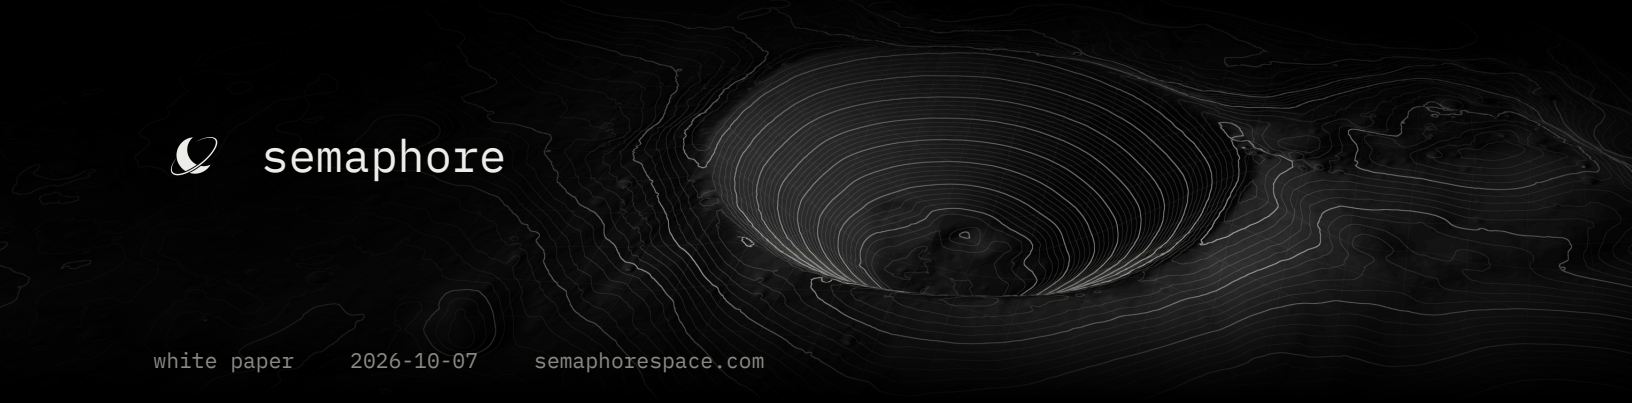


semaphore

white paper 2026-10-07 semaphorespace.com

Physical AI will make Lunar Missions Safer, Faster, and Cheaper

Co-Founder, Chief Hardware Systems Engineer Co-Founder, Chief Software Systems Engineer

Traffic at the Moon is set to rise sharply. Analysts count 116 lunar missions planned for the next ten years [1], and the European Space Agency (ESA) counts more than 400 over the next twenty years [2]. Yet every lunar mission today is run from Earth by its own team, over deep-space antennas such as NASA's, which are already oversubscribed [3]. Its rovers are mostly driven the same way, a few meters at a time [4, 5]. Run that way, the coming traffic would need far more antenna time than exists, and a ground team for every machine.

Machines wait for Earth because they lack current information about the Moon and the vehicles around them. Semaphore builds physical artificial intelligence (AI) infrastructure for the Moon. It supplies that information from below the surface to orbit and turns it into answers. Our orbiters and surface systems will collect operational and environmental intelligence at the Moon. The same orbiters will look out into space and warn of impacts before they occur. Each observation will update our digital twin of the Moon. An early version of the twin runs in the cloud today, and our first orbiters launch in 2028.

We train our physical AI models on the twin. They return typed answers with calibrated probabilities to lunar missions and their machines, in time to act. The models run in the cloud and at the Moon, on our spacecraft and on vehicles that carry them. Answers computed at the Moon still reach vehicles while Earth is out of view. With those answers, lunar operations can continue around the clock and on the far side, as vehicles work toward goals that people set.

Operators and their autonomous machines can choose landing sites, navigate by terrain, and plan around sunlight and radiation. They can track what has changed and screen every move against other missions and known hazards. They can test against the twin before they commit and act on the Moon's latest state instead of a script written far in advance. Semaphore plans to screen lunar orbits as a public service.

The same systems serve users beyond the missions. They supply on-demand imagery of the surface and of spacecraft, and resource prospecting data for in-situ resource utilization. Space insurers receive mission risk data and independent observation of mission execution, and agencies and defense receive independent tracking and space safety data. The result is fewer lost spacecraft, less work and dependence on the ground, and a Moon where many missions can operate quickly but safely side by side.

