# A Vision for Humanity's Future in Space

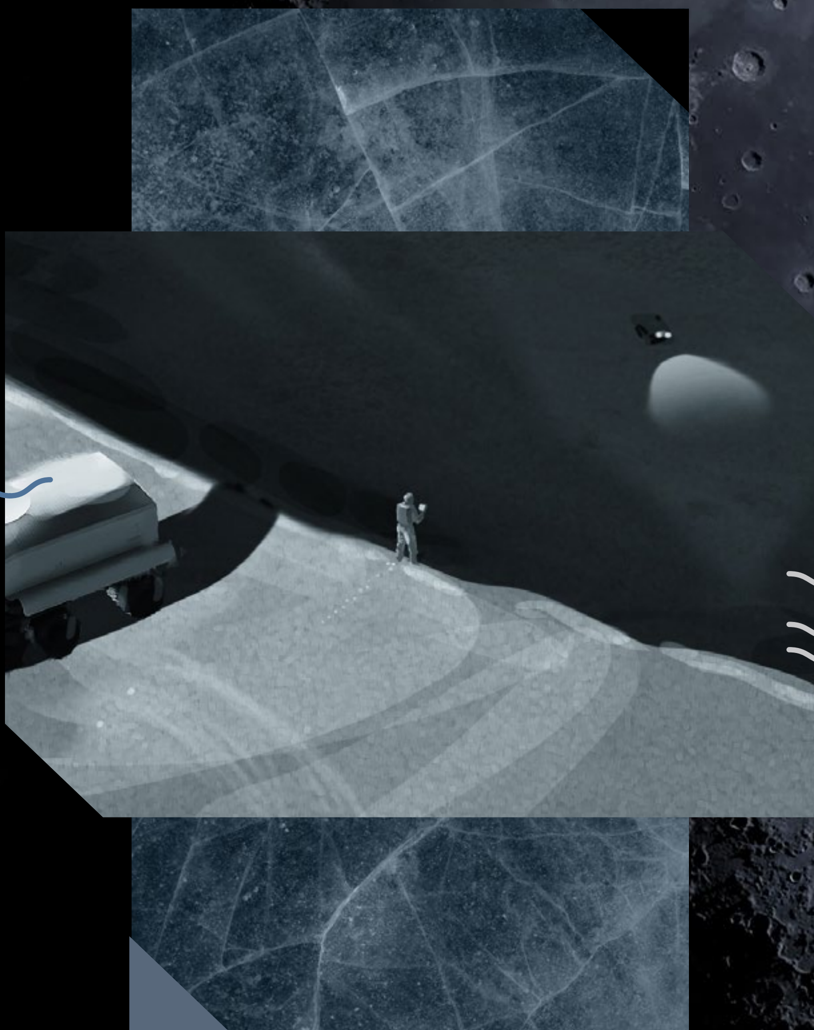

Lockheed Martin's
Water-Based
Lunar Architecture

LOCKHEED MARTIN

# Humanity is going to the Moon and on to Mars.

NASA has established a credible, well-constructed Moon to Mars architecture. With the successful launch of Artemis I, the first phase of that plan is underway. Through international agreements like the Artemis Accords, space agencies and commercial companies around the world are working on projects supporting this architecture, from the Orion crew vehicle and spacesuits to lunar landers and surface habitats.

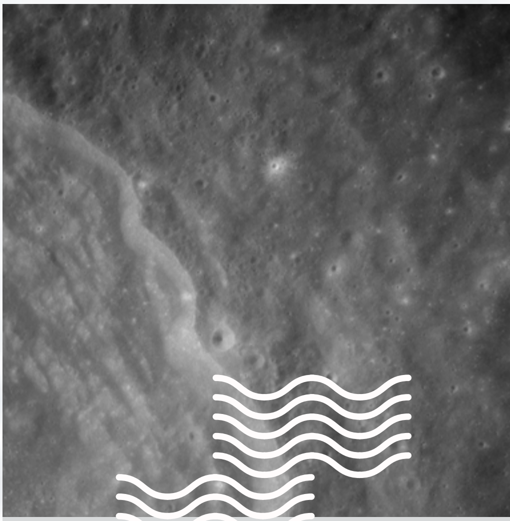

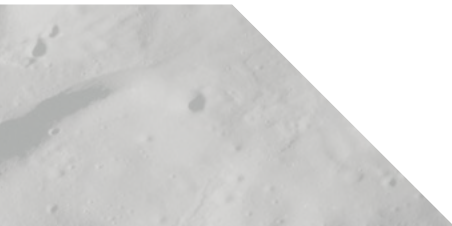

**This is a vision for humanity's water-based, nuclear-enabled future in space.**

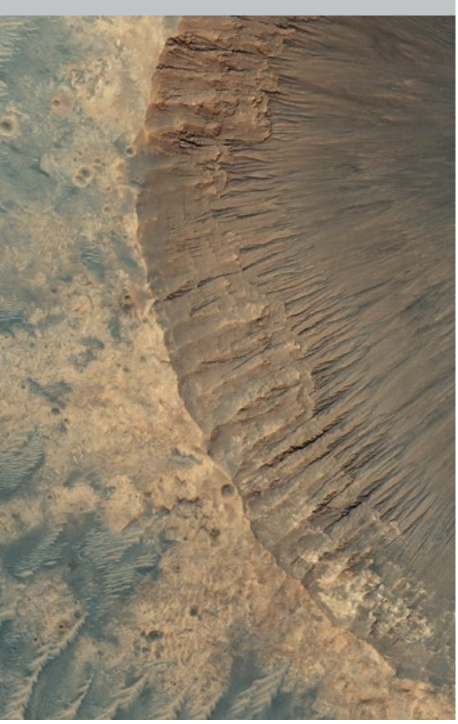

An aerospace leader since 1912, Lockheed Martin has a unique history of robotic and human space exploration. In this document, we describe a vision for how cislunar development could play out into the 2040s.

It is imperative that such a vision be sustainable—technically, politically, and economically. According to this guiding principle, the best lunar architecture over the long term is one that is water-based, nuclear-enabled, and commercially invested.

This technically credible vision is grounded in the physics of rocketry, humanity's scientific knowledge of the solar system, and the principles of engineering complex space hardware. It leverages insights gained from our company's long experience of spaceflight success, including having built more interplanetary spacecraft than all other U.S. companies combined. It also synthesizes emerging results from cutting edge research programs, architectural studies, and new technologies and platforms that we are developing and investing in ourselves.

Though its authorship by Lockheed Martin lends it an authentic sense of realism, this vision is not focused on our company. Nor does it belong to any one company or nation. This is a reference vision for all nations, all companies, all partners—a snapshot of the future that all of humankind can, and should, strive to achieve.

With humanity embarked on the first steps of the Moon to Mars journey, now is the right time to share this vision. It will take different forms for different audiences, from websites to interactive conference presentations to technical papers.

The novella white paper is intended to immerse you in this vision by placing you in the driver's seat. Walk a mile in the dusty boots of a commercial lunarcontractor. Kick the tires of your rover. Explore the construction site of a lunar megaproject. Reflect on how and why lunar exploration will have developed over the next two decades to achieve this vision. Imagine how you might fit into this rugged lunar community. What is life like at the very edge of the human experience?

Welcome to Artemis Base Camp. We hope you enjoy this glimpse into humanity's future.

## The Moon in the 2040s: A Vision of a Developing Human Frontier

Settling back into the driver's seat, you punch in the final coordinates into the display. The electric motor hums quietly to life and you're off again. The stark black and white vistas of the lunar surface roll quietly by as your rover ferries you to the last site on the list.

The drive between maintenance locations is by far the easiest part of the job, and maybe the most enjoyable. Long, quiet trips across this magnificent desolation are a great way to spend a good chunk of your six-month contract. After all, if it wasn't the best seat in the house, the recruiting brochure wouldn't have featured so many beautiful snapshots from the front seat of the rover.

These moments of rest between service locations are a well-earned reprieve. You've got a busy job as one of just a few dozen humans keeping the whole Moon running. By the 2040s, there is a thriving international community on Earth's nearest celestial neighbor. This is because the Moon is the perfect testing ground and resource hub to support humanity's expansion into the solar system.

Here, groundbreaking science furthers the pursuit of knowledge and promises answers to the most fundamental questions, like who we are, where we came from, and what might be in store. The Moon's alien beauty masks a hostile environment of harsh radiation, vacuum, extreme temperatures, and abrasive dust. It turns out that the life-or-death challenge of lunar survival makes our nearest celestial neighbor an excellent environment to test out hardware and procedures for Mars, an even more challenging, remote, and unforgiving destination.

The Moon is an important source of material wealth outside Earth's gravity well. Here, we can develop the resources we need to operate in deep space more sustainably, helping us break free from the long and expensive logistics train of Earth launch. We have even reached the point where lunar resource development is beginning to flow value back to Earth orbit. For all these reasons, it's only natural that there exists a growing community here as humanity takes its first steps out of its Earthly cradle.

Your fellow lunar travelers are a diverse bunch. Many are scientist-astronauts from the world's various space agencies. Intelligent, highly trained, and level-headed, they are out on the frontier performing trailblazing science. There are also commercial contractors like you, working for companies that provide and maintain cutting edge infrastructure that enable those science missions and other, more routine lunar operations.

You work for Aquarius, one of the logistics support companies that keep the Moon's infrastructure humming. Its customers comprise organizations across every segment of lunar activity: in-space transportation, communications, navigation, surface mobility, habitation, power,

**Earth is the cradle of humanity, but one cannot live in a cradle forever.**

Konstantin Tsiolkovsky (1857–1935)



resource production and distribution, and construction. These systems help turn the barren surface into a more hospitable ecosystem. It’s your responsibility to know each piece of this infrastructure backwards and forwards: how it works, how it can break, and how to fix it if it does. And boy, is there a lot of infrastructure.

Some of this infrastructure is positioned on orbit above the Moon, in a near-rectilinear halo orbit, or NRHO, and at the first and second Lagrange points. This infrastructure includes NASA's lunar Gateway, commercial habitats, a large cryogenic propellant refinery and depot, and more recently, a staging area for lunar-derived metals and construction materials on their way downwell to Earth. These orbiting bases represent the entry ramp of an interplanetary highway that reaches all the way to Martian orbit. Interplanetary transit vehicles arrive here from the Moon every 26 months, refueling their enormous hydrogen tanks and taking on crew and cargo for the next Martian mission. In orbit above Mars is a small depot, powered by sunlight, where supplies of water are turned into propellants. This Martian gas station allows for landers to refuel between missions to and from the Martian surface where, after a half-century of robotic exploration, the beginnings of a small permanent base are being established.

By far the busiest hub remains the lunar south pole. You remember the view from the lander when you arrived, burning its way down the long arc of the landing trajectory. How the blackness of space gave way to a blinding white horizon as the lander leveled out and the south polar surface raced up to meet you. There, spread out in all its glory, was the triumphal sprawl of the Artemis Base Camp. The lunar south pole in the 2040s is a stunning vision of the human frontier in development.

The cluster of landing pads and spaceport infrastructure where you arrived are a bustle of activity. Some landers are in the middle of being loaded or unloaded by support robots, while others are in the process of being refueled for their return to orbit. The landing pads lie on the outskirts of the Artemis Base Camp, far enough away where there’s less of a risk of being impacted by any ejecta that gets blown past the pads’ protective berms. Rovers dot the scene, crisscrossing the Moon from pole to equator. A community of inflatable habitats and workstations occupies the center. An ever-expanding power grid of fission reactors and solar farms provides lifegiving energy to all of this infrastructure.

Over towards the lip of a yawning crater, resource production and storage facilities keep the growing lunar community alive with water and life support gases, while also supplying the clean, high-efficiency fuels that sustain trade routes linking the Earth-Moon-Mars system. These trade routes are the paths that travelers and uncrewed systems use to carry goods from one place to another, where those goods will have a higher value based on their usefulness and scarcity at the destination. But as a commercial contractor, you know that a trade route is more than just a way to get from point A to B.

Fundamentally, a trade route is a trajectory whose function is repeatable profit. That trajectory becomes commercially viable when the delivered goods return more value than the journey cost. It becomes commercially sustainable when that returned value is repeatable, so that the trader can keep expecting value after they make the delivery a hundred times. Trade routes turn a single commercial "mission success" into a sustainable architecture.

A good trade route finds the path of least resistance between the source and destination. A better one creates a cycle of trade, where the return journey can profitably carry new, unique, valuable goods from the destination back to the route's origin.

By the 2040s, this is how the inner solar system's commercial transportation system has evolved to work. It stretches from Earth orbit to Mars with its hub at the Moon. This path is fueled by water-based resources, powered by safe and reliable fission, and buoyed by commercial-government collaboration that keeps the value of a clean economy flowing back to Earth. This path is paved with infrastructure, and it's maintained by contractors like you.

## A Water-Based Lunar Economy

The inner solar system runs on water. For drinking, radiation shielding, a universal solvent, a source of breathable life support gases—water's many uses make it the inner solar system's resource of choice. And it's ubiquitous. We tend to think of space as dry and barren compared to the oceanic blue marble we call home. But the inner solar system is replete with ice. It can be found in significant quantities in

**The lunar south pole in the 2040s is a stunning vision of the human frontier in development.**

lunar craters, frozen into asteroids, and even in huge quantities at the Martian poles and in subsurface glaciers.

And then there's the most critical family of commodities: water-based propellants. Some of the long-haul asteroid mining spacecraft use water in its original form, simply heating it into steam for the journey back to Earth orbit. On the lower-efficiency end, hydrogen and oxygen gases and storable high-test peroxide make for readily available attitude control or main engine propellants. But you're a production infrastructure contractor. To you, these are mostly derivative markets. By fraction of sales, and by the number of hours you spend poring over the relevant diagnostics and production data, the most important commodities are liquid hydrogen and oxygen. These cryogens—LH2 and LOX—are the primary products of your company's systems and the main commercial reason you're here.

As anyone familiar with Tsiolkovsky's rocket equation knows, one of the most important factors in space transportation is the choice of propellant. According to physics and chemistry, that choice is clear. Liquid hydrogen and oxygen constitute the most efficient chemical rocket propellants we know of. What's more, when you put liquid hydrogen through a nuclear thermal propulsion engine, the efficiency doubles.

The bottom line for the company's mission planners was that hydrogen and oxygen offered the smallest propellant mass of any chemical combination required to move a given payload

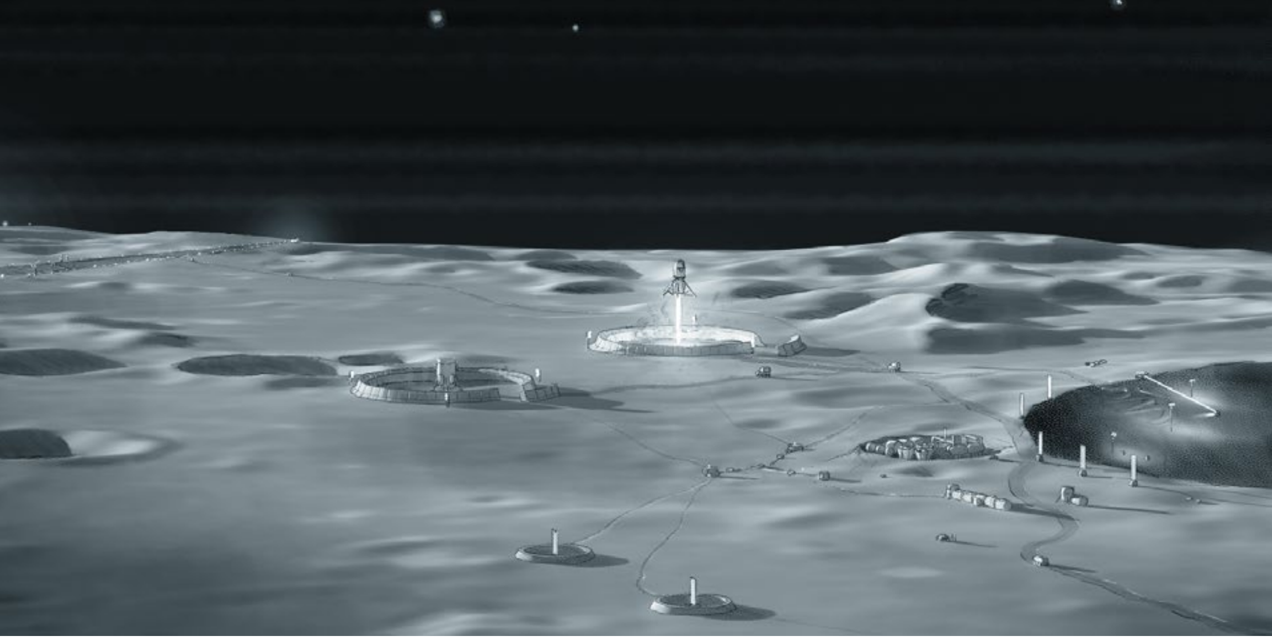

mass through space. Most of the mass of an in-space transportation vehicle, like a rocket's second stage or a dedicated space tug, is in its propellant, so for the same payload, lighter LH2-LOX upper stages also required smaller first-stage rockets.

This advantage increases when you launch empty and refuel on orbit. Eventually, performing that refueling using propellants already outside Earth's gravity well, produced from lunar water, avoids the Earth launch costs of delivering that propellant altogether.

This is not to say the approach was cheap. Hydrogen storage is tricky, as you know well. It has an extremely low density—around one fourteenth that of water or liquid oxygen—which requires enormous storage vessels. It requires sophisticated fluid management technologies for long-term storage, boiloff minimization, and transfer from one vehicle to another. Such systems are expensive. The excavation and purification systems required to produce water from icy regolith are challenging and power-hungry, as are the electrolysis and liquefaction systems that turn that water into cryogens. None of these technologies existed in a lunar-ready state at the start of the Artemis program. They had to evolve from their terrestrial versions, and their complexity and difficulty meant significant non-recurrent engineering costs.

A refueling economy must be built on sound business principles. Companies looking to develop cryogenic storage technologies in the 2020s had to ask themselves—would the advantages of a refueling economy built on water-based propellants be worth the significant up-front investment?

Over the short term, no. Frankly, the lowest-cost lunar exploration architecture looks something like Apollo, where a primary goal was to temporarily visit the Moon as soon as possible and just a handful of times, using minimalistic, expendable vehicles slimmed down to launch on a single rocket.

The Artemis program is far broader in its scope than Apollo. Through Artemis, humanity was going to the Moon to stay. We were heading there to put down roots, and to extend ourselves further into the solar system, for good. And unlike Apollo, Artemis does not benefit from the most generous budgeting NASA has ever received, or ever will. We needed a solution that was the most cost-effective over the long term.

Back in the '20s, your company figured out that, on a time scale of decades, the key to financially sustainable space development was to lower the barriers to entry for lunar access while reducing recurring costs. This meant establishing cheap and routine lunar transportation—as cheap as space travel could be, anyway. The biggest bang for their buck was going to come from bringing down recurring costs through reusability. And the most critical part of reusability is refuelability.

With enough missions headed to the Moon, there emerged a clear point in time beyond which the cost curve of a refuelable architecture grew slower than that of a direct, disposable one. That's because launch is still the highest unavoidable cost, even now in the 2040s, so the best savings to the customer come from reducing non-payload mass. Because the systems are refuelable, they can perform dozens of missions before the end of their life, for just one launch cost.

In addition to refuelability, Aquarius' early founders also understood the most cost-effective architectures were based on supply chains that could one day become Earth-independent. The solution also had to support Earth's clean-energy future. Because hydrogen-oxygen propellants are made from water, and because water is found just about everywhere in the inner solar system, LH2 and LOX sourced in space offered a pathway for weening the lunar economy off costly Earth-launched logistics.

Now, after a decade of large-scale LH2 transport, storage, and refueling under its belt, your company understands how to do it well and repeatably. Its systems are mature and many of the known risks are figured out.

High-performance cryocooling, zero boiloff, and no-vent fill are just some of the clever techniques from a long list of cryo fluid management know-how, many of which your company has pioneered over the last 15 years. It offers logistics companies the highest proportion of payload mass to vehicle weight possible thanks to the performance and efficiency of nuclear-hydrogen engines.

With the developmental challenges of hydrogen management solved, water-based propellant architectures remain much more attractive than methane-based alternatives. Water is found throughout the inner solar system. Methane, however, can be sourced from only two locations: the surface of Earth or the surface of Mars, at the bottoms of two deep gravity wells many millions of miles apart, with no other viable propellant source in between. Martian in-situ resource utilization (ISRU) isn't slated to become commercial-scale for a long time. This makes for an overextended supply chain, and one that is unforgiving in the event of an emergency.

Even when Martian ISRU is mature, you have to imagine it will still be more productive to make water-based propellants from ice caps and subterranean ice sheets than trying to wring methane out of Mars' notoriously thin atmosphere. You can see why the cost curve for your company's hydrogen-based architecture has flattened more optimally since the late 2030s and 2040s compared to methane.

That is not to say these architectures are incompatible. Oxygen makes up the majority of the reactive mass of methane-burning systems. It can be sourced in a few ways, including from the very same water molecules that give your company its hydrogen. In fact, LOX-LH2 production naturally generates a surplus of unusable oxygen thanks to a difference in ratios of what is electrolyzed from water versus what is burned by LOX-LH2 engines. This surplus can be sold to operators of methane-based architectures.

As for liquid hydrogen, there is very little being launched from Earth these days. Each year hundreds of tons of propellant are made locally in space, outside of Earth's deep gravity well, avoiding the high transportation costs and environmental impacts of launching all that propellant up through Earth's atmosphere.

These savings, achieved over the long term and enabled by space resources utilization, are thanks to the infrastructure you maintain. So many moving parts, and so many things that need to survive for decades in the harsh lunar environment.

The orbital gas stations that keep the interplanetary superhighway moving at low cost and high reliability are supplied by an enormous and complex pipeline stretching from the bottom of craters here on the Moon, all the way out to the depot at Mars, and back to Earth orbit.

This pipeline pulls ice out of the ground, turns it into purified water, cracks that water into life support, and cryocools those gases down into the two most powerful chemical propellants in existence. Downstream of production are surface storage depots, launch infrastructure, on-orbit receiver or "catcher" elements, and the tugs and other transportation vehicles needed to move those commodities around the inner solar system's trade routes.

Pipeline can be a dirty word on Earth. It conjures images of pollution, exploitation, and armed conflict. Today, in space as on Earth, humanity is leveraging the technologies and operational experience gained from extractive industries for a cleaner future. With this hydrogen system—the backbone of a new energy economy in space—humanity has abandoned hydrocarbons for lifegiving water to make powerful, clean-burning fuels.

**The Moon is a testing ground and resource hub to support humanity's expansion into the solar system.**

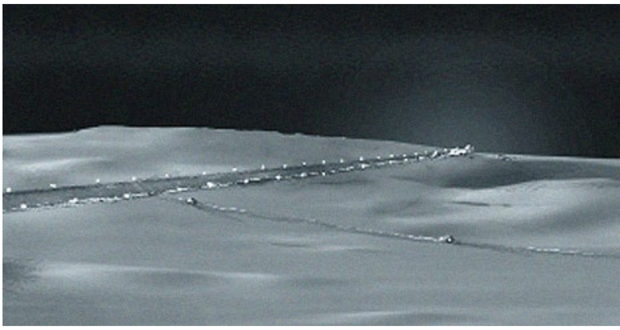

## Electromagnetic Launch

As impressive as the vast architectural web of this pipeline is when it works right, it is also daunting to consider all the things that can go wrong. This is *your* pipeline for these six months, you reflect, the responsibility sobering you as you pull up the master diagram on the dash. Scrolling through it, you zoom in on the highlighted piece of infrastructure that marks out the last location you'll visit on this sortie.

Schematics and timetables spill across the tempered glass. The rover goes dark for a moment as it passes through the long shadow of a boulder a few hundred meters away. The display splashes the cab in a calming blue. The digital representation of a construction site rotates slowly in front of you: one of the company's newest projects, poised to revolutionize the capability and efficiency of the entire cislunar transportation ecosystem.

It's an electromagnetic (EM) catapult, you remind yourself, almost with incredulity. As if you needed more proof this job was like something out of a Heinlein book. It's hard to believe that almost 80 years later we're finally making it a reality, and that you're the lucky contractor that gets to build it. Well, technically, the robots are bearing the brunt of the actual construction. And you'll be back on Earth years before the job is finished. But for these six months, you're the one on the ground keeping the bots running like clockwork.

When the project is finally finished, this launcher will use electromagnetic force to accelerate canisters of water, cryogens, or other resources down a long track until they reach orbital velocity. From there, they soar out over the surface of the south pole, gaining altitude until reaching apogee where a catcher element aggregates each canister. Getting resources from there to the depot waiting in the Gateway NRHO orbit happens over a whirling series of crisscrossing, looping trajectories and one small but well-timed burn to tip it over into the right orbit. You remember the mission designer who found this orbital needle in the haystack. You once had the opportunity at the company's headquarters to hear their explanation of that work, with them throwing out terms like chaos theory and Poincaré manifolds. The orbital mechanics are black magic to you. But the promised results were clear enough.

The beauty of the EM launcher is that it allows for the next revolutionary increase in system-wide efficiency for delivering mass to lunar orbit. You run through the production numbers in your head: annual demand; delivery and replenishment timetables; thermal cycling and boiloff rates; each month's launched liquid inventory numbers; production rates; percentage ice concentrations by weight. You know it all by heart. The numbers are pretty good. By this point, these systems are just about as tightly optimized as the company can design them. You've met the brilliant engineers who designed this architecture, and if anyone can squeeze out another gram-per-kilowatt, it's them. But even they can't beat physics. Staring you in the face in the little display is that giant efficiency bottleneck: burning—wasting—most of that precious propellant just to launch to NRHO the water and fuel the customer needs.

Once it becomes operational, the EM launcher changes the paradigm. Instead of burning more propellant than what's delivered, commodities can be launched directly to orbit using electricity. The propellant saved will no longer need to be produced and stored, removing the majority of the power costs of the surface portion of the propellant pipeline. Instead, some of that power would be rerouted to the EM launcher. It represents a complete step up in terms of kilograms delivered per kilowatt consumed, system-wide. Revolutionary.

You think of how many more Mars Base Camp missions could be sent to the Red Planet. Or how much more construction material could be dropped into Earth orbit where tomorrow's orbital megastructures are taking place. How many Earth-orbiting solar power beaming stations could we assemble then? Would the cost of a prototype O'Neill cylinder finally drop within reach of the billionaire-philanthropists?

**A trade route is a trajectory whose function is repeatable profit.**

With today's energy markets the way they are, and climate remediation proceeding at its current pace, relief valves will be needed somewhere. Projects like these give you hope for a brighter future.

When you think of what other miracles such an increase in efficiency could deliver, you get the sense you've only just scratched the surface, given how much humanity has already accomplished with today's fleet of transportation vehicles. You feel your eyes drawn skyward through the rover's sunroof where, hundreds of kilometers above you, the cislunar superhighway is bustling with the activity of that fleet.

## Evolution of Cislunar Refueling

What was once an empty sky is now populated with artificial satellites. Communications and navigation satellites move in a slow-looping arc overhead. A fleet of cargo tugs and propellant tankers converge on the service depot in Gateway orbit. Commercial prospecting, mapping, and situational awareness monitors peer into the surface regolith from on high. Closer to the horizon, a thin strand of frozen steam catches the light, revealing the trajectory of a descending crew lander before its flickering engines disappear behind the next ridge.

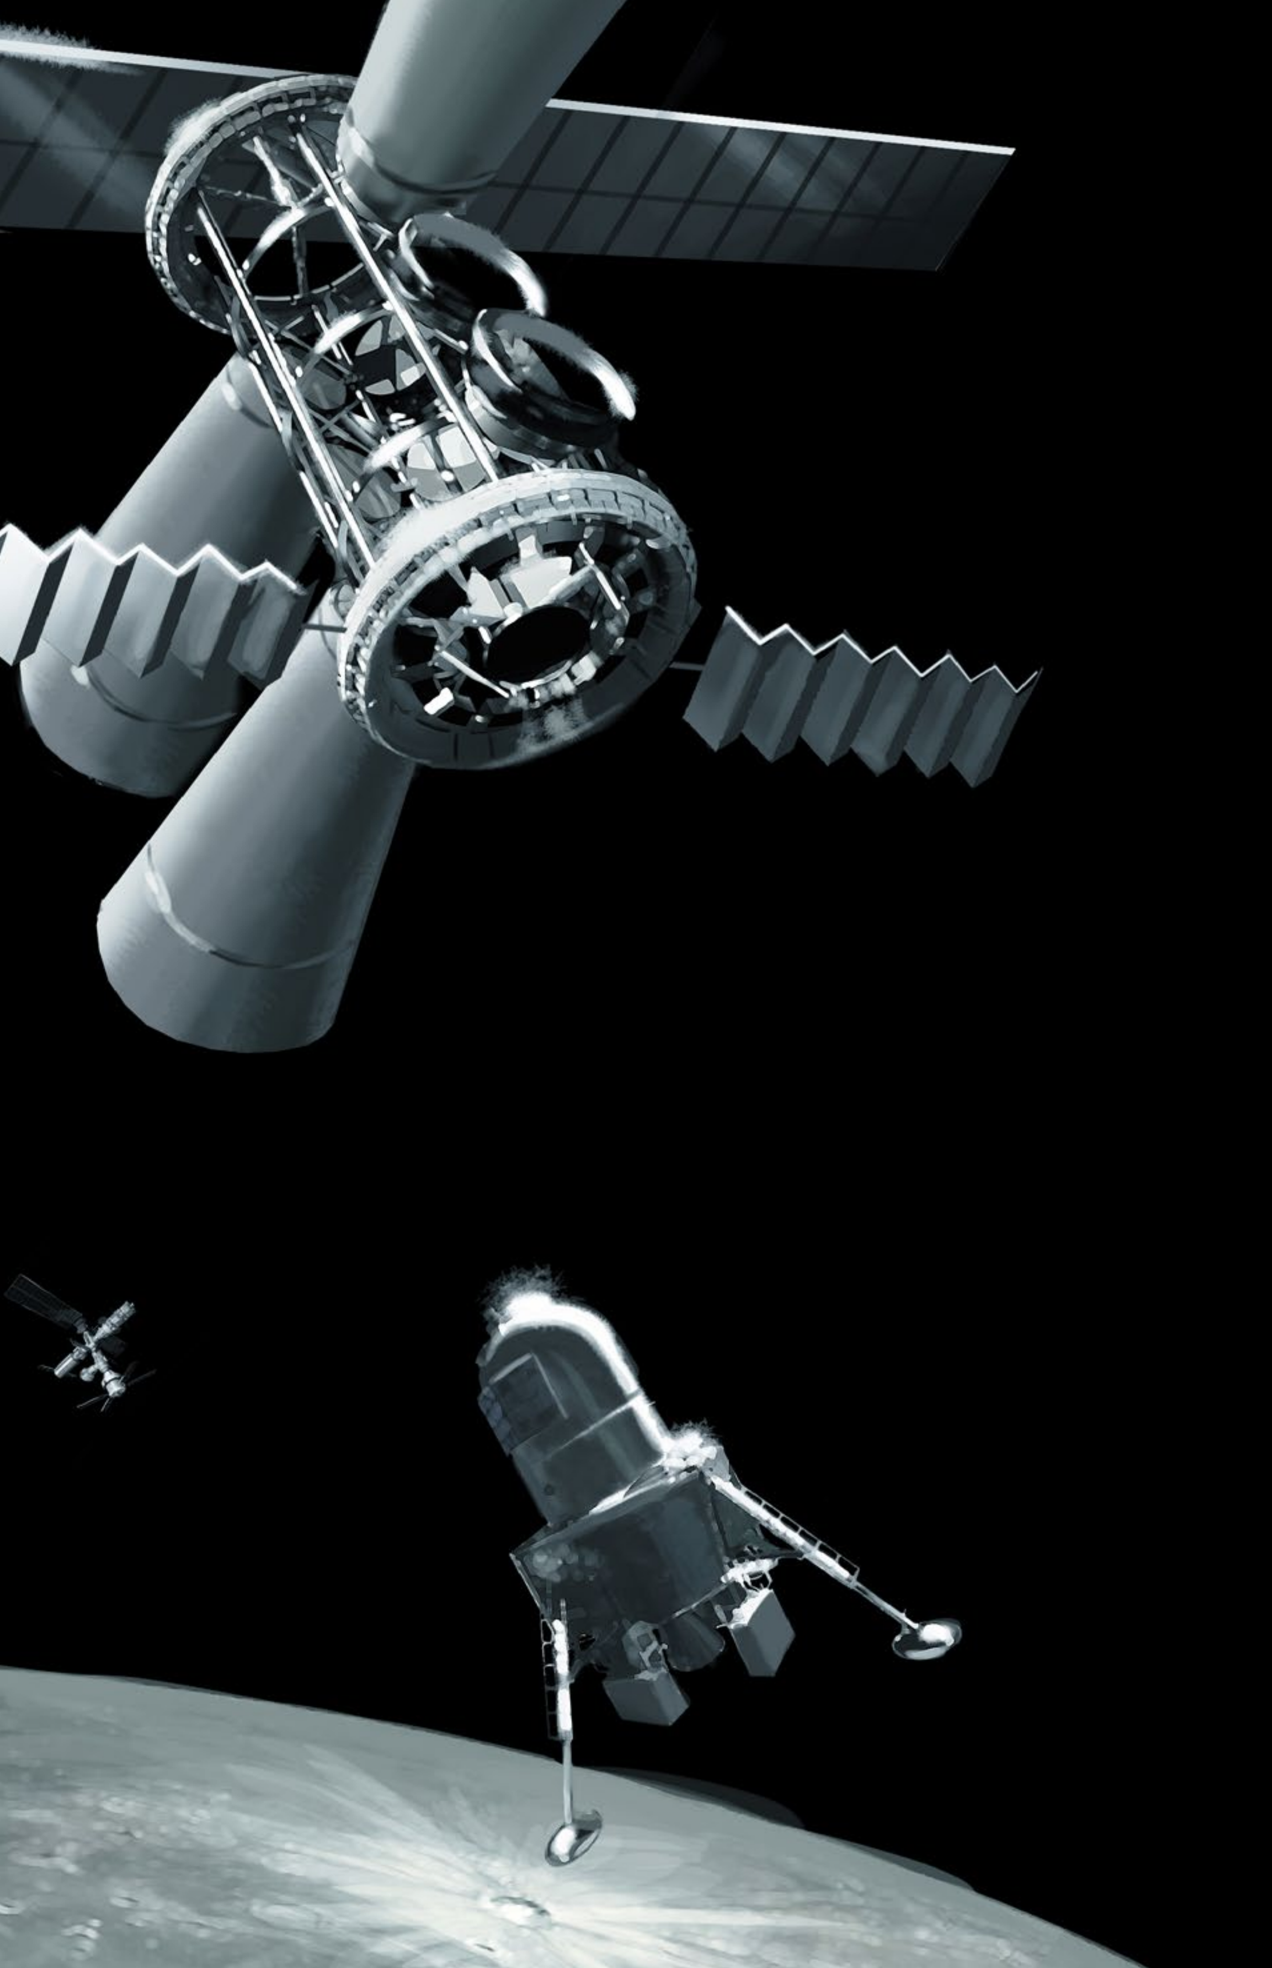

The Moon's south pole is busier than Earth's these days, but it wasn't long ago that NASA was only landing one mission per year. The U.S. didn't even field cryo depot demonstrations or refuelable LOX-LH2 delivery vehicles until the late 2020s. But things started coming together over the last decade. You remember how exciting the livestream was of the first-ever lunar hydrogen refueling, when a Cislunar Transporter tanked up the lunar lander for its sortie to the Moon—and then how quickly repeat performances became routine.

After enough missions, lunar lander refueling services almost became boring. But boring is good. For you, it means predictable operations. For the company, it means everything is going right. The risks have been proven out. Disasters have been avoided. The development costs are being steadily amortized.

With a higher lunar mission cadence and larger amounts of propellant sitting in space for longer durations, the next logical step for bringing down costs was to place a depot at the Moon. Hydrogen is a pain to keep cold for long durations. Early systems that had to perform dual roles of keeping propellant cryo-cooled and transporting it around, and were naturally unable to optimize for either. But the next generation of systems could focus on particular roles.

Fixed infrastructure elements like the depot are able to devote all of their mass and power to keeping propellants cold over long durations and pumping it to client vehicles in need of refueling. Separately, transporter systems are optimized for quickly and efficiently moving cryogens from source to depot with as little boiloff as possible. This offloads much of the penalties of long-term cryocooling dry mass, from insulation to backup cryocoolers, onto the depot, while keeping each transporter light and burning as little fuel as possible.

The depot was wisely designed to accommodate additional tanks and buses in a modular fashion as the scale and frequency of lunar


refueling increased. A couple dozen tons of hydrogen per year might have been sufficient for the early 2030s, when lunar landers and a nuclear thermal propulsion (NTP) tug provided only a small number of missions each year. Now that we're into the 2040s, there's one Mars Base Camp vehicle departing for Mars every 26 months. In gross tonnage, the depot must be sized to dispense as much as 230 tons of hydrogen at one time. Stockpiling that much hydrogen means a monthly quota of over a dozen tons each month, on top of roughly as much commercial propellant needed for everyday customers.

The numbers always hit you as a huge amount, but you remember what we're able to do with that much delta-v. Every 26 months, we cycle 150 tons of reusable dry mass to Mars: orbital outposts, Phobos-Deimos exploration kits, surface power systems, rovers, habitats, and small-scale water and oxygen production plants. A small fleet of atmospheric LOX- and LH2-burning landers ferries crew and their surface infrastructure down to Mars and brings them back to orbit with troves of science. With each return to orbit, the landers refuel at a preplaced orbital depot and ready themselves for the next sortie.

In the near future you expect Martian ISRU will come online at a scale sufficient to establish

**Martian exploration would be years behind if not for the high cadence and performance of NTP vehicles.**

a contingency propellant reserve. This allows for greater mission flexibility should the lander require an emergency fuel replenishment on the surface, or if a topoff would allow a greater treasure of samples. Until then, though, the orbital depot will continue to be nominally replenished by high-efficiency solar electric propulsion tankers, which carry water from the Moon that the depot electrolyzes into fuel for the landers.

Martian exploration would be years behind if not for the high cadence and performance of each supply mission. The choice of nuclear thermal propulsion and its water-based liquid hydrogen propellant for Mars transport has been a fundamental enabler.

Hydrogen-fueled NTP boasts a specific impulse double that of water-based oxygen and hydrogen—already the most efficient chemical combination and 150% better than methane-oxygen. Its power and efficiency provide a range of breakthrough benefits, from faster transit times and mission flexibility to delivering more mission hardware using less propellant and dry mass. NTP's significant thrust gets crew to Mars in less time than any other engine, meaning less radiation exposure and less mass devoted to crew consumables. As a spacefarer, you appreciate the greater abort and early-return capability NTP provides, increasing the range of options available to crew in the event of mission critical events, up to and including throwing on the brakes and bringing crew back to Earth in the event of a mid-mission self-rescue.

NTP also allows for simpler spacecraft design compared to other nuclear alternatives, with fewer failure modes and fewer subsystems and components to integrate. In the event of a Martian depot failure or other refueling emergency, an NTP tug can hurry a cargo of lander propellants out to Mars in record time. NTP burns are also shorter in duration. All these benefits increase reliability and safety for spacefarers like you.

Excluding NTP and large-scale hydrogen storage, the technologies required for this type of Mars transit vehicle were largely available back in the 2010s and early 2020s. They included NASA's Orion spacecraft, which for years was the only crew-rated vehicle capable of surviving the harsh environment of deep space and returning safely from those vast distances at incredibly high speeds. Today's Mars transit vehicle has two redundant modules, each composed of an Orion atop an NTP tug, that function as the vehicle's excursion vehicle, command deck, and escape pod. While there is a diversity of commercially provided, crew-rated vehicles, in-space transports, and landers, Orion has been the crew vessel that ensured the Artemis program's success and now functions as the common capstone of several safe, reliable, high-performance crew delivery and return architectures.

Thinking back to the lunar propellant depot, you know modularity has been as key to its design as it was for Mars Base Camp. This allowed storage capacity to be built up incrementally

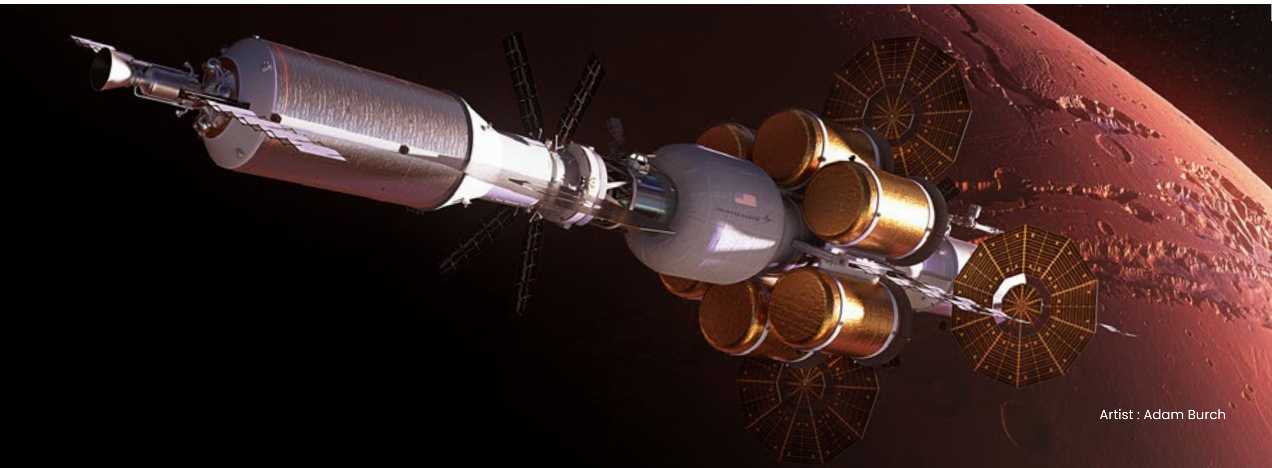

Artist : Adam Burch

and economically, from one cislunar tug per year to frequent interplanetary missions.

It was difficult to forecast exactly how quickly the cislunar refueling market would grow, as well as which types of propellant customers would buy beyond the first NASA contracts. Some prognosticators anticipated that cislunar refueling would be comprised entirely of LOX and methane. Indeed, the easier storage of methane-based systems over hydrogen has kept some companies flying LOX-methane tugs and landers through today. Many of those companies found value in prepositioning methane at the lunar depot. In other forecasts, hydrogen predominated. Still others were bearish on propellant storage at all.

In either case, there was uncertainty. It would have been risky to build one large, bespoke depot with tanks designed and sized for specific propellant types and amounts, based entirely on a guess of how the refueling needs of a developing lunar economy would play out years later.

Instead, a modular approach was the better choice. This gave the depot the ability to add new storage tanks and support buses as needed. Common standard interfaces allowed for tanks with a variety of sizes and designs from different manufacturers to be aggregated. If additional cryocooling power was needed, another power bus could be added. If more pressurants were needed, a gaseous bottle supply bus could be added. This offered a flexible approach. It allowed for the tankage of different propellant types and quantities to be supplied as needed according to the pace of demand of the refueling market as different propulsion systems emerged.

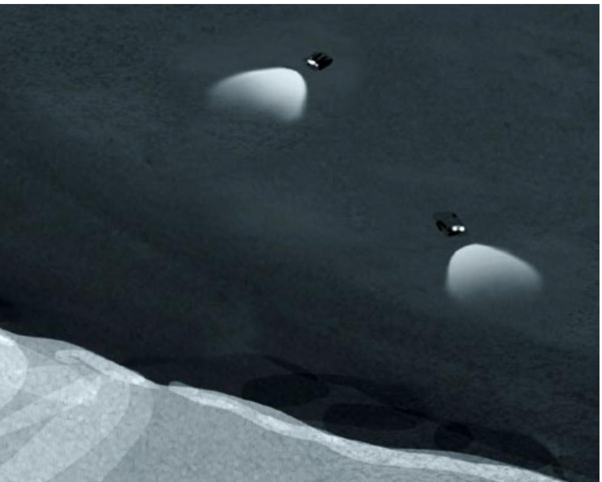

As refuelable lunar missions increased in frequency and the yearly demand in deliverable commodities increased, the value of local, just-in-time refueling increased accordingly. Thus the total storage capacity at the depot increased in scale. The demand was clearly there. But the supply still involved launching commodities upwell from Earth at a premium.

Around the same time as the first depot modules were emplaced, the potential economic viability of lunar-sourced propellant was becoming a reality. Data coming out of south pole prospecting missions in the 2020s and early 2030s confirmed the presence and economic extractability of ice in the shallow subsurface. With enough borehole data of good quality, accurate block models of the lunar subsurface could be built that offered three-dimensional maps of ice concentration and depth. These models quantified how much water could be feasibly extracted, the number and cost of the systems required to excavate and produce it, and the revenue it would bring in. In other words, it became possible to precisely model the value of a mine's cash flow.

Equally important was the greater confidence in the presence and extractability of water provided by the data. In mineralogical terms, the resource companies were able to take what had merely been inferred resources and eventually reclassify them as proven reserves. This allowed them to put a dollar value on their claimed ice deposits and list them as assets against their balance sheet. It was this market capitalization, along with projected cash flows, that gave companies the data they needed to secure investment and kickstart the lunar water resource economy.

In the late 2030s, these factors began to converge. On-orbit refueling at the Moon was becoming routine and the demand for propellant delivery was growing. An economically viable source of propellant had been demonstrated on the lunar surface. For the first time, a clear opportunity emerged for those that could offer a lower-cost alternative to sourcing propellants from Earth.

That's when Aquarius was born. What finally closed the business case for the investors was NASA's up-front guarantee, years in advance, to purchase the first 230-ton hydrogen load that would propel Mars Base Camp to the Red Planet and back.

You tip your hat to their courage and foresight—NASA, the company, its investors, and all the commercial partners involved in each stage of the ecosystem—for the role they played in creating this future. Without infrastructure on the surface, science missions would be shorter, less well equipped, and less frequent. Access to the lunar surface would still be reserved for a few astronauts for a few weeks each year.

## Remotely Operated Robotics

The rover starts to slow as you approach the construction site. Like a sculptor's block of marble, the long, gentle slope of a ridge provides the perfect natural foundation upon which to build a ramp to the sky. A flat launcher would be better, as it would reduce the delta-v corrections needed at apogee. But given the topography of this region, an incline was needed to clear the horizon. The bottom portion has already been reworked by the earthmoving robots, picked clean of boulders and leveled with uncanny precision to the nearest fraction of a degree. You marvel at the unnatural precision of the geometry. The smoothly sintered strip forms a polygon of jarring regularity against the rugged randomness of the rest of the natural slope, with all its pocks and blemishes and cratered shadows. Alien is the word that comes to mind, as out of place next to its surroundings as the obelisk from 2001: A Space Odyssey. But of course—humans are the alien intelligence here, reforming the surface of a foreign world through our robotic assistants and according to our programming.

Taking manual control, you maneuver the rover up to a tiny feature: a small, solitary hill, smooth and unremarkable. The only clue belying the presence of the fission reactor buried here is the communications mast protruding from the regolith and the red strobe on its control panel.

This reactor's job is to supply electricity to the robotic workers and construction rigs, but as construction progresses, you know that larger, more powerful and efficient reactors will be delivered nearby to create the launcher's power generation installation. Rather than using the chemical energy bound up in tanks of rocket fuel to achieve the delta-v of launch, the miracle of the EM launcher lies in its use of electricity: massless energy, delivered at the speed of electrons, powered by megawatt-scale fission reactors, and fine-tuned by the harnessing, converters, controllers, and megabatteries of a sophisticated lunar grid.

Today, the robots are preparing the system to connect the line from this reactor to the local hub. This is where the next reactor, which is due for delivery soon, will be plugged into the grid. The company will likely have you drive back here in person again to supervise that installation—from a safe location nearby, of course, until the delivery lander has departed and the reactor is buried. Today, though, you're on site as a "human in the loop," monitoring and approving the execution of each step of the process.

While these days the robots are increasingly capable of performing most operations autonomously, they could still make a mistake while deadheading the line, fitting the harness, or performing any other tricky operations with low tolerance for error. As on Earth, we've used advances in artificial intelligence to automate most of the mundane monitoring and controls for the lunar power grid.

You think that sometime soon, contractors standing by to telerobotically take control and execute the operations remotely will be cleared

to do so from a greater distance from the job site, perhaps from the comfort of a workstation at a surface habitat or even orbital station. But the Infrastructure Directorate hasn't yet approved the company for full autonomy for the most critical interfacing operations. Installing critical hardware like a new reactor fall into this bucket, as does today's task of preparing those interfaces and hardware for future integrations. So today, you're monitoring from the job site, commanding certain key sequences, and standing by to take remote control of each step, or even suit up and head out to haul cables yourself if needed.

You wait for your rover to establish a local wireless connection to the reactor's terminal. This will let you issue stop commands over Wi-Fi, just in case radio comms go down at precisely the wrong time. In another moment, you're connected to the robotic mechanics as well. With the session live, you tick through your manual operator's checklist. Everything looks good so far. Your suit is ready in case you need to step out and manually take over a task from one of the robots. You radio contractor communications and let them know you're green. They hand over control and you signal receipt. Here goes. You issue the command to execute.

One of your dashboard screens follows the robot's point of view, both in LiDAR point clouds and via the 4K camera. The dust kicked up by the robot's wheels billows in the robot's floodlight, morphing slowly in one-sixth gravity like mud clouds disturbed by a sea floor exploration capsule. Keeping an eye on the visuals, you cycle the other screen through various modes, including operator flowcharts and software packages executing line-by-line in real time. Checklists in hand, you approve the execution of each major operation and radio your every move to contractor comms.

Your company has had on-site contractors operating in the loop for a good number of years now, and every rotation advances the ability of the company, the Infrastructure Directorate, and humanity to perform tough space operations routinely and without error. This experience is essential as we build ever more ambitious projects, like the EM launcher under construction nearby, or the expanding branches of the power grid—themselves mostly self-sufficient microgrids—at key locations on the lunar surface.

## The Lunar Power Grid

The establishment of the grid is nearly legendary among your fellow contractors. Power is essential for every operation on the lunar surface, and it is critical to the launcher's success. But generated power comes in many forms: at different locations, times, qualities, and voltage and current levels. Though most suppliers have adopted 28V or 120V as standard voltages, their supplies are all at different currents and cleanliness qualities. Especially today, with the barriers to entry lower than ever, it seems like every month a new company launches their own take on the next best power generation system.

They vary by location, too. Near the poles, the mountaintops are dotted with solar masts tracking the eternal Sun as it circles the horizon. They contribute solar energy to the grid via high-voltage cables over long distances, or else through optical beaming systems where the terrain is too rugged or the distances are too long to lay heavy cable. In deep, freezing lunar craters, especially the large ones where beaming efficiencies drop off, fission reactors deployed on the crater floor perform best.

At the lower latitudes, including the science stations in the mare of the lunar nearside, the two-week lunar night requires either the constancy of fission or an energy storage system with two weeks capacity. Modern large-capacity batteries, developed for Earth's heavy industry decarbonization, can store enough energy for a small microgrid. The specific power of those storage products is fairly limited, however, and each additional megawatt-hour of total power capacity requires a linear increase in battery mass, meaning more hardware launched and delivered from Earth.

Above a certain energy capacity, it becomes more cost-effective to use systems where the reactive mass (like the energy-creating portion of a fuel cell) is relatively small compared to a much larger storage mass, like tanks of fluid. These systems include fuel cells, which take advantage of the cislunar water economy to supply their hydrogen and oxygen reactants. They also include emerging technologies like flow batteries, which have a landed reactor that works like an electrochemical battery, but with pipes connected to a tank farm constructed from local materials and filled with water and chemical reagents produced locally. In the case of fuel cells and flow batteries, adding additional storage capacity requires only the tankage and plumbing to be sent from Earth, while the reactive media—by far the largest portion of a system's mass—are made locally from water, other volatiles, and regolith.

The grid is what connects these disparate power sources, storage points, and loads across the Moon into self-sufficient networks. It rectifies the noise of different signals according to high-quality standards, carries them across hundreds of kilometers of wire from source to sink, beams them over rough terrain or down into craters, steps the power up and down to bridge those distances more efficiently, monitors where supply is generated and where loads are needed, shunts the current to storage devices when supply exceeds demand, and routes that stored energy to peak-usage terminals if demand ever exceeds supply. Most importantly for you, it guarantees your company can keep its production infrastructure fed around the clock with a smooth supply of high-quality power.

That said, this doesn't mean constant perfection for everyone, everywhere, at all times. The current ultra-high voltage (UHV) cable laying project should be complete in

the next couple years, linking the microgrids together into one singular network. By 2050, you expect the long distance UHV lines might reach all the way to the equator to link the south pole with the lunar equatorial solar farm. But until then, they need to function with relative independence. The south pole grid, such as it is, is in reality still just a segmented collection of smaller local microgrids constructed at the key locations on the south pole. Though each one has local power management and distribution capabilities, some are more reliable than others.

## Highly Reliable Nuclear Power

The power grid has become the heartbeat by which you and your fellow contractors measure the health of the lunar base. It's in our nature as humans, you believe, that when a community gathers in comfort at the end of a long day, we need a staple subject to talk about. For ages, that something has been the weather. You're pretty sure this truism reaches from the pub table back through prehistory to the campfires of hunter-gatherer tribes. It makes sense. Your back and brain get tired from a long day of work. You need something to talk about, and topics that work best are those that are familiar, easy to discuss, and widely experienced across the community. The weather is perfect. Everyone experiences it. The whole community is saddened by too much rain or uplifted by a sunny day. It's noteworthy when it changes, and it brings a comforting constancy when it doesn't. It can be an inconvenience, a disaster, or a delight.

But what does a community of astronauts and contractors talk about on the Moon, where all that changes about the sky is what direction the shadows come from this week? With the direct impact it has on every lunar activity, power has become that staple subject, mundane and ever-present, that underpins the mess hall's shift change conversation.

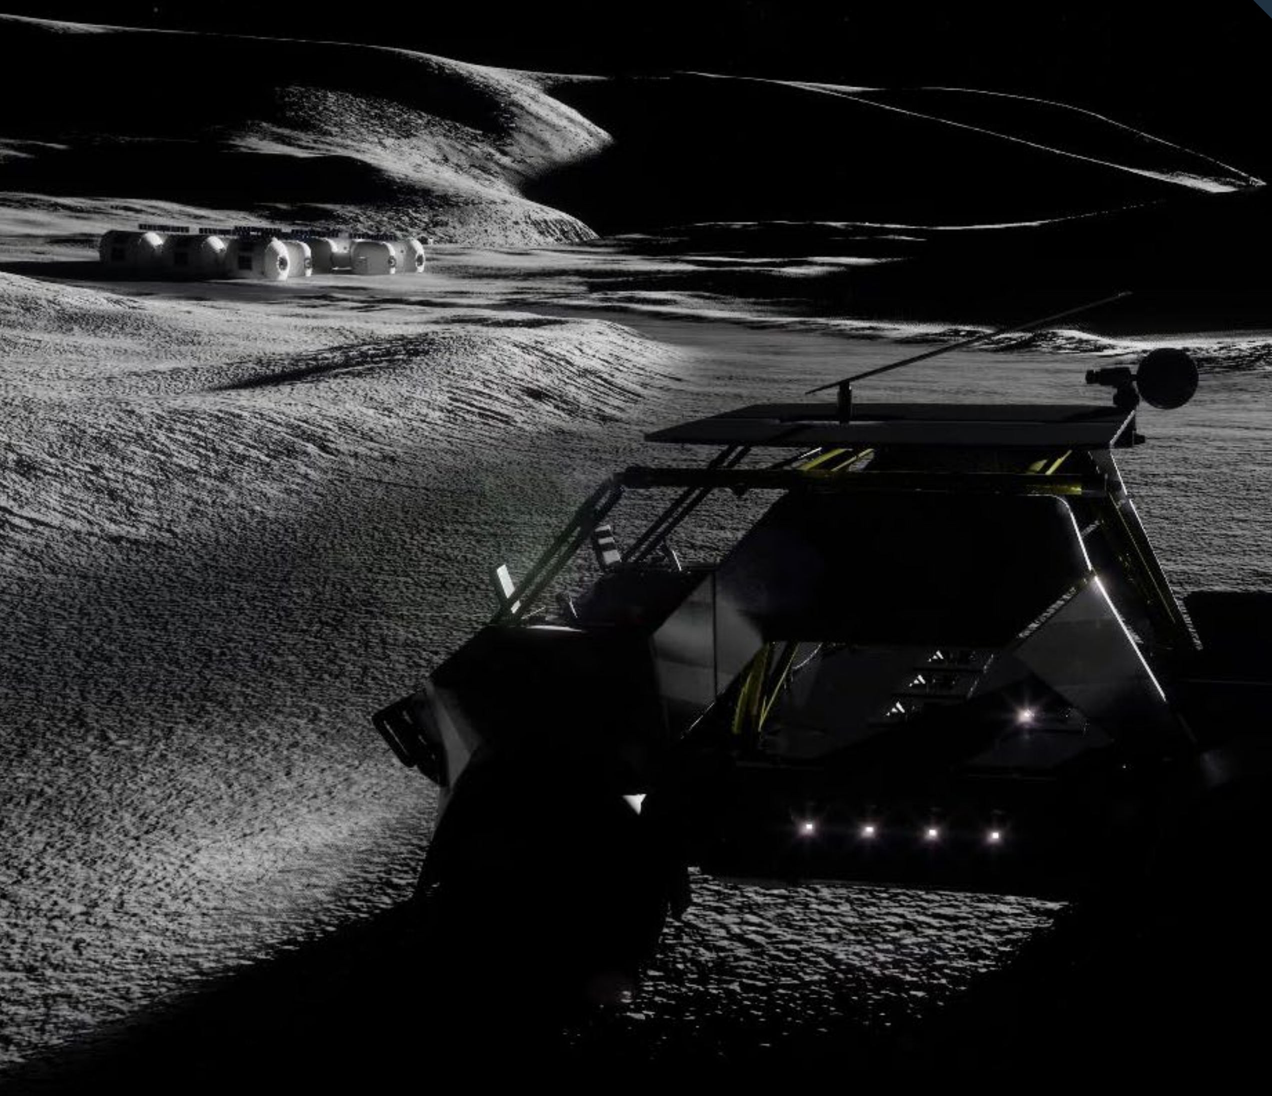

Technically, you and the other contractors aren't allowed to talk shop. Sensitivity of each company's operations and all that. But below the friendly competition is a kind of unwritten code. You're all humans, united against an insidious space environment that tries its best each second to creep into the systems that keep you alive, and that you were all hired to keep running. Power is one of the most essential lifegiving utilities. Its supply, and your ability to count on it, is what enables you and your companions to survive in the harshness of the lunar poles. How could you not discuss it? The hiccups and peccadillos of each different local microgrid ensure there's always something new to discuss each shift.

*How's the feed over at de Gerlache?*

*Spotty. Think we're going to have to throttle back quiescent ops for the first week of the night and see how it goes.*

*Yeah, that sounds about right. Bad dust occultation on half the arrays, during your predecessor's rotation. Bunch of her electrostatic sweepers froze overnight, right before a resupply hit the pad.*

*I remember that. Really cut into our bots' uptime. She was a hero back at headquarters when she got us back on schedule. Wish me luck this time.*

For you and your company, things are good on the power front thanks to Aquarius' nuclear-based strategy. Enabled by initial solar-powered systems, the company chose a decade ago to put its chips on power and propulsion technologies based on nuclear fission. The startup costs were significantly more than solar, and the skeptics and detractors were many. Their ranks even included some of the power consumer community. This was justified: solar has many benefits, from its low startup costs to its light weight and modularity. It's a great choice for bootstrapping small-scale mobile or temporary outposts. But for established stations, commercial activity, and at the scale of large, fixed infrastructure, the benefits of fission in the outyears became clear, especially in a time before we had the ability to manufacture solar arrays from lunar materials at large scale.

For solar, dust mitigation techniques can only keep so much dust off the panels and out of the turntables. The bigger a solar array, the more likely it is to be hit by debris or micrometeoroids. And the arrays naturally degrade over time, with the efficiency of some of the first-generation Artemis arrays inching down around 90 or even 85% today in 2040.

The biggest challenge is availability. Solar works best in continuous sunlight. The portion of the Moon's surface that remains in the sun for more than half of the lunar day is incredibly small and relegated to polar highlands that present access issues without capable mobility vehicles. As your colleagues so often lament, many of these peaks—Shackleton Connecting Ridge, de Gerlache-Kocher Massif, Malapert Massif further north—are usually some distance from resource operations or science surveys. You can only imagine how treacherous it was to emplace heavy cable down those slopes. Over long distances, that downmass adds up quickly, increasing the broader system-wide cost. Power beaming is a gamechanger when you can avoid cabling over rough terrain, but becomes more challenging over long distances.

Fission, on the other hand, can generate power anywhere it is deployed, and has safe High-Assay, Low Enriched Uranium (HALEU) fuel, multiplex control redundancy, and the ability to shield reactors in regolith. With the space environment trying to destroy your functionality every chance it gets, the reliability of fission can't be understated.

You wonder what the mess hall conversations of previous generations were like, and how the gripe du jour evolved over time. The advent of industrial scale operations and significant commercial activity, and the growing pains of a grid approaching global access, is a relatively recent development.

It took a decade for the grid to evolve into its current form. In the Artemis era, the power "ecosystem," such as it was, constituted little more than a solar array with a charging port for mobile vehicles. But with deliberate architecting by leading space infrastructure companies, sites were chosen that combined the best science potential, best solar power generation, and the best future potential for resource production operations. The right digital tools and sophisticated understanding of the surface environment allowed planners to optimize, to the nearest ten meters, the locations of each element of these “microgrids”: solar arrays, connective cable routing, beaming transmitters and receivers, and habitation, science, and ISRU loads.

At first, just about everything was on wheels. Each outpost’s microgrid moved with it, floating islands of power in an otherwise sterile landscape. A mobile architecture in the early years is what gave us the balance of flexibility and capability to achieve many of the goals of a permanent base, while allowing us to pursue the question of where the best science was to be found. With more of the surface explored, faster, it became clear where the best sites were to put down roots and emplace permanent infrastructure. We were able to settle down confidently in a location we knew held an intersection of huge scientific value, abundant power availability, and proximity to recoverable resources. Viewed with your privileged position of hindsight, this ended up being the most cost-effective way to proceed, taken overall as an entire mission campaign.

As for fusion and its promise of limitless power, experimental testbeds have made impressive progress. The experts expect a functional municipal-scale reactor in the next few years on Earth, and you expect it won’t be long until the first experimental reactor arrives at the Moon.

Those reactors are largely deuterium-tritium based, meaning the long-sought helium-3 mining operations for fusion fuel haven't yet come to pass. The part of you that grew up on Ray Bradbury and Arthur C. Clarke is a bit disappointed by that. The challenge with helium-3 has always been the difficulty of mining it economically, given how thinly it’s dispersed across the lunar surface. But investor interest remains. Demand gets more encouraging each year for niche markets like border security or medical imaging, keeping each new generation of He-3 mining ventures coming back to the table with better systems and better business cases. If it proves commercially valuable, the EM launcher is capable of accommodating the export of those commodities to Earth. For now, though, the only nuclear power production on the Moon is fission.

## Habitation and Construction

During your periodic sweep of the instrument panel, you find your eyes keep naturally drifting to the radiation monitor. There’s something about a spacefarer’s instinctive vigilance against radiation—natural and induced—that is hard to overcome. You remind yourself you’ve followed all the protocols, and not once in your rotation has one of the company’s many radiation failsafes triggered. This time, like the last time and the time before that, the monitor doesn’t register anything above background levels. The regolith mound over the reactor is doing its job admirably, and the early warning sensors for solar events provide enough time to not get surprised by environmental radiation effects either.

In the early years, NASA kept a wide radiation safety zone around its first-generation fission reactors, installed as they were on landed platforms that stood a dozen or more feet above the surface. For habitats, where crew needed to safely spend an extended amount of time in a shirtsleeve environment, the radiation protection was built into the structure of the vessel.

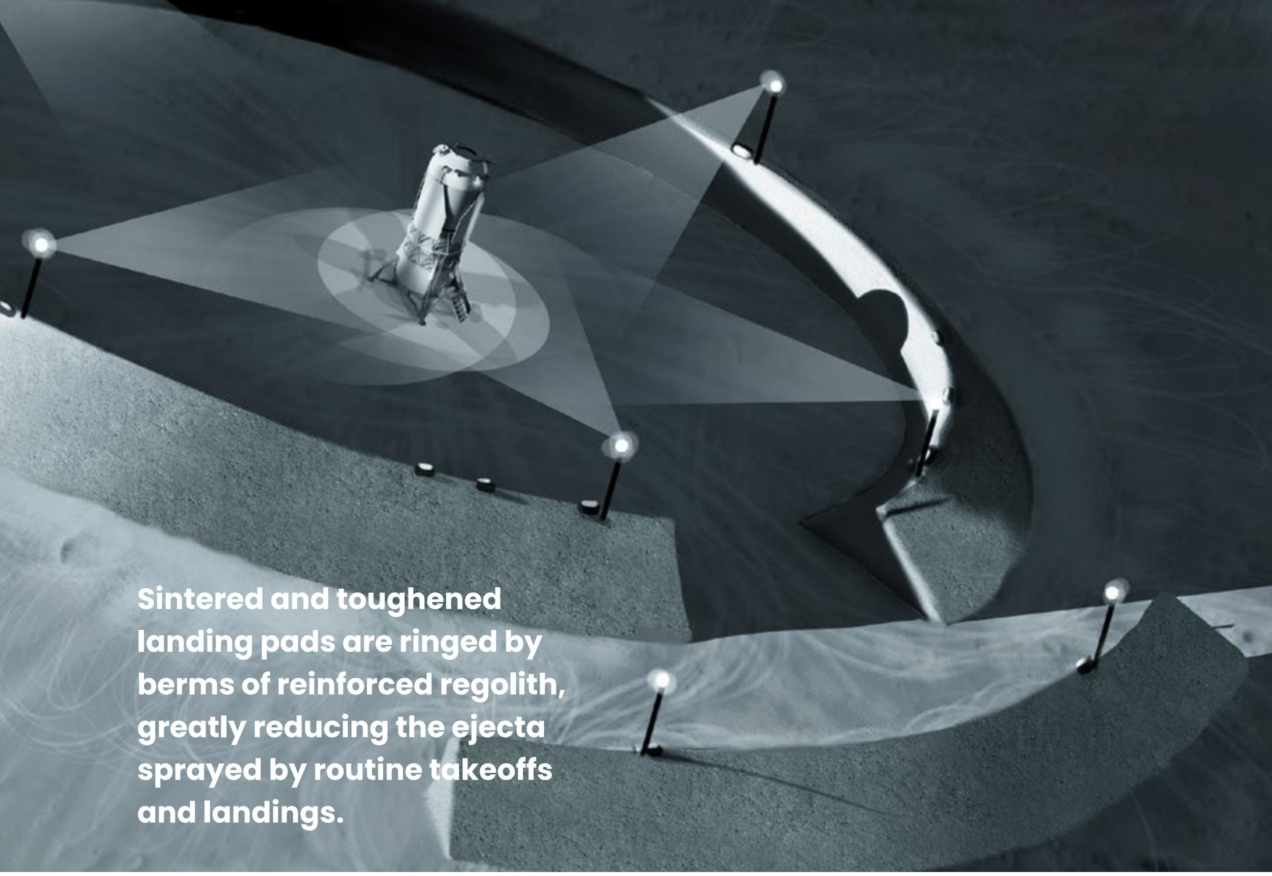


The original Artemis habitats—not including the landers astronauts arrived in—were lander-delivered hardshell habitats evocative of the ISS, Gateway, and some early pressurized rovers. The most effective habitats, however, are those based on inflatable softgoods technologies. Pound for pound, the woven material is tougher than the legacy metallic hardshell modules, and the multi-layer inflatables also provide more protection from micrometeoroids, orbital debris, and radiation at lighter masses than those legacy metallics. This affords NASA valuable mass savings. Perhaps the biggest benefit is the volume saved. A softgoods habitat could be launched in a tightly packed configuration to better fit inside a launch fairing, expanding to three or four times its packed size after being deployed on the surface. These mass and volume savings offer unique abilities, such as mounting habitats on mobility platforms with separate airlocks to provide better lunar dust mitigation.

Inflatable softgoods technologies were just as effective on orbit. In the decade after NASA deorbited the ISS, a number of private firms stepped up to provide Earth-orbiting space habitats commercially. These floating outposts provided multiple spaces where orbital science could continue, in-space manufacturing could be tested and improved, paying tourists could live and play, and a host of other nascent activities developed.

With the arrival of lunar excavators, it became feasible to safely nestle habitats, reactors, and other important infrastructure in regolith. Carving and backfilling trenches allowed for power and communications cables to be safely buried beneath the surface of the Moon, away from micrometeors, radiation, and sharp rover wheels. Later generations of habitats have been explored using different combinations of softgoods environmental bladders and tough regolith superstructures. Techniques evaluated include burying them and covering them in regolith, placing them inside pre-existing caves and lava tubes, and using any of the myriad 3D-printed skeleton approaches demonstrated by successive generations of regolith additive manufacturing technologies.

Earthmoving and regolith preparation technologies were used to toughen landing pads through sintering and ringing them with berms of reinforced regolith. This greatly reduced the ejecta sprayed by routine takeoffs and landings. The subsurface offered even better advantages. Huge underground tanks, serviced via surface ports, allow for large-scale, impact-resistant, temperature-stable propellant storage within arm's length of launchpads. As you know well, every meter and every minute shaved off of the refueling workflow reduces the boiloff losses from handling and transport. A kilogram saved is a kilogram earned.

## Commerce Supports Science

Your radio crackles just as the robots are wrapping up their operation. Contractor comms alerts you to a message that has just arrived on your dashboard. The mission planners have added a new waypoint, subject to your acceptance. Given how smoothly the installation went, you think you have enough schedule to pencil in one last stop before heading back to Artemis Base Camp.

You check the map and recognize the location without having to zoom in. You know it well. It's the same crater you've been to a couple times during your rotation, in order to check on the company's propellant production infrastructure there. You skim the accompanying text. It's a science services request. An international space agency would like you to pick up a science parcel dropped off by their robotic in-situ sample collector. It looks like a simple task. An hour of your time at most, well within the remaining timeline margin for this sortie. The planners likely wouldn't have offered it to you otherwise.

You skim through the terms and conditions. Scrawling your signature on the dashboard, you hit send and radio your acknowledgment of the requested contract modification.

You're surprised there's that much left to survey. You're under the impression that after a decade and a half, anything of note below 85° S has been pretty thoroughly explored, sampled, and mapped. The hotbeds of scientific exploration these days are all in the South Pole Aitkin Basin, or the equatorial nearside stations in the mare. There's also the far side where they're assembling an enormous radio telescope. Down here at the south pole, so close to the lunar base, exploration had mostly been turned over to commercial prospectors looking to sharpen the pencil on the already-known.

Contractor comms get back to you in under a minute. Cleared to proceed. Years ago, this sort of thing might have taken months or even years to work into the schedule, with lots of deliberation by the mission planners. One-off contract modifications used to be even tougher. There was a time when picking up another country's science parcel would have legally constituted an act of space piracy without a lengthy and complicated chain of approvals and ownership transfers.

Now, with global access, routine operations, highly capable vehicles, and with the legal and political processes streamlined, it was a breeze to rework a sortie to accommodate on-the-go service requests or added maintenance stops.

With the advent of commercial companies operating on the Moon, space exploration agencies were eager to take the most routine and mundane tasks off the plates of scientist-astronauts and their highly optimized

**Each time we can routinize an advancement, we can spend our collective human effort on the next groundbreaking thing.**

exploration robots. NASA, for example, was very keen to get multiple science campaigns up and running in collaboration with their international partners. Their scientist-astronauts and deep space spacecraft are the primary actors of that story, expanding the frontiers of deep space and answering some of the most fundamental questions about our role in the universe: the origins of the solar system; the search for extraterrestrial life; new planets around distant stars which humans might one day call home; peering through lunar farside telescopes into that empty primordial cradle into which were born the first stars. A few Nobel Prizes have already been won for this work.

Given the importance of these missions, and the perennial challenges of limited funding levels, space agencies were eager to contract out the more routine tasks. While scientist-astronauts expand the frontiers of science and the outer edge of known space, their commercial partners like you are there to provide the services that back them up. Contractors make sure that they have what they need to do their job safely and sustainably today, while at the same time building up the infrastructure for a more capable, more populous lunar and Martian future tomorrow.

The tasks are many: picking up samples, dropping off replacement parts, ferrying payloads from A to B, delivering supplies, or taking unplanned measurements en route. Some might call the science support services jobs—and contractor services in general—mundane, tedious, or unglamorous. But that's exactly the point. It's the job of contractors like you to take what used to be cutting edge yesterday and make it repeatable today, so that tomorrow it can become routine.

Each time science expands our horizons, commercial and industrial development fill in the newly discovered space with human progress. Each time we can routinize an advancement, we can spend our collective human effort on the next groundbreaking thing. The job for the present, then, was continuing to set ourselves up to provide that commercially supported end state. As a spacefaring community, we do that by opening up demand, creating novel solutions, increasing mission cadence, routinizing operations to reduce costs, and preparing ourselves for the next step-change in capability.

## Autonomy, Communication and Navigation

This new last stop today is in service to that greater mission. For this task, you'll need to employ the dedicated robotic assistant that has been tagging along behind your rover during today's sortie. The pressurized rovers these days are extremely capable, but permanently shadowed regions (PSR) are still pushing the limit. You get shivers down your spine peering down into the freezing primordial blackness of some of these craters. Even if the mission planners and contract manager lost their minds and cleared you to drop into a PSR, you'd be content to let the robots go fetch.

That's what your Mobile Utility and Logistics Element, or MULE, is for. Solo contractors work independently, and are often isolated, but you're never alone. The majority of the work these days is offloaded onto ever more capable autonomous sidekicks, especially when that work is simple, dangerous, or monotonous. These workers were always on, always connected to mission control, and—outside of scheduled recharging cycles or unexpected maintenance—always ready to support.

As if to signal its readiness to take on the task, a dashboard beep alerts you to the notification that your MULE, following dutifully behind you this entire sortie, has also connected to the orbiting communications constellation. It switches from navigating locally off of your vehicle to its own global navigation based on orbital and surface waypoints.

These pressurized long-distance rovers have been getting more comfortable and capable every few years. The autonomy has improved by leaps and bounds, and when you couple powerful compute with better surface mapping from each new survey, you realize you've shaved almost an hour off of the last route.

There is quite a bit of capable autonomy these days. Lunar edge computing, where powerful computers process data locally on the surface before beaming the refined information back to Earth, is showing promise even in its infancy. With the Moon's prospecting, surface mapping, and science activities generating petabytes each day, edge computing enables the best use of the bandwidth of the growing but still limited cislunar communications pipeline.

Navigation capabilities are making good use of all that data. Your rover is well served down here on the pole, with the heavy coverage of a network of mast-mounted ground beacons and relays to the orbital LunaNet-compatible network. We're not quite at the level of a full lunar version of Earth's GPS system, but the capability is increasing in leaps and bounds, and it's certainly enough to do your job. Each year the services improve, and the cost becomes more competitive thanks to the regular appearance of new entrants with increasingly useful systems. LunaNet is working and everything is running smoothly, with discussions in place to transition the management from NASA to a more internationally and commercially supported NGO.

**You're part of an elite cadre up here—every activity that occurs on the Moon is due to the combined intelligence, effort, and ingenuity of entire nations.**

## International Coordination

A tone issues from the dashboard. It's contractor communications again. They've rerun the path planning tool for this science parcel pickup and have found a more direct way to reach that crater. The catch: it will take you through the edge of the another country's safety zone in this sector. After some back-and-forth at corporate, they've decided it's worth the effort to navigate those permissions. They sound fairly confident the request will be approved this time, so they tell you to alter course and follow the new waypoints. It'll be a moment before the Infrastructure Directorate secures the right approvals.

Some newspapers that come across your feed make a big deal about international competition, typically painting it as something more serious. You're not so sure of that. There have been two notable international crew rescue missions now, and easily a dozen more interventions to rescue the autonomous assets of other nations. If you're being honest, you think things have progressed down the most positive of the likely timelines. Full collaboration might not be possible yet, but open and accessible communication helps to keep things civil in space.

We went to the Moon for scientific curiosity, inspiration, and economic success, but we also stayed for something more. Pride in national posture is one of the key elements that has propelled everyone to their current level of lunar development. With so many countries operating on the Moon, the fact that we respected each other's integrity and kept the phone lines open is the biggest cause for celebration. After all, we are all humans facing the same

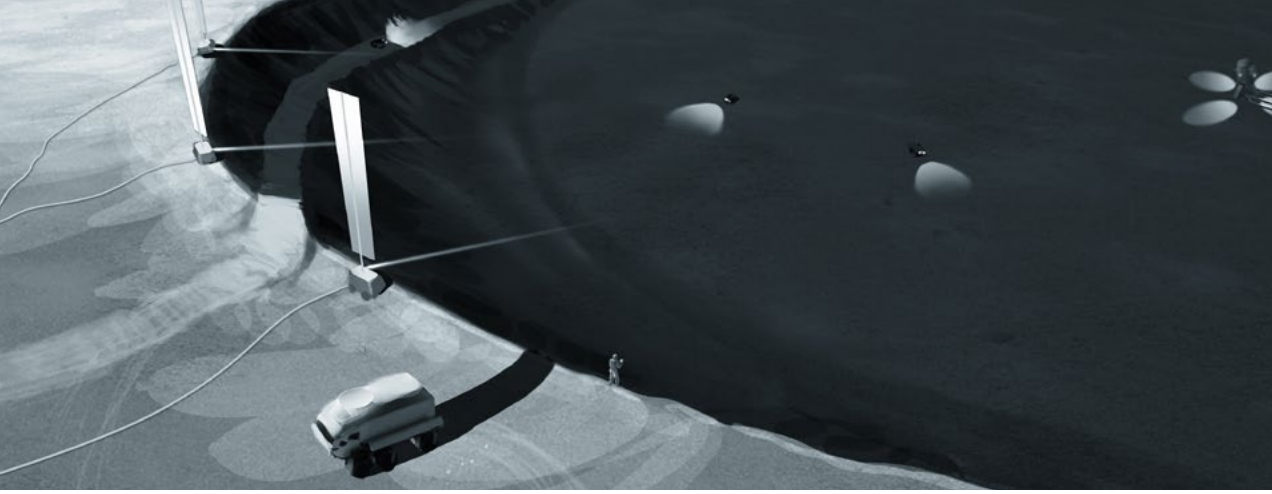

challenge on the Moon: a hostile environment of vacuum, radiation, dust, and extreme heat and cold.

As if to mirror your optimism, the contractor communicator comes back on the line. The Infrastructure Directorate has secured the go-ahead to traverse the other country's safety zone. The following waypoints are approved, you hear, as they're added to your screen. That's not bad at all, you think. The approval process has gotten so much smoother, even just in your first few months here. With the pace of development increasing, and the need for traverses through each other's safety zones, it is becoming more necessary to shake hands and tip hats as we traipse around each other's digital fences. That's not to say that there are borders, or that anyone has claimed territory, per se—the Outer Space Treaty forbids that in no uncertain terms, and the Artemis Accords reaffirms it. But each year we establish new mines, infrastructure safe-operating zones, hazard ellipses downrange of launch and landing trajectories, and other regions where you don't want to be caught without a comms link to the local authority. In that way, it's no different than navigating passage through terrestrial airspaces and receiving the necessary approvals from air traffic control.

Each space program requires the best its country has to offer. You're part of an elite cadre up here—every activity that occurs on the Moon is due to the combined intelligence, effort, and ingenuity of entire nations. Though you don't know it, you suspect that every human on the Moon shares that feeling. Each nation brings its best.

Regardless of which activities are being carried out in which safety zones by which actors, the development taking place represents the absolute cutting edge of what the human race has to offer. Regardless of the flag painted on the side of a vehicle, the mere fact that it survives up here is a testament to humanity's triumph over the vast emptiness of space as we take our first intrepid steps out of the cradle of Earth and into the inner solar system.

## A Diverse Commercial Ecosystem

This collaboration isn't restricted to international or interagency partnerships, either. Even in a competitive market, the value chains of the cislunar economy require handshakes and partnerships at many junctures.

Completely vertically integrated companies are the exception, not the rule. Most of us can't keep our systems running without the cooperation of many other technology and service providers. Like a terrestrial biological ecosystem, each player fills different roles.

Companies like Aquarius provide the high-reliability, high-redundancy, high-performance hardware that makes up the core of the most critical infrastructure, like with the power grid or the resource pipelines of the most important strategic resources. These are things that we just can't get wrong, for the good of everyone on and off the Moon. It's hard to put too high a price on the reliability of the entire ecosystem.

Then, with the resilience of that core infrastructure guaranteed—the sturdy roots and trunk of a tree—many opportunities arise

at different branches for more value-optimized firms, whose customers are willing to tolerate temporary interruptions in return for lower costs.

When this ecosystem functions successfully, the barriers to entry for lunar surface access drop away, allowing anchor providers to work hand-in-hand with the bootstrapped, risk-taking pioneers to add additional exciting functionality at successively lower price points.

Although, as with Earth ecosystems, survival is hardly guaranteed. As the harsh environment and the needs of the market put each system to the test, winning products emerge. Sound architectural management, by leading providers and in collaboration with a community of commercial, civil, and international partners, helps keep individual products interoperable with the rest of the ecosystem.

## Humanity's Future in Space

Another dashboard beep—you're approaching the crater. You punch in some commands to the utility MULE, commanding it to make ready its payload interface tool. The tool is designed for commercial interoperability thanks to the establishment of common, open-sourced interfaces and standards. It's like how a USB plug functions as a universally interoperable interface for power and data for all sorts of different equipment interfaces. There could be a hundred science parcels, each with any number and shape and weight of sample containers, but the same tool could handle every one because their common interfaces are all identical.

You perform a quick pre-operation check. The MULE is ready to go. This, like so many other activities, is a crewed operation with robot support. You'll be taking your pressurized rover down into shadow, just beyond the lip of the gentle sloping crater rim. It is the MULE's job to perform the dull, dirty, and dangerous work of grabbing the container and storing it inside your rover's external payload bay. You'll remain in comfort in your pressurized rover. You might not be deep into the crater, but you're still in the

dark, and you're just as leery as the Infrastructure Directorate when it comes to extravehicular activities in shadowed terrain. That's what the robot assistant is there for, after all.

Topographically, the crater rim is indistinguishable as you draw near. A gentle slope fades from white to gray to black as it falls away into the shadow of the crater rim. A lake of blackness stretches in front of you, 10 kilometers from end to end. As you approach, it begins to cover you in darkness. First the ground goes dark, then the wheels. The shadow advances up the outline of your windshield until your rover is completely dark. Your head swims. You experience this phenomenon each time you visit a crater: bottomless pit below and starless sky above, with only the thin slice of illuminated rim visible. The lack of features makes the scene stretch away into infinity above and below you.

Then, after those few initial seconds, your eyes adjust. Pinpricks of light twinkle at the bottom of the crater, the bustling hub of a water production plant coming into view. Red strobes mark the tops of production towers. Below them, floodlights illuminate utility pads and access stations. Robotic excavation vehicles crawl like ants, tracing paths from the water separation plant out to the excavation site and back. Other vehicles are busy moving processed regolith away from the plant and shuttling volatiles between tank farms. At the center, a distillation tower stands above the crater floor. Spotlights strike its side where a mosaic of logos represents the diverse coalition of companies involved. The dust kicked up by dozens of roving vehicles shrouds the scene in a haze, like a light at the bottom of the ocean.

Another light appears, much closer to you now, as the MULE rounds a boulder on its way back from fetching the science parcel from a deeper part of the crater. Its LiDAR sensors expertly find your rover's external payload bay, into which its robotic arm inserts the samples deftly and with care. It's doing its duty to preserve the core samples' delicate stratigraphy. A digital tone signals the successful completion of the task. Your final job of the day is complete. It's time to take these science samples back to the Artemis Base Camp.

The final hours of your sortie pass quickly. You store the samples, dismiss your MULE at the charging station, park your pressurized rover at the airlock, step into the habitat, and return to your quarters. Hours later, you find yourself relaxing in your bunk with a faded, dog-eared copy of your favorite book.

The day's hard work is done. Your routine tour of Aquarius' plants and infrastructure elements confirmed that the lunar ecosystem is running smoothly. You visited a construction site where an electromagnetic launcher is taking shape, an element whose completion will be a revolution in terms of mass launched to orbit and value returned to Earth. You oversaw the robotic preparation of a fission reactor, one of many that will one day power this launcher, for connection to the larger lunar power grid. You navigated not only lunar terrain but international relations from the seat of your rover, furthering another nation's science efforts while collaborating with a third by effectively coordinating a traverse through their safety zone.

You put the book down for a moment. Your mind drifts back to the striking scene of the crater floor and the busy hum of robotic activity. It occurs to you that the full arc of human exploration to date was represented in that scene. The pursuit of knowledge and prosperity in that crater continues the journey of those who crossed oceans, built flying machines to soar above the clouds, and ventured into orbit to see our whole world suspended in space. We achieved spaceflight, then early crewed missions, and then decades of humans living and studying in space. Our systems became reusable, then refuelable on-orbit, using water-based propellants to reach further and faster into deep space. And now the infrastructure in that crater is reducing the dependence on an earthbound logistical chain.

Alongside this technological evolution came the growth of true commercial demand, supplanting the traditional government monopsony and sustaining a life of its own. The emplacement of depots and other supporting infrastructure followed, serving this growing market more efficiently and permanently. In parallel, we had been preparing for the first use of propellant made offworld, which over time would steadily replace Earth-sourced commodities with those made locally. Even now we are replicating that evolution at Mars.

Looking into the next decade, you're excited to reverse the equation and return energy and material back to Earth orbit. Thanks to the efforts of you, of Aquarius, and of all space agencies and companies, humanity is able to reach higher on Maslow's hierarchy of needs. We're one stage further in Tsiolkovsky's plan. One step closer to O'Neill's dream.

Humanity's path to Mars and beyond runs past the Moon. This path is water-based and nuclear-enabled. It's uncovered by scientist-astronauts, paved with infrastructure, and maintained by you—as well as the tens of thousands of others who helped get you there. That path, though, ultimately leads back to Earth, where billions of humans stand to enjoy the benefits returned by this great odyssey.

And with that thought in mind, you turn off the lights.

Many thanks to this novella white paper's author Adam Marcinkowski and the team of Lockheed Martin engineers who developed this architecture.



Surface imagery of the Moon and Mars courtesy of NASA.

Learn more about Lockheed Martin's water-based lunar architecture vision here.


LOCKHEED MARTIN