 3 m stop is comfortable if the terminal discipline exists
Operator practice as a published standard	West Virginia requires a minimum of **100 hours of practice** S-WVU-SAR2026	The single most useful planning number in this sub-mission
A detector that works indoors and fails in the field	Missouri S&T could not locate a specific rock through its cameras in 2023 S-MST-NEWS2023	Every detector is unproven until it runs outdoors, in sun, at range

The design insight worth acting on

The rules hand over geometry in three places. The astronaut's GNSS point is provided for the 5-point drive, so the approach is a waypoint drive and visual acquisition only has to work over the last 25 m, exactly as it did in 2011–2016. Fetch! and Come! guarantee that the astronaut is at least 20 m away and in the field of view, so the rover has a known bearing to look along. Because the astronaut may carry electronics but may not deploy anything S-URC27-QA AUQ26, and because a two-way astronaut-to-C2 link is encouraged S-URC27-QA AUQ5, a **suit-carried GNSS beacon** would turn Come! and the initial drive into waypoint problems, leaving person detection needed only for Follow!.

That is a large reduction in perception risk for 5–10 points, at the cost of a small radio and a battery in the suit. Two things have to be answered before it can be designed in: whether URC accepts a beacon-derived position as satisfying the autonomous drive, and how the beacon is declared in the communications documentation, which is due **2027-04-30** S-URC27-RULES §2.c.ii. The beacon is a hypothesis to test, not a settled design.

Follow! is a ladder, not a modality

Because the four methods are alternatives and the rules permit switching mid-mission S-URC27-QA AUQ6, they form a fallback ladder rather than a choice. The ladder attempts gesture first, falls back to audio, then to a sign or a device button. A rung that fails cannot subtract points, so the worst case is the 5-point device and the best is 15.

RUNG	POINTS	EFFORT FOR A FIRST BUILD	IMPLEMENTATION
Device button	5	About 1.5 person-weeks	Wearable button on the suit, received by the rover
Printed sign	5	About 1 person-week	ArUco sign, reusing the detection needed for route posts
Audio	10	About 3.5 person-weeks	Suit microphone to C2 to rover command
Gesture	15	About 6–10 person-weeks	Visual recognition, promoted only after an outdoor test with zero false positives

The effort figures are planning estimates from the astronaut-interaction study, by inference; no URC source measures effort.

Which subsystem options serve it

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Suit	**Coverall, cooling vest, analogue 5.8 GHz video pod, LoRa button and a printed ArUco sign**, which secures the 5 suit points and a 5-point Follow!	An IP-camera suit pod when the C2 already runs an IP video stack and someone owns the Linux integration
Follow! modality	**Gesture first**, only after 90 % true-accept and zero false positives at 6–8 m in sun on the real suit	**Audio** when the gesture test has not been run, and never a device alone
Audio path	**Suit microphone to C2 to rover command** as the primary route	A direct microphone only if the astronaut is within a few metres and the rover is quiet
Person perception	**A YOLO-class detector with depth** for range and bearing, and pose estimation for gesture	The astronaut's own camera, which AUQ24 permits, when the rover's view is occluded
Fetch grasp	**A wide jaw of 100 mm or more, or a hook on the shaft**, with a wrist camera and a vertical lift of 100 mm before retracting	A scoop or under-slide only if the hammer lies on flat ground, which it may not
Stay! implementation	**An explicit hold state that disables the planner** and asserts zero velocity	A GNSS hold point only if it is proven not to creep on a slope
Arrival and mode signalling	**An LED driver on the authoritative mode topic**, with a sunshade, tested at noon	A C2 indication in addition, which is recommended but not required `S-URC27-QA` AUQ18

Minimum viable versus stretch

- **The minimum that scores:** a suited team member with a working camera and microphone (5) plus the four short behaviours of drive to the GNSS point (5), Stay! (5), Come! (5) and Give! (5). None of them needs person-following or a hammer grasp, but Come! still needs the astronaut's position at 20 m, which is either a beacon or a person detector. With a sign or device Follow! (5) the total is **30 of 50**.
- **The credible level:** add Fetch! (10) for **40 of 50**, which needs hammer detection and a grasp, both 2026-class capabilities.
- **The stretch: 50 of 50**, adding gesture Follow! for 15 and earning the full margin.

Risks to retire early

RISK	WHY IT IS REAL	THE TEST THAT RETIRES IT	LEAD TIME
The wrong person is followed	A judge or spectator can stand near the astronaut, and AUQ10 allows the astronaut to stand close to the rover	A fiducial panel on the suit and a requirement that the team's own marker is present before closing distance: zero wrong-person locks in 20 runs with a bystander crossing	4 months
A gesture false positive	Natural arm swing read as a command, or a bystander waving	Two-of-three confirmation plus a rest-pose gate: 90 % true-accept and zero false positives in 10 minutes of field noise	4 months
A command is not recognised	Wind at a June mean of 7.3 m/s, distance, a small sign or a gloved button	The fallback ladder with timeouts, oversized buttons and a microphone near the mouth: every rung passes its own bench test and Follow! completes in 3 minutes	3 months
A command is recognised twice	A repeated trigger re-arms the follow controller after Stay!	Command identifiers with a latch that ignores duplicates within 2 seconds: 20 of 20 single-transition runs	3 months
The hammer is not detected or the grip fails	0.91 kg and low profile on gravel, with Missouri S&T's rock-detection failure as the cautionary case	Fifty automated grasps on five ground poses: 45 of 50, and no drop over a 20 m carry	4 months
Stay! creeps	Suspension wind-up or a planner tick	A hold state that disables the planner, with an encoder zero-motion check over the full walk	3 months
Suit heat and astronaut distress	June conditions of 33–41 °C under a coverall	A cooling vest, hydration and a timed suit trial in sun: 30 minutes without distress and a helmet removable in under 10 seconds	3 months
The suit feed is not on screen	Five points depend on the C2 display rather than on a recording	An always-visible window and a check with the C2 judge before the clock starts	1 month
A hands-off violation	A reflex grab during a near-collision S-URC27-QA AUQ21	The autonomy key out of reach and fault-injection drills: zero control touches across 10 injected near-misses	2 months
Communications documentation omission	The suit pod, the command device and any repeater are all communications devices S-URC27-RULES §2.c.ii	A device list owned by the communications lead: every device listed with its licence and power	2027-04-30

What is not published

- Whether a suit-carried GNSS beacon is an acceptable basis for the autonomous drive to the astronaut.
- Whether voice guidance from the astronaut to C2 counts as human assistance while driving.
- Whether an inertial gesture is legal, since the rule says visual recognition and the Q&A permission names cameras.
- Any measured gesture success rate, stopping distance or speech-recognition false-accept rate at URC.

Further reading

- Astronaut interaction and the EVA suit works the suit, the gestures and the safety envelope.
- Software stack and autonomy covers the behaviour tree and the hold state.
- Autonomy: route finding shares the 40-minute clock and the same localisation.
- Missouri S&T, West Virginia and Monash carry the perception and operator-practice evidence.

MISSION BACKFIT · CHAPTER 6.5

Autonomy: route finding

Route finding is the other 50 points of the Autonomy mission and the part where autonomy either works or does not. Two targets sit at exact coordinates. One is on a hill where not every approach direction is navigable; the other is behind the hill and out of radio line of sight. Each target is worth 25 points and each has a 1 m stop tolerance.

50

POINTS FOR ROUTE FINDING

Two targets at 25 points each

2027

S-URC27-RULES

1 m

STOP TOLERANCE

Absorbs GNSS, map and control error together

2027

S-URC27-RULES

20 × 20 cm

AR POST FACES

0.5–1.5 m up, ArUco 4x4_50, 2.5 cm cells

2027

S-URC27-RULES

41 %

AUTONOMY ZERO RATE

67 exact zeros in 162 team-years

2022–2026

S-URC22-SCORES

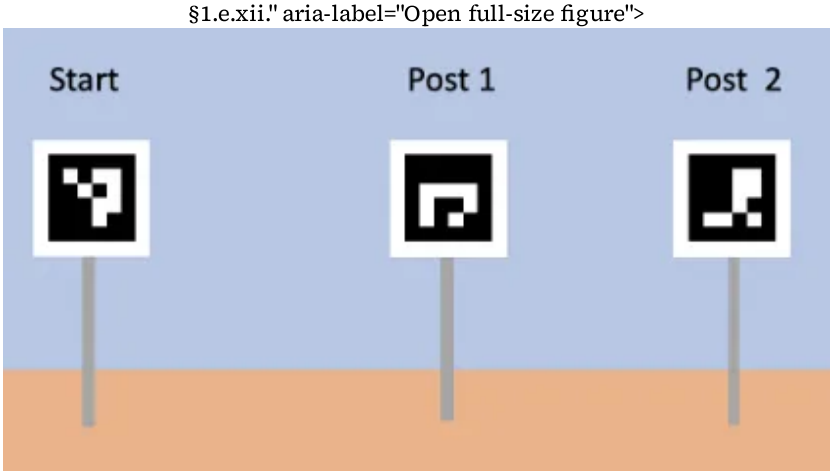
Key takeaways

- **This is the all-or-nothing half.** Across 162 team-years, 41 % of Autonomy scores are exactly zero, the highest zero rate of any mission S-URC22-SCORES S-URC26-SCORES .
- **Exact coordinates make the driving easier and the stop harder.** The coordinates remove the need for visual search, while the stop tolerance tightens from 2–3 m in 2026 to **1 m** in 2027 S-URC27-QA S-URC27-RULES §1.e.xvii AUQ16.
- **The cleanest experiment in the record is Monash.** A ground-up rebuild scored **0** in 2024 and **14** in 2025; removing the sophisticated planner and returning to NavFn scored **86** in 2026 S-MONASH-SAR2026 .
- **The radio link never gates the drive.** The second target is designed to be out of line of sight, so one link-status check is a 25-point failure.

What the mission physically demands

DEMAND	NUMBER	BASIS
Targets	Two, in any order, with no return required, at **exact** coordinates	S-URC27-RULES S-URC27-QA §1.e.xii AUQ16
Stop tolerance	Within 1 m of the target, plus an indication of arrival	S-URC27-RULES §1.e.xvii
Terrain	Hilly natural terrain; target 1 on a hill where not all approach directions are navigable; target 2 behind the hill and out of radio line of sight	S-URC27-RULES §1.e.xv §1.e.xvi
Markers	Three-sided AR posts with 20 × 20 cm faces, 0.5–1.5 m above ground, ArUco 4x4_50 with 2.5 cm cells and a one-cell white border	S-URC27-RULES §1.e.xii
Arrival indication	A rear LED flashing green; a C2 indication is strongly recommended but not required	S-URC27-RULES S-URC27-QA §1.e.ii AUQ18
Map data	URC will attempt a digital elevation model at 1 m resolution, likely GeoTIFF, not guaranteed; satellite imagery is encouraged	S-URC27-RULES S-URC27-QA §1.e.xiii AUQ20
Time	Inside the same 40 minutes as astronaut assistance	S-URC27-RULES §1.e.i
Penalties	20 % for an operator stop, 50 % for teleoperation back over explored ground, capped at 50 % per task; a physical intervention is uncapped, and teleoperation may not enter unexplored ground	S-URC27-RULES §1.e.xviii
Stop budget implied	About 0.6 m of the 1 m tolerance allocated to the stop itself	S-URC27-RULES §1.e.xvii, by inference
Approach kinematics	1 m/s at 0.3 m/s² needs about 1.7 m to stop	By inference

§1.e.xii." aria-label="Open full-size figure">



The 4x4_50 tag geometry is unchanged from 2026: three 20 × 20 cm faces, one tag, readable from any side. The post is the one feature on the course that a camera can confirm at close range, so the final approach can be visual even though the arrival signal itself is an LED `S-URC27-RULES` §1.e.xii.

The Mars Society / URC, 2026. Competition rulebook; reproduced for reference.

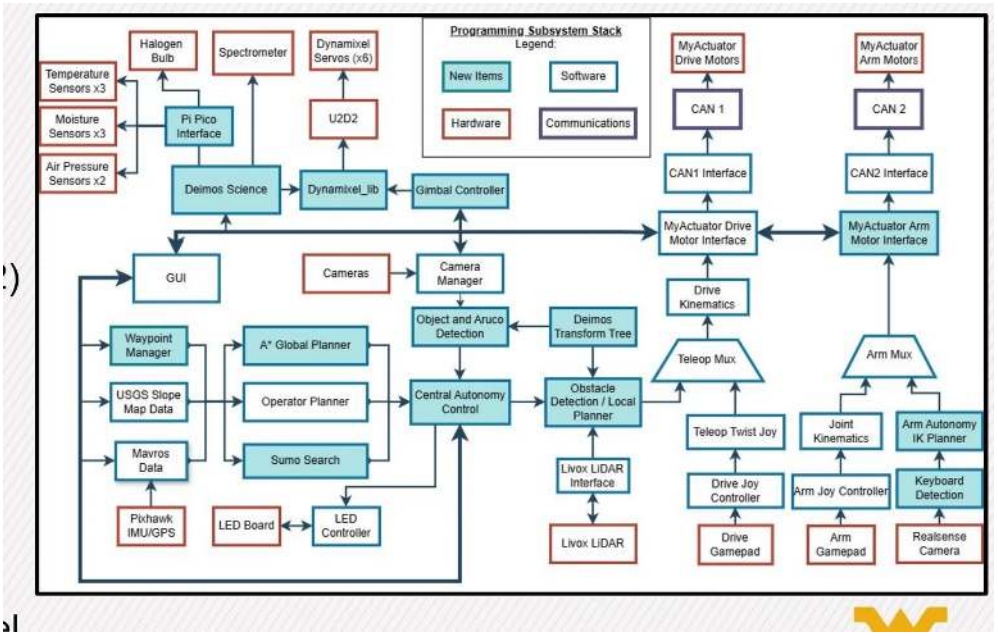
Original source *Full resolution*

What the best teams did

YEAR	TEAMS	MEAN	MEDIAN	EXACT 0	MAXIMUM
2022	26	23.9	11.2	12	84
2023	31	26.4	16.8	15	96 (Michigan)
2024	34	23.9	5.6	17	72
2025	35	31.5	28.0	7	86 (BYU)
2026	36	26.7	21.0	16	86 (five teams)

Derived from the official Autonomous Navigation column in the five result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. *Michigan's 96 in 2023 is the window maximum.*

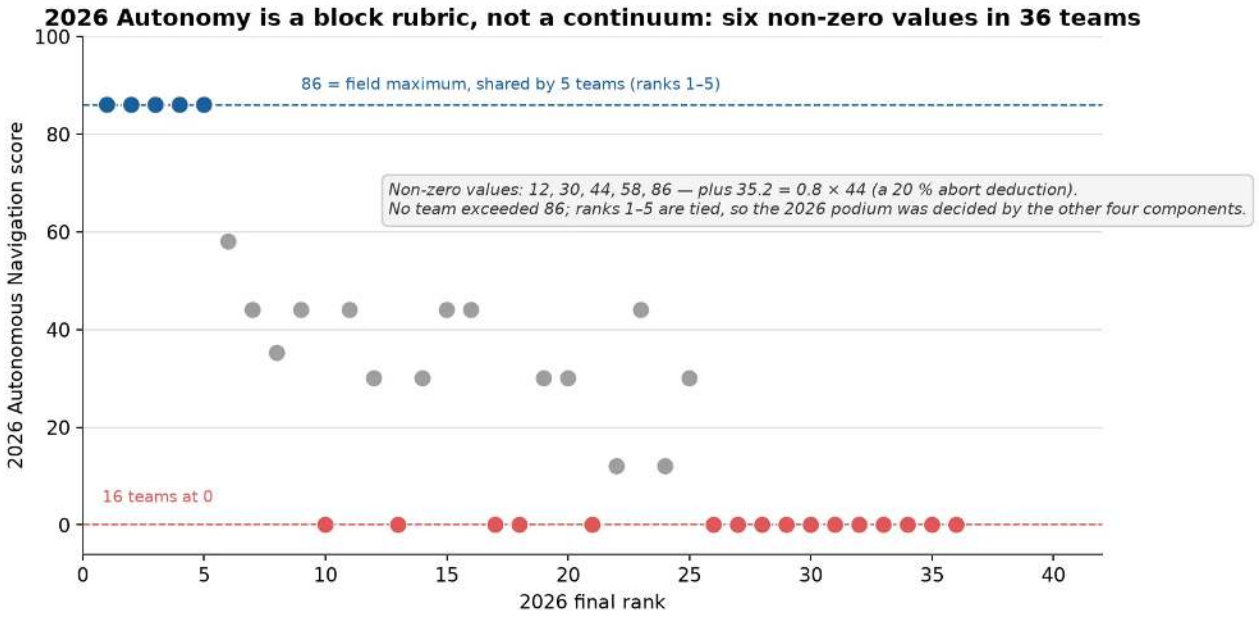
- **Missouri S&T** sat at 56 for three seasons, then reached **86** in 2026 with satellite terrain data and LiDAR planning `S-MST-SITE-ATHENA` `S-URC26-SCORES`.
- **Monash** went **0 (2024), 14 (2025), 86 (2026).** The 2026 fix was removing a state-lattice planner in favour of NavFn with a voxel costmap, after a ground-up rebuild had scored zero `S-MONASH-SAR2024` `S-MONASH-SAR2026`.
- **West Virginia** reached **86** in 2026 with A* over USGS elevation data, a Livox Mid-360 with FAST-LIO, and local traversability planning `S-WVU-SAR2026`.
- **Michigan** reached **96** in 2023 after a microservice rewrite, then scored 56, 56 and 86 `S-URC23-SCORES` `S-URC26-SCORES`.



The 86-point 2026 stack in the team's own words: a GUI, a waypoint manager, A over USGS slope data, a local planner, ArUco and object detection, and an LED board. Nothing in the diagram is exotic; the value is that every path is connected and tested* `S-WVU-SAR2026` .

Logan Gold, Ian Long, Parker Vance, Izaak Whetsell (WVU Team Mountaineers), 2026.

`S-WVU-INCOSE2026` *Full resolution*



The 2026 field. The five teams at exactly 86 are the top five ranks, sixteen teams sit at zero, and the rest fill a gap between 12 and 58. The 86 is the rubric ceiling, not the best achievable capability `S-URC26-SCORES` .

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC26-SCORES` *Full resolution*

The same values as an interactive chart:

The transferable lesson. Each of the three teams that document their planning stack in detail used a terrain map, from an elevation model or from LiDAR, with a conservative planner. The team that moved from sophistication to conservatism gained 72

points. The map data is worth buying, the planner is worth keeping boring, and the effort belongs on the terminal stop.

Does 2027 invalidate it

Partly. The AR-post geometry is unchanged `S-URC26-RULES` `S-URC27-RULES` §1.f.iv §1.e.xii, and exact coordinates make the navigation easier, but three things tighten at once.

ELEMENT	STATUS	VERDICT
AR-post detection and terminal servoing	Same tag library and cell size	**Carries forward**
GNSS waypoint navigation	Still the backbone, with coordinates now exact	**Carries forward, easier**
Stop tolerance	From 3 m at a GNSS point and 2 m at a post in 2026 to **1 m**	**Raised bar**
C2 arrival display	The LED is now sufficient and the C2 banner optional	**Easier**
Hill route where not all approaches work	New emphasis, with no published inclinations	**Raised bar**
Second target out of radio line of sight	New, and not deliberately tested in 2025 or 2026	**New**
Penalty exposure	20 % for an operator stop, 50 % for teleoperation, capped at 50 % per task	**Raised bar**
Mission structure	30 minutes, a 2 km cap and five target classes became 40 minutes and two targets	**Invalidated as a mission, valid as a stack**

Which subsystem options serve it

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Localisation	**Dual-antenna RTK with a dual Kalman filter, wheel odometry and a magnetometer**, because heading error dominates once the rover is more than 10 m from the last fix	Single-antenna RTK with inertial heading; accept 2–4° of error and drive slower near targets
Map	**A USGS 3DEP elevation model at 1 m resolution turned into a slope costmap**, which is available now; treat the URC model as a bonus	Onboard voxel obstacle avoidance where the model is missing or rocks are smaller than 1 m
Global planner	**A lattice planner over the elevation costmap** for a non-holonomic skid-steer rover	NavFn when compute is limited, which is what Monash used in its 86-point year
Local planner and speed	**Regulated pure pursuit at 1.0 m/s or less**, a 2.5 m lookahead, and a terminal approach at 0.3 m/s or less	A sampling controller only with a GPU and time to tune it
LiDAR	**A Livox Mid-360, or none**, added once GNSS and stereo already work	No LiDAR in the first season if the stack is not end to end, because extrinsics, time sync and mapping are the failure modes a first build cannot absorb
Visual mapping	**Skipped in the first season**, because the targets are exact coordinates	Added only where the rover must cross a long GNSS shadow
Communications for the shadow	**A 900 MHz recovery band** carrying a low-rate heartbeat and drive commands	A rover or drone relay, which is allowed but unproven by any team in the sample
Behaviour	**A behaviour tree with three to five custom nodes** for the mode LED and task start and end	A hand-written state machine, written as an explicit table rather than nested conditions
Arrival	**An LED on the authoritative mode topic**, flashing green at 2–4 Hz, with a sunshade and a noon test	A C2 banner in addition, which is recommended but not required

Minimum viable versus stretch

- **The minimum that scores:** one target within 1 m. With exact coordinates, a good RTK fix and a terminal-stop discipline, that is achievable without terrain intelligence, provided the chosen target's approach is navigable. Any order is allowed, so the easier target can go first.
- **Partial credit is the safety net and it is unpriced.** The rule gives partial points for autonomously completing part of the route, at a rate that is not published. A rover that stops and holds may earn something; a rover teleoperated back earns 50 % at best and may enter no unexplored ground. The failure behaviour should therefore be stop and hold, not call for help.
- **The stretch:** both targets, including the one out of sight, inside the shared 40 minutes and with the astronaut sub-mission already complete.

Risks to retire early

RISK	WHY IT IS REAL	THE TEST THAT RETIRES IT	LEAD TIME
Overshoot past 1 m	1 m/s at 0.3 m/s² needs about 1.7 m to stop, and no published hit rate exists	Thirty runs around a surveyed point with a terminal approach at 0.3 m/s: 90 % within 1 m and 100 % within 1.5 m	4 months
The link-status check gates the drive	The second target is designed to be out of line of sight S-URC27-RULES §1.e.xvi	Two distinct fault policies, hold for the astronaut and continue and buffer for the route; the rover reaches the target with the radio off	4 months
The planner refuses the only gap	Inflation radius against gap width, for example Monash at 1.28 m S-MONASH-SAR2026	Inflation at 0.4 times the narrowest intended gap, with a route confirmed on the elevation model rather than in the planning software	3 months
The hill route is not traversable	URC publishes no slope limits S-URC27-QA AUQ17	An instrumented slope test at 10, 15, 20 and 25° on local analog soil, setting the slope limit from measured slip margin	4 months
A GNSS base offset puts everything 1.1 m off	The base station sets every coordinate the rover uses	A 30–60 minute survey-in on a known point, one base file all week, and logged fix type: repeatability within 2 cm over three sessions	3 months
The judge does not see the arrival LED	The LED washes out in sun S-URC27-QA AUQ3	A daylight test at 20 m with a sunshade and a 2–4 Hz flash, confirmed by an observer at noon	2 months
An operator aborts to be safe	An operator stop costs 20 % on that task	The rover's own abort as the default and an operator abort as a logged two-step action: zero operator aborts in 10 rehearsals	2 months
Hands on during autonomous driving	A reflex is the documented failure S-URC27-QA AUQ21	The autonomy key out of reach and fault-injection drills: zero touches across 10 injected near-misses	2 months
Time runs out across both sub-missions	40 minutes shared, and the astronaut sequence alone takes about 14 minutes	A timed combined rehearsal with at least 5 minutes of margin, three runs in a row	3 months

What is not published

- The partial-credit rate for a partly completed route.
- Whether URC will actually release the 1 m elevation model, which is promised as an attempt only.
- Any slope or step limit for the hill route.
- Whether voice guidance or a beacon-derived position is acceptable during the autonomous drive.

Further reading

- Software stack and autonomy covers the localisation, planner and behaviour choices in more detail.
- Compute, sensors and perception covers RTK, cameras and LiDAR.
- Autonomy: astronaut assistance shares the 40-minute clock and the same localisation.
- Missouri S&T, Monash, West Virginia and Michigan carry the route-finding evidence.

MISSION BACKFIT · CHAPTER 6.6

System Acceptance Review

The System Acceptance Review is the mission that decides whether there is a field season at all. A written report and a five-minute video, due **2027-03-05**, are worth 100 of the 500 points and are the only route to the other 400, because only the teams with the highest scores are invited. It is also the least discriminating score in the competition.

100

SAR POINTS

A fifth of the maximum score before penalties

2027

S-URC27-RULES

2027-03-05

SUBMISSION DEADLINE

The gate to the four field missions

2027

S-URC27-RULES

36

TEAMS INVITED

In practice the field expanded to 37 and 38

2023–2026

S-URC23-SCORES

2.8 %

SHARE OF FINALS VARIANCE

Against 30.6 % for Autonomy

2022–2026

S-URC22-SCORES

Key takeaways

- **The gate matters more than the score.** The review is 20 % of the maximum, and only the top 36 scores reach the field. A good package cannot win the competition; a weak one ends the season at a desk.
- **It is the least discriminating element.** A standard deviation of 4.5 on a 100-point score, and 2.8 % of finals-total variance, against 16–28 points of spread on each field mission S-URC22-SCORES S-URC26-SCORES .
- **The cut sits within a few points of the median.** In 2023–2025 the admitted set stopped between **3.4 and 5.8 points** below the field median. The 2026 numbers are not comparable, because the lowest admitted score fell about nine points in the year the rubric was rewritten.
- **The strongest teams run a production line.** West Virginia averaged **94.61** and Missouri S&T **94.23** over five seasons, while mid-field teams swung 16–17 points.

What the review demands

Two components, both mandatory: a written report and a video, submitted as one package, with original content and full attribution. URC takes plagiarism seriously enough to make it grounds for disqualification S-URC26-SAR . The 2027 rulebook fixes the character of the review: “The SAR will focus on demonstration of the capabilities and customization of the rover systems/sub-systems to perform each of the missions.” The review is competitive against other teams, and the 36 highest scores are invited S-URC27-RULES §3.a.iii.

The 2027 rubric itself is not published, and the page that will carry it promises a release in early 2027 S-URC27-SAR . The preparation that survives an unpublished rubric is evidence plus traceability, not narrative.

How 100 points enter the final score

Each mission and the review are worth 100 points for a total of 500 S-URC27-RULES §1.a.i. Penalties are additive and a mission cannot score below zero. So the review is 20 % of the maximum, and it is the only part of the score that can be earned before the rover works.

What the cut actually looked like

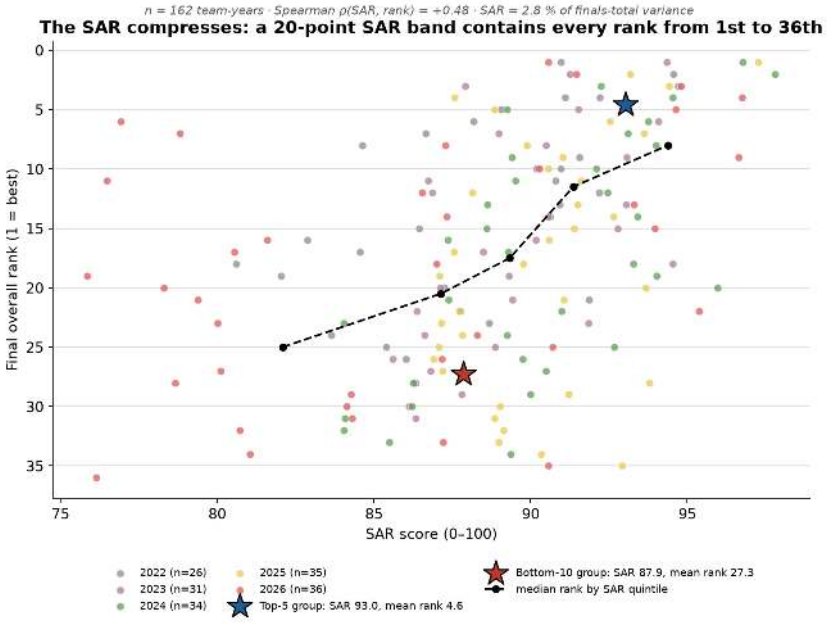
The official result sheets carry a review score for every finalist, which makes the lowest admitted score visible even though rejected packages are not published. That gives an upper bound on the cut line.

YEAR	LOWEST AMONG TEAMS THAT COMPETED	MEDIAN	HIGHEST	STANDARD DEVIATION
2022	80.60	88.42	94.71	3.71
2023	85.61	89.31	94.54	2.71
2024	84.04	89.87	97.80	3.57
2025	86.90	90.33	97.27	2.54
2026	75.86	86.77	96.75	6.39

Derived from the review column in the five result sheets `S-URC22-SCORES` `S-URC23-SCORES` `S-URC24-SCORES` `S-URC25-SCORES` `S-URC26-SCORES`. *Among all 37 teams admitted in 2023, the lowest was 85.10, for a team that was admitted and did not compete.*

Three readings matter for planning.

- **The cut line sat 3.4 to 5.8 points below the median on the 2023–2025 rubrics.** A first build needs to be close to the field median to be safe, by inference from the table.
- **2026 is not comparable.** The lowest admitted score fell about nine points in the year the requirements were described as significantly updated `S-URC26-SAR`, and the standard deviation nearly doubled. A 2027 package should be benchmarked against its own rubric coverage, not against 2025's 86.90.
- **The nominal field is 36 and the real field is larger.** The field was expanded to 37 in 2023 and 38 in 2024, 2025 and 2026, and there is no stand-by list for withdrawals. The design target is a cut at 36, without relying on the expansion.



Every rank from first to 36th sits inside a roughly 20-point band of review scores. The review moves a team from one tier to another only at the margin, which is exactly what a gate looks like in the data `S-URC22-SCORES` `S-URC26-SCORES`.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-URC-RESULTS` *Full resolution*

What the field rewards

The 2026 rubric is the best available model of what 2027 will reward, and its weights are published.

COMPONENT	WEIGHT	SUB-CRITERION	SUB-WEIGHT
Report	**40 %**	System overview and architecture	8 %
		Integration and testing	6 %
		Project management, including a Gantt chart and a budget table	11 %
		Science plan, one page	15 %
Video	**48 %**	Rover mobility and terrain	8 %
		Arm and end effector	8 %
		Autonomous navigation	6 %
		Science payload	6 %
		Overall integration	12 %
		Communications, C2 and operations	8 %
Judges' assessment	**12 %**	Readiness to compete by late May	4 %
		Ruggedness for the MDRS environment	4 %
		Novelty and scope of work	4 %

2026 rubric `S-URC26-SAR`, the only published model; the 2027 rubric is promised for early 2027. The weights total 100 %.

The consequence is structural. Forty-eight of the 100 points are won or lost in a five-minute video, the written report is worth 40, and a subjective judges' assessment is worth 12. The video is also where the review is easiest to lose, because of a single published rule: if something is not specifically shown, the judges cannot award points for it `S-URC26-SAR`. A zoomed shot of an end effector with no rover in frame earns nothing, and a separate demonstration of each subsystem does not show integration.

The 2026 format requirements are a usable template for 2027.

- **The report is seven pages**, with pages 1–4 free, page 5 the Gantt chart only, page 6 the budget table only and page 7 the science plan only. Nothing after page seven is considered.
- **The video is five minutes or less**, publicly accessible, with a month and year caption on every clip. Footage from before September of the current cycle is treated as weak evidence of maturity, and a clip of a superseded architecture is treated as inapplicable.
- **Audio quality matters more than video quality**, in the judges' own words, and music is discouraged.
- **The science page is read by scientists and the rest by engineers.** The judges state that science judges may only skim the non-science pages and non-science judges may only skim the science page, so the video should show the systems that execute the science rather than explain the science `S-URC26-SAR`.

