

MINING · AUTONOMOUS OPERATIONS · CONNECTIVITY

The Network Sets the Ceiling on the Autonomous Mine

Why connectivity — not the trucks or the software — is the binding constraint on autonomy at scale: the three networks you are really building, and which decisions strand you in three years

AUDIENCE

Mining IT/OT, innovation
& automation leaders

FOCUS

Resilient connectivity for
autonomous mining

REGION

LATAM

Contents

01	The Decision That Caps Autonomy	3
02	What “The Network” Actually Means in a Pit	4
03	The Network That Worked Until the Telemetry Came On	5
04	Three Networks, Not One	7
05	For the Access Layer, Build on What Your OEM Validated — Mesh First, 5G Rising	8
06	The Backbone Is a Bandwidth Cannon	10
07	The Third Network: Redundancy, or You Leave Production to Luck	11
08	When Geography Stops Mattering	12
09	The Question to Ask Before You Sign	13
10	The Asymmetry of Waiting	14
11	Bets That Age Badly	15
12	What to Verify in Your Own Operation	16
13	Key Insights	18
14	Sources	19

01

THE DECISION THAT CAPS AUTONOMY

The autonomous truck is not the decision that determines whether your mine becomes autonomous. Neither is the fleet-management software. The decision that sets the ceiling is the communications network, and most mining operations choose it for the use case in front of them today, which condemns them to rebuild it inside three years. At NTT DATA we work alongside mining operations on exactly this decision — designing the network the mine's autonomy will demand, not the pilot that gets signed — and the pattern we see repeat is the one this piece takes apart.

Be precise about the claim, because it is easy to oversell. The network does not, on its own, make a mine autonomous — the trucks, the fleet-management layer, the people, and the operating discipline all have to be right. What the network does is set the *ceiling*: it is the binding constraint that silently decides how far the rest can scale. A brilliant autonomy program on an under-built network stalls at a fraction of its design. The network rarely gets credit when autonomy works, and it is almost always the thing that caps autonomy when it doesn't.

That ceiling is uncomfortable to talk about, because the network is the line item nobody fights over. The autonomous haulage system, or AHS¹, gets a steering committee. The dispatch and fleet-management software gets a proof of concept and a vendor bake-off. The network gets a coverage map, a price per square kilometre, and a signature. It is treated as plumbing. Plumbing is exactly what it is not.

Autonomy is a real-time control problem distributed across a moving pit. Every autonomous machine is a node that must be continuously seen, commanded, and corrected. Every sensor that feeds a safety interlock, every camera that lets a remote operator take control, every voice channel that connects a person in the pit to a person in a control room, rides the same physical infrastructure. The trucks and the software are bought from a catalogue. The network is the thing that decides how many of them you can actually run at once, from how far away, and whether the whole operation survives the day a single fibre is cut.

¹AHS — Autonomous Haulage System.

KEY INSIGHT — The network is not what makes a mine autonomous; it is the binding constraint that sets the ceiling on how far autonomy can scale. Choosing it for today's use case strands tomorrow.

This piece argues one thing, and returns to it in every section: the network is the binding constraint on autonomy at scale, and it is **three** networks, not one. A reader could collapse that into "buy more bandwidth" or "buy private 5G." That is not the argument. The argument is architectural. It is about which layers you are committing to: whether your access layer matches what your autonomy vendor has actually validated, whether your backbone has a migration path to optical transport before tomorrow's sensor density arrives, and whether that backbone has a second, independent road home for the day the first one fails. Bandwidth is a number. This is a decision about the shape — and the survivability — of the operation.

The mines that get this wrong do not get it wrong because they bought cheap equipment. They get it wrong because they sized the network to the pilot and signed before anyone asked what the mine would need when autonomy stopped being a pilot.

02

WHAT "THE NETWORK" ACTUALLY MEANS IN A PIT

Start with the substrate, because "the network" is an abstraction that hides the real engineering problem.

Picture an open-pit copper operation. The pit floor drops several hundred metres below the crest. Benches step down in tiers, and the working faces move every shift as the mine plan advances. Haul roads switch back up the walls. There is no fixed geometry: the place you needed coverage last month is mined out, and the place you need it next month does not exist yet. Dust, blast events, diesel particulate, and steel everywhere degrade radio frequency propagation. Line of sight is a luxury that the terrain revokes constantly.

Onto this surface you are asked to place a control-grade network. Not a network that lets someone check email near the workshop. A network that carries safety-relevant command and control for machines that weigh hundreds of tonnes and move without a person in the cab. When an autonomous truck loses its link, it does not coast. It enters a safe state and stops. A truck that stops in the wrong place stops the trucks behind it, and a stalled haul cycle is measured in tonnes per hour of lost production. Coverage is not a convenience metric. It is an availability metric, and availability is production.

Now add the second reality that procurement usually misses. The machines in the pit are only the first half of the system. The other half is everything the pit sends somewhere else: to a control room, increasingly to a remote operations center, or ROC², that may sit in a regional city tens or hundreds of kilometres away. High-definition video so a remote operator

²ROC — Remote Operations Center.

can supervise and intervene. Sensor streams feeding analytics and AI inference. Voice. Telemetry for the historians and the planners. The pit network gathers all of it. Something else has to carry it out.

And then add the reality that procurement misses entirely, until the day it cannot: that “something else” is a single road, and roads get cut. Fibre burns in a wildfire, gets severed by an excavator, or is stolen for the metal. When the road home goes down, every autonomous machine in the pit is blind to the people supervising it at once. The third question — what happens when the backbone fails — is the one nobody asks at signing and everybody asks at 2 a.m.

KEY INSIGHT — There are three networks, not one: the in-pit access network, the backbone to the remote operations center, and a redundant, independent path for that backbone. Most operations build the first, under-build the second, and forget the third.

Hold those three in mind, because the entire argument of this piece lives in how they differ. The first is the network the machines touch. The second is the network the data travels on after the machines are done with it. The third is the network that exists for the day the second one dies. They have different physics, different vendors, and different failure modes, and they are almost never bought by the same person with the same horizon in view.

03

THE NETWORK THAT WORKED UNTIL THE TELEMETRY CAME ON

At a copper operation in northern Chile that we worked alongside, we watched this sequence up close. The details below are composited and anonymised, but the pattern is one we have seen repeat, and one that connectivity leads across the region will recognise immediately.

The operation was running autonomous haulage at scale — a full fleet, well north of fifty trucks, with a fleet-management layer and a self-healing wireless mesh across the active benches. It worked. The mesh held the machines’ links through the switchbacks, and availability targets were met. The point that matters is what autonomous haulage actually needs to move: position, dispatch, traffic-management commands, basic health telemetry. A few hundred kilobits per machine — light, bursty, undemanding. The network had been sized and priced for exactly that, plus a short hop to a control room at the edge of the pit. It performed precisely as specified, and the business case closed.

Then the operation reached for the next tier of value, and two things changed at once. First, the control room moved — out of a cabin at the pit edge and into a remote operations center in the regional city, roughly 120 kilometres away, where it was safer, easier to staff, and cheaper to keep specialists. Second, and decisively, the data changed character. Real-time

telemetry and live high-definition video were switched on across the machines that generate the most of it — the drill rigs above all, whose penetration data, vibration signatures, and multi-camera feeds run continuous, high-rate, and unforgiving. This was no longer a few hundred kilobits of commands per machine. It was sustained, aggregated, latency-sensitive media aimed at the remote center in the city.

The in-pit mesh absorbed all of it. Add machines, add nodes, and the access layer scaled as engineered. The half that had not been engineered was the road out.

It surfaced, as these things do, in the middle of a night shift. A Tuesday in late March 2024, somewhere around 2 a.m., end of the southern-hemisphere summer, the dusty season not yet broken. The full telemetry-and-video load was live for the first time. In the remote center, the supervisory video began to stutter. Frames dropped. Latency on the link climbed and spiked. When a remote operator tried to take manual control of a machine flagged for an obstacle, the takeover arrived too late — past the threshold the safety case tolerates — and the system did what it is designed to do. It throttled. Machines eased into safe states. Production for that shift was reconciled the next morning against a number nobody wanted to present.

The pit access mesh was not the problem. It had capacity to spare. The problem was the backbone to the city: a legacy microwave path that had been perfectly adequate for commands and a nearby control room. Commands and position fit in a fraction of its capacity. Real-time telemetry and drill video, aggregated across the whole operation and carried the 120 kilometres into the city, asked it overnight for many times more than that path had ever carried, sustained and without pause³. The path could not do it. And here is the point a VP needs to hear: no configuration, no tuning, no vendor could fix it without rebuilding it. The ceiling was physical.

We summed it up bluntly at the time, and we repeat it here because it captures the whole pattern: the network ran the trucks fine — it fell over the day real-time telemetry and live video came on across the drill rigs.

The remediation was not a tuning exercise. It was a re-engineering of the backbone: fibre where it could be trenched, and a transport layer designed for aggregated video and inference rather than for commands — a multi-year capital line that had never appeared in the original autonomy business case. The network was chosen correctly for the use case of the day it was bought; it was wrong for the operation it had to become.

KEY INSIGHT — The backbone runs basic autonomy just fine and falls over the day you switch on real-time telemetry and video from the drill rigs: that is the data that does not fit.

That is the recognition moment this piece is built around. The failure was not incompetence. Every individual decision was defensible in its moment. The failure was a horizon problem, and a horizon problem dressed as a procurement decision is the most expensive kind, because it looks finished the day you sign it.

³Concretely: the legacy microwave link carried a few hundred megabits; the aggregated real-time telemetry and video load demanded well over a gigabit per second sustained, with low jitter, around the clock — an order-of-magnitude jump no tuning closes.

04

THREE NETWORKS, NOT ONE

The pattern generalises, and naming it precisely is the first practical step.

Every connected mine pursuing autonomy is building three distinct networks, even when it believes it is building one. The first two are functional layers; the third is what keeps the second one alive.

The first is the **in-pit access network**. This is the layer the machines, the people, and the sensors actually touch. Its job is mobile, control-grade coverage across a terrain that changes every shift, with self-healing behaviour as nodes move and links break and reform. Its physics are radio-frequency physics: propagation, interference, handover, latency at the edge. Its dominant constraint is reliable coverage under motion. This is the “last mile” of the mine — except the last mile here is mobile, and it moves every shift.

The second is the **backbone to the remote operations center**. This is the layer that carries the aggregated output of the pit to wherever the people and the compute live. Its job is to move large, sustained, low-jitter streams over distance: high-definition video from many machines at once, AI inference traffic, voice, telemetry, consolidated and pointed at the ROC. Its physics are transport physics: capacity, latency over distance, jitter. Its dominant constraint is aggregate throughput at bounded latency.

The third is the **redundant backbone path**: a second, physically independent road home, so that a single cut on the primary transport does not take the entire operation off the air. Its job is not capacity or coverage — it is survival. Its physics are the physics of *diversity*: two paths that share no pole line, no road corridor, no conduit, and ideally no medium, so that whatever kills the first cannot kill the second. Its dominant constraint is independence. This network does nothing on a good day. It is the only one that matters on the worst one.

Figure 1 · The three-network model: an in-pit access layer (autonomous machines, people, sensors) feeding an aggregation point; a backbone transport layer carrying consolidated video, inference, and voice to the remote operations center; and a redundant, physically diverse second path protecting that backbone. Most procurement specifies the access layer in full, under-sizes the backbone, and omits the third path entirely.

The reason this distinction matters is that the three layers fail in opposite ways and on opposite schedules. The access layer degrades gracefully and visibly: a coverage gap on a new bench is obvious, and the fix is another node. The backbone fails invisibly and then catastrophically on capacity: it carries commands and basic telemetry fine, and then saturates the moment real-time telemetry and live video come on, with no incremental fix available. And the backbone fails *totally* on a single physical event — a fire, a cut, a theft — unless a second path is already in the ground and already carrying nothing but readiness.

Procurement criteria make all of this worse. The standard questions asked of a mine network are coverage and cost: how much of the pit does it reach, and what does it cost per square kilometre. Both questions interrogate the access layer. Neither question interrogates whether the backbone can carry real-time telemetry and video three years from now, and neither

asks what happens to the operation when the backbone is severed. The criteria optimise the one layer that is easy to fix and ignore the two that are impossible to fix cheaply once you have hit the wall.

KEY INSIGHT — Not every autonomy ceiling is the network. Sometimes the bottleneck is the fleet scheduler, or an organisational gap wearing a network costume. Before you spend, one question separates them: when autonomy stalls, is the backbone saturated or sitting at thirty percent?

Before going further, one honest caveat about this model. Not every autonomy stall is a network ceiling. Sometimes the fleet-management scheduler is the bottleneck and the link is barely utilised. Sometimes the problem is an organisational or skills gap wearing a network costume. There is a one-question test that separates the real constraint from the look-alikes, and it is this: when autonomy stalls, is the backbone link saturated, or is it sitting at thirty percent? If the link is saturated, you have the structural problem this piece is about. If it is idle, your ceiling is somewhere else, and rebuilding the network will not move it. Diagnose before you spend.

05

FOR THE ACCESS LAYER, BUILD ON WHAT YOUR OEM VALIDATED — MESH FIRST, 5G RISING

The access-network choice is the one place in this entire decision where you do not have to rely on opinion, vendor positioning, or your own benchmarking. The autonomy original equipment manufacturers, or OEMs⁴, have already done the validation, and they have published it.

This matters because the access network is not a generic connectivity purchase. It is a control-grade safety system that the autonomous fleet depends on for command, correction, and safe-stop behaviour. The OEM that builds the autonomous machines has a direct stake in which wireless technology carries that control traffic, and the serious ones have certified specific vendors against their autonomy platforms. The signal to read is not a vendor's brochure. It is which vendor your machine's manufacturer has put its name behind.

The record is specific:

⁴OEM — Original Equipment Manufacturer.

Table 1 Table 1 · What the autonomy OEMs have actually validated. Komatsu certified Rajant Kinetic Mesh for FrontRunner AHS (2022) and has run FrontRunner over dedicated private LTE since 2019; Caterpillar completed in-house validation of Rajant wireless for Cat MineStar Command for Hauling (2023), with mine-site production validation as the announced next step; Epiroc built connectivity capability by acquiring Radlink, and Rajant is validated for Epiroc’s autonomous drills, including at Los Bronces in Chile. Sources in the Sources block.

AUTONOMY OEM	AUTONOMY PLATFORM	VALIDATED CONNECTIVITY
Komatsu	FrontRunner AHS	Rajant Kinetic Mesh (certified 2022); private LTE / 5G (Nokia/Sedna) qualified path
Caterpillar	Cat MineStar Command for Hauling	Rajant Kinetic Mesh, validated by Caterpillar (2023)
Epiroc	Agnostic autonomous haulage & drilling	Radlink (acquired); Rajant validated for autonomous drills

Read that table carefully, because there is a trap in it. Rajant, Radlink, and RADWIN are three different companies. The names rhyme; the roles do not. Rajant is the mesh vendor certified by Komatsu and validated by Caterpillar for autonomous haulage. Radlink is the connectivity capability Epiroc acquired to support its agnostic autonomy. RADWIN is a separate company offering point-to-multipoint, or PtMP⁵, broadband wireless. Conflating them is the kind of error that ends up in a procurement document, and it matters because the OEM validation attaches to specific vendors, not to the category.

And here is where 5G enters — honestly, as a first-class player, not as a footnote. Private cellular is no longer the future of the access layer; it is arriving now. Komatsu has run FrontRunner AHS over dedicated private LTE since a year-long qualification in 2019, and in 2025 a Komatsu FrontRunner deployment at Nevada Gold Mines became the first U.S. production deployment of FrontRunner AHS, supported by a purpose-built private 5G network (Sedna and Nokia). Newmont, after a private-5G trial with Ericsson at its Cadia operation, green-lit private 5G across its Tier-1 sites globally — a rollout spanning operations from Australia and Papua New Guinea to Mexico, Suriname and Ghana. The standards make the direction credible: the latest generation of industrial 5G already targets response times and a reliability strict enough that a heavy machine can accept control commands over cellular with no room to hesitate⁶. What this means for the operation is concrete: private 5G stops being just general connectivity and starts qualifying for the control traffic that only certified mesh could carry before.

State the status precisely, because a specialist will. Today, mesh remains the *predominant validated* control network for safety-critical autonomous haulage, and private 5G is the *rising, increasingly qualified* alternative — already in production for some platforms, co-deployed for general connectivity, video, and IIoT on others. This is not “5G has replaced the mesh.” It is “the validated set is broadening from mesh to mesh-and-private-cellular, and the smart operation tracks both.” The integration view is the right one: autonomy at scale is a marriage

⁵PtMP — Point-to-Multipoint (a broadband wireless topology in which one hub serves many endpoints).

⁶URLLC — Ultra-Reliable Low-Latency Communication, the 3GPP feature set (from Release 16) targeting ~1 ms user-plane latency at up to “six-nines” reliability; its Time-Sensitive Networking (TSN) integration is the path that lets 5G bridge into deterministic industrial control.

of technologies — certified mesh where control traffic demands it, private LTE/5G where capacity and device density demand it — and the access layer increasingly runs both.

So the access-layer rule is narrow and firm. For the control-grade wireless that your autonomous machines depend on, start from your autonomy OEM's validated list — mesh today, and the private-cellular path your OEM has qualified — and treat departures from it as risks you must justify, not defaults you may assume. The validation is the most expensive engineering you will never have to pay for. Use it.

06

THE BACKBONE IS A BANDWIDTH CANNON

If the access layer is a solved problem with a published answer, the backbone is the unsolved one, and it is where the ceiling of autonomy is actually set.

Think of the backbone as a cannon. The pit aggregates everything the autonomous operation produces and fires it at the remote operations center: high-definition video from dozens of machines simultaneously, the inference traffic from analytics and computer vision, voice from everyone connected to the operation, and telemetry on top. This is not the modest, bursty telemetry of a SCADA⁷ world, where a sensor reports a value every few seconds. This is sustained, high-volume, latency-sensitive media, many streams at once, all pointed in the same direction at the same time.

The arithmetic is unforgiving, and it is worth translating into the only unit that matters in a steering committee: a single machine under remote supervision carries several cameras, and at high fidelity one machine alone demands more sustained uplink than an entire operation consumed in the era of commands⁸. Multiply that by the drill rigs and a haulage fleet of many tens of machines, add fixed cameras, analytics, and inference, and the backbone goes from carrying a trickle to carrying a torrent, without pause, around the clock. Basic autonomy asks for a fraction of that. That is exactly the trap: a backbone sized for commands looks sufficient right up until the day the telemetry and video come on, and then it is short by an order of magnitude.

Figure 2 · The bandwidth cannon: dozens of high-definition video, inference, and voice streams aggregate in the pit and are fired down the backbone toward the remote operations center. The backbone's sustained-throughput-at-bounded-latency ceiling, not the access mesh, is what caps how far autonomy can scale.

This is why the urgency is real and not rhetorical. Sensor density and video fidelity are both rising, and they are rising faster than telemetry ever did. Higher-resolution cameras, more cameras per machine, more machines under remote supervision, richer on-board perception feeding centralised inference. Every one of those trends loads the backbone, not the access

⁷SCADA — Supervisory Control and Data Acquisition.

⁸At broadcast-grade fidelity, a single machine can demand on the order of twenty to thirty megabits per second of sustained uplink for video alone, before telemetry and inference. Basic autonomy — commands and position — asks for a few hundred kilobits per machine.

mesh. The operation that sizes its transport for today's telemetry is sizing it against the one variable that is guaranteed to grow.

At gigabit-and-rising sustained loads over tens or hundreds of kilometres, with hard latency and jitter limits because a remote operator must be able to take control in time, legacy microwave backhaul runs out of room. The honest engineering answer is optical transport: fibre and, where the future is being built deliberately, photonics. The ceiling of autonomous operations is not set in the last wireless hop to the truck. It is set in the road out of the pit.

07

THE THIRD NETWORK: REDUNDANCY, OR YOU LEAVE PRODUCTION TO LUCK

A backbone sized for the bandwidth cannon solves capacity. It does nothing for survival. A single fibre path carrying a gigabit of supervisory video is still a single fibre path — and a single path is a single point of failure for the entire autonomous operation. This is the network that procurement omits, and it is the one experienced connectivity engineers treat as non-negotiable: the backbone needs a second, independent road home. The access layer does not. The asymmetry is the whole point.

Why the backbone and not the last mile? Because of blast radius. A coverage gap on a bench affects one machine on one face; you add a node and move on. A severed backbone affects one hundred percent of everything at once: AHS supervisory control, SCADA, video, voice, personnel tracking, planning — all of it, simultaneously, for the entire operation, for as long as the cut takes to repair. The backbone is the single point through which the whole mine talks to the people who run it. Concentration of risk is concentration of consequence. You harden what fails hardest.

And backbones in mining fail in ways that are neither rare nor exotic. They fail by fire: a great deal of long-haul fibre is strung aerially on pole lines, and aerial fibre in a fire corridor is not degraded by heat, it is destroyed — the polymer coatings and jacketing that hold a cable together fail at temperatures a wildfire flame front exceeds many times over. In February 2024, the Valparaíso-region wildfires in Chile knocked 226 communications antennas out of service in that region alone (per SUBTEL), through a mix of direct fire damage and power loss; operators were still restoring service days later. They fail by theft: cable is stolen for its metal, and crews stripping copper routinely sever the fibre bundled beside it — an attack pattern that has climbed steeply enough that the industry now tracks it as a category of its own. And they fail, most commonly of all, by the blunt accident of an excavator bucket through a buried run. Fire, theft, and the backhoe: three ways to lose the road home, none of which a bigger pipe prevents.

KEY INSIGHT — The backbone needs redundancy; the last mile does not. A single backbone cut blinds the entire operation at once, and aerial fibre is exposed to fire, cable theft, and the excavator. Without a physically diverse second path, you are running production on luck.

Redundancy that works has one rule, and it is unforgiving: **the two paths must share no fate**. Two fibres in the same trench, on the same pole line, or down the same road corridor are not redundancy — they are one path that costs twice as much, because the fire or the excavator that takes one takes both. Real diversity means a geographically separate route and, where possible, a different medium, so that the event that kills the primary cannot reach the secondary. This is where the technology integration becomes concrete: the primary backbone is optical fibre; the redundant path can be a second diverse fibre route, a licensed microwave link on a different geography, a private-5G or fixed-wireless hop, or — increasingly viable in the remote Andes and Atacama where no terrestrial route reaches — a low-earth-orbit, or LEO⁹, satellite link. LEO services now deliver round-trip latency in the low tens of milliseconds and hundreds of megabits per terminal, which is ample for supervisory video, SCADA, and management traffic to keep the operation visible and safely controllable during a primary outage, even if it is not the path for sub-ten-millisecond control. In Chile and Peru, carriers already package LEO connectivity for exactly this remote-operations role.

Engineered properly, the failover is invisible to the operation. Traffic jumps to the surviving path faster than a remote operator can notice — faster than the latency margin a supervisory takeover already tolerates¹⁰ — so a backbone cut becomes a logged event rather than a stopped shift. The point is not that nothing ever breaks. The point is that when the primary road burns or is cut or is stolen, the autonomous operation keeps talking to its remote center on the second road, and production does not go dark while a repair crew drives two hundred kilometres of desert to find the break.

08

WHEN GEOGRAPHY STOPS MATTERING

Optical fibre raises the backbone ceiling. Photonics changes what the backbone is for. The distinction is worth a section, because it reframes the entire decision from “how much bandwidth do I provision” to “where does the operation actually have to be.”

NTT’s IOWN initiative is built around an All-Photonics Network, or APN¹¹: an architecture that pushes photonics end to end rather than converting between optical and electronic domains at every hop. The published 2030 targets are deliberately aggressive: 100 times

⁹LEO — Low Earth Orbit, the satellite-constellation class (e.g. Starlink) offering low-tens-of-milliseconds latency, suitable for supervisory and management traffic.

¹⁰On ring topologies, protection-switching schemes move traffic to the surviving path in well under a tenth of a second — below the latency budget a supervisory takeover already tolerates.

¹¹APN — All-Photonics Network, the photonics-end-to-end architecture at the core of NTT’s IOWN initiative.

the power efficiency, 125 times the transmission capacity, and one two-hundredth of the end-to-end latency of current infrastructure. State the epistemic status clearly, because the specialist will: these are 2030 targets, a direction of travel with numbers attached, not measurements you can benchmark in a pit today. What is not a target is availability. The All-Photonics Network has been commercial since March 2023 through NTT East and West. The direction is real, and it has shipped.

What makes this relevant to a mine is latency, not headline capacity. When end-to-end latency collapses toward the physical limits, the question of where data and inference physically sit stops mattering. Remote compute begins to behave as if it were local. That is not a slogan; it is being demonstrated. In an IOWN APN trial, video from a remote-controlled inspection robot was transmitted roughly 700 kilometres, Okayama to Tokyo, within 500 milliseconds of latency. Treat it as what it is: a demonstration, a directional signal of what photonic transport makes feasible, not a guaranteed production service level for your specific link. The direction it points is what makes it worth acting on.

For a mine, geography-free compute is not an abstraction. It means the remote operations center does not have to be near the pit, and neither does the inference. You can place AI workloads where compute is cheapest, where power is greenest, where governance is cleanest, or where your specialists already live, and still supervise machines in the pit as if the compute were sitting at the edge. The backbone stops being a cost you minimise and becomes an option you exercise: the operation gets to decide where its intelligence lives, instead of being told by the speed of light over copper.

This is also where the adjacent claim gets cleanly separated. A backbone built for the bandwidth cannon is necessary, and fibre solves the capacity problem. A redundant path solves the survival problem. Photonic transport solves a third, different problem: it makes the *location* of compute a choice rather than a constraint. The first keeps autonomy from stalling. The second keeps it from going dark. The third decides how cheaply and how well you can run the intelligence behind it for the next decade.

09

THE QUESTION TO ASK BEFORE YOU SIGN

Everything above reduces to a single diagnostic the reader can run before committing to any network purchase. It is deliberately uncomfortable, because the comfortable version of the question is the one that strands operations in three years.

DIAGNOSTIC QUESTION *Did we size the backbone for tomorrow's sensor and video density at full-density autonomy with the control room in the city — and give it a physically diverse second path — or did we size it for today's telemetry, on a single road, with the pilot's control room across the pit?*

The reason this question is decisive is causal, and it is worth naming the type of failure precisely, because the diagnosis determines the cure. The three-year rebuild is not, at root, a knowledge failure that better information dissolves. It is structural, rooted in the incentives of how mine networks are procured. Capital is justified per use case. The pilot justifies a pilot-sized network. Coverage and cost are the criteria because they are the criteria the procurement process knows how to evaluate. The horizon is short because the business case that releases the money is short. Those incentives do not change because someone reads a whitepaper.

But the failure is structural only up to the moment of decision. At the moment of signing, it is fully corrigible, because the buyer can reframe the criteria before the contract closes. After signing, the structural physics takes over and the rebuild becomes inevitable. That is the asymmetry the diagnostic exposes: a question that costs nothing to ask today prevents a capital program that costs years to execute tomorrow. The window in which the problem is cheap to fix is the window before the signature.

So the questions to bring to the integrator on Monday are not “does this network cover the pit.” Of course it covers the pit; that is table stakes. The questions are the diagnostic above, and the follow-ups that make it concrete: show me the aggregate sustained uplink this design supports at full-density autonomy; show me the migration path to optical and photonic transport; and show me the second, physically diverse path home and how fast it takes over when the first one is cut. If those answers are not in the proposal, the proposal is sized for today, on a single road.

10

THE ASYMMETRY OF WAITING

Treat the decision as the asymmetric bet it is, because framed correctly the choice answers itself.

The cost of building for tomorrow now is an incremental premium on a network you are buying anyway: transport designed for aggregated media rather than telemetry, an access layer drawn from your OEM's validated list of mesh and private cellular, a physically diverse second path for the backbone, and a contracted migration route to private 5G and to a photonic backbone. It is a larger line item in a capital plan you are already approving. It is bounded, it is known at signing, and it is spent once.

The cost of building for today and rebuilding later is unbounded in the ways that matter. It is the multi-year capital program the northern Chile operation walked into: a backbone re-engineered under production pressure, after the wall has already been hit, after a night shift has already been reconciled against a number nobody wanted to present. It is the autonomy program stalled at a fraction of its designed fleet while the road out of the pit is rebuilt. It is the shift lost to a single fibre cut that a second path would have made a non-event. It is the strategic option you never got to exercise, because the operation that cannot scale autonomy on the network it has is the operation that watches a competitor scale on the network it built.

KEY INSIGHT — Building the backbone for tomorrow (capacity, redundancy, and a migration path) is a bounded premium spent once; rebuilding it later is a multi-year program spent under production pressure.

The asymmetry is the entire argument for acting now rather than later. A premium paid once, at signing, against a rebuild paid in years, under pressure, with production on the line. The mine that treats the network as plumbing pays the second invoice. The mine that treats it as the infrastructure decision that caps autonomy pays the first, and gets to keep its options.

11

BETS THAT AGE BADLY

Reinforce the argument by negation. There are four connectivity bets that look prudent at signing and age into obsolescence, and they are distinct mechanisms, not variations on one.

The first is **the closed bet: single-vendor lock-in**. Committing the whole network, access and backbone, to one vendor's proprietary stack feels like simplicity and a single throat to choke. It ages badly because autonomy is multi-vendor by nature. Your haulage OEM, your drilling OEM, and your analytics layer will not all standardise on one connectivity vendor's roadmap, and the day you need to integrate equipment the vendor did not anticipate, the lock-in that simplified procurement becomes the constraint that caps the operation. The mechanism is dependency: you have outsourced your ceiling to someone else's product strategy.

The second is **the borrowed bet: depending only on unlicensed spectrum**. Unlicensed bands are free, available, and perfectly serviceable until the pit gets crowded. As machine density rises and more systems contend for the same air, interference is not a possibility but a trajectory, and you do not control it because you do not own the spectrum. The mechanism is contention: you have built control-grade autonomy on a resource that becomes less

reliable precisely as you use more of it. A serious autonomy network plans for licensed or private spectrum, including private 5G, before the unlicensed bands saturate.

The third is **the single-road bet: a backbone with no diverse path**. One transport route home, however fat the pipe, is one fire, one theft, and one excavator away from taking the whole operation dark. The mechanism is concentration: you have routed the entire mine's connection to its people through a single physical thing that the desert, the climate, and the metal thieves are actively trying to break. Redundancy is not gold-plating; it is the difference between a logged event and a stopped shift.

The fourth is **the short bet: sizing to today with no way forward**. Sizing the network to today's telemetry instead of tomorrow's sensor and video density is one face; having no migration path to private 5G or to a photonic backbone is the other. They are the same mechanism: a planning horizon shorter than the asset's life. A network with no migration path is a network with an expiry date you did not write down.

Notice that none of these four is a technology mistake. Each is a horizon mistake wearing a technology costume. The closed bet trades future flexibility for present simplicity. The borrowed bet trades future reliability for present cost. The single-road bet trades survival for a smaller invoice. The short bet trades the asset's full life for the pilot's business case. Avoid all four, and you have avoided the rebuild.

12

WHAT TO VERIFY IN YOUR OWN OPERATION

This is a gate, not a summary. Before signing any network decision tied to autonomy, run these six checks. Each has a binary outcome, and any single failure means the design is sized for today.

1. **Project the demand.** Calculate peak aggregate sustained uplink at full planned autonomy: machines times video fidelity times streams per machine, plus drills, fixed cameras, inference, and voice. Use the three-year fleet plan, not the pilot. *Confirmed if the number is written down. Disconfirmed if nobody has it.*
2. **Measure the ceiling.** Compare that demand against the backbone's actual sustained capacity to the remote operations center, at the latency and jitter your safety case requires. *Confirmed if headroom is at least double projected demand. Disconfirmed if it is under that.*
3. **Check the access list.** Confirm the control-grade access vendor appears on your autonomy OEM's validated list — mesh or the OEM-qualified private-cellular path — and that you have not confused Rajant, Radlink, and RADWIN. *Confirmed if the vendor is validated for your platform. Disconfirmed if the choice rests on a generic coverage map.*
4. **Test the spectrum.** Confirm the design does not depend solely on unlicensed spectrum at full machine density, and that a licensed or private-5G path exists. *Confirmed if a licensed path is contracted or planned. Disconfirmed if the whole bet is on unlicensed bands.*

5. **Prove the second road.** Confirm the backbone has a physically diverse redundant path — different route, ideally different medium, sharing no pole line, conduit, or corridor with the primary — with a defined, tested failover time. *Confirmed if the diverse path and its failover are documented and tested. Disconfirmed if “redundancy” is two fibres in one trench, or a verbal assurance.*
6. **Find the path forward.** Confirm the contract contains an explicit migration path to private 5G and to an optical and photonic backbone. *Confirmed if it is in writing. Disconfirmed if it is a verbal assurance.*

Run those six Monday morning. If every answer comes back confirmed, you are buying the autonomous mine. If even one comes back disconfirmed, you are buying the pilot, and you are buying the rebuild that follows it.

■ *The autonomous mine is not bought one truck at a time. It is bought one network at a time — and on more than one road — and the network you let yourself ignore today is the one you will pay to rebuild in three years.*

13

KEY INSIGHTS

The core claims of this paper, each able to stand on its own — a one-block-per-insight memory aid.

1 The network is not what makes a mine autonomous; it is the binding constraint that sets the ceiling on how far autonomy can scale. Choosing it for today's use case strands tomorrow.

2 There are three networks, not one: the in-pit access network, the backbone to the remote operations center, and a physically diverse redundant path for that backbone. Most operations build the first, under-build the second, and forget the third.

3 The backbone carries the basics of autonomy — positioning, dispatch, commands — just fine. It falls over the day you switch on real-time telemetry and live video from the drill rigs: that is the data that does not fit.

4 The right access network is the one your autonomy OEM already validated: Komatsu certified Rajant mesh, Caterpillar validated Rajant, Epiroc acquired Radlink and validated Rajant for its drills. And beware: Rajant, Radlink, and RADWIN are three different companies; conflating them is a procurement error.

5 The backbone needs redundancy; the last mile does not. A single cut blinds the whole operation at once, and aerial fibre is exposed to fire, cable theft, and the excavator. Without a physically diverse second path — sharing no route or medium — you run production on luck.

6 NTT's IOWN All-Photonics Network targets 100x power efficiency, 125x capacity, and 1/200 latency by 2030, and has been commercial since March 2023.

7 Building the backbone for tomorrow — capacity, redundancy, and a migration path — is a bounded premium spent once; rebuilding it later is a multi-year program spent under production pressure.

14

SOURCES

The claims in this paper that rest on public, verifiable evidence, with their sources. Vendor performance figures are labelled as vendor claims; NTT IOWN figures are 2030 targets except where commercial availability is stated.

- Rajant Kinetic Mesh certified for Komatsu FrontRunner AHS (2022) — International Mining; Canadian Mining Journal.
- Komatsu FrontRunner AHS over dedicated private LTE (2019) — Komatsu Newsroom.
- Caterpillar validates Rajant wireless for Cat MineStar Command for Hauling (2023) — Caterpillar press release; International Mining.
- Rajant validated for Epiroc autonomous surface drilling, incl. Los Bronces (Chile) (2022) — GlobeNewswire.
- Komatsu FrontRunner AHS on private 5G (Sedna and Nokia) at Nevada Gold Mines — first U.S. production deployment (2025) — International Mining; Barrick press release.
- Newmont green-lights Ericsson private 5G across global sites after Cadia trial (2024) — RCR Wireless.
- 3GPP Release 16 URLLC and TSN integration (2020) — Ericsson Technology Review: 5G NR evolution. The ~1 ms user-plane latency and “six-nines” (99.9999%) reliability are the 3GPP URLLC design targets (3GPP TS 22.261 / TR 38.913), not figures from the Ericsson article.
- RADWIN mining solutions — vendor positioning, fibre-class PtMP (vendor claim) — RADWIN.
- Chile Valparaíso-region wildfires impact on communications infrastructure (SUBTEL: 226 antennas out of service in the Valparaíso region, mixed fire damage and power loss) (Feb 2024) — SUBTEL situation report; Emol.
- Attacks on communications infrastructure / cable theft trend and economic impact — NCTA Vandalism 2025 Report (PDF).
- LEO satellite (Starlink) for remote business connectivity in Chile and Peru — Entel-Starlink partnership.
- NTT IOWN All-Photonics Network — architecture and APN functions — NTT IOWN APN; 2030 targets (100x power efficiency, 125x capacity, 1/200 latency) — NTT DATA: what makes IOWN impressive; IOWN APN trial — ~700 km Okayama→Tokyo inspection-robot video within 500 ms (Jan 2025) — NTT DATA press release.

For more on this topic, contact: **Rodrigo Pinto**, Director TMT - Network - IOWN, NTT DATA · rodrigo.pinto.padilla@emeal.nttdata.com.