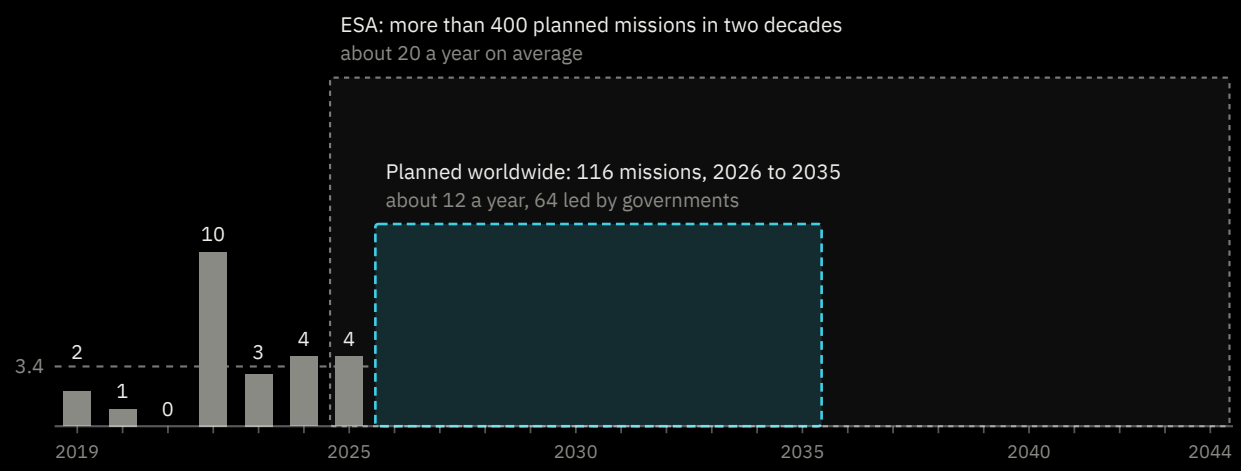
introduction

Every robotic rover that has traveled far on the Moon has been driven from Earth. China's Yutu-2 stops at each waypoint while a team on Earth reconstructs the terrain from stereo images and plans the next leg [4]. The National Aeronautics and Space Administration (NASA) plans to drive its Volatiles Investigating Polar Exploration Rover (VIPER) the same way [5, 6]. Ten drivers will work three eight-hour shifts around the clock, and the driver on console will send a new command every 2 to 8 m [5, 7]. VIPER can drive at 20 cm/s, but commanding, localization, and hazard checks bring its average to about 1 cm/s. That is roughly the speed at which shadows sweep across the polar terrain [8]. At that rate its planned traverse of up to 20 km takes about 23 days of driving [8, 9].

That way of working will not scale to the traffic ahead. Of the 116 missions planned for 2026 to 2035, 64 are led by governments, and only about half come from North America [1]. China and Russia are developing their own lunar research station in parallel with the United States effort [1]. NASA's Moon base plan alone calls for up to 25 missions by 2029, including 21 landings, and up to 24 more landings by 2032 [10]. Planned missions come at more than three times the recent pace, and ESA's count over two decades implies nearly six times (Figure 1) [1, 2, 11].

Planned lunar missions run at more than three times the recent pace, and ESA's two-decade count at nearly six times

Lunar missions launched worldwide each year, missions planned for 2026 to 2035, and ESA's count of planned missions to 2044



2022 counts Artemis I and the small spacecraft it carried. Launched counts are a floor.
The dashed grey line marks the recent average, 3.4 launched missions a year. Each band spreads its total evenly.
The two counts come from different sources and methods.

Figure 1. Lunar missions launched each year [11], missions planned worldwide for 2026 to 2035 [1], and ESA's count of planned missions over two decades [2]

The traffic includes crews, power, construction, and tourism. In May 2026 NASA awarded Astrolab $219 million and Lunar Outpost $220 million to build crewed lunar rovers. It awarded Blue Origin $188 million to land them, with mobility systems on the surface by 2028 [12]. A 2025 NASA directive calls for a reactor of at least 100 kilowatts electric, ready to launch early in fiscal year 2030 [13]. Lunar Forge plans to sinter regolith into housing and shielding for such reactors [14]. GRU Space plans to land a lunar hotel for up to

four guests in 2032 [15]. Each plan depends on where its site lies, when it is lit, what the terrain permits, and what else moves nearby.

Every vehicle driven from Earth needs a team, time on the large dish antennas that reach deep space, and a line of sight to Earth or to a relay satellite. NASA's Deep Space Network (DSN) has been oversubscribed for more than 30 years, at times by 40 percent. NASA expects demand to grow tenfold by the early 2030s [3]. On Mars, Perseverance waited for its next commands at least seven times in 21 months because no antenna was free [3]. A landing site near the south pole sees Earth directly for only 17 to 43 percent of a year [16]. Relay satellites such as NASA's Lunar Communications Relay and Navigation Systems and ESA's Moonlight will raise that coverage and add positioning [2, 16]. Terrain hazards, illumination, and the location of other vehicles are outside their scope.

