

Quali Whitepaper

Federated Infrastructure: The New Economics of IT Infrastructure

Turning infrastructure from a fixed cost into a reusable asset that can be shared, extended, valued and monetized.

By Lior Koriat, Quali CEO and Edan Evantal, Quali CTO

Contents

1. The New Economics of IT Infrastructure	3
2. Federated Infrastructure Requires a Common Language	4
3. From Technical Complexity to Business Value	5
4. Following the Money: The Economics of Federated Infrastructure	7
5. The Federated Infrastructure Economy	9
6. The Strategic Cost of Standing Still	12
7. From Vision to Reality	17

1. The New Economics of IT Infrastructure

For most organizations, IT infrastructure is the cost of doing business. Compute, networks, storage, cloud services, GPU capacity and increasingly complex technology stacks are acquired to support applications, users and business services. The value of that infrastructure is usually measured indirectly, by what it runs, and the business outcomes associated with the applications it supports.

Meanwhile, the infrastructure itself costs money to acquire, operate, update, and maintain. Capacity is distributed across technology and organizational silos. Utilization varies. Specialist skills are required to manage different platforms, and the infrastructure stack that one organization has integrated and validated is difficult to make available to another organization that needs the same capability.

Federated Infrastructure changes that economic model.

Infrastructure can become something an organization not only consumes but provides. Available capacity, specialized infrastructure, and complete validated environments can be made available to another organization. Applications, services, operational expertise, and intellectual property can be added, increasing their value. That capability can be consumed directly, extended with additional services and made available again, creating new opportunities for infrastructure owners, technology providers and consumers alike.

None of this requires an organization to redesign its infrastructure estate or suddenly become an infrastructure provider. Federation can begin much more pragmatically. An enterprise might start by making a validated GPU environment available to another internal team, allowing a business unit to consume infrastructure operated elsewhere in the organization, extending a proven environment to a partner, or consuming a specialized capability from a provider rather than building it internally. The underlying infrastructure remains where it is owned and operated by the teams that understand it.

What changes is the way those resources are defined, governed and consumed. Instead of requiring every new relationship to introduce another set of tools, integrations, credentials, processes, and specialist skills, the differences underneath are absorbed into a consistent operating model. Organizations can therefore federate a specific capability first, prove the operational and economic value, and extend the model as additional providers; consumers and infrastructure become relevant.

Federation is not an objective that requires an enterprise-wide transformation before any value can be realized. It is an operating model that can be introduced incrementally, using infrastructure and organizational relationships that already exist. The first benefit may be as simple as making an expensive resource easier to use across the business. The same foundation can subsequently support partner delivery, external consumption, value-added services and monetization.

For some organizations, this creates the opportunity to move infrastructure from being purely a cost center toward a source of revenue. For others, the return comes from higher utilization, avoiding duplicated investment, gaining temporary access to scarce resources, or reusing specialized infrastructure and expertise across business units, partners and customers.

The critical change is that the value of infrastructure is no longer measured only by what it supports. It can also be measured by the value the infrastructure itself can deliver.

WHEN FEDERATED INFRASTRUCTURE BECOMES STRATEGIC

Different business drivers can make infrastructure federation a strategic priority.
The same operating model can address multiple needs.




Figure 1. Avenues to a Federated Infrastructure

2. Federated Infrastructure Requires a Common Language

The infrastructure underneath a federation will never be uniform. It may include public and private clouds, physical infrastructure, Kubernetes environments, GPU platforms, networks, storage, edge systems, and specialized technology stacks. Each has different APIs, credentials, architectures, terminology, operational models and lifecycle requirements.

Federation cannot depend on eliminating those differences. Instead, those technologies must be made consistently usable.

That requires an infrastructure control plane capable of translating heterogeneous infrastructure into a common operating model through which infrastructure can be defined, governed, delivered, consumed, and operated. The objective is not to standardize the infrastructure. It is to standardize the way infrastructure becomes usable.

A data scientist should not need to understand the infrastructure architecture behind every GPU environment they consume. An application team should not need a new operating model because an environment moves from an internal data center to a partner infrastructure. Developers, application owners, infrastructure teams, and external consumers may have different reasons for using infrastructure, but they need to be able to interact with it through the same language.

The complexity underneath does not disappear. It moves into the control plane. The control plane absorbs the diversity of the underlying infrastructure and turns its capabilities into consistent, understandable and governed building blocks. Policy, identity, security, lifecycle, capacity, isolation, metering and operational control become part of the common model rather than something every provider and consumer must reconstruct for each relationship.

