SOLUTION OVERVIEW

How to Securely Connect Your Clouds and Networks

F5 Distributed Cloud Multi-Cloud Transit delivers high-performance networking and security across public and private clouds.

KEY BENEFITS

Significantly faster deployments
Accelerate cloud migration or adoption of a new cloud provider using a consolidated service that exposes the same API and networking and security capabilities across any cloud provider.

Substantial cost savings
Reduce TCO by 70% with our SaaS-based service, usage-based model, automated lifecycle management, and ability to connect sites using F5's high-performance global backbone.

Extreme productivity gains
Speed up migration to Infrastructure as Code with built-in automation assistance and lifecycle management. Improve collaboration by delivering self-service to DevOps and SecOps teams.

SOLUTIONS FROM CLOUD PROVIDERS AND OTHERS AREN'T HELPING—THEY'RE SIMPLY MAKING THINGS MORE COMPLICATED.

Organizations driven by business transformation have tasked their IT teams to accelerate their migration to one or more public clouds to gain the simplicity and automation promised. As these enterprises migrate their IT environments to the cloud, they are faced with everyday challenges to acquire new skills in cloud networking as well as the operational complexity of managing disparate products.

While the benefits of cloud migration—cost savings, more efficient use of IT resources, greater ability to expand and scale—are known, there is still a need to address the challenges that infrastructure and operations teams face with networking and security in public clouds. These include:

- **Slower deployments:** Network Operations (NetOps) teams find that even a single public cloud's constructs are different from what they are used to on-premises. They need help with orchestrating and managing those constructs (e.g., TGW, security groups, policies to steer traffic). Even if they are working with a single cloud, many begin planning for multi-cloud scenarios to enable different public clouds to work together in delivering applications. Many enterprises face multi-cloud networking challenges due to mergers and acquisitions where the combining companies are in different clouds. They want a consistent policy and operational model across clouds. However, every public cloud's construct is different. All of this significantly slows down deployments and impacts time to revenue for a new service.
- **Operational complexity:** While a public cloud, powered by automation and Infrastructure as Code (IaC), provides agility and speed, NetOps seek enterprise-class Day-2 visibility, control, regulatory compliance, and the diagnostics that they would get on-premises. When dealing with multi-cloud, there is no single-screen observability. This results in significantly higher operational complexity to build and operate across multiple clouds. It could also mean additional security risk due to the challenges of potentially inconsistent enforcement of organizational policies across multiple clouds.
- **Lack of advanced networking and security constructs:** NetOps teams need the same level of advanced networking and security constructs that they are used to on-premises. While public clouds have begun to provide granular controls (such as policy-based VPNs), configuring and operating these policy controls is complex at scale. For example, security groups are configured per VPC, which makes it cumbersome across hundreds of VPCs. To achieve advanced networking/security control, enterprises augment public cloud networking/security with multiple point products, such as virtual routers and firewalls. However, these appliances have different operational models and policy configurations, resulting in inconsistent policies and adding to the operational complexity.

KEY FEATURES

Consolidated networking and security
Leverage consolidated networking, security, and load balancing in any cloud.

Advanced application policies and API security
Configure service policies at the individual API and API method levels, and leverage app-layer WAF, API security, and a network firewall to enable security controls at L3-L7.

Secure global backbone
Rely on F5 Global Network to provide secure and high-performance connectivity across nodes deployed in multiple cloud providers, data centers, and edge locations.

Private connectivity
Implement F5 Distributed Cloud Private Link, a last-mile high-speed physical connection from the customer premises to the F5 global backbone and to public clouds.

Centralized observability and diagnostics
Get centralized visibility and insights into network, security, apps, and users, eliminating the need to gather data from multiple sources, plus real-time log streaming and analytics.

Multi-tenancy and self-services
Use a shared multi-tenant platform that allows developers, DevOps, NetOps and SecOps to collaborate while maintaining separation of duties.

