



Unlocking the Cloud Operating Model

Achieving the fastest path to value in a modern, multi-cloud world



Executive Summary

To thrive in an era of multi-cloud architecture, driven by digital transformation, Enterprise IT must evolve from ITIL-based gatekeeping to enabling shared self-service processes for DevOps excellence.

For most enterprises, 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 the 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, processes, and tools to it.

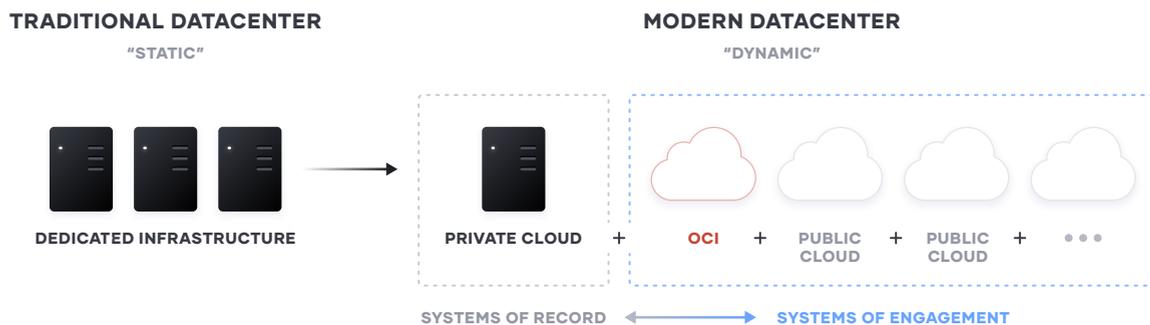
In this white paper, HashiCorp and Oracle Cloud Infrastructure (OCI) look at the implications of the cloud operating model, and present solutions for enterprises to adopt this model across infrastructure, security, networking, and application delivery.

Motivating Changes and Business Drivers for Cloud Adoption

- Increase Business Agility
- Reduce Capital Costs
- Reduce Operational Costs
- Reduce Business Complexity

Transitioning to a Multi-Cloud Datacenter

The transition to cloud, and multi-cloud, environments is a generational transition for IT. This transition means shifting from largely dedicated servers in a private datacenter to a pool of compute capacity available on demand. While most enterprises began with one cloud provider, there are good reasons to use services from others and inevitably most Global 2000 organizations will use more than one, either by design or through mergers and acquisitions.



The cloud presents an opportunity for speed and scale optimization for new “systems of engagement” -- the applications built to engage customers and users. These new apps are the primary interface for the customer to engage with a business, and are ideally suited for delivery in the cloud as they tend to:

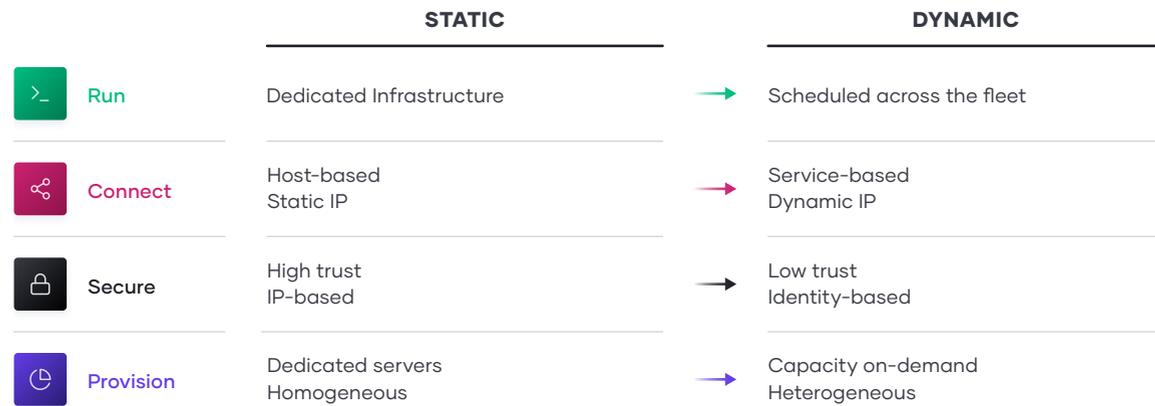
- Have dynamic usage characteristics, needing to scale loads up and down by orders of magnitude during short time periods.
- Be under pressure to quickly build and iterate. Many of these new systems may be ephemeral in nature, delivering a specific user experience around an event or campaign.

There are situations where customers will require highly regulated or security-focused businesses to meet their demanding latency and data residency requirements. This is something that a public cloud may not be able to provide unless it's in the same data center as your infrastructure. If that is not the case, Oracle Cloud brings its complete portfolio of public cloud services and Oracle Fusion SaaS applications into your data center so you can reduce data center costs, upgrade legacy applications using modern services, and meet your most demanding requirements. You will connect to Oracle Dedicated Region Cloud@Customer just like you do today to a public region using Oracle tools, SDKs, and the management console.

As a result, enterprises end up with a true public cloud with two flavors, depending on your requirements. The Oracle Dedicated Region Cloud@Customer brings best-in-class price-performance and security to mission-critical workloads that are unlikely to move to the public cloud for several years.

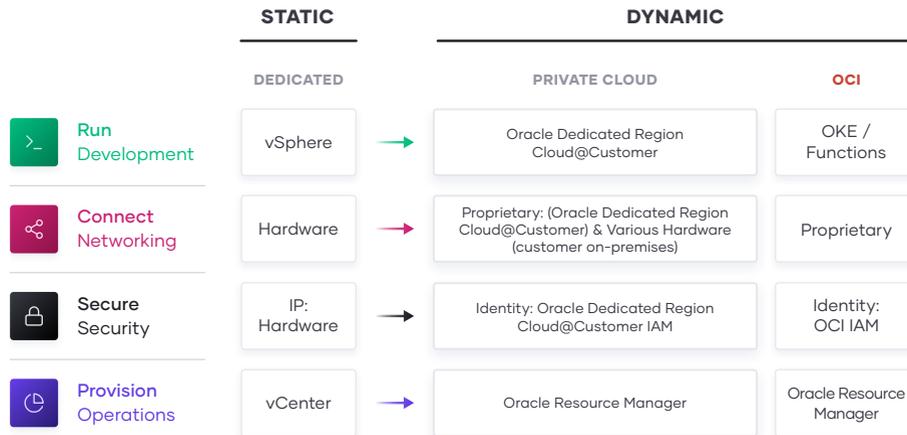
Implications of the Cloud Operating Model

The essential implication of the transition to the cloud is the shift from “static” infrastructure to “dynamic” infrastructure: from a focus on configuration, and management of a static fleet of IT resources, to provisioning, securing, connecting, and running dynamic resources on demand.



