HashiCorp Consul

Unlocking the Cloud Operating Model: Networking

Contents

Overview

For most enterprises, the goals of digital transformation efforts mean delivering new business and customer value more quickly, and at a very large scale. The implication for Enterprise IT then is a shift from cost optimization to speed optimization. The cloud is an inevitable part of this shift as it presents the opportunity to rapidly deploy on-demand services with limitless scale.

To unlock the fastest path to value of the cloud, enterprises must consider how to industrialize the application delivery process across each layer of the cloud: embracing the cloud operating model, and tuning people, process, and tools to it.

In this white paper, we look at the implications of the cloud operating model, and present solutions for IT teams to adopt this model at the networking layer.

Challenges with multi-cloud service networking

In a traditional static network, monolithic applications are connected and secured by a fleet of network infrastructure, such as load balancers and firewalls. Companies focus on managing the traffic flowing in and out of the private datacenter (also known as north-south traffic). The network is protected by coarse-grained network segments. The network operation processes are normally ticket-based and manual. For example, it could easily take days or weeks to update a chain of network infrastructure devices in order to launch a new application.

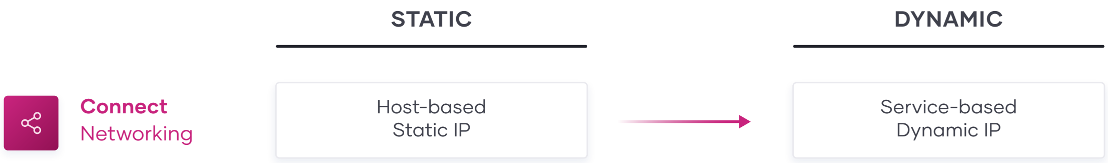

Figure 1: Static to Dynamic infrastructure in the network layer

The trend of application modernization, compounding with cloud adoption has shifted the underlying networking layer from host-based static networking to service-based dynamic networking.

In today's multi-platform, multi-cloud network:

- Microservices are running on ephemeral containers with dynamic IPs.
- Environments with duplicate IP Address spaces have grown in frequency.
- The explosion of service-to-service communication (also known as east-west traffic) results in complex traffic patterns and potential vulnerabilities from lateral movement attacks.
- The public clouds blur the boundary of network perimeters, presenting a new set of security risks.

As the network becomes larger and more dynamic, it puts pressure on networking and security teams to keep pace with more changes to existing network infrastructure. These teams often cannot keep up with the increasing demand, and their manual approach to network operations undermines the agility of application teams.

HashiCorp Consul: Multi-cloud service networking simplified

Designed to solve the rising networking challenges for dynamic infrastructure and enable progressive delivery practices, Consul provides a multi-cloud software networking platform to discover, connect, and secure services.

Consul takes a different approach to networking by:

- Abstracting the operation from IP address to service level based on identity.
- Providing a central shared service workflow to locate the dynamic location of services.
- Enabling service discovery to allow services to register, discover, and connect with each other directly.
- Driving automation to eliminate the operational burden of manually updating network infrastructure.
- Provides a service mesh solution to simplify service networking by shifting naming abstraction, routing, authorization and other networking functionality from a central networking infrastructure layer to the endpoints.
- Being platform-agnostic, which allows IT teams to maintain a single networking workflow to networking across different runtime platforms, public cloud or private clouds.