This distinction is fundamental. Connecting multiple clouds, data centers, or infrastructure platforms does not create federation. Infrastructure crossing technology and organizational boundaries must carry consistent governance and operational controls across those boundaries as well.

An infrastructure provider must be able to determine what it makes available, who can consume it, under what conditions and with what capacity. Its credentials, intellectual property, and underlying implementation must remain protected. Consumers must be isolated from one another, while policy, security and lifecycle behavior remain predictable regardless of who owns the infrastructure or where it executes.

Capacity must be understood and, when necessary, reserved. Consumption must be measured and attributed. Environments must remain governed beyond their initial deployment, and the common operating model must remain reconciled with the infrastructure delivering the service.

These are the foundations that make federation credible. Heterogeneous infrastructure can remain heterogeneous, independently owned and operated, while becoming consistently consumable across organizational and technology boundaries.

The result is more than simplified infrastructure management. It creates the technical foundation upon which infrastructure can be shared, extended, and ultimately treated as a transferable business asset.




Figure 2 - Federated Infrastructure Architecture

3. From Technical Complexity to Business Value

The significance of Federated Infrastructure becomes clearer when viewed as a progression of three transformations. Each builds on the previous one, moving infrastructure from a collection of independently managed technical resources toward a model in which infrastructure capabilities can be consumed, exchanged and valued across organizational boundaries.

The First Transformation: Simplifying Technical Complexity

Infrastructure is heterogeneous, fragmented, and complex. Different technologies require different skills, tools, integrations, and operational knowledge. This disparity makes infrastructure difficult to consume consistently and even harder to share safely between organizations.

An infrastructure control plane changes that relationship by absorbing the differences underneath and presenting a common language above them. The infrastructure does not become less sophisticated, and the technologies underneath do not become the same. Instead, unnecessary complexity is neutralized for the people and organizations consuming it.

The result is a consistent way to define, deliver and operate infrastructure regardless of its underlying technology, ownership, or location. Infrastructure providers retain control of their resources, while consumers gain access to capabilities without having to understand or reproduce the complexity behind them.

The Second Transformation: Changing the Operating Model

Technical simplification enables a more significant operational change. Infrastructure that previously required specialist knowledge, bespoke integration, and dedicated operational processes can become a reusable business capability.

Instead of every team or organization having to acquire, understand, integrate and operate infrastructure independently, capabilities can be packaged into consistently defined, governed and consumable Offerings.

An Offering establishes what is being delivered, the infrastructure on which it can execute, the policies governing its use and the conditions under which it can be consumed. The consumer receives a defined capability rather than unrestricted access to the infrastructure behind it.

This changes the relationship between infrastructure owners and consumers. A provider can make a validated environment available to another business unit, partner or customer while retaining control of its underlying resources. The consumer can use that capability without reproducing the provider's infrastructure, expertise, or operational processes.

Infrastructure has moved beyond being a technical resource that must be individually acquired and managed. It has become something that can be packaged, delivered, and reused across organizational boundaries.

The Third Transformation: Changing the Economics

The first two transformations establish something fundamentally different from traditional infrastructure management. Once infrastructure can be consistently defined, governed and consumed across organizational boundaries, its economic value is no longer confined to the organization that owns it or the applications it was originally purchased to support.

Available capacity can become a commercial service. A validated architecture can become a product. Specialized infrastructure can be consumed without being purchased. Operational expertise, software, applications, and intellectual property can be combined with existing infrastructure to create increasingly valuable Offerings.

More importantly, the value does not have to stop with the first transaction.

An organization consuming a federated Offering can add its own capabilities and make the resulting Offering available to another consumer. Infrastructure supplied by one provider can support software from another, managed services from a third, and applications from a fourth. Each participant contributes additional value without necessarily owning or operating the infrastructure underneath.

The result is a model in which infrastructure value can accumulate as capabilities move between providers and consumers. The original infrastructure owner can derive economic value from resources already deployed, while other participants create additional commercial value through the capabilities they contribute.

Federation does not make infrastructure valuable. It makes infrastructure value transferable.

That distinction changes the economics of infrastructure. Resources, architectures, expertise and services that were previously tied to individual organizations can become components of a broader system of economic exchange.

But making value transferable is only the beginning. For this model to operate commercially, the value contributed by each participant must be measurable, attributable, and capable of supporting independent commercial relationships.