Decomposing this implication, and working up the stack, various changes of approach are implied:

- **Provision.** The infrastructure layer transitions from running dedicated servers at 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.
- **Secure.** 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.
- **Connect.** 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 does not have the same control over the network, or the physical locations of compute resources, and must think about service-based connectivity.
- **Run.** The runtime layer shifts from deploying artifacts to a static application server to deploying applications with a scheduler atop a pool of infrastructure which is provisioned on-demand. In addition, new applications have become collections of services that are dynamically provisioned, and packaged in multiple ways: from virtual machines to containers.

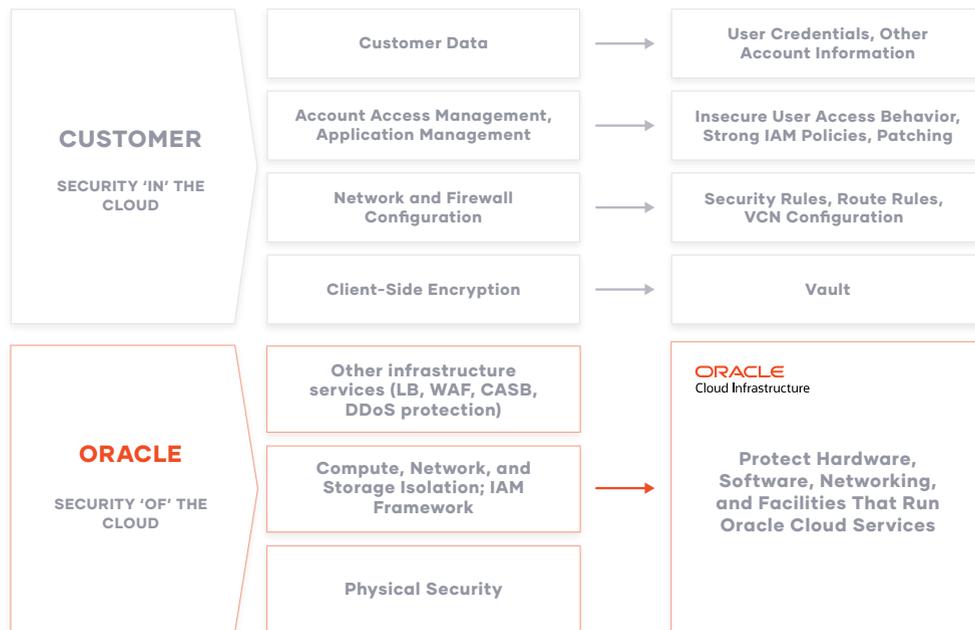


To address these challenges those teams must ask the following questions:

- **People.** How can we enable a team for a multi-cloud reality, where skills can be applied consistently regardless of target environment?
- **Process.** How do we position central IT services as a self-service enabler of speed, versus a ticket-based gatekeeper of control, while retaining compliance and governance?
- **Tools.** How do we best unlock the value of the available capabilities of the cloud providers in pursuit of better customer and business value?

OCI Shared Responsibility Model

A secure journey to the cloud requires that organizations develop a core competency around understanding the cloud security shared responsibility model. The dynamic nature of the cloud means such a competency must also account for change. The following graphic illustrates Oracle Cloud shared security responsibility model:



Oracle is solely responsible for all aspects of the physical security of the availability domains and fault domains in each region. Both Oracle and the customer are responsible for the infrastructure security of hardware, software, and the associated logical configurations and controls.

As a customer, your security responsibilities encompass the following:

- The platform you create on top of Oracle Cloud.
- The applications that you deploy.
- The data that you store and use.
- The overall governance, risk, and security of your workloads.

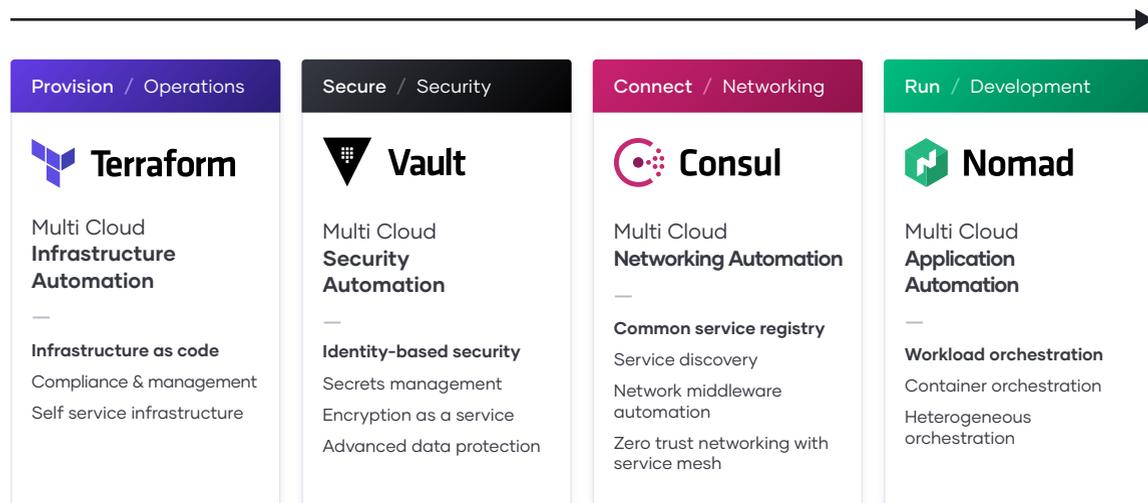
The shared responsibility extends across different domains including identity management, access control, workload security, data classification and compliance, infrastructure security, and network security.

Unlocking the Cloud Operating Model on OCI

As the implications of the cloud operating model impact teams across infrastructure, security, networking, and applications, we see a repeating pattern amongst enterprises of establishing central shared services - centers of excellence - to deliver the dynamic infrastructure necessary at each layer for successful application delivery.

