

Observability in the Cloud Operating Model

Fast-track multi-cloud
success with Splunk and HashiCorp

Whitepaper

Contents

Executive summary 03

Implications of multi-cloud requirements 05

Multi-cloud provisioning and monitoring as code 10

Multi-cloud security and compliance 12

Multi-cloud networking and system-wide observability 15

Conclusion 19

Executive summary

Digital market shifts are propelling organizations in every industry to adopt cloud. As organizations strive to improve customer experience, increase innovation velocity and rapidly scale digital applications, the cloud becomes a key initiative. As explained in the [Cloud Operating Model](#), moving to cloud provides opportunities to rapidly deploy dynamic services, achieve practically unlimited scale, improve developer agility, and bolster application resilience.

Despite the promise of transformational results, enterprises struggle to gain the full value of multi-cloud environments. In fact, in some cases, siloed approaches in the hybrid, multi-cloud world may result in more complexity, slower innovation and higher costs than it did before cloud migration. The challenge begins with a lack of deep visibility into applications, cloud infrastructure and services across multi-cloud and hybrid environments before, during and after cloud migrations.

To unlock the fastest path to value in the cloud, organizations must have a strategy to unify and automate Observability across the cloud stack. Enterprises must:

1. Achieve unified, real-time Observability across the entire cloud stack and cloud operations
2. Adopt automation in provisioning Observability alongside cloud infrastructure and services
3. Get visibility into identity and resource access control; establish audit trails to meet compliance requirements
4. Leverage higher-level abstractions such as Kubernetes that can provide application portability across multiple hybrid clouds; consider service mesh that can decouple applications from cloud-provider specific network services and act as a system-wide source for Observability data

Observability

Observability is a measurement of the quality of software, services, platforms and products that allows us to understand how systems are behaving. Observability makes investigating and diagnosing problems easier; the more observable a system, the quicker we can understand why it's acting up and fix it.

Fundamentally, **Observability is about everything 'data'**. DevOps teams determine what data to generate, collect, aggregate, summarize and analyze to gain meaningful and actionable insights. Observability solutions use metrics, traces and logs as data types to understand and debug distributed systems. However, the mere availability of data doesn't deliver an enterprise-grade Observability solution.

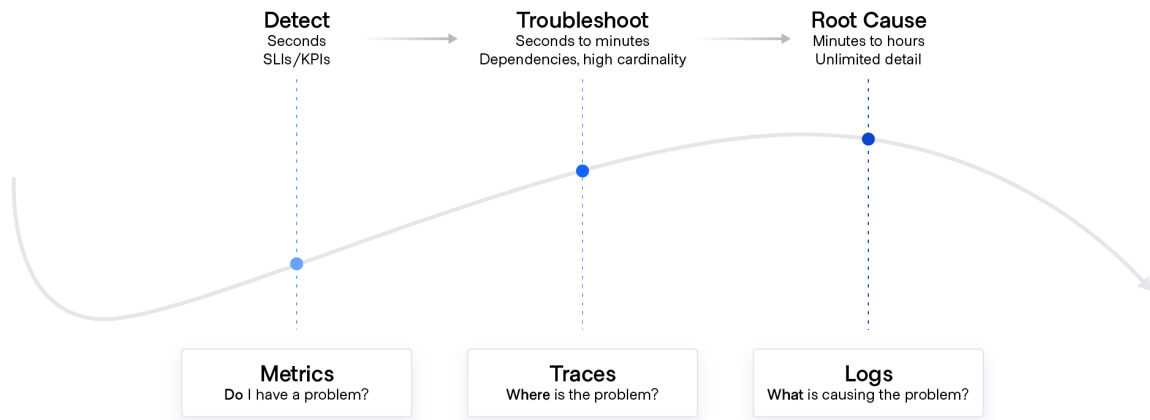


Fig: Three pillars of observability

In the dynamic, multi-cloud world, siloed and point monitoring tools transition to unified and ubiquitous Observability across the entire stack — cloud infrastructure, runtime orchestration platforms such as Kubernetes, cloud managed services such as Amazon RDS or Google Cloud Pub/Sub, etc., and business applications. This unification of Observability allows teams to understand the interdependencies of cloud services and components.

Observability becomes an integral part of the DevOps toolchain that enables developers and operations teams to quickly detect, triage, troubleshoot and resolve performance issues – before such issues can impact end-users.

Implications of multi-cloud requirements:

The [HashiCorp Cloud Operating Model](#) is a blueprint for how organizations migrate to, and address the challenges of, a multi-cloud reality. While there may be a variety of technologies that offer the capabilities of deploying applications in the cloud, organizations should understand and plan for how to address the five key workflows: provisioning, security, networking and observability.

Provisioning:

The foundation for adopting the cloud is provisioning – infrastructure, cloud services, and observability. In the static world, monitoring was treated as an afterthought. In dynamic cloud environments, the infrastructure layer transitions from running dedicated servers at a limited scale to a dynamic environment where organizations can easily adjust to increased demand by spinning up thousands of servers and scaling them down when not in use. As architectures and services become more distributed, the sheer volume of compute nodes increases significantly. Automating provisioning across every layer of the stack becomes a critical capability.