NASA's Mars 2020 mission shows another way. Its Perseverance rover has driven itself for more than 90 percent of the distance it has covered, and once drove 699.9 m without human review [17, 18]. Its autonomous navigation drives at about 110 m an hour, against 15 to 18 m an hour for Curiosity [19].

Distance forced that change. A signal takes up to about 22 minutes each way between Earth and Mars [20]. At the Moon, the time from sending a command to seeing its result is 6 to 10 seconds. That is short enough for NASA to drive VIPER from Earth in near real time [5]. NASA gives a second reason for keeping people in the loop. AI techniques "require a lot of training data," and VIPER is the first robotic rover NASA will drive remotely on the Moon [21].

The number of machines will force the same change at the Moon. The first steps are under way. In 2024 two small robots from the Japan Aerospace Exploration Agency (JAXA) became the first rovers to explore the lunar surface fully autonomously [22]. NASA's Endurance rover concept would cross about 2,000 km of the South Pole-Aitken basin in under four years and names autonomy as its key enabling technology [23]. NASA's Moon base plan counts on "a fleet of crewed and autonomous rovers" to keep the base "supplied and functioning continuously" [10].

Landers already perform powered descent autonomously and correct their guidance on board in real time [24]. Descent is where most recent landers were lost. Of thirteen landing missions since 2019, five came down upright and worked as planned. Most of the eight that crashed or tipped over failed in the hardware or software for navigation, altimetry, or guidance (Figure 2) [25–37].

NASA states the need:

> "Autonomous systems allow spacecraft, rovers, and robots to operate without relying on constant contact with astronauts or mission control." NASA Space Technology Mission Directorate, 2024 [38]

> "Before hardware is trusted to operate independently on location, which for Artemis missions includes the Moon, it must be tested virtually." NASA Space Technology Mission Directorate, 2024 [38]

In two of the eight failures the operators traced part of the cause to the data the lander was tested and navigated against. ispace changed its landing site after its critical design review, and its simulations "did not adequately incorporate the lunar environment on the navigation route" [32]. Intuitive Machines reported that IM-2's optical navigation relied on orbital images taken from 100 km, which could not account for how craters appear at low altitude in south-polar light [36]. Where the reference map was adequate, the method

worked. Japan's Smart Lander for Investigating Moon (SLIM) matched craters to its onboard map [39]. At 50 m altitude it was within about 10 m of its target, possibly 3 to 4 m, against a goal of 100 m [34].

Closing that gap takes lunar intelligence gathered for the whole Moon and the space around it. Operational intelligence is the state, path, and plan of each vehicle and crew working there. Environmental intelligence is the world they work in: terrain, light, temperature, dust, gravity, resources, radiation, and space weather. It is historical, current, and projected. Kept current, it supplies the models' training data and the record against which each of their answers is graded. It is the basis for checking an orbit, a landing, or a drive against other missions. Operators and governments now ask for that check [40, 41].