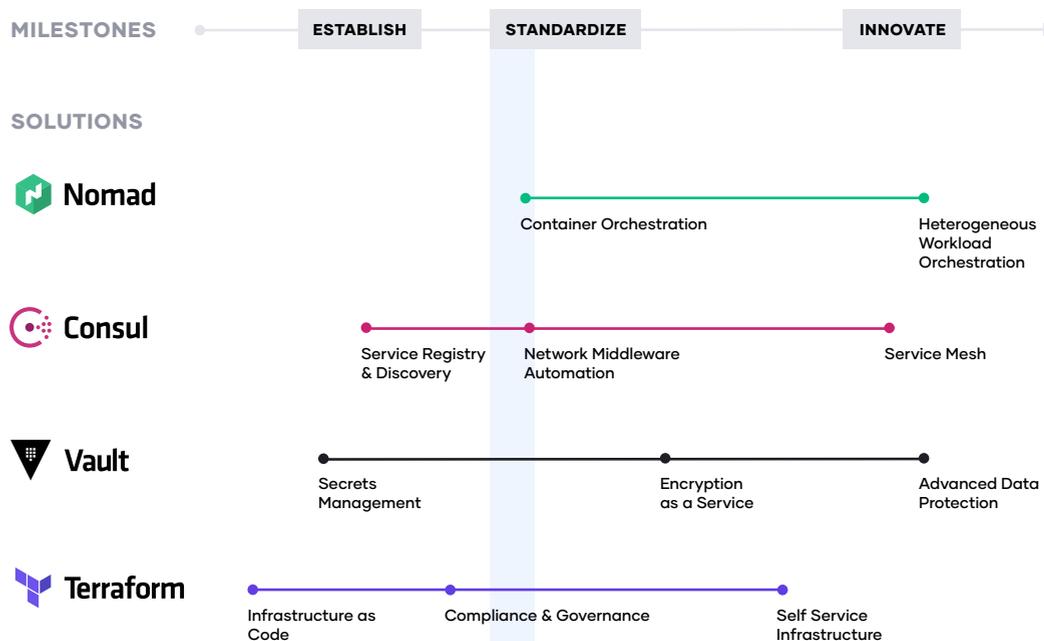
As teams deliver on each shared service for the cloud operating model, IT velocity increases. The greater cloud maturity an organization has, the faster its velocity.

EXPANDING USE OF THE HASHICORP STACK INCREASES MATURITY AND VELOCITY FOR OUR CUSTOMERS



The typical journey we have seen customers adopt, as they unlock the cloud operating model, involves three major milestones

- 1. Establish the cloud essentials** – As you begin your journey to the cloud, the immediate requirements are provisioning the cloud infrastructure typically by adopting infrastructure as code and ensuring it is secure with a secrets management solution. These are the bare necessities that will allow you to build a scalable and truly dynamic cloud architecture that is futureproof.
- 2. Standardize on a set of shared services** – As cloud consumption starts to pick up, you will need to implement and standardize on a set of shared services so as to take full advantage of what the cloud has to offer. This also introduces challenges around governance and compliance as the need for setting access control rules and tracking requirements become increasingly important.
- 3. Innovate using a common logical architecture** – As you fully embrace the cloud and depend on cloud services and applications as the primary systems of engagement, there will be a need to create a common logical architecture. This requires a control plane that connects with the extended ecosystem of cloud solutions and inherently provides advanced security and orchestration across services and multiple clouds.

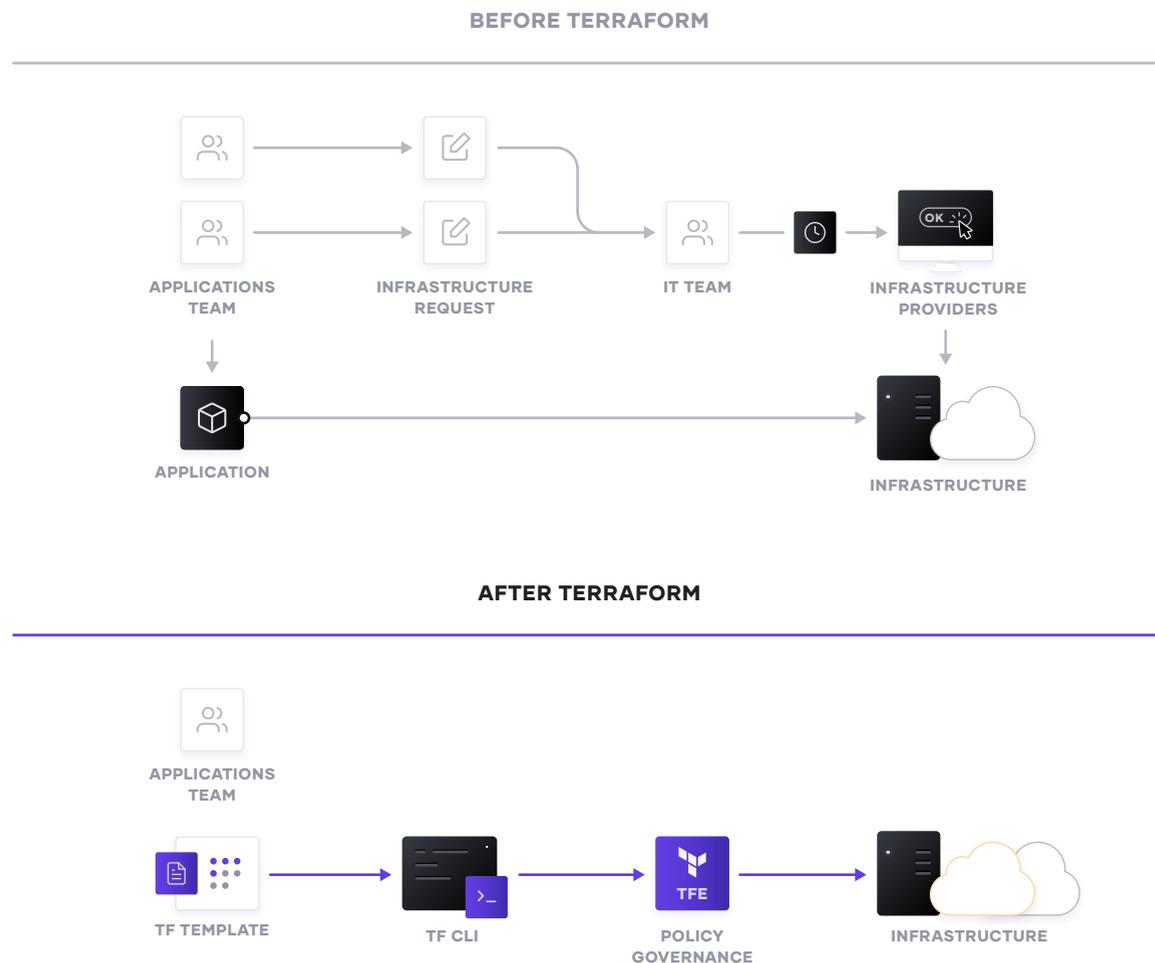


What follows is the step-by-step journey that we have seen organizations adopt successfully.

Step 1: Multi-Cloud Infrastructure Provisioning

The foundation for adopting the cloud is infrastructure provisioning. 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. Oracle Cloud Infrastructure offers a fully managed service, inOCI Resource Manager which integrates with the industry-standard of HashiCorp Terraform. To provision and manage OCI-based infrastructure, customers have multiple options: OCI Resource Manager, Terraform CLI, or Terraform Cloud.