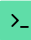



Security:

The security layer transitions from a fundamentally “high-trust” world enforced by a strong perimeter and firewall to a “low-trust” or “zero-trust” environment with no clear or static perimeter. As a result, the foundational assumption for security shifts from being IP-based to using identity-based access to resources. This shift is highly disruptive to traditional security models. Enterprises must have visibility and control over access to resources.

Networking:

The networking layer transitions from being heavily dependent on the physical location and IP address of services and applications to using a dynamic registry of services for discovery, segmentation and composition. An enterprise IT team doesn’t have the same control over the network or the physical locations of compute resources, and must think about service-based connectivity in order to achieve system-wide Observability by leveraging service mesh.

Each cloud provider has their own solution to these challenges. For enterprise IT teams, these shifts in approach are compounded by the realities of running on hybrid and multi-cloud infrastructures and the varying tools each technology provides.

		Static	Dynamic			
		DEDICATED	PRIVATE CLOUD	AWS	AZURE	GCP
	Run Deployment	vSphere	→ vSphere	EKS / ECS Lambda	AKS / ACS Azure Functions	GKE Cloud Functions
	Connect Networking	Hardware	→ Various Hardware	CloudMap AppMesh	Proprietary	Google Istio
	Secure Security	IP: Hardware	→ Identity: AD/LDAP	Identity: AWS IAM	Identity: Azure AD	Identity: GCP IAM
	Provision Operations	vCenter	→ Terraform	CloudFormation	Resource Manager	Cloud Deployment Manager

Observability

The key requirements of a modern observability platform – one fitted for the cloud – are outlined below:

Real-time Problem Detection:

Because modern software-defined infrastructure spins up and down within minutes, every second counts in delivering a flawless end-user experience. The ability to instantly discover and accurately alert on anomalies is now imperative.

Auto-Discovery and Integrations:

The observability platform should be able to automatically discover applications or cloud services and provide pre-built curated visualizations. Out-of-the-box monitoring of every layer of infrastructure allows DevOps teams to correlate metrics across systems to understand the interdependencies between services.

Accurate Alerting:

Even traditional and simplistic tools can fire off an alert when a metric crosses a static threshold. But such static alerting approaches are not adequate in constantly changing, ephemeral cloud environments. An advanced observability approach provides sophisticated alerting capabilities including dynamic baselines, automatic outlier detection and sudden changes leveraging AI and ML.

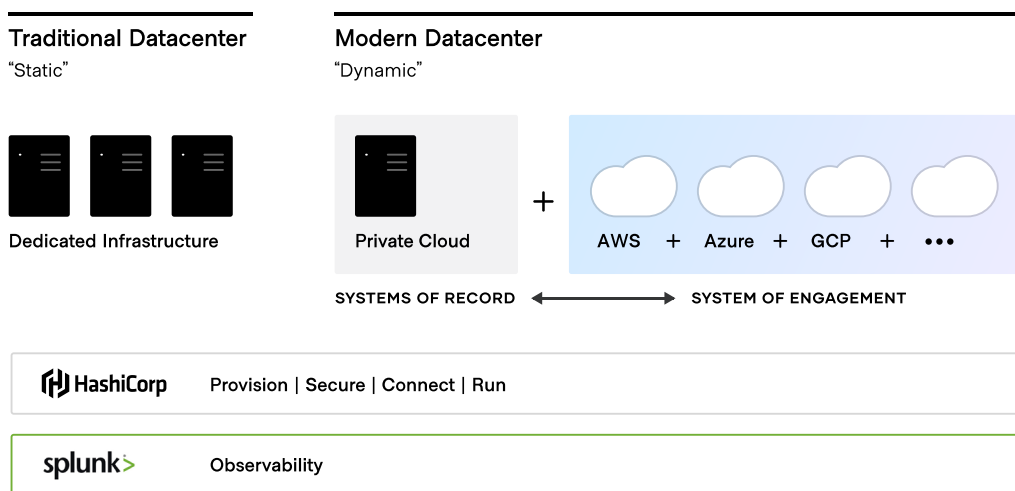
High Cardinality Analytics:

Ephemeral cloud infrastructure and the distributed nature of cloud-native applications exponentially increase the cardinality of performance metrics. Tagging and labels are the basic tenets in cloud architectures to aggregate, segment, slice-and-dice data and get meaningful insights. To act on actionable insights, high-cardinality analytics in real-time is a key requirement of a modern observability platform.

Observability-as-a-Service:

The DevOps paradigm of “you build it, you own it” boosts agility in part by decentralizing operational responsibility to individual teams. More people across the organization now need access to observability, and this decentralization can easily lead to fragmented tools and data. Fragmentation can lead to higher costs and, even worse, highly inefficient operations. Modern observability platforms provide centralized management so teams and users have access controls, and can gain transparency and control over consumption- allowing for better collaboration.

As explained in the HashiCorp Cloud Operating Model, The essential implication of the transition to multi-cloud is the shift from “static” services to “dynamic” cloud services.