That is where Federated Infrastructure moves beyond a new operating model and begins to establish the foundations of a new infrastructure economy.

A consistent operating model turns heterogeneous infrastructure into reusable, governed capabilities that can be consumed, extended and monetized.



Figure 3. From Technical Complexity to Business Value.

4. Following the Money: The Economics of Federated Infrastructure

The economic significance of Federated Infrastructure becomes clearer when we follow a capability from its original infrastructure owner through the organizations that add value to it, and ultimately to the customer who consumes it.

Consider a sovereign cloud provider operating a substantial investment in GPU infrastructure. The provider has acquired the hardware, built the data center, established security and compliance controls, and developed the operational expertise required to deliver reliable compute capacity. Yet utilization fluctuates. Capacity that is valuable and expensive to maintain may remain unused between customer workloads.

Through federation, the provider can make available a defined portion of that capacity as an Offering. Rather than exposing its infrastructure directly, it publishes a governed capability with specified GPU resources, availability, reservation conditions, security requirements and measurable consumption. The provider retains ownership and operational control while making its infrastructure commercially consumable by another organization.

A technology partner consumes that Offering and adds a validated AI software stack, including the necessary frameworks, configurations and software licenses. The partner has now created something more valuable than the underlying compute capacity: a ready-to-use AI environment that another organization can deploy without independently assembling and validating the technology.

A managed service provider can extend the same capability further, adding monitoring, security operations, service assurance, compliance management and technical support. The resulting Offering is no longer simply GPU capacity or an AI software environment. It is a managed AI capability assembled from the contributions of several independent organizations.

The final customer consumes the complete Offering. The customer pays for the delivered capability under the agreed commercial terms, without having to establish separate infrastructure and operational relationships with every organization involved in delivering it.

The Offering Becomes the Product

The fundamental commercial unit in this model is the Offering.

An Offering defines what the consumer receives, the resources and services required to deliver it, the conditions governing its use and the basis on which consumption can be measured. It separates the product being consumed from the underlying infrastructure and organizations contributing to its delivery.

This distinction allows infrastructure to participate in a commercial value chain without requiring every participant to own the complete technology stack.

The sovereign cloud provider supplies GPU capacity. The technology partner supplies software, licensing and validated configurations. The managed service provider supplies operational expertise and service commitments. Each contributes to the resulting Offering and establishes the commercial terms under which its contribution is made available.

The customer purchases the accumulated value, not simply the underlying hardware.

How Value Is Measured and Money Flows

For this economic model to function, consumption must be measurable and attributable to the capabilities supplied by each participant.

The underlying infrastructure provider may charge for reserved GPU capacity, actual GPU-hours consumed or another agreed measure of resource availability. The software partner may apply license entitlements, subscription charges or usage-based fees. The managed service provider may charge for operational services, service levels or the duration of the managed environment.

These commercial models do not have to be identical. What matters is that the consumption associated with an Offering can be understood and attributed to the participating providers.

In our example, the customer pays the managed service provider for the complete environment. The managed service provider compensates the technology partner under their commercial agreement, while the technology partner pays the sovereign cloud provider for the infrastructure capacity consumed or reserved. Each organization earns revenue from the value it contributes, while retaining responsibility for its own resources, services and commercial relationships.

The federation provides the common operational foundation required to deliver and measure the capability across these organizational boundaries. Commercial agreements and billing systems determine how the resulting charges are calculated and settled.

The significance is that each participant can act as both consumer and provider. An organization can consume an existing Offering, add value, and publish a new Offering incorporating those capabilities. The infrastructure remains independently owned, but its economic value can extend through multiple commercial relationships.

From Infrastructure Cost to Economic Asset

This changes the economics for every participant.

The sovereign cloud provider can generate revenue from capacity that might otherwise remain underutilized. The technology partner can commercialize its validated architecture and intellectual property without owning the GPUs. The managed service provider can deliver a differentiated service without acquiring the entire infrastructure stack. The customer gains access to a complete capability without reproducing the capital investment, integration work and operational expertise behind it.

More importantly, the economic opportunity is not limited to a single provider-consumer transaction.

The same infrastructure capacity can support different Offerings, each incorporating different software, services, operating models and commercial terms. A validated environment can be reused across

customers. Specialized expertise can be commercialized repeatedly. Additional participants can contribute new capabilities without rebuilding the underlying infrastructure.

Federated Infrastructure therefore creates the conditions for infrastructure, software and services to participate in a shared system of economic value creation.