- **Decrease in productivity:** Application architecture is shifting from monolithic to microservices. DevOps teams now need to operate clusters with back-end load balancers, API gateway, and service mesh technologies. DevOps must operate many clusters per team, per microservice, and per app. They cannot wait on NetOps to provision and connect their clusters as it takes too long, with multiple iterations, negating the service velocity benefits of moving to microservices. NetOps teams, meanwhile, must have self-service capabilities where only NetOps can be responsible for configuring organization-wide policies while providing DevOps the ability to configure the app-specific policies on their own.

Cloud to Cloud Transit: A Key Use Case for F5 Distributed Cloud Mesh

F5® Distributed Cloud Mesh has been designed for the multi-cloud transit scenarios predominant in migrations to one or more clouds. Distributed Cloud Mesh provides consistency and controls across multiple public and private clouds, with unified policies and single-pane-of-glass management for traffic, workloads, and more. The result is reduced operational complexity and greater efficiency and cost savings.

The SaaS-based Distributed Cloud Mesh delivers seamless and secure networking into any cloud. It connects multiple clouds securely via site-to-site over the Internet, over your own private backbone, or via the F5 Global Network. F5® Distributed Cloud Console, an accompanying solution, is a single pane of glass for managing distributed apps and infrastructure across multi-cloud, providing early alerts and generating actionable business insights.

Distributed Cloud Mesh delivers consolidated networking, security, and load balancing/application delivery control in any cloud. It also automates the configuration of cloud resources such as AWS Transit Gateway, Microsoft Azure, and Google Cloud Platform networking to reduce the automation burden on NetOps.

The F5 Distributed Cloud Platform offers multi-tenancy, enabling NetOps to deliver self-service to DevOps and reduce the need to manage additional services per app cluster. Optionally, the F5 Global Network can be leveraged to provide secure and high-speed connectivity across nodes deployed in multiple cloud providers, data centers, or enterprise edge locations.

SECURITY TEAMS HAVE THE RESPONSIBILITY OF PROTECTING YOUR APPS, YOUR CUSTOMERS, AND YOUR DATA FROM AN INCREASING RANGE OF THREATS.

For a cloud to cloud transit use case, the key features of Distributed Cloud Mesh that differentiate it and the F5 Distributed Cloud Platform from competitors include:

- Consolidated L3-L7 networking and security services, with unified, cloud-agnostic policies
- Centralized management plane with distributed control and data plane across cloud or edge
- Full multi-tenancy/self-service for collaboration across DevOps, NetOps, and SecOps
- SaaS-based operations with single pane of glass for policy, lifecycle management, and end-to-end observability
- F5's Terraform provider and public APIs that deliver to the automation needs of app teams
- Support for tools such as Opsgenie and Slack for alerting and Splunk and Datadog for SIEM, simplifying life for DevOps and SecOps
- F5 Global Network, a high-capacity global physical backbone, available optionally with F5® Distributed Cloud Private Link, a last-mile physical connection from the global network to the customer's premises