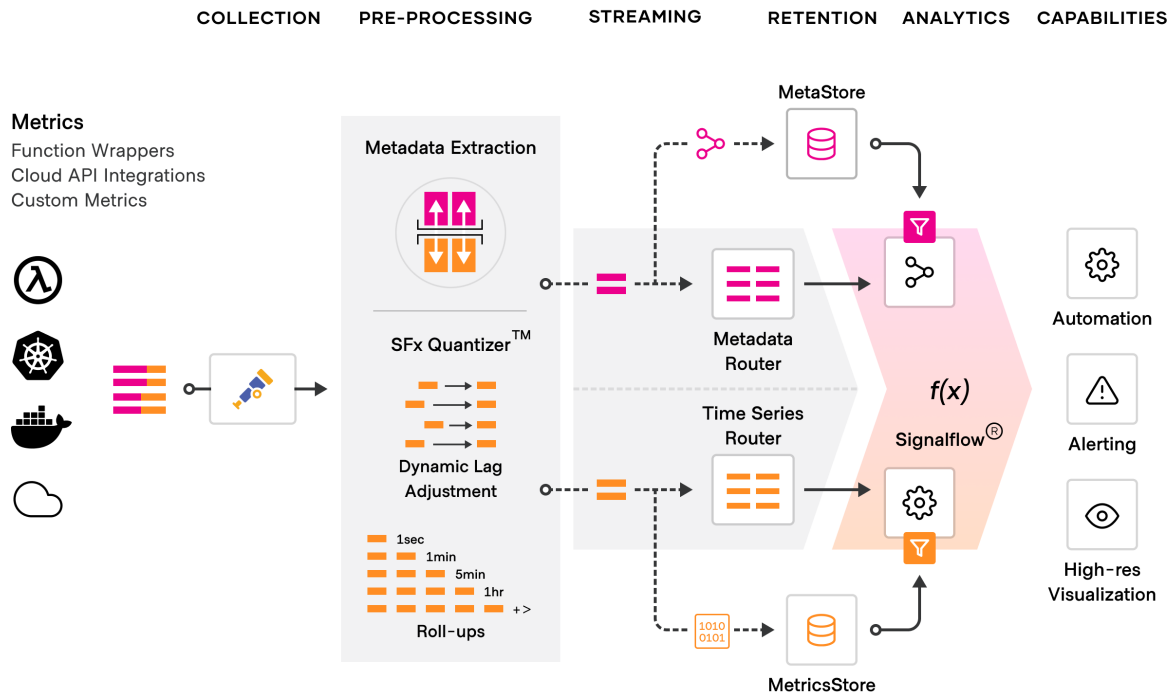
In this white paper, we look at the implications of multi-cloud requirements and recipes to achieve the full potential of multi-cloud with Splunk and HashiCorp, with a focus on Observability across multiple clouds.

Splunk and HashiCorp: Enabling the Cloud Operating Model

Organizations across every industry and around the world use Splunk to accelerate their digital business initiatives and cloud migration, remove organizational silos through better collaboration between development and operations teams, increase developer productivity, deliver better customer experiences and reduce time to resolution using Splunk observability solutions:

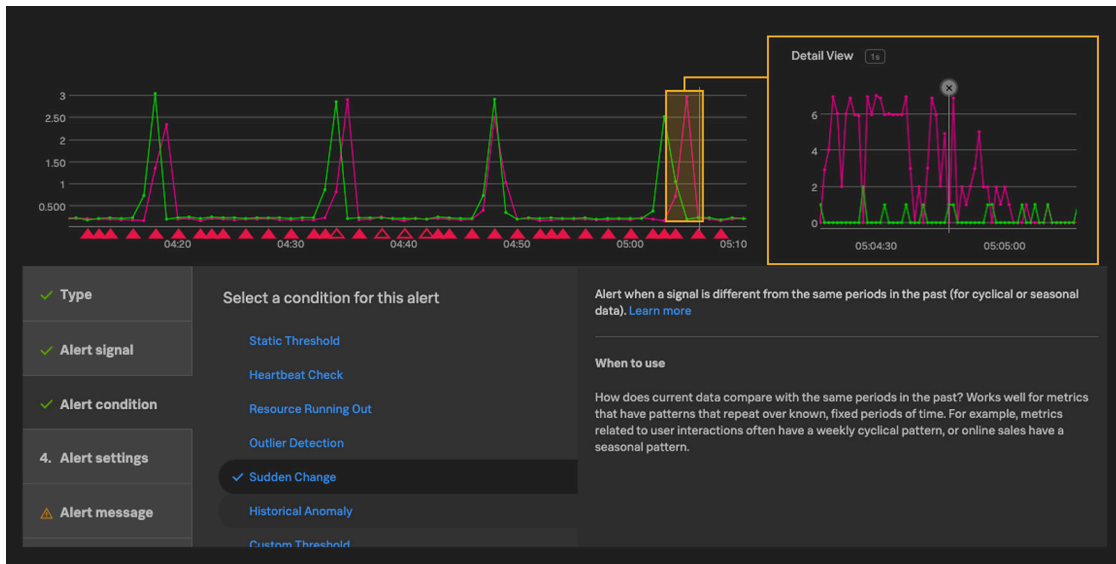
Splunk Infrastructure Monitoring: The only real-time metrics-based monitoring platform for cloud infrastructure and services. It seamlessly integrates with cloud providers and automates observability with hundreds of pre-built integrations into most popular technologies.