The first commercial relationship may be straightforward: one provider, one partner and one customer. But as more organizations participate, the relationships can become considerably more sophisticated. Consumption of one Offering may generate demand for several underlying Offerings, each supplied by a different organization and governed by its own licensing, metering and commercial arrangements.

At that point, the economic model begins to extend beyond a conventional infrastructure supply chain. It becomes the foundation for a distributed infrastructure economy.

Multiple participants contribute to a single offering and are compensated based on the value they provide.




Figure 4. Following the Money Flow: How a Federated Infrastructure Economy Works

5. The Federated Infrastructure Economy: When Value Moves Across the Mesh

The commercial relationships described so far follow a relatively simple progression. An infrastructure provider makes capacity available. A technology partner adds software and expertise. A service provider adds operational capabilities. A customer consumes the resulting Offering, and each participant earns from the value it contributes.

But there is no architectural reason for federation to stop at a linear chain of providers and consumers. Once infrastructure capabilities can be consistently defined, governed, consumed and measured across organizational boundaries, the same model can extend across many independent providers, each participating in multiple commercial relationships. An organization can be a consumer of one Offering, a provider of another and a contributor to several others simultaneously.

This changes the scale of the opportunity. Federation is no longer simply a way to make existing infrastructure more accessible or generate additional revenue from unused capacity. It becomes a framework through which infrastructure, software, services and intellectual property can participate in a distributed economy.

From Supply Chains to an Infrastructure Mesh

Consider an enterprise consuming an AI application delivered as a managed service. The application uses an inference platform supplied by one provider, a validated software stack from another, GPU

capacity operated by a sovereign cloud, storage from a regional infrastructure provider and security services supplied by a specialist partner.

The customer may have a commercial relationship with only the application provider. Yet delivering that application depends on capabilities contributed by several independent organizations. Each of those organizations may itself consume capabilities from other participants in the federation. The inference provider might use GPU capacity from several infrastructure operators. The application provider might combine inference services with licensed data services and specialized AI models. The managed service provider might deliver the same application across different sovereign regions to satisfy customer requirements.

The result is not a fixed hierarchy of infrastructure ownership. It is a network of independently operated capabilities assembled according to the requirements of the service being delivered. An Offering becomes a building block that can participate in other Offerings. Its value is no longer limited to the original provider-consumer relationship. It can become part of multiple services, commercial arrangements and business models without requiring the underlying infrastructure to change ownership or location.

Independent providers, software vendors and service providers combine capabilities into composite Offerings that can be consumed, extended and re-used across organizations.



Figure 5. The Federated Infrastructure Economy

When Consumption Creates a Cascade of Economic Value

The implications become particularly significant when consumption is measured across the complete chain.

A customer interacting with an AI application may generate a request measured in tokens. Fulfilling that request consumes inference capacity, which in turn requires GPU resources, memory, networking, storage and supporting software services. Some resources may be charged according to actual usage. Others may be reserved, licensed, subscribed to or priced according to service-level commitments. These measurements are different, but they are economically connected.

A single customer transaction can therefore generate multiple consumption events and commercial obligations across the federation. The application provider may charge for tokens or completed transactions. The inference provider may charge for inference capacity. The infrastructure provider may charge for GPU-hours or reserved capacity. Software vendors may apply licensing or usage entitlements. Managed service providers may charge for operations, security and availability.

The customer does not necessarily need to understand or manage these underlying relationships. The commercial model can present a single consumable service while the contributing providers retain their own pricing structures, contractual arrangements and revenue streams.

For this to operate at scale, the infrastructure economy requires more than a record of what was deployed. It needs a consistent relationship between the Offering being consumed, the underlying capabilities contributing to its delivery and the consumption attributable to each participant. That relationship provides the basis for commercial systems to calculate charges, apply entitlements and reconcile obligations between providers.

The economic significance is substantial. Revenue generated by one service can support multiple independent businesses contributing to its delivery, while the same underlying infrastructure and software capabilities can participate in many different services.

Infrastructure Becomes Part of a Distributed Economic System

As participation expands, the federation can evolve beyond bilateral infrastructure-sharing arrangements into a network of reusable commercial capabilities.

A sovereign infrastructure provider could make GPU capacity available to several AI platform operators. Those platforms could support hundreds of application providers. Managed service organizations could combine those applications with regional infrastructure, security and compliance services. Enterprises could consume the resulting capabilities directly or extend them with their own intellectual property and make them available to others.