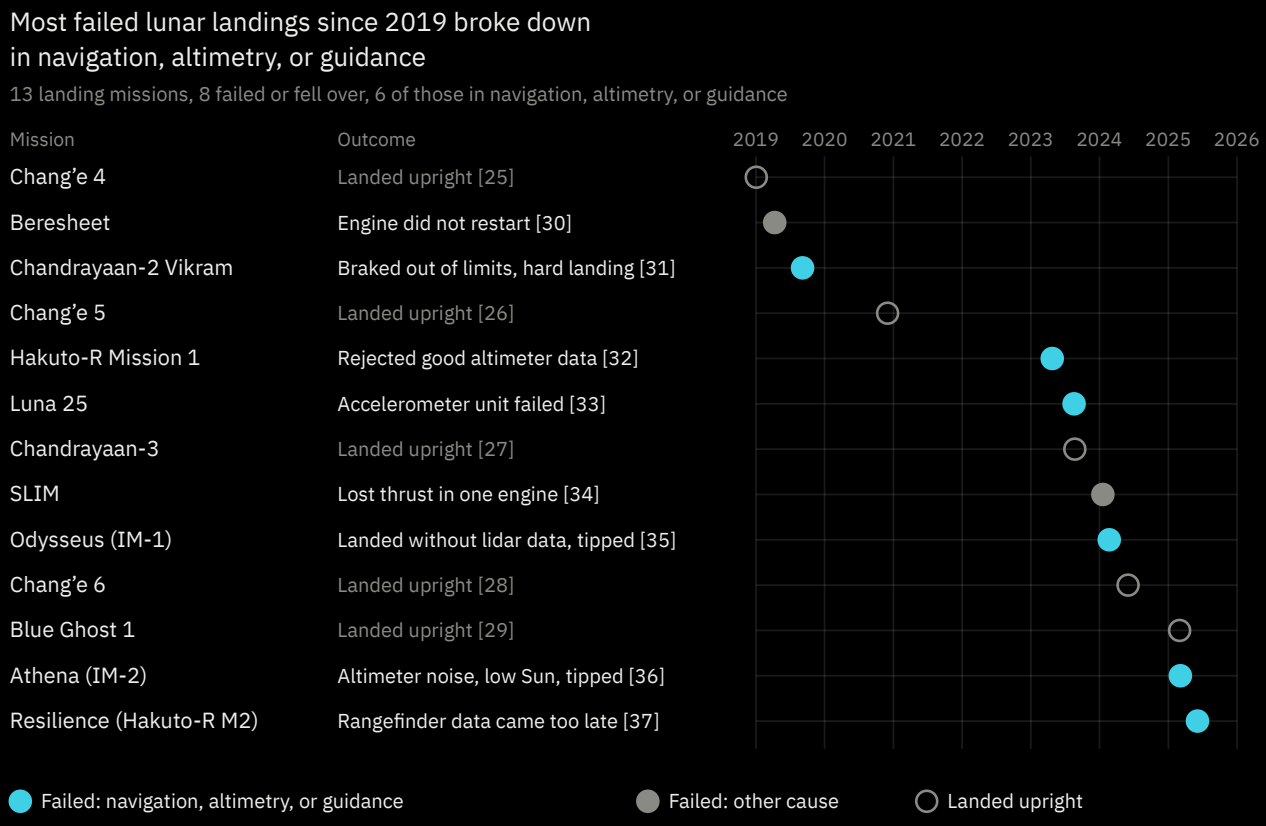


Figure 2. Lunar landing missions since 2019 and the cause of each loss [25–37]

why shared lunar intelligence?

Today each mission assembles its own site evidence and plans from public maps, staffs a driver team, and books antenna time separately [3, 5, 42]. Shared intelligence makes missions safer, faster, and cheaper. Its value grows with traffic.

safer

Two of the recent losses trace in part to the data the lander was tested and navigated against (Figure 2) [32, 36]. Shared intelligence gives each lander the same current account of terrain, light, and traffic, ren-

dered for its hour. It grades every answer against what happened. Screening adds a check across missions: whether a plan is safe for its neighbors [40, 41].

faster

Today a lunar machine stops whenever a person on Earth must approve its next step. With current answers, it checks that step itself against the ground, the light, and its neighbors, and keeps working. Those answers are broadcast by our spacecraft, sent from the cloud, or computed by a model the machine carries. NASA's 2026 ranking of civil space technology shortfalls places onboard computing third of 32 [43]. On Mars, a rover that drove itself crossed a field of boulders in about a third of the time earlier NASA rovers would have needed [18].

The work now funded and planned for the Moon puts many machines on one site. Crewed rovers and the landers that deliver them share it with a power plant and machines that build with regolith [12–14]. When each move waits on Earth, every machine on the site waits with it. When the machines know the site and each other's plans, they can work side by side at their own pace. Construction and science then go on at the same time.

cheaper

Lunar missions pay for the same work many times. When NASA moved commercial deliveries to new sites, the changes cost NASA an extra $46.6 million and delayed one mission by more than a year. The changes were largely to adapt landers to conditions at the new sites [42]. A team that has current evidence on a site's terrain, light, and temperature before it designs a lander can plan for those conditions from the start. Waiting on Earth costs staff time. Onboard autonomy on NASA's Earth Observing-1 satellite cut its operations cost by about $1 million a year [44]. A vehicle that reads its own ground and sends home only its answers needs fewer hours on deep-space antennas [3].

Today each lander and rover carries every capability it needs. The Defense Advanced Research Projects Agency (DARPA) began its 10-Year Lunar Architecture study to end the arrangement "whereby each lunar lander or activity must organically support all required resources such as survival power, communications, and data storage" [45]. NASA's relays will establish "a GPS-like capability at the Moon," and a commercial power service for the surface is in development [16, 46]. Shared intelligence applies the same idea to perception. A vehicle in range of our orbiters and surface systems receives hazard maps, position fixes, and advisories that its own sensors would otherwise have to produce. The sensing that keeps it safe when links drop stays on board. Designers can then leave out redundant mapping hardware and spare compute.