Driven by the patented streaming architecture, its approach to ingest, store, and retrieve data is fundamentally different from traditional batch and query-based solutions. As metric data streams into Splunk, metadata is separated from metric value data to scale these datastores separately and provide high-cardinality analytics. The streaming architecture allows users to get insights and take action in real-time – dashboards refresh, alerts fire, and automation tasks trigger all within seconds – whereas other solutions take much longer.



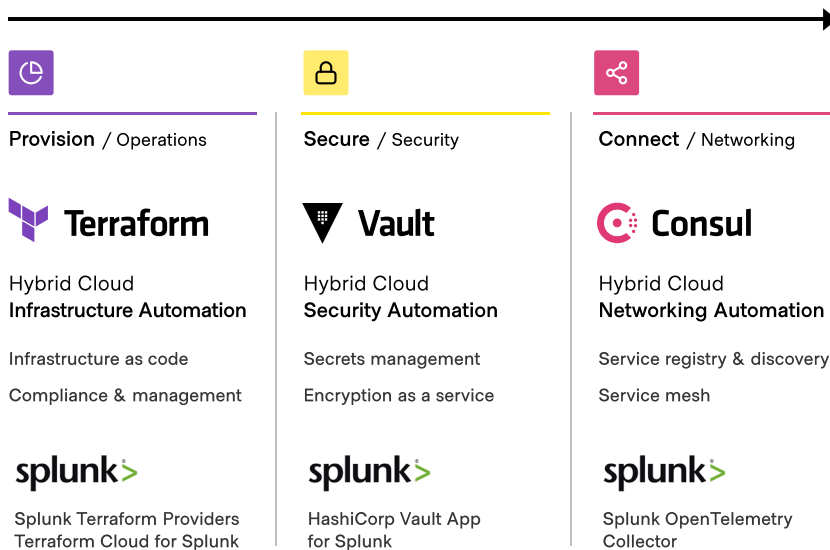
Splunk Infrastructure Monitoring provides point-and-click machine learning capabilities to create accurate alerts and avoid alert storms. Sophisticated algorithms such as Historical Anomaly, Sudden Change, and Resource Running Out go beyond alerting on static thresholds and alert when there is a true anomaly. DevOps teams can validate and fine-tune the accuracy of alert conditions by backtesting against historical data.

In addition, there are more than 20 built-in statistical functions to ensure alert accuracy by not only looking at the raw metric data but also evaluating trends and patterns.



Splunk Enterprise / Splunk Cloud: The most comprehensive, flexible and scalable platform to easily investigate, monitor, analyze and act on structured or unstructured data.

To enable organizations with Observability in the Cloud Operating Model, Splunk is partnering with HashiCorp, the leader in cloud infrastructure automation software. HashiCorp tools are designed to help organizations provision, secure, connect and run any infrastructure for their applications and help them along their digital transformation journey.

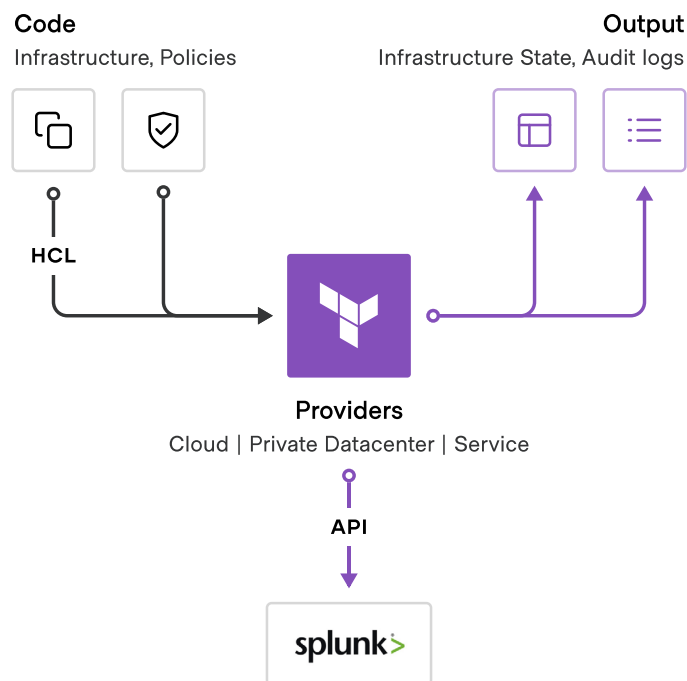


Organizations have adopted the following best practices to achieve success in the cloud with HashiCorp and Splunk products:

Multi-cloud provisioning and monitoring as code

Provisioning is the foundational layer of enabling application deployments in the cloud. Organizations need to have a consistent way to manage multiple services across multiple regions and clouds. Leveraging infrastructure as code enables operators to create and manage infrastructure on-demand, improving efficiency and reducing the time to deploy. Achieving monitoring as code should be a critical consideration so that Observability isn't treated as an afterthought and the best practices are built into the code and delivered via the CI/CD process each and every time an application environment is provisioned.