Each participant contributes something distinct. Each can retain control of its resources, technologies and commercial terms. And each can derive value from capabilities that extend beyond its own organizational boundaries.

The network does not require every provider to use the same infrastructure, adopt the same pricing model or participate in a single marketplace. What it requires is a common operating model through which capabilities can be defined, governed and consumed, together with the metering, attribution and commercial mechanisms necessary to support the relationships between them.

The infrastructure control plane provides the technical foundation for this economy. It maintains consistency across heterogeneous infrastructure, governs the lifecycle of consumable capabilities and provides the operational context through which consumption can be understood. Commercial systems can build upon that foundation to manage pricing, licensing, billing and settlement.

As those capabilities develop, infrastructure providers gain access to demand beyond their existing customer relationships. Software vendors gain new ways to distribute and monetize validated technology. Service providers can assemble differentiated offerings without owning every underlying resource. Enterprises can consume complete capabilities assembled from infrastructure and expertise distributed across multiple organizations.

The result is a fundamental change in how infrastructure participates in the creation and distribution of economic value.

A New Economic Model for IT Infrastructure

For decades, infrastructure economics have largely been shaped by ownership, capacity planning, utilization and the cost of delivering services. Cloud computing changed how infrastructure could be acquired and consumed, but the commercial relationship has generally remained centered on the provider delivering the service and the customer purchasing it.

Federated Infrastructure introduces a different possibility.

Infrastructure can remain independently owned and operated while becoming part of a broader network of consumable capabilities. Software, services and intellectual property can be added by organizations that do not own the underlying resources. Consumption can be measured across the contributing layers, allowing commercial value to be attributed beyond the organization delivering the final service.

This is not simply infrastructure sharing at a larger scale. It changes the relationship between infrastructure ownership, service delivery and economic participation.

An organization no longer needs to own the complete technology stack to create and commercialize a valuable infrastructure-based service. An infrastructure provider no longer needs a direct relationship with every customer ultimately benefiting from its resources. A validated capability can become part of multiple products, delivered through different providers, in different locations, under different commercial arrangements.

The ultimate expression of Federated Infrastructure is therefore not a marketplace of spare compute capacity. It is a distributed infrastructure economy in which independently owned resources, software and services can be assembled into consumable products, with value and revenue extending across the organizations that contribute to them.

The customer consumes one capability. An entire infrastructure economy can participate in delivering it.

6. The Strategic Cost of Standing Still

The transition toward a federated infrastructure economy will be uneven. Some organizations will build the capabilities and partnerships to shape it. Others will wait for commercial proof while preserving the option to enter. Still others will continue to invest primarily in the established model of owning, integrating and selling infrastructure within familiar boundaries. Each choice carries a different combination of opportunity, cost and risk.

For infrastructure providers, cloud operators, software vendors, managed service providers and enterprises, the question is no longer only whether federation improves utilization or simplifies operations. It is whether participation in a broader network of consumable capabilities becomes a competitive advantage.

That advantage may come from distribution, differentiated services, partner relationships and the ability to combine independently supplied capabilities, not simply the cost or performance of the underlying technology.

If federation becomes an established commercial model, organizations will compete both on what they own and on the economic relationships in which their capabilities can participate.

Three Strategic Responses to an Emerging Economy

Organizations are likely to approach this transition from three broad strategic positions.

The Architects: Shape the Economy

The first group will actively develop the technical, operational and commercial foundations needed to participate in federation.

For infrastructure providers, this may mean making capacity available through standardized, governed Offerings that can be consumed by multiple partners. For software vendors, it may mean packaging validated technologies that can operate across different infrastructure providers. Managed service organizations may develop services that combine independently supplied infrastructure, software and operational expertise into differentiated products.

These organizations are not simply adopting a new technology architecture. They are building the relationships and capabilities required to participate in an expanding economic ecosystem. Their advantage is the opportunity to establish early distribution channels, develop partner relationships, influence operating conventions and gain practical experience before the market matures.

An infrastructure provider that participates in several federated relationships may gain access to demand generated by software vendors, AI platform operators and service providers that would otherwise be beyond its direct commercial reach.

A software vendor may expand distribution without having to operate the underlying infrastructure. A managed service provider may deliver a broader portfolio without financing every component of the technology stack.

The potential reward is access to new revenue streams, improved capital efficiency and greater commercial reach. The corresponding risk is that early investment may precede market demand, and some operating models or partnerships may not deliver the expected returns.