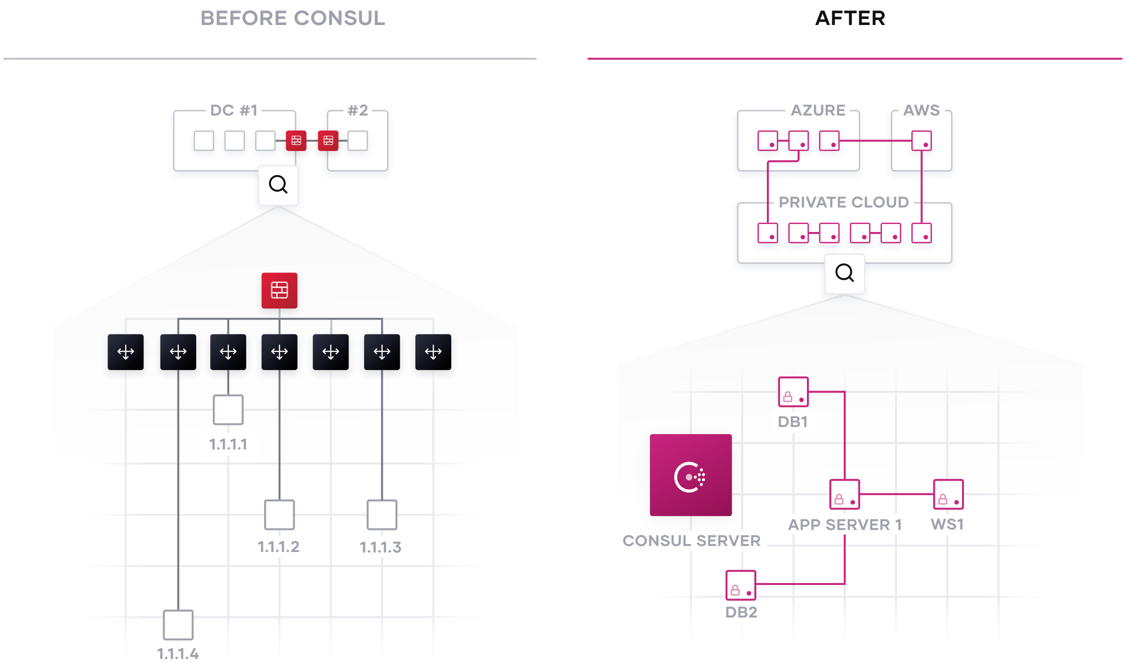

Figure 2: From traditional networking to modern services networking with Consul

Consul's crawl, walk, run journey

Adoption of HashiCorp Consul typically follows a three-stage pattern for most organizations. First, organizations can establish a foundation of services networking with Consul using a common service registry and by providing service discovery across their multi-cloud environments. Second, organizations can standardize on delivering a shared services approach across various operational teams to automate networking processes that were formerly manual and ticket based. Third, they can start to migrate and innovate their service networking to a service mesh architecture to improve agility, security, performance and reduce cost.

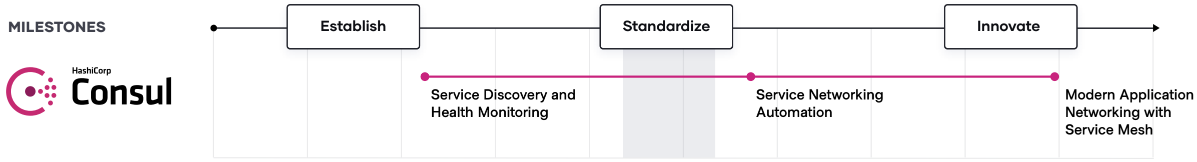
MILESTONES
HashiCorp Consul
Establish
Standardize
Innovate
Service Discovery and Health Monitoring
Service Networking Automation
Modern Application Networking with Service Mesh

Figure 3: Evolution of technology and process in the network layer

Crawl: Central Service Registry, Discovery, and Health Checks

For organizations at the beginning of their cloud journey, finding a way to track and manage the explosion of services across multiple subnets, data centers, and cloud regions is their first hurdle.

The traditional approach relies on load balancers to provide a naming abstraction to represent a service with static virtual IPs. The process to track the network location of services often takes the form of spreadsheets, load balancer dashboards, or configuration files, all of which are disjointed, manual processes that are not ideal.

The starting point for networking in the new cloud operating model is the deployment of a common service registry, which provides a real-time “directory” of what services are running, where they are, and their current health status.

A light-weight Consul agent is running along with a service instance, allowing this service instance to programmatically register with the registry as well as discover and connect to other registered services, automatically through DNS or API interfaces. Additional tooling allows Consul to register and monitor appliances or devices that cannot run an agent.