[HashiCorp Terraform](#) is the world's most widely used cloud provisioning product and can be used to provision infrastructure for any application using an array of providers for any target platform. Enterprises can either run it in the cloud with [Terraform Cloud](#) or manage it on their own with [Terraform Enterprise](#). Since monitoring – dashboards, alerts and more – is a part of your infrastructure, it's helpful to manage them in a similar way. By using monitoring as code, enterprises get automation, visibility and shareable assets that can be used across the organization to create consistent workflows for provisioning observability.

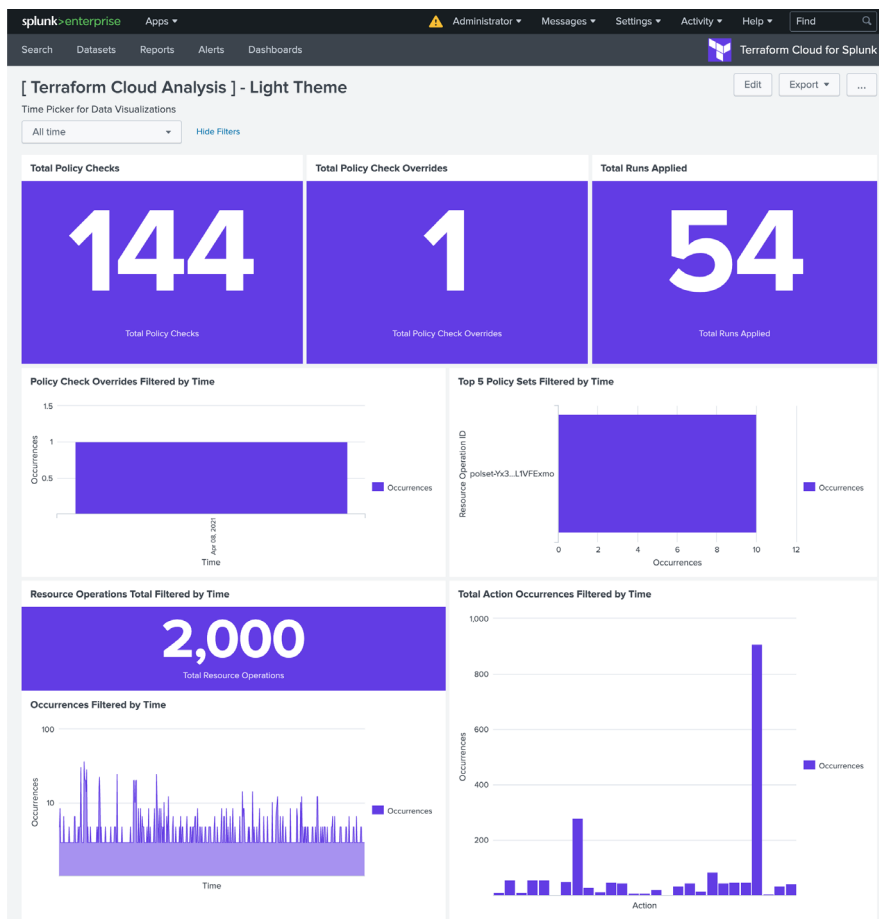


Terraform providers are available for [Splunk Enterprise / Splunk Cloud](#) and [Splunk Infrastructure Monitoring](#) from the official [Terraform Registry](#).

Establishing this consistent approach ensures that Observability is top of mind at the base layer of your application deployments. Operators and developers aren't bound by ticket-based systems and security teams can trust that infrastructure is consistently provisioned within compliance of organizational policies.

Operationalizing Terraform Cloud

HashiCorp Terraform Cloud customers can integrate with Splunk using the official Terraform [Cloud for Splunk](#) app to understand Terraform Cloud operations. Audit logs from Terraform Cloud are pulled into Splunk, giving immediate visibility into key platform events within the predefined dashboards. Customers can identify the most active policies, significant changes in resource operations, or filter actions by specific users within your organization. The app can be used with Splunk Cloud and Splunk Enterprise.



Multi-cloud security and compliance

Historically, datacenters were static infrastructure and had dedicated servers and IP addresses with a clear network perimeter. Security took on a “castle and moat” approach where you could secure everything by controlling the entry and exit points and establish a secure perimeter.