---

beyond the missions

The same intelligence serves four groups of users who do not operate missions.

imagery of the surface and of spacecraft

Independent images confirm what happened at the Moon. Today they come mostly from NASA's Lunar Reconnaissance Orbiter (LRO). It imaged the IM-1 lander about 44 hours after it landed and Blue Ghost the same day [47, 48]. LRO took 87 days to find India's Vikram lander, in part because its first images of the

site were poorly lit [49]. Imaging another spacecraft is rarer still. LRO's image of Danuri needed precisely timed pointing and still smeared the orbiter to ten times its size [50]. NASA states that LRO is beyond its design lifetime, and its proposed successor is a concept in early study [51].

Our orbiters will image the surface and the spacecraft around the Moon on request. Each image will enter the twin with its geometry and lighting, so a landing can be confirmed or a lost vehicle found at the next pass.

risk data for space insurers

Space insurers price risk on heritage, and lunar landers have little [52]. In 2023 the space insurance market as a whole faced $1.43 billion in claims against about $550 million in premiums. Aon measures the market's full capacity against risks "with demonstrated reliability and heritage" [52]. The insurer that wrote the first lunar policy now ends cover earlier in the mission. ispace's Mission 1 policy covered launch through landing, while its Mission 2 policy ended once the lander reached a 100 km lunar orbit [53, 54]. The Mission 1 payout came 114 days after the crash because the landing could not be confirmed [53]. Parametric insurance pays on an index that an independent third party reports [55]. An independent record of each mission phase gives insurers the evidence to cover descent and landing, to settle claims, and to write parametric cover with a lunar trigger.

prospecting data for in-situ resource utilization

In-situ resource utilization (ISRU) needs maps that do not exist yet. NASA's Moon to Mars objectives call for industrial-scale ISRU, and the first need under its lunar ISRU shortfall is to "locate, characterize, and map" useful resources [43, 56]. NASA's own measurement study found a single ground-truth point for buried polar ice: the 2009 impact of the Lunar Crater Observation and Sensing Satellite. It found orbital data too coarse to define a reserve, and it asked whether a bank would "front the loan to start the mine" [57]. Only about 3.5 percent of polar cold traps show ice at the surface [58]. The missions meant to close the gap have been lost, cut short, or delayed. Lunar Trailblazer was lost after launch, and the lander carrying NASA's Polar Resources Ice Mining Experiment came to rest on its side. VIPER has yet to launch [59–61].

The twin will join orbital estimates to landed ground truth. It will show where water ice and other resources are likely to be and at what depth, and update as each mission reports. Each estimate will carry a calibrated probability and its evidence. Developers, agencies, and lenders will use it to choose where to prospect and how to size a plant.

tracking and space safety data for agencies and defense

No catalog of independently tracked objects exists at the Moon. NASA states that no Department of Defense catalog covers lunar space and that passive tracking is not available beyond geosynchronous orbit [62]. NASA's Jet Propulsion Laboratory (JPL) reports that not every operator shares its ephemerides and that debris at the Moon cannot be tracked [63]. When a rocket body struck the Moon in 2022, its origin was still uncertain months later [64]. The 2022 National Cislunar Science and Technology Strategy called for an integrated catalog of cislunar objects that includes human-made objects on the surface and is built with private partners [65]. NASA's 2026 technology priorities still include extending U.S. space situational aware-

ness to cislunar space [43]. A 2026 paper by the co-chairs of a United Nations action team finds awareness of non-active objects at the Moon "limited or non-existent" [66].

The catalog will draw on our own observations of spacecraft, landers, rovers, and debris. It will include objects whose operators do not report them. Agencies will receive independent orbit states, confirmations of landings and impacts, and records of activity near declared safety zones. These will support the notification and coordination the Artemis Accords call for [41].

The Moon's surface changes mostly when something strikes it from space [67]. In 2024 NASA's Scout system predicted where and when a small asteroid would hit Earth's atmosphere, 95 minutes before it did [68]. Our orbiters will look out into space and give the same kind of warning at the Moon.

design principles

Each of our models is scoped to one operational question, such as the traversability of terrain ahead of a rover or the probability of collision between two spacecraft. A narrow scope keeps a model small. It trains faster and at lower cost, fits within the memory and power budget of a flight computer, and can be verified and validated across its operating envelope. Each model is calibrated, and its evaluation is published with reliability diagrams and calibration error measured against recorded outcomes [69]. Calibration lets an operator use a model's probability directly in a risk decision.

We build on proven open-source flight and ground software and add to it only where our mission's concept of operations requires. We contribute changes back where they help other users. On board our spacecraft, our own code centers on the applications that run our sensors and models. The flight software underneath is NASA's Core Flight System (cFS), built at NASA's Goddard Space Flight Center. It has run more than 40 NASA missions and was on board IM-1, the first commercial lunar landing [70, 71]. NASA engineering studies rate cFS at technology readiness level (TRL) 9, the top of NASA's scale. It has run in lunar orbit on LRO since 2009 [72, 73]. Our links follow the open standards that NASA's lunar interoperability specifications name, so other systems can read what ours send [74, 75].

data architecture

Our models improve in a loop. We build a dataset, train a model on it, and test it on real data it has never seen. The model then answers questions, and we compare each answer with what actually happened. Those results go into the next dataset.

