HashiCorp

A leadership guide to cloud success for educational organizations

Achieving the fastest path to optimize modern, hybrid-cloud datacenters

WHITEPAPER

Table of contents

Executive summary

The IT environment for educational organizations is changing dramatically. To thrive in the modern digital world, IT leaders must evolve from ITIL-based gatekeeping to enabling shared self-service processes for improved digital experiences.

Cloud is quickly becoming the default choice for organizations to deliver new value to their students and faculty. Forward-looking state and local organizations must adopt a cloud operating model — a framework for adopting cloud services — to maximize their agility, reliability, and security in order to deliver superior user experiences and outcomes.

Digital experiences are the primary interface between students and the organizations that support them. Modern digital interactions are responsively designed and built cloud-first to provide rich, personalized experiences informed by large-scale data processing and intelligence as quickly as possible.

For most educational organizations, these digital transformation efforts mean delivering new solutions and constituent value faster, often at a very large scale. Higher education organizations undergoing a digital transformation inevitably put pressure on the teams delivering and supporting software applications. The cloud is an inevitable part of this shift as it presents the opportunity to rapidly deploy on-demand services with limitless scale. In addition, the implementation of automated, self-service platforms is essential to help these organizations overcome widespread skills gaps and free up critical IT resources to focus on end-user value.

This white paper looks at the implications of the cloud operating model for educational organizations and presents solutions for IT teams to adopt this model to drive successful cloud adoption.

Introduction

A cloud operating model is essential for educational organizations to succeed with cloud adoption and thrive in modern, hybrid cloud environments. This white paper explains the components of that approach by focusing on proven patterns for standardizing how people work, the processes they follow, and the tools they consume.

At educational organizations, student services, registrations, classes, research, activities, and more are quickly becoming digital or hybrid experiences. Administrators, faculty, parents, and increasingly tech-savvy students demand that these activities be easy to use, fast, and efficient, or they will quickly lose traction. Additionally, implementation of digital systems is essential to help education institutions optimize resources and save money by moving away from existing manual processes that waste time, money, and resources.

In many ways, educational organizations are like small governments, managing huge swaths of data they should be leveraging to better serve the needs of their students and keep up with the ever-changing educational landscape. Educational institutions need tools to easily collect, store, secure, and analyze this data to help set organizational priorities and investments.

3 key cloud initiatives for educational institutions

As a whole, educational organizations must focus on three main cloud initiatives as they manage their digital transformation:

1. **Education institutions must accelerate transitions to the cloud to adapt to the changing needs of students and faculty.** By leveraging the cloud, organizations can become more agile, be better able to address their capital funding challenges and provide users with digital access to needed services.
2. **Educational organizations must focus on modernizing and optimizing their IT infrastructure as part of this cloud transformation.** With more and more workloads going digital, infrastructure will need more capacity and increased performance to keep up. This will enable schools to optimize their IT costs while supporting these new digital tools and systems.
3. **As security threats continue to grow, educational institutions must double down on their efforts to protect their systems.** To accomplish this, they need to secure all layers of their cloud stack. There is no single solution to ensure proper protection for organizations. Instead, they must make security a core tenet of everything they do in the cloud.

The implementation of a cloud operating model promises to help the teams of engineers who provision, run, and manage cloud infrastructure as they drive their organizations' cloud transformation. Using the right tools enables these teams to create and operate highly automated platforms available on demand across the organization.

Developers can access these platform capabilities via self-service processes, making it easy for them to quickly create new environments and new service instances. IT teams have the responsibility to keep the platform stable, resilient, performant, and secure. Critically, bolstered by a reliable baseline of services, application development teams across the organization can use these platforms to create and release new capabilities to constituents faster.

To understand how this works, let's look through the four layers of the cloud operating model to see how it supports core educational institutions challenges and initiatives.

Transitioning to the cloud

The transition to cloud environments is a generational shift for IT departments. The transition means moving from largely dedicated servers in a private datacenter to a pool of compute capacity available on demand.