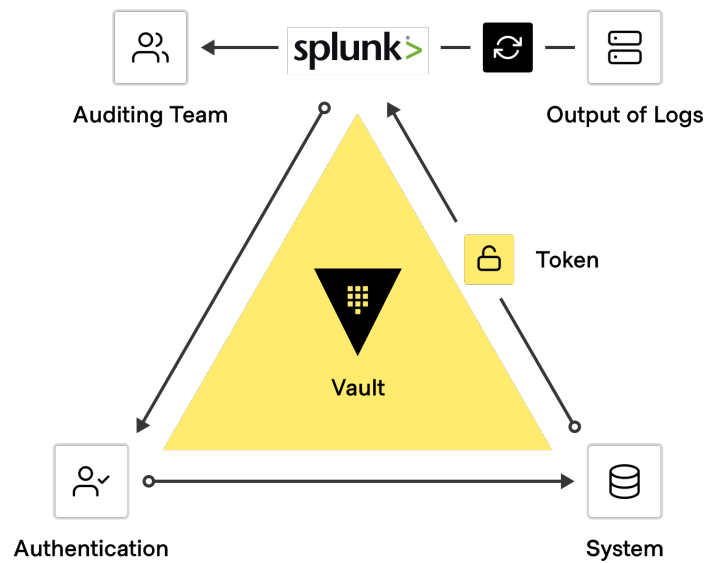
Private networks were inside the “castle” and assumed high trust and integrity. As companies move to the cloud, the measures they took to secure their private datacenters start to disappear. IP-based perimeters and access are replaced by ephemeral IP-addresses and a constantly changing workforce with the need to access shared resources. Managing access and IPs at scale becomes brittle and complex.

Securing infrastructure, data, and access becomes increasingly difficult across clouds and on-premises datacenters, requiring increased overhead and expertise. The shift to cloud and multi-cloud requires a “[zero-trust](#)” approach to security. Zero-trust means that nothing is trusted and everything is authenticated and authorized.

With [HashiCorp Vault](#), users can manage secrets and protect sensitive data. Vault enables teams to securely store and tightly control access to tokens, passwords, certificates and encryption keys for protecting machines and applications. This provides a comprehensive secrets management solution and, when coupled with [HashiCorp Consul](#) and [HashiCorp Boundary](#), a comprehensive zero trust approach to security. Enterprises have a number of ways to consume Vault: self-managed (with [Vault Enterprise](#)), and as a managed service on HashiCorp Cloud Platform ([HCP Vault](#)).

Having granular visibility into access patterns to credentials and secrets enable DevSecOps teams to answer the following questions:

- What happened?
- When did it happen?
- Who initiated it?
- On what resource did it happen?
- Where was it observed?
- From where was it initiated?



Correlating Vault audit logs and other security events enables a radar-like system. Without one, enterprise IT is flying blind. In the static world, enterprises caught and logged isolated anomalous behavior. But, in today's dynamic multi-cloud world, most serious threats are distributed across multiple systems.

Without sophisticated security event analysis solutions, attacks are allowed to germinate and grow into catastrophic incidents. Splunk Enterprise Security (ES) is the nerve center of the security ecosystem, giving DevSecOps teams the insight to quickly detect and respond to internal and external attacks, simplify threat management and minimize risk. ES helps teams gain organization-wide visibility and security intelligence for continuous monitoring, incident response, SOC operations and provides executives a window into business risk.

Splunk seamlessly ingests audit logs from Vault through file-based logs in JSON format ingested via Splunk Universal Forwarder.

Operationalizing Vault Enterprise

Vault Enterprise customers can use the Splunk app to get out-of-the-box operational and security visibility. This app provides pre-built dashboards and reports that span various monitoring use cases. Refer to the detailed, [step-by-step guide](#) that includes recommendations on the most important metrics to monitor, why they are important to monitor, thresholds to be watched out for, and more insights on telemetry and audit for Vault.

