Case study

Integrating a multi-supplier lunar program

How lateralworks rebuilt one integrated master schedule and a weekly refresh cadence to put a commercial lunar lander program back on a converging path.

FTTM engagement series.

A commercial lunar exploration company, roughly a year behind its first customer target, with critical work spread across international subcontractors, a prime contractor, and a government sponsor — and a schedule that had stopped converging.

Prepared by
lateralworks
Program management

Engagement
Lunar lander program, 2023
Engagement review, 2026 edition

Online
lateralworks.com
FTTM engagement series

Engagement summary

What was at stake, and what was achieved

A commercial lunar program was roughly a year behind its first customer commitment, and the gap was widening. The work was distributed across international subcontractors — the structural bus, the navigation system, propulsion, avionics, and flight software — but the risk of it failing to come together was held in one place, by the company. The program carried a negative margin, so every month of slip converted directly into financial exposure. In an industry where fixed-price lunar contracts have pushed more than one company into bankruptcy, an open-ended schedule gap was an existential risk, not a planning inconvenience.

404 days
Initial critical-path gap to the first customer target

1 schedule
All suppliers' work integrated and refreshed weekly

Every target
Critical path re-identified and pulled in each week

In-house
Scheduling capability transferred to the client's own team

lateralworks rebuilt how the program was planned and run. The many disconnected schedules were replaced by a single integrated master schedule that every supplier's work lived inside, refreshed in one room every week. At each refresh the critical path to every target was re-identified, the cause of every change was recorded, and pull-in actions were opened and tracked until the trend to each date turned.

By the close of the engagement, the program had the machinery to converge. The first critical-path study had put the finish more than four hundred days past the first customer target; the gap was now quantified, owned, and being driven down through continuous pull-in rather than absorbed. Supplier schedules were folded into the integrated master schedule, with suppliers helping build it for the first time. A core scheduling team was seeded inside the company and the skills were transferred to its own managers. The schedule had stopped being a reporting artifact and become the place the program was run.

The system to close the gap existed; using it was now the company's work. Supplier integration had started but was not complete, the core team was a beginning rather than a finished structure, and two strategic questions — the size of the payload, and whether to fly the mission incrementally — were left deliberately open for the company to resolve.

Program at a glance

One integrated schedule across every functional element

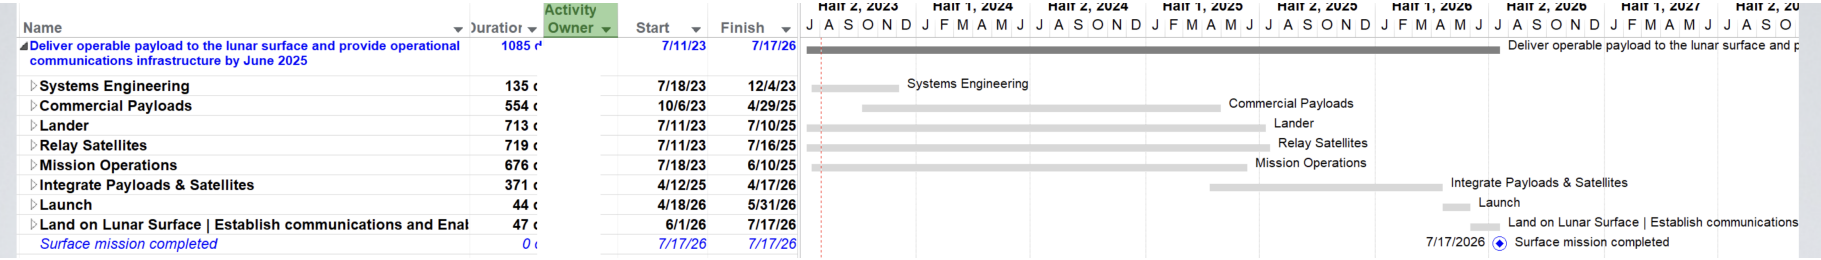

Name	Duration	Activity Owner	Start	Finish
Deliver operable payload to the lunar surface and provide operational communications infrastructure by June 2025	1085 d		7/11/23	7/17/26
Systems Engineering	135 d		7/18/23	12/4/23
Commercial Payloads	554 d		10/6/23	4/29/25
Lander	713 d		7/11/23	7/10/25
Relay Satellites	719 d		7/11/23	7/16/25
Mission Operations	676 d		7/18/23	6/10/25
Integrate Payloads & Satellites	371 d		4/12/25	4/17/26
Launch	44 d		4/18/26	5/31/26
Land on Lunar Surface | Establish communications and Enal	47 d		6/1/26	7/17/26
Surface mission completed	*0 d*		*7/17/26*	*7/17/26*

The high-level integrated master schedule. Every functional element and subcontract — the structural bus, the navigation system, propulsion, commercial payloads, relay satellites, and mission operations — was laterally integrated into one system schedule with a single objective, baseline, and critical path, rather than tracked as separate supplier plans. This is the full scope the weekly refresh kept driving toward convergence.

01

The program
What was at stake

The client was a commercial space company building autonomous systems to land payloads on the lunar surface. Its near-term goal was concrete: land a set of customer experiments on the Moon and establish communications by mid-2025. Reaching it meant finishing a lander, qualifying its avionics, integrating flight software, and proving the whole stack through a sequence of reviews. No single organization held that work end to end.

Most of it sat with subcontractors. Switch cards, avionics boxes, guidance and navigation, flight software, and the structural bus were each developed by different suppliers, on different schedules, against different assumptions. A prime contractor sat between the company and a government sponsor. The company carried the integration risk for all of it.

When lateralworks was brought in, the program was roughly a year behind the customer target, and the gap was not closing. With significant technical unknowns still open, the schedule carried real financial exposure. Every month of slip widened an already negative margin. The question put to the engagement was direct: could the program be run differently so the gap shrank, without compromising what the first customer, the shareholders, and the supplier partners each needed?

The honest answer started with a diagnosis. The program did not have one schedule. It had many, loosely stapled together, with no cadence that forced them to converge.

02

The as-found state

Why the schedule had stopped converging

An integrated master schedule existed, but it was not being used to run the program. It had been built in functional detail to satisfy the prime contractor's reporting requests, not to manage the critical path to a date. Supplier schedules lived outside it. Slips arrived week to week with no agreed cause and no recovery plan, and they were absorbed rather than fought.

The deeper issues were behavioral, and the engagement recorded them plainly. The team tended to justify a slip rather than solve it. Time and money were treated as if they were abundant. Attention went to accounting for schedule margin rather than to protecting it, a habit that tends to make late finishes self-fulfilling. Decisions deferred to what the prime and the sponsor wanted instead of the company taking control of its own schedule. Program-management maturity was thin, ownership of outcomes was low, and meetings ran without focus.

Concerns also ran above the program. Observers pointed to heavy capital intensity with no clear path to profitability, leadership gaps at the corporate level, and a design lead whose insistence on full detail blocked the “roughly right” thinking the schedule needed. None of this was hidden. The first job was simply to make the schedule gap visible, name its causes, and make closing it the team’s problem rather than something that happened to them.

03

The approach

One integrated schedule, refreshed every week

The fix was structural before it was cultural. lateralworks rebuilt the integrated master schedule so it managed the program rather than described it. The detailed functional view that served the prime's reporting was set aside. In its place went a schedule organized around the critical path to each target date, at a level a program could actually be steered with.

A new baseline and a new set of targets were established and, importantly, accepted by the prime contractor as the going-forward plan. That acceptance mattered. It gave the team one set of dates to converge on instead of two.

Supplier schedules were then pulled into the same model. For the first time, key subcontractors helped build the IMS rather than reporting against a plan handed to them. Each supplier’s deliverables entered the schedule as linked activities on the critical path, so a slip inside one supplier showed up immediately as movement on a target the whole program could see.

The cadence that held it together was a weekly Refresh. Once a week the team, and increasingly the suppliers, updated the schedule together in one room. The Refresh did three things every time. It re-identified the critical path to each target. It captured the cause of every change, so a slip was never anonymous. And it drove pull-in actions on the paths that mattered, with the team asked what it would take to recover, not whether recovery was possible.

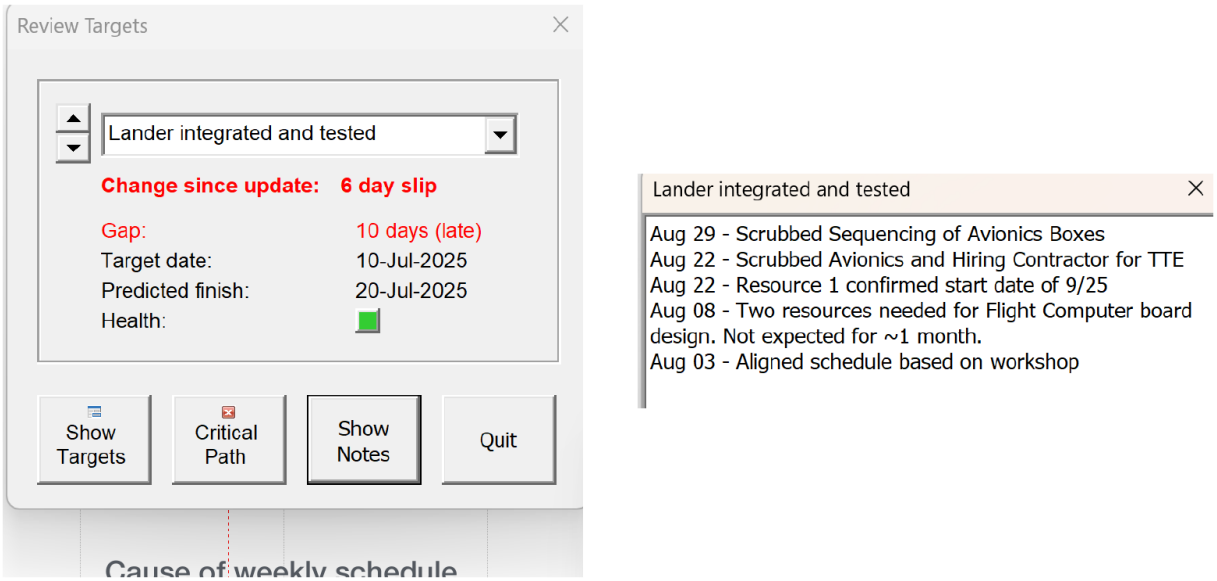

Each target carried a quantified gap, a health state, and a predicted finish (left). Every weekly change was logged with its cause, so movement was never anonymous (right).

Trends to each target were tracked from week to week, so the program could see not just where a date stood but which way it was moving, and why.

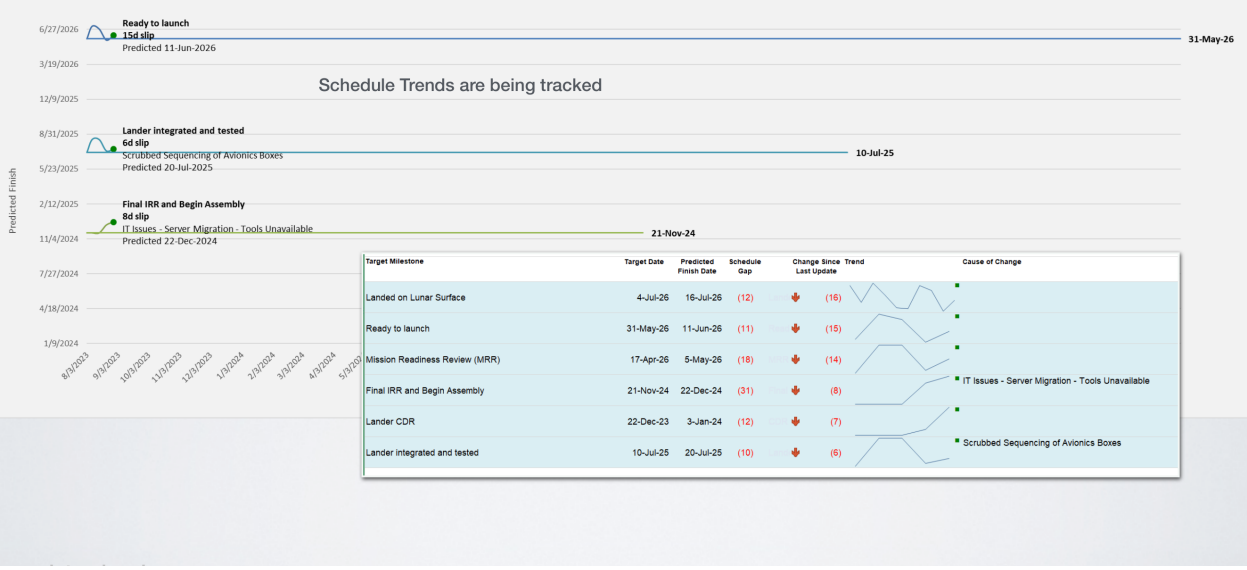

Predicted finish for each target, tracked weekly against the committed date. Trend, schedule gap, and the cause of the latest change sit side by side.

This is the part worth dwelling on, because it is the transferable mechanism. Integration was not achieved by a meeting or a tool. It was achieved by putting every supplier's work into one schedule, refreshing that schedule in one room every week, and tracking the trend to each target until the causes of slip had names and owners.

04

Finding time
Challenging the assumptions behind the gap

Rebuilding the schedule exposed the gap. Closing it needed time the plan did not obviously contain, so the engagement ran the team through lateralworks' FTTM method for finding it.

The method starts from a simple permission: to ask, of any part of the plan, why it is the way it is, and whether there is another way to do it. The team takes an area of focus — here, landing customer experiments and establishing communications by the target date — writes down the assumptions baked into the current approach, and challenges the ones that move the schedule most. Some assumptions survive the test and the plan stands. Others give way and open new ideas. Even when the answer is unchanged, the exercise tends to surface options the team could not see before.

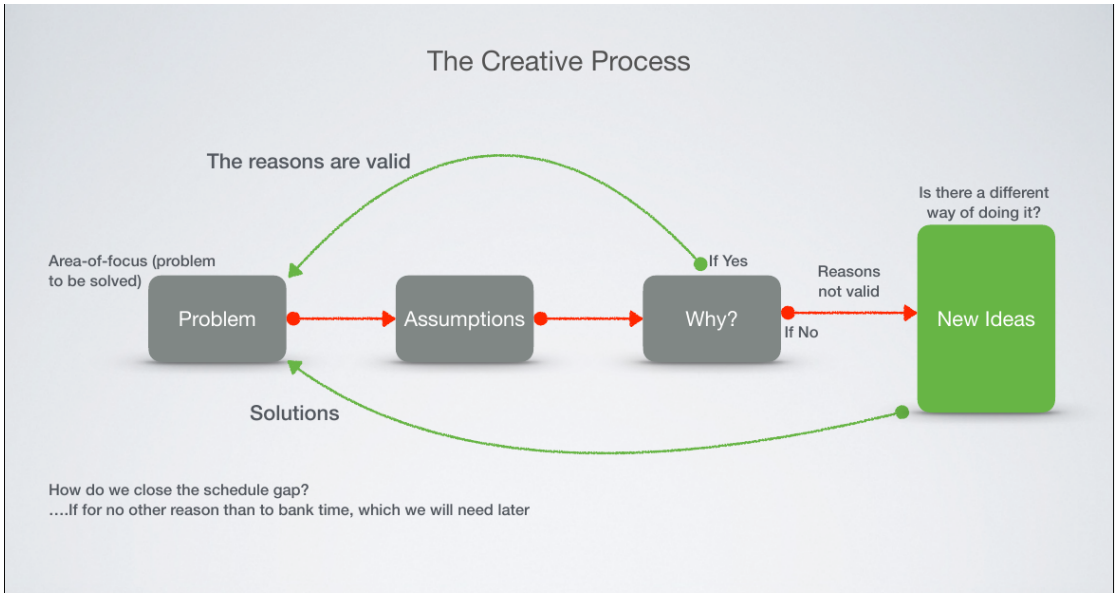

The assumption-challenging loop: take the problem, surface the assumptions, ask why, and either validate the plan or generate new ideas — banking time the program will need later.

Underneath the method sits a model of what a healthy program holds in balance: the right product, the right time, and the right team, continuously defined and planned inside an environment that has to support all three. The engagement used it to frame what the program was missing — clearer product definition, a faster near-term schedule, deeper supplier integration, and senior talent on the team.

The three rights — product, time, and team — held in balance and continuously re-planned. The frame the engagement used to locate what the program was missing.

The sharpest assumption was the one no one had answered: what is the payload? A lander is sized and designed around what it carries. With the payload undefined, the team was designing a large, expensive vehicle against requirements it could not yet state. The engagement pushed the question hard, and paired it with a second. Did the mission have to be attempted in one step? A national space agency had recently reached the Moon by deliberately limiting payload mass and scope, succeeding precisely because it narrowed its risk. The same logic suggested an incremental path: prove a smaller capability first, then scale, rather than betting everything on a single oversized increment.

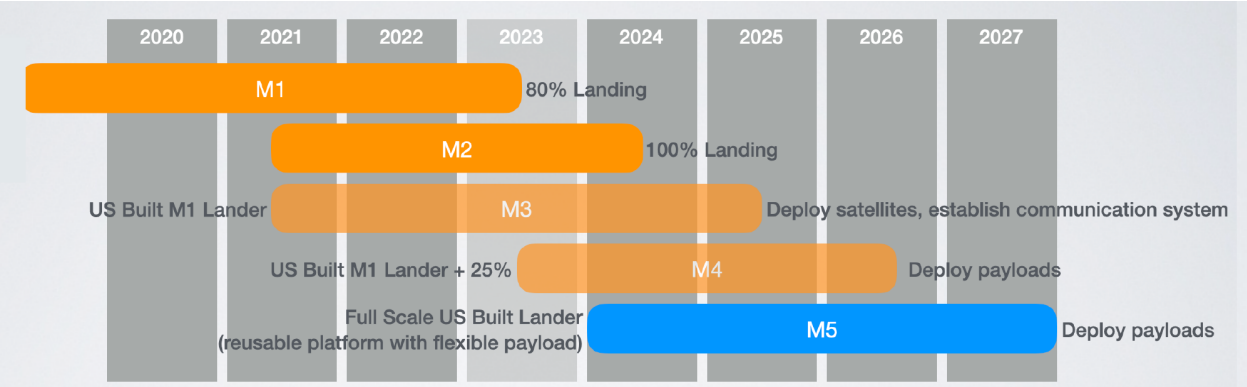

An incremental path: prove landing and communications with a smaller vehicle, then scale capability and payload across later increments, distributing risk instead of concentrating it.

Integrating the suppliers was three-quarters of the problem. A single schedule, refreshed in one room every week, was the answer.

lateralworks engagement review

05

Results
What the engagement put in place

By the close of the engagement, the program had the machinery to converge, even though it had not yet converged. What changed was concrete.

- The integrated master schedule was restructured and was being used to run the program, not to satisfy a reporting request.
- A new baseline and targets were established and accepted by the prime contractor as the forward plan.
- The schedule gap was known and quantified, and trends to each target were tracked week to week.
- Key supplier schedules were integrated into the IMS, with suppliers helping build it for the first time.
- The weekly Refresh was running, with the schedule updated from live team inputs.
- Core scheduling skills were transferred to the program's own managers.
- A starting FTTM core team was in place, and the engagement had fed concrete inputs into the company's organizational design.

The team had also begun to feel ownership of the schedule rather than treating it as someone else's number. That shift, from absorbing slips to arguing about how to recover them, was the precondition for everything else.

The engagement was equally clear about what these gains were not. Supplier participation had started but was far from complete. The core team was a start, not a finished structure. The schedule was visible and managed, but the gap was still open. The system to close it existed. Using it was now the company's work.

06

The open items
What remained to improve

The engagement left a set of recommendations that were less about scheduling mechanics than about the conditions a converging program needs.

Organization came first. The program needed a heavyweight program manager with a deputy who managed suppliers directly, freeing procurement to focus on sourcing and contracts. It needed a single technical authority, a chief engineer, and a single system architect, rather than design decisions diffused across the team. And it needed senior, experienced talent added to raise the program's overall maturity.

Supplier integration had to go the rest of the way. The engagement was blunt that this was about three-quarters of the whole problem. Each supplier needed a direct line to the deputy PM and the technical leads, more detailed schedules built jointly and folded into the IMS, weekly status into the same Refresh, and visibility into the whole program rather than only its own slice of it.

Autonomy and focus were the harder asks. The recommendation was to accelerate toward an empowered set of operations that owned its program and its budget, escalating to corporate only when a deviation crossed a set threshold. Alongside it: act rather than wait for permission, change the default from acceptance to action, and resist taking on new programs that would pull people off work that was already under-resourced.

Two strategic questions were left deliberately open, because the program could not answer them and needed to. What is the payload, and therefore how large does the lander actually need to be? And should the mission be pursued incrementally, proving a smaller capability first, to distribute risk rather than concentrate it in one oversized step?

07

The takeaway

An integration model others can reuse

The aerospace specifics are particular to this program. The integration model is not.

Any hardware program split across multiple subcontractors faces the same structural problem this one did. The work is distributed, but the risk of it not coming together is held in one place. The instinct is to manage that risk with reporting: each supplier sends status, someone assembles it, and the assembly is mistaken for control. It is not. Status describes the past. It does not force the future to converge.

What forces convergence is a single integrated master schedule that every supplier's work actually lives inside, refreshed on a fixed weekly cadence in one room, with the critical path to each target re-identified every time, the cause of every change recorded, and the trend to each date tracked until slips have names and owners. The schedule stops being a document and becomes the place the program is run.

That is the reusable result from this engagement. The lander, the avionics, and the lunar target are specific. One schedule, refreshed every week, with suppliers inside it and trends tracked to every target, is how a program made of many companies becomes one program.

I can’t thank you enough for all the help you gave us. I felt the company was on a path toward disaster, and you made a huge course correction toward success.

— Chief executive officer, the client

A

Appendix — market context

The commercial race to the Moon

This engagement took place inside one of the hardest and youngest industries in modern aerospace. The context matters, because it explains why a year of schedule slip was not a planning inconvenience. It was an existential risk.

A wave of private companies is now attempting what only national programs once did: land working spacecraft on the Moon and operate them on the surface. The model is largely a creature of public policy. A national space agency buys delivery to the lunar surface from private vendors under fixed-price contracts, betting that competition will break the cost of access that once ran to roughly a billion dollars a mission. The early results show how unforgiving that bet is.

The technical wall. Landing autonomously on the Moon is brutally hard. There is no atmosphere to brake against, so a lander must shed orbital velocity — well over a kilometer per second — on its engines alone, throttling precisely enough to neither slam into the surface nor bounce off it. Earth is too far away to fly the final descent by remote control, so the vehicle has to sense hazards and choose a safe touchdown point on its own, in real time. The most scientifically valuable sites, near the south pole, are also the most treacherous, and more than one recent lander has touched down only to tip over on uneven ground. Those that land then face a thermal cycle that swings from roughly 120 degrees Celsius in daylight to below minus 130 at night, on a clock set by a fourteen-day lunar day that ends most unheated missions at the first sunset.

The public record is sobering. Across the first years of commercial attempts, more missions have failed than succeeded. Landers have been lost to propellant leaks before reaching the Moon, to descent-software and sensor faults in the final minutes, and to tip-overs that cut short otherwise successful touchdowns. The first fully successful commercial soft landing did not come until early 2025. Commercial teams face the same physics as government programs, but with smaller budgets, thinner staffing, and little heritage hardware, and most are meeting these problems for the first time.

The financial wall. The economics are as punishing as the physics. Building, testing, and flying lunar hardware costs hundreds of millions of dollars, with any return pushed years into the future, which leaves companies crossing a long “valley of death” on raised capital before a market exists. The fixed-price contracts that opened the field also transfer cost and schedule risk onto small, relatively young vendors. The sponsoring agency’s own inspector general found that across early commercial lunar task orders, costs rose by more than two hundred million dollars and missions slipped by an average of well over a year, with the real time from award to launch running closer to forty-four months than the thirty originally planned. That pressure has already pushed one awardee into bankruptcy and driven consolidation among the survivors.

Underneath all of it sits a single dominant customer. Government demand still drives the lunar economy, and it arrives one mission at a time, which makes revenue lumpy and forecasting hard; a vendor's survival can hinge on the timing of the next task order. In that environment, schedule is not a reporting metric. Every month of slip burns fixed capital against a fixed price, and on a program already carrying a negative margin, slippage compounds directly into financial exposure.

This is the backdrop to the engagement. A program running about a year behind its first customer commitment, with critical work spread across international suppliers and a negative margin, was not facing an ordinary planning problem. It was facing the central risk of its industry — that the schedule, and the money, run out before the vehicle is ready. Bringing the schedule under control was, in the most direct sense, a matter of the company's survival.

B

Appendix — the system in practice
Anonymized scheduling artifacts

The figures below show the actual scheduling system the program ran each week, with identifying details removed. They are the anonymized form of the machinery described in the approach and the results: one integrated schedule, reviewed on a fixed cadence, with the critical path to every target re-identified and the cause of every change recorded.

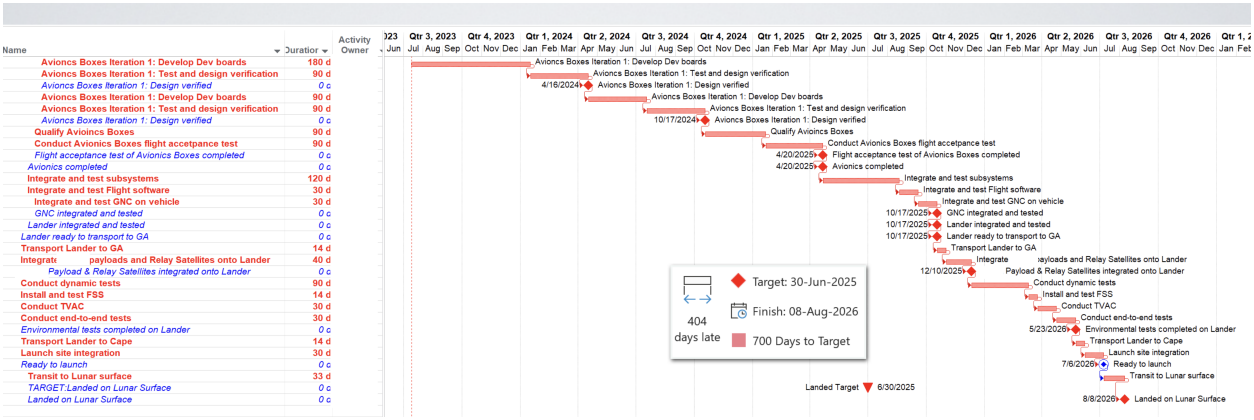

The first critical-path study. Read end to end, the driving path put the landing finish more than four hundred days past the first customer target — the gap the engagement set out to close. The number got worse before it got better: each weekly refresh re-identified this path and opened pull-in actions against it, and the trend to the target was driven back over the months that followed.

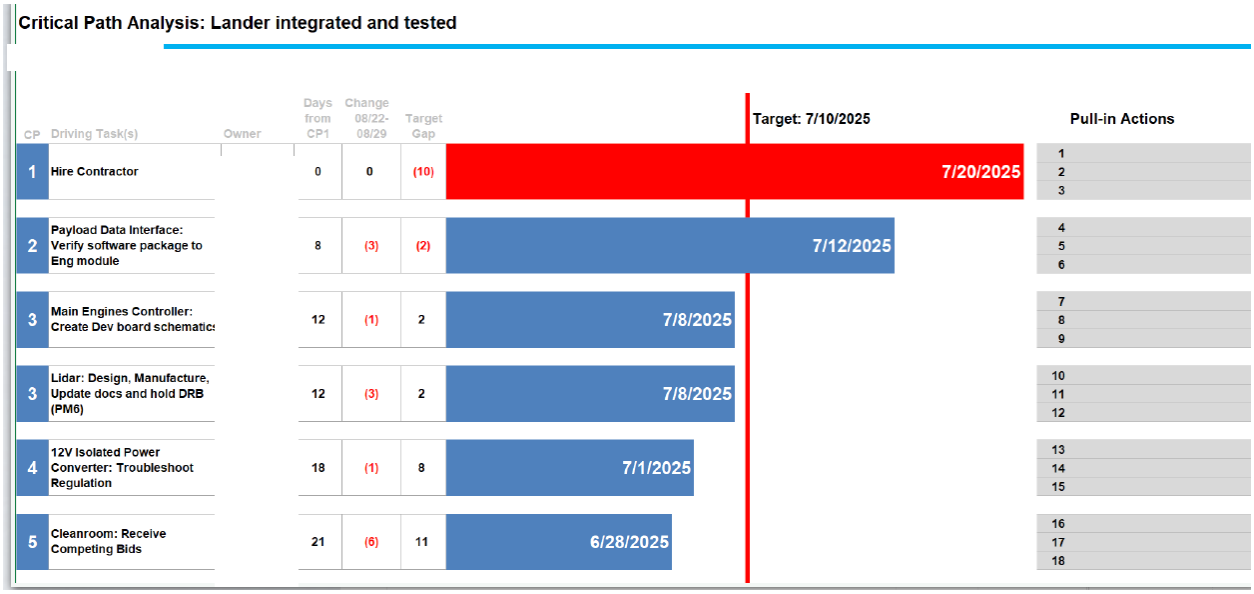

Critical path to a target, ranked. Each driving task is shown with its days from the critical path, its change since the last update, and its gap to the target date, alongside the pull-in actions opened to recover it. The same analysis was produced for every target the program tracked.

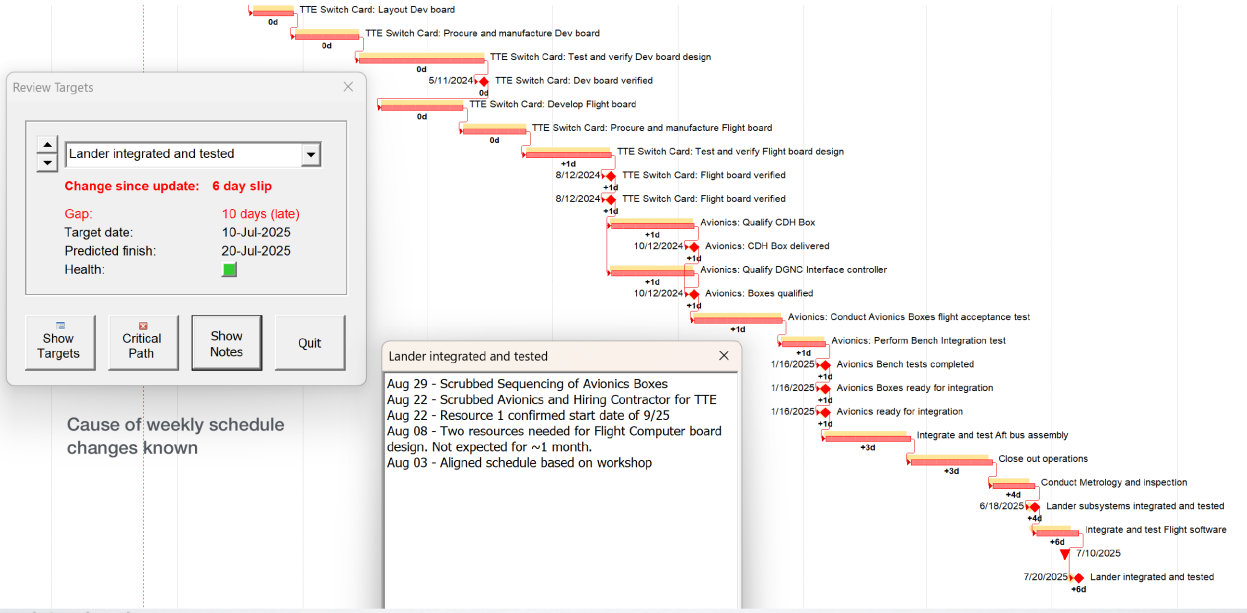

The weekly review in use. Each target carries a quantified gap, a health state, and a predicted finish (left), and every change to the schedule is logged with its cause (lower right), so movement on a date is never anonymous.

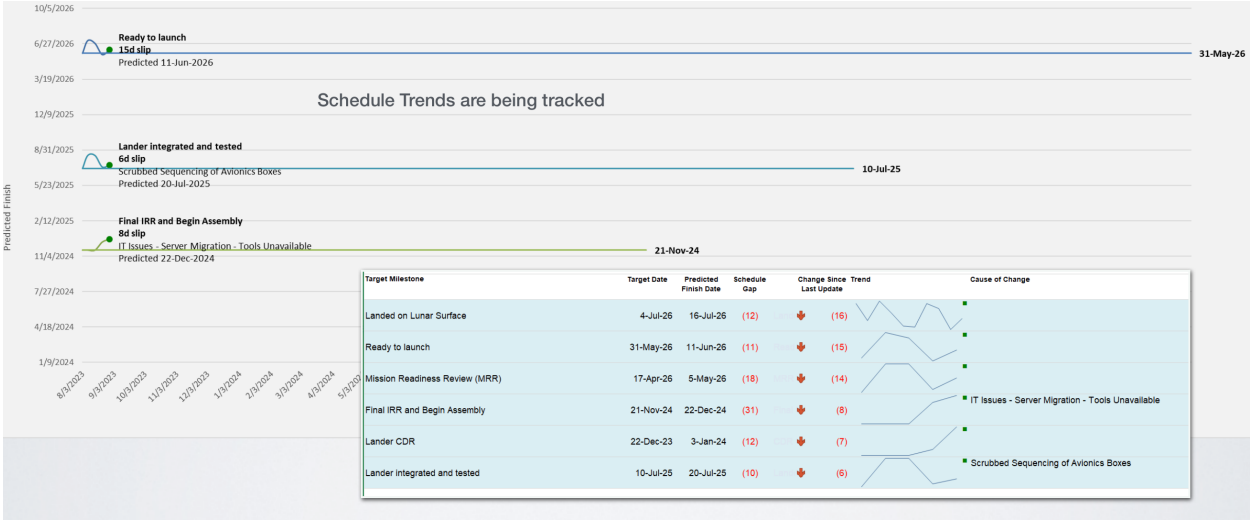

Predicted finish for each target, tracked over time against its committed date. Read together, the trends show not just where each date stands but which way it is moving, and why — the signal the program used to drive recovery.

Sources

• U.S. government inspector-general audit of commercial lunar delivery services (2024–2025): cost growth exceeding two hundred million dollars across early task orders, average schedule slips of more than a year, and award-to-launch timelines closer to forty-four months than the thirty originally planned.

• The Conversation, on the technical difficulty of autonomous lunar landing (2025): airless high-velocity descent, real-time hazard avoidance, south-pole terrain hazards, and the thermal swing of the fourteen-day lunar day.

• IEEE Spectrum, Payload, and NASASpaceflight reporting on the commercial lunar landing record (2024–2025): the first fully successful commercial soft landing in early 2025, alongside missions lost to propellant, sensor, or descent-software faults or to post-touchdown tip-overs.

• New Space Economy and related industry analyses (2022–2025): the capital intensity and “valley of death” facing commercial lunar ventures, a fixed-price awardee’s bankruptcy, and the consolidation that followed.