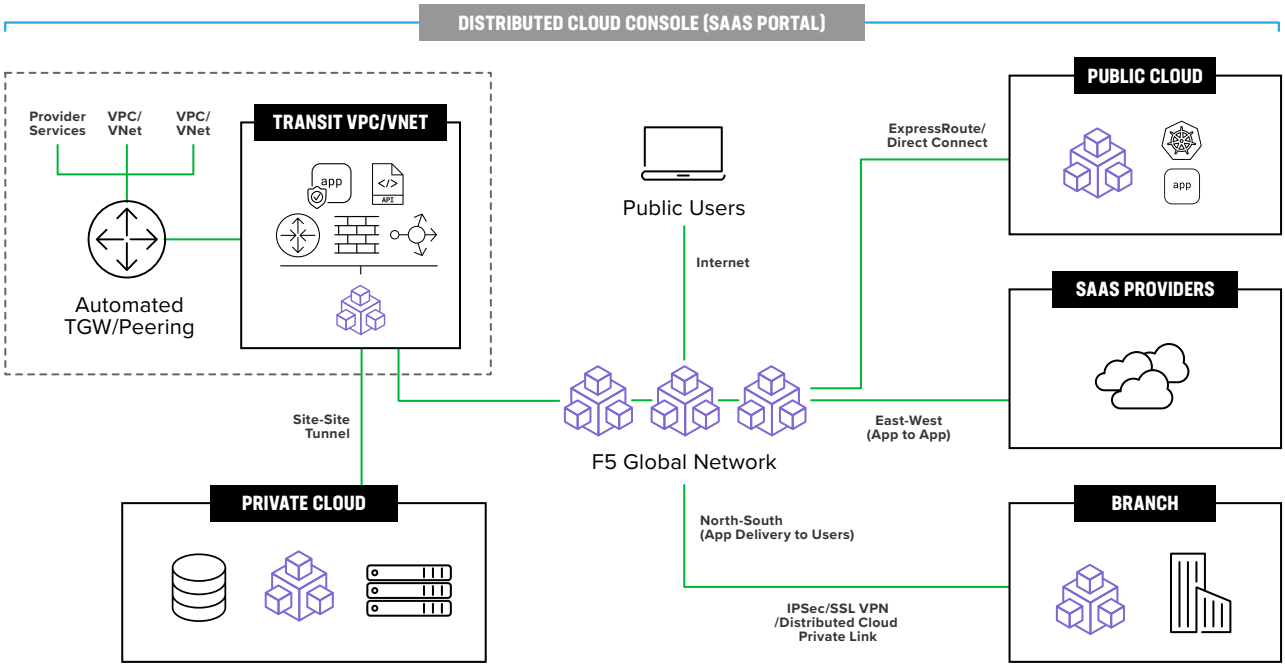

Figure 1: A visual look at how Multi-Cloud Transit works inside Distributed Cloud Mesh

F5'S CONSOLIDATED NETWORKING AND SECURITY SOFTWARE STACK AND F5 GLOBAL NETWORK REDUCE TOTAL COST OF OWNERSHIP AND THE OPERATIONAL BURDEN FOR NETOPS.

A Real-World Example of How Multi-Cloud Transit Provides Value

A large media company acquired a new streaming service and began managing multiple clouds. The media company's network previously had been on-premises and in AWS and the streaming service's network on-premises and in Azure. The combined company sought to converge environments to leverage shared services—including user databases across the two networks to promote cross-selling to both customer bases.

A major hurdle was the networks' overlapping IP address space, which prevented the network consolidation. This overlap prevented reliable delivery of shared services across the two public clouds. Disparate operational models between the two clouds also created significant complication and troubleshooting challenges.

The media company considered deploying multiple point solutions—a router, site-to-site VPN, network firewall, load balancer, and application firewall. But the impact would be:

- Lengthy deployments of new services due to managing and configuring multiple point products in multiple clouds
- Dynamic network address translation (NAT) of the traffic based on the relationship of the client to the fully qualified domain name (FQDN) of the shared service
- Higher mean time to recovery (MTTR) due to operational complexity in different clouds
- Increased costs due to having separate teams manage products at each point; they now have four locations, two private clouds, and two public clouds (AWS and Azure)

Distributed Cloud Mesh addresses the issue of IP address overlap natively by presenting the shared services to each cloud individually without requiring any renumbering or reconfiguration. It also provides the customer with:

- A single intent-driven, repeatable policy across multiple clouds
- One pane of glass for managing multiple layers of the stack and multiple clouds
- The flexibility of multiple options for connectivity—including the F5 Global Network, a customer-provided backbone, site-to-site Internet, and F5 Distributed Cloud Private Link to the customer's premises
- Faster troubleshooting due to a single cloud-agnostic platform serving multiple teams
- Single-team management by having just one operational model, resulting in lower costs

More information

Interested in talking to an F5 Distributed Cloud sales specialist? Contact sales@f5.com today.

Cloud to cloud transit is available as part of F5 Distributed Cloud Mesh. Learn more or test it yourself by signing up for a free trial.