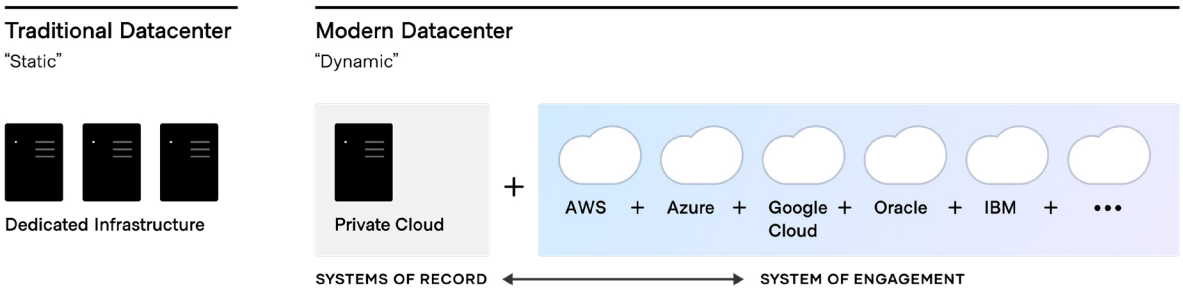

The cloud presents an opportunity to optimize the speed and scale of the new **systems of engagement** — the applications built to engage users. These new apps are the primary interface for constituents to engage with state and local organizations, and are ideally suited for delivery in the cloud as they tend to:

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

For most organizations though, these systems of engagement must also connect to existing **systems of record** — the core databases and internal applications, which often continue to reside on infrastructure in existing datacenters. As a result, many organizations end up with a hybrid — a mix of multiple public and private cloud environments.

The challenge for most educational organizations is to consistently deliver these applications to the cloud with the least possible friction across the various development teams. Compounding this challenge, the underlying workflows have changed from manipulating virtual machines in a self-contained environment to manipulating cloud resources in a shared environment.

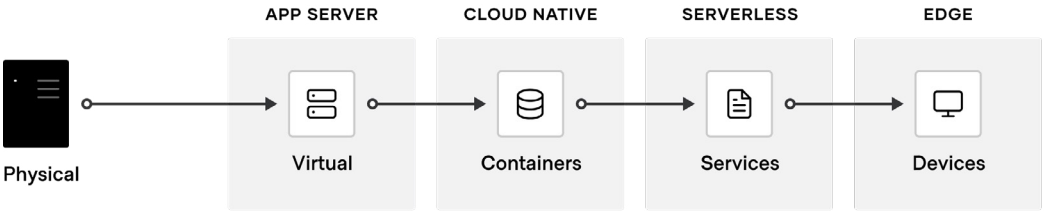

For cloud computing to work, organizations need a platform of consistent workflows that can be reused at scale throughout the institution. This requires:

- Consistent instruction sets for provisioning
- Identity for security and network connections
- Privileges and rights so applications can be deployed and run

Treat your platform as a product

Ideally, educational institutions should build and run their cloud platform as a product, a key principle of user-centered design. When building a cloud platform "product," the goal is to understand the needs of the organization and the teams that are building services to run atop the platform. A platform should bring demonstrable value to promote adoption and success.

Building a platform as a product is unlike running a traditional IT project. Rather than discrete projects with start and end dates, platforms require iterative product development. IT teams should be constantly ingesting feedback from organization and technical stakeholders, and regularly shipping new features and improvements based on this feedback.

Standardize workflows with a cloud operating model

A cloud operating model impacts teams across all layers of the stack. Educational organizations in the earlier stages of cloud migration and transformation must focus on three key layers: infrastructure, security, and networking. The greater cloud maturity an organization has achieved, the faster its velocity. That's why it is important to develop a cloud platform to deliver the dynamic services necessary at each layer to facilitate rapid infrastructure deployment, a strong security posture, and greater operational efficiency.

Based on what HashiCorp has seen at successful state and local organizations, there are several best practices that easily translate for educational institutions to leverage for adopting a cloud operating model and applying it to their platform at the infrastructure, security, and networking layers:

Standardize infrastructure provisioning with HashiCorp Terraform

The foundation for running a cloud platform is infrastructure provisioning. HashiCorp Terraform is the world's most popular cloud provisioning product. Terraform is used to provision infrastructure for any service using an array of providers for any public cloud or popular software application.