The integrated health check will monitor a service’s health status so the IT team can triage the availability of each instance and Consul can help prevent routing traffic to unhealthy service instances. Consul will automatically remove services that are unhealthy, re-adding them when service health is restored.

Walk: Service Networking Automation

The next step is to unlock agility with existing networking architectures through a standard process using a shared service approach and network automation . Instead of a manual, ticket-based process to reconfigure the traditional network middleware every time there is a change in service network location or configurations; Consul provides a publisher/subscriber (PubSub) like service to automate these operations by updating configuration changes of the network devices. Terraform can also be used to enable rapid day zero operations of the resources used when provisioning new infrastructure.

The newly added service instances will automatically “publish” their location information with the service registry. The network infrastructure can subscribe to service changes from the service registry (by using tools like Consul Template or native integration). This enables a publish/subscribe style of automation that can handle highly dynamic infrastructure and scale much higher.

For example, F5 BIG-IP devices are integrated with Consul. Instead of being configured to balance traffic to a static pool of backend services, BIG-IP can be configured to reroute to a dynamic set of backends that are provided by Consul. The moment an additional backend service is provisioned and registered, BIG-IP will automatically update the set of backend services and begin routing traffic to it instantly, without manual intervention. The same approach applies to firewalls and other critical network infrastructure as well. This publish/subscribe model significantly accelerates the productivity of the network and IT operation teams by automating provisioning based on changes in the infrastructure environment.

It also decouples the workflow between teams. As organizations start to scale up the number of services they are running, it becomes more difficult for a single operator to manage the many teams who are developing and deploying them. Consul namespaces allow global operators to create isolated management environments in a shared cluster environment and enables global admins to delegate per-tenant admins the capability to manage policies for their own services. With namespaces, global operators can independently deploy applications and publish to Consul, while network operations teams can subscribe to handle downstream automation.

Run: Service mesh

As organizations continue to scale with microservices-based or cloud-native applications, the underlying infrastructure becomes larger and more dynamic. Modular services need to communicate with each other to compose business logic and functionality, leading to an explosion of east-west traffic.

Existing networking approaches with network appliances cannot effectively handle east-west traffic in dynamic settings. They cause a proliferation of expensive network infrastructure, introduce single points of failure all over the system and add significant operational overhead to IT teams.

Furthermore, application based networking has driven significantly more complex requirements on traditional network teams than existed before. With significant amounts of workloads becoming ephemeral, as well as being highly distributed as microservices, the ability to successfully route and lifecycle application traffic across the network without downtime becomes critical to organizations.

A distributed service mesh pushes routing, authorization and other networking functionalities to the endpoints in the network, rather than imposing them through a central point in the infrastructure. This makes the network topology simpler and easier to manage, it removes the need for expensive central infrastructure within east-west traffic paths, and it makes service-to-service communication much more reliable and scalable because of the network's decentralized nature. Additionally, it removes the dependency for development teams to incorporate routing and authorization rules directly in application code.

Consul provides an API driven control plane, which integrates with sidecar proxies alongside each service instance (such as Envoy, HAProxy, and Nginx) that provide the distributed data plane.

The service mesh approach allows critical functionality like naming, segmentation and authorization, traffic management and observability to be configured through policies in the central registry and to be enforced by proxies at the endpoint where an individual service is running.

Consul enables a zero-trust network model by securing service-to-service communication with automatic mutual TLS encryption and identity-based authorization. Network operation and security teams can define the security policies through intentions with logical services rather than IP addresses. For example, allowing web services to communicate with databases, instead of IP1 to IP2. Proxies will enforce security consistently regardless of how services scale up and down or migrate to other platforms.

The security benefits of a service mesh based approach are several fold. For most organizations, traffic within a network zone (such as production or PCI) is relatively flat. This means a compromise of a single service would allow an attacker to move laterally to other systems in the same zone. Consul enables a much more fine grained authorization of access to avoid this.

Consul can also be integrated with HashiCorp Vault and AWS PCA for centralized PKI and certificate management.