Data enters from three directions. Public archives come first. They hold decades of record, such as 17 years of LRO maps in NASA's Planetary Data System and radiation records from the National Oceanic and Atmospheric Administration (NOAA) [76, 77]. Our own original data products come next. Once our orbiters and surface systems launch, they add observations that are not in public archives or operator reports [78]. Missions contribute their plans and states when they submit them for screening. Each dataset traces every sample to its source and states the terms of its use. An operator's data enters training only when that operator offers it.

We build the twin from those sources one layer at a time: terrain, light, temperature, traffic, and radiation. Each layer is checked against its sources and against independent references, and it lists the checks it passes [79].

We render synthetic training images from the twin's terrain, and each image comes with exact ground-truth labels. Real images close the sim-to-real gap, the difference between rendered images and real ones. Models that work on other sensor data, such as radiation, train on measured time series.

Each answer is kept with what happened next: the measured orbit after a close approach, the images after a landing, the dose after a storm. We score every model against those outcomes.

A model runs in one of three places.

1. On the ground it answers through our software interfaces beside a mission's own systems, for teams that still decide from Earth. Screening starts here.
2. On our own orbiters and surface systems it watches independently and broadcasts advisories such as a close approach, a landing plume, or a blocked path. Any vehicle in range can receive them, whether or not it carries the model.
3. On a user's lander, rover, or orbiter it runs as an application beside that vehicle's own software.

Wherever it runs, a model gives the same answers. Those answers are shared across missions. It does not replace a vehicle's own guidance.

---

how we get there

We retire the largest risks first, with the fewest resources. Our path runs in five steps (Figure 3).

1. We develop a digital twin of the Moon from public NASA data and our original data products. An early version runs in the cloud today with an operations console. Orbit screening comes first, as a public service [79, 80].
2. We demonstrate end-to-end integration between the twin and a simulated spacecraft. We reported this on 2026-10-04. NASA's cFS and our instrument application took three timed observations over Tranquility Base. The ground then published the orbit back to the twin. That closed the loop in simulation from the twin to the flight software to the ground and back [75].
3. We demonstrate the same integration with flight hardware in the loop: a flight computer and sensing instruments. This is where the models join the flight software. Until then they train on synthetic data from the twin and on real labeled data, and are tested against real NASA data. At this step they run on flight hardware and take in data from real sensing instruments for the first time.
4. In 2028 we launch two small orbiters into low Earth orbit (LEO). The spacecraft are mostly off the shelf, and the applications and models on board are ours. What they test is our physical AI infrastructure, from sensing through the twin and the models to the answers. Answers follow the paths they will take at the Moon: between our orbiters, to other operators' satellites that listen to our broadcasts, and through the cloud. We grade those answers against true values from satellite navigation and the public catalog of Earth orbit. The results, with the orbiters' own data, train the next models. Each test retires a risk the lunar mission would otherwise carry, and in LEO a fault costs far less to fix than at the Moon.

5. With the system proven in LEO, we launch at least two small orbiters of our own to lunar orbit, followed by our surface systems. They observe the ground and the spacecraft around it independently of what operators report, grade the models, and carry their answers to every vehicle in range [78].

Users do not have to wait for our spacecraft. We offer each model through our software interfaces as soon as it passes its tests on real data.