The Observers: Preserve Optionality

A second group will recognize the opportunity but wait for stronger evidence of commercial adoption.

This can be a rational strategy. Infrastructure investments are expensive, commercial models take time to mature, and organizations have legitimate reasons to avoid committing resources before the economic benefits are established.

However, there is an important distinction between waiting to make a major investment and waiting to develop the ability to participate.

Organizations that preserve optionality can begin standardizing infrastructure definitions, improving lifecycle automation, establishing metering and attribution capabilities, and evaluating potential partnerships without immediately changing their business models. Those that simply defer the discussion may discover that the technical foundations, commercial relationships and market conventions have developed without them.

By the time the economic model is proven, competitors may already have established preferred relationships with infrastructure providers, software vendors and distribution partners. The risk is not necessarily that late entrants will be excluded. It is that they may enter with less influence, higher integration costs and fewer opportunities to shape the commercial relationships on which they depend.

The Defenders: Protect the Existing Model

The third group will continue relying primarily on established approaches to infrastructure acquisition, integration, ownership and commercialization. For some organizations, dedicated infrastructure will remain appropriate. Security, sovereignty, performance, regulatory requirements and predictable utilization can all justify direct ownership and tightly controlled operating environments.

The strategic risk arises when an organization assumes that this must remain the only economically viable model. Traditional infrastructure strategies often require organizations to finance capacity, integrate technologies, maintain specialized operational skills and establish direct commercial relationships with customers.

Competitors participating in a federation may be able to assemble equivalent or more specialized capabilities from multiple providers without reproducing the complete infrastructure investment. For an infrastructure provider, the challenge may be particularly significant. Owning high-quality infrastructure does not guarantee access to the broadest market for that infrastructure.

A provider with limited distribution may find itself competing against another provider whose capacity participates in numerous independently delivered services.

Both may own comparable technology. Both may have similar operating costs. Yet one can potentially generate demand through an ecosystem of commercial relationships while the other remains dependent on its own sales channels.

The competitive distinction begins to shift from who owns the infrastructure to whose infrastructure can participate in the greatest number of valuable commercial relationships.

The Business Consequences of Strategic Choice

These three approaches create different economic possibilities.

Organizations will take different approaches to the federated infrastructure economy, each with distinct opportunities and risks.



Figure 6A. Three Strategic Positions: Architects, Observers and Defenders.

The business implications extend across several dimensions:

Business dimension	Architects	Observers	Defenders
Revenue opportunity	Develop new Offerings, partnerships and revenue channels	Enter once demand is demonstrated	Depend primarily on existing products and customers
Capital efficiency	Access or monetize capacity across organizational boundaries	Make selective investments while evaluating the model	Continue funding and managing dedicated capacity
Market access	Extend distribution through participating providers and partners	Join more established networks later	Rely heavily on direct commercial relationships
Differentiation	Combine specialized infrastructure, software and services into new products	Adopt validated approaches as they mature	Compete through existing capabilities and service models
Strategic influence	Help establish relationships and operating conventions	Adapt to emerging conventions	Risk dependence on intermediaries controlling customer access
Principal risk	Premature investment and execution uncertainty	Lost learning, relationships and market position	Structural cost and distribution disadvantages if the market changes

These are potential outcomes, not predictions. Commercial demand, interoperability, governance, security and economic incentives will determine where federation is viable.

The distinction is economic as much as technical: participants may be able to create and capture value beyond the boundaries of what they directly own.

Early Evidence: The Market Is Already Moving

The complete federated infrastructure economy described here is not yet an established industry-wide model. But several of its constituent commercial practices are already visible.

These examples illustrate different pieces of the opportunity—not proof of a unified, automated multi-party settlement system.

- **NVIDIA DGX Cloud Lepton** illustrates the growing importance of aggregating independently operated GPU infrastructure and making it accessible through a broader commercial channel. By connecting AI developers with participating infrastructure providers, NVIDIA demonstrates how infrastructure demand can extend beyond an individual provider's existing customer relationships. The significance is not simply the ability to locate available GPUs. It is the commercial opportunity created when infrastructure providers can participate in a larger network of demand.
- **AWS Marketplace and Channel Partner Private Offers** demonstrate another important component: commercial transactions involving multiple participating organizations. Software vendors, channel partners and customers can participate in structured commercial relationships in which value is distributed between parties.

This is not an open, cross-provider infrastructure federation. It does, however, establish that multi-party commercial participation, intermediary distribution and coordinated financial transactions are already practical business models.

