



AVENTUS PERSPECTIVE / COMPUTE

Compute Is Becoming Critical Infrastructure. It Should Be More Open.

As demand for compute grows, the question is no longer only how much capacity the world can build, but how that capacity is accessed, verified and shared.

Every digital product eventually comes down to a physical question: Where does the work actually happen?

Software may appear weightless, but every transaction, AI response, industrial calculation and digital service depends on machines somewhere allocating processing power, memory, storage, networking and energy.

That underlying capacity is compute.

Access to it helps determine which ideas become working products, how quickly they can grow and what they cost to operate.

The scale of future demand is becoming difficult to overlook. In its 2025 Energy and AI report, the International Energy Agency projected that global electricity consumption by data centres will more than double between 2024 and 2030, reaching around 945 terawatt-hours. AI is expected to be the largest driver of that increase, alongside growing demand for other digital services.

Meeting that demand will require enormous investment in physical infrastructure.

It also raises a broader question.

As software takes on more responsibility, how much of the infrastructure beneath it should depend on a small number of operators, and where would wider participation, independent verification and greater portability create real value?

At Aventus, that question is increasingly important to us.

We have spent almost a decade working on systems where organisations, applications and participants need to coordinate without simply taking one another's word for everything.

Aventus Cloud is the next stage of that work.

The case for wider participation

Cloud computing transformed the way businesses build software.

Instead of purchasing and maintaining all of their own infrastructure, organisations could access computing resources when needed, scale them as demand changed and concentrate more of their effort on the applications their customers actually use.

Those advantages remain substantial.

But convenience can create dependency.

The UK's Competition and Markets Authority concluded its cloud services market investigation in 2025 after examining the difficulties customers face when switching between providers or using several providers together. Its findings identified technical and commercial barriers, including interoperability restrictions and data egress costs, that can limit customer choice.

There is room for more ways to participate in the compute economy.

Open standards, portable applications, multiple infrastructure providers and distributed networks can all contribute.

Decentralised compute introduces another model: machines operated by different participants contribute resources under common network rules. The network coordinates participation, distributes work and establishes how the resulting output is accepted.

For some workloads, separate participants can also verify the result.

That independence becomes particularly useful when several organisations rely on the same outcome.

Imagine a shared business process in which a calculation determines a payment. Everyone involved needs confidence in the inputs, the rules applied and the final result. If one participant exclusively controls both the calculation and the evidence supporting it, every other participant inherits that dependency.

Independent execution or verification can provide a stronger basis for agreement when the workload justifies it.

There is also a resilience benefit. When work can be distributed among genuinely separate operators, infrastructure has more routes available when individual participants become unavailable.

None of that happens simply because a network has a large node count.

Thousands of machines provide limited decentralisation if they ultimately depend on the same operator, critical service or point of control. Useful decentralised infrastructure requires clear participation rules, reliable coordination and mechanisms for detecting incorrect or unavailable work.

The objective should therefore be more precise than decentralising compute for its own sake.

Use decentralisation where independence, verification or distributed participation improves the service.

The opportunity reaches beyond crypto

Distributed computing is not a new idea, nor is its usefulness limited to digital assets.

BOINC allows people to contribute spare computing capacity to scientific research across fields including astronomy, climate science and biomedicine. Folding@home has coordinated distributed computing for biomedical research since 2000.

Their architecture, economics and purpose differ substantially from blockchain networks, but they demonstrate something fundamental: useful computing work can be coordinated across machines that are not sitting inside one data centre.

The potential applications extend into ordinary business infrastructure.

Software teams can distribute suitable testing workloads across different machines and environments. Organisations relying on a shared calculation can introduce independent verification. Researchers can divide appropriate jobs into batches and distribute them across available resources.

Industrial systems create another opportunity.

Processing information closer to the devices producing it can reduce some of the latency and data-management challenges created by sending everything to a distant cloud. NIST's work on fog computing has documented this model of placing distributed compute and analytics throughout a network rather than exclusively at its centre.

AI expands the conversation further, although it also demonstrates why different workloads require different infrastructure.

Training very large models generally requires tightly connected accelerators exchanging large quantities of information. Research such as DiLoCo explores ways to reduce those communication requirements, showing both how difficult distributed AI training can be and how rapidly the field continues to develop.

A network of ordinary computers should therefore be judged by the workloads it can perform well, not by pretending it is interchangeable with a purpose-built GPU cluster.

That distinction matters to us.

The strongest decentralised compute products will not succeed because the word decentralised appears in their description. They will succeed when customers can submit useful work, receive a dependable result and obtain the level of evidence the task requires.

The underlying architecture may be highly distributed.

The customer experience should not have to be complicated.

What nearly a decade of building has taught us

Our journey started in 2017 with blockchain-based ticketing.

The problem was practical. We wanted to improve the relationship between ticket issuers and ticket owners through better oversight, communication, standardisation and lower operational friction.

That work led to Aventus Network and, over time, infrastructure serving a much broader range of enterprise applications.

We have since worked across rewards and loyalty with Enigmatic Smile, energy infrastructure with Energy Web, aviation and logistics, supply chains and prediction markets.

Each deployment brought different operational requirements, users and commercial expectations.

It also taught us something important about infrastructure.

The technology has to accommodate the business, not the other way around.

Costs need to make sense at transaction level.

Integrations need to work with systems companies already use.

Software needs to remain maintainable after launch.

And the people relying on the service need to know what happens when something fails.

Working through those problems shaped the way we build today.

It also exposed the limits of any single layer.

A blockchain can establish agreement about the state recorded on that blockchain. Useful applications frequently depend on much more.

They need information originating elsewhere.

They need rules applied to that information.

They need results checked.

They need different systems to coordinate.

As applications become more sophisticated, more of the work surrounding the chain becomes important.

That progression is what led us to Aventus Cloud.

What Aventus Cloud adds

We launched Aventus Cloud on mainnet in April 2026.

Today, more than 2,500 nodes are live, and the network is designed to scale towards 40,000.

We are deliberately focused on workloads where distributed execution or verification has a clear purpose.

Aventus Cloud currently supports three core categories of work.

Decentralised oracle services allow an application to request information it cannot establish onchain itself. That might include an asset price, a prediction-market outcome or another external event. Nodes can inspect the relevant information and return a consensus result.

Summary verification allows Cloud nodes to check blockchain state summaries before another system relies on them. We use this mechanism today to verify the Aventus Parachain's weekly summary to Ethereum.

Custom business logic execution allows an application to register specialised logic that Cloud nodes can execute when required, rather than placing every piece of application logic inside the chain or behind one centrally operated backend.

The model gives developers access to a distributed operator network without requiring them to recruit and coordinate an entirely new community themselves.

Any compatible blockchain or system that can submit a work request and read the returned result can use the network.

That is important because building an application is already difficult.

Building the application, creating a separate compute network, recruiting operators, designing incentives and maintaining all of that infrastructure is a very different undertaking.

Shared infrastructure allows teams to concentrate on the product while using a network that already exists.

Participation is deliberately lightweight. Our node software can run on ordinary computers and laptops across macOS, Windows and Linux, with Docker available for server environments.

We want participation to be accessible without pretending every computer is suited to every workload.

As Aventus Cloud develops, matching the right workloads to the capabilities of the network will be one of the most important parts of the job.

One ecosystem, different infrastructure needs

There will never be one architecture suitable for every application.

Some organisations need a dedicated blockchain in which business logic, governance and economics are designed around the application itself.

Our Enterprise Appchains provide that option through configurable Substrate-based infrastructure and can connect directly to Aventus Cloud when decentralised compute, oracle services or summary verification are required.

Other teams may already have their application or blockchain and only need a particular piece of work independently executed or checked.

They should not need to rebuild an entire infrastructure stack to obtain that capability.

This is where our experience becomes useful.

We have spent years working at both levels of the problem: the applications businesses actually want to operate and the blockchain infrastructure beneath them.

That gives us a practical framework for deciding where decentralisation adds value, which workloads the network should support and what developers need in order to integrate them.

The next phase of Aventus Cloud should ultimately be measured by something much more meaningful than integrations announced.

Repeat work.

An integration tells us somebody was willing to connect.

Repeated workloads tell us the service is useful enough to keep using.

That is the standard that matters.

The people and economics behind the network

Distributed infrastructure only exists when people choose to operate it.

Node operators contribute machines, availability and time. The network needs clear rules for coordinating those contributions and an economic system capable of sustaining them.

AVT has defined roles across our ecosystem.

It is used for transaction fees on Aventus Network and for participation in governance. Active Aventus Cloud nodes receive AVT through the network's reward model, while operators can optionally stake AVT to increase node reward weight.

Those functions connect AVT to the operation and governance of the infrastructure.

Their long-term significance will depend on what the infrastructure itself accomplishes.

Reward mechanisms can help establish network capacity.

Demand gives that capacity a purpose.

We should therefore judge progress by whether developers and businesses repeatedly choose to use the network, whether operators can participate on understandable terms and whether the economics remain transparent as the ecosystem grows.

Governance belongs in the same conversation.

Decisions about infrastructure affect the people contributing to it. Changes to incentives, treasury allocations and development priorities deserve scrutiny.

Our governance framework allows community proposals and onchain voting, including changes to parameters governing Aventus Cloud and AVT.

For governance to have substance, participants need enough information to make informed decisions.

Proposals should state the problem being addressed.

Expected benefits, costs and risks should be visible.

Progress should be reported in terms the community can evaluate.

Changes should come with an explanation.

Our community contributes by operating nodes, scrutinising proposals, developing ideas and challenging assumptions. Long-standing members also carry institutional knowledge about previous decisions and commitments that can easily disappear as a project evolves.

That participation is valuable.

An infrastructure network built around distributed participation should be willing to apply the same principle to its own direction.

The future we are building towards

The future of compute will not belong to one architecture.

Large data centres will remain essential for demanding workloads.

Private infrastructure will remain appropriate when control, regulation or confidentiality requires it.

Edge systems will continue moving computation closer to the devices and data that need it.

Distributed networks can create another layer for workloads that benefit from wider participation, independent execution or separate verification.

Good infrastructure will put each job in the environment that suits it.

Security, privacy, latency, energy use, communication overhead, cost and verifiability all belong in that decision.

This is the future we are building towards at Aventus.

Nearly a decade of enterprise blockchain work has given us experience turning distributed systems into infrastructure used by real organisations. Aventus Cloud gives us a live network on which we can extend that experience into decentralised compute. AVT and community governance give participants ways to contribute to the operation and direction of the ecosystem.

The opportunity now is to make that infrastructure increasingly useful.

That requires disciplined development, straightforward integrations and close attention to the organisations and developers using what we build.

Every useful workload gives someone another reason to use Aventus Cloud.

Every successful deployment teaches us something that can improve the next one.

We started in 2017 with a practical question about how technology could create greater confidence in a ticket and the transactions surrounding it.

Nearly a decade later, the scale of the question has changed:

How can organisations rely on computation performed across infrastructure they do not fully control?

Our answer is not that every workload should be decentralised.

It is that when independence, verification and wider participation genuinely improve the outcome, the infrastructure to support them should be available.

Building that infrastructure is the next chapter of Aventus.

Aventus builds enterprise blockchain infrastructure and decentralised compute for the next generation of onchain applications.