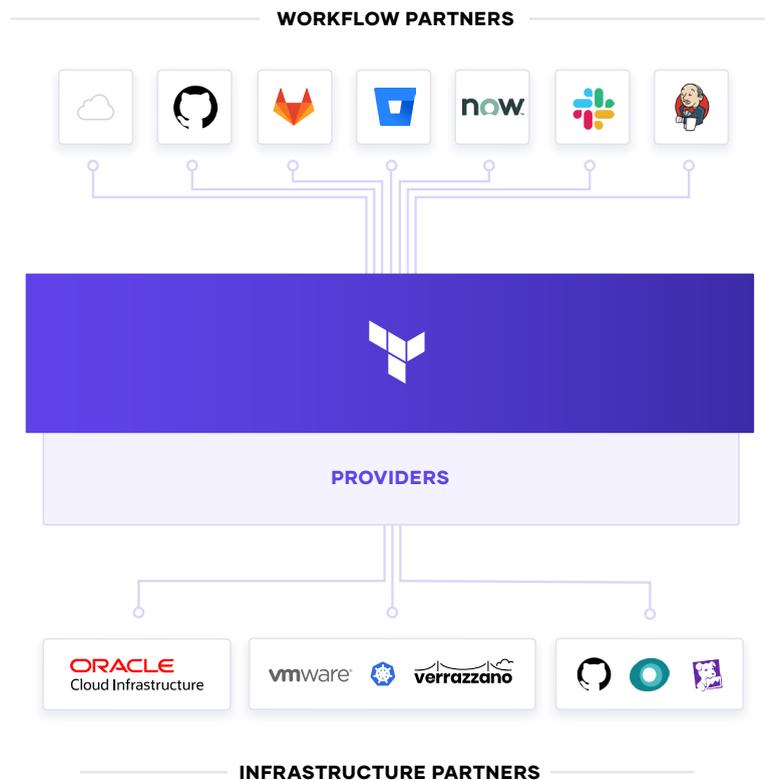
To achieve shared services for infrastructure provisioning, IT teams should start by implementing reproducible infrastructure as code practices, and then layering compliance and governance workflows to ensure appropriate controls.



Reproducible Infrastructure as Code

The first goal of a shared service for infrastructure provisioning is to enable the delivery of reproducible infrastructure as code, providing DevOps teams a way to plan and provision resources inside CI/CD workflows using familiar tools throughout.

DevOps teams can create Terraform modules that express the configuration of services from one or more cloud platforms. Terraform integrates with all major configuration management tools to allow fine-grained provisioning to be handled following the provisioning of the underlying resources. Finally, templates can be extended with services from many other ISV providers to include monitoring agents, application performance monitoring (APM) systems, security tooling, DNS, and Content Delivery Networks, and more. Once defined, the templates can be provisioned as required in an automated way. In doing so, Terraform becomes the lingua franca and common workflow for teams provisioning resources across public and private cloud.



But, what happens when you already have infrastructure deployed in the cloud and need to build Terraform templates from existing applications? Oracle Resource Manager and OCI Terraform provide an easy way to discover resources in your environment and export them to Terraform configuration and state files. You can now standardize all your environments using Terraform.

For self-service IT, the decoupling of the template-creation process and the provisioning process greatly reduces the time taken for any application to go live since developers no longer need to wait for operations approval, as long as they use a pre-approved template.

Compliance and Management

For most teams, there is also a need to enforce policies on the type of infrastructure created, how it is used, and which teams get to use it. HashiCorp's Sentinel policy as code framework provides compliance and governance without requiring a shift in the overall team workflow, and is defined as code too, enabling collaboration and comprehension for DevSecOps.

Without policy as code, organizations resort to using a ticket-based review process to approve changes. This results in developers waiting weeks or longer to provision infrastructure and becomes a bottleneck. Policy as code allows us to solve this by splitting the definition of the policy from the execution of the policy.

Centralized teams codify policies enforcing security, compliance, and operational best practices across all cloud provisioning. Automated enforcement of policies ensures changes are in compliance without creating a manual review bottleneck.

Step 2: Multi-Cloud Security

Dynamic cloud infrastructure means a shift from host-based identity to application-based identity, with low- or zero-trust networks across multiple clouds without a clear network perimeter.

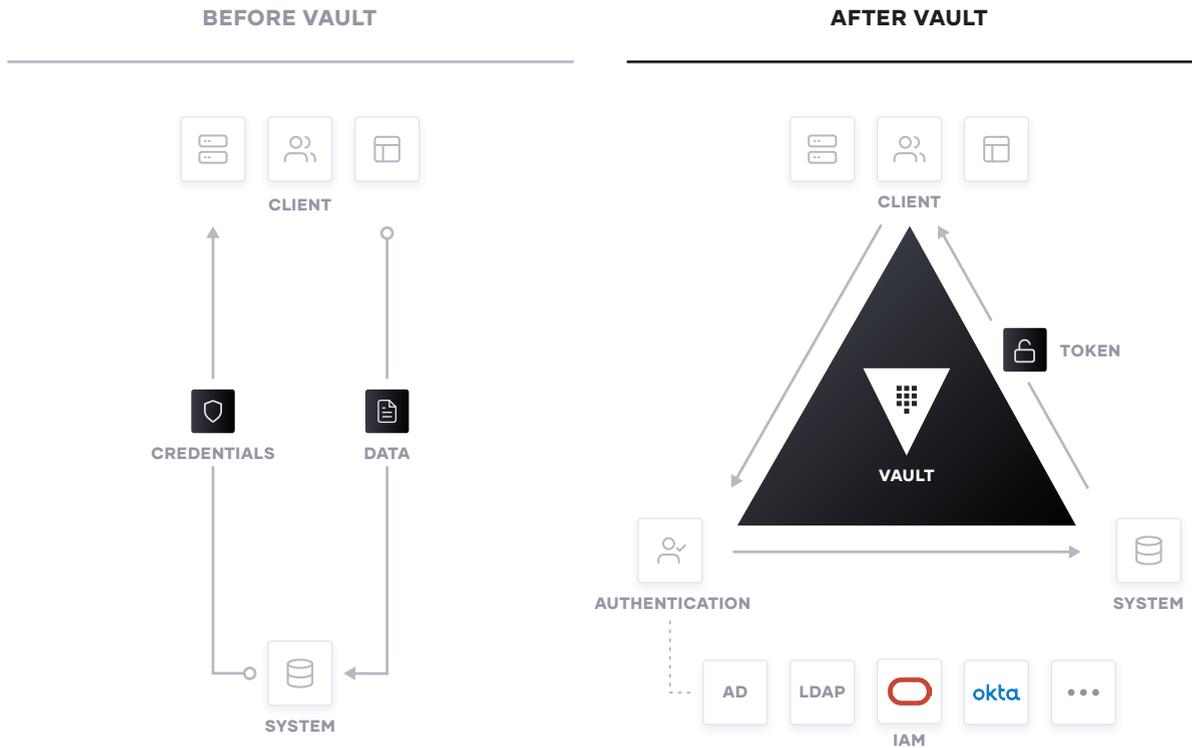
In the traditional security world, we assumed high trust internal networks, which resulted in a hard shell and soft interior. With the modern “zero trust” approach, we work to harden the inside as well. This requires that applications be explicitly authenticated, authorized to fetch secrets and perform sensitive operations, and tightly audited.

