

Top 5 Ways to Accelerate Multicloud Success

How enterprises can improve SaaS experience, connect distributed clouds and AI workloads, secure all traffic, and assure performance from user to application.



Table of Contents

The Five Use Cases	3
Keep SaaS Apps Fast for Every User, Everywhere	4
Connect Apps and Data Across Clouds Without Rebuilding the Network	5
Give Critical Cloud and AI Workloads a Predictable Path	6
Secure the Right Traffic in the Right Place	7
Know Where Experiences Breaks Before Users Feel It	8
What These Five Changes Deliver	9
Building the Foundation for What Comes Next	10

Accelerating Cloud Success Starts with the Network

For most enterprises, getting to the cloud was never the hard part. Making cloud work consistently, securely, and at the pace the business now expects is.

The cloud is not a singular destination. It is SaaS applications, hyperscaler regions, private data centers, colocation hubs, edge locations, AI data centers, GPU-as-a-service platforms, and emerging neocloud providers. Users, applications, and data move across all of them, often in the same business workflow.

That evolution changes the role of the network. Beyond serving as a path to the cloud, the network acts as a combined security enforcement point and operating model. This foundation ultimately determines whether multicloud accelerates your business or introduces new friction.

When the network is not designed for this reality, the symptoms show up quickly:

- Users feel it first: video calls stutter, CRM pages lag, and SaaS performance varies by location.
- Application teams wait: every new cloud, region, or provider connection becomes a one-off networking project.
- AI teams hit friction: inference, training, and data retrieval depend on fast, secure access to distributed compute and data sources.
- Security teams face blunt choices: inspect everything and add latency, or bypass controls and accept risk.
- Operations teams work reactively: they only see the part of the path they own, while user experience depends on networks they do not control.

The enterprises pulling ahead are not simply using more clouds. They are building a cloud networking model that makes distributed environments perform as one operating fabric.

The Five Use Cases

Multicloud success is easier to understand when it is mapped to the outcomes organizations are trying to achieve. These five use cases show how the network can improve experience, accelerate cloud and AI initiatives, strengthen security, and give IT better control across environments.

	Customer use case	The customer reality	Business outcome
1	Keep SaaS apps fast for every user, everywhere	Users in branches, campuses, and remote sites need collaboration and business apps to feel local.	Better SaaS experience, fewer performance tickets, and more consistent productivity.
2	Connect apps and data across clouds without rebuilding the network	Application teams use AWS, Azure, Google Cloud, private data centers, and AI environments at the same time.	Repeatable cloud onboarding, consistent segmentation, and simpler cloud-to-cloud connectivity.
3	Give critical cloud and AI workloads a predictable path	AI, analytics, databases, and cross-cloud services cannot depend on best-effort internet behavior.	Private, software-defined paths for workloads that need consistent latency, bandwidth, and control.
4	Secure the right traffic in the right place	Security teams need to inspect sensitive traffic without forcing every app through the same path.	Selective service insertion, integrated security, and policy that follows applications and data.
5	Know where experiences break before users feel it	Problems often span ISPs, SaaS providers, CDNs, cloud services, and enterprise networks.	Faster root cause, proactive assurance, and fewer blind spots across owned and unowned networks.

1. Keep SaaS Apps Fast for Every User, Everywhere

Customer moment

A branch employee joins a Microsoft Teams call before a customer meeting. Another employee opens Salesforce to pull the latest opportunity record. Both applications are healthy, but the experience depends on the path the traffic takes through local circuits, internet providers, security controls, and SaaS front doors.

When SaaS traffic is backhauled through a data center or sent over a path that is congested, lossy, or high-jitter, the user sees the problem immediately. IT may not. The application dashboard says the service is up. The branch router is up. But the experience is still poor.

What needs to change

SaaS performance has to be managed as a live user experience, not a static routing decision. The network should continuously evaluate available paths, understand which application is using them, and steer traffic based on real-time quality indicators such as latency, loss, and jitter.

That lets collaboration traffic use the path with the lowest jitter, business applications use the path with the best response time, and trusted SaaS traffic avoid unnecessary detours when direct access is the better experience.