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

Before Terraform

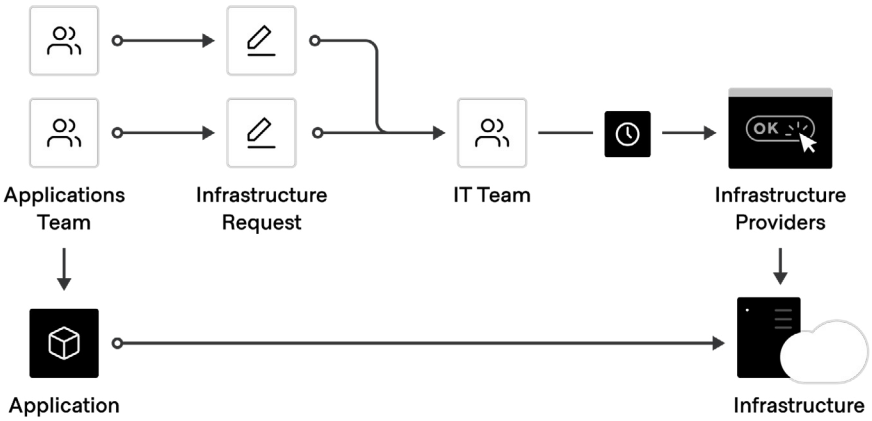

After Terraform

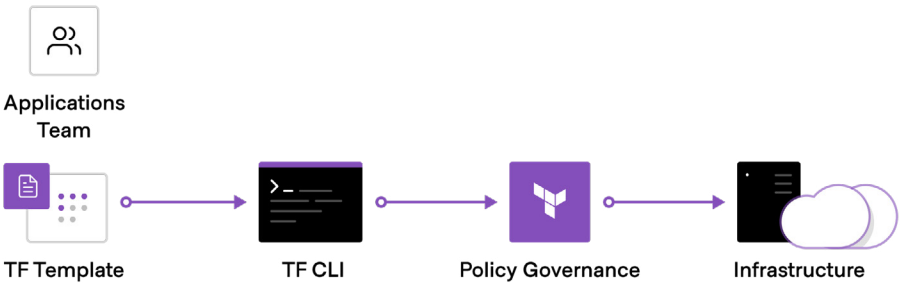

Without an automated provisioning workflow, requests for resources must be reviewed and approved manually. With HashiCorp Terraform, IT teams can automatically route all requests through pre-approved templates to deliver on-demand access to infrastructure for application teams.

Reproducible Infrastructure as Code

The first goal of a shared service for infrastructure provisioning is to enable the delivery of reproducible infrastructure as code, giving teams a way to plan and provision resources inside CI/CD workflows using familiar tools. Ideally, this provisioning just works, and is abstracted away from development teams.

IT teams can create Terraform modules, which act as templates that express the configuration of services from one or more cloud platforms and on-premises environments. 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 software vendors to include monitoring agents, application performance monitoring (APM) systems, security tooling, DNS, databases, 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 to help the platform scale and extend its capabilities.