Design implications for a first build

DECISION	OPTION A APPLIES WHEN	OPTION B APPLIES WHEN
Documentation framework	**One spreadsheet and one page of traceability** for the first season	A model-based systems framework only if the team already uses one; nothing larger fits the page budget
Evidence strategy	**Test reports, failure logs and mission footage generated during the build**	A last-month writing sprint, which the 2026 revision punished for teams with no artefacts
Science content	**About one page of science plan plus video time**, expanding into the May plan S-URC27-QA AQ1	A full plan inside the review, which is wasted effort because it is rewritten in May
Modularity story	**One integrated rover with a two-bolt end-effector swap**	A front-half payload swap, only with a large team and a spare chassis
Mass evidence	**Allocation at concept, a CAD roll-up at the preliminary review and a scale at integration**, with the 1.2 m box built early	Weighing only at the end

A final Rover Expense Report in work-breakdown form is due **2027-05-26**, and the fielded rover is capped at **$26,000** S-URC27-RULES §3.c.ii §3.c.iii. The report is a different document from the project budget table inside the review, and it covers only the rover as the rules define it.

Minimum viable versus stretch

- **The minimum that clears the gate:** a package that maps one to one onto the four missions, shows at least one measured test result per mission family, carries the one-page science plan, and states honestly what does not yet work.
- **The stretch:** a traceability figure, a mass budget with margin, a test matrix, failure logs with corrective actions, and footage of every mission attempted on analog terrain. That is the West Virginia and Missouri S&T pattern.
- **The timing trap:** the May science plan expands the one-page version submitted in March, so the writing starts twice and the second version is the scored one in Astrobiology.

Risks to retire early

RISK	WHY IT IS REAL	THE COUNTERMEASURE
The rubric is replaced	The 2026 standard deviation of 6.39 and the nine-point drop in the cut show what a rewrite does to a field	Preparation against evidence of capability, which every version of the rubric has rewarded
A strong rover, a weak package	The cut sits inside a ±3-point band, so a strong team can miss the field with 400 points of rover on the floor	Evidence built during the build, with the package reviewed monthly against the rubric
The science plan owner cannot speak at the mission debrief	AQ8 limits the mission debrief to members who stayed in C2	The plan owner and the debrief speaker are two different people, and the plan owner stays in C2 all day
The video shows CAD and talk	The rubric rewards demonstrated capability and the judges discount renders	Footage on analog terrain, with the rover integrated and every clip captioned with month and year
Format penalties	Page limits, page roles and video length are hard limits	A format check against the published template before submission

Further reading

- Review, judging and awards carries the full rubric, the format rules and the award structure.
- Systems engineering and integration covers the artefacts the review asks for.

- 155 engineering requirements is the traceability table the report is meant to show.
- What 2027 invalidates explains why two new requirements must be demonstrated in the review by March.
- Deadlines and milestones places the review among the fixed dates for the season.

MISSION BACKFIT · CHAPTER 6.7

What 2027 invalidates

Most of what the field learned between 2022 and 2026 still applies to 2027. Three changes genuinely break earlier work: the mandatory bioassay, the key-and-lock and hose hardware, and the commanded astronaut interaction. Everything else carries forward with a raised bar, and the schedule consequence is larger than any single task change.

3

GENUINELY INVALIDATING CHANGES

The bioassay definition, the Equipment hardware, the astronaut commands

2027

S-URC27-RULES

2027-03-05

WHEN TWO OF THEM MUST BE SHOWN

The bioassay payload and the astronaut interaction

2027

S-URC27-RULES

94

THE 2026 SCIENCE SCORE THAT FAILS 2027

Amino-acid chemistry alone does not meet AQ9

2026

S-WVU-2026-SCIPLAN

29–35

THE EQUIPMENT MEAN THAT CARRIES FORWARD

The five-season base rate for that mission

2022–2026

S-URC22-SCORES

Key takeaways

- **The changes are narrower than they look.** Three requirements are genuinely new. The rest raise a bar on capabilities the field already exercises.
- **A documented counter-example settles the payload question.** A payload that scored **94/100 in 2026** fails the 2027 bioassay definition without a hardware change, while the payload that scored **96 in 2022** complies S-WVU-2026-SCIPLAN S-URC27-QA AQ9.
- **Three tasks disappear and free time**, and one window appears: the 10-minute transition between Autonomy and Equipment.
- **The schedule is the real invalidation.** The bioassay and the astronaut interaction must be demonstrated in the review by 2027-03-05, three months before mission rehearsals can start.

The invalidation matrix

2027 CHANGE	WHAT IT INVALIDATES	WHAT IT LEAVES VALID	CONFIDENCE
Two life detections, at least one a bioassay S-URC27-RULES S-URC27-QA §1.b.v AQ9	Amino-acid-only payloads and mineral spectroscopy as the sole life detection, including the West Virginia 2026 ninhydrin channel	Site documentation, depth and mass sampling, the cache, the debrief and wet biomolecule chemistry	High: a documented 94-point counter-example exists
The stratigraphic profile is removed	The 2024–2026 habit of spending site time on a hillside photograph	Everything else in documentation	High
The Delivery marked-path task is removed	Path-following rehearsals	Terrain capability, search, grasp and the GNSS approach	Medium: the terrain list is unchanged
Delivery stage times and the task list are pending	Rehearsals timed to the 2026 sequence of 35 and 20 minutes with a 40-point gate	The capability set	High for planning, none for design
A key is added, USB is removed and the hose gasket is halved	Fixtures built for USB insertion and any 2026 tuning of the hose fitting	Everything else, while the arm precision requirement rises	High
Autonomy is restructured into astronaut and route sub-missions	The 2026 mission sequence, and navigation-only autonomy as a scoring strategy	AR-post navigation, RTK waypoints, object detection and the LED protocol, which is relaxed	High
A 10-minute transition between Autonomy and Equipment S-URC27-RULES §1.a.iv	Any plan that changes mass or configuration between the two	A single configuration that needs no re-weigh	High

The three that genuinely break earlier work

1. The bioassay definition. Two life detections are required, and at least one must be a bioassay that detects a molecule formed only by biology or in its presence S-URC27-RULES S-URC27-QA §1.b.v AQ9. Amino-acid chemistry alone does not satisfy it, and neither does a mineral spectrum. A payload that scored 94 in 2026 therefore fails in 2027 without a hardware change S-WVU-2026-SCIPLAN , while a three-channel wet-chemistry payload from 2022 was already compliant.

2. The Equipment hardware. The key in a tubular pin-tumbler lock has no URC precedent, and its 1 × 1 cm tags are the smallest targets in the rulebook. The hose moves from a fitting with the gasket removed to a 1.5-inch cam lock with a half-thickness gasket, which is a stiffer push than 2026's. The USB Type-C task is gone S-URC27-RULES S-URC26-RULES §1.d.ii §1.e.ii.

3. The commanded astronaut interaction. Follow!, Stay!, Fetch!, Come! and Give! are new as a command loop S-URC27-RULES §1.e.iv §1.e.x. The geometry around them has nine years of precedent, but no earlier rulebook contains a command gesture or a follow clause.

What carries forward

COMPONENT	2022–2026 APPROACH	STATUS IN 2027
Site documentation	Panorama, close-up and GNSS workflow	**Carries forward**, slightly easier without the stratigraphic profile
Sample acquisition	Auger or drill with depth evidence	**Carries forward**, unchanged since 2024
Cache and handover	A sealed tube, removable by hand	**Carries forward**, unchanged
Debrief	Includes the stratigraphy discussion	**Raised bar on the biology**, with less geology to hide behind
Terrain capability	Six passive wheels, printed treads, skid steer	**Carries forward**, still needed to reach the objects and the sites
Search and grasp	Wrist camera, object detector, wide jaw	**Carries forward**
Autonomy stack	RTK, waypoints, a terrain map, a conservative planner	**Carries forward and gets easier**, because the targets are exact coordinates
Arrival and hands-off discipline	LED and no operator input during autonomy	**Carries forward and is relaxed**, because the LED is now sufficient
Equipment routine tasks	Buttons, switches, knobs, valve, latch, drawer, cache	**Carries forward**
One configuration across missions	A rover that does not change between missions	**Becomes more valuable**, because of the 10-minute transition

What the removals give back

- **The stratigraphic profile** leaves documentation and debrief, which frees one to three minutes of site time and one debrief bullet S-URC26-RULES S-URC27-RULES §1.b.iii §1.b.iii.
- **The Delivery marked-path task** removes one path-following family and its rehearsals. It does not reduce the terrain requirement.
- **The USB Type-C task** deletes one bankable Equipment task and all the fixtures and poses built for it.
- **The 10-minute transition** is not a removal but a constraint: any change in mass or configuration between Autonomy and Equipment forces a re-weigh inside the window S-URC27-QA EQ4.

What looks invalidated and is not

- **A 2022 payload is not obsolete.** The wet three-channel payload that scored 96 in 2022 was closer to the 2027 clause than the Raman instrument that replaced it in 2024.
- **Mission scores are still evidence of capability.** They transfer as evidence of documentation, sampling and debrief skill, even where the payload behind them does not transfer.
- **The 2026 autonomy stack survives its mission.** The 2026 mission structure is invalidated, and the stack that scored 86 in it is not.

- **The 2026 discipline carries over.** The LED and the hands-off rule remain, and the C2 banner becomes optional rather than required `S-URC27-QA AUQ18`.
- **The Equipment base rate is unchanged.** The five-season mean of 29–35 remains the best estimate for 2027, with the key task as extra downside and the removed USB task as a small offset.
- **Terrain capability still decides Delivery.** The marked path is gone, but the objects still sit across sand, gravel, boulders and steep loose slopes.

What this means for the rover

1. **The design follows the score shape, not the mean.** The guaranteed-floor work, which is documentation, sample, cache and debrief in Astrobiology and the suit plus short behaviours in Autonomy, is protected and never traded for 1 m route targets, gesture Follow! or the key.
2. **The two 2027-specific payloads are frozen before the robotics.** The bioassay pair and the astronaut interaction both need review evidence by 2027-03-05, while the drive and the arm can be tuned until June.
3. **"Two life detections" is a payload-architecture decision.** Two orthogonal biomolecule assays share one detector, with the AQ9 justification written into the Science Plan.
4. **The geometry the rules give away is worth taking.** The astronaut's GNSS point, the 20 m standoff with the astronaut in view, and exact route coordinates convert hard perception into easy localisation.
5. **Follow! is built as a ladder, not a modality.** Gesture, then audio, then a sign or a device, so the floor is 5 points and the ceiling is 15.
6. **Equipment is a precision problem with a time budget.** The cheap tasks come first, because any order is allowed, then the drawer and cache, then typing, the key and the hose.
7. **Autonomy stays boring and the stop discipline excellent.** The best-documented 2026 gain came from replacing a sophisticated planner with a simpler one.
8. **A link check never gates autonomous driving.** The second route target is deliberately out of radio sight, and one link check is a 25-point failure.
9. **The transition and the weigh-in are rehearsed as mission tasks.** The 10-minute window, the 1.2 m box and the 50 kg limit can cost more than a weak arm, and they are cheap to practise.
10. **Documentation is generated as the system is tested.** The review is 100 points and the only route to the field, and the teams that average 94 have the artefacts in hand when the deadline arrives.

Further reading

- Mission-by-mission backfit carries the master task table and the score shapes behind each verdict here.
- Astrobiology mission works the payload change in detail.
- Equipment servicing works the key, the hose and the removed USB task.
- Autonomy: astronaut assistance works the command interaction.
- System Acceptance Review explains why the March deadline governs the build order.
- 110 engineering problems turns each item above into a task with an acceptance test.

MISSION BACKFIT · CHAPTER 6.8

Coupling, conflicts and twelve trade studies

Most 2027 decisions belong to one subsystem and one group of people. A smaller set cannot be taken inside any single subsystem, because the number that settles one question sets the budget for another. Arm mass chooses the drivetrain actuator class, wheel diameter sets the stow envelope, the centre of gravity and the torque demand at the same time, and the battery is sized by a peak current that appears only when the arm stalls on a hill. This page works the twelve trades that cross those boundaries, gives the recommended starting point for each, and names the measurement that would overturn it.

12
TRADE STUDIES
Options, consequences, a recommendation and a flip condition

16
NAMED CONFLICTS
Each mapped to one trade and one owning group

46.0 kg
PREDICTED WEIGH-IN
A 40.0 kg allocation plus 15 % growth

5 %
MASS PENALTY PER KILOGRAM
Of the points scored, per event, above 50 kg

Key takeaways

- The 50 kg cap is consumed by single decisions, not by their sum. One arm class adds **6.0 kg**, one wheel construction adds **7.5 kg**, and a change of cell chemistry adds **3.5 kg**, against a reserve of **4.0 kg**.
- Wheel diameter is the least reversible number in the rover. Moving from 250 mm to 510 mm costs **12.9°** of static tip angle, against **5.6°** for moving from a 4.4 kg arm to a 14 kg arm.
- A 300 mm wheel is not 30 mm of free clearance. At 45 kg and 30° it needs **9.31 N·m per wheel**, three per cent over the **9.0 N·m** rating of the common actuator class, before any shock or slip margin.
- A teleoperation-only rover caps the season at **405 of 500 points**. The first **15** autonomy points are built from hardware the rover needs anyway; the last **50** are a software programme.

How a trade is written

Each trade follows the same five parts: the options, the consequences with their numbers, a recommended starting point, the condition under which the recommendation flips, and the test that settles it. A recommendation without a flip condition is a preference, and a preference is what a design review turns into an argument.

The arithmetic that makes mass the dominant coupling is published in the rules. The mass penalty is **5 % of the points scored per kilogram over 50 kg, per event** S-URC27-RULES §2.a.iii. One kilogram over costs **4 points in an 80-point event**, in every event of the season. The oversize penalty is **40 % of an event's points**, which is 32 points at 80 and is equivalent to eight kilograms overweight. Weight rounds down, so 50.9 kg records as 50 kg and 51.0 kg costs the full 5 %, which is why the design target is **48.5 kg** rather than 50 kg. A human intervention costs **20 % of a mission's points**, and a hand battery swap is one S-URC27-RULES §3.e.vii.

Finishing beats dieting, and the stow envelope beats both per unit of violation. The trades below are ordered by how much they cost to reverse, not by subsystem.

The sixteen conflicts

ID	CONFLICT	TRADE	OWNING GROUPS	FROZEN BY
CF-01	Arm mass class against drivetrain margin	T-01, T-02	Arm, mobility, systems	PDR
CF-02	Arm reach against vertical stow budget and typing precision	T-02	Arm, systems, perception, software	PDR and SAR
CF-03	Wheel diameter against the cube, the centre of gravity and torque	T-03	Mobility, systems, power	**PDR, irreversible**
CF-04	Mast height against camera stability and the cube	T-06	Communications, perception, systems	PDR
CF-05	Camera count against link budget and operator workload	T-11	Perception, communications, operations	PDR
CF-06	Compute power against thermal load, volume and the battery bay	T-05	Perception, power, environment, systems	PDR and SAR
CF-07	Bioassay against arm stiffness and the 25-minute clock	T-09	Science, arm, environment, operations	**December 2026**
CF-08	Modularity for 70 kg against the 10-minute transition and staffing	T-08	Systems, operations, arm	PDR
CF-09	E-stop scope against communications and logging	T-12	Power, communications, software, operations	PDR and SAR
CF-10	Drone mass and accounting against 70 kg and the E-stop reading	T-07, T-12	Delivery, systems, power	PDR
CF-11	Whether the Equipment Servicing cache is the team's own Astrobiology cache	T-09	Science, arm	December
CF-12	Astronaut link topology	T-11, T-06	Astronaut, communications, operations	PDR
CF-13	GNSS base survey-in against the 15-minute setup	T-06, T-11	Perception, operations	SAR
CF-14	900 MHz sub-band agility against the three-link architecture	T-11	Communications, delivery, software, astronaut	2027-04-30
CF-15	Battery peak current against energy and mass	T-04	Power, mobility, arm	CDR
CF-16	Teleoperation-only score ceiling against autonomy effort	T-10	Software, perception, operations	PDR

Mass allocation for a 50 kg cap

T-01 is the decision that the other eleven spend. The 50 kg cap, the 70 kg fielded-parts total and the 1.2 m cube are three budgets, not one, and no team in the public record publishes a mass budget. The allocation is therefore a design act rather than a measurement, and it is the first thing the PDR freezes.

WORK PACKAGE	ALLOCATION (KG)	STRETCH (KG)	WHAT IT BUYS
Mobility: chassis, suspension, six wheels, six motors, controllers	15.0	18.0	250–300 mm wheels, six driven, no steering actuators
Arm: 5–6 degrees of freedom, about 1 m, gripper, wrist camera, harness	8.0	14.0	5 kg payload, 1.5 m reach, millimetre typing with visual servo
Science: drill or auger, fluidics, two assays, cache	5.0	6.0	Two assays on one detector, body-mounted drill
Power: pack, BMS, distribution, converters, harness	5.0	6.0	8S4P pack of about 500 Wh with contactor and precharge
Compute and perception: computer, cameras, IMU, GNSS	2.5	3.5	15–25 W compute, four cameras, RTK, no LiDAR
Communications and mast: two radios, antennas, masts, C2 antennas	1.5	2.5	2.4 GHz plus 900 MHz on a two-mast architecture
Structure, fasteners, cable, miscellany	2.0	3.0	Three to five per cent of the rover, budgeted rather than discovered
Deployables: repeater, remote camera	1.0	2.0	A 900 MHz repeater; the drone is excluded from the 50 kg S-URC27-RULES §2.b.vii
Subtotal	**40.0**		
Growth at 15 %	6.0		Unbuilt structure, brackets, parts not yet designed
Predicted weigh-in	**46.0**		
Reserve to 50 kg	4.0		The last six weeks of brackets, tape and a second E-stop
Design target, not to exceed	48.5		1.5 kg for scale, dust and rounding

One change at a time shows how thin the reserve is. The table below starts from the 46.0 kg prediction and moves one decision at a time; the penalty column is the cost at an 80-point event.

CHANGE FROM THE 46.0 KG PREDICTION	CHANGE (KG)	PREDICTED (KG)	OVER 50 (KG)	PENALTY PER EVENT	FOUR EVENTS
Arm 8.0 to 14.0 kg, the ARES class S-ARES-PAPER-2025	+6.0	52.0	2.0	8.0	32.0
The same, plus the 2.7 kg actuator step at the 45° cliff	+8.7	54.7	4.7	18.8	75.2
Wheels 0.6 to 1.85 kg each, six of them S-MONASH-GH-NOVA	+7.5	53.5	3.5	14.0	56.0
Lithium iron phosphate at 840 Wh	+3.5	49.5	none	none	none
Science 5.0 to 8.0 kg	+3.0	49.0	none	none	none
Thermal stack: shell, shade, fan	+1.5	47.5	none	none	none
Arm 8.0 to 4.4 kg, the RSX class S-RSX-ARM-URDF	−3.6	42.4	none	none	none

The arm class also interacts with the drivetrain, which is why CF-01 owns two trades. A fixed-point model that closes the published arm classes at a dry mass of 25.0 kg plus the arm plus a 1.229 kg pack lands at **30.6 kg, 35.7 kg and 40.3 kg** for arms of 4.4, 9.5 and 14.0 kg. The 14 kg arm therefore sits **0.3 kg below the 40.6 kg** at which six 9.0 N·m actuators stop being adequate on a 45° loose slope. One extra bracket then forces the larger actuator class, **2.7 kg for six**, and the rover is at 43.0 kg with 7.0 kg of headroom instead of 9.7 kg. The arm decision silently consumes the drivetrain margin.

Three of those single changes cross 50 kg on their own, and each one is a decision taken inside a different subsystem. That is why the arm class is settled before the motor class, and why the allocation is a design decision rather than a sum.

The field explains why the target is 48.5 kg and not 50 kg. Median published competition mass is **49.1 kg**, and five of seventeen configurations sit between 49.1 and 50.0 kg, three of them exactly at 50.0 kg. That is a pile-up against the line, not an optimum. The decisive outlier is WVU, which won in 2023 at **33 kg** with a 5-DoF arm and 425.35 points S-WVU-SAR2023 S-URC23-SCORES ; fifty kilograms is one and a half championship rovers.

> **Recommended allocation**
>
> Allocate **40.0 kg**, predict **46.0 kg** after 15 % growth, hold the not-to-exceed value at **48.5 kg** and keep **4.0 kg** under the cap. Weigh the bare chassis, each subsystem and the integrated rover at three points, and weigh the heaviest legal configuration. USU's 2025 rover is the failure mode: 46.9 kg in its science configuration and 51.8 kg in its arm configuration, a 4.9 kg spread with one configuration over the line S-USU-2025-POSTER .

Flips when the predicted mass at PDR exceeds 48.5 kg. Cut in this order: the deployables and the second radio, then the arm from 9.5 kg to 5 kg or less, then composite wheels to printed wheels, then two wet assays to one wet and one optical. The thermal stack and the reserve are not cut, because both are paid for later at a worse rate. The drone sits outside the 50 kg; how it counts against the 70 kg total is not published. **Settled by:** a scale, a spreadsheet and three weigh dates.

Arm length against stow and typing precision

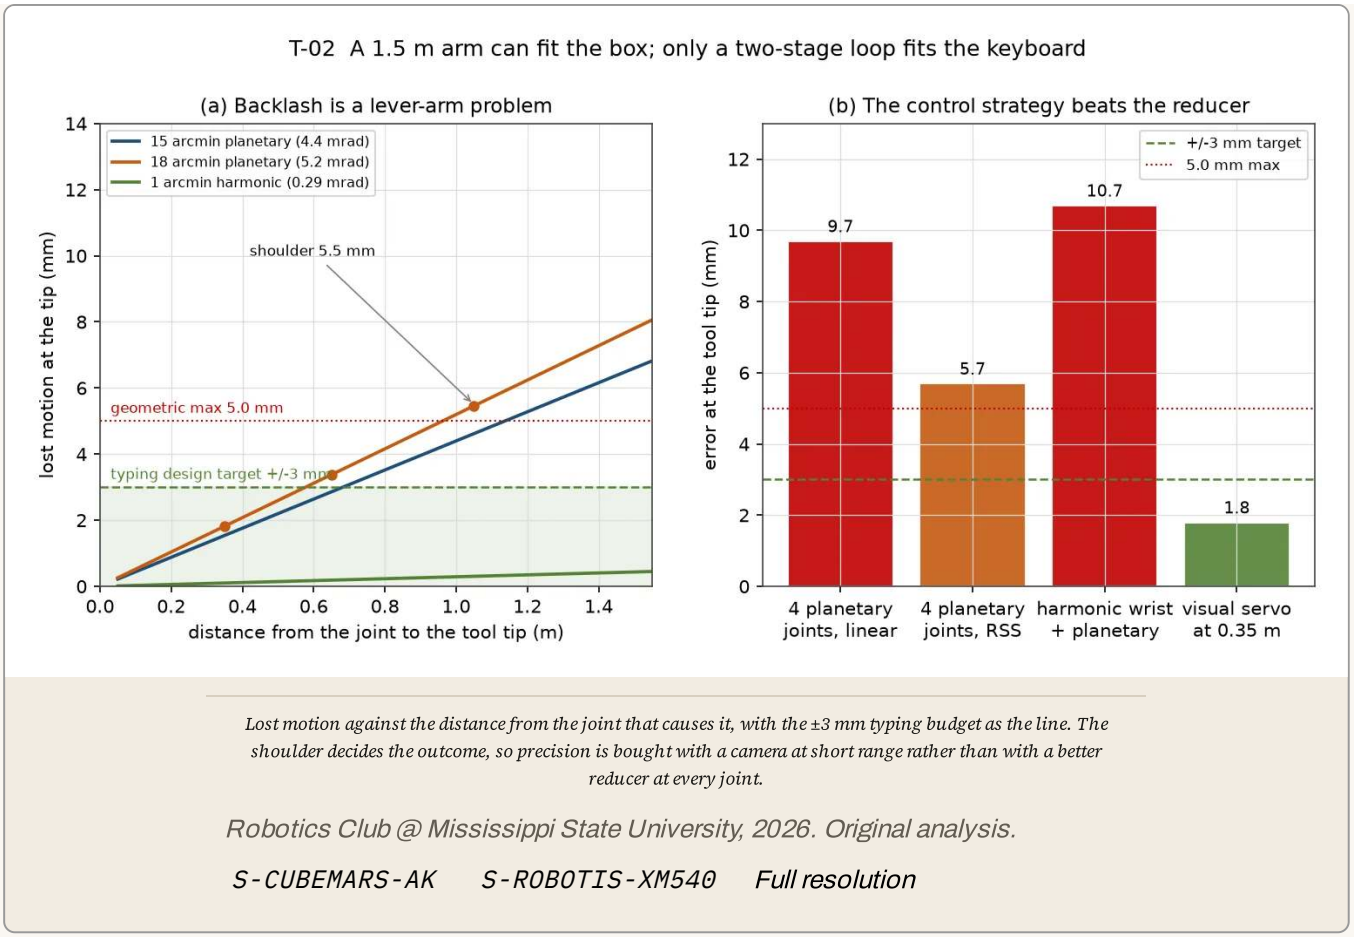
T-02 resolves the conflict between a target that is **1.5 m above ground** S-URC27-RULES §1.d.i REQ ARM 001 and a cube of **1.2 m** S-URC27-RULES §2.a.ii. Horizontally the two are compatible. Vertically they are not, and the fold that solves the vertical problem is the mechanism that destroys the precision the keyboard needs.

With 250 mm wheels the deck sits at **0.41 m**, so the arm must cover **1.09 m above the deck** and **0.41 m below it**; the shoulder therefore belongs near the front face. A 1.05 m arm folds into a **0.55 m** plan footprint, and the dimension that actually binds is the parked vertical thickness of **0.12–0.15 m**. The stow problem is flat, and every end-effector that hangs below the forearm breaks the cube.

Backlash converts into tip error through the lever arm, and the shoulder owns the lever. Integrated actuators quote **15–18 arcmin**, or **4.4–5.2 mrad**, in their datasheets S-CUBEMARS-AK S-ROBOTIS-XM540 .

JOINT, FIVE-AXIS ARM	LEVER TO TIP (M)	ERROR AT 4.4 MRAD (MM)	ERROR AT 5.2 MRAD (MM)	ERROR AT 1 ARCMIN (MM)
Shoulder	1.05	4.6	5.5	0.30
Elbow	0.65	2.9	3.4	0.19
Wrist pitch	0.35	1.5	1.8	0.10
Wrist roll	0.15	0.7	0.8	0.04
Linear sum		9.7	11.4	0.6
Root sum square		5.7	6.7	0.4

The typing budget is **±3 mm design and 5.0 mm geometric maximum**, from a 13–14 mm keycap top and a 3 mm tip on a keyboard at 19.05 mm pitch S-URC27-TASKHW . Open-loop planetary joints are two to four times over that budget before any vision error is counted, and fitting a harmonic wrist alone still leaves 10.7 mm because the shoulder dominates. Closing the last **30 mm** with a wrist camera at a 0.35 m standoff turns 5.2 mrad into **1.8 mm**, inside the budget. Missouri S&T aims a laser and presses with a linear axis, BYU uses a visual-servo library, and ARES detects the keyboard with a learned model S-MST-GH-ROVECOMM S-BYU-SAR26 S-ARES-SAR26 .



Lost motion against the distance from the joint that causes it, with the ±3 mm typing budget as the line. The shoulder decides the outcome, so precision is bought with a camera at short range rather than with a better reducer at every joint.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-CUBEMARS-AK *S-ROBOTIS-XM540* *Full resolution*

The structure pays for reach as well. The shoulder passes **176 N·m peak and 78 N·m static** into the chassis, so a folding elbow must carry that through a lock rather than through a continuous member. A stalled joint draws **20 A and 140 W** and reaches 60 K above ambient in four to five minutes, which is a thermal event inside a task that already has a clock on it.

Recommended starting point

A **6R serial arm with 0.95–1.05 m links**, the shoulder at 0.45–0.55 m from the front face, one fold at the shoulder and elbow laying the arm along the deck, a **body-mounted drill**, and a **wrist camera** running a two-stage loop: move to 20–30 mm from the key, then servo the last 30 mm and press with compliance. No telescoping arm and no body lift in a first build.

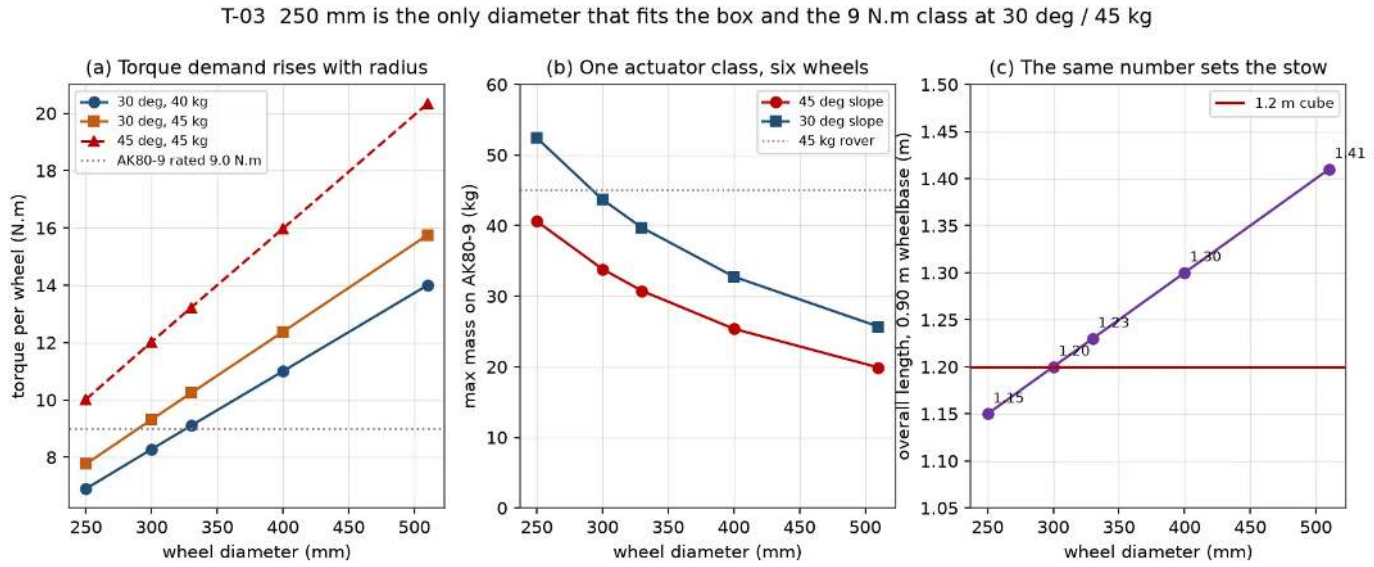
Flips when the task must be typed open loop, which puts a harmonic or cycloid drive at the shoulder where the 1.05 m lever is; when the parked budget falls below **0.12 m**, which removes anything hanging below the forearm; or when the arm must be 5 kg or less, which is the RSX class at **4.42 kg and 0.87 m** with reduced payload. **Settled by:** a full-scale lander panel, all nine Equipment poses, and the folded envelope measured with a tape.

Wheel diameter, the four-way coupling

T-03 is the only number that changes obstacle climbing, centre of gravity, stowed size and drivetrain torque at once, and it does so in conflicting directions. Once the wheels and gear ratios are made it is **irreversible**, which is why it is frozen at the PDR rather than at the critical design review.