HashiCorp 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. Beyond that, Vault helps protect data at rest and data in transit. Vault exposes a high level API for cryptography for developers to secure sensitive data without exposing encryption keys. Vault can also act like a certificate authority, to provide dynamic short lived certificates to secure communications with SSL/TLS. Lastly, Vault enables a brokering of identity between different platforms, such as Active Directory on premises and OCI IAM to allow applications to work across platform boundaries.

Vault is widely used including across stock exchanges, large financial organizations, hotel chains, and everything in between to provide security in the cloud operating model.

To achieve shared services for security, IT teams should enable centralized secrets management services, and then use that service to deliver more sophisticated encryption-as-a-service use cases such as certificate and key rotations, and encryption of data in transit and at rest.

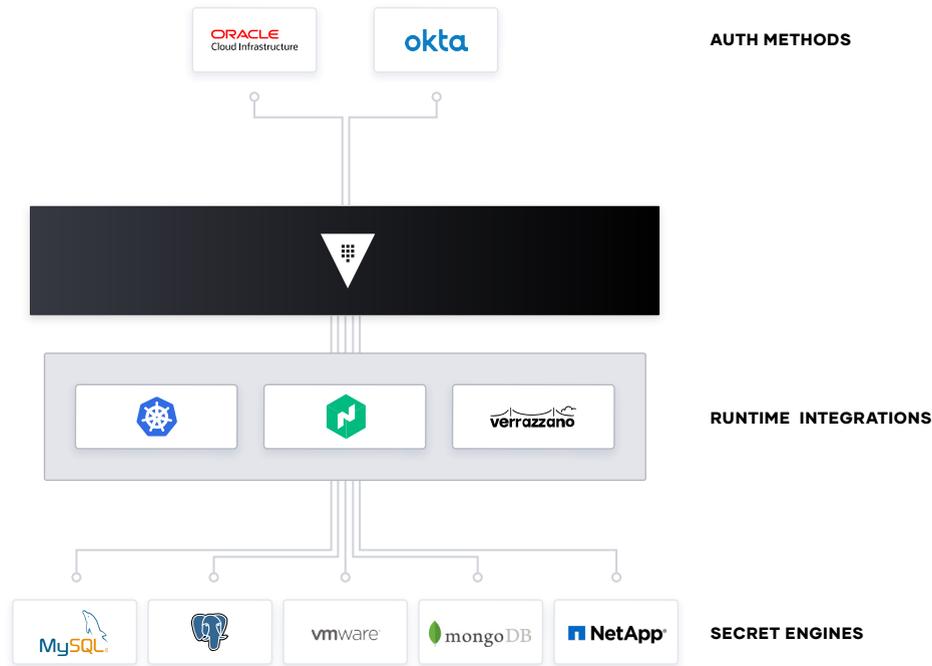
Oracle Cloud Infrastructure provides IT teams a managed service that enables them to manage and control AES symmetric keys used to encrypt their data-at-rest. The OCI Vault service is integrated with identity and audit services. IT teams can monitor the key lifecycle with Oracle Audit to meet enhanced compliance requirements and control permissions for individual keys and vaults using OCI IAM.



Secrets Management

The first step in cloud security is typically secrets management: the central storage, access control, and distribution of dynamic secrets. Instead of depending on static IP addresses, integrating with identity-based access systems such as OCI IAM and Azure AAD to authenticate and access services and resources is crucial. Oracle Cloud Infrastructure provides a unified identity and access management, via a unified single sign-on experience and automated user provisioning, to manage resources across Azure and Oracle Cloud.

Vault uses policies to codify how applications authenticate, which credentials they are authorized to use, and how auditing should be performed. It can integrate with an array of trusted identity providers such as cloud identity and access management (IAM) platforms, Kubernetes, Active Directory, and other SAML-based systems for authentication. Vault then centrally manages and enforces access to secrets and systems based on trusted sources of application and user identity.



Enterprise IT teams should build a shared service which enables the request of secrets for any system through a consistent, audited, and secured workflow.

Encryption as a Service

Additionally, enterprises need to encrypt application data at rest and in transit. Vault can provide encryption-as-a-service to provide a consistent API for key management and cryptography. This allows developers to perform a single integration and then protect data across multiple environments.

Using Vault as a basis for encryption-as-a-service solves difficult problems faced by security teams such as certificate and key rotation. Vault enables centralized key management to simplify encrypting data in transit and at rest across clouds and datacenters. This helps reduce costs around expensive Hardware Security Modules (HSM) and increases productivity with consistent security workflows and cryptographic standards across the organization.

While many organizations provide a mandate for developers to encrypt data, they don't often provide the "how" which leaves developers to build custom solutions without an adequate understanding of cryptography. Vault provides developers a simple encryption as a service API that can be easily used, while giving central security teams the policy controls and lifecycle management APIs they need.

Advanced Data Protection

Organizations moving to the cloud or spanning hybrid environments still maintain and support on-premise services and applications that need to perform cryptographic operations, such as data encryption for storage at rest. These services do not necessarily want to implement the logic around managing these cryptographic keys, and thus seek to delegate the task of key management to external providers. Advanced Data Protection allows organizations to securely connect, control, and integrate advanced encryption keys, operations, and management between infrastructure and Vault Enterprise, including automatically protecting data in MySQL, MongoDB, PostgreSQL, and Oracle databases using transparent data encryption (TDE).

For organizations that have high security requirements for data compliance (PCI-DSS, HIPAA, etc), protecting data, and cryptographically-protecting anonymity for personally identifiable information (or PII), Advanced Data Protection provides organizations with functionality for data tokenization, such as data masking, to protect sensitive data, such as credit cards, sensitive personal information, bank numbers, etc.