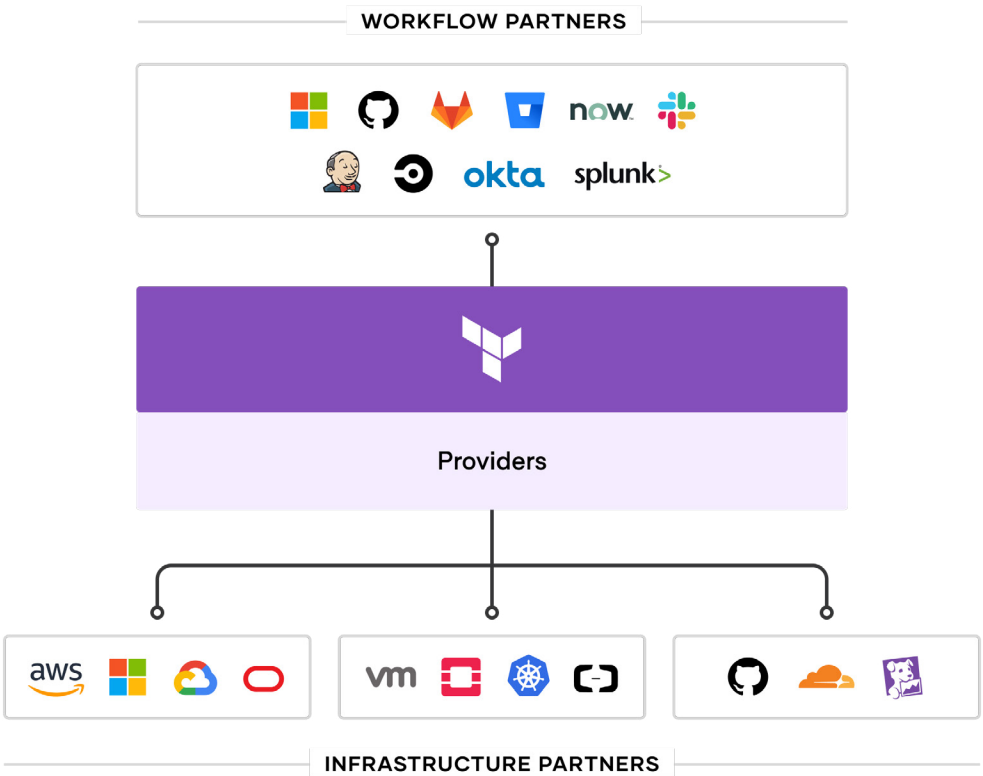

The cloud ecosystem has standardized on HashiCorp Terraform: more than 2,000 Terraform providers help teams adopt consistent provisioning workflows

For teams intent on delivering self-service infrastructure, the decoupling of the template creation process from the provisioning process reduces how long it takes for an application to go live since developers can use a pre-approved template instead of waiting for operations approval.

Compliance and management

Most IT teams also need to enforce policies on the type of infrastructure created, how it is used, and which teams get to use it. This is especially true for educational institutions that must conform to complex privacy rules and regulations. HashiCorp's Sentinel policy as code framework provides compliance and governance without requiring a shift in the overall team workflow. Sentinel is defined as code too, making its policies easy to share and understand as part of the larger platform.

Without policy as code, organizations resort to using a ticket-based review process to approve changes. This can become a bottleneck, making developers wait weeks or longer to provision infrastructure. Policy as code lets state and local organizations solve this by abstracting the definition of the policy from the execution of the policy.

IT teams should codify policies enforcing security, compliance, and operational best practices across all service provisioning. A shift-left approach, which requires code testing and validation early in the development process, automates enforcement, assuring that changes are in regulatory compliance without creating a manual review bottleneck.

Build automated images with HashiCorp Packer

HashiCorp Packer is an open source tool that enables IT teams to create identical machine images for multiple clouds from a single source template. A common use case is creating golden images that teams use to help standardize cloud infrastructure.

Packer automates the creation of any type of machine image, including Docker images and images for use with cloud service providers. Often, images created with Packer are inputs that start a provisioning workflow with Terraform.

Educational organizations IT teams that desire more automation can use the HCP Packer registry. HCP Packer is a cloud service that bridges the gap between Packer's image factories and Terraform's image deployments. This toolchain lets security and IT teams work together to create, manage, and consume images in a centralized way.