WHEEL DIAMETER	LENGTH AT 0.90 M WHEELBASE	CENTRE OF GRAVITY (M)	STATIC TIP ANGLE	SIX PRINTED WHEELS (KG)	TORQUE PER WHEEL AT 40 KG, 30° (N·M)	MAXIMUM MASS ON SIX 9.0 N·M ACTUATORS, 30° AND 45°	VER
250 mm	**1.15 m**	0.258	**61.5°**	**6.0**	6.9	52.4 and 40.6 kg	**Reco**
300 mm	1.20 m	0.289	58.8°	8.6	8.3	43.7 and 33.8 kg	Only diet c actua
330 mm	1.23 m	0.307	57.1°	10.5	9.1	39.7 and 30.8 kg	Fails 0.90
400 mm	1.30 m	0.350	53.6°	15.4	11.0	32.8 and 25.4 kg	Need short
510 mm	1.41 m	0.418	**48.6°**	**25.0**	**14.0**	25.7 and 19.9 kg	Only wheg

The 45 kg example decides it. At 40 kg a 250 mm wheel needs **6.9 N·m**; at 45 kg it needs **7.76 N·m**, inside the 9.0 N·m rating. At 300 mm and 45 kg the demand is **9.31 N·m**, three per cent over the rating before shock, slip or efficiency margin. A 300 mm wheel is therefore either a different drivetrain class, at **2.7 kg for six** actuators, or a rover at 43.7 kg or less. Top speed is never the constraint: 250 mm already gives **5.1 m/s**, five times the 1.0–1.5 m/s target.



Torque per wheel against the 9.0 N·m rating, the mass six actuators can carry at 30° and 45°, and the stowed length against the 1.2 m cube. The 250 mm and 300 mm rows sit either side of the actuator rating, which is the real trade rather than the step height.

Robotics Club @ Mississippi State University, 2026. Original analysis.

`S-CUBEMARS-AK80-9` `S-URC27-RULES` *Full resolution*

The best-documented teams moved smaller for the same reason. WVU ran 20-inch whegs in 2024, 24-inch folding whegs in 2025 and returned to **20-inch wheels with no folding mechanism in 2026** `S-WVU-SAR2025` `S-WVU-SAR2026` ; Missouri S&T independently built a compact suspension that no longer needs to fold `S-MST-SITE-ATHENA` . Geometry beat mechanisms in both programmes in the same year.

Recommended starting point

250 mm printed TPU wheels with 10–15 mm V-chevron grousers, a **0.85–0.95 m wheelbase**, six wheels and no fold, on the 9.0 N·m actuator class. Choose the diameter from the box, then the wheelbase from the terrain, then the motor from the torque.

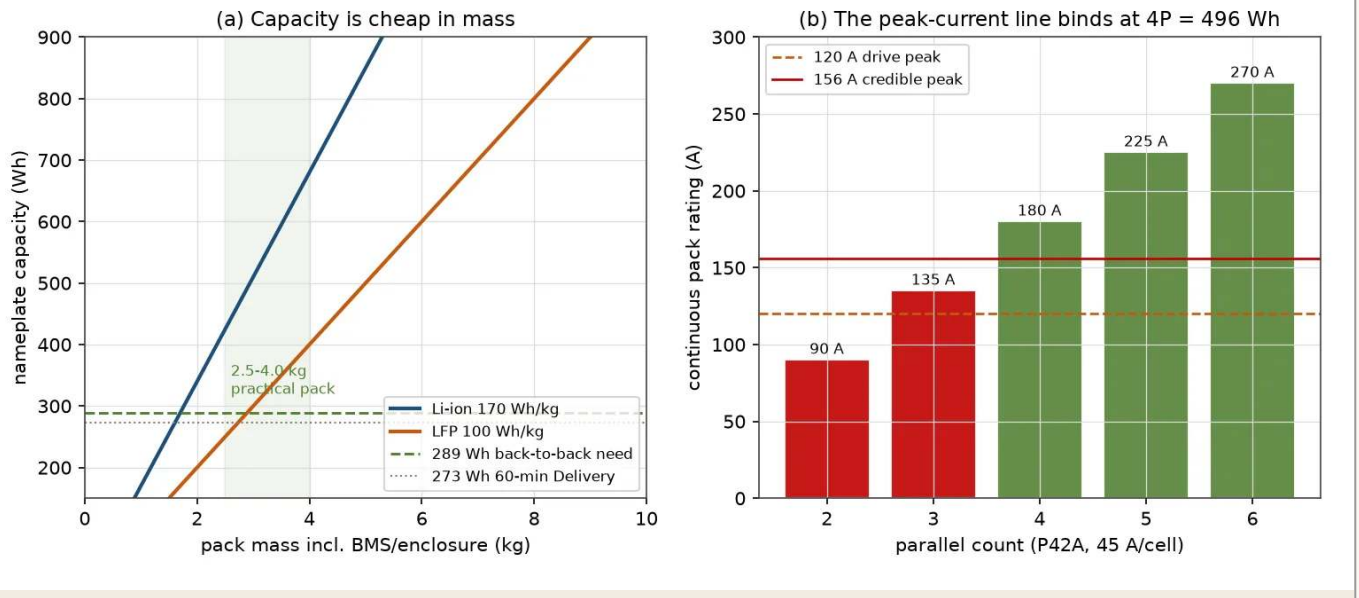
Flips when 200 mm steps or vertical faces are confirmed as the Delivery bottleneck, which points to 330–400 mm, a shorter wheelbase, a larger actuator and a fold or a diagonal stow; when a tilted or diagonal weigh-in is proven with a timber box and accepted, which stops the box setting the diameter; or when the weigh-in exceeds 50 kg, at which point 250 mm is the only diameter that saves mass without changing the drivetrain. **Settled by:** the 1.2 m box first, then a half-scale rocker rig over a rock pile, measuring step climb at 250, 300 and 400 mm before a wheel is printed.

Battery capacity, mass and the peak-current trap

T-04 is a new conflict in 2027. Energy is cheap in kilograms, peak current is expensive in parallel cells, and both land on the same pack. The mission needs are **273 Wh** for the 60-minute Delivery at 50 kg, **289 Wh** for a 40-minute autonomy run followed by a 30-minute servicing run, and **410 Wh** for a 90-minute sensitivity case. The credible simultaneous peak is **156 A:** 120 A of drive, 26 A of arm and 10 A of hotel load.

PACK	CAPACITY	CONTINUOUS RATING	WHAT IT BUYS AND COSTS
8S4P 21700	**496 Wh**	180 A	2.9 kg of cells plus 0.5–1.0 kg of BMS and enclosure; covers 289 Wh with 70 % headroom and the 156 A peak, and is the recommendation
8S3P 21700	372 Wh	135 A	About 0.6 kg less pack and 0.5 kg less BMS; below the 156 A peak unless a drive and arm interlock is proven
8S6P 21700	744 Wh	270 A	4.4 kg of cells, 1.5 kg more than the 8S4P, for energy that no 2027 mission needs
Lithium iron phosphate, 24 V	480–840 Wh	pack-dependent	Safety and deep discharge, at 1.7 times the mass per watt-hour: 8.4 kg at 840 Wh
Two 40 V 4 Ah packs	320 Wh	pack-dependent	A between-mission swap is free; 320 Wh covers 289 Wh with 11 % headroom, which is thin

The energy loop does not bind. Adding 100 Wh adds **0.59 kg** of pack, 1.4 W of drive load and 0.01 kg of pack, a gain of **0.0137 kg of battery per kilogram of rover**, so capacity is almost linear in mass. What binds is a continuous rating above **156 A.** Three parallel strings cover the 120 A drive peak but not the credible coincidence of a climb and a stalled arm; four is the honest minimum, and four lands at **496 Wh**, the same place the energy requirement lands. Missouri S&T's 80 A distribution limit with a 10 ms BMS trip fired on a legitimate peak, and their design auto-restored after three seconds `S-MST-2023-POWER`.



Capacity against mass for two chemistries, with the 273 Wh and 289 Wh needs marked, and continuous rating against parallel count against the 120 A drive and 156 A credible peaks. The two requirements meet at four parallel strings.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-VENDOR-MOLICEL-P42A-DS *Full resolution*

Recommended starting point

8S4P 21700 lithium-ion, about 496 Wh, 2.5–4.0 kg including the BMS and enclosure, with a BMS rated **150 A or more** and a **1–3 s trip delay**, precharge, a coordinated fuse, and the pack in its own compartment. A second pack travels as a spare, because spares do not count against the 50 kg S-URC27-RULES §2.a.iii.

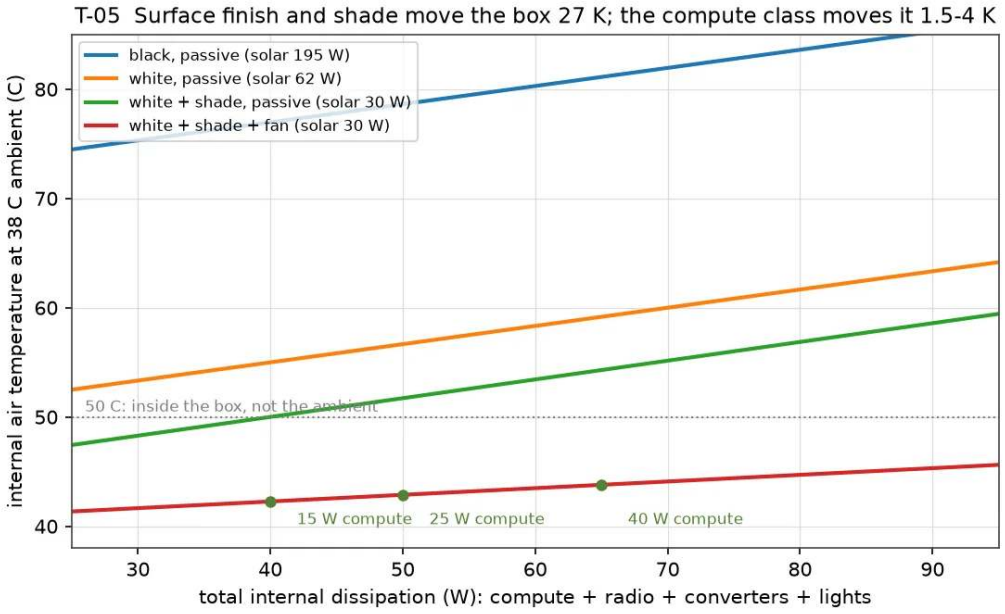
Flips when a tested drive-and-arm interlock and a stall timeout make **3P at 372 Wh** adequate, saving about 0.6 kg; when the worst case is confirmed at **90 minutes**, which points to 6P at 744 Wh and reopens T-01 by 1.5 kg; or when chemistry is chosen for safety, which adds **2.1 kg at 500 Wh** or **3.5 kg at 840 Wh**. **Settled by:** a bench discharge on the real profile and a current log of a hill climb with the arm stalling. No team publishes a measured watt-hour per mission.

Compute power against thermal load

T-05 is the trade where the cheap quantity and the expensive one are different. Compute power is cheap at 15–25 W average. Compute **heat** is expensive, and it is paid in enclosure volume, filter maintenance and a separate battery compartment rather than in watts. The box is 0.4 × 0.3 m with 0.52 m² of surface and 0.12 m² of top, in 38 °C ambient air.

TREATMENT	SOLAR ABSORBED (W)	INTERNAL AIR AT 40 W TOTAL	INTERNAL AIR AT 15, 25 AND 40 W OF COMPUTE, PLUS A 25 W BALANCE	MASS
Black, passive	195	**77 °C**	77.0, 78.6 and 81.1 °C	none
White, passive	62	**55 °C**	55.0, 56.7 and 59.2 °C	paint
White plus sunshade	30	**50 °C**	50.0, 51.7 and 54.3 °C	0.35 kg
White, shade and filtered fan	30	**42 °C**	42.3, 42.9 and 43.8 °C	0.2 kg of fan and duct

The compute class moves the box by **1.5 to 4 K**. Finish and shade move it by **27 K**. The second-order cost is the one that catches teams: if the computer shares the battery compartment, the pack sits at 50–77 °C, and a lithium-ion cell stored at 40 °C and full charge retains about **65 % of its capacity after a year** S-BU-808 . The thermal design therefore forces a separate battery compartment, which is a volume claim inside the cube. The fan itself needs **13 CFM at 1,309 m** plus 30–50 % for filter back-pressure, so 20 CFM free delivery is the design point, and the whole thermal stack is about **1.5 kg and 2 W**.



Internal air temperature for four treatments across the three compute classes. The vertical distance between the black and white curves is the argument: the finish and the shade decide the outcome, and the computer class is a small term.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-NVIDIA-JETSON-ORIN *Full resolution*

Recommended starting point

A **15–25 W Jetson-class computer**, a **white enclosure with a sunshade and a filtered fan**, heat-sunk to the skin, the **battery in its own compartment**, **1.5 kg and 2 W** budgeted for thermal hardware, and a commanded degrade mode that drops visual odometry, a camera or a network before the box derates on its own.

Flips when the enclosure must be sealed and fanless, which caps compute at **15 W** and moves heavy perception to the base station, where a white passive box sits at 50 °C; when onboard power passes **40 W**, which points to a larger box and a second fan; or when a hot-day test holds **45 °C or less** passively, which removes the fan and the filter burden. **Settled by:** a 38 °C sun test with a thermocouple in the box and a 30-minute full-stack run, repeated with the filter half-clogged.

Mast height, stow and link quality

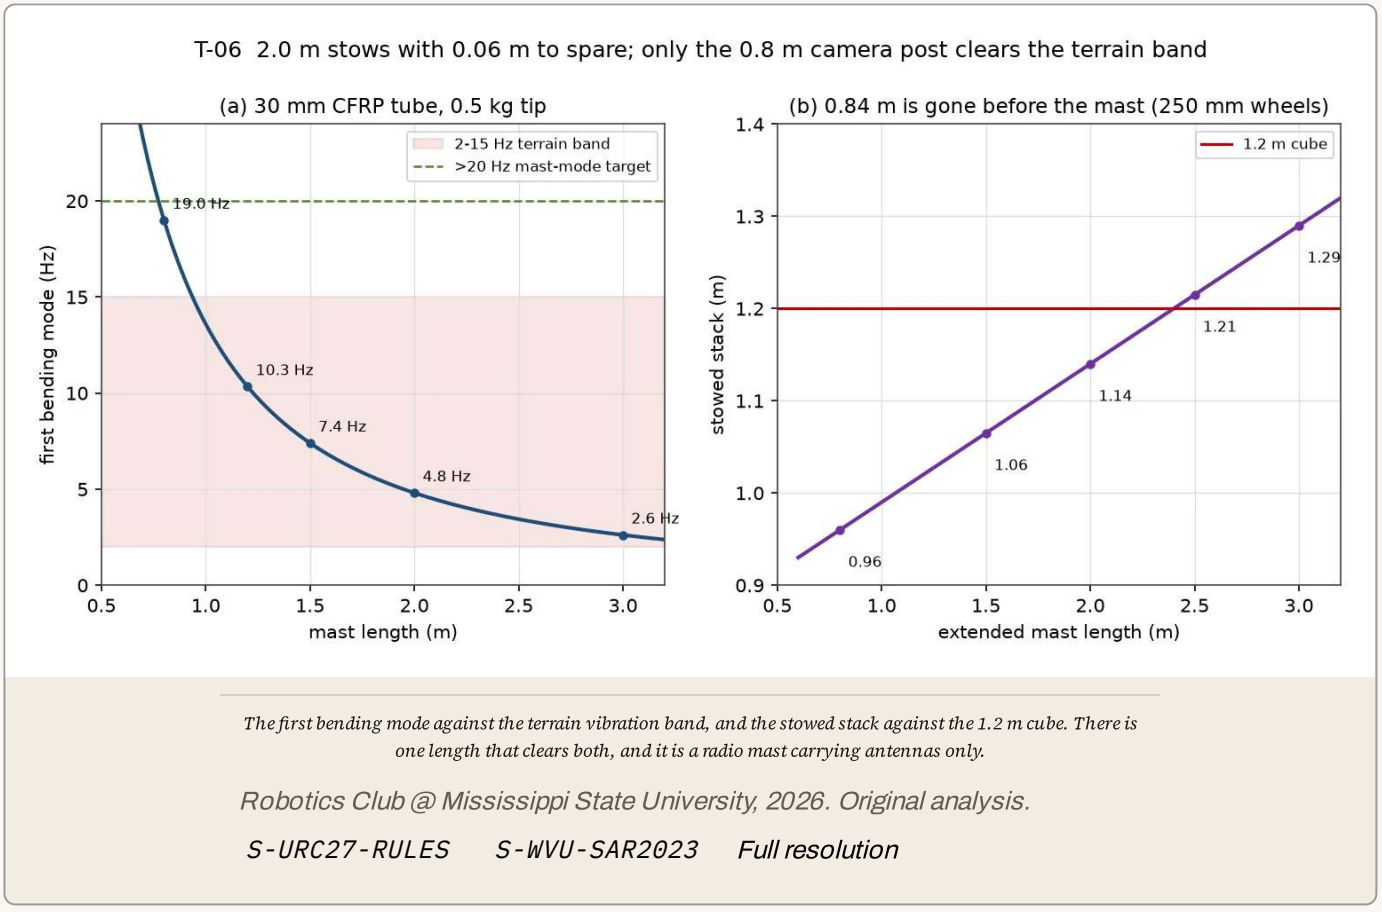
T-06 begins by correcting a rule that is widely misread. The **3 m limit applies to the base-station antenna**, which must be within 5 m of the C2 station, with guys within 10 m and at least 25 m of cable S-URC27-RULES §2.c.iv REQ COM 009. The rover mast is bounded by the **1.2 m cube**. A 2.0 m rover mast is a design choice; a 3.0 m base antenna is a rule.

The stow stack before the mast is 0.84 m: 0.255 m to the axle, 0.255 m of wheel, 0.16 m of deck, 0.12 m of parked arm and 0.05 m of antenna and gimbal. With a five-segment mast collapsing to 15 % of its extended length, the numbers run as follows.

EXTENDED MAST	0.8 M	1.5 M	2.0 M	2.5 M	3.0 M
Stowed stack	0.96 m	1.06 m	**1.14 m**	1.22 m	1.29 m
First bending mode	**19.0 Hz**	7.4 Hz	**4.8 Hz**	3.4 Hz	**2.6 Hz**
Tip deflection at 5 N	0.7 mm	4.6 mm	10.9 mm	21.4 mm	36.9 mm

Two things are visible in that table. A 2.0 m mast fits the cube with 0.06 m to spare but sits at **4.8 Hz**, inside the 2–15 Hz band that terrain excites, and a 3.0 m mast fails the cube before it helps the link. A 20 mm tube is worse still, at 9 Hz for a 0.8 m post.

Height is not the whole link argument either. At 1 km a 915 MHz path wants **5.5 m of Fresnel clearance** that no rover mast provides, while 900 MHz carries **28.2 dB of non-line-of-sight margin** against **minus 7.6 dB for 2.4 GHz.** The wind case is the last constraint: a 65 km/h gust gives a dynamic pressure of about **155 Pa**, which is 23 N·m on a 0.15 m² antenna at a 1 m lever and **78 N·m** on a 3 m mast with a 0.5 m² dish, enough to bend a 25 mm tube and to sway the tip about 170 mm.



The first bending mode against the terrain vibration band, and the stowed stack against the 1.2 m cube. There is one length that clears both, and it is a radio mast carrying antennas only.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *S-WVU-SAR2023* *Full resolution*

Recommended starting point

Two masts. A **0.8 m, 30 mm carbon-fibre camera post** at 19 Hz, and a **1.5–2.0 m telescoping radio mast** carrying antennas only, collapsing to 0.30 m and stowing at 1.06–1.14 m. The base antenna stands on its own tripod of 3 m or less at the C2 station, with at least 25 m of cable, surveyed the day before.

Flips when a 1 km non-line-of-sight test passes with the 900 MHz plane at 1.5 m, which frees 0.14 m of stow margin; when perception needs an eye height above 1.0 m, which points to a 1.2 m camera post at 10.3 Hz with isolators and a single-axis gimbal; or when the stow margin falls below 0.03 m, which drops the radio mast to 1.5 m before anyone considers a tilt. **Settled by:** a range test at 1 km line of sight and 300 m non-line of sight with the real antennas. WVU measured its own envelope at **750 m line of sight and 164 m non-line of sight** in 2026 `S-WVU-SAR2026` , so the test envelope, not the link budget, is the binding constraint.

Drone or no drone

T-07 is unusual because the answer is a null recommendation, and the argument is arithmetic rather than capability. A drone is optional and scores no points directly: full points are possible without one S-URC27-RULES §2.b.i. It buys time in Delivery and can relay a link, and it costs a C2 seat, a field spotter, a registration and a landing step that the Autonomy mission makes mandatory before the rover may proceed.

2027 CONSTRAINT	REQUIREMENT	CONSEQUENCE
Mass	5 kg or less, including a dummy mass equal to the battery S-URC27-RULES §2.b.v	Airframe and payload are 5 kg minus two battery masses
Rover limits	Drones are excluded from the rover's weight and size restrictions S-URC27-RULES §2.b.vii	Outside the 50 kg and the 1.2 m cube
70 kg total	Whether a drone counts among fielded rover parts is not published	Design as if it does not, plan as if it might
Missions	No drone in Astrobiology or Equipment Servicing S-URC27-QA RQ9	The two missions where scouting helps most are closed
Autonomy	Land at the designated pad before the rover proceeds S-URC27-RULES §1.e.xiv	Reconnaissance only during astronaut assistance
Crew	A pilot in C2 plus a designated field spotter in visual line of sight S-URC27-RULES §2.b.ii	One C2 seat and one field body out of six seats and fourteen people a day
Law	Registration above 250 g, remote identification, below 120 m, class G airspace S-URC27-RULES §2.b.vii	Paperwork and a trained pilot

The record is the argument. WVU's 2023 custom quadrotor was the primary autonomy vehicle, had **nine minutes** of flight and was in early testing at the SAR; the rover scored **56** in Autonomy S-WVU-SAR2023 S-URC23-SCORES. WVU then moved autonomy to the rover and kept a commercial quadcopter for Delivery scouting. Missouri S&T's first drone never flew at the competition, and its 2026 aircraft scouts and relays S-MST-SITE-HYPERION S-MST-SITE-ATHENA. No team publishes a drone mass, and the 2027 limit is a weighing rule.

Recommended starting point

No drone in the 2027 first build. Spend the 1.0 kg deployables line on a 900 MHz repeater if the link needs one, and spend the pilot's C2 seat on a second rover operator. The drone buys time in one mission, not points, and the scarcest resource in a first build is integration hours.

Flips when the Delivery first-stage search is confirmed as the bottleneck and rover detection range is insufficient, which points to a commercial airframe under 5 kg at the PDR rather than a custom one; when the 900 MHz non-line-of-sight link fails at 1 km, which makes a relay drone a communications asset that helps Delivery only; or when the team has eight field-ready people and a registered pilot by January. **Settled by:** a timed Delivery rehearsal with and without a spotter, and a written answer on how the drone counts against 70 kg.

Modular or single configuration

T-08 asks whether the **70 kg fielded-parts total** should become an architecture. The rule gives a 50 kg rover **20 kg** of headroom for a second configuration, and the best-documented team moved away from using it over three years.

QUANTITY	VALUE	BASIS
Per-mission cap	50 kg per configuration	S-URC27-RULES §2.a.iii
Fielded-parts total	70 kg, so 20 kg of headroom over a 50 kg rover	S-URC27-RULES §2.a.iii
Second configuration	Must itself be 50 kg or less	S-URC27-RULES §2.a.iii
Setup	15 minutes including weigh-in and all communications	S-URC27-RULES §1.a.ii
Autonomy to Equipment	As short as 10 minutes, including a possible re-weigh	S-URC27-QA EQ4
C2	Six students; 24 registered, two advisors, 14 people a day	S-URC27-RULES S-URC-LOGISTICS §3.b.ii
Runner rule	A member who leaves C2 to intervene cannot return to operate that mission	S-URC27-RULES §3.e.i
The failure the rule creates	USU 2025: 46.9 kg science configuration, 51.8 kg arm configuration	S-USU-2025-POSTER

WVU's three-year migration is the evidence. A front-half chassis swap took second place in 2024 with 391.80 points, payload inserts through the chassis took second again in 2025 with 396.17, and a single integrated chassis with a two-bolt end-effector change took fourth in 2026 with 376.70, while the architecture simplified each year S-WVU-SAR2024 S-WVU-SAR2025 S-WVU-SAR2026 S-URC24-SCORES S-URC25-SCORES S-URC26-SCORES.

The arithmetic of a swap inside ten minutes includes the swap, a possible re-weigh of unmeasured duration, a re-point of the antennas, and two operators unavailable for the next mission if they leave C2 under the runner rule. With fourteen people a day and six seats, there are about two spare bodies. A whole-chassis swap is a schedule with no failure budget.

Recommended starting point

One integrated rover plus a two-bolt, one-connector end-effector change. Use the 70 kg allowance as a reserve rather than an architecture: keep a 5 kg science-module option inside the 20 kg headroom, but do not design a mission day around changing it. The SAR splits 40 report, 48 video and 12 judges S-URC26-SAR, and 48 video points need an integrated rover by February.

Flips when the single configuration cannot close at 48.5 kg or less by the PDR after the T-01 cuts, which points to a two-bolt science-module change of 5 kg or less rehearsed to eight minutes with two named people including a re-weigh; or when the Equipment cache question, CF-11, resolves to a different container, which is a change of end-effector rather than of module. **Settled by:** ten rehearsals with the real crew and rover against a ten-minute clock, weighing the heaviest configuration at every one.

Science payload against arm payload

T-09 separates two costs that look like one. Science mass is cheap: three kilograms costs about **0.5 N·m per wheel** and **0.04 kg of battery**. Science **integration** is not cheap, because the same decision forces a reaction-torque path, a sealed heated bay and a **10.5-minute block of the 25-minute clock**.

LINE	VALUE	EFFECT ON THE ROVER
Science bay with body-mounted drill	3.5–4.5 kg	0.5 N·m per wheel at 30°, 0.04 kg of pack, no cliff
Cache: 50 mL tube and holster	0.15–1.0 kg	Handled by the arm; 5 kg or less by rule
Arm payload requirement	5 kg, 7.5 cm grasp, 40 cm cube	Sets shoulder torque at practical reach
Arm design torques	Wrist 11 N·m, elbow 53 N·m, shoulder 176 N·m peak	A wrist drill loads the weakest joint
Drill reaction torque	Not published	Body mounting routes it into the chassis
Science power	30–100 W drill, 10–20 W heater, 5–10 Wh per mission	The peak coincides with drive and arm
Chemistry storage	38 °C or less; one amplification method needs 65 °C for 30 minutes	Sealing against dust fights venting a heater
Mission clock	25 minutes, of which acquisition and documentation are 10.5	The drill is the critical path, not the chemistry
Second assay	About 3.5 minutes to load and 2 minutes to read	A 30-minute incubation cannot be the second capability
Fluids	12 mL of liquid, 10 mL of waste	Contained, under a no-spill rule

The wrist is the wrong place for the drill, and the numbers say so: the wrist is designed for **11 N·m**, the elbow for 53 and the shoulder for 176. A 30 cm drill on a compliant wrist wanders, and it also spends 1.5–2.5 kg at the worst lever, cutting usable payload to about 2.5–3.5 kg. Drilling costs about **8 Wh** for five minutes at 100 W, six per cent of a 125 Wh mission, but it lands on the same bus as a 120 A drive peak and a 1.1 kW arm peak, so the science rail is fused and isolated. WVU solved the same problem with a dedicated drill on a linear actuator `S-WVU-SAR2025`.

> Recommended starting point
>
> A **body-mounted auger or drill on a 200 mm linear stage** with a graduated clear liner and a scoop fallback; **3.5–4.5 kg** of science in a sealed bay with desiccant and a heater isolated from the electronics; a **1 kg cache**; an **8 kg arm class with 5 kg of payload**; and **two assays on one detector** with the pair chosen by December, so the assay appears in the SAR video by 2027-03-05 `S-URC27-RULES` §1.b.i.

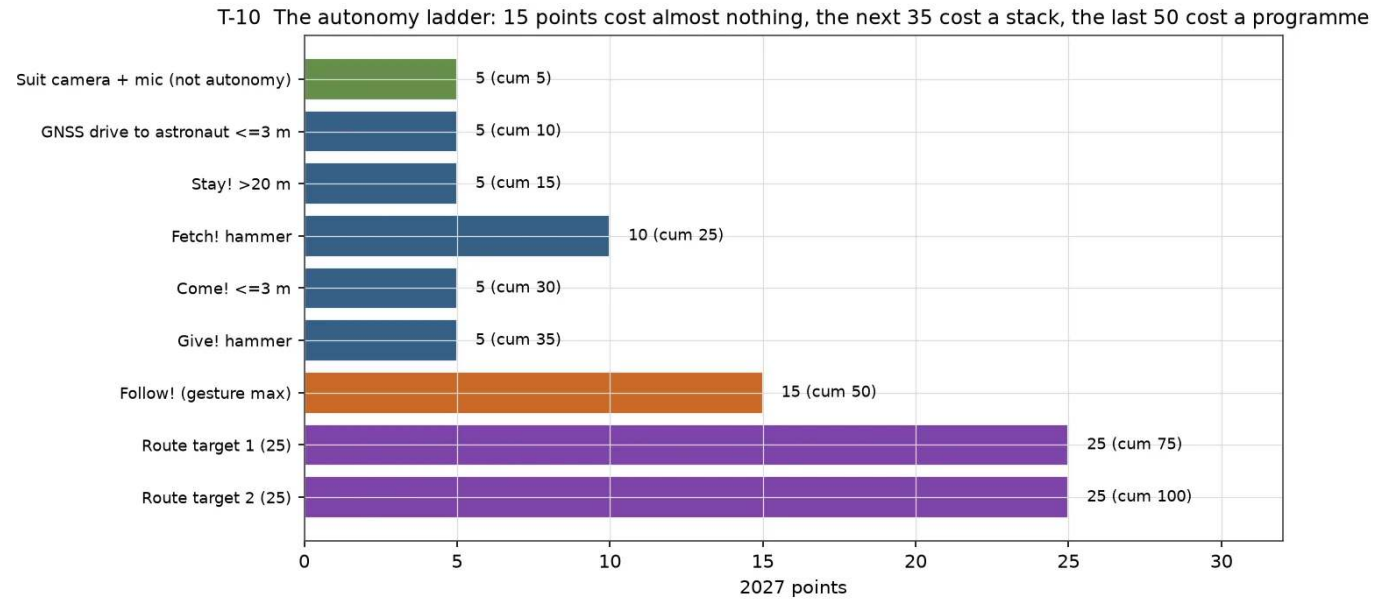
Flips when two assays need more than 5 kg or 5 L, which points to a shared detector or an optical second capability with a written answer from URC; when hard or rocky soil defeats the auger, which adds a percussive drill of about 1 kg and reopens `T-01`; or when a wrist drill is unavoidable for reach, which steps the elbow to 53 N·m and moves to the 14 kg arm class, reopening `T-01` and `T-03` together. **Settled by:** the cam-lock, latch and drawer mocks, a 10 cm soil column, and a 25-minute dry run that measures reaction, depth provenance and the two-assay clock.

Teleoperation against autonomy investment

`T-10` is the highest-variance decision in the season. Autonomy is the mission with the widest spread in the record, at a **mean of 26.7 and a standard deviation of 27.8**, with **41 % of team-years scoring zero** across 162 finals appearances from 2022 to 2026. It is also the only 2027 mission whose points are fully published, which makes its value calculable rather than assumed.

ELEMENT	POINTS	CUMULATIVE	WHAT IT NEEDS
Suit camera and microphone, not an autonomy element	5	5	A coverall, a 5.8 GHz pod and a C2 display
GNSS drive to the astronaut and stop within 3 m	5	10	RTK and a waypoint controller
Hold while the astronaut moves more than 20 m	5	15	A latching hold and a disabled planner
Fetch the hammer	10	25	Person and hammer detection plus an autonomous grasp
Come within 3 m	5	30	Detection at 20 m or more, or a suit beacon
Give the hammer	5	35	An autonomous place or drop
Follow, gesture at maximum	15	50	Person tracking on natural terrain with a 3 m stop
Route target 1	25	75	Hill route, 1 m stop and a visible state light
Route target 2	25	100	No link, an onboard decision and a visible state light

The effort curve is not linear. The first **15 points** need RTK, a state machine and a light, which is hardware the rover already carries. The next **35** need person detection, an autonomous grasp and a safety case with a person inside the motion envelope. The last **50** need a planning stack that works with **no link behind a hill** and stops within a metre, and 41 % of the field has scored zero on that family. The 2026 ceiling of **86** is a property of the rubric, not a capability limit.