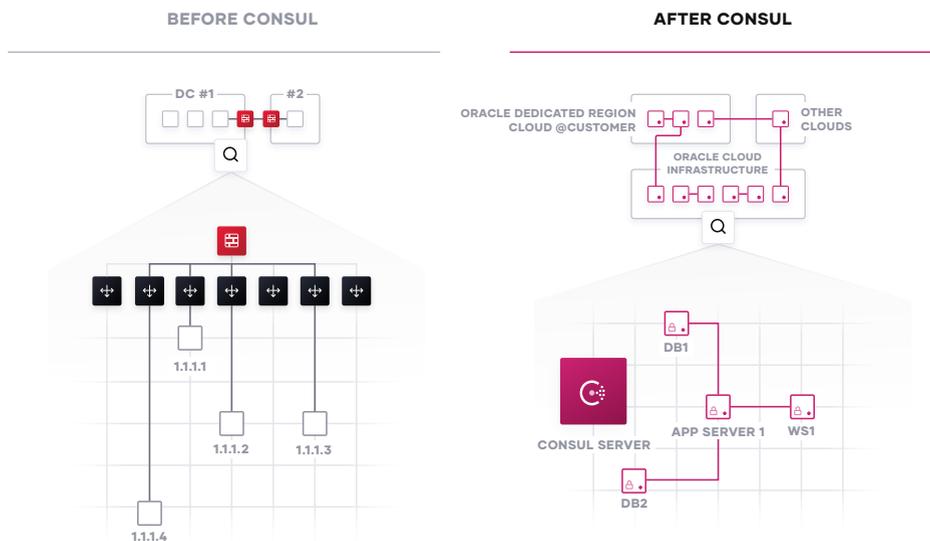
Step 3: Multi-Cloud Service Networking

The challenges of networking in the cloud are often one of the most difficult aspects of adopting the cloud operating model for enterprises. The combination of dynamic IP addresses, a significant growth in east-west traffic as the microservices pattern is adopted, and the lack of a clear network perimeter is a formidable challenge.

HashiCorp Consul provides a multi-cloud service networking layer to connect and secure services. Consul is a widely deployed product, with many customers running significantly greater than 100,000 nodes in their environments.

Networking services should be provided centrally, where by IT teams provide service registry and service discovery capabilities. Having a common registry provides a “map” of what services are running, where they are, and their current health status. The registry can be queried programmatically to enable service discovery or drive network automation of API gateways, load balancers, firewalls, and other critical middleware components. These middleware components can be moved out of the network by using a service mesh approach, where proxies run on the edge to provide equivalent functionality. Service mesh approaches allow the network topology to be simplified, especially for multi-cloud and multi-datacenter topologies.

Using HashiCorp Consul, customers can manage Oracle Cloud Infrastructure and other cloud environments to ensure data, users and clouds stay connected. This helps mitigate a customer’s pain by delivering a centralized, easy-to-deploy suite of provisioning and monitoring networking tools that are running in Oracle Cloud Infrastructure.



Service Discovery

The starting point for networking in the cloud operating model is typically a common service registry, which provides a real-time directory of what services are running, where they are, and their current health status. Traditional approaches to networking rely on load balancers and virtual IPs to provide a naming abstraction to represent a service with a static IP. 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.

For Consul, each service is programmatically registered and DNS and API interfaces are provided to enable any service to be discovered by other services. The integrated health check will monitor each service instance'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 can be integrated with other services that manage existing north-south traffic such as a traditional load balancers, and distributed application platforms such as Kubernetes, to provide a consistent registry and discovery service across multi-datacenter, cloud, and platform environments.

Network Infrastructure Automation

The next step is to reduce operational complexity with existing networking middleware through network automation. Instead of a manual, ticket-based process to reconfigure load balancers and firewalls every time there is a change in service network locations or configurations, Consul can automate these network operations. This is achieved by enabling network middleware devices to subscribe to service changes from the service registry, enabling highly dynamic infrastructure that can scale significantly higher than static-based approaches.

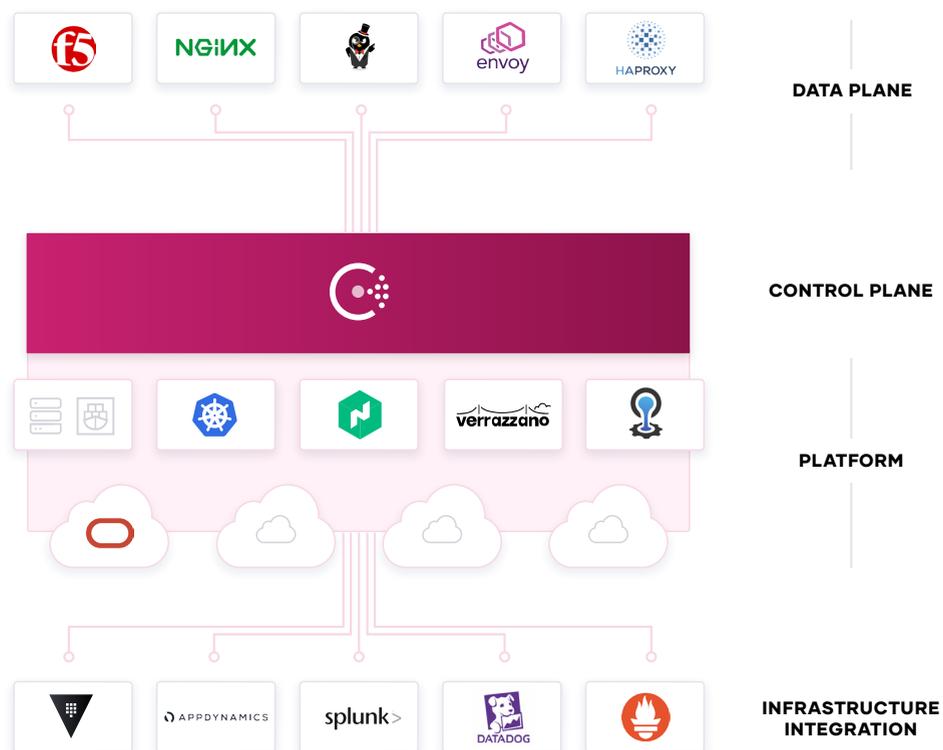
This decouples the workflow between teams, as operators can independently deploy applications and publish to Consul, while NetOps teams can subscribe to Consul to handle the downstream automation.

Zero Trust Networking with Service Mesh

As organizations continue to scale with microservices-based or cloud-native applications, the underlying infrastructure becomes larger and more dynamic with an explosion of east-west traffic. This causes a proliferation of expensive network middleware with single points of failure and significant operational overhead exposed to IT teams.

Consul provides a distributed service mesh that pushes routing, authorization, and other networking functionalities to the endpoints in the network, rather than imposing them through middleware. This makes the network topology simpler and easier to manage, removes the need for expensive middleware within east-west traffic paths, and it makes service-to-service communication much more reliable and scalable.