Where Cisco fits

Cisco SD-WAN Cloud OnRamp makes SaaS performance measurable, giving network teams the visibility to optimize and explain application experience. It delivers real-time analytics, deeper Microsoft 365 and Webex telemetry, and secure SaaS access through SSE tunnels when needed.

Customer outcome

Users get a more consistent SaaS experience from every location, while IT gains evidence of what happened on the path between the user and the application.

2. Connect Apps and Data Across Clouds Without Rebuilding the Network

Customer moment

A new digital service runs in AWS. Analytics sits in Google Cloud. Identity and collaboration depend on Microsoft Azure and Microsoft 365. Sensitive financial data remains in a private data center. Now an AI use case adds inference in a neocloud or GPU-as-a-service environment.

The business expects one workflow. The network team sees different cloud constructs, routing models, security controls, and operating tools. Every new environment introduces another set of decisions: how to connect, how to segment, how to secure, how to monitor, and how to keep the experience consistent.

What needs to change

Multicloud success requires a repeatable way to connect applications and data across environments while still respecting the native constructs of each provider. The goal is not to hide every cloud difference. It is to give teams a consistent operating model for routing, segmentation, policy, and visibility across them.

This becomes even more important in the AI era. AI applications rarely live in one place. A user-facing app may call an inference service in one environment, retrieve enterprise data from a private data center, and enrich results with data from a public cloud. The network has to make those connections secure, performant, and operationally manageable.

Where Cisco fits

Cisco SD-WAN Cloud OnRamp for Multicloud extends the SD-WAN fabric into major public clouds such as AWS, Microsoft Azure, and Google Cloud. It helps teams apply consistent routing, segmentation, and security policy while using cloud-native constructs where they make sense.

The result is a more universal operating experience across providers, including cloud-to-cloud connectivity and policy consistency inside the cloud.

Customer outcome

Cloud onboarding shifts from a custom integration project to a repeatable connectivity service that supports site-to-cloud, cloud-to-cloud, and AI data access patterns.

3. Give Critical Cloud and AI Workloads a Predictable Path

Customer moment

An AI team needs to move data to a GPU cluster for model training. A manufacturing site needs reliable access to a cloud-hosted control application. A financial analytics workload depends on low-latency access between regions. In each case, best-effort internet routing may work sometimes, but inconsistency becomes a business risk.

Traditional private circuits can deliver predictability, but they often come with long lead times, rigid capacity planning, and economics that do not match the way cloud teams work. Enterprises need private-connectivity behavior with software-defined agility.

What needs to change

Not every workload needs the same path. Some traffic can use the internet. Other traffic needs private, SLA-backed connectivity. And AI workloads may need on-demand access to neocloud providers, AI data centers, or GPU-as-a-service platforms. The network should give teams a way to choose the right path based on performance, security, and cost requirements.

The architecture behind that outcome may include software-defined cloud interconnect, provider backbones, and private transport options. For the customer, the value is predictable access to critical workloads across a distributed landscape.

Where Cisco fits

Cisco SD-WAN Cloud Interconnect integrates with software-defined cloud interconnect providers such as Equinix and Megaport so teams can provision private connectivity through Cisco SD-WAN Manager. Cisco also supports cloud and backbone integrations such as AWS Cloud WAN, Microsoft Azure Virtual WAN, and Google Cloud Network Connectivity Center for high-performance site-to-cloud, site-to-site, and cloud-to-cloud paths.

For AI, Cisco's integration with Megaport AI Exchange extends this model to AI data centers, neocloud providers, and GPU-as-a-service platforms. Customers can reach distributed AI infrastructure on demand without turning every connection into a months-long networking project.

Customer outcome

Critical cloud and AI workloads get more predictable latency, bandwidth, and control, while IT can provision connectivity faster and plan private interconnect capacity more intentionally.

4. Secure the Right Traffic in the Right Place

Customer moment

A security team wants sensitive data flows inspected before they reach a cloud workload. At the same time, the collaboration team wants trusted Microsoft 365 traffic to avoid unnecessary latency. The AI team needs private access to data and compute, but not every packet should traverse the same stack of security services.

This is where many cloud networking projects become uncomfortable. If everything is forced through one inspection point, performance suffers. And when too much traffic goes direct, security teams lose control. The right answer depends on the application, data sensitivity, user context, and destination.