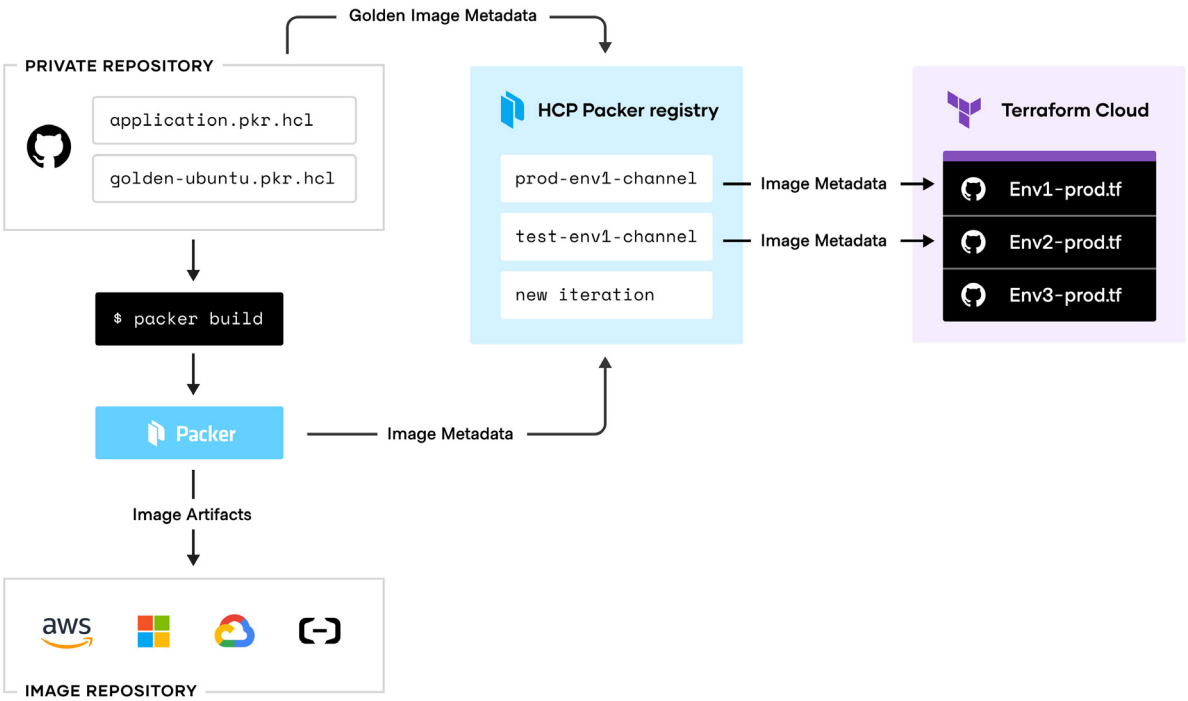

Manage secrets and protect data with HashiCorp Vault

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

In the traditional security world, we assumed high trust within 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 both humans and applications be explicitly authenticated, specifically authorized to fetch secrets and perform sensitive operations, and tightly audited.

HashiCorp Vault enables IT teams to securely store and tightly control access to tokens, passwords, certificates, and encryption keys. This provides a comprehensive secrets management solution for machines and applications. 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 also can act as a certificate authority to provide dynamic, short-lived certificates to secure communications with Secure Sockets Layer (SSL)/Transport Layer Security (TLS). Lastly, Vault enables a brokering of identity between different platforms, such as Azure Active Directory (AAD) and AWS Identity and Access Management (AWS IAM), to allow applications to work across platform boundaries.

To achieve shared services for security, IT teams should enable centralized secrets management services, and then use that foundation to address more advanced encryption as a service use cases such as certificate and key rotations, and encryption of data in transit and at rest. This embeds security considerations within the platform, so development teams need only plug into the provided APIs to ensure their service meets organizational security standards.