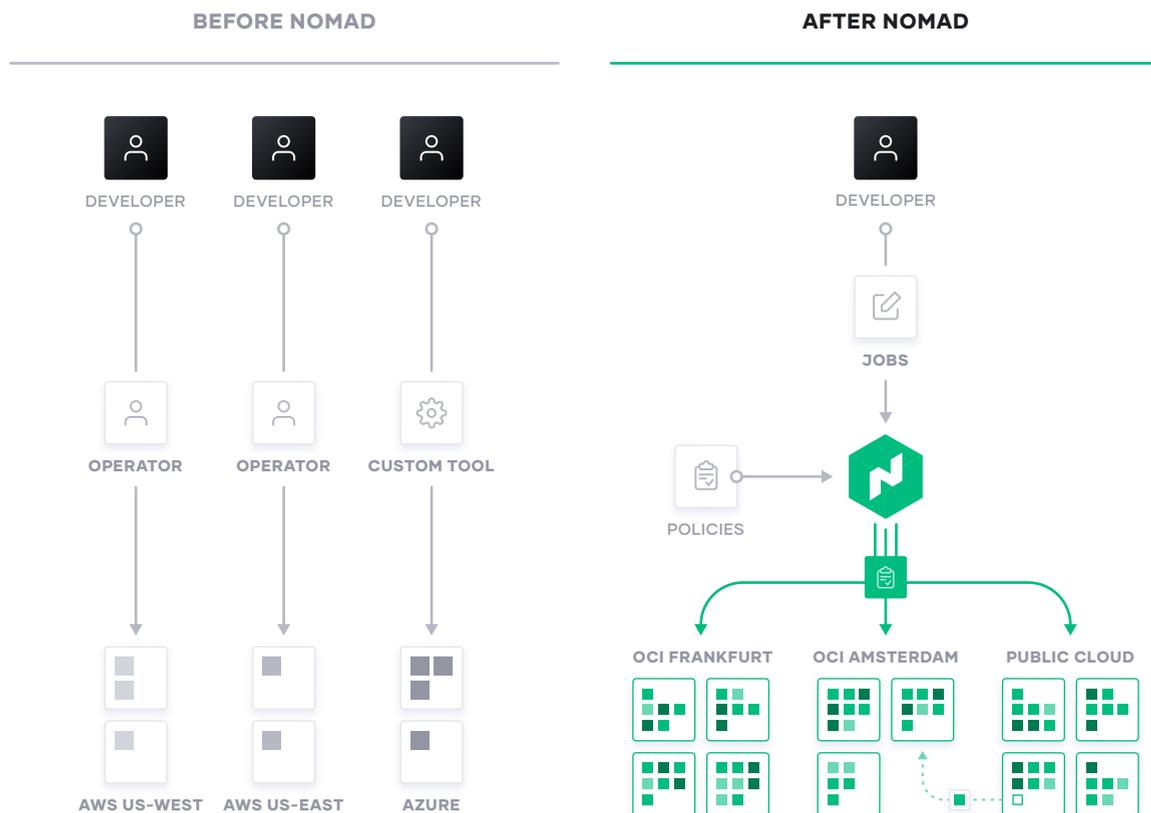
Consul is an API-driven control plane that integrates with sidecar proxies alongside each service instance (proxies such as Envoy, HAProxy, and NGINX). These proxies provide the distributed data plane. Together, these two planes enable a zero trust network model that secures service-to-service communication with automatic 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.



Step 4: Multi-Cloud Application Delivery

Finally, at the application layer, new apps are increasingly distributed while legacy apps also need to be managed more flexibly. HashiCorp Nomad provides a flexible orchestrator to deploy and manage legacy and modern applications, for all types of workloads: from long running services, to short lived batches, to system agents.

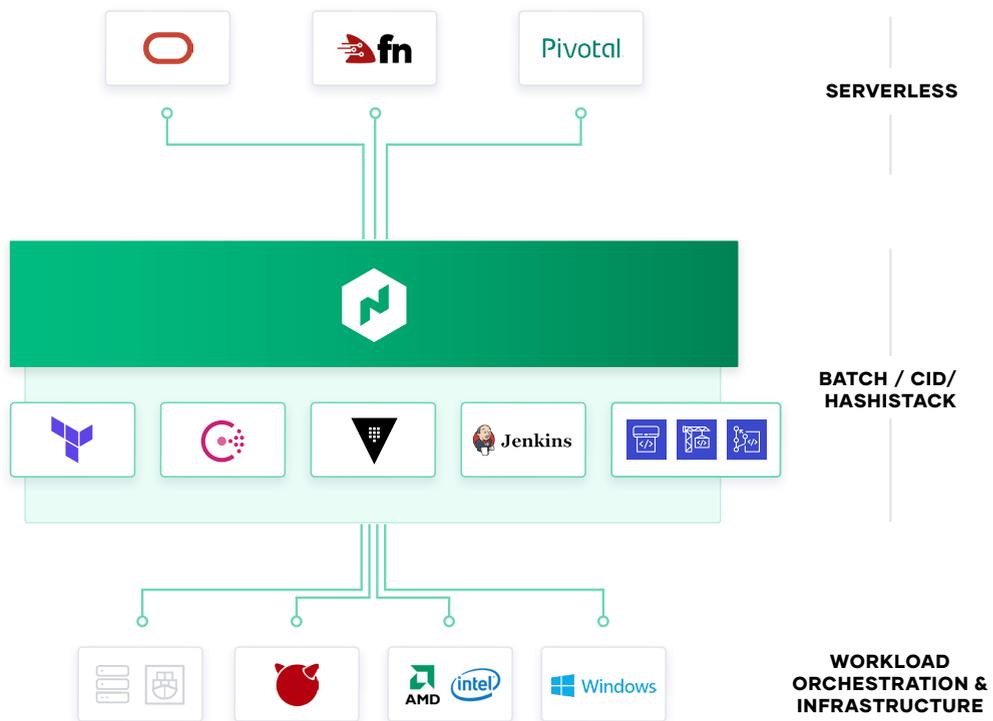
To achieve shared services for application delivery, IT teams should use Nomad in concert with Terraform, Vault, and Consul to enable the consistent delivery of applications on cloud infrastructure, incorporating necessary compliance, security, and networking requirements, as well as workload orchestration and scheduling.



Mixed Workload Orchestration

Many new workloads are developed with container packaging with the intent to deploy to Kubernetes or other container management platforms. But many legacy workloads will not be moved onto those platforms, nor will future serverless applications. Nomad provides a consistent process for deployment of all workloads from virtual machines, through standalone binaries, and containers, and provides core orchestration benefits across those workloads such as release automation, multiple upgrade strategies, bin packing, and resilience.