What needs to change

Security need to be selective and policy driven. Some traffic should go through security services. Trusted traffic should go directly to a VPC, SaaS service, private data center, or cloud workload. AI data flows often need private connectivity, encryption, segmentation, and strict policy enforcement end to end.

The use case is not just 'add security.' It is applying the right security inspection and policies to the correct traffic without making the network harder to operate.

Where Cisco fits

Cisco Catalyst SD-WAN supports segmentation and service insertion so specific traffic can be steered through the services it requires. Cisco SD-WAN can also work with cloud-delivered SSE/SASE architectures, on-box security capabilities, and encrypted transport to help enforce policy consistently across distributed environments.

For customers, the value is selective security: inspect what needs inspection, optimize what can go direct, and keep policy consistent as environments change.

Customer outcome

Security is strengthened by risk and business context. Teams can inspect sensitive flows, preserve performance for trusted applications, and maintain consistent policy across cloud, branch, and AI environments.

5. Know Where Experiences Breaks Before Users Feel It

Customer moment

A payroll application slows down on the last day of the month. The SaaS provider is healthy, the branch network is operating normally, and the ISP reports no issues. Yet, employees are still waiting and the helpdesk queue keeps growing.

This is the reality of cloud operations. User experience depends on a chain of networks and services: enterprise WAN, local ISP, internet backbone, CDN, SaaS provider, cloud provider, VPC, API gateway, load balancer, and application tier. Traditional monitoring sees only a slice of that chain.

What needs to change

Operations teams require assurance from the user to the application, including the parts of the path they do not own. They need to know whether the problem is inside the enterprise, across the internet, within a cloud region, inside a SaaS service, or tied to a recent cloud configuration change.

Visibility alone is not enough. The next step is insight: correlating measurements, surfacing anomalies, understanding impact, and recommending remediation before users feel the problem at scale.

Where Cisco fits

As cloud networking becomes more distributed, customers need visibility beyond devices. They need to assure the digital experience end-to-end. Cisco ThousandEyes Assurance provides visibility from users to applications across owned and unowned networks, helping teams understand the full path of every connection. Cloud Insights maps cloud dependencies and traffic flows, while Traffic Insights correlates network telemetry with synthetic measurements to identify where performance issues occur. Cisco AI Assistant and AI-powered insights help teams move faster from symptoms to likely causes and remediation.

Customer outcome

IT moves from reactive troubleshooting to proactive assurance, with a clearer view of where experience breaks and what action to take next.

What These Five Changes Deliver

Each use case delivers value on its own. Together, they create a cloud networking operating model that is easier for IT to run and the business to rely on.

Outcome

01

User experience

02

Cloud agility

03

AI readiness

04

Security control

05

Operational assurance

What changes for the business

SaaS apps feel more consistent across branches, campuses, and remote locations.

New cloud, region, and workload connections become repeatable instead of bespoke.

Teams can reach distributed AI data, compute, and inference environments with lower friction.

Sensitive traffic can receive the right inspection without slowing every application path.

IT can see, diagnose, and act across cloud paths it does not directly own.

Building the Foundation for What Comes Next

Multicloud already stretched legacy network designs. AI will push them further. AI workloads introduce higher bandwidth needs, lower tolerance for delay, more dynamic access patterns, and greater sensitivity around where data moves and how it is secured.

The AI landscape is also distributed by design. Applications, inference services, enterprise data, hyperscaler services, private environments, and neocloud infrastructure may all participate in one workflow. That makes the network a strategic dependency for AI adoption, not an implementation detail.

The five use cases in this ebook are not a one-time checklist. They are an operating model for cloud networking:

- Optimize the experience users actually feel.
- Make cloud and AI connectivity repeatable across environments.
- Give critical workloads the right path based on performance, security, and cost.
- Apply security selectively and consistently.
- Assure the full path from user to application, including networks you do not own.

Cisco SD-WAN Cloud OnRamp, Cisco Catalyst SD-WAN, Cisco secure access capabilities, software-defined cloud interconnect integrations, and Cisco ThousandEyes help enterprises turn these principles into a practical architecture for multicloud success.

Explore how Cisco SD-WAN Cloud OnRamp can help automate, secure, and optimize connectivity across SaaS, multicloud, AI, and distributed enterprise environments.