Distributed compute marketplaces, including Akash Network, offer another indication of demand for infrastructure supplied by independent operators. Their architectures and commercial mechanisms differ from the enterprise federation model discussed here, but they illustrate the broader movement toward making distributed compute resources commercially accessible beyond their original owners.

Taken together, these examples show infrastructure aggregation, intermediary distribution and independent resource supply. None by itself establishes the cross-provider governance, attribution and commercial settlement model described in this paper.

The Cost of Waiting Is Not Always Visible

The cost of waiting is difficult to see because it rarely appears as an immediate line item in a financial statement. Existing customers still buy, infrastructure still runs and capital budgets still follow familiar cycles. Meanwhile, competitive positions can shift.

- **Lost distribution opportunities.** Providers that establish early relationships with software vendors, AI platforms and managed service organizations may become preferred sources of infrastructure capacity. Competitors entering later may need to displace established relationships or accept less attractive commercial terms.
- **Higher integration costs.** Organizations that continue building highly customized, isolated infrastructure environments may accumulate technical and operational dependencies that make later participation more expensive. Standardizing infrastructure capabilities and automating their lifecycle becomes harder when every environment has developed independently.
- **Reduced commercial influence.** Early participants can help define how capabilities are packaged, measured, governed and commercially consumed. Later participants may need to accommodate operating conventions and commercial arrangements established by others.
- **Capital allocation risk.** Infrastructure investments made today may remain productive for years. If future demand increasingly flows through federated relationships, organizations that cannot expose their infrastructure through those channels may face a narrower market for assets they have already financed.
- **Loss of customer influence.** Perhaps the most consequential risk is that the organization supplying the underlying infrastructure may no longer control how that infrastructure reaches the customer.

As infrastructure capabilities become components of larger Offerings, the organization assembling and delivering the final service may own the customer relationship, determine the commercial proposition and influence which underlying providers participate.

Infrastructure ownership remains valuable. What changes is the importance of distribution and the ability to participate in services sold by others.

When Comparable Infrastructure Produces Different Business Results

Consider two regional cloud providers with comparable GPU infrastructure and operating standards.

Both have invested in compute, networking, facilities and specialized teams. Both must earn a return on those assets. The first provider continues selling infrastructure capacity directly to customers through its established sales organization.

The second provider retains the same ownership and operational control but also makes validated capacity available through federated Offerings. AI platform operators, software vendors and managed service providers can incorporate that capacity into services sold to their own customers.

The first provider's addressable demand depends largely on its direct commercial reach. The second can also reach demand generated by partners that package its capacity into their own services.

If those channels attract customers, the same class of infrastructure asset could produce different utilization, acquisition costs and returns. That outcome is possible, not guaranteed: partner margins, integration costs and commercial terms still matter.

The principle extends beyond GPU providers.

A software vendor that can distribute validated technology through multiple infrastructure partners may reach customers it could not serve independently. A managed service provider that can assemble capabilities from several sources may expand its portfolio without equivalent capital investment. An enterprise that can consume governed capabilities across multiple providers may reduce the need to reproduce specialized infrastructure in every business unit or geography.

In each case, the economic advantage comes from separating the ability to create and deliver value from the requirement to own every component involved.

Preparing for Federation Without Betting the Business

The strategic implications do not require organizations to abandon existing infrastructure investments or commit immediately to an unproven commercial model.

There is a practical middle ground between full participation and passive observation.

Organizations can begin by identifying infrastructure and software capabilities that could become reusable Offerings. They can standardize how those capabilities are defined and deployed, establish consistent governance and lifecycle management, and improve their ability to measure consumption and attribute resource usage.

They can evaluate where external infrastructure or partner-provided capabilities could improve service delivery, and where their own technology or operational expertise might become commercially valuable to others.

These preparations can improve operational consistency and flexibility today, even if broad commercial federation develops slowly. An organization whose infrastructure remains fragmented, manually integrated and difficult to expose beyond its own operating boundaries may find that the decision has effectively been made for it.

Different approaches to the federated infrastructure economy lead to very different outcomes.




Figure 6B. Three Strategic Paths: Illustrative Economic Opportunity Scenarios Over Time

Adoption will be uneven. Some workloads will remain better suited to dedicated infrastructure, and early participants will face genuine execution and investment risk. The trajectories in Figure 6B illustrate possible outcomes, not forecasts.

The strategic question is whether today's infrastructure and technology investments preserve the ability to participate in tomorrow's commercial relationships.

An organization can choose when to participate. It cannot assume that its competitors, partners and customers will wait for it.

7. From Vision to Reality: Building the Federated Infrastructure Economy

The case for federation is ultimately a case for making infrastructure economically useful beyond the boundaries of its owner. That requires operational foundations capable of turning complex environments into dependable, governed Offerings, and commercial relationships capable of rewarding the organizations that contribute to them.

The first part of that transition is no longer purely theoretical.

Infrastructure automation is moving from isolated provisioning tasks toward the repeatable delivery of complete infrastructure and software environments. The challenge now is to make those environments consistent enough to be consumed and governed across different teams, technologies and providers. Without that repeatability, each new provider relationship becomes another bespoke integration project.

Establishing the Technical Foundation

Quali's Torque platform was developed around this principle: transforming complex infrastructure and software environments into standardized, reusable Offerings that can be provisioned, operated and governed through a common control plane.

Rather than requiring every consuming organization to reproduce the underlying integration and operational expertise, the Offering captures the infrastructure definition, associated software components, configuration and lifecycle processes needed to deliver a complete environment.

This approach enables organizations to manage heterogeneous infrastructure through a consistent operating model while retaining the underlying technologies, operational boundaries and governance requirements.

This is more than a provisioning improvement. It makes specialized environments repeatable and consumable without requiring every user to reproduce the integration work behind them.

A reusable Offering creates a foundation for making infrastructure capabilities available to additional consumers, teams and potentially organizations. Consistent lifecycle management helps providers maintain operational control. Governance and policy enforcement support the requirements of enterprise consumption. Consumption visibility and attribution establish the operational information needed to understand how resources are being used.

That repeatability is a prerequisite for federation, though not a substitute for the commercial systems needed to operate it.

Industry Adoption Provides an Important Signal

Cisco's selection of Quali technology and its introduction of Stack Automation provide a practical example of this operating model moving into broader industry adoption.

Stack Automation applies the concept of reusable, validated infrastructure and software environments to the delivery of complex technology stacks. It reflects the need to make sophisticated infrastructure capabilities easier to deploy, manage and consume consistently rather than requiring each organization to assemble and integrate the complete environment independently.

Cisco's adoption should not be confused with delivery of the complete federated economic model. Its significance is narrower and more concrete: a major infrastructure provider has brought a reusable-stack automation approach into its portfolio.

The next challenge is to extend such operating foundations across independent provider-consumer relationships, without requiring common ownership of the resources or a single commercial model.

Connecting the Operating Model to the Economic Model

Technical federation and commercial federation are related, but distinct. A control plane can define and govern an Offering, automate its lifecycle and provide operational context for consumption. It does not, by itself, establish multi-party licensing, pricing, invoicing or settlement.

Those obligations require agreements and commercial systems that translate measured consumption into charges, entitlements and payments. The distinction matters: the infrastructure foundation can be real and useful before the complete economic network exists.

The Next Chapter of Infrastructure Economics

Virtualization improved how infrastructure resources were used. Cloud changed how they were purchased. Automation changed how complex environments could be delivered. Federation points toward a further shift: changing how independently owned capabilities can combine to create, and share in, economic value.

In that model, a provider does not need to own the customer relationship to contribute infrastructure. A service provider does not need to own every component to deliver a differentiated product. And a customer does not need to contract separately with every organization whose capabilities help deliver the service.

The prize is not a larger marketplace for spare capacity. It is a broader economic role for infrastructure, software and expertise: reusable contributions to products assembled and consumed across organizational boundaries.

Realizing that prize will take interoperable operations, trusted governance, accurate attribution and commercial arrangements that work for every participant. The market examples in this paper show pieces of the journey; they do not remove the work still required.

But the strategic choice is already present. Organizations can build the foundations for participation, evaluate where their capabilities could create value in someone else's Offering, and decide which relationships they want to shape rather than inherit.

Quali's contribution is to help make complex infrastructure consumable and governable through reusable Offerings. Cisco's adoption provides a tangible signal that this foundation matters. The economic opportunity built on top of it is larger than any one platform or provider.

The decisive change will come when the value of infrastructure is measured not only by what its owner can deploy or sell directly, but by the services and businesses it can help enable across a network.

The next era of infrastructure will not be defined simply by who owns the resources. It will be defined by who can turn those resources into value, who can participate in delivering that value, and how broadly that value can move.