For modern applications - typically built in containers - Nomad provides the same consistent workflow at scale in any environment. Nomad is focused on simplicity and effectiveness at orchestration and scheduling, and avoids the complexity of platforms such as Kubernetes that require specialist skills to operate and solve only for container workloads.



Nomad integrates into existing CI/CD workflows to provide fast, automatic application deployments for legacy and modern workloads.

High Performance Compute

Nomad is designed to schedule applications with low latency across very large clusters. This is critical for customers with large batch jobs, as is common with High Performance Computing (HPC) workloads. OCI offers tremendous HPC performance. In an OpenFOAM 24M Cell Model benchmark, OCI HPC showed performance consistent with some of Japan's most powerful supercomputers, including the Research Institute for Information Technology (ITO), Kyushu University and Osaka University Cybermedia Center Over-Petascale Universal Supercomputer. Using identical infrastructure (processors, memory per core, etc.), OCI HPC achieved similar performance while delivering substantial cost savings!

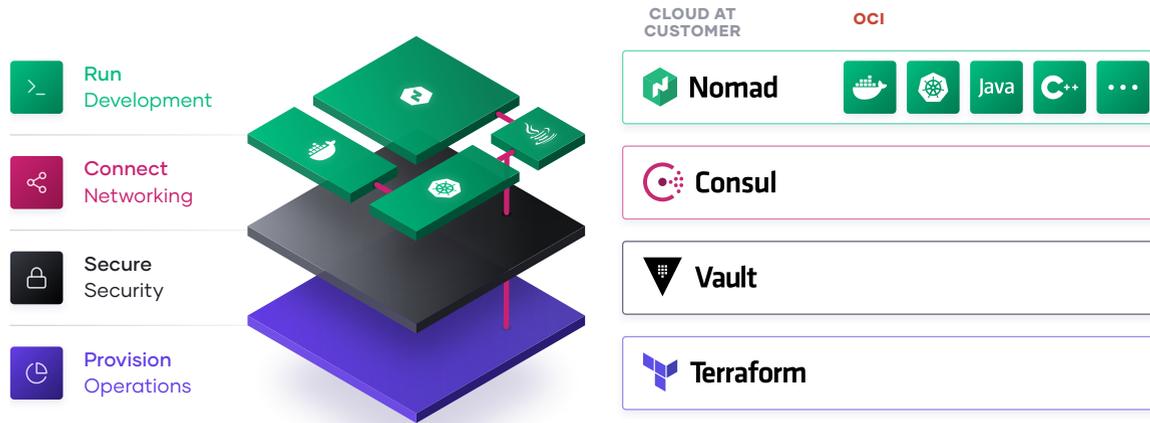
Nomad enables high performance applications to easily use an API to consume capacity dynamically, enabling efficient sharing of resources for data analytics applications like Spark. The low latency scheduling ensures results are available in a timely manner and minimizes wasted idle resources.

Multi-Datacenter Workload Orchestration

Nomad is multi-region and multi-cloud by design, with a consistent workflow to deploy any workload. As teams roll out global applications in multiple data centers, or across cloud boundaries, Nomad provides orchestrating and scheduling for those applications, supported by the infrastructure, security, and networking resources and policies to ensure the application is successfully deployed.

Step 5: Industrialized Application Delivery Process

Ultimately, these shared services across infrastructure, security, networking, and application runtime present an industrialized process for application delivery, all while taking advantage of the dynamic nature of each layer of the cloud. Embracing the cloud operating model enables self-service IT, that is fully compliant and governed, for teams to deliver applications at increasing speed.



OCI Best Practices Framework

The OCI Best Practice Framework has been developed to help cloud architects plan their journey to the cloud and ensure their applications workloads are designed in a secure, resilient, highly available and efficient way.

Based on five pillars – operational efficiency, security, reliability, performance and cost optimization – the Framework provides guidance on Oracle recommended best practices to design and operate cloud topologies that deliver business value.

Operational Efficiency. The operational efficiency pillar focuses on running and monitoring your applications and infrastructure resources to deliver the maximum business value. Key topics include, day to day operations, responding to events, and adopting standards to successfully support your application workloads.

Security. The security pillar focuses on protecting your systems and data. Key topics include identity and access management, securing your infrastructure resources, managing Oracle Audit services and defining controls to detect security events and mitigating risks.

Reliability. The reliability pillar focuses on the ability to prevent, and quickly recover from failures to meet business and customer demand. Key topics include project requirements, recovery planning, and how to manage changes on demand.

Performance. The performance pillar focuses on using infrastructure resources efficiently. Key topics include selecting the right resource types and sizes using Oracle Flexible shapes, monitoring performance, and running predictable high-performance and low latency infrastructure as business evolves.

Cost Optimization. Cost Optimization focuses on removing unneeded costs. Key topics include understanding spending patterns, monitoring consumption, analyzing your bill, and reducing spending through Oracle Cloud Infrastructure Cost Analysis, Budgets, and Usage Reports.

[Best practices framework for Oracle Cloud Infrastructure](#) 

Conclusion

In conclusion, a common cloud operating model is an inevitable shift for enterprises aiming to maximize their digital transformation efforts. The HashiCorp suite of tools and Oracle Cloud Infrastructure seeks to provide solutions for each layer of the cloud to enable enterprises to make this shift to the cloud operating model.

Enterprise IT needs to evolve away from ITIL-based control points with its focus on cost optimization, toward becoming self-service enablers focused on speed optimization. It can do this by delivering shared services across each layer of the cloud, designed to assist teams in delivering new business and customer value at speed.

Oracle Cloud Infrastructure provides a robust, scalable, secure, highly available and cost-effective cloud platform to meet the needs of the modern enterprise. The technologies provided by Oracle Cloud Infrastructure, including the Dedicated Cloud@Customer region and the Hashicorp stack give developers and IT broad support for migrating existing workloads as well as creating new cloud-native applications.

About HashiCorp

HashiCorp is the leader in multi-cloud infrastructure automation software. The HashiCorp software suite enables organizations to adopt consistent workflows to provision, secure, connect, and run any infrastructure for any application. HashiCorp open source tools Vagrant, Packer, Terraform, Vault, Consul, and Nomad are downloaded tens of millions of times each year and are broadly adopted by the Global 2000. Enterprise versions of these products enhance the open source tools with features that promote collaboration, operations, governance, and multi-data center functionality. The company is headquartered in San Francisco and backed by Mayfield, GGV Capital, Redpoint Ventures, True Ventures, IVP, and Bessemer Venture Partners. For more information, visit www.hashicorp.com or follow HashiCorp on Twitter [@HashiCorp](https://twitter.com/HashiCorp).