Every autonomy scoring element in the order a first build can earn it, with the cumulative total on the right. The first fifteen points are nearly free and the last fifty are a programme, which is what the tiering in this trade follows.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *Full resolution*

WVU's investment is the counter-example to the software-only view. It logs a **minimum of 100 practice hours per operator**, and its systems-engineering presentation reports that 50 hours or more of testing found many mechanical, electrical and software faults `S-WVU-SAR2026` `S-WVU-INCOSE2026`. The hardware is unremarkable: a 2.4 GHz primary link, a lower-band recovery link and GNSS with LiDAR navigation. The hours are the investment.

Recommended starting point

Secure the 15-point floor in the first build, build toward the 35-point astronaut tier by the SAR, and attempt the 50 route points only if RTK with a terrain model completes three of five full runs by January. The rules guarantee the geometry of the astronaut tasks: a published stand-off, an exact target coordinate and a required hold. Do not buy a sixth camera or a second computer before the route gate is cleared.

Flips when fewer than two dedicated software people are available by December, or no RTK fix is available by February, in which case stop at 15–20 points and put the hours into Delivery and Equipment Servicing; or when route autonomy runs three of five by the SAR, which is the largest single block in the competition. **Settled by:** five full mock runs on hilly terrain with the real link, scored by someone who did not write the code. One run is a demonstration; five are a capability.

Camera count against link budget

T-11 is a communications decision disguised as a perception decision, and it converts into points through the speed at which an operator can drive safely. The arithmetic is unforgiving. Six cameras at 720p30 and 3 Mbit/s each is **18 Mbit/s**; the 900 MHz link carries **200 kbps** at 1 km, which is **1.1 %** of that demand. The 2.4 GHz radio carries **6 Mbit/s at 2.62 km** or 54 Mbit/s at 0.42 km, and 5.8 GHz with the radio at the mast carries **6 Mbit/s at 21.4 km** or its top modulation at 1.90 km. Dropping the top modulation to 6 Mbit/s buys about **nine times the range**. Latency of 100–300 ms means a rover at 1.0 m/s travels **0.25 m** between a command and the image of it.

OPTION	STREAMS	DEMAND	WHERE IT SURVIVES
Drive and wrist only	2	4–6 Mbit/s	A 900 MHz-only day
Drive, wrist, science still and rear	4	8–12 Mbit/s, adaptive	All four missions on 2.4 GHz close in and 5.8 GHz at range
Six streams	6	18 Mbit/s	Close in only; operator overload becomes the limit
One stream plus stored stills	1	2–3 Mbit/s	The worst link, with no live arm view

Recommended starting point

Four cameras at 720p30 (drive, wrist, science still and rear) at **2–4 Mbit/s adaptive**, plus a high-resolution still path. Freeze the count against the link rather than against the sensor catalogue, and keep the **900 MHz sub-band plan** for every link, including drone and astronaut, in the same interface document, which is due **2027-04-30** S-URC27-RULES §2.c.v.

Flips when a 5.8 GHz test with the radio at the mast shows 20 Mbit/s or more at 1 km, which makes six streams feasible subject to operator load; when the day is 900 MHz only, which leaves one or two streams plus stored stills; or when autonomy earns the route points, which reduces the number of streams an operator needs. **Settled by:** a measured bitrate per camera at the intended resolution, and one full route driven on the final antennas.

E-stop scope against communications and logging

T-12 is the trade with a pass-or-fail rule behind it. One button must cease all power draw from **any and all batteries**, including compute, communications and payload S-URC27-RULES §2.a.v. The radio dies with everything else, so the C2 station sees nothing and the field judge reads the rover's state from its own light.

OPTION	DESCRIPTION	COST	VERDICT
Single normally-open contactor loop	One button opens a loop that drops the main contactor; every rail is downstream	$$	**Recommended**; one action kills everything and can be tested with a current probe
Distributed contactors	One per battery or rail, all driven by the same button	$$$	Needed only when two packs run at different voltages
Software stop	Firmware stops the motors and keeps the computer alive	$	**Non-compliant** with §2.a.v; no partial credit

Five design rules follow. One precharge and contactor loop, with the state-light driver on a branch inside the loop but independent of the computer. No software in the kill path. Mode and task transitions written to non-volatile storage immediately, because the E-stop kills the logger. A passive mechanical state flag so the judge can read the state after the link is gone. And a bench test with a current probe: press the button and verify zero amps at the pack, including the computer and the radio.

> Recommended starting point
>
> **A single normally-open loop**, with the state light on its own branch inside the loop, a precharge, a fuse coordinated below the contactor rating, and a documented kill-chain test in the SAR video.

Flips when two packs at different voltages appear, which needs distributed contactors on one loop; or when a judge explicitly requires post-kill telemetry, for which no ruling was found, and only then a non-battery last-gasp indicator. **Settled by:** 50 presses under load with a voltmeter on every rail, photographed for the review package.

The decision chain

The trades are not independent. An edge in the chain below means the downstream number cannot be frozen until the upstream one is.

DECISION	TRADE	FREEZES AT	WAITS ON
D01 Mission targets, including autonomy tier and drone	T-07, T-10	2026-11-15	nothing
D02 Configuration strategy, single or modular	T-08	2026-11-15	nothing
D03 Mass allocation at 40.0 kg plus 15 %	T-01	PDR	D01, D02
D04 Wheel diameter and wheelbase at 250 mm and 0.90 m	T-03	PDR	D03
D05 Actuator class, 9 N·m or the next step	T-03	PDR	D04
D06 Bus voltage, 24 V or 48 V	T-04, T-12	2026-11-30	D05
D07 Arm reach and stow, 6R at 1.05 m with a body drill	T-02	PDR	D03, D04, D09
D08 Arm precision strategy, wrist camera and visual servo	T-02	PDR	D07
D09 Science assay pair	T-09	December 2026	D03
D10 Battery capacity and parallel count at 8S4P	T-04	CDR	D05, D06
D11 Compute class and thermal treatment	T-05	PDR	D01, D10
D12 E-stop topology, single loop	T-12	2026-11-30	D06, D10
D13 Camera count and frequency bands	T-11	2027-04-30	D01, D11
D14 Mast architecture, two masts	T-06	PDR	D11, D13
D15 Interface document version 1 and procurement	all twelve	2026-11-30, orders 2026-12-15	D07, D08, D09, D11, D12, D14
D16 Build, integrate and test, power-on	all twelve	2027-01-31	D15

The critical path is **D01, D03, D04, D05, D06, D10, D11, D13, D14, D15, D16.** Nothing on it can be compressed by working harder, only started earlier. D04 has no upstream dependency other than the allocation and is irreversible after tooling, so freezing it early is worth more than optimising it. D09 is the only December decision off the PDR path, and it feeds the SAR video, which makes it both the longest-lead science item and the most likely to slip. D06 and D12 must be frozen before any harness exists. There is deliberately no edge from the drone to the mass allocation, because the drone sits outside the 50 kg and the 1.2 m cube; it feeds the 70 kg question carried inside D01.

The dates that bind these trades

DATE	EVENT	DECISIONS FROZEN
2026-10-30	Registration closes S-URC27-RULES §3.a.i	Team scope and mission targets, D01
2026-11-15	Internal	Autonomy tier and drone, D01; single or modular, D02
2026-11-30	Internal	Interface document version 1, D15: bus voltage, connectors, mounts; E-stop topology, D12
2026-12-02	**PDR** S-URC27-RULES §3.a.ii	Mass allocation, D03; wheel diameter, D04; arm concept, D07; configuration, D02
December 2026	Internal	Science assay pair, D09; the SAR video must show it S-URC27-RULES §1.b.i
2026-12-15	Internal	Long-lead orders released: motors, gearboxes, boards, composite stock
2027-01-31	Internal	Power-on of the integrated rover; compute and thermal, D11
2027-02-14	Internal	Full system integrated; camera count, D13, proven
2027-02-01 to 02-20	Internal	SAR filming, current architecture only
2027-03-05	**SAR** S-URC27-RULES §3.a.iii	Frozen mass, stow, E-stop evidence, assay footage
2027-04-10	Internal	Software feature freeze
2027-04-24	Internal	Mechanical configuration freeze
2027-04-30	**Communications documentation** S-URC27-RULES §2.c.ii	Bands, antennas, drone and astronaut links, licence basis
2027-05-21	**Science Plan** S-URC27-RULES §1.b.i	Assay pair, sites, interpretation
2027-05-26	Expense Report S-URC27-RULES §3.c.ii	Work-package expenses
2027-06-02 to 06-05	**Finals**	Repairs only

What a first build does first

- **Take the irreversible decisions first:** mass allocation, wheel diameter and bus voltage. Each is cheap to decide now and expensive to change later, and everything downstream becomes a re-test rather than a redesign.
- **Design to 46 kg with a 4 kg reserve, not to 50 kg.** The field is pinned at 49–50 kg because the cap decides there, not because 50 kg is optimal. A 33 kg rover won in 2023 S-WVU-SAR2023 .
- **Buy typing precision with a camera, not with joints.** The shoulder lever defeats a reducer budget; at a 0.35 m standoff, 5.2 mrad becomes 1.8 mm.
- **Choose the wheel from the box, the wheelbase from the terrain, and the motor from the torque**, in that order.
- **Size the pack against the 60-minute Delivery and the back-to-back autonomy and servicing case**, and check the 156 A peak before choosing the BMS. Both answers land at four parallel strings.
- **Treat thermal design as a mass and volume item.** A white shell, a shade and a filtered fan, with the battery in its own compartment; the computer class is the small term.
- **Use two masts, and remember that 3 m is the base-station rule.** A 3 m rover mast fails the cube before it helps the link.
- **Do not fly a drone in the first build**, and use the 70 kg allowance as a reserve rather than an architecture.

- **Measure the six numbers nobody publishes:** mass, watt-hours per mission, weigh-in duration, setup time, teleop ground speed, and arm backlash under load at 38 °C. Each converts an estimate into a measured value.

What would change a recommendation

GAP	WHY IT MATTERS HERE	WHAT CLOSES IT
Whether a drone counts against the 70 kg total	Changes the headroom in T-08 and the accounting in T-07	A written question to the URC Director before the PDR
Whether the Equipment Servicing cache is the team's own Astrobiology cache	Changes the mass, thermal and end-effector answer in T-09	A written question before December
Weigh-in duration	Decides whether a module change fits the 10-minute window	Ten weigh-ins timed at a mock event
Teleop ground speed on real terrain	The 1 km and 30-minute arithmetic in T-10 and T-11	A field log at an analogue site
Arm backlash under load at 38 °C	The datasheet figure is nominal, and T-02's 1.8 mm depends on it	An output-encoder lost-motion test
Drill reaction torque in MDRS-type soil	Decides wrist or body mounting, which is T-09	A load cell on the auger in three soils
Effective rolling resistance on the site terrain	The dominant assumption in T-03 and T-04	A tow test at 40 kg on sand and gravel
The 2027 Autonomy rubric and any ceiling	The investment target in T-10	Re-read the rules and the Q&A monthly from December
Whether the 3 m antenna limit applies to a separate GNSS base	Affects T-06 and the survey-in plan behind CF-13	A written question, asked with the December batch
Drone masses for WVU, Monash and Missouri S&T	The 5 kg rule is a weighing rule and no team publishes a number	Ask, and weigh one at the 2027 finals

Further reading

- Mass and power derives the 249 W average, the 3.7 kW peak and the pack sizing that T-01 and T-04 use.
- Interfaces and couplings works the 66 interfaces, of which 37 are high criticality.
- Mobility and arm and end-effectors carry the full comparison tables behind T-02 and T-03.
- The 13-week plan and the first 90 days put these freeze dates on a calendar.
- 110 engineering problems turns each recommendation into a row with an experiment and a success criterion.

THE BUILD PLAN · CHAPTER 7

From analysis to build

The earlier parts of this book end in three lists: the problems a 2027 rover creates, the ways a season can fail, and the questions nobody on a first build can answer yet. This part turns those lists into a plan. It carries **110 engineering problems** in five dependency waves, **70 risks** with a stated likelihood and impact, and **77 open questions** grouped by who can answer them. The plan runs from 2026-10-26 to the finals on 2027-06-02, and both of its hard gates fall in the first half of that window.

110

ENGINEERING PROBLEMS

Each with an experiment and a written success criterion

2026-2027

253

DEPENDENCY EDGES

Between the problems, which is what sets the order

2026-2027

70

RISKS

Six categories, scored by likelihood and impact

2026-2027

77

OPEN QUESTIONS

16 to URC, 22 measurable, 27 need hardware, 12 team decisions

2026-2027

Key takeaways

- **The PDR is due 2026-12-02, five weeks into the season** S-URC27-RULES §3.a.ii, and the **SAR is due 2027-03-05, ten days before the integration and test wave opens** S-URC27-RULES §3.a.iii. Both are gates that close a wave, not reviews that happen at the end.
- **The first four things to do are measurements.** The numbers that size the rover are not published, so the season is won by replacing them with measured values before the design is committed.
- **The register is ordered by dependency, not by subsystem.** Thirteen rows have five or more dependents; they decide when everything else becomes possible.
- **Nine risks carry the argument**, and six of the nine are about process rather than hardware: the review, the build target, integration time, depth of the team, scope and mass tracking.

Three registers and one calendar

PART OF THE PLAN	ROWS	THE QUESTION IT ANSWERS	WHERE IT IS SET OUT
Engineering problem register	110	What has to be learned, decided or tested, in what order, and what closes each item	Five waves, 253 dependency edges, 155 traced requirements
Risk register	70	What can go wrong, how likely it is on the evidence, and what it costs if it happens	Six categories, a likelihood and impact band per row
Open questions	77	What is unknown, who can answer it, and what the team assumes until they do	Four groups: URC, measurement, hardware test, team decision
The first 90 days	13 weeks	Which work unblocks the most other work, and in what order	The register's dependencies, read forwards
The 13-week plan	13 weeks	Which decision is taken in which week, and who owns the row	A dated calendar with four gates

The registers are not a schedule and they are not a wish list. They make no component recommendation of their own and they carry no cost ranking: priority is set by rule consequence and dependency, never by price. The design recommendations live in the subsystem and trade pages.

The five waves

WAVE	WEEKS	CALENDAR	WHAT IT IS FOR	GATE
W0 Measure	0–4	2026-10-26 to 11-22	Bench, rig and field measurements that replace the rulebook's missing numbers; a judge-like task mock	none
W1 Architecture	4–12	2026-11-23 to 2027-01-17	The irreversible or expensive-to-reverse choices, with a freeze date named for each	**PDR 2026-12-02**, week 5
W2 Detailed design	12–20	2027-01-18 to 03-14	Interface documents, sizing analyses, procedures, and the February evidence runs the SAR needs	**SAR 2027-03-05**, week 19
W3 Integration and test	20–27	2027-03-15 to 05-02	The rover assembled and the four missions run against a clock	test evidence for the finals
W4 Review and rehearsal	27–31	2027-05-03 to 06-01	Review responses, frozen procedures, rehearsal, and the final documents	**Science Plan 2027-05-21**, **finals 2027-06-02 to 06-05**

The communications documentation is due **2027-04-30**, in the last week of W3 rather than in the review wave, which is a detail that changes where the work sits.

The two gates inside the first half

The PDR falls **five weeks** into the season, inside the architecture wave, at a point where most first builds have not finished measuring. The SAR falls **ten days before** the integration and test wave opens, so the rover has to be filmed at partial fidelity while it is still being assembled. The consequence is structural: the February evidence runs are a scheduled part of the detailed-design wave, not an afterthought, and a team that discovers the ordering in March has already lost the submission.

The same two dates explain the shape of the risk register. The highest-ranked risk is missing the SAR cut, where two of three registrants never reach a score sheet, and the second is building to June instead of building to March.

Registration closes **2 days** `2026-10-30`

PDR submission **35 days** `2026-12-02`

How to read a register row

Every problem row is read left to right, and the structure is the same for all 110 rows.

FIELD	WHAT IT GIVES	WHY IT EXISTS
ID	A stable identifier for the life of the project	The risk register, the reviews and the team tracker all reference it
Problem	One testable sentence with no solution inside it	A problem that contains its answer is a decision already taken
First experiment	The smallest thing that produces the missing information, sized in days	If it needs the whole rover, it belongs in a later wave
Success criterion	A number or a binary, written before the experiment runs	Rule numbers are quoted; a design target is labelled as derived
Unlocks	How many later problems depend on this one	The working measure of downstream design unlocked

Priority is deliberately coarse: **29 Critical**, **72 High**, **8 Medium** and **1 Low**. The fine ordering is the unlock count, and the two are read together. A High row with eight dependents is better value than a Critical row with none.

The questions, in four groups

GROUP	COUNT	WHO ANSWERS	HOW LONG IT TAKES
Questions to URC by email	16	The URC Director, with answers mirrored onto the public Q&A page	Weeks, and the page has changed twice in a month after publication
Measurements the team can take	22	The team, on a bench or with a load cell	Days
Tests that need hardware or a field day	27	The team, with a prototype or an analogue site	Weeks, and they need a date
Decisions the team takes	12	The team, in writing and dated	One meeting each, and the risk is not deciding

Each question carries a **meanwhile** assumption, so an unanswered question never stops work. The sixteen judge questions go out first, in one numbered email in November 2026, because they take longest to come back.

The calendar in three views

VIEW	SPAN	WHAT IT IS FOR
The 90 days	2026-10-26 to the end of January 2027	The measurement wave, the PDR, and most of the architecture decisions
The 13 weeks	2026-10-26 to 2027-01-23	The same window as a dated calendar with one decision per week and four gates
The season	31 weeks to 2027-06-05	Five waves, four review submissions and the test window after the SAR

After the SAR there are **89 days** to the finals, and they are a test window rather than a build window `S-URC27-RULES §3.a.iv`. That number is the reason the gate dates matter more than any individual subsystem deadline.

Where to start

- **A new team, first week:** read the problem register for the five waves, then the first 90 days for the order of work. The first four rows to start are a mass frame, a task mock, a link site booking and a kill-chain breadboard.
- **A team that has a design:** read the risk register first, and check the nine high-likelihood, high-impact rows against the current plan.
- **A team in review season:** read the open questions, send the sixteen judge questions, and write the meanwhile assumptions onto the affected rows.
- **Following the argument in order:** coupling and the trade studies supplies the recommendations that each register row exists to test.

What the plan assumes

The plan holds while five things are true: the 2027 rules stay as published, with the 50 kg and 70 kg limits and the 1.2 m cube; the team can build and measure on a bench and at two analogue sites before June; at least six people can be trained to hold the C2 seats; long-lead parts are ordered against numbers frozen in weeks 6 to 9; and the questions are answered or their fallbacks accepted by their dates.

If one of the first four fails, the dates move and the order does not. A register that is read once is a document; one that is reviewed weekly is a management system, and the problem register describes the 45-minute weekly review that keeps it current.

Further reading

- 110 engineering problems sets out every row with its experiment and criterion.
- 70 risks scores each row and names the change that removes it.
- 77 open questions lists the questions, the decision each blocks and the meanwhile assumption.
- The first 90 days reads the register's dependencies forwards.
- The 13-week plan turns the same window into a dated calendar with four gates.
- Coupling, conflicts and twelve trade studies supplies the decisions the register tests.

THE BUILD PLAN · CHAPTER 7.1

110 engineering problems

This is the season's engineering register: 110 numbered problems, each with the requirement that creates it, the experiment that produces the missing number, and the criterion that closes it. Every problem traces to at least one 2027 requirement, and all 155 requirements are covered by at least one row. The rows are ordered by dependency wave rather than by subsystem, because the order in which the work becomes possible is not the order in which subsystems are usually written down.

110

NUMBERED PROBLEMS

In five dependency waves, each with a success criterion

2026-2027

253

DEPENDENCY EDGES

Between the rows, 13 roots and 21 leaves

2026-2027

155

REQUIREMENTS TRACED

Every 2027 requirement carried by at least one row

2027

29

CRITICAL ROWS

With 72 High, 8 Medium and 1 Low

2026-2027

Key takeaways

- **The first four things to do are measurements, and none of them resembles a design-review deliverable.** The numbers that size the rover are unpublished, and the season is won by replacing them before the design is committed.
- **Thirteen rows have five or more dependents.** They are the spine of the season, and only they should be scheduled without an argument.
- **The arm decision silently consumes the drivetrain margin.** An arm in the 14 kg class lands the rover 0.3 kg below the mass at which the chosen actuator class stops being adequate, and the next bracket costs 2.7 kg for six larger actuators.
- **The SAR forces a February evidence programme.** The submission is due ten days before the test wave opens, so four missions have to be filmed at partial fidelity while the rover is still being integrated.

What a row contains

FIELD	WHAT IT GIVES	WHY IT EXISTS
ID	A stable identifier for the life of the project	The risk register, the reviews and the team tracker all reference it
Problem	One testable sentence with no solution inside it	A problem that contains its answer is a decision already taken
First experiment	The smallest thing that produces the missing information, sized in days	If it needs the whole rover, it belongs in a later wave
Success criterion	A number or a binary, written before the experiment runs	Rule numbers are quoted; a design target is labelled as derived
Unlocks	How many later problems depend on this one	The working measure of downstream design unlocked
Traces to	The requirement and the rule or ruling behind it	Every problem exists because the 2027 rules require something

Priority is deliberately coarse. **29 rows are Critical**: open at their gate, with the season at risk through a pass-or-fail rule, a milestone, or a measurement that three waves depend on. **72 are High**: they block three or more later problems or a whole mission capability. **8 are Medium** and **1 is Low**, recoverable later without a rule breach.

The fine ordering is the unlock count, and the two are meant to be read together. A High row with eight dependents is better value than a Critical row with none, which is the case for a May submission deadline.

The five waves

WAVE	PROBLEMS	WEEKS	CALENDAR	WHAT IT IS FOR	GATE
W0 Measure	23	0–4	2026-10-26 to 11-22	Bench, rig and field measurements that replace the rulebook's missing numbers; the judge-like task mock	none
W1 Architecture	26	4–12	2026-11-23 to 2027-01-17	The irreversible or expensive-to-reverse choices, each with a freeze date	**PDR 2026-12-02**, week 5
W2 Detailed design	28	12–20	2027-01-18 to 03-14	Interface documents, sizing analyses, procedures and the February evidence runs	**SAR 2027-03-05**, week 19
W3 Integration and test	20	20–27	2027-03-15 to 05-02	The rover assembled and the four missions run against a clock	test evidence for the finals
W4 Review and rehearsal	13	27–31	2027-05-03 to 06-01	Review responses, frozen procedures, rehearsal and the final documents	**Science Plan 2027-05-21**, **finals 2027-06-02 to 06-05**

The five waves sum to 110. The PDR is due **2026-12-02** S-URC27-RULES §3.a.ii, the SAR **2027-03-05** S-URC27-RULES §3.a.iii, and the finals run **2027-06-02 to 06-05** S-URC27-RULES §3.a.iv. The communications documentation is due 2027-04-30, in the last week of W3, so its substance has to be finished before that week opens.

All 110 problems

The table is searchable and sortable, and it filters by wave and by priority. Dependency counts are in the Unlocks column; a row with a high count is a row that frees design work elsewhere.

ID	WAVE	PROBLEM	PRIORITY	TRACES TO	DEPENDS ON	SUCCESS CRITERION
W0-01	0	Establish the mass budget frame, the weigh-in procedure and the 1.2 m cube gauge before any hardware is committed	Critical	REQ-SYS-001, -002, -003, -005, -006, -007; §2.a.ii–iii; §1.a.iv	Not applicable	Allocation with growth current mock fits the g
W0-02	0	Build a judge-like Equipment Servicing mock, drawer, latch, valve, buttons/switches/knobs, two tubular pin-tumbler locks, 1.5-inch cam-lock hose port, keyboard and display mounts, because no task geometry is published	Critical	REQ-ARM-003, -004, -005, -006, -007, -008, -009, -010, -011; §1.d.ii; EQ1–EQ3	Not applicable	Every one of the nine E within the rulebook's d second person; every
W0-03	0	Measure the forces, travels and torques of every Equipment sub-task, the numbers the rulebook does not give	Critical	REQ-ARM-002, -003, -005, -006, -007, -008, -009, -010, -020; EQ1, EQ2	W0-02	Curve per task, repeat factor recorded; every
W0-04	0	Measure the typing target: key actuation force and travel, key pitch, and whether a compliant tip can press one key without touching its neighbours, from the wrist-camera standoff the arm will actually use	High	REQ-ARM-011, -019; CMP-003; SFT-014, -015; EQ3	W0-02	≥95 % single-key actu measured force and tra standoff
W0-05	0	Measure whether the specified rock-pick hammer can be grasped and lifted repeatably from the ground	High	REQ-ARM-013, -014; SFT-006; CMP-005; AUQ8	W0-02	≥45/50 grasps; zero dr a recovery procedure f
W0-06	0	Measure the cache interface: handle grasp, insertion clearance, lid closure and lock turn, and hand-removal time	High	REQ-ARM-004, -017; SCI-010, -011; §1.b.vii; §1.d.ii	W0-02	Both variants: mass ga handle grasp repeatabl
W0-07	0	Measure the force and time to acquire ≥5 g from below 10 cm in the soil classes the rules describe, including "very hard and difficult to excavate"	Critical	REQ-SCI-004, -005, -017; ARM-015; AQ6, AQ10	W0-02	≥5 g below 10 cm in ≤4 and torque recorded; a
W0-08	0	Establish whether the candidate bioassays work on real soil extract at all, recovery, interference, blanks, and how long they take	Critical	REQ-SCI-001, -002, -003; AQ9	W0-07	ATP recovery ≥50 % in 120 %; blank and steril temperature
W0-09	0	Measure the wheel–soil traction and rolling resistance that the torque and slope design will rest on	High	REQ-MOB-001, -008; §1.c.ii; AQ6	Not applicable	μ and Crr per soil with a by ≥20 % traction diffe
W0-10	0	Establish what obstacle the rover must climb, in the absence of a published number, and what wheel diameter buys it	High	REQ-MOB-007; §1.c.ii; DQ1 (incl. vertical faces), AUQ17 (inclinations not provided)	Not applicable	Climb success ≥8/10 a documented maximum requirements as a deri
W0-11	0	Build the link budget model and run the first real link test at ≥1 km with terrain in the way, because the 2027 course guarantees non-line-of-sight	Critical	REQ-COM-001, -004, -006, -008; §1.c.ii, §3.d.ii; RQ13	Not applicable	Measured link at 1 km a 900 MHz path that ca prediction vs measure
W0-12	0	Quantify the 900 MHz sub-band compliance problem before a radio is chosen	High	REQ-COM-010, -011; §2.c.v	W0-11	All three sub-bands de compliant; a written pr
W0-13	0	Establish the achievable RTK accuracy and heading with the team's own base station, including survey-in time, because 1 m navigation and 3 m astronaut stops are scored	High	REQ-CMP-001, -002; §1.e.v, §1.e.xvii; RQ11	Not applicable	Horizontal repeatability survey-in; a document

W0-14	0	Establish the real detection range for every tag and sign the 2027 rules use, in full sun, with candidate cameras	High	REQ-CMP-003, -004, -005, -006, -009, -010; §1.d.ii, §1.e.xii	Not applicable	Decode rate ≥95 % at t measured max range p practice task
W0-15	0	Establish the internal temperature a sealed electronics box reaches in the sun at MRS-like conditions, and whether white paint, a shade and a filtered fan are enough	High	REQ-PWR-005, -009; MOB-011; §3.d.iii	Not applicable	Internal air ≤50 °C at 38 ΔT per treatment; fan f
W0-16	0	Establish whether candidate sealing (connectors, glands, lens windows, fan filter) survives dust and light rain without losing function	Medium	REQ-PWR-005, -009; MOB-011; §3.d.iii	Not applicable	No functional loss afte interval derived; a writt
W0-17	0	Prove a single-action E-stop removes all power from every battery, including compute, comms and payload, and choose a topology that cannot weld	Critical	REQ-SAF-001, -002, -003; PWR-004; §2.a.v	Not applicable	0 V at every rail in ≤100 LED array remains pov
W0-18	0	Establish the pack's real sag and the BMS trip margins, using pulses that represent a hill climb plus a stalled arm	High	REQ-PWR-002, -007, -008; §1.c.i	Not applicable	No trip and no converte curve; BMS threshold s
W0-19	0	Time the C2 station: 15-minute setup including weigh-in and measurement, 10-minute teardown, with the real kit	High	REQ-SYS-005, -008, -009, -010; COM-009; OPS-012	W0-01	Worst-of-five setup ≤1 teardown ≤8 min; weig fastest run
W0-20	0	Select and secure analogue test sites that can stand in for MDRS, because testing at MDRS is prohibited	High	REQ-OPS-007; §3.d.iv	Not applicable	Two sites with measur shadow; a test-day ch
W0-21	0	Capture the transport load case: vibration and shock the rover will actually see on the road to Hanksville	Medium	REQ-SYS-001, -003; environment §3.3; judges' "half hour of continuous shaking" advice	Not applicable	A load case (peak g, do retention test are desig
W0-22	0	Establish the Astrobiology 25-minute sequence on a bench, including which assay starts first and whether two results fit	High	REQ-SCI-006, -014; §1.b.ii; AQ10	W0-07, W0-08	Both results logged by card with the slowest a
W0-23	0	Define the test-record schema and the evidence plan before the first measurement, so every Wave-0 result is usable in the PDR and SAR	High	REQ-OPS-001, -002; §3.a.ii–iii	W0-01	Every Wave-0 measur SAR writer can lift; zer
W1-01	1	Fix the rover's mass class, including the arm's mass class and the drivetrain actuator class it silently buys	Critical	REQ-SYS-002, -003, -006; §2.a.iii; CF-01	W0-01	A chosen arm mass cl is ≤50 kg and whose w rating at the steepest r
W1-02	1	Choose the mobility architecture: 4 or 6 wheels, rocker-bogie or wheel-legs, skid or active steering	High	REQ-MOB-001, -007, -008; §1.c.ii; DQ1	W0-01, W0-09, W0-10	A chosen architecture slope, a documented c that fits the cube gaug
W1-03	1	Fix wheel diameter, wheelbase and track, the number that re-opens seven to nine interfaces if it changes later	Critical	REQ-SYS-001; MOB-001, -007; CF-03; IF-004	W0-01, W0-09, W0-10, W1-01, W1-02	A diameter frozen in C/ static tip angle ≥55°, a
W1-04	1	Choose the bus voltage and the drivetrain actuator class (integrated 48 V modules vs motor+gearbox on 24 V)	High	REQ-PWR-001, -007; IF-011/IF-012; §2.a.i	W1-05	A voltage/actuator pair margin, keeps harness documented E-stop ra
W1-05	1	Size the drivetrain: gear ratio, torque margin at the measured slope, and the continuous thermal duty	High	REQ-MOB-001, -007, -009; §1.c.ii; §1.c.iii	W0-09, W0-10, W1-02, W1-03	≥20 % torque margin a a 20-minute climb at 3(