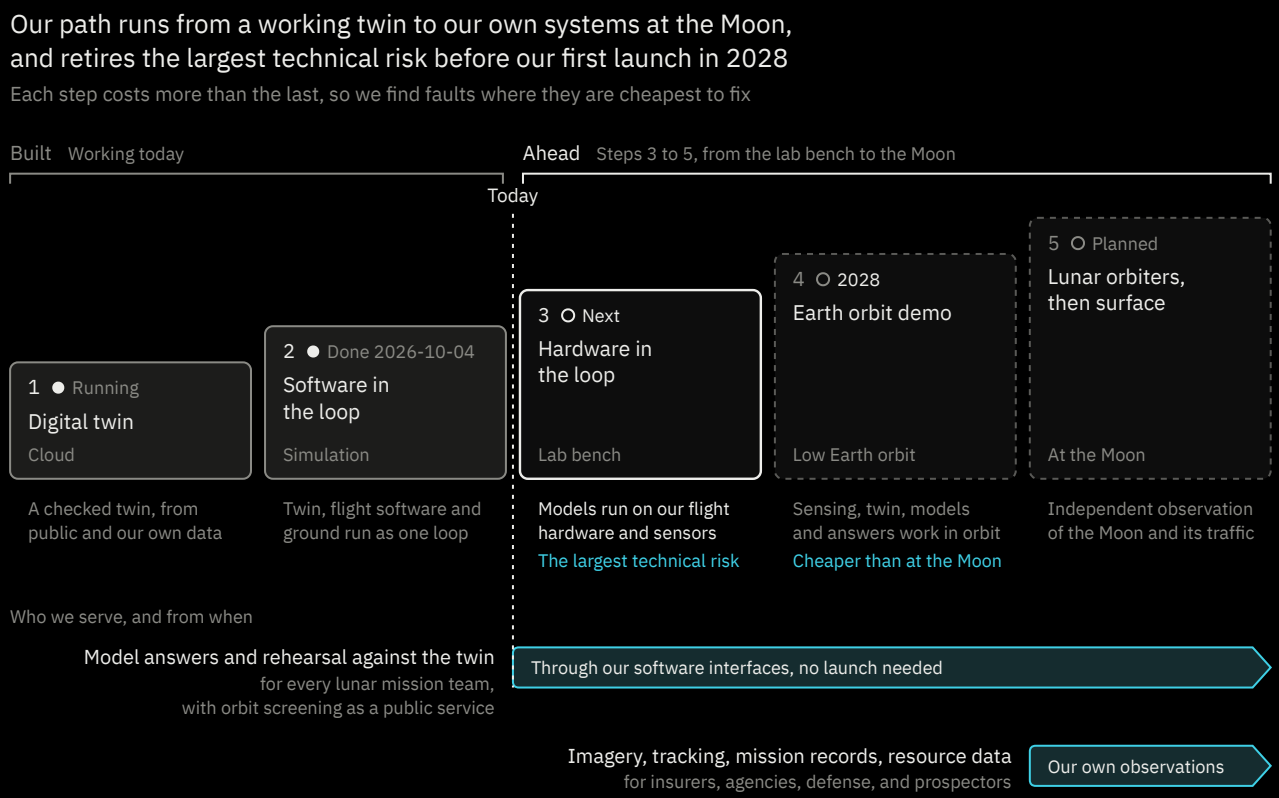


Figure 3. The path to the Moon: what each step proves, and who it serves

We invite every lunar operator to screen their orbits with us, and every mission team, agency, and researcher to build on the twin. Write to us at contact@semaphorespace.com.

references

1. Analysys Mason, Moon base plans will push back the lunar market opportunity (August 2026)
2. ESA, ESA launches Moonlight to establish lunar communications and navigation infrastructure (October 2024)
3. NASA Office of Inspector General, IG-23-016, audit of the Deep Space Network (2023)
4. J. Wang et al., Computer vision in the teleoperation of the Yutu-2 rover, ISPRS Annals (2020)
5. NASA Science, VIPER lunar operations
6. NASA, NASA selects Blue Origin to deliver VIPER rover to Moon's south pole (September 2025)
7. NASA Ames, lunar science real-time operations and mission systems integration on NASA's VIPER mission, SpaceOps (2025)
8. M. Shirley et al., NASA Ames, VIPER traverse planning, Lunar and Planetary Science Conference (2022)
9. K. Ennico-Smith et al., NASA Ames, VIPER mission traverse planning (2023)
10. NASA, Building the Moon Base (March 2026)
11. NASA Space Science Data Coordinated Archive, lunar exploration timeline
12. NASA, NASA provides update on Moon base rovers, landers, missions (May 2026)
13. NASA, directive on fission surface power (August 2025)
14. Payload, Lunar Forge unveils plans to build lunar infrastructure out of dust (July 2026)
15. GRU Space, lunar hotel plans
16. NASA Goddard, lunar communications relay and navigation systems (2025)
17. NASA JPL, NASA's Perseverance Mars rover ready to roll for miles in years ahead (December 2025)
18. NASA JPL, autonomous systems help NASA's Perseverance do more science on Mars (September 2023)
19. NASA JPL Robotics, Mars 2020 rover mobility
20. NASA, communication delays, disruptions, and blackouts for crewed Mars missions, ASCEND (2022)
21. NASA, artificial intelligence and NASA's first robotic lunar rover (December 2023)
22. JAXA, LEV-1 and LEV-2 lunar surface exploration (June 2026)
23. J. T. Keane et al., NASA, Endurance: lunar South Pole-Aitken basin traverse and sample return rover (2022)
24. C. I. Restrepo et al., NASA, next generation NASA hazard detection system development (2020)
25. China National Space Administration, Chang'e-4 landing on the Moon's far side (January 2019)
26. China National Space Administration, Chang'e 5 lands on the Moon (December 2020)
27. Indian Space Research Organisation, Chandrayaan-3 (August 2023)
28. China National Space Administration, Chang'e-6 lands on the Moon's far side (June 2024)
29. NASA, cameras on Blue Ghost capture first-of-its-kind Moon landing footage (March 2025)
30. The Times of Israel, Beresheet engine restart investigation (April 2019)
31. ThePrint, government confirms Chandrayaan-2 Vikram hard landing (November 2019)
32. ispace, results of the HAKUTO-R Mission 1 lunar landing (May 2023)
33. TASS, cause of the Luna 25 loss (October 2023)
34. JAXA, SLIM landing results (January 2024)
35. SpaceNews, IM-1 lunar lander mission winds down (February 2024)
36. Spaceflight Now, issues that prevented an upright touchdown on the IM-2 landing (May 2025)
37. ispace, technical cause analysis for HAKUTO-R Mission 2 (June 2025)
38. NASA Space Technology Mission Directorate, Student Lunar Autonomy Challenge (2024)
39. JAXA Institute of Space and Astronautical Science, 20 minutes of terror: SLIM will attempt a pinpoint landing (2023)
40. D. Berry, NASA JPL, conjunction assessment process for lunar orbiters (2022)
41. NASA, The Artemis Accords (2020)
42. NASA Office of Inspector General, IG-24-013, NASA's Commercial Lunar Payload Services initiative (2024)
43. NASA, FY2026 civil space technology shortfall prioritization (2026)
44. R. Sherwood, S. Chien et al., NASA JPL and Goddard, enhancing science and automating operations using onboard autonomy (2006)
45. DARPA, 10-Year Lunar Architecture (LunA-10) capability study
46. Astrobotic, LunaGrid-Lite completes critical design review, flight model underway (August 2025)
47. NASA, NASA's LRO images Intuitive Machines' Odysseus lander (February 2024)