Before Vault

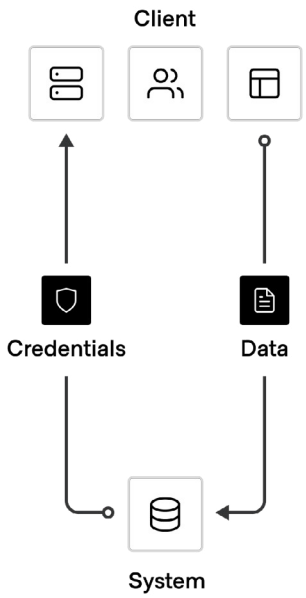

After Vault

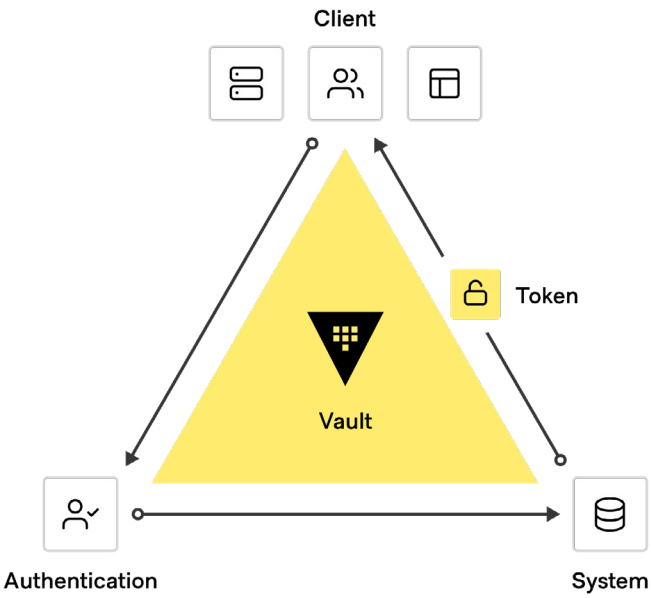

Traditional high-trust deployments do not have a programmatic way to protect passwords and other sensitive information. With HashiCorp Vault, teams implement an automated workflow for both people and machines to centrally manage access to credentials.

Secrets management

The first step in cloud security is secrets management: the centralized storage, access control, and distribution of dynamic secrets. Instead of depending on static IP addresses, it's vital to integrate with identity-based access systems such as AAD and AWS IAM to authenticate and access services and resources.

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 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.

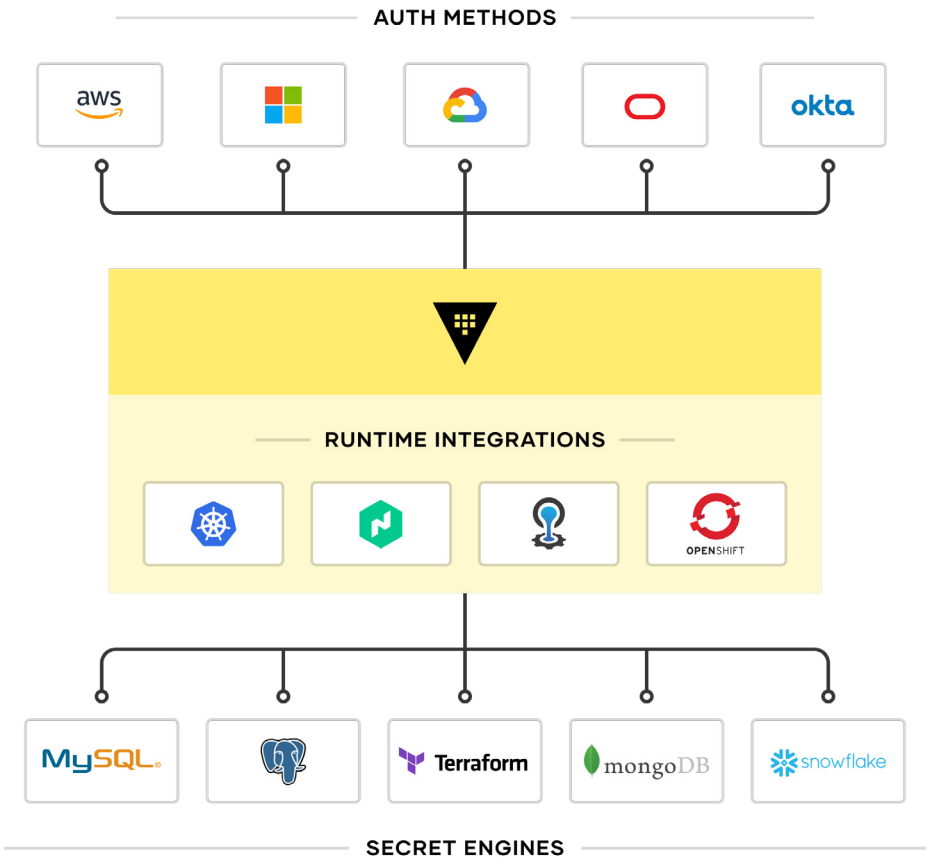

HashiCorp Vault works with common sources of identity to be a trusted identity broker at scale

Encryption as a service

Educational organizations 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 IT teams to perform a single integration and then protect data across multiple environments.

While many organizations provide a mandate for developers to encrypt data, they don't often supply the "how," which leaves developers to build custom solutions without an adequate understanding of cryptography. Educational institutions IT teams use Vault to give developers a simple API, while adjacent security teams can use policy controls and lifecycle management APIs as needed.

Secure and manage access with HashiCorp Boundary

The modern security principle of identity extends beyond the machine-to-machine access workflows addressed by Vault. Identity is equally central to securing human-to-machine interactions.

Traditional solutions for safeguarding user access require distributing and managing SSH keys, VPN credentials, and bastion hosts. This approach creates risks around credential sprawl and requires substantial manual effort to maintain. For example, it's all too common for educational organizations to reduce toil by setting their SSH keys to never expire, which gives users indefinite access to entire networks and systems.

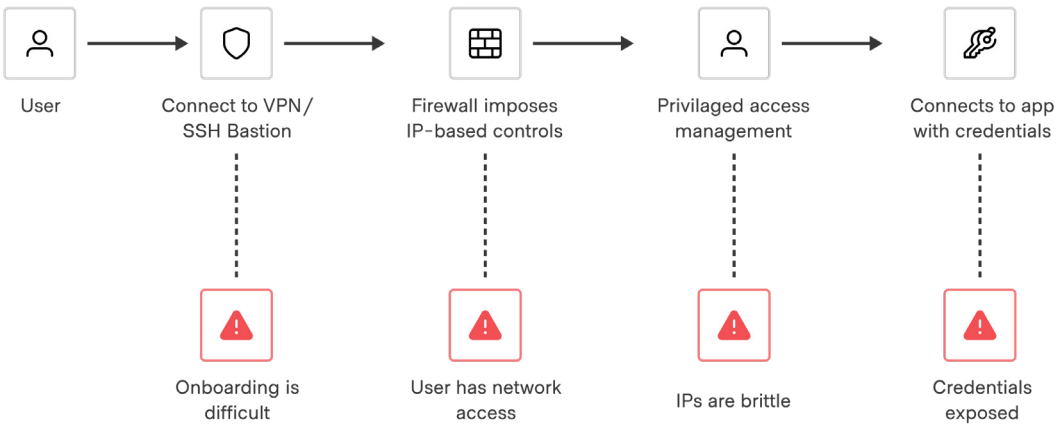

Traditional access workflows are manual and introduce multiple points of vulnerability to the system

Consider this common scenario: a developer trying to access a production app, such as a constituent payment portal, to troubleshoot an issue. In a traditional access model, the app is likely deployed to a private network (or VPC) on a set of well-known virtual IP addresses. Access is configured at the IP level.

This philosophy breaks down in the world of cloud infrastructure, due to the nature of ephemeral cloud resources and dynamic IP addresses. A modern approach for human-to-machine access relies on verifying the identity of the user. Further, access and provisioning of permissions to the end system is automated.

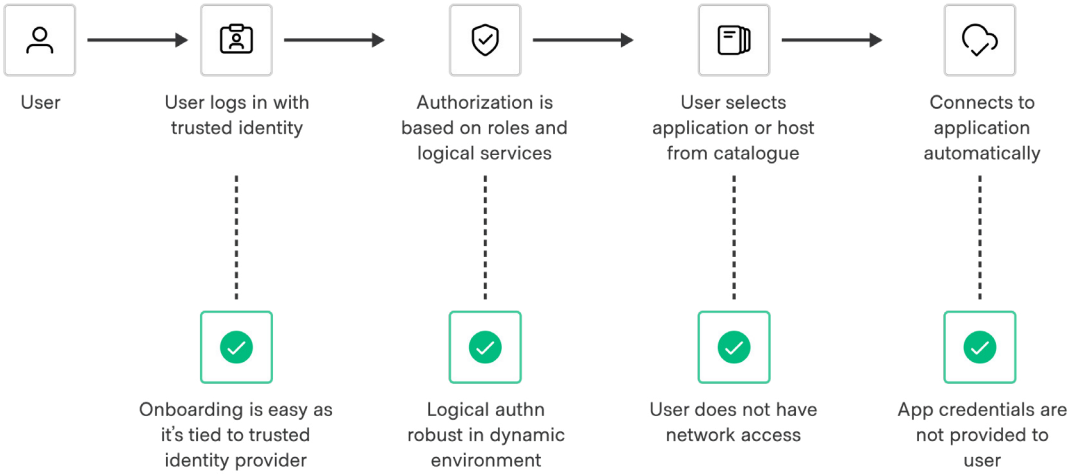

Modern access workflows are automated and use identity as their foundation

HashiCorp Boundary solves modern access challenges in a cloud operating model. Boundary is a secure remote access solution that provides an easy way to safeguard access to applications and critical systems with fine-grained authorizations based on trusted identities. It governs access across clouds, local datacenters, and low-trust networks, without exposing the underlying network.

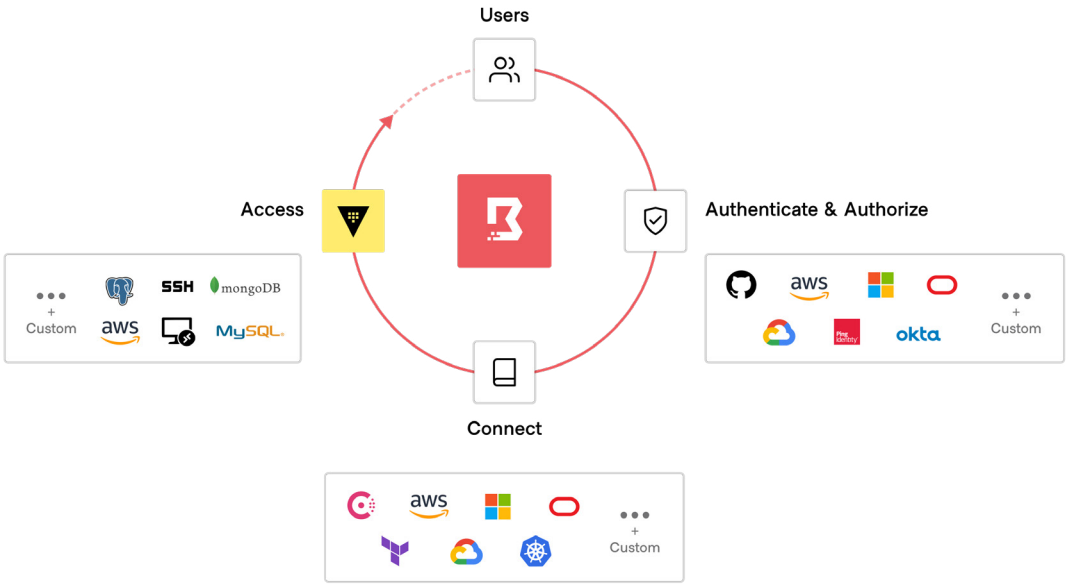

HashiCorp Boundary delivers simple and secure remote access to any system anywhere, based on user identity

Users can authenticate to Boundary using their identity provider of choice (AAD, Okta, AWS IAM, etc.). From there, each user can be tightly authorized to perform specific actions on a dynamic set of targets, and be granted just-in-time access to connect to those targets via credentials provided by Vault or another credential management solution.

This model fits the larger cloud operating model approach of shared services because it provides secure access to dynamic infrastructure with:

- **Identity-based access controls:** Engineers can streamline just-in-time access to privileged sessions (e.g. TCP, SSH, RDP) for users and applications. Teams can control access permissions with extensible role-based access controls.
- **Access automation:** IT teams can define the perimeter of resources, identities, and access controls as code through Boundary's fully instrumented Terraform provider, REST API, CLI, and SDK. They can automate the discovery of new resources and enