Autonomous Infrastructure

The Missing Layer of Physical AI

Ascent Systems Technologies

White Paper | 2026

1. Beyond the Autonomous Machine

Physical AI is extending artificial intelligence beyond the digital domain and into systems that perceive, decide and act in the physical world. Robots, uncrewed vehicles, intelligent sensors and other autonomous systems can increasingly navigate complex environments, interpret what they observe and perform useful tasks with limited human intervention.

Much of the progress—and much of the attention—has naturally focused on the intelligence and autonomy of the machines themselves. Better perception, more capable AI models, improved navigation and increasingly sophisticated decision-making continue to expand what an individual autonomous system can accomplish.

But the machine does not operate in isolation.

It requires energy. It must communicate and exchange information. It may depend on positioning, external sensing and edge or cloud computing. It must withstand its operating environment. Mobile systems eventually need to recharge, refuel, exchange payloads or data, and in many cases be recovered or serviced.

In established environments, much of this supporting infrastructure is simply assumed to exist. Power is available. Communications networks provide connectivity. Computing resources can be reached. Roads, facilities and people provide logistics and maintenance.

Move the same autonomous system into an Arctic region, a remote coastline, a disaster area or a contested environment, and those assumptions begin to break down.

This exposes an important distinction:

Platform autonomy is not operational autonomy.

A machine may be highly autonomous in performing its task while remaining dependent on infrastructure and human logistics to sustain its operation. Increasing the intelligence of the machine does not, by itself, remove that dependency.

As Physical AI moves farther from established infrastructure, operates for longer periods and assumes more consequential roles, the infrastructure supporting it must evolve as well.

We call this missing layer **Autonomous Infrastructure**: a distributed network of self-sustaining nodes that provide energy, communications, edge intelligence and mission services where conventional infrastructure is unavailable, unreliable or impractical.

The objective is not merely to make individual machines autonomous. It is to create an operating environment in which autonomous systems—and the people and sensors working with them—can remain effective with progressively less dependence on continuous external support.

2. From Platform Autonomy to Operational Autonomy

Consider an autonomous UAV capable of planning its route, avoiding obstacles, detecting objects of interest and adapting its behaviour without continuous operator control. By conventional measures, it may possess a high degree of autonomy.

Yet someone may still have to transport it to the operating area, establish its communications, provide electrical power, launch it, recharge it after every sortie, retrieve its data and maintain the equipment supporting it.

The aircraft is autonomous. The operation is not.

The same distinction applies beyond UAVs. An intelligent sensor may perform sophisticated processing locally while depending on externally supplied power and a fixed communications network. A ground robot may navigate and perform tasks without an operator but return regularly to a staffed facility for charging and servicing. A distributed sensing system may generate valuable information yet become ineffective when its connection to a centralized network is interrupted.

These are not deficiencies in the autonomous platforms themselves. They are **system-level dependencies**.

For short-duration operations close to established infrastructure, those dependencies may be entirely acceptable. They become more consequential as distance, duration, environmental severity and operational complexity increase.

A system operating for hours presents one problem. A network of intelligent physical systems expected to operate for weeks or months across a large, sparsely supported region presents another.

At that scale, autonomy cannot be considered solely as an attribute of an individual platform. It becomes an attribute of the complete operational system.

That system must provide several persistent functions: energy for sensors, communications, computing and mobile assets; connectivity across available terrestrial and non-terrestrial networks; edge processing when remote computing is unavailable or undesirable; situational and environmental awareness; support for mobile autonomous systems; and protection against the conditions in which the system must operate.

Today these functions are often provided separately. Power systems, communications equipment, computing, sensors and robotic support infrastructure may each require their own installation, protection, maintenance and logistical chain.

This invites an irony: increasingly autonomous machines supported by an increasingly complicated collection of **non-autonomous infrastructure**.

Autonomous Infrastructure approaches the problem differently. Instead of treating supporting infrastructure as a collection of passive utilities, it treats infrastructure as part of the autonomous system itself.

A self-sustaining infrastructure node can generate and manage energy, maintain multiple communications paths, process information locally, monitor its environment and its own condition, support modular payloads, and provide services to mobile autonomous systems. It can adapt to changing resources and conditions and continue performing essential functions during periods of disrupted or limited connectivity.

Distributed nodes extend this capability across an operating area. Each can function independently, while collaborating with other nodes, mobile systems and users when connectivity permits. Intelligence, sensing and communications therefore need not depend on a single fixed installation or continuous connection to a remote command centre or cloud service.

The result is a shift in the boundary of autonomy.

Instead of asking only:

How autonomous is the machine?

we can begin asking:

How autonomous is the operation?

3. Autonomous Infrastructure: A Systems Architecture

Autonomous Infrastructure is not defined by a particular platform, technology or configuration. It is a systems architecture in which distributed, self-sustaining nodes provide the physical and digital services required to support persistent operations.

The architecture begins with the **node**.

Each node establishes a local operating capability by combining several functions that conventional infrastructure often provides separately: energy generation and storage, communications, edge computing, sensing and situational awareness, and support for modular payloads and autonomous systems. The combination will vary according to application, but the underlying principle remains the same: essential services are available at the point of operation without assuming continuous access to fixed infrastructure.

Energy is fundamental. Physical AI ultimately consumes physical power. Sensors, processors, communications equipment and robotic systems cannot remain operational on intelligence alone. An autonomous infrastructure node must therefore be capable of generating, storing and managing energy locally while continuously balancing available resources against operational demand.

Communications must be similarly adaptive. Rather than depending on a single communications path, nodes can use available terrestrial and non-terrestrial networks, local wireless links and node-to-node connectivity. The objective is not continuous connectivity at any cost, but the ability to maintain essential services, select available paths and continue useful operation when external communications are degraded or interrupted.

Edge intelligence complements this communications resilience. Processing information at or near the point of collection reduces bandwidth requirements and dependence on centralized computing. Sensor data can be filtered, interpreted and fused locally; routine decisions can be made at the edge; and relevant information can be exchanged with other nodes or users when connectivity permits.

Sensing gives the infrastructure awareness of both its surroundings and its own state. Environmental conditions, energy availability, communications status, payload health and other operational parameters can inform how resources are allocated and how the node configures itself. This creates a feedback loop between perception, decision and physical adaptation—the same principle underlying Physical AI, applied to infrastructure itself.

Mobile autonomous assets extend the reach of each node. They can perform sensing, inspection, communications relay, logistics or other tasks beyond the node's immediate location. The infrastructure, in turn, can provide them with energy, connectivity, computing, data exchange and other services required for sustained operation. As launch, recovery and charging become increasingly automated, robots can operate from autonomous infrastructure nodes rather than relying on permanent structures and personnel.

The architecture becomes more capable as nodes collaborate.

Distributed across an operating area, nodes can share information, extend communications, coordinate sensing and support mobile systems operating between them. Different nodes may host different payloads or perform different roles, allowing capability to be distributed rather than concentrated in a single installation.

Importantly, collaboration does not require dependence. Each node should retain sufficient local capability to continue essential functions when isolated. When communications are restored, information and system state can be exchanged and coordinated operation resumed. This makes the architecture inherently suited to disconnected, disrupted, intermittent and limited-bandwidth (DDIL) environments.

Autonomous Infrastructure therefore combines two characteristics that can appear contradictory: **independence and collaboration**. Individual nodes remain useful on their own; collectively, they create capabilities greater than those of any single node.

The resulting architecture is not simply infrastructure *for* autonomous systems. The infrastructure itself becomes an active participant in the autonomous operation—sensing, adapting, allocating resources, supporting mobile systems and collaborating across a distributed physical environment.

4. From Architecture to Reality

Autonomous Infrastructure is an architectural concept, not a particular product. Its principles can be implemented in different forms and scaled according to the environment, application and services required.

But the architecture is not hypothetical.

Ascent Systems Technologies has been developing and fielding systems that embody these principles through the **Autonomous Expeditionary Mission System (AEMS)**. AEMS provides a practical implementation of an autonomous infrastructure node: a self-contained, auto-deployable platform capable of establishing mission services supported by persistent power generated in situ, where conventional infrastructure is limited or unavailable.

At its foundation is a hybrid energy system combining solar PV generation, battery storage and fuel-cell power. Intelligent energy management continuously balances generation, storage and consumption, allowing the node to sustain its own operation while supporting communications, computing, sensing and other payloads.

AEMS integrates multiple communications options, including terrestrial and non-terrestrial connectivity, with local networking to support both external communications and nearby systems. Edge computing enables data processing and decision support at the node, reducing dependence on continuous remote connectivity and allowing essential functions to continue in DDIL conditions.

The platform can accommodate modular sensing, communications and other payloads and is capable of supporting uncrewed aerial (UAV) and ground (UGV) vehicles as well as other autonomous and robotic systems. This allows the node to serve not simply as a stationary installation, but as an operating base from which mobile autonomous systems can extend sensing, communications and other capabilities across the surrounding area.

Physical autonomy also matters at the infrastructure level. AEMS can deploy and stow its major systems automatically and manage its configuration in response to operational and environmental conditions. This reduces the need for personnel to establish and continuously support infrastructure at the deployment site and enables the system to transition between operating and protected configurations when conditions require it.

These capabilities have moved beyond laboratory trials. AEMS has been deployed in extended field evaluations in Canada, including operation in support of remote monitoring and current deployment on Vancouver Island as part of the **COAST project**, where the system is being evaluated in a maritime/coastal environment. Earlier development and testing have also included work under Canadian Department of National Defence and Transport Canada projects.

The significance of these deployments is not that AEMS represents the only possible form of Autonomous Infrastructure. Rather, they demonstrate that capabilities often treated as separate infrastructure functions—energy, communications, edge computing, sensing, environmental adaptation and support for autonomous systems—can be integrated into a self-sustaining physical node and operated outside conventional infrastructure.

That distinction is important.

AEMS is a node of Autonomous Infrastructure. It is not Autonomous Infrastructure itself.

The larger architecture emerges when independently capable nodes collaborate with other nodes, sensors, robotic systems and users across an operating area. Individual implementations may differ substantially in size, configuration and purpose while adhering to the same architectural principles.

This also makes the concept scalable. Autonomous Infrastructure may begin with a single node establishing persistent capability at one location, expand into a cluster of collaborating nodes across an operating area, or ultimately form distributed networks spanning multiple regions.

The objective at every scale remains the same: to move essential infrastructure closer to the systems that need it while reducing dependence on permanent structures, continuous connectivity and personnel at the point of operation.

5. From Autonomous Machines to Persistent Autonomous Operations

The evolution of Physical AI invites a broader question about what we mean by autonomy.

Today, autonomy is most often attributed to individual machines. A robot can navigate autonomously. A UAV can plan and execute a route. An intelligent sensor can classify what it observes. These capabilities are increasingly sophisticated, but operational outcomes rarely depend on one autonomous machine acting alone.

Real-world operations involve systems of systems.

Sensors generate information. Communications move it. Computing turns data into knowledge. Energy sustains them all. Mobile systems extend physical reach. People establish objectives, exercise authority and intervene when their judgement is required. Infrastructure connects these elements and allows them to function over time.

As these systems become more distributed, the **persistence of the operation** becomes as important as the autonomy of any individual component.

This suggests a progression in Physical AI:
from autonomous machines, to autonomous systems, to persistent autonomous operations.

At the latter stage, autonomy becomes an ecosystem property. Individual assets may enter or leave the network. Communications paths may change. Energy availability may fluctuate. A robot may return to a node for charging and subsequently redeploy. Sensor information may be processed locally when connectivity is constrained and shared more broadly when bandwidth becomes available. Different nodes may assume different roles according to their location, resources and payloads.

The infrastructure provides continuity through these changes.

This does not imply removing people from the system. Operational autonomy should reduce the need for people to perform routine infrastructure functions while preserving human involvement where judgement, authority, creativity or accountability are required. The objective is not a human-free operating environment, but one in which human effort is directed toward the purpose of the operation rather than sustaining the infrastructure behind it.

The resulting operational model is fundamentally different from deploying autonomous machines from conventional bases. Instead, autonomous and robotic systems operate with and from distributed infrastructure that possesses autonomy of its own.

That shift can create operational advantage in several ways. Persistent local energy extends duration. Edge intelligence reduces dependence on remote computing and constrained communications. Multiple communications paths improve resilience. Distributed sensing expands situational awareness. Autonomous deployment and environmental adaptation reduce the logistical burden of establishing and maintaining capability. Support for mobile systems extends the physical reach of each infrastructure node.

The value becomes particularly significant where permanent infrastructure is difficult, expensive, vulnerable or undesirable: Arctic and remote regions, maritime and coastal environments, emergency and disaster response, critical infrastructure, persistent monitoring, and defence and security operations.

In these environments, persistence itself becomes a capability.

Physical AI has made remarkable progress in teaching machines to perceive, reason and act. The next challenge is to create operating environments capable of sustaining those machines—and the people and systems working with them—for as long as the purpose requires.

Autonomous Infrastructure addresses that challenge by extending autonomy beyond the individual machine to the physical and digital foundation that sustains the operation.

The next step in Physical AI may therefore not simply be more autonomous machines.

True Autonomy may be **autonomy that can persist.**