To address the application networking concerns, layer 7 routing and traffic policy management is provided by Consul and enforced by routing traffic based many possible conditions (HTTP header, path based routing, etc...) to support use cases such as canary, A/B testing and gradual application deployment rollouts, and application lifecycle efforts. These practices have become the foundation of progressive delivery for applications in the enterprise and can only be effectively achieved leveraging a service mesh. For cross-cloud communications, Consul's Mesh Gateway feature routes traffic to the correct endpoint on a private network without requiring expensive IPSec or MPLS connectivity.

Adopting and scaling service mesh for modern application networking

For many organizations, adopting service mesh is not a simple or binary solution. Isolated environments, cross-cloud networking, and a shifting application landscape create challenges for operators who are trying to adapt to complex networking requirements. Often they may lead to the perception that operators should only deploy service mesh in greenfield environments to support new, modular applications primarily running on Kubernetes rather than tying into an organization's broader IT workflow. As a result, legacy applications located on-premises are siloed, using traditional infrastructure and point solutions, as well as requiring manual updates of VPNs, firewalls, and load balancers as workloads scale.

Organizations require services networking to become more application centric, and allow organizations to bring old and new applications together in a heterogenous way, rather than requiring an all or nothing approach involving complex cloud migration and application refactoring. Organizations want their service mesh to be consistent across all of their workload platforms, runtimes, and environments.

Consul's approach provides operators an on-ramp to adopting and scaling service mesh, not just for cloud-native workloads, but for traditional and adjacent workloads as well. Expanded gateway types in Consul lower the barrier to entry for the adoption of service mesh across a broad set of applications and cloud environments. This results in higher operational efficiencies and the establishment of a consistent network and application platform that functions across multiple runtimes. In practice, organizations can adopt a service mesh across their entire environment, without having to re-platform existing applications and workloads, unlike other service mesh solutions.

The first step when onboarding modular applications into a service mesh is to establish a north-south connection to the upstream service. With Consul ingress gateways, external inbound requests can be routed directly to the service mesh without the use of more standalone API gateways. Connections can be made using predefined layer 7 traffic rules and is automatically secured using mTLS.

Due to the central role of proxies in service meshes, users have typically needed to deploy them alongside all machines to take full advantage of the mesh. With terminating gateways, organizations can enable the mesh services to communicate with external services outside the mesh, where sidecar proxies can't be deployed. The terminating gateway acts as an egress point for the service mesh which terminates mTLS connections, and allows the enforcement of Consul's service based intentions as well as communication with external services in the desired (often on-premises) environment.

The WAN Federation enhancements to Mesh Gateway provides a simplified and standardized way to connect environments across cloud datacenters and regions where networking has become even more complex. These environments may be highly secured, or leveraging overlapping IP addresses.

WAN Federation allows traffic to be proxied securely through the mesh gateway with only the gateways exposed to the WAN. Traffic is automatically encrypted leveraging mTLS and doesn't need to be decrypted. Consul leverages the SNI header of the TLS protocol to determine the destination of the traffic and then forwards it to the destination, establishing a point-to-point connection.

Summary

In a multi-cloud datacenter, networking services should be provided centrally, whereby IT teams start with providing service registry and service discovery capabilities. This common registry establishes a foundation by providing a "map" of what services are running, where they are, and their current health status.

As the environment grows, the registry can be queried programmatically to drive network automation of API gateways, load balancers, firewalls, and other critical network infrastructure components used for north-south traffic. Network infrastructure components used for east-west traffic can be automated to reduce the need for manual ticket based systems.

Network topologies and complex application centric traffic can be easier to manage by using Consul's service mesh, where proxies run on the edge with greater scalability and security. Consul's expanded gateways provide cloud and on-premise interoperability,lowering the barrier to entry for the adoption of service mesh across a broad set of applications and environments. This results in higher operational efficiencies and faster time to market.

Consul's crawl, walk, run journey is a simple path that organizations can follow towards a more simplified and scalable networking approach to support distributed applications with multi-cloud and multi-datacenter topologies.

HashiCorp