W1-06	1	Choose the stow and deploy concept: geometry-first, or folds; one mast or two; whether a diagonal placement is claimed	Critical	REQ-SYS-001, -007; ARM-018; IF-003/IF-005; trap 2 and trap 3	W0-01, W1-03	Fits the 1.2 m gauge in with the fold count stat
W1-07	1	Choose the arm architecture: degrees of freedom, serial vs prismatic-X, mount location and reach class	High	REQ-ARM-001, -002, -018; §1.d.i; CF-02	W0-03, W0-06, W1-01	Every 2027 Equipment stowable in the measu
W1-08	1	Choose the joint/backlash class and the gravity-holding strategy	High	REQ-ARM-011, -019, -020; trap 5	W0-04, W1-07	A joint class whose su vision-servo strategy t or counterbalance
W1-09	1	Choose the end-effector architecture and tip set: parallel gripper plus swappable tips, or a tool changer	High	REQ-ARM-002, -003, -007, -008, -009, -010, -020; §1.c.iii, §1.d.ii	W0-02, W0-03, W0-05, W0-06, W1-07	All nine sub-tasks perf force and jaw opening
W1-10	1	Choose the sample-acquisition architecture and the depth-provenance method	High	REQ-SCI-004, -005, -016; ARM-015; AQ6, AQ10; CF-07	W0-07, W1-01	≥5 g below 10 cm in ≤4 accept without a discu refuses
W1-11	1	Choose the science payload architecture: the two life-detection capabilities, the shared detector, the fluidic path and the waste containment	Critical	REQ-SCI-001, -002, -003, -012, -013, -018, -020; §1.b.v, §1.b.ix; AQ7, AQ9; CF-07	W0-08, W0-22	Both biomolecules det sequence; a written AC when inverted
W1-12	1	Choose the power architecture: chemistry, cell configuration, capacity, BMS behaviour and charging	High	REQ-PWR-001, -002, -007, -008; §1.a.ii, §1.c.i; IF-011/IF-012	W0-18, W1-01	Full 60-minute worst-c BMS trip at the measu
W1-13	1	Choose the E-stop and distribution architecture, including the precharge, the independent LED driver and the last-gasp indication	Critical	REQ-SAF-001, -002, -003; PWR-004; §2.a.v; IF-016/IF-017; CF-09	W0-17, W1-12	A single action remove LED state is judge-visi procedure
W1-14	1	Choose the communications architecture: how many planes, which bands, whether a repeater or drone relay is in the reference, and who owns the mast	High	REQ-COM-001, -002, -004, -005, -006, -008, -010, -013; §1.e.xvi, §3.d.ii; CF-04, CF-14; IF-005	W0-11, W0-12, W1-16	A two-plane architectu rehearsed behaviour o including the astronaut
W1-15	1	Choose the compute and camera architecture against the link budget, not before it	High	REQ-CMP-003, -004, -010, -011; SFT-003; AUQ2; CF-05, CF-06; IF-019/IF-020	W0-11, W0-14, W1-14	A camera count and co headroom; the percept
W1-16	1	Choose the mast architecture: height, one mast or two, folding, guying and isolation	High	REQ-COM-008; CMP-009; SYS-001; IF-005/IF-006; CF-04	W0-11, W0-14, W1-03	First mode ≥19 Hz; sta can deploy in ≤60 s; a s
W1-17	1	Decide how much autonomy the team is really building, and in which direction the fallback runs	High	REQ-SFT-001, -002, -003, -004, -005, -006, -007, -008, -009; §1.e; AUQ2, AUQ21	W0-09, W0-10, W0-13, W0-14, W1-15	A chosen stack that co documented fallback; a
W1-18	1	Choose the GNSS base-station architecture and solve the survey-in versus 15-minute-setup conflict	High	REQ-CMP-002; SCI-009;	W0-13, Q-03	≤0.15 m horizontal rep procedure that fits the

				§3.d.v; RQ11; CF-13; IF-026		
W1-19	1	Choose the modularity and configuration plan, including the 70 kg accounting and the 10-minute Autonomy → Equipment transition	High	REQ-SYS-003, -006, -011; OPS-012; §1.a.iv; EQ4; CF-08; IF-035; SYS-013	W0-01, W0-19	Every configuration ≤5 10 minutes including a
W1-20	1	Decide whether the team flies a drone at all, and if so in which role and under what mass accounting	Medium	REQ-DRN-001…-011; §2.b; RQ9, RQ14, RQ15; CF-10	W0-11, W1-01, Q-04	A go/no-go with a meas drone ≤5 kg with the du worked
W1-21	1	Choose the chassis/frame architecture and the manufacturing process	High	REQ-SYS-001, -002, -014; MOB-007; IF-001/IF-007	W0-09, W0-10, W0-21, W1-03	Measured stiffness an survives 1.5× reaction instructions
W1-22	1	Choose the thermal architecture: sealed or filtered, white or shaded, and whether the pack gets its own compartment	High	REQ-PWR-005, -009; MOB-011; IF-028…IF-033; CF-06; SYS-012	W0-15, W0-16, W1-15	Chosen architecture k ambient with the full w
W1-23	1	Choose the astronaut-interaction architecture: what the suit carries, which command ladder is built, and how the suit feed reaches C2	High	REQ-SAF-005; COM-015; CMP-006, -007, -008; SFT-004, -007; §1.e.iv–xi; AUQ5, AUQ10–AUQ12, AUQ15, AUQ24, AUQ26; CF-12	W0-14, Q-12, Q-19	A ladder that scores ≥1 min, and a suit feed on
W1-24	1	Choose the software stack, the simulation path and the freeze policy	High	REQ-SFT-001…-019; AUQ2, AUQ4, AUQ6, AUQ13, AUQ21; OPS-001	W1-15	One command path fro on hardware; a dated f
W1-25	1	Choose the C2/operator architecture: seats, per-role pages, logging and the authoritative clock	High	REQ-OPS-003, -005, -006; COM-003; CMP-014; IF-020/IF-025/IF-034/IF-041	W0-19	Six seats each with a p leaving it; a single auth CSV
W1-26	1	Produce the PDR submission package from the evidence the above rows generate	Critical	REQ-OPS-001; §3.a.ii	W1-01…W1-25	Every requirement eith owner and a date; a ma no capability claimed w
W2-01	2	Define and freeze the mass/CG ICD: a weighed value, a CG height and a mounting interface per subsystem, plus the stability analysis on the resulting geometry	High	REQ-SYS-001, -002, -003, -006; MOB-007; IF-001, IF-002, IF-003	W1-01, W1-03, W1-06, W1-19	A mass/CG ICD whos number (roll and pitch) table
W2-02	2	Define and freeze the power ICD: rails, currents, fusing, wire gauges, connector families, grounds and the kill-path interfaces	Critical	REQ-PWR-001, -004, -007; SAF-001, -002; IF-011… IF-017	W1-04, W1-12, W1-13	Every rail documented deliberate short blows
W2-03	2	Define and freeze the data ICD: CAN and Ethernet topology, message schema, telemetry packet, one authoritative mode enum, time sync and logging format	High	REQ-CMP-001, -013, -014; SFT-008, -016, -018; IF-018…IF-025	W1-15, W1-24, W1-25	Every interface has a s the session; the LED is

W2-04	2	Define and freeze the thermal ICD: dissipation per box, air path, filter access, surface finish and the temperature telemetry the software must consume	High	REQ-PWR-005, -009; IF-028…IF-033; CMP-014	W1-22	A thermal ICD with limi that would let an opera
W2-05	2	Define and freeze the timing/operations ICD: setup, teardown, weigh-in, the 10-minute transition, mission clocks and who announces what	High	REQ-SYS-008, -009, -010, -011; OPS-003, -005, -012; EQ4; IF-034…IF-041	W1-19, W1-25	Every phase has an ow transition all fit with ma
W2-06	2	Close the drivetrain analysis on the measured μ: torque per wheel, skid-scrub torque in a spin turn, gear/belt backlash and the predicted slip on the steepest slope	High	REQ-MOB-001, -007, -009; IF-011	W0-09, W1-05	Predicted slip ≤20 % a continuous rating; a do
W2-07	2	Design the suspension: link geometry, articulation range, wheel travel, hardpoints and the folding (or non-folding) structure	High	REQ-SYS-001; MOB-001, -007; IF-004	W1-02, W1-03, W1-06	Contact fraction ≥X on inside the measured en
W2-08	2	Close the structural analysis: drop and shock cases, arm shoulder reaction, mast bending, drill reaction and fastener retention	High	REQ-SYS-001, -002, -014; MOB-007; IF-001, IF-005, IF-007, IF-010	W1-03, W1-06, W1-21, W0-21	No permanent set at 1 critical fastener witnes
W2-09	2	Produce the arm kinematics and prove reach and stow against all nine Equipment poses plus the hammer pickup	High	REQ-ARM-001, -002, -011, -018; §1.d.i; CF-02	W1-06, W1-07	Every pose reachable within the vertical budg
W2-10	2	Size every arm joint for torque, speed, thermal duty and holding, and set the current limits and brake/counterbalance policy	High	REQ-ARM-002, -020; IF-012; IF-029	W1-05, W1-07, W1-08	Every joint delivers the survivable for ≥60 s; ho
W2-11	2	Design the end effector: jaws, tips, compliance, force sensing, cable routing and the wrist-camera mount	High	REQ-ARM-002…-011, -019, -020; IF-021/IF-022	W0-03, W1-09	≥95 % task success p within the measured w
W2-12	2	Design the typing routine: two-stage visual servo, latency budget, error recovery with backspace, no auto-repeat, and a judge-reset display	High	REQ-ARM-011, -012, -019; SFT-014, -015; CMP-003; EQ2, EQ3	W0-04, W1-08, W2-09, W2-11	≥3 of 4 key lengths typ key without abort; no d
W2-13	2	Design the drill/auger: reaction-torque path, slip clutch, stroke, graduated liner, sample transfer and the scoop fallback	High	REQ-SCI-004, -005, -017; ARM-015, -016; AQ6, AQ10; ARM-016	W0-07, W1-10	≥5 g below 10 cm in ≤4 chassis reaction budge sites
W2-14	2	Design the fluidic path: pumps, metering, cuvettes, waste containment, inversion sealing, carry-over control and the heater decision	High	REQ-SCI-001, -002, -012, -018; AQ7, AQ9	W0-08, W1-11	Metering within ±5 %; when inverted for 10 m
W2-15	2	Produce the power budget v2 and freeze the pack, fuse, contactor, precharge and BMS threshold numbers on measured loads	Critical	REQ-PWR-002, -004, -007, -008; SAF-002; IF-011, IF-012, IF-013; PWR-003	W0-18, W1-12	No brownout at the coi limit; every fuse and th
W2-16	2	Design the thermal system v2: enclosure, fan/filter, sunshade, altitude derate and the battery compartment	High	REQ-PWR-005, -009; MOB-011; IF-028, IF-033	W0-15, W0-16, W1-22	Sustained workload wi a filter service interval
W2-17	2	Produce the link budget v2 and the frequency plan from measured radio data, and site the antennas	High	REQ-COM-004, -006, -008, -010, -011, -013; §1.c.ii, §3.d.ii; MOB-003	W0-11, W0-12, W1-14	A margin figure per link §2.c.v for every link; a
W2-18	2	Design perception: camera and lens choice per task, field of view, calibration, exposure strategy and sun shading	High	REQ-CMP-003, -004, -005, -006, -009, -010; IF-018, IF-022	W0-14, W1-15	≥95 % decode at the ra px; a documented exp

W2-19	2	Design localisation: RTK configuration, dual-antenna heading, EKF tuning, the GNSS-shadow behaviour and the error budget for the 1 m stop	High	REQ-CMP-001, -002; MOB-005, -006; SFT-001, -012; IF-023/IF-026	W0-13, W1-18	A documented error bu after fix loss that break
W2-20	2	Design navigation: DEM costmap, slope limit from measured data, planner choice, arrival logic at 0.6 m and the progress monitor with bounded recovery	High	REQ-SFT-001, -002, -003; MOB-006; CMP-012; §1.e.xv–xviii; AUQ17; CMP-015	W0-09, W0-13, W1-17	≥80 % of the driven rou routes around DEM art
W2-21	2	Define the mode/state machine: the authoritative autonomy state, the LED drive, the hands-off lockout, and the log that proves both	High	REQ-CMP-013, -014; SFT-008, -009, -018; §1.e.ii–iii; AUQ3, AUQ9, AUQ18, AUQ21; IF-025	W1-17, W1-24	The LED follows the au mission; zero control t
W2-22	2	Design the astronaut command pipeline and the stop-distance budget for the 3 m rules	High	REQ-MOB-004, -005; SFT-004, -005, -010, -011; CMP-006, -007, -008; §1.e.vi–xi; AUQ7, AUQ10, AUQ13, AUQ15; ARM-021	W0-14, W1-23	Every stop ≤3 m with ≥ person locks; a docum
W2-23	2	Write operations procedures v1: setup, teardown, weigh-in, four mission run sheets, the abort ladder, intervention and the runner rule	High	REQ-OPS-003, -005, -006, -008, -009, -011, -012; SYS-008…-011; §3.b, §3.e; COM-014, MOB-010	W0-19, W1-19, W1-25	Each sheet runs end-t every penalty trigger is discussion
W2-24	2	Write the safety case and the E-stop verification plan, including the battery-fire response and the astronaut separation rules	Critical	REQ-SAF-001…-008; PWR-008; §2.a.v; §3.e.ii; AUQ1	W0-17, W1-13, W1-12	Every safety requirem anyone on the team ca run sheet
W2-25	2	Assemble the SAR package and video storyboard from the test record, not from renders	Critical	REQ-OPS-001, -002; §3.a.iii; SCI-015	W0-23, W1-11, W1-14, W1-26, W2-28	Every requirement in § yet"; the video shows f
W2-26	2	Set the configuration freeze plan: dated freezes, a change log, a known-good reference and a rollback image	High	REQ-SYS-014; SFT-017; OPS-001	W1-19, W1-24	A freeze date per interf change log that a revie
W2-27	2	Design the drone if the team is flying one: airframe, dummy mass, relay payload, launch/land cycle and the pilot/spotter procedure	Low	REQ-DRN-001…-011; §2.b; RQ9, RQ14, RQ15; AUQ22, AUQ25; OPS-010	W0-11, W1-20	Endurance covers the when the rover is behir rehearsed
W2-28	2	Run filmable pre-SAR evidence for all four missions in weeks 15–17, because the SAR video is due 2027-03-05 and the Wave-3 test programme starts after it	Critical	REQ-OPS-002, SCI-015; §3.a.iii	W0-14, W2-12, W2-14, W2-23	Every mission's core c test record; the SAR p not displace a freeze
W3-01	3	Integrate the electrical system on a bench replica: power, harness, kill loop, LED, compute, radios and one actuator class	High	REQ-SAF-001, -002; PWR-004, -007; SYS-004	W2-02, W2-03, W2-15, W2-24	No rail live after the kill from a deliberate brow
W3-02	3	Bring up the drivetrain on stands and then on the ground: current, temperature, backlash, wheel retention and the post-assembly fastener review	High	REQ-MOB-001, -007, -009; SYS-002	W2-06, W2-07, W2-08	Current within 10 % of inside the ICD; wheels

W3-03	3	Run the first integrated field day: drive, link, camera and telemetry for 30 continuous minutes at the analogue site	High	REQ-COM-003, -004, -006; SFT-019; SYS-008	W0-20, W1-15, W2-03, W2-17	Zero reboots in 30 min driving task without lea
W3-04	3	Run the full Equipment mission against the mock lander, timed, with all nine sub-tasks attempted in the order the points suggest	High	REQ-ARM-003…-011; MOB-002; SYS-011; §1.d.i	W2-09, W2-11, W2-23	≥6 of 9 sub-tasks comp recorded points-per-mi evidence says so
W3-05	3	Accept the typing routine end-to-end: 3–6 letters, backspace corrections, no auto-repeat, judge-style display reset, no contact with the keyboard frame or display	High	REQ-ARM-011, -012, -019; SFT-014, -015; EQ2, EQ3	W0-04, W2-12	≥15/20 correct with no frame/display contacts
W3-06	3	Accept the key/lock task end-to-end, including a tethered key and recovery from a dropped or mis-inserted key	High	REQ-ARM-007, -019, -020; §1.d.ii	W0-03, W2-11	≥25/30 successful tran without a human; the 1
W3-07	3	Verify the task forces on the real hardware against the W0-03 measurements, and set current/force limits from the result	Medium	REQ-ARM-008, -009, -010, -020; §1.d.ii	W0-03, W2-11	Every task force within documented per task
W3-08	3	Run the full Astrobiology mission on the clock: two sites, documentation, sampling, two assays, cache seal and hand-over, then the debrief	Critical	REQ-SCI-003…-014, -019; §1.b.ii–x; AQ8, AQ10	W0-22, W1-11, W2-13, W2-14, W2-23	Both life-detection resu handed over inside 5 m numbers
W3-09	3	Accept the sampling system in three soils, 10 runs, including the scoop fallback when the drill refuses	Medium	REQ-SCI-004, -005, -016, -017; ARM-015; AQ6, AQ10	W0-07, W2-13	≥9/10 runs deliver ≥5 g fallback recovers a fail
W3-10	3	Run the autonomy trials: a 30-run 1 m stop distribution, the hill route, the out-of-LOS second target and a judge-visible LED at 20 m	Critical	REQ-MOB-005, -006; SFT-001, -002, -003, -013; CMP-004, -013; §1.e.xv–xviii; AUQ18; ARM-021	W0-13, W0-14, W1-17, W2-19, W2-20	≥27/30 stops inside 1.0 completed; 5/5 NLOS at noon
W3-11	3	Rehearse the astronaut sub-mission with a suited stand-in, running all five commands in order	Critical	REQ-MOB-004; SFT-004…-008, -010; CMP-006, -007, -008; COM-015; §1.e.iv–xi; AUQ5, AUQ7, AUQ13	W1-23, W2-22	≥40/50 twice in a row; suit feed confirmed bet
W3-12	3	Accept communications at 1 km with the real rover: three sub-bands, all links, link-loss behaviour and a crowded-spectrum day	Critical	REQ-COM-001…-011, -013; §1.c.ii, §2.c.v, §3.d.ii	W0-11, W1-14, W2-17	≥10 dB margin at 1 km the rover holds or retur
W3-13	3	Run the thermal soak: 38–40 °C ambient, full workload, sun on the enclosure, fan and filter service, and a fan-failure test	High	REQ-PWR-005, -009; MOB-011; §3.d.iii	W1-22, W2-15, W2-16	No thermal throttle on failure alarm the operat
W3-14	3	Run the dust and light-rain test, then review every connector, encoder and calibration it touched	High	REQ-MOB-011; PWR-005; §3.d.iii	W0-16, W2-16, W2-18	No functional loss; cali procedure that fits a m
W3-15	3	Test the packaged rover: shake it, transport it, power it on, and check every connector, fastener and calibration afterwards	High	REQ-SYS-001, -003; environment §3.3; judges' transport advice	W0-21, W2-08, W3-02	Every subsystem work the checklist catches
W3-16	3	Verify the E-stop under load, the last-gasp indication and the judge-visible LED, and run the battery-fire drill	Critical	REQ-SAF-001, -002, -003, -007; PWR-004, -008; §2.a.v	W0-17, W1-13, W2-24	0 V on every rail after c completed from the wr

W3-17	3	Rehearse the weigh-in, cube check and 10-minute transition with the real rover in every mission configuration	Critical	REQ-SYS-001…-007, -011; OPS-012; EQ4	W0-01, W1-19, W2-01, W2-05	Every configuration ≤5 check ≤3 min; the trans
W3-18	3	Run a fault-injection day and prove the safety behaviours: link loss, low battery, sensor fault, arm stall, brownout, camera loss	Critical	REQ-SFT-010, -017; SAF-008; §2.a.v; §3.e.ii; AUQ1	W2-21, W2-23, W3-10	Each fault recovered in understood; zero contr each event
W3-19	3	Prove power endurance: a full 60-minute Delivery profile on one pack, the regen case, and a charge from a 120 V circuit	High	REQ-PWR-002, -006, -008; §1.c.i, §3.d.i	W0-18, W1-12, W2-15	Mission completed wit inside the between-mis
W3-20	3	Cut and prove the software release candidate: CI, rollback image, replay review, calibration and dataset reference	High	REQ-SFT-016, -017, -018; OPS-001	W1-24, W2-03, W2-26	Release boots from a c the calibration referenc
W4-01	4	Close the SAR review response: record every reviewer question, answer it with an artefact, and fix the evidence gaps it exposes	High	REQ-OPS-001, -002; §3.a.iii	W2-25, W3-18	Every review item clos requires hardware that
W4-02	4	Submit the communications documentation package: FCC/licence basis, device list, EIRP table, sub-band plan and the antenna siting rules	Critical	REQ-COM-007, -009, -010, -011, -012; §2.c.ii–vi; RQ6–RQ8	W0-12, W1-14, W2-17	Submitted by 2027-04- and a procedure for ea
W4-03	4	Write and submit the Science Plan (~5 pages plus background science)	Critical	REQ-SCI-015; §1.b.i; AQ1, AQ9	W1-11, W2-14, W3-08	Submitted by 2027-05- W3-08 runs; the AQ9 ju
W4-04	4	Run the four-mission dress rehearsal on unfamiliar ground with judge role-play, a full six-seat crew and the field crew	Critical	All mission REQs; §1.a–e	W0-20, W3-04, W3-05, W3-08, W3-10, W3-11, W3-12	All four missions attem demonstrated at least
W4-05	4	Freeze the mission cards, abort thresholds, penalty ladder and intervention plan	High	REQ-SFT-010, -011, -012; OPS-005, -009; §1.e.xi, §1.e.xviii, §3.e; COM-014, PWR-003	W2-23, W3-16, W3-18	Every card answers "w thresholds are number
W4-06	4	Assemble the spares kit against the failure list and drill field repair to the target time	High	REQ-SYS-014; §3.e; environment §3.6	W3-02, W3-06, W3-14, W3-15	Ten drills, every repair every safety/RF-path
W4-07	4	Finalise transport, travel and the post-travel re-test protocol, because testing at MDRS is prohibited	High	REQ-OPS-007; §3.d.iv; environment §3.5	W0-21, W3-15	Crate loads in ≤30 min step requires a tool tha
W4-08	4	Build the finals-week plan: day-by-day roles, attendance limits, the unknown mission order and the drone/spare-battery decisions	High	REQ-OPS-003, -004, -011; CQ4; S-URC-LOGISTICS; OPS-010	W1-20, W2-05, W3-12	Every mission has a na booked across a 10-mi order
W4-09	4	Enforce post-SAR change control: repairs only, every change tested, and the known-good tag protected	High	REQ-SYS-014; failure synthesis FMEA-31; judges' design-stability advice	W2-26, W4-01	No change reaches the known-good tag is unc
W4-10	4	Prepare the judge-briefing pack: autonomy announcements, LED semantics, the suit-feed confirmation, drone permission protocol and the mission script	Medium	REQ-SFT-008, -009; CMP-013, -014; COM-015; §1.e.ii–iv; DRN-005; AUQ22, AUQ25	W1-23, W2-21, W3-16	The judge can state the launch permission and ambiguity in the pack
W4-11	4	Train alternates and rehearse the after-action review: replay every log, name every fault, rotate the seats	Medium	REQ-OPS-003, -005; SFT-018	W2-23, W3-20	Every seat has two tra with owners; the log al

W4-12	4	Decide and rehearse the fallback configurations: teleop-only, no-drone, no-LiDAR, and the reduced science payload	Medium	REQ-SFT-019; DRN-011; CMP-011; SCI-001	W1-17, W1-20, W3-18	Each fallback scores a event without retrainin
W4-13	4	Run the final weigh-in dry run, freeze the configurations and publish the weigh card	Critical	REQ-SYS-001…-007, -011; OPS-012	W0-01, W1-19, W3-17	A weigh card per confi the cube result; ≥1.0 k

The dependency spine

110 problems connect through **253 dependency edges**. Thirteen rows have no prerequisite, because they are pure measurements, and twenty-one are leaves, because they are submissions, rehearsals or drills. Thirteen rows carry five or more dependents, and they are the rows that decide when the rest becomes possible.

RANK	ID	PROBLEM	PRIORITY	UNLOCKS	WHY IT IS LOAD-BEARING
1	W0-01	Mass budget frame, weigh-in procedure and the 1.2 m gauge	Critical	**9**	Feeds the arm class, wheel size, stow concept, configuration plan and weigh-in
2	W0-11	Link budget model and the first 1 km non-line-of-sight test	Critical	**8**	Decides radios, bands, camera count, mast, drone value and the licence package
3	W0-14	Real tag and sign detection ranges in full sun	High	**8**	Sets every camera and lens, the arrival strategy and the gesture range
4	W0-09	Wheel and soil traction, and rolling resistance	High	**7**	The physical basis of torque, slope limit and route planning
5	W1-01	Rover mass class, arm mass class and the actuator cliff	Critical	**7**	The arm class quietly chooses the drivetrain actuator class
6	W1-03	Wheel diameter, wheelbase, track and stow geometry	Critical	**7**	The least reversible number in the project
7	W0-02	A judge-like Equipment Servicing mock	Critical	**6**	Makes every task-force measurement possible
8	W1-15	Compute and camera architecture against the link budget	High	**6**	Sets link load, thermal load and the software target
9	W1-19	Modularity and configuration plan	High	**6**	The 70 kg rule, the 10-minute transition and the operator roster in one decision
10	W2-23	Operations procedures, first version	High	**6**	Turns a test programme into repeatable evidence
11	W0-03	Task force and torque measurement campaign	Critical	**5**	Sizes the arm, the gripper and the compliance
12	W0-07	Soil and drilling force measurement	Critical	**5**	Decides the sampling architecture and the chassis reaction path
13	W0-10	Step and obstacle rig, climb against wheel diameter	Critical	**5**	Sets the wheel diameter, the torque demand and the tilt limit

Six of the thirteen rows are in the measurement wave. That is the register’s central claim about a first build: the numbers that size the machine are not published, so the early season is spent producing them.

Working the register every week

A register that is read once is a document. One that is reviewed weekly is a management system, and the minimum useful ritual is **45 minutes, five questions**, in this order.

1. **What closed this week, and what artefact proves it?** A row closes only when its success criterion is met and the record is filed where the SAR evidence can be found. A measurement either exists or the row is still open.
2. **What is now unblocked?** Every closed row carries the number of later problems it unlocks. Those are the candidates for next week, and they are the only rows that should be scheduled without an argument.
3. **Which Critical row is nearest its gate?** The PDR gate set and the SAR gate set are fixed, so anything in them that cannot close in the remaining weeks is a scope decision to take now.
4. **What did a measurement contradict?** A number that disagrees with a subsystem study is a finding: record it in the row and in the study, and re-run the rows that used the old number.
5. **Which question now blocks more than one row?** If a written question has grown dependents, escalate the email rather than guessing in two places.

A worked row

The wheel-diameter row, `W1-03`, shows what a well-formed entry looks like. The problem is stated as a decision rather than a preference: fix wheel diameter, wheelbase and track. It depends on four measurements: the mass frame (`W0-01`), traction (`W0-09`), the step rig (`W0-10`) and the arm mass class (`W1-01`). Its first experiment is a packaging study of 250, 280 and 330 mm wheels inside the cube gauge, plus a tip-angle calculation at each diameter. Its success criterion is a diameter that fits the cube with 20 mm per side, keeps the static tip angle at or above 55° as a derived target, and has a bearing and shaft design at the measured torque. The artefact is a signed envelope drawing, the packaging calculation and a photograph in the gauge. It unlocks seven rows: gear ratio, suspension, mast height, chassis, the mass and centre-of-gravity document, kinematics, and the structural load cases.

That shape is the test for every row: a decision, its inputs, one experiment, one number, one artefact, and the list of what it frees. A row with no number in its success criterion is a topic, not a task.

Coverage of the 155 requirements

Every one of the **155** requirements is traced by at least one problem: 155 cited, zero uncovered, zero unknown identifiers. The distribution is itself a check on the analysis.

DOMAIN	REQUIREMENTS	COVERED BY A ROW
Arm	21	21
Science	20	20
Software	19	19
Communications	15	15
Compute and perception	15	15
Systems	14	14
Operations	12	12
Mobility	11	11
Delivery and drones	11	11
Power	9	9
Safety	8	8
Total	**155**	**155**

The drone is the thinnest area, at eleven requirements carried by four problems and one of them Low priority, which matches the absence of any strong drone evidence in the record. Safety has eight requirements and eight problems, three of which are tests, the right ratio for a pass-or-fail clause.

Twenty-one requirements are carried by exactly one row, and twelve of those are mandatory. Two are worth watching, because a single scheduled item is load-bearing for something scored. The panorama and close-up requirements, REQ-SCI-007 and REQ-SCI-008, both rest on the Astrobiology rehearsal W3-08, so a documentation-only dry run is added earlier. REQ-SAF-004, that every mobile deployable carries its own E-stop, rests on W2-24 and is correct only while nothing is deployed.

Twenty questions the register gates

Twenty questions in the register are the ones no amount of engineering can settle. Each is a written question to the URC Director, each gates named rows, each has an answer date, and each has a conservative fallback so that an unanswered question never stops work. The four with the largest design consequences are below; the open questions page carries the wider set of 77.

QUESTION	WHAT IT DECIDES	ASK BY	FALLBACK IF UNANSWERED
Does one instrument with two readouts count as two life-detection capabilities?	Payload architecture and the SAR science page	2026-11-15	Build two independent channels on one detector
Do the antenna height and base-distance limits apply to a separate GNSS base antenna?	The RTK geometry and the 1 m stop	2026-11-15	Apply 3 m and 5 m to all base antennas
Must a spare drone battery carried but not connected sit inside the E-stop rule, and does it count in 70 kg?	Battery architecture, the drone plan and the weigh-in	2026-12-01	Carry no spare drone battery on the rover
What provenance record must accompany the cache?	The cache design and the hand-over procedure	2026-12-01	Print a provenance card, seal it and photograph it in place

Each question goes as one numbered email that quotes the clause, states the interpretation the team intends, and asks for confirmation or correction, with a request that the answer be mirrored onto the public Q&A page, which is itself a binding rule source. The reply is recorded against the rows it gates.

What would make this register wrong

- **The 2027 numbers move.** The 50 kg, 70 kg and 1.2 m limits, the five milestone dates and the mission structure are confirmed for 2027. A rules revision would force every row that cites the mass, the timeline or the mission structure to be re-read before the next review.
- **A task weight changes the order, not the content.** The Astrobiology, Delivery and Equipment Servicing weights are unpublished. When they arrive, the sequence inside the last two waves changes, because the highest-points task is attacked first; the problems themselves do not.
- **A measurement can overrule a study.** Every value carried over from a subsystem study is an estimate until the measuring row closes. Where a measurement disagrees, the measurement wins and the row records the difference.
- **The plan assumes one rover.** Two mission configurations with a real payload change would grow W1-19 and W2-05 and turn the 10-minute transition into a rehearsal item with its own test.

Further reading

- 70 risks points every risk at a problem row, so a risk with no row is a missing row here.
- 77 open questions carries the full question set and the meanwhile assumption for each.
- The first 90 days reads these dependencies forwards, week by week.
- 155 engineering requirements is the table every row traces into.
- Coupling, conflicts and twelve trade studies settles the design decisions the architecture wave tests.
- Interfaces and couplings supplies the interface rows the detailed-design documents must contain.

THE BUILD PLAN · CHAPTER 8.1

70 risks

The register holds 70 scored rows in six categories. Each row carries a likelihood and an impact, and both come from the competition rather than from engineering intuition. Likelihood is anchored to a base rate in the public record wherever one exists, and is a stated judgement where none does. Impact is measured in the scoring architecture, so that a row either costs points inside a mission, zeroes one or more missions, or ends the field entry altogether.

70

SCORED RISKS

Six categories, each with a likelihood and an impact

2026-2027

9

HIGH AGAINST HIGH

The rows that carry the argument of the register

2026-2027

36

HIGH IMPACT

Removes the field entry, zeroes two missions or risks a person

2022-2026

162

TEAM-YEARS BEHIND THE BASE RATES

Finals appearances, 2022 to 2026

2022-2026

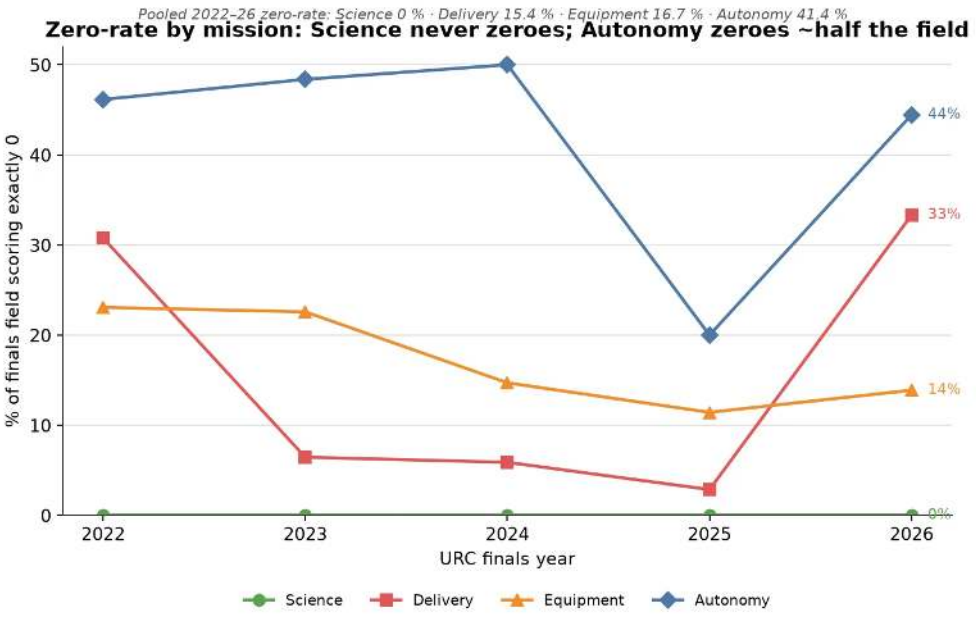
Key takeaways

- **Nine rows are high likelihood and high impact, and six of the nine are about process rather than hardware:** how the review is assembled, when the build is aimed, whether integration time exists, how deep the team is, what the scope is, and whether mass is tracked.
- **The top row is the review gate.** Two of three registrants never reach a score sheet, and the published cut has sat at 84 to 87 points on recent rubrics.
- **One safety rule changed in 2026, not 2027.** A team copying a 2024 or 2025 kill path can be compliant with the rule it read and non-compliant with the rule it faces.
- **A connector is the highest-ranked technical risk.** The judges' own advice names it, and the best-documented instance cost a top team about 45 points.

How likelihood and impact are set

BAND	LIKELIHOOD	IMPACT
High	A named base rate at 50 % or above	Removes the field entry, zeroes two or more missions, or puts a person at risk
Medium	10 to 50 %	Zeroes one mission or a block of 20 points or more
Low	Below 10 %	Costs points inside a mission
No base rate	The band is a stated judgement	The impact definitions above still apply

The bands are what make the register arguable. A base rate can be disputed, and that is the intended argument. The ranking rule is **likelihood times impact first, then gate before score, then upstream before downstream**, and it produces the order used below. The evidence behind each row is what to argue with, not the ordering.



The zero rate by mission and year, over 162 finals appearances. These rates are the base rates behind the high-likelihood bands: a mission that four teams in ten fail outright is a high-likelihood risk, not an unlucky one.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC-RESULTS *Full resolution*

The register in one matrix

	IMPACT HIGH	IMPACT MEDIUM	IMPACT LOW	TOTAL
Likelihood high	**9**	9	1	**19**
Likelihood medium	21	20	2	**43**
Likelihood low	6	2	0	**8**
Total	**36**	**31**	**3**	**70**

Six categories produce that matrix. Five categories would be the obvious division, but the five risks a first build must carry in 2027, which are a first-time programme, knowledge concentration, attrition, over-scope and an unowned calendar, fit none of technical, schedule, operational, safety or rules without distortion. They are therefore grouped as **team**, and they roll up into schedule in any strict five-category view.

CATEGORY	ROWS	WHAT IT COVERS
Rules and qualification	14	The review gate, the weigh-in, the cube, the communications filing, the chemical plan, the expense report
Technical	22	Margins that vanish under load, and behaviour that works on a bench and not in the field
Schedule	9	The calendar built as a sum of subsystem durations
Operational	11	The six-seat station, the two-configuration change, the drone crew
Safety	9	One element that has to work first time
Team	5	First-time programme, depth, attrition, scope, calendar ownership

All 70 risks

The table is searchable and sortable, and it filters by category. Each row carries the mechanism that produces the risk and the change that removes it, rather than a way to hide the symptom.

ID	RISK	CATEGORY	LIKELIHOOD
R-R01	The team misses the top-36 SAR cut and has no field competition at all	rules	H, 33–37 % of registrants appear in a score sheet (99→36 … 116→38); 2027 raises the document bar on both sides of the gate
R-R02	Invited, but the SAR scores low and gives away tens of points on the desk	rules	M, first-timer with no prior video; 2027 rubric unpublished
R-R03	Mass over 50 kg in a mission configuration at weigh-in	rules	M, the 50 kg limit has not moved since 2007 while the 2027 payload requirement grew (two life detections incl. a bioassay); FMEA-29 score first build 6/10
R-R04	Rover does not fit the 1.2 m cube at weigh-in	rules	L, no oversize event found in the aggregated 2022–26 data
R-R05	Fielded total exceeds 70 kg across all missions and parts	rules	L, most teams field one configuration, and per-configuration mass is the number everyone tracks
R-R06	900 MHz non-compliance: wrong sub-band at the wrong time, >8 MHz occupied, or a link the rule does not permit	rules	M, the required three sub-bands are 902–910 / 911–919 / 920–928, while least one candidate radio documents only two selectable bands (902–91 915–928)
R-R07	E-stop non-compliant with the 2026/2027 text, it stops the motors but does not switch off *all systems on every battery*	rules	M, the sentence "All systems on the rover must have the power switche by this single E-stop" entered in the 2026 rulebook (§3.a.v) and is uncha in 2027 (§2.a.v); kill chains published by 2024–25 teams predate it
R-R08	Hazardous-chemical plan rejected, or inconsistent with the SAR science page	rules	M, 2027 makes wet chemistry mandatory and the plan must cover transportation, usage, disposal, precautions and accidents; most 2026 corrections were "cradle-to-grave" tracking
R-R09	Communications documentation not submitted by 2027-04-30	rules	M, device and licence details are due 2027-04-30 and any change must notified immediately; the astronaut link, drone link and any repeater all c
R-R10	Rover Expense Report late, or missing → disqualification	rules	L, one document, one date, warned twice
R-R11	No-spill violation, bioassay fluids reach the ground	rules	M, every team flying a bioassay flies fluids, and the no-spill clause appli water and reaction products too; the desert has "very limited water supp and no quick access to emergency medical care"
R-R12	Drone cannot be flown legally or safely at the site	rules	L, Remote ID, registration if >250 g, TRUST proof, <400 ft AGL, class G National Park flight, ~5 miles from Hanksville airport
R-R13	Registration deadline 2026-10-30 missed	rules	L, five days out, nothing to submit but team details
R-R14	Attendance limits exceeded at the finals	rules	M, the rulebook defers the numbers to the website; the logistics page gi 24 students / 2 advisors / 14 individuals per day [S-URC-LOGISTICS]; a first-time 24-person team plans to bring more
R-T01	The link collapses behind terrain at 1 km, or at the intentionally out-of-LOS second target	technical	H, measured twice: 850 m LOS → 300 m NLOS, and 750 m LOS → 164 NLOS in a high-noise environment; the 2027 route target is *designed* t out of LOS
R-T02	A connector or harness fails after the transport leg	technical	H, the judges call it the most common field failure ("works perfectly in th motel parking lot and then not at all because of a loose wire or connecto database entries incl. one damaged USB-C that took out several camer and two motors
R-T03	Spectrum congestion from 35 other teams collapses video or telemetry mid-mission	technical	H, 36–38 teams on the same bands; FMEA-05 scores L 8/10, RPN 336
R-T04	Pack sag browns out the compute on a climb	technical	M, an 8S4P P42A pack sags ≈4.6 V (15.6 %) at 145 A; FMEA-06 RPN 3
R-T05	BMS trips on a legitimate two-joint arm stall	technical	M, a published BMS trips at 80 A for >10 ms with a 3 s auto-restore, a d that would trip on a legal arm hold
R-T06	Electronics or battery overheat at 37.8 °C in a sealed box	technical	H, the environment is a rule, not a guess; 40 W in a sealed 300×250×15 mm box at 38 °C ambient implies ~27 K over ambient, and one team retrofitted fans for exactly this
R-T07	Arm precision and backlash make the fine tasks impossible (typing, 1 cm key tags, keyway alignment)	technical	H, planetary joint backlash of 15–18 arcmin is 4.4–5.2 mm at 1 m, large than a ±1–2 mm keyway and comparable to a keycap pitch of 19.05 mn

R-T08	The key/lock task fails outright	technical	H, no URC precedent at all; the 1 × 1 cm tags are the smallest targets i rulebook
R-T09	The bioassay cannot finish inside the 20–30 minute window	technical	M, the clock is the rule, and the constructed time budget already spend 17.5 of 20 min with only ~2.5 min of slack
R-T10	The payload satisfies the words but not the intent: a chemistry demo that is not a bioassay under AQ9	technical	M, AQ9 says a chemical method alone does not satisfy the requirement amino acids alone are insufficient; a payload that scored 94/100 in 2026 2027 without a hardware change
R-T11	Autonomous person-following does not work: no detection at ≥20 m, or a false positive that drives the rover at the wrong thing	technical	H, no URC precedent for a commanded follow; the modality ladder fixes reward at 5/5/10/15
R-T12	The 1 m route stop is not repeatable	technical	M, 25 points per target and the tolerance is 1 m; no team publishes a hit FMEA-20 RPN 294
R-T13	RTK base-coordinate error puts both targets ~1 m off	technical	M, no CORS and no internet in the field, so the team runs its own base; nobody on the team has surveyed a base
R-T14	Dust ingress blinds optics, bridges contacts, loads the fan	technical	M, zero sourced URC mission failures are attributed to dust, which mak *unquantified*, not unlikely; every Phase-4 study recommends IP-rated connectors and lens protection
R-T15	Transport damage found on arrival	technical	M, the judges name a half hour of continuous shaking; one 2026 team fo parts damaged in transport and arrived days early specifically to find the
R-T16	A 5 kg object is dropped, or slips out of the grasp during the carry	technical	M, the envelope is 5 kg with a ≤7.5 cm grasp feature; the field's evidenc that retention, not strength, is the problem
R-T17	The sample is not ≥5 g from below 10 cm in hard soil	technical	M, soil ranges "reasonably easy to very hard and difficult to excavate" w no published force]
R-T18	Drone integration: over 5 kg with the dummy mass, or a launch the rules no longer allow	technical	M, ≤5 kg including a dummy mass equal to the battery mass, no high-w waiver, pad-only launch
R-T19	The autonomy stack gates driving on a link flag, or the planner stalls and burns the mission	technical	M, a single link_ok dependency stops the rover behind the hill; planner c is the field's most expensive software decision (0 → 14 → 86 by *simplifying*)
R-T20	A repeater or drone relay is planned as the answer to NLOS and is not proven	technical	M, repeaters are explicitly allowed and encouraged; no team in the sam documents a working relay at URC; anything carried counts toward ma
R-T21	Charging fails at the site (230 V-only charger, wrong connector, dead unit)	technical	L, one dual-voltage charger and a spare removes it; the site is 120 V 60
R-T22	Mass and complexity grow during integration, eating the 50 kg, the power budget and the thermal margin	technical	H, the pattern that generates the largest single-year losses: three of the biggest in this study follow a whole-system change; FMEA-31 RPN 294
R-S01	Build-to-June instead of build-to-March: the rover is assembled for the competition, not demonstrated for the SAR	schedule	H, the SAR is due 2027-03-05 and needs demonstration footage; the 20 rubric put 48 of 100 review points in the video; the PDR asks judges to assess readiness
R-S02	The PDR is treated as a formality	schedule	M, 5 pages, non-competitive, judged on "likelihood of being ready"; a firs time team writes its first one in week 6
R-S03	The December 2026 assay decision slips, so no bioassay exists in the SAR video, or in June	schedule	M, chemistry, reagents, fluidics, a chemical plan and a timed rehearsal downstream of one choice, and the SAR is 13 weeks after the PDR
R-S04	Long-lead hardware arrives late and slips the arm or the comms	schedule	M, one documented case: new arm gearboxes late enough that a team carried a heavy, mature arm instead; custom PCBs and machined joints the usual long poles
R-S05	The freeze date is not set, or it slips	schedule	H, T12 (design-process) appears at every level of the field, from the bes programmes to a team that lost every mission
R-S06	Integration and test time is not reserved	schedule	H, a first build's plan is a list of subsystems, and the sum of their duratio is the whole calendar; FMEA-31's mechanism is exactly this
R-S07	No field test campaign, no site, no date, or too late	schedule	M, testing at MDRS is prohibited before, during and after the event, and road testing is restricted to approved sites or town
R-S08	Capacity collapse over the university winter break	schedule	H, the 90-day window spans late December and early January
R-S09	Documentation effort collapses as build effort rises	schedule	M, one team's own SAR score fell 6.7 points while its total rose 57
R-O01	The six-seat C2 is over-subscribed, with no bench depth and the person with the clock also driving	operational	H, no more than 6 students in C2 at any time, advisors excluded; a first team of 24 has more willing operators than seats
R-O02	Runner forfeit: a team member leaves C2 to help and cannot come back; a spectator intervenes and terminates the mission	operational	M, "will NOT be permitted to return"; an unrequested intervention "count an emergency stop ... the current mission will be considered terminated

R-O03	The two-configuration change plus re-weigh does not fit in 10 minutes	operational	M, Equipment may start as soon as 10 minutes after Autonomy, same C adjacent course, and a re-weigh is required inside it unless the configura is unchanged and compliant or the same penalty is accepted
R-O04	Drone pilot/spotter discipline fails	operational	M, the pilot is in C2 with no direct view, a spotter maintains visual line of sight, and while the rover is autonomous only the drone pilot may touch controls
R-O05	Setup (15 min), teardown (10) or retrieval (10) overruns	operational	H, three clocks, all with judges on them, and the antenna mast plus 25 r cable are the slow items
R-O06	The mission order and day are unknown until about a week before, and differ between teams	operational	H, advisory but explicit: a schedule released about one week ahead, pe team task order
R-O07	A hand touches a control while the rover is autonomous	operational	M, hands-off is a mandatory rule enforced by observation; the documen failure mode is human reflex
R-O08	The astronaut E-stops the rover, or a spectator intervenes	operational	M, URC "would sympathize with an astronaut that E-stops a rover that getting too close"; an unrequested intervention terminates the mission
R-O09	Delivery Stage-1 gate time-out converts a good run into 15/100	operational	M, 14 of 31 teams scored exactly 15 in 2023; one team's own report is " out of time in Stage 1"; the 2027 stage durations and gate are unpublishe
R-O10	Autonomy time is misallocated across two sub-missions in 40 minutes	operational	M, 40 min for astronaut assistance + route finding, second target behin hill
R-O12	The LED is not visible in daylight, or shows the wrong colour, so autonomy credit is lost although the task worked	operational	M, red/blue/flashing green, daylight-visible; no photometric spec exists teams verify themselves
R-X01	Battery fire, in the rover or in a hot parked vehicle	safety	M, the rule names "a battery fire" as the motivating emergency; no URC rover fire was found in any source read, an absence of reports, not an absence of risk
R-X02	The E-stop does not remove power, welded contactor, failed MOSFET, or an opened kill loop	safety	M, contactor welding without precharge is standard contactor physics; burned kill-switch PCB and a rear axle with no kill path at all were found transport on a 2017 rover
R-X03	Arm injury: crush, pinch, or a dropped 5 kg mass / 0.91 kg steel hammer	safety	M, the system lifts 5 kg and a 0.91 kg 300 mm all-steel hammer; no forc limits are published and the Q&A only warns that lander fixtures are deli and rovers are not
R-X04	Rover–astronaut contact during the follow/come sub-mission	safety	M, the rover is asked to approach a human autonomously and stop with m; every T10 failure mode (wrong person, false-positive gesture, creep, overshoot) becomes a safety event
R-X05	Field heat illness, operators in a metal C2, the astronaut in a suit to the neck	safety	M, full sun, 37.8 °C, and the astronaut wears a suit covering to the neck wrists and ankles; the rules require the helmet be removable if they bec hot
R-X06	Chemical or fluid exposure from the bioassay payload	safety	L, the reagents are bench-scale and the quantities small; the site has no quick medical access
R-X07	Tip-over or runaway on a slope	safety	M, the judges state rovers have been "driven unintentionally into gulleys and even flipped over"
R-X08	Drone propeller injury or loss of control over people	safety	L, no URC drone injury found in any source read; the airframe is up to 5
R-X09	Manual handling injury lifting a ~50 kg rover inside the 10-minute transition	safety	M, the transition is timed and the rover is at the mass limit
R-M01	First-time team: no institutional template, no prior SAR or video, twelve subsystems learned at once	team	H, this is the team's first URC attempt
R-M02	Knowledge concentration: one or two people hold the planner, the arm's calibration, or the assay chemistry	team	H, a first-year club of 24 typically has one person per hard subsystem
R-M03	Attrition and capacity, spring semester, graduation, competing commitments	team	M, an 11-month project across a North-American academic year
R-M04	Over-scope on the four new 2027 capabilities, bioassay, astronaut interaction, 1 m route finding, key/lock, build all four, finish none	team	H, 2027 is the largest single-year rewrite since the four-mission format settled, and a team meeting a discontinuity in its rookie year has no recovery year
R-M05	No owner for the administrative calendar, the milestones that are not engineering	team	M, six hard dates between 2026-10-30 and 2027-06-02 and a first-time t has no one whose job is the calendar

The nine that carry the argument

	ROW	IN ONE LINE
1	R-R01	Miss the top-36 review cut, so there is no field entry at all
2	R-M04	Attempt all four new 2027 capabilities and finish none of them
3	R-T02	A connector fails on the transport leg and the rover is motionless at the start line
4	R-T01	The link collapses behind terrain, by design, at the second route target
5	R-T22	Mass and complexity grow during integration and eat the 50 kg, the power margin and the thermal margin
6	R-S01	Build to June instead of building to March, when the review needs demonstration footage
7	R-T07	Arm backlash is 4 to 5 mm at 1 m, which is wider than the keyway
8	R-M02	One person holds each hard subsystem
9	R-S06	Integration and test time is never reserved

The shape of that list is not a coincidence. Six of the nine are process rows, and three are technical, of which one is a connector. That is the same shape the failure analysis found from the other direction: rovers fail where a single element is load-bearing and untested under field loads, and the field under-tests exactly those elements, because testing them is slow and happens in a car park.

The one pass-or-fail row

R-R07 is the row that can end a season on a technicality. The requirement is that one button ceases all power draw from **any and all** batteries, and that all systems on the rover have their power switched off by that single E-stop `S-URC27-RULES` §2.a.v. The second sentence entered the rulebook in **2026**, not in 2027. A team copying a 2024 or 2025 design can therefore be compliant with the rule it read and non-compliant with the rule it faces, and the failure is silent until a judge asks for the current probe.

What removes it: a hardwired normally-closed loop to a direct-current-rated contactor with precharge, one connector at most in the safety path, and 50 presses under load with a voltmeter on every rail. What creates it: a software stop, a drive-bus-only kill path, or a design copied from a year whose wording was different.

The rows with the strongest evidence

R-R01, the review cut. The package is scored against more than a hundred others and lands outside the top 36, and there is no stand-by list. The published cut sits at 84 to 87 on the 2023 to 2025 rubrics, where the lowest admitted scores were 85.10, 84.04, 86.90 and 75.86. Two of three registrants never reach a score sheet, and the cut is not a quality filter: 25 team-years from 2022 to 2026 hold a review score with no finals result, including a 92.85 that would have been third best among that year's finalists. Better prose does not help; December test evidence does.

R-T01, the link behind a hill. Range falls from hundreds of metres to 160 or 300 m when terrain intervenes: two teams measured 850 m falling to 300 m, and 750 m falling to 164 m. In 2027 the second route target is *designed* to be out of sight. Congestion from a full field is the same failure mode from the other side. What removes it: the radio at the mast rather than at the end of 25 m of cable, a 900 MHz recovery plane, and a rover that keeps planning when the link flag is false.

R-T02, the connector. The rover works in the car park and is motionless at the start line. This is the judges' own headline `S-URC-ADVICE`, and the best-documented instance is a damaged cable that cost a strong team about 45 points `S-WVU-NEWS2026`. What removes it: latched connectors, service loops, one connector at most in any safety path, and a packaged shake test followed by a full power-on, which almost nobody in the record runs.

R-T22, mass growth. Every interface grows one part, and together they eat the 50 kg, the power margin and the thermal margin. Growth surfaces at the weigh-in, at the fielded total and at the stowage cube. The largest single-year score losses in the record follow whole-system changes, and the mass limit has not moved in twenty years while the payload requirement grew. What removes it: a weekly weigh-in from first assembly, a mass owner who can refuse a part, and a design target well below the cap.

R-S01 and R-S06, the calendar. The review is due 2027-03-05 and needs demonstration footage; on the only rubric available, the 2026 one used as a proxy, **48 of 100 points were in the video** `S-URC26-SAR`. A rover that first runs in April cannot be filmed in March.

The related row is that the plan is a list of subsystem durations added together, so integration receives nothing and every mission becomes a first attempt at the event. After the review there are **89 days**, and they are a test window rather than a build window.

What would change this register

The register is written so that it can be falsified. These events should delete or re-score a row rather than be reinterpreted around it.

EVENT	ROWS IT CHANGES
A first-time-specific review failure rate is computed from the 25 review-only team-years	R-R01, R-S01
URC publishes a non-finish or intervention log	R-T02, R-T05, R-X02, R-R07
The 2027 review rubric is published	R-R01, R-R02, R-S09
The 2027 Delivery sequence is published	R-O09, R-S06
URC answers that one combined assay suffices	R-T09, R-T10, R-R08, and team decision D5
URC answers that the antenna limits apply to all base antennas	R-T12, R-T13
The first field day shows signal collapse between 300 and 700 m	R-T01, R-T03, R-T20
The arm measurement shows 5 mm or worse repeatability	R-T07, R-T08, and decision D10
The team signs a deferral list naming two or more capabilities	R-M04, R-R01, R-S06, R-T22
Any 2027 team publishes measured link, thermal or bus data	R-T04, R-T06

Further reading

- 77 open questions carries the questions whose answers move these rows.
- 110 engineering problems is where each risk finds its row; a risk with no row is a missing row there.
- How rovers fail is the evidence base behind the technical rows.
- What separates the top from the rest computes the zero rates and variance shares the likelihood bands rest on.
- The 13-week plan assigns the first thirteen weeks of these rows to dated weeks.

77 open questions

A first build carries more unknowns than it can close, and the useful response is not to close them all but to sort them by who can answer them. The 77 questions below are grouped that way: **16 go to URC by email**, **22 the team can measure on a bench**, **27 need a prototype or a field day**, and **12 are the team's own decisions**. Each carries a meanwhile assumption, so an unanswered question never stops work.

77
OPEN QUESTIONS
Sorted by who can answer them, with a fallback for each
2026-2027

16
TO URC BY EMAIL
One numbered letter in November, answers on the public page
2026-2027

22
MEASURABLE NOW
A bench, a load cell and one to three days each
2026-2027

27
NEED HARDWARE OR A FIELD DAY
Each with its pass criterion written in advance
2026-2027

Key takeaways

- **The judge questions go out first, in one numbered email in November 2026**, because their answers take longest and the public Q&A page has changed twice in a month after publication.
- **A week of measurement usually settles a decision before an answer arrives.** The 22 measurable questions are cheaper than the email that asks about them.
- **Every question has a meanwhile assumption**, and the assumption is the conservative one. An unanswered question moves a decision to the fallback; it does not stop the season.
- **Five documents are unpublished as of 2026-10-25:** the 2027 review rubric, the Science Plan format, the Delivery sequence and its stage gates, the attendance numbers and the partial-credit scale for route finding.

Four groups, four clocks

GROUP	COUNT	WHO ANSWERS	WHAT IT COSTS TO GET AN ANSWER
Ask URC	16	The URC Director, with the answer mirrored onto the public Q&A page	A written letter and weeks of waiting
Measure it	22	The team, on a bench or with a load cell	One to three days per question
Test on hardware	27	The team, with a prototype or an analogue site	A field day and a written pass criterion
Team decision	12	The team, in writing and dated	One meeting, and the real risk is not deciding

The 77 questions run on different clocks, which is the point of the grouping. The judge questions go out first, the measurements run in parallel, the hardware tests need a calendar slot, and the team decisions need a date and a signature.

All 77 questions

The table is searchable and sortable, and it filters by group. Each row carries the decision the question blocks, what settles it, who owns it and the meanwhile assumption.

ID	GROUP	QUESTION	DECISION IT BLOCKS
A1	Measure it	What is the insertion force and turning torque of a tubular pin-tumbler lock, and how much free play does the key have?	End-effector torque budget; whether a torque-limited wrist is enough
A2	Measure it	What push force seats a 1.5 in cam-lock with a half-thickness gasket, and at what depth does it seal?	End-effector design; whether the rover must brace against the lander
A3	Measure it	What are the Redragon K552 actuation force and travel per key, and how much lateral repeatability is needed?	Typing-tip design and the servo strategy
A4	Measure it	What are the drawer pull force, latch force, valve torque, and button/switch forces?	Task ordering and compliance design
A5	Measure it	What are the hammer's mass, length, handle diameter and CoM?	Fetch grasp geometry
A6	Measure it	What are the sample tube dimensions and the cache opening torque?	Cache/arm interface and the 5-minute handover
A7	Measure it	What is the joint backlash under load and at 38 °C?	Whether open-loop task execution is possible at all
A8	Measure it	What is the camera-to-joint-command latency?	Whether visual servoing can close the loop over the link
A9	Measure it	What is the mass and CoM of the assembled arm?	Stability analysis and wheel/motor sizing
A10	Measure it	What is the real peak current and pack sag during a simultaneous climb and arm stall?	Pack sizing, fuse, contactor, harness
A11	Measure it	Where do the BMS trip thresholds sit relative to a legitimate two-joint stall?	BMS choice and firmware current limits
A12	Measure it	What is the enclosure air and pack temperature on a 38 °C day in the sun?	Cooling design and the sealing decision
A13	Measure it	What throughput and latency does a two-hop repeater link give?	Whether a relay is part of the architecture
A14	Measure it	What is the encoded bitrate of each camera stream?	How many cameras the link can carry
A15	Measure it	What is the dust size distribution and concentration at rover height?	Filter specification and lens protection

A16	Measure it	What is the UV index and dose at the site?	Polymer and cable-jacket selection
A17	Measure it	What is the lid surface temperature in the sun?	Whether the box can be passive
A18	Measure it	What is the fan-filter service life in clay dust?	Filter maintenance interval in the field schedule
A19	Measure it	What vibration spectrum reaches the rover's mounts on the Hanksville road?	Crate design and fastener retention policy
A20	Measure it	What is the C2 power capacity, amps, circuits, UPS?	C2 kit and whether a UPS is needed
A21	Measure it	What is the C2 table geometry, and may we bring furniture or a second screen?	C2 layout and the six-seat design
A22	Measure it	What is the real distribution of setup times?	Whether 15 minutes is achievable at all
J1	Ask URC	Does one combined instrument, a biochemical assay that detects a biomolecule named in AQ9, satisfy both required life-detection capabilities, or must the two capabilities be physically independent?	Payload architecture: one instrument or two. Changing it late is ~1 kg, a fluidic channel and a second detector [subsystems/03 §8 #1]
J2	Ask URC	Does the E-stop's "any and all batteries" (§2.a.v) include a spare drone battery carried on the rover but not connected?	Battery architecture and the drone's spare-battery plan; a human swap is also an intervention (RQ14)
J3	Ask URC	Do the ≤3 m height / ≤5 m from C2 antenna limits (§2.c.iv) apply to a separate GNSS base-station antenna, not just the communications antenna?	RTK reference geometry and multipath at the C2; the 1 m stop depends on it [A29 in 2027-missions-and-scoring §9]
J4	Ask URC	Is the Equipment Servicing cache the team's own Astrobiology cache, or a judge-provided container?	Arm interface and whether one cache design serves both missions [A12; REQ-ARM-017]
J5	Ask URC	Is there a points penalty for exceeding 70 kg fielded total, or is it simply not fieldable?	Whether the 70 kg limit is a design constraint or a penalty line [A1; RQ2 says strict]
J6	Ask URC	Does the 40 % oversize penalty (§2.a.ii) apply per event or once?	Whether an oversize rover can still compete meaningfully
J7	Ask URC	Is a non-camera gesture sensor (e.g. an IMU on the astronaut's sleeve) compliant with the "visual recognition of a beckoning gesture" requirement?	The 15-point Follow path, and the astronaut's carried electronics [AUQ24, AUQ26; subsystems/12 §8]
J8	Ask URC	Does a repeater or drone relay count as a "mobile deployable" requiring its own E-stop, or does a fuse suffice because it cannot move on its own?	Whether the relay needs a second kill path and a second inspection [§2.a.vi; 2027-qa-digest §7 item 2]
J9	Ask URC	Is the astronaut speaking tactical directions to C2 ("turn left now") "human assistance" during autonomy?	Operator discipline and whether voice may be used tactically [AUQ2/AUQ21 vs AUQ15; subsystems/12 §8]
J10	Ask URC	If the astronaut presses the rover's E-stop, is the mission terminated and are points already earned preserved?	The follow-approach protocol and how close we allow the rover to be [§3.e.ii; AUQ1]

J11	Ask URC	Does the 5-point EVA-suit camera/mic award require the operators to act on the feed, or only that it is available and monitored?	Suit electronics complexity and the C2 display layout [§1.e.iv; subsystems/12 §8]
J12	Ask URC	Is there a published partial-credit scale for route finding ("partial points will be given"), and what is it?	Route strategy: attempt the far target or bank the near one [§1.e.xvii; A17-class gap; F7 in 05-trends §5.2]
J13	Ask URC	May the 3–6 letter launch key contain digits?	Typing routine's character set [subsystems/02 §5.7]
J14	Ask URC	Does a compressor or vacuum pump need a Mars-feasibility exception under §2.a.iv / RQ12?	Sample handling and pneumatic options [A10]
J15	Ask URC	What are the penalties for a 900 MHz sub-band violation, late communications documentation, or a no-spill violation?	Whether these are compliance gates or priced risks [no formula in §2.c, §3.e; A1-class gap]
J16	Ask URC	May the science-debrief pair watch live field video during the mission, or does that count as "seeing the site"?	Staffing: whether two C2-only people must stay blind [AQ8; subsystems/09 §8 #7]
C1	Test on hardware	Does the link hold ≥1 km in a terrain shadow on the 900 MHz plane alone?	Comms architecture; whether a relay is needed
C2	Test on hardware	Does the radio support all three 900 MHz sub-bands, with ≤8 MHz occupied?	Radio selection and the compliance sheet
C3	Test on hardware	Does the rover continue autonomously and signal arrival with the link down?	Whether the second route target is reachable at all
C4	Test on hardware	Can a 20 cm 4x4_50 tag be decoded at 10–20 m in full sun, and a 1 cm tag at task range?	Camera/lens selection and the arrival strategy
C5	Test on hardware	Is the base-coordinate absolute accuracy enough to guarantee a 1 m stop?	Localisation and the survey-in procedure
C6	Test on hardware	Is the 1 m stop repeatable?	Route-finding strategy and internal target
C7	Test on hardware	Does the follow controller hold stand-off at the astronaut's real walking speed?	Follow design and the safety case
C8	Test on hardware	What is the gesture recognition rate and false-positive rate at 5–20 m in sun?	Whether gesture (15 pts) is attempted at all
C9	Test on hardware	Does the audio channel survive wind and motor noise?	Audio (10 pts) as the primary Follow modality
C10	Test on hardware	Is the hammer grasp repeatable on the real product?	Fetch (10 pts)
C11	Test on hardware	Does the arm survive driving loads?	Arm mounting, carry pose, and whether the arm can be rigid

C12	Test on hardware	Do the wheels stay on under impact?	Wheel retention design
C13	Test on hardware	What slope can the rover actually climb on loose soil?	DEM slope limit for the hill route
C14	Test on hardware	Does a full mission complete on the clock?	The whole rehearsal schedule
C15	Test on hardware	Is the Autonomy → Equipment transition achievable in 10 minutes?	Configuration strategy (R-O03)
C16	Test on hardware	Does the bioassay answer the life question inside the window, with controls?	The payload clause (mandatory)
C17	Test on hardware	What are the extraction efficiency and false-positive rates in MDRS-type soil?	Which assay survives contact with clay and humics
C18	Test on hardware	Do the reagents survive 38 °C for five days?	Reagent choice and packing
C19	Test on hardware	Does the cache seal survive inversion and gloved removal?	Cache design (mandatory)
C20	Test on hardware	Can we get ≥5 g from below 10 cm in hard soil, in time?	Sampling architecture
C21	Test on hardware	Does the 1 m DEM (and the free LiDAR) resolve the navigable gullies?	Route planning approach
C22	Test on hardware	Does the autonomy stack complete ten cold-boot full runs?	Whether autonomy is a capability or a demo
C23	Test on hardware	Does the E-stop remove all power from every battery, repeatedly?	Eligibility (R-R07)
C24	Test on hardware	Does the rover survive 30 minutes of packaged shaking?	Transport design (R-T02, R-T15)
C25	Test on hardware	Does the electronics box survive 40 min at 40 °C with the real workload?	Cooling design (R-T06)
C26	Test on hardware	Can the team diagnose and repair in 20 minutes outdoors?	Spares kit and the field-repair plan
C27	Test on hardware	Is the LED visible at 20 m at local noon in both colours?	Autonomy credit (R-O12)
D1	Team decision	Which of the four new 2027 capabilities we attempt, and which we explicitly defer	It is the mechanism behind R-M04, and every week of indecision costs two weeks of integration
D2	Team decision	One configuration for all four missions, or modular configurations	Sets the 50/70 kg budget and whether R-O03's 10-minute re-weigh exists at all
D3	Team decision	Drone: fly it or leave it home	Needs a pilot and a spotter and two backups; full points are possible without it
D4	Team decision	Autonomy command modality (device 5 / sign 5 / audio 10 / gesture 15)	Determines perception effort and the astronaut's carried electronics
D5	Team decision	Two independent life-detection methods, or one instrument + one chemistry channel	Mass, contamination control and J1's answer

D6	Team decision	The freeze date and the known-good tag	Without it, R-S05 is certain
D7	Team decision	The six C2 seats, the commander, and who never touches a control	Named seats are the mitigation for R-O01, R-O02, R-O07
D8	Team decision	Whether to field the bioassay at all if J1 says "two independent" and mass is binding	A late "we can't fit it" is a rule breach, not a trade
D9	Team decision	The autonomy abandonment minute and the sub-mission split	Prevents R-O10 from eating the second sub-mission
D10	Team decision	Whether the key/lock task is attempted at all	It is the highest-uncertainty task and its points are unpublished
D11	Team decision	Which field sites we will use, given no MDRS testing	R-S07; the rule prohibits testing at MDRS and restricts off-road testing
D12	Team decision	Whether we keep and publish our own failure log	The single largest gap in the field's knowledge, and the one asset no competitor has

The sixteen questions for URC

These go as one numbered letter in November 2026, each quoting its clause and stating the interpretation the team intends. Answers bind every team once they reach the public page.

#	QUESTION	WHAT IT DECIDES	MEANWHILE
J1	Does one instrument with two readouts satisfy both required life-detection capabilities, or must they be independent?	One instrument or two; about 1 kg and a second fluidic channel	Two capabilities on a shared detector
J2	Do “any and all batteries” include a spare drone battery carried but not connected S-URC27-RULES §2.a.v?	Battery architecture and the drone spare-pack plan	Yes; every carried pack sits behind the kill path
J3	Do the antenna height and base-distance limits apply to a separate GNSS base antenna?	The RTK geometry and the 1 m stop	The limits apply to all base antennas
J4	Is the Equipment Servicing cache the team’s own Astrobiology cache, or a container provided by the judges?	The arm interface and whether one cache design serves both	One cache design serves both
J5	Is there a points penalty for exceeding 70 kg, or is the rover simply not fieldable?	Whether 70 kg is a trade or a gate	A gate
J6	Does the oversize penalty apply per event or once?	Whether an oversize rover can still compete meaningfully	Sized to fit; the penalty is not planned as absorbable
J7	Is a non-camera gesture sensor, such as an inertial unit on a sleeve, compliant with visual gesture recognition?	The 15-point Follow path	No; vision is built and audio is the reliable 10
J8	Does a repeater or drone relay need its own E-stop as a mobile deployable, or does a fuse suffice?	Relay mass and inspection	A fuse is assumed to suffice
J9	Is an astronaut giving tactical directions to the C2 station “human assistance” during autonomy?	Operator discipline and whether voice may be used tactically	Voice for safety and situation reports only
J10	If the astronaut presses the rover’s E-stop, is the mission terminated and are earned points kept?	The approach protocol and the stand-off	The plan keeps it from happening
J11	Does the suit camera and microphone award require operators to act on the feed, or only that it is shown?	Suit electronics and the C2 layout	The feed is always shown and the judge is told it is live
J12	Is there a published partial-credit scale for route finding?	Whether to attempt the far target or bank the near one	Partial credit is assumed small
J13	May the 3 to 6 letter launch key contain digits?	The typing character set	Letters and digits are both supported
J14	Does a compressor or vacuum pump need a Mars-feasibility exception?	Sample handling and pneumatic options	Compressors are avoided
J15	What are the penalties for a sub-band violation, a late communications package or a no-spill breach?	Whether these are gates or priced risks	All three are treated as gates
J16	May the science debrief pair watch live field video, or does that count as seeing the site?	Two C2-only people kept away from imagery	Two people see no imagery until the debrief

The measurements worth starting first

The 22 measurable questions are the cheapest in the set, and six of them unblock more than one subsystem. Each is a day’s work with a load cell, a stopwatch or a rig.

- **The lock:** insertion force, turning torque and key free play on a pin-tumbler lock of the same class. It sets the end-effector torque budget.
- **The keyboard:** actuation force and travel per key, and the lateral repeatability needed at the camera standoff. It decides whether open-loop typing is possible at all.
- **Backlash under load at 38 °C:** the datasheet figure is nominal, and every precision claim in the arm design rests on it.
- **Pack sag at the real peak:** a 150 A pulse for the legitimate peak and a 200 A pulse for the trip, at low state of charge.
- **Setup time:** ten setups timed with the mast and 25 m of cable. The 15-minute rule is a measured quantity or an assumption S-URC27-RULES §1.a.ii.
- **The cache:** tube dimensions, opening torque and a gloved removal, which together decide the 5-minute hand-over.

What a hardware test looks like

The 27 hardware questions are written as pass criteria before the test runs, which is what separates them from demonstrations. Six of them carry most of the risk.

TEST	PASS CRITERION
Link at 1 km in a terrain shadow, on the lower band alone	Drivable at the worst shadow point; recovery inside 10 s
1 m stop repeatability	90 % of runs inside 1.0 m, and 100 % inside 1.5 m
Follow stand-off at the astronaut's walking speed	2.5 to 4.5 m for 90 % of the run, and never closer than 1.5 m
The bioassay inside the mission window	Both results by T+23:00, with controls behaving
E-stop removes all power, repeatedly	Zero volts at every rail, 50 times out of 50
The packaged rover shaken for 30 minutes	Every peripheral works afterwards

The twelve decisions

These are not facts to discover but choices, and the risk is not taking them. Each is written down, dated and signed, because an undecided scope is the mechanism behind the over-scope risk.

#	DECISION	RECOMMENDED CALL
D1	Which of the four new 2027 capabilities the team attempts, and which it defers	Bioassay yes, because it is mandatory; route finding yes; astronaut interaction partial; key and lock last
D2	One configuration for all four missions, or modular configurations	One configuration if mass allows, which removes the 10-minute re-weigh
D3	Drone: fly it or leave it home	Flying only with a dedicated pilot and spotter pair by January; otherwise it stays home
D4	Autonomy command modality: device, sign, audio or gesture	Device and audio; gesture as an experiment that can fail safely
D5	Two independent life-detection methods, or one instrument with two channels	Two methods on one detector, so one survives the other's failure
D6	The freeze date and the known-good tag	2027-04-16; after it, fixes only
D7	The six C2 seats, the commander, and who never touches a control	The commander never drives; two operators per station
D8	Whether to field the bioassay if the answer to J1 requires two independent methods and mass is binding	Mass is deleted elsewhere to fit it, not the bioassay
D9	The autonomy abandonment minute and the split between sub-missions	20 minutes or fewer for the astronaut tasks, the remainder for route finding, with a hard abandon minute
D10	Whether the key and lock task is attempted at all	Cheap tasks banked first; the key attempted only if the measurements pass by week 13
D11	Which field sites the team uses, given that testing at the site is prohibited	Two legal sites by week 8
D12	Whether the team keeps and publishes its own failure log	Yes, from week 1, one paragraph per failure, with the mechanism

What an unanswered question costs

Nothing, if the meanwhile assumption is written first. The cost appears when a team argues the same question twice, or when two rows are built on opposite guesses. That is why the register records the assumption against the affected problem rows and why the letter asks for every answer to be mirrored onto the public page: one answer, filed once, used everywhere.

The five unpublished documents are the exceptions worth watching. A missing review rubric cannot be guessed, only proxied against the 2026 one and labelled as a proxy. A missing Delivery sequence cannot be rehearsed, only ordered by points per minute when it arrives. Both are re-read monthly from December.

Further reading

- 70 risks shows which rows each answer moves.
- 110 engineering problems carries the rows these questions gate.
- The first 90 days places the letter and the measurements on a calendar.
- The Q&A page collects the rulings that already bind every team.
- 155 engineering requirements is where a question usually begins, because a requirement with no clause behind it is an open question in disguise.

THE BUILD PLAN · CHAPTER 7.2

The first 90 days

Ninety days is the measurement wave, the PDR and most of the architecture wave. The sequence below is the problem register's dependency order read forwards, so it is a work sequence rather than a schedule with names attached; ownership belongs to the team. Four gates sit inside the window, and four early warnings mark a season that is already slipping.

90

DAYS

From 2026-10-26 to the end of January 2027

2026-2027

4

GATES INSIDE THE WINDOW

Mass, task geometry, link and payload, safety

2026-2027

5

WEEK OF THE PDR

The register counts weeks from 0, so this is week 6 elsewhere

2026-12-02

6

LOAD-BEARING ROWS STARTED

Six of the thirteen high-dependency rows are measurements

2026-2027

Key takeaways

- **The first four weeks replace the numbers the rulebook does not publish.** They are bench and field measurements, and every later wave spends their results.
- **The PDR falls five weeks into the season**, at 2026-12-02, so the measurements have to start on day one rather than after the team is organised.
- **Four gates decide whether the rest of the register is executable:** a weighed mass allocation, a task panel with measured force curves, a measured link with terrain in the way, and a payload pair with a proven kill chain.
- **If those four are false at the end of January, no amount of design work recovers the season.** Each has a fallback, and the fallback is accepted in writing.

The sequence

The register numbers its rows from week 0, which is why its week 5 is the sixth numbered week in the risk calendar later in this part. The dates are the same dates. The PDR is due **2026-12-02** S-URC27-RULES §3.a.ii, and the SAR is due **2027-03-05** S-URC27-RULES §3.a.iii, forty-one days after the window closes.

WEEKS	DATES	WHAT GETS DONE	REGISTER ROWS	WHAT IT DECIDES
0–1	2026-10-26 to 11-08	Weigh every part; cut the timber 1.2 m gauge; start the judge-like task panel; breadboard the kill chain; write and send the question letter	W0-01, W0-02, W0-17, and the twenty questions	The mass frame and the task geometry become measured facts, and the questions start their clock
2–3	11-09 to 11-22	Force and torque campaign on the panel; keyboard and hammer trials; soil and drilling rig; traction and step rigs; the 1 km link test; RTK survey-in trial; tag range sweep; enclosure lamp test	W0-03 to W0-11, W0-13 to W0-15	Every unpublished force and range in the arm, mobility and communications designs is replaced
4	11-23 to 11-29	Close the measurements; fix the arm mass class; fix the stow concept; choose the assay pair in working form; choose the E-stop and distribution architecture; time the C2 station	W1-01, W1-06, W1-11, W1-13, W0-19, W0-20	The four PDR claims: mass, stow, payload and safety
5	11-30 to 12-06	PDR submitted on 2026-12-02; freeze the numbers that tooling depends on	W1-26 and the PDR gate set	The design is credible enough to keep building
6–9	12-07 to 2027-01-03	Mobility and drivetrain: four or six wheels, wheel diameter, torque, voltage, chassis; arm architecture and backlash class; power architecture	W1-02 to W1-05, W1-07 to W1-09, W1-12, W1-21	Tooling can start and long-lead parts can be ordered against a frozen number
10–13	2027-01-04 to the end of January	Communications, compute and mast; autonomy stack; GNSS base strategy; modularity plan; astronaut ladder; C2 architecture; the first interface documents	W1-14 to W1-20, W1-22 to W1-25, W2-01 to W2-05	The detailed-design wave begins with numbers already frozen rather than with assumptions

The same window as a dated calendar with one decision per week is on the 13-week plan page. This page gives the order of the work; that page gives the decision dates and the owner of each row.

The four gates inside 90 days

1. **Mass.** A weighed allocation at 48.5 kg or less with growth allowed for, and a concept that passes the cube gauge. The row is W0-01 and it unlocks nine later problems.
2. **Task geometry.** A task panel with measured force curves for all nine Equipment sub-tasks, built before any arm hardware is bought. The row is W0-02 and W0-03.
3. **The link.** A link measured at 1 km with terrain in the way, RTK accuracy measured, and tag ranges measured in full sun. The rows are W0-11, W0-13 and W0-14.
4. **Payload and safety.** An assay pair chosen with bench data and a written justification, and an E-stop circuit proven on the bench with a current probe. The rows are W0-08, W1-11 and W0-17.

A gate is not a meeting. Each is a number with an artefact behind it, and each is closed by a document that the review package can later quote. These four are technical gates; the four management gates G1 to G4, which track ownership and decision dates, are on the 13-week plan page.

Four early warnings

These are the observable signs that the schedule is already behind, and each has the same remedy: run the measurement, or accept the fallback in writing and move on.

- **The task panel is not built by the end of week 2.** Every arm and end-effector decision is now being taken on assumed forces, and the arm's open questions are all open at once.
- **There is no measured link by the end of week 3.** The radio, the camera count and the mast are being chosen on a spreadsheet model, and the communications package becomes a discovery exercise in April.
- **The assay pair is not chosen by the PDR.** The review video cannot show a mandatory capability, and the science payload is re-architected after the review, which is the most disruptive change this project can make.
- **The question letter has no answer by the end of November.** Three of the four December decisions are then taken on the fallback, which is the conservative choice and may cost capability.

Four commitments the 90 days must not break

- **The arm is not bought before the mass class is fixed**, because the arm class consumes the drivetrain margin and that coupling is invisible until the motor class changes.
- **The chassis is not built before the wheel diameter is frozen**, because that one number reopens seven to nine interfaces.
- **Radios are not chosen before the link test, and cameras are not chosen before the tag sweep.** Both are measurements, and both are cheap in week 3 and expensive in April.
- **The question letter is not deferred.** Answers take weeks, and three of the four December decisions depend on them.

What the window produces

By the end of January the team holds a weighed mass frame, a task panel with force curves, soil and traction data, a measured link and tag range, a lamp-tested enclosure, a proven kill chain, a measured pack sag, a timed station, two analogue sites and a recording discipline that makes every later measurement usable as review evidence. That list is the deliverable, and it is worth more than a design that is one revision further along.

Further reading

- 110 engineering problems is the full register these rows come from.
- The 13-week plan dates each decision and names the four gates G1 to G4.
- 77 open questions is the question letter and the meanwhile assumptions.
- 70 risks shows what a missed gate costs later.
- Coupling, conflicts and twelve trade studies supplies the decisions the architecture weeks freeze.

THE BUILD PLAN · CHAPTER 8.3

The 13-week plan

The thirteen weeks run from 2026-10-26 to 2027-01-23, one week per row below. Registration closes on 2026-10-30, the PDR is due **2026-12-02** in week 6 S-URC27-RULES §3.a.ii, and the SAR is due **2027-03-05**, **41 days after the window closes** S-URC27-RULES §3.a.iii. That last number is the design constraint on the whole window: the thirteen weeks have to end with something filmable rather than with a design review.

13
WEEKS
One decision per week, four gates, one owner per row
2026-2027

4
GATES
Registered, designable, irreversible choices made, filmable
2026-2027

2026-12-02
PDR
In week 6 of this calendar
2026

41
DAYS TO THE SAR
The review is due 2027-03-05
2027

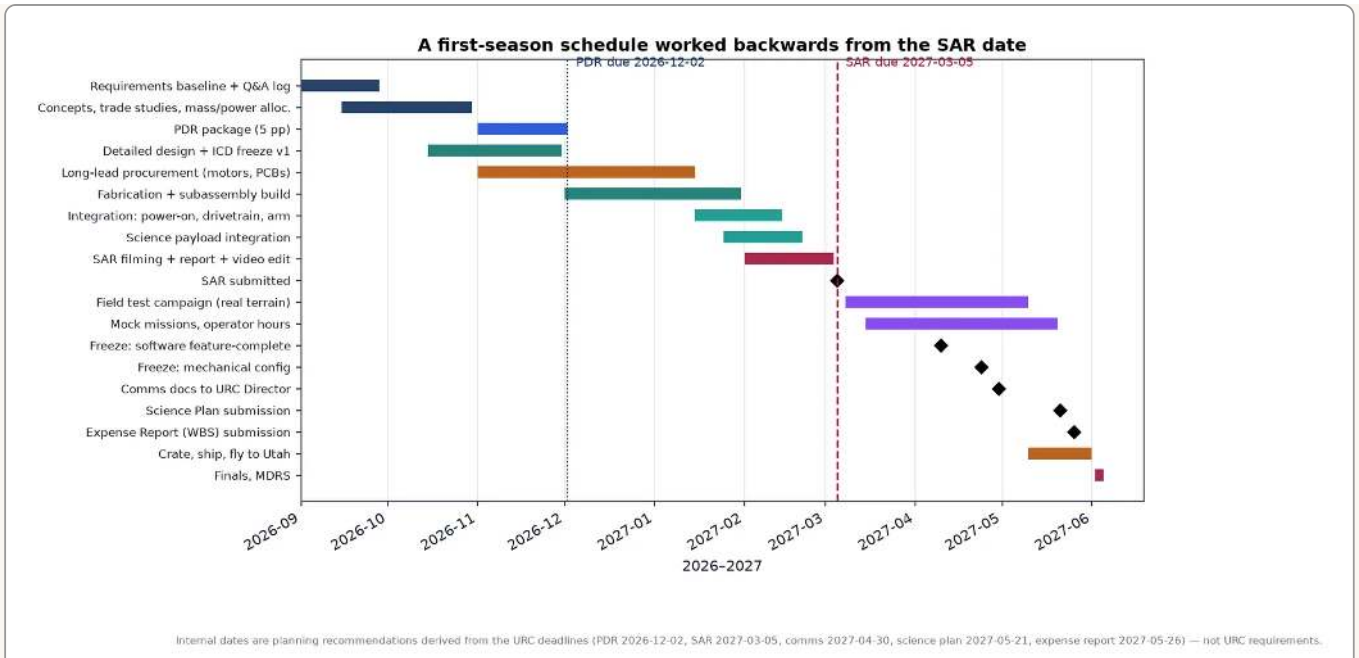
Key takeaways

- **Dates come before engineering.** Registration and the review calendar are closed in week 1, before any hardware is ordered.
- **Week 2 is task hardware, week 3 is the weigh-in, week 4 is the link.** Those three weeks decide the arm, the mass and the radios, and all three are cheap in November and expensive in February.
- **The irreversible science decision is taken before the winter break**, in weeks 6 to 8, so the fluidics has something to be designed around in January.
- **Week 13 ends with footage, not with a drawing.** If the integrated rover has not run on the clock by then, the review becomes documents only and the risk of missing the finals rises materially.

The four ordering rules

- **Dates before engineering.** Registration, the expense template and one named owner for every row come first, because an unowned calendar is itself one of the risks.
- **Measurements aimed at more than one subsystem.** The task-hardware mock-ups close six open questions in the requirements and four open arm problems with cardboard and a load cell, so they are worth more than any single subsystem's design work in the same week.
- **The two highest-uncertainty subsystems are prototyped before the architecture freezes:** the link behind a hill, and arm precision.
- **The irreversible decision, the chemistry, is taken before the winter break**, because it is the longest-lead science item and the one the review video depends on.

Where the thirteen weeks sit in the season



The season built backwards from the 2027-03-05 review deadline. The thirteen weeks on this page are the left-hand block: they end forty-one days before the review, which is why the last week is dedicated to footage rather than to design.

Robotics Club @ Mississippi State University, 2026. Original analysis.

S-URC27-RULES *Full resolution*

PDR submission **35 days** 2026-12-02

SAR submission **128 days** 2027-03-05

The thirteen weeks

WEEK	DATES	WHAT IS BUILT OR MEASURED	REGISTER ROWS	DECISION IT UNLOCKS
1	10-26 to 11-01	Registration on Friday 10-30; every lead and a calendar owner named; the public Q&A page captured; the sixteen judge questions sent as one email; the expense template agreed	R-R13 closed, R-M05, owners for R-R09 and R-R10	Every other row has a name attached
2	11-02 to 11-08	Task-hardware mock-ups built and measured: keyboard rig, tubular lock, cam-lock with a half gasket, cache, drawer, latch and valve panel, rock pick	R-T07, a reference for R-T08, and questions A1 to A5	The arm torque and force budget, and the end-effector architecture
3	11-09 to 11-15	The 1.2 m gauge and the first weigh-in; a mass register per configuration and in aggregate, including drone, suit and repeaters; a 45 kg allocation; the PDR started	R-R03, R-R04 and R-R05 made visible; question A22	Mass and configuration strategy, decision D2
4	11-16 to 11-22	Communications bench and first field day behind a hill: sub-band change, occupied bandwidth, 25 m of cable and the mast, signal logging, the setup-time spectrum scan	R-T01, R-R06, tests C1 and C2	The communications architecture: band, recovery plane, relay or not
5	11-23 to 11-29	PDR assembly: architecture version 1, a schedule with an integration bar, this register, the test plan; every item with a lead time over four weeks ordered	R-S02, R-S04 and the integration row R-S06	Architecture version 1 and the long-lead order list
6	11-30 to 12-06	**PDR submitted on Wednesday 12-02.** The December gate opens: two life-detection methods shortlisted, one bioassay named with its biomolecule and its positive and negative controls	R-S02 closed, R-T10 approached	The payload concept, decision D5
7	12-07 to 12-13	Chemistry decided in writing; the hazardous-chemical plan and the SAR science page written from one table; chemicals ordered	R-S03, R-R08	Payload frozen enough to design the fluidics
8	12-14 to 12-20	The deferral list published with the scoring cost of each deferral; two legal field sites identified; bioassay bench work: spike and recover, blanks, stability at 38 °C	R-M04, R-S07	Scope, and where the team tests
9	12-21 to 12-27	Arm bench week, portable over the winter break: backlash under load, 100 typing presses at ±3 mm, 50 key grasps, camera latency	R-T07 and R-T08 quantified	Whether open-loop fine work is possible, decision D10
10	12-28 to 01-03	RTK week, portable: survey-in at 30, 60 and 120 minutes, the 30-run stop distribution, and day-to-day base repeatability	R-T12, R-T13	The localisation architecture and the internal stop target
11	01-04 to 01-10	Autonomy smoke test: one planner, ten cold-boot runs, a forced link-down run, the state-light check in daylight	R-T19, R-012	Route finding alone, or with astronaut interaction: decisions D4 and D9
12	01-11 to 01-17	Integration demonstration and the first SAR footage: drive, arm, link and camera on the clock, with two operators in a mock station that has no view. The session is filmed	R-S01, R-S06, R-M02	Whether the review storyboard is achievable
13	01-18 to 01-24	Gate review: every row re-scored, the ninety-day report signed, the storyboard fixed against footage that exists, the freeze date published	every row re-scored	The build plan, taken backwards from the review

Two rows carry more weight than the rest. Week 3 makes the mass a measured quantity, and week 4 makes the link a measured quantity; every decision after them is cheaper because of it. Week 12 is the only week whose output is footage, and its result decides whether the review can be assembled from evidence or only from drawings.

The four gates

GATE	CLOSES AT	PASS CONDITION
G1	End of week 1	Registered, and every row has an owner
G2	End of week 6	Designable: an architecture, a mass allocation and a measured force budget
G3	End of week 8	The irreversible choices are made: chemistry, the chemical plan, the deferral list and the field sites
G4	End of week 13	Filmable: an integrated rover on the clock and on film

If G4 fails, the review becomes documents only and the chance of missing the field entry rises materially. That is the one gate whose failure cannot be recovered by working harder in February, because the footage has to show a rover that already runs.

These four are management gates, tracking ownership and decision dates. The four technical gates, which close on a measured number and an artefact, are set out on the first 90 days page.

What is deliberately absent

The window is defined as much by what it does not attempt. There is no flight rover, no chassis fabrication, no second autonomy planner, no gesture recognition before its hardware test, no attempt at the key and lock task before the measurements, and no drone flight without a named pilot and spotter. Nothing in the thirteen weeks depends on an unpublished URC document.

That list is a scope decision, and it is the same decision the risk register asks for in week 8. A first build that tries to do these things and the measurements in the same window does neither.

What the window must hand the review

- Footage of an integrated rover doing something measurable, filmed in week 12.
- A measured force and torque table for every Equipment interface.
- A real weigh-in behind a mass table, with the heaviest configuration included.
- A repeatability **distribution**, not a best run, for the arm and for the 1 m stop.
- A link log with range, shadow point and sub-band evidence.
- A signed deferral list, and the answers to the judge questions on file.
- The register, re-scored, as the systems-engineering evidence the review looks for.

A note on week numbering

The problem register counts its rows from week 0, so the PDR appears in its week 5 and in week 6 of this calendar. Both refer to the six days ending 2026-12-06. The dates bind; the week labels differ only in where the count starts.

Further reading

- The first 90 days gives the same window as a work sequence ordered by dependency.
- 110 engineering problems is the source of every row referenced above.
- 70 risks describes what each gate protects against.
- 77 open questions is the question letter sent in week 1.
- System Acceptance Review explains what the review is scored on.

APPENDICES · A

Team dossier index

The book profiles 51 teams: eight Tier A, 24 Tier B and 19 Tier C. The profile is where a team's rover, its results and the evidence behind both are set out, and this page indexes all of them by university, country, years at the finals and source count. The tier rule and the 68-team census are on the census page.

51

TEAM PROFILES

8 Tier A, 24 Tier B, 19 Tier C

2022–2026

68

TEAMS IN THE CENSUS

49 carry a profile of their own

2022–2026

S-CENSUS

187

TEAM-YEARS OF RESULTS

Five official score sheets

2022–2026

S-URC-RESULTS

19

CENSUS TEAMS WITHOUT A PROFILE

They publish almost nothing; the absence is the finding

2022–2026

S-CENSUS

Key takeaways

- **The tier is a result, not a reputation.** It comes from the official scores of 2022 to 2026 plus a documentation test, and documentation can move a team one band but cannot buy a tier on its own.
- **Eight teams carry the design evidence.** The Tier A profiles draw on design reviews, public code and review films; most Tier C profiles rest on a website and a score row.
- **Two profiles sit outside the census.** Alberta and Waterloo registered each year without reaching the finals sheet and are profiled anyway, because their public engineering record is unusually good.

How the tiers are set

A team's tier is derived from its finals record over 2022 to 2026 and from a tested documentation flag, never from its reputation S-CENSUS . Tier A is a podium finish in any year, a top-six average across at least two finals, or an exceptional public record with a best finish inside six. Tier B is a consistent top-20 average, a best finish inside twelve, or a documented team with two finals. Tier C is everyone else, including any team that recorded a review score but no mission scores.

The documentation test is deliberately narrow. A full public report set, meaning design-review documents or an open documentation archive, counts for more than a project page, and a project page counts for more than a website or a university news item. That is why two teams with similar results can sit in different tiers: the tier answers how much evidence exists as well as how well the team finished.

All profiles

The table lists every profile once, in tier order, and each team name links to its page.

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
A	AGH Space Systems (AGH University of Krakow)	AGH University of Krakow	Poland	4	23
A	BYU Mars Rover (Brigham Young University)	Brigham Young University	United States	5	14
A	Michigan Mars Rover Team (MRover)	University of Michigan	United States	5	58
A	Missouri S&T Mars Rover Design Team (MRDT)	Missouri University of Science and Technology	United States	5	53
A	Monash Nova Rover	Monash University	Australia	5	41
A	Northeastern University Mars Rover Team ("NURover")	Northeastern University	United States	5	20
A	UIU Mars Rover Team	United International University	Bangladesh	5	32
A	West Virginia University: Team Mountaineers	West Virginia University	United States	5	68
B	A.S.T.R.A. (University of Alabama in Huntsville)	University of Alabama in Huntsville	United States	5	20
B	ARES Project (Tohoku / Keio, Japan)	Multi-institution (Tohoku, Keio,...)	International (multi-country)	3	15
B	BRACU Mongol-Tori	BRAC University	Bangladesh	5	28
B	DAM Robotics (Oregon State	Oregon State University	United States	2	14

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
	University)				
B	Husky Robotics (University of Washington)	University of Washington	United States	2	13
B	IITB Mars Rover Team (IIT Bombay)	Indian Institute of Technology Bombay	India	3	15
B	IMPULS (Kielce University of Technology)	Kielce University of Technology	Poland	3	16
B	KNR Rover Team (Warsaw University of Technology)	Warsaw University of Technology	Poland	2	14
B	Legendary Rover (Rzeszów University of Technology)	Rzeszow University of Technology	Poland	2	15
B	Mars Rover Manipal (Manipal Institute of Technology)	Manipal Institute of Technology	India	3	15
B	MAVRIC (Iowa State University)	Iowa State University	United States	2	15
B	MIST Mongol Barota	Military Institute of Science and Technology	Bangladesh	5	20
B	New Haven Robotics (University of New Haven)	University of New Haven	United States	3	11
B	PCz Rover Team (Czestochowa University of Technology)	Czestochowa University of Technology	Poland	2	15
B	Project Scorpio (Wrocław	Wroclaw University of Science and	Poland	3	16

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
	University of Science and Technology)	Technology			
B	QSET: Queen's Space Engineering Team (Queen's University)	Queen's University	Canada	5	13
B	Quantum Robotics (Instituto Tecnológico y de Estudios Superiores de Monterrey)	Instituto Tecnológico y de Estudios Superiores de Monterrey	Mexico	4	17
B	Robotics for Space Exploration (RSX), University of Toronto	University of Toronto	Canada	3	22
B	SC Robotics (Saddleback College / Irvine Valley College)	Saddleback College	United States	4	13
B	SKA Robotics (Warsaw University of Technology)	Warsaw University of Technology	Poland	2	14
B	Sooner Rover Team (University of Oklahoma)	University of Oklahoma	United States	4	15
B	Toronto MetRobotics (TMR)	Toronto Metropolitan University	Canada	4	15
B	Wisconsin Robotics	University of Wisconsin-Madison	United States	4	21
B	Yonder Dynamics (UC	University of California, San Diego	United States	5	11

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
	San Diego)				
C	Ad Astra (Sairam Institutions)	Sairam Institutions	India	2	10
C	Binghamton University Rover Team (SUNY Binghamton)	Binghamton University (SUNY)	United States	3	10
C	Carleton Planetary Robotics Team (CPRT) (Carleton University)	Carleton University	Canada	2	13
C	Ceres (Universidad de los Andes)	Universidad de los Andes	Colombia	1	10
C	Cornell Mars Rover (CMR)	Cornell University	United States	4	22
C	ITU Rover Team (Istanbul Technical University)	Istanbul Technical University	Turkey	4	14
C	KARURA (international: Texas A&M + Japanese universities)	Multi-institution (Texas A&M + Japanese universities)	International (multi-country)	3	13
C	Mars Rover UdeG Space (Universidad de Guadalajara)	Universidad de Guadalajara	Mexico	3	10
C	Project Altair (Islamic University of Technology)	Islamic University of Technology	Bangladesh	2	9
C	Robocol (Universidad de	Universidad de los Andes	Colombia	2	15

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
	los Andes, Colombia)				
C	Space Concordia Robotics Division (Concordia University)	Concordia University	Canada	2	10
C	Space Exploration Alberta Robotics (SPEAR) (University of Alberta)	University of Alberta	Canada	2022-2026 registrations; no finals	13
C	Team Anveshak (IIT Madras)	Indian Institute of Technology Madras	India	2	15
C	Team Interplanetar (BUET)	Bangladesh University of Engineering and Technology (BUET)	Bangladesh	2	10
C	Team RoSE: Robotic Space Exploration (University of Hawai‘i at Mānoa)	University of Hawaii at Manoa	United States	3	11
C	UBC Rover (University of British Columbia)	University of British Columbia	Canada	1	11
C	UMD Loop (University of Maryland, College Park)	University of Maryland, College Park	United States	1	11
C	UW Robotics / UWRT (University of Waterloo)	University of Waterloo	Canada	2022-2026 registrations; no finals	15

TIER	TEAM	UNIVERSITY	COUNTRY	YEARS	SOURCES
C	YILDIZ ROVER (Yıldız Technical University)	Yıldız Technical University	Turkey	5	13

How to read the columns

- **Years** is the number of editions from 2022 to 2026 in which the team recorded a finals result. For the two teams that registered each year without reaching the finals, the column gives the registration span instead.
- **Sources** is the number of distinct source documents the profile cites. It measures breadth of evidence, not quality of design, and a profile can be short and still be well sourced.
- **Country** and **University** follow the URC team-info pages for the years the team competed.

Teams without a profile

Nineteen of the 68 census teams have no page of their own. They appear in the census with their results, and most publish no engineering material beyond a website, so a profile would rest on a score row and a photograph. The absence is recorded rather than filled, because a thin page that reads like a full one would mislead a team choosing which machines to study.

The two teams profiled outside the census are the reverse case. Alberta and Waterloo never recorded a finals result between 2022 and 2026, but both publish enough engineering material to be worth a page, so they carry a Tier C profile and a note in the table that they reached registration only.

> **What a source count measures**
>
> The count is of distinct documents, not of claims. A single review film can support a dozen statements about a suspension, and a team with a public repository will always out-count a team with a good rover and a quiet website. The source index explains how those documents are distributed.

APPENDICES · C

Source index

The bibliography holds 1,276 entries: 1,251 rows in the main registry and 25 team records cited by the profiles that are not in the merged table. This page shows how they are distributed by group and what each group contributes. The full list, with every identifier, date and original link, is the bibliography.

1,276
BIBLIOGRAPHY ENTRIES
Every identifier used in the book
2027 edition

1,251
ROWS IN THE MAIN REGISTRY
The stable source list
2027 edition

1,009
TEAM MATERIAL
Team sites, reports, code and films
2022–2026

93
OFFICIAL URC DOCUMENTS
Rulebooks, Q&A snapshots, score sheets and pages
2007–2027

Key takeaways

- **Four entries in five come from the teams themselves.** The 1,009 team sources supply the architecture evidence; 93 URC documents supply every limit, ruling, deadline and score.
- **The official record is small and decisive.** One current rulebook, one Q&A snapshot and five score sheets settle what the competition asks and what the field scored, and all three are also kept as archived copies.
- **A source can be cited without endorsing its conclusion.** The registry records vendor datasheets, archived pages and documents that were searched for and not found, because a blank result is evidence about the record.

By group

Every entry is filed in one of five groups. The team material dominates by count because a profile needs a team's own site, report, repository and review film to describe a rover.

GROUP	COUNT
Team material	1,009
News and university	101
Official URC	93
Vendor documentation	63
Archived and other	10

What each group contributes

GROUP	WHAT IT SUPPLIES	WHERE IT CARRIES THE ARGUMENT
Team material	Team sites, design-review documents, papers, code and review films	Team profiles and every subsystem comparison
News and university	University releases, sponsor case studies and competition write-ups	Facts a team does not publish itself, such as a mass in a news item
Official URC	Rulebooks, Q&A snapshots, official score sheets and team-info pages	Rules, deadlines, rulings, the census and every score statistic
Vendor documentation	Datasheets and product pages	Motor, pack, sensor and radio specifications, with list prices where published
Archived and other	Saved pages and archived snapshots	Documents that change or disappear, kept as the version this book cites

The groups are not equal in weight. A single official document can settle a rule that a hundred team pages only illustrate, and a team's review film can be the only public account of a suspension that worked.

How a citation resolves

An identifier such as `S-URC27-RULES` is stable and appears in the same form everywhere it is used. In the text it is a small link that opens its row in the bibliography, which gives the title, author or publisher, date, group and a link to the original document where one exists. Each page also ends with a list of the sources it used, so the evidence for one page can be read without leaving it.

Two related reference systems work the same way. A rule or Q&A reference such as `§2.a.ii` or RQ11 links to the rules page or the Q&A rulings, and a requirement identifier links to the requirements table. The primary documents themselves, from the 2007 rulebook to the 2027 Q&A snapshot, are kept on the primary documents page.

> **What the counts do not measure**
>
> A source is counted once, by its registry row. A single team page can support many claims, so a team with one excellent review film can outrank a team with a long but shallow site on evidence, and the count alone will not show it. A cited source is not an endorsement either: the registry records what was available, including documents that partly contradict the analysis. The bibliography is the authoritative list, and this page is only its shape.

APPENDICES · D

Evidence coverage and known limits

The analysis can only be as good as what URC and the teams publish, and that record is uneven. This appendix states, topic by topic, where the evidence is strong enough to design against and where it is thin enough that a measurement has to replace a citation. It also sets out what the 2027 rules leave unpublished, because those omissions decide the order of the team's own test work.

162

FINALS TEAM-YEARS

Five official score sheets, 2022–2026

2022–2026

S-URC-RESULTS

68

TEAMS IN THE CENSUS

49 carry a profile, 19 have no page

2022–2026

S-CENSUS

17

ROVER CONFIGURATIONS WITH A MASS

13 published, one weighed by a judge

2021–2026

77

OPEN QUESTIONS

16 to URC, 22 measurable, 27 need hardware, 12 team decisions

2027

Key takeaways

- **The rule and result record is close to complete.** The rulebooks from 2007 to 2027, the 2022–2026 score sheets and the 2027 Q&A snapshot are kept as archived copies, so a limit, a score or a ruling can be checked against the document that stated it.
- **Team architecture is deep at the top and shallow below it.** The eight Tier A profiles draw on design reviews, public code and review films; most Tier C profiles rest on a website and a result.
- **Mass, centre of gravity and energy are the weakest numbers.** Fourteen of 187 team-years publish a total mass, one mass was read from a judge's scale, one team publishes a subsystem breakdown, and no team publishes a measured watt-hour figure for a mission or a centre-of-gravity table.
- **Three parts of the 2027 record do not exist yet:** task weights for Astrobiology, Delivery and Equipment Servicing, the SAR and Science Plan requirements, and any published mass, dimension or force for the objects a servicing task uses.

What the record covers well

The rules are the strongest part of the record. The 2027 rulebook was published in September 2026 and the Q&A was current to 11 October 2026, and both can be read against the archived copies S-URC27-RULES S-URC27-QA . Every hard limit, mission task and deadline in this book traces to a clause, and the 155 requirements give each trace a line with the check that confirms it S-URC27-REQ . The Q&A is a living document that changed twice in its first month, so a ruling carries the date it was read, and the same questions have to be re-read before the PDR on 2 December 2026 and again before the SAR on 5 March 2027 S-URC27-QA S-URC27-RULES §3.a.ii–iii.

The results are nearly complete. The five official score sheets for 2022 to 2026 hold 187 team-years, 162 of them with finals scores, and every mean, zero rate and variance share in the statistics part is computed from them rather than estimated S-URC-RESULTS S-CENSUS . Two limits survive even here: a rank is not a column in a score sheet but a derivation from row order, and a team with a review score and no mission scores has no rank at all.

Team architecture is the third strong area, and its depth depends on the team. The 51 profiles are built from published reports, repositories, review films and papers, and the Tier A pages can follow a suspension or an arm down to a part S-CENSUS . The tier is a documentation test as well as a form guide: Tier A holds eight teams with a podium finish or an exceptional public record, Tier B holds 24 with a consistent finals presence, and Tier C holds the remaining 19 profiles, most of which rest on a website and a result.

Where the evidence thins out

TOPIC	WHAT THE RECORD LACKS	HOW THE BOOK HANDLES IT
Mass and centre of gravity	14 of 187 team-years publish a total mass; one mass was judge-weighed; no centre-of-gravity table exists	A 17-configuration model, labelled as a model
Energy in use	No team publishes a measured watt-hour figure for a mission	A bottom-up load model with its duty cycle stated
Task objects and forces	No published mass, dimension or force for a lock, cache, valve, connector, button, hammer or suit	Design targets taken from comparable hardware, labelled as targets
Course geometry	No published slope, step or drop limit, and no figures when the Q&A was asked	A self-chosen obstacle set, measured against a local field
Architecture below Tier B	Most Tier C teams publish a website and a film rather than a report	Short profiles that state the result and the absence
2027 scoring detail	Task weights unpublished for three missions; SAR and Science Plan not posted	The 2026 rubric as a labelled proxy
Delivery and drones	No comparable published account and no delivery-specific end-effector comparison	A synthesis, with the scope limit stated on its own page

Mass, centre of gravity and energy

Mass is the number that most constrains a first build and the number the field publishes least. The coverage counts are the finding.

QUANTITY	PUBLISHED IN THE FIELD	WHAT THE ANALYSIS USES INSTEAD
Total mass of a competition rover	14 of 187 team-years (7.5 %): UIU 5, BRACU 2, WVU 2, Northeastern 2, Utah State 1, MIST 1, Toronto Met 1	17 configurations from eight teams, 2021–2026 S-CENSUS
Mass read from a judge's scale	1: MIST Phoenix, virtual URC 2021, 42.1 kg S-MIST-VURC2021	The other 16 rows are team-published, arithmetic from a stated reduction, or a team or university statement
Subsystem mass breakdown	1: Monash Banksia 2026, rolling chassis and arm S-MONASH-GH-NOVA	One CAD set plus two allocations, each labelled
Battery chemistry, voltage or capacity	4 teams: UIU, WVU, Missouri S&T and Monash	Pack sizing from published cells and duty-cycle arithmetic
Measured energy for a mission	0 teams	A bottom-up load model, table F.6
Centre-of-gravity table	0 teams	Tip angles estimated, with the basis stated
Champions 2022–2026 with a published mass	1 of 5: WVU 2023, 33.0 kg S-WVU-SAR2023	The lightest known winning rover, not the typical one

Across the 17 configurations the minimum is 33.0 kg, the maximum 53.0 kg, the mean 47.9 kg and the median 49.1 kg, and four of them sit within 0.2 kg of the limit, one at 49.8 kg and three at exactly 50.0 kg. Those numbers size a mass budget well. They do not predict a weigh-in, because the one measurement in the record comes from a 2021 virtual event, and the rest are the teams' own statements or deductions from them.

Task objects, forces and the course

The servicing tasks are described in words and scored on completion, but URC publishes no mass, dimension or force for the objects involved: the sample tube, the cache, the key and lock, the hose fitting, the valve, the buttons, switches and knobs, the hammer, or the astronaut's suit S-URC27-RULES S-URC27-QA §1.d.i–ii. A gripper torque budget, a lift rating and a typing force therefore rest on a mock that the team builds, not on a published value, and every such target in this book says so.

The course is described by terrain type rather than by numbers. The Q&A gives no design value for slope or drop: inclinations vary by route and are not provided S-URC27-QA S-URC27-QA AUQ17 D1. Sign-reading distance and camera resolution are left to

performance rather than stated as distances. The response in this book is a self-chosen obstacle set, measured against a local rock pile and the previous season's course film, and the result is a design target rather than a rule.

The 2027 rubrics and weights

Three 2027 documents are promised and not yet available. The SAR page says the 2027 requirements will be posted in early 2027 S-URC27-SAR , the Science Plan page says the same S-URC27-SCIENCEPLAN , and the Delivery course description is promised before the competition rather than with the rulebook S-URC27-RULES §1.c.i.

The weights are the second half of the problem. Each mission carries 100 points, but per-task values are published only for the Autonomy Mission S-URC27-RULES §1.a.i §1.e.iv–xviii. Astrobiology names four components without weights S-URC27-RULES §1.b.xi, and Equipment Servicing scores each task without publishing its value S-URC27-RULES §1.d.i–ii. A comparison between two sub-tasks in those three missions is therefore an ordering, not a price.

Where the book needs a price it uses the 2026 rubric, which is the most recent published one, and labels it as a proxy S-URC26-SAR . That rubric is a floor rather than a guarantee: 2027 adds an astronaut-assistance autonomy mission, a fine-manipulation key task and a display interface, and none of them is priced in the 2026 document.

Delivery and drones: one recorded scope limit

Delivery is the only subsystem without a dedicated study of its own, and the reason lies in the record rather than the mission. No comparable published account of the 2027 Delivery course, its stage tasks or its drone use exists, so the chapter is a synthesis from the Delivery rules S-URC27-RULES §1.c, the drone rules S-URC27-RULES §2.b.v–vii, the Delivery Q&A and the arm, astronaut and communications material. Container and stowage detail is thinner than in the other eleven chapters, and there is no delivery-specific end-effector comparison table.

Two rule questions in this area remain open and each one changes a design: whether a drone counts toward the 70 kg fielded total, and whether the Equipment Servicing cache is the team's own astrobiology cache S-URC27-RULES §2.b.vii. Neither is settled by the rulebook, and both are worth a written answer from the judges before the preliminary design review.

From a gap to a measurement

A gap is only useful if it names the test that closes it. The 77 open questions collect the ones that change a design decision, grouped by what would settle each.

GROUP	QUESTIONS	WHAT CLOSES IT
Ask URC	16	A written answer from the URC Director, quoted with its date
Measure it	22	A bench measurement of a part or a mock
Test on hardware	27	A test on the built rover or a prototype
Team decision	12	A choice the team records, with its reason

The 110 engineering problems carry the same gaps into the build plan, each with an acceptance criterion and a place in the dependency order. A handful of them close in a week and turn several rows on this page from a limit into a number: a weigh-in of the completed rover, one logged discharge under a measured load, a lock torque on a matching cylinder, and one climb test on a 30° slope of known soil.

What stays open after the measurement

- **Rank is a derivation.** A score sheet is sorted, not ranked, so a one-place difference between two teams is not a published fact.
- **A photograph is not a datasheet.** A motor, gearbox or camera identified from a review film is recorded as an inference, with the visual basis stated, until a document confirms the part.
- **A winning design does not transfer.** Where this book describes what worked for a champion, it describes that team's choice and the result it produced, not a requirement for anyone else.

- **The commanded astronaut interaction is new.** Follow, Stay, Fetch, Come and Give, with a beckoning gesture, have no predecessor task, so that part rests on the rules, the Q&A and first principles rather than on observed practice.

Further reading

- About the evidence sets out the four grades, the citation form and the way a number is read.
- Team dossier index lists the 51 profiles; the census holds all 68 teams and the tier rule.
- Mass and power data carries the 103 data rows behind the mass figures quoted here.
- Mass and power shows what the model concludes and where it is sensitive.
- 77 open questions turns each limit into a test, a measurement or a question for the judges.
- Delivery, payload handling and drones is the chapter with the recorded scope limit.

APPENDICES · F

Mass and power data

These are the seven tables behind the mass and power analysis: rover mass totals, component masses, battery packs, electrical loads, team scores, model outputs and coverage counts. A row reads left to right as team or quantity, year, value with its unit, the evidence basis and the source. The model that turns the rows into a budget is Mass and power.

17
ROVER CONFIGURATIONS
Thirteen with a published mass
2021–2026

103
DATA ROWS
Seven tables, F.1 to F.7
2027

14
TEAM-YEARS WITH A PUBLISHED MASS
7.5 % of the 187 team-years in the window
2022–2026

118
ROWS IN THE PRINTED APPENDIX
Including the table headers and separators
2027

Key takeaways

- **Every row carries its own evidence basis.** A mass is team-published, arithmetic from a stated reduction, a CAD sum, a university report, or the single judge-weighed figure in the record.
- **The mass rows cover 17 configurations from eight teams, and four of them sit within 0.2 kg of the limit**, one at 49.8 kg and three at exactly 50.0 kg.
- **Table F.7 sets the limit on the rest.** It counts what the field publishes, and those counts are the reason every mass and energy figure in this book is a model rather than a measurement of one rover.

The seven tables

TABLE	WHAT IT HOLDS	ROWS
F.1	Rover mass totals: 17 configurations from eight teams, 2021–2026	17
F.2	Partial and component masses: chassis, suspension, wheels, arm and pack items	19
F.3	Energy storage: chemistry, configuration, voltage, capacity and energy for each pack	22
F.4	Electrical loads: the datapoints used in the load model	12
F.5	Scores: the teams in the mass database, with their mission scores and totals	12
F.6	Model outputs: power, energy and nameplate for a 35 kg and a 50 kg rover	12
F.7	Coverage statistics: what the field publishes, and how often	9
	Total	**103**

The full table

The table opens on one of the seven and switches between them with the selector at the top; the search box filters the visible table by team, year, item, basis or note.

#	TEAM	YEAR	ROVER	CONFIGURATION	MASS (KG)	EVIDENCE BASIS	SOURCE ID	BASIS / N
1	UIU Mars Rover	2022	Maven	competition	49.8	team-published	S-UIU-ROVER-PAGES-EARLIER-GE-03	team rov page spe table
2	UIU Mars Rover	2023	Telos	competition	47.8	team-published	S-UIU-ROVER-PAGES-EARLIER-GE-02	team rov page spe table
3	UIU Mars Rover	2024	Yggdrasil	competition	49.2	team-published	S-UIU-ROVER-PAGES-EARLIER-GE	team rov page spe table
4	UIU Mars Rover	2025	Axios	competition	48.5	team-published	S-UIU-AXIOS-2025	team rov page spe table
5	UIU Mars Rover	2026	Aurion	competition	47.5	team-published	S-UIU-ROVER-PAGE-AURION	team rov page spe table; pa also stat 15 kg pay capacity
6	BRACU Mongol-Tori	2026	Taurus	competition	50.0	team-published	S-BRACU-ROVER-PAGE-TAURUS-2026	team spe table; ex at the 50 limit
7	WVU Team Mountaineers	2023	Wanderer	competition incl. 5-DoF arm	33.0	team-published	S-WVU-SAR2023	SAR p.1 verbatim “With this arm equipped Wanderer weighs 3 resulting 20 kg reduction from its predeces
8	WVU Team Mountaineers	2022	predecessor of Wanderer	competition	≈53.0	estimated from the team’s stated reduction	S-WVU-SAR2023	33 kg + stated 20 reduction a measu statemer
9	Northeastern University	2026	unnamed 2026 rover	competition	≈46.0	reported by the university	S-NEU-NEWS-2026	universit news: 10 “10 % lig than 202
10	Northeastern University	2025	unnamed 2025 rover	competition	≈51.0	estimated from the later figure and the team’s stated reduction	S-NEU-NEWS-2026	46 kg / 0.
11	Monash Nova Rover	2021	Wombat	ARC competition (no URC in 2021)	49.1	team-published	S-MONASH-ROVERS	team rov page
12	Utah State University	2025	Sophia	science configuration	46.9	team-published	S-USU-2025-POSTER	final-repo poster, measure
13	Utah State University	2025	Sophia	arm configuration	51.8	team-published	S-USU-2025-POSTER	final-repo poster, measure 1.8 kg ov the limit
14	BRACU Mongol-Tori	2025	Hypersonic	competition	48.0	team-published	S-BRACU-SITE-HYPERSONIC	team rov page
15	Toronto MetRobotics	2026	Ingenuity	competition	50.0	team-published	S-TMR-METROBOTS	team spe page; no URC 202 score (di compete
16	MIST Mongol Barota	2021	Phoenix	competition, virtual URC 2021	42.1	judge-weighed at the event	S-MIST-VURC2021	the only judge-weighed mass in t

#	TEAM	YEAR	ROVER	CONFIGURATION	MASS (KG)	EVIDENCE BASIS	SOURCE ID	BASIS / N
								record: official U stream, 1 the team own card says 45 k
17	MIST Mongol Barota	2025	Aurora X	competition	50.0	reported by the team	S-MIST-SITE-ROVERS	team rov card

What the mass rows show

Across the 17 configurations the minimum is 33.0 kg, the maximum 53.0 kg, the mean 47.9 kg and the median 49.1 kg. Four configurations sit within 0.2 kg of the limit: UIU's 2022 Maven at 49.8 kg, and three rovers listed at exactly 50.0 kg, BRACU's Taurus in 2026, Toronto Met's Ingenuity in 2026 and MIST's Aurora X in 2025. That pile-up under the line is a design choice rather than a coincidence: the limit is 50 kg per deployed configuration, with a penalty of 5 % of that event's points per kilogram over it S-URC27-RULES §2.a.ii–iii. The lightest rover in the set, WVU's 33.0 kg Wanderer in 2023, won the competition with 425.35 points S-WVU-SAR2023 S-URC-RESULTS.

The distribution is a record of teams that chose to publish. Fourteen of 187 team-years in the 2022–2026 window carry a total mass, which is 7.5 % of the field, and exactly one mass in the whole record was read from a judge's scale: MIST Phoenix, 42.1 kg, at the virtual event in 2021 S-MIST-VURC2021. The rest are the teams' own statements, arithmetic from a stated change, or a sum of CAD link masses.

Reading an evidence basis

BASIS	WHAT IT MEANS	HOW TO USE IT
Team-published	The team's own page, report or specification table	Good enough to size a budget; not a weigh-in
Judge-weighed	Measured on a scale at the event, with the source in the row	The only measurement class in the table
Arithmetic from a stated reduction	Derived from a later or earlier figure and a percentage the team gave	Treat as an estimate with a stated chain
CAD sum	Added from the link masses in a published model	A design intent, not a built mass; open links understate it
Reported by a university	A news release or similar secondary account	Context only until a team document confirms it

The F.2 rows show why the basis matters. The Monash Banksia chassis, suspension, ankles, wheels and differential bar add to 24.42 kg from the published robot model, while the arm in a separate model file sums to 6.54 kg with several links at near zero, so the real arm is likely heavier S-MONASH-GH-NOVA. A CAD total is a floor on the mass of the parts it describes.

What the coverage rows say

Table F.7 is short and it constrains everything else on this page. One team publishes a subsystem breakdown, four teams publish battery chemistry, voltage or capacity, no team publishes a measured energy figure for a mission, and no team publishes a centre-of-gravity table. One of the five champions from 2022 to 2026 has a published mass.

Those limits are set out topic by topic on Evidence coverage and known limits. The practical consequence is that the power figures in table F.6 are model outputs: the 249 W average and 3.7 kW peak of the analysis come from the duty cycle and the load lines, not from a logged discharge, and the pack sizes of 273 to 289 Wh follow from that model.

> **How to read the mass figures**
>
> Masses are the teams' own published figures unless the basis column says otherwise. A row marked as arithmetic is derived from a stated reduction, not weighed. The only judge-weighed mass in the record is MIST Phoenix at 42.1 kg. The 50 kg limit applies to each deployed configuration, not to the total of all fielded parts, which is capped separately at 70 kg.

Download and reuse

The seven tables are available as one CSV, which is the same data the table above displays: mass and power data. The official score sheets remain the authority for any result, and the source identifier in each row resolves in the bibliography.

The data may be reused with attribution to the Robotics Club at Mississippi State University; the terms are on Credits and reuse.

Further reading

- Mass and power is the analysis these rows support, with the 50 kg budget and the pack arithmetic.
- Mass states the headline mass statistics with their sample.
- Power and energy states the draw and energy figures with their duty cycle.
- Evidence coverage and known limits explains why the field publishes so little.

About the evidence

This is a reference book, not a promotional one. Its usefulness depends on a reader being able to see where a claim came from, how strong it is and what it does not cover. This page explains the evidence grades, the citation system, the way numbers are presented, and the known limits of the record.

1,251
REGISTERED SOURCES
Each with a title, publisher, date and archived copy
2026

307
MEDIA FILES
Every figure resolves to an author, a year and a source
2026

155
TRACED REQUIREMENTS
Each tied to a rule or a Q&A clause
2027
S-URC27-REQ

162
FINALS TEAM-YEARS
The sample behind the score statistics
2022-2026

What the book is

The book analyses the 2027 University Rover Challenge for a university team building a Mars rover for the first time. It covers the competition rules and scoring, the teams and their machines, twelve subsystems, the cross-cutting problems of mass, power, interfaces and failure, the 2027 mission backfit, and a build plan. The print edition is the same text in a fixed A4 layout; this edition adds the full source list, the asset record and the interactive tables.

The unit of evidence is a source, not a paragraph. Every non-trivial factual claim carries a source identifier. Where the record does not contain an answer, the text says so rather than filling the space with a guess.

Evidence grades

Four grades appear in the deeper pages and in the data tables. They are visible as small pills.

GRADE	WHAT IT MEANS	HOW IT IS USED
PRIMARY	Stated in a primary source: the team's own report, site, code or video, an official URC document, official results, or a manufacturer's datasheet.	The strongest grade. It still depends on the source being accurate and current.
REPORTED	A credible secondary source: a news article, a university release, a sponsor case study or a competition write-up.	Useful for context and for facts a team does not publish itself. It is one step removed from the team.
INFERENCE	A reasoned deduction from images, video, partial data or comparable systems.	The text says what the inference rests on. An inference is not a measurement.
NOT PUBLISHED	No public figure exists, and the text says where it was looked for.	The honest answer. It becomes a measurement task or an open question, not a silent omission.

A statement without a grade is either a direct quotation, a description of the book's own analysis, or a definition. A grade is not a score: a well-sourced claim can still be wrong if the source is wrong, and an inference can be more useful than a weak primary figure.

How a claim is worded

The reading text rarely prints a grade beside a sentence. The wording carries it instead, and the wording is deliberate.

READING	WHAT IT TELLS A READER
A source identifier with nothing beside it	A primary source states it: a team document, an official URC document, a datasheet or an official score sheet
“the team reports”, “according to the team”	A credible secondary account, one step removed from the team’s own record
“estimated”, “by inference”, “target derived rather than published”	The book’s own reasoning from stated parameters, images or comparable systems
“not published”, “not stated in the rules”, “unmeasured”	The record has no answer, and the text says where it was looked for

A value taken directly from a source always carries that source’s identifier. A value that the book derived is stated with its inputs, so the arithmetic can be repeated and argued with.

How sources are cited

Every source has a stable identifier of the form `S-TOPIC-SHORT`, for example `S-URC27-RULES` for the 2027 rulebook or `S-MST-2026-SAR` for Missouri S&T’s 2026 review film. The identifier is the same everywhere it appears, so a claim can be traced from a page, to the bibliography, to the original document.

In the text, an identifier is a small link. It opens the entry in the bibliography, which gives the title, author or publisher, date and an original link where one exists. A page also ends with a list of the sources it used, so the evidence for a single page can be read without leaving it.

Rule and Q&A references work the same way. A clause such as `<a class="clause" href="../competition/rules.html#clause-2a">§2.a.ii</a>` links to the rules page; a Q&A identifier such as `RQ11` links to the Q&A page. Requirement identifiers such as `<a class="ref" href="../competition/requirements.html#req-req-mob-004">REQ-MOB-004</a>` link to the requirements table.

The bibliography is not a reading list. It includes vendor datasheets, team code, archived pages and searches that found nothing. A source can be cited because it is the best available evidence for a fact, not because it endorses the conclusion.

How to read the numbers

Five questions settle most disagreements about a number:

1. **What is the sample?** A score statistic may cover 162 finals team-years from 2022 to 2026, 36 teams in one year, or 17 rover configurations from eight teams. The page states the sample next to the number.
2. **What is the period?** Rules, scores and designs change. A 2024 figure is labelled 2024, and a claim about 2027 is labelled 2027.
3. **Who owns the value?** A number belongs to the analysis that derives it, not to the page that quotes it. The key numbers page gives the owner and the source for each headline value.
4. **Is it measured, published or inferred?** A team-published mass is not a judge-weighed mass. A CAD sum is not a scale reading. The basis is stated.
5. **What is the unit and the rounding?** `12 %` has a space, ranges are written 2022-2026, and a figure is not rounded from `47.3 N·m` to `47 N·m`. A value that is an estimate says so.

Two places need special care. First, the 2027 task weights for Astrobiology, Delivery and Equipment Servicing are not published, and the Autonomy Mission is the only one whose sub-task points are given `S-URC27-RULES §1.e.iv`. Second, the 2027 review and Science Plan specifications were not published when this edition was prepared, so the 2026 documents are used only where they are labelled as such `S-URC27-SAR` `S-URC27-SCIENCEPLAN`.

What the evidence covers well

The rules, the official scores and the team-published material are strong. The 2027 rulebook, the 2022 to 2026 score sheets and the 2027 Q&A snapshot are kept as archived copies. The 2022 to 2026 results are complete enough to support distributions, zero rates and variance decomposition. The top teams publish enough video, code and documentation to reconstruct their architecture at the subsystem level.

The record is weaker in four places:

- **Mass and centre of gravity.** Most teams publish no mass, and fewer publish a centre of gravity. The mass model rests on 17 configurations, 13 of them published masses, plus arithmetic from stated reductions and CAD sums. It is a good model of the field, not a measurement of any one rover.
- **Task-level forces.** The 2027 rules describe the equipment servicing tasks but do not publish the forces, travels or torques. Those numbers have to be measured on a mock. The book gives design targets and says they are targets.
- **Delivery and drones.** The delivery study is a synthesis from the rules, the Q&A and team evidence rather than a dedicated study, and the page says so.
- **The qualification funnel.** Nineteen of the 68 census teams have no profile because they publish almost nothing. They are listed in the census and represented by their score row. Two further teams that registered but never reached the finals sheet have profiles because their public engineering record is unusually good.

A page-by-page account of what is missing, measured or inferred is on evidence coverage.

What the book does not do

The book does not design the rover. It gives the requirements, the evidence, the options and the trades, and it recommends a first-build shortlist. The detail design, the test campaign and the build are the team's work.

It does not reason about the team's budget, funding or ability to pay. Cost appears only as a relative property in a comparison table, from `$` to `$$$$$`, and as sourced list prices with a date. The competition's own $26,000 fielded-equipment cap is recorded as a competition rule, not as a constraint on the analysis `S-URC27-RULES §3.c`.

It does not claim that a choice which worked for a top team will work for a first build. The dossiers show what teams did and what it cost them. The subsystem pages separate the evidence from the recommendation.

Corrections and reuse

The original analysis and the charts made for this book may be reused with attribution to the Robotics Club at Mississippi State University. Quotations, photographs, team documents and vendor images remain the property of their owners and are credited with a link to the original. Where a licence is known it is shown; where it is not, the credit names the author and the source. The credits page states the terms in full.

Corrections are welcome. A correction should name the page, the claim, the source and the replacement value. The standard is the same as the one used here: a number needs an owner, a sample and a date.

> **The one rule**
>
> A claim may be wrong, but it may not be unsourced. Where the record has no answer, the book says so.