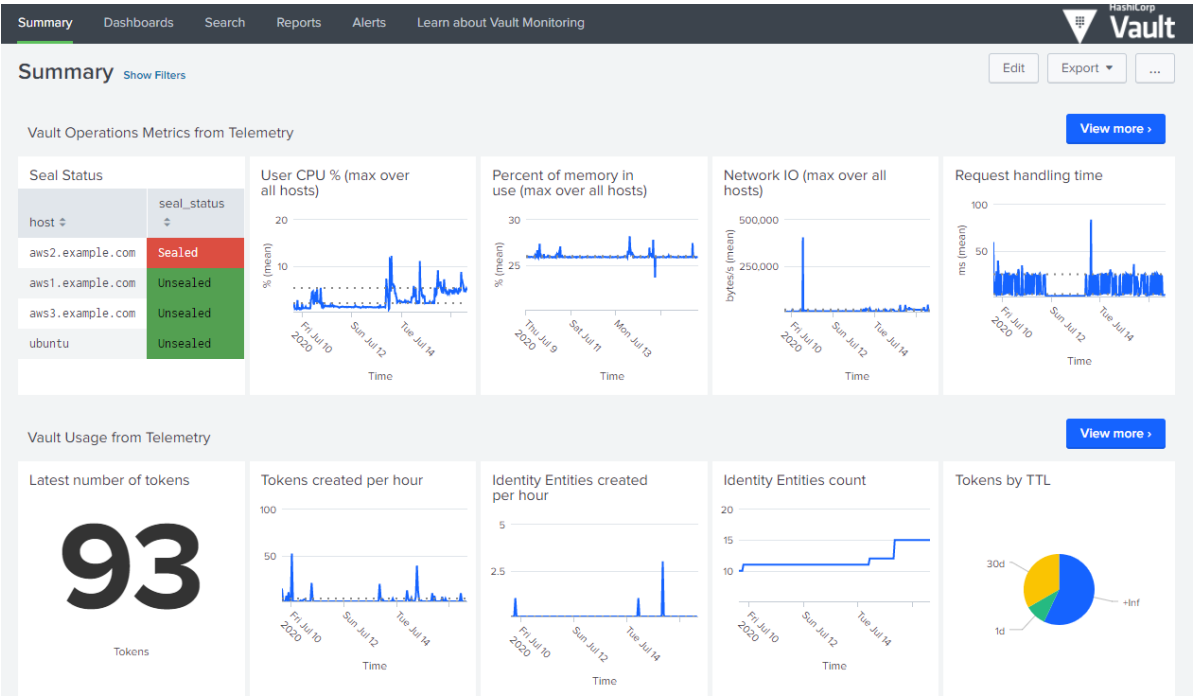


Fig: Vault Cluster Summary Dashboard

Multi-cloud networking and system-wide observability

Microservices deployed in the cloud spread the application logic across multiple services where every service adds a network dependency. The successful execution of the application logic is now embedded in the network data flow between these services.


The distributed nature of microservices across multiple availability zones, regions or cloud providers makes monitoring and debugging workflows challenging:

- How do these services discover and communicate with each other?
- Like the butterfly effect in chaos theory, a minor change to an individual service could have a catastrophic effect on the performance of other services. In those cases, where do teams begin the troubleshooting efforts and how to determine the root cause?
- Every service may be calling multiple other services and may sit on the execution path of requests initiated by other services. What is the service dependency graph? Who is calling whom?
- How well is the overall application performing? How do teams make the overall application more resilient during failures encountered by individual services?


These questions are easy to answer in a static, monolithic world, but cloud-deployed microservices architectures can bring an SRE team to a halt without a new Observability approach.

[HashiCorp Consul](#) enables such an approach. Consul is a service networking solution to connect and secure services and provide Observability data across any runtime platform and public or private cloud: service mesh. Fundamentally, a service mesh is a policy-driven proxy layer that channels all communication between microservices. Service meshes such as Consul incorporate a sidecar proxy with each instance of a microservice application. Each application communicates only with its local sidecar proxy and, in turn, the proxies communicate amongst themselves to form a mesh of services. Enterprises have a number of ways to consume Consul: self-managed (with [Consul Enterprise](#)), as a managed service on Azure ([HCS](#)), and as a managed service on HashiCorp Cloud Platform ([HCP Consul](#)).

Before Consul & Splunk



After Consul & Splunk



Consul Enterprise helps teams solve the challenge of discovering and automating the connections between services. Consul's service discovery capabilities keep an active log of the services running in multiple environments and its service mesh capabilities leverage mTLS for maintaining secure connections.

Splunk Infrastructure Monitoring can seamlessly ingest microservices performance data from Consul Enterprise to give you a complete view of the system's performance from a single platform.

Installation

Splunk Infrastructure Monitoring provides out-of-the-box integration with Consul and the integration is enabled in one single, one-time step by telling Splunk OpenTelemetry Collector to start collecting and ingesting Consul metrics and events.

When deploying in Kubernetes environments, simply add the following configuration while deploying [Splunk OpenTelemetry Collector](#). Refer to detailed [documentation](#) for a list of metrics and configurations options.

```
otelAgent:
  config:
    receivers:
      smartagent/consul:
        type: collectd/consul
        host: localhost
        port: 8500
        enhancedMetrics: true
```

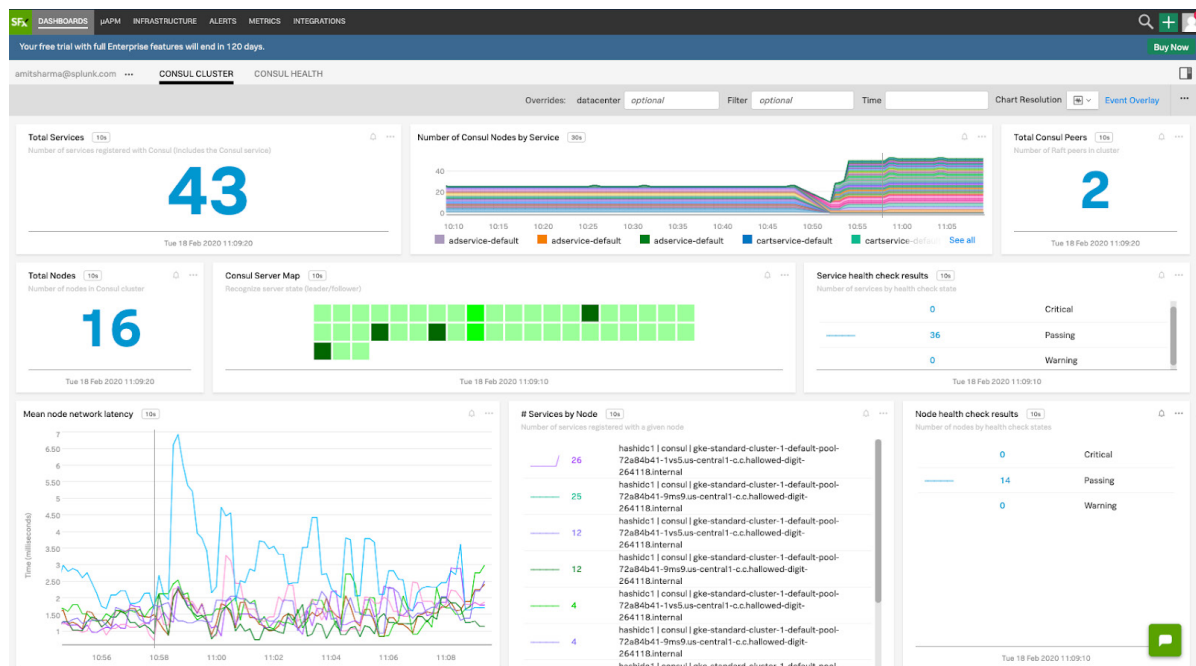

When running as a [service mesh](#), Envoy's metrics can be seamlessly ingested by Prometheus exporter:

```
smartagent/prometheus-exporter:  
  type: prometheus-exporter  
  host: 0.0.0.0  
  port: 9102
```

That's it! Splunk Infrastructure Monitoring will start collecting all relevant data instantly.


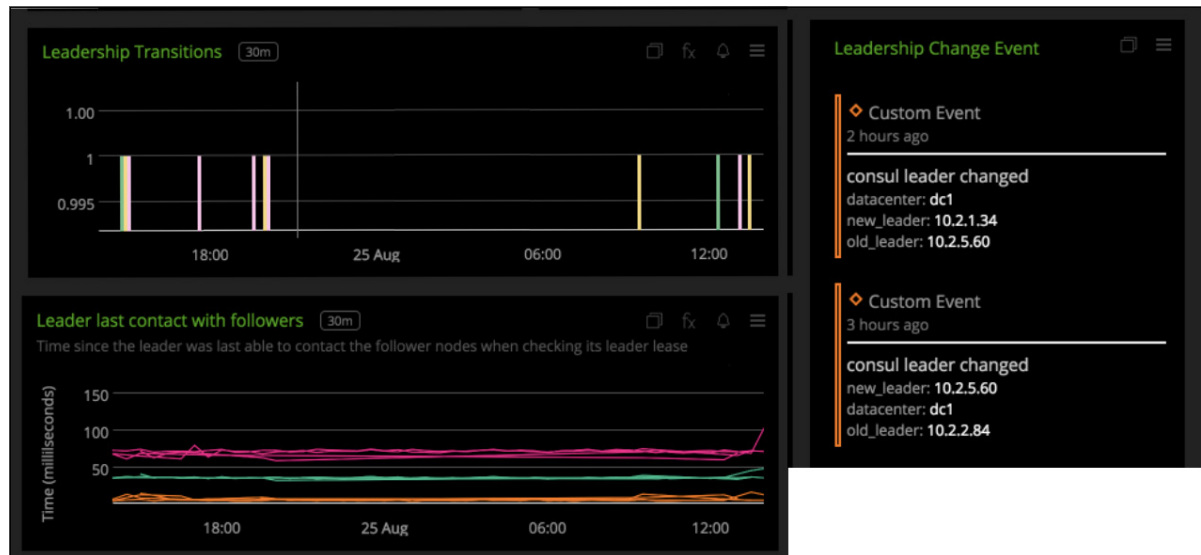
Real-time Visualization of Consul Cluster

Splunk Infrastructure Monitoring provides an out-of-the-box dashboard to visualize the health and performance of the entire Consul environment. This dashboard can be cloned and easily customized to correlate Consul performance with other technologies in your stack, for example, cloud infrastructure, Kubernetes or application-level metrics.



Monitoring the Health of Consul Cluster

Splunk Infrastructure Monitoring provides an out-of-the-box view into the performance of Consul cluster with pre-built dashboards. Key metrics to track are leadership transitions, Raft lifecycle related metrics and Gossip metrics related to node visibility to the cluster. A full list of available metrics is listed in our [documentation](#) for Consul integration.



Conclusion

In the preceding sections, we presented recipes to unlock the fastest path to value in modern multi-cloud environments by achieving real-time Observability with Splunk and adopting a common Cloud Operating Model by HashiCorp. The combination of Splunk and HashiCorp portfolio delivers:

- Monitoring as code to provision monitoring assets alongside infrastructure provisioning with HashiCorp Terraform
- Complete visibility and compliance audits into authentication authorization and credential access requests with HashiCorp Vault and Splunk
- System-wide Observability with HashiCorp Consul and Splunk

Ready to learn how we can accelerate your multicloud success with the cloud operating model? Sign-up for a free trial of [Splunk Infrastructure Monitoring](#) and get real-time visibility into your cloud infrastructure and services. Future-proof your observability investment with a proven solution trusted by thousands of enterprises globally.

Learn more:

Refer to the following blogs and tutorials to start your observability journey in the cloud with HashiCorp and Splunk:

[Implement Observability as Code with HashiCorp and Splunk](#)

[Manage Your Splunk Infrastructure as Code using Terraform](#)

[HashiCorp Terraform Cloud Audit Logging with Splunk](#)

[Operationalizing HashiCorp Vault: Introducing a New Splunkbase App To Monitor Vault](#)

[Monitor Telemetry and Audit Device Log Data with Splunk](#)