48. LROC team, NASA LRO, Blue Ghost seen obliquely (March 2025)
49. LROC team, NASA LRO, Vikram lander found (December 2019)
50. NASA, NASA's LRO finds photo-op as it zips past South Korea's Danuri Moon orbiter (April 2024)
51. NASA Science, Lunar Exploration and Science Orbiter (LExSO)
52. Aon, Space Insurance Market Report, Q1 2026
53. ispace, extraordinary report on receipt of lunar insurance proceeds, filed with the Kanto Local Finance Bureau (August 2023, in Japanese)
54. ispace, Mission 2 update and lunar insurance (December 2024, in Japanese)
55. Swiss Re, parametric solutions as a contributor to closing protection gaps (2023)
56. NASA, Moon to Mars objectives (2022)
57. J. Kleinhenz et al., NASA Glenn, Lunar Water ISRU Measurement Study, NASA/TM-20205008626 (2020)
58. S. Li et al., Direct evidence of surface exposed water ice in the lunar polar regions, PNAS (2018)
59. NASA JPL, NASA's Lunar Trailblazer Moon mission ends (August 2025)
60. NASA, NASA's lunar drill technology passes tests on the Moon (April 2025)
61. NASA Science, VIPER
62. L. Newman, NASA, conjunction assessment briefing (2025)
63. Z. B. Tarzi and D. S. Berry, NASA JPL, international collaboration in extraterrestrial conjunction assessment (2025)
64. NASA, NASA's Lunar Reconnaissance Orbiter spots rocket impact site on Moon (June 2022)
65. National Science and Technology Council, National Cislunar Science and Technology Strategy (2022)
66. United Nations Office for Outer Space Affairs, A/AC.105/2026/CRP.39, Action Team on Lunar Activities Consultation (2026)
67. NASA Science, Moon facts
68. NASA, NASA system predicts impact of a very small asteroid over Germany (January 2024)
69. C. Guo, G. Pleiss, Y. Sun, K. Weinberger, On calibration of modern neural networks, ICML (2017)
70. NASA Goddard Space Flight Center, Core Flight System
71. NASA Goddard, Intuitive Machines marks a new lunar era with help from NASA's Core Flight System (February 2024)
72. Origins Space Telescope Study Team, NASA, Origins Space Telescope mission concept study report (2019)
73. L. Prokop and J. Wilmot, NASA, NASA's Core Flight Software: a reusable real-time framework (2014)
74. NASA, International Communication System Interoperability Standards (2024)
75. Semaphore, software in the loop, end to end (October 2026)
76. NASA Planetary Data System, LRO Camera reduced data records (2009 to present)
77. NOAA National Centers for Environmental Information, GOES-R space environment data
78. Semaphore, hardware
79. Semaphore, layers of the digital twin (2026)
80. Semaphore, introducing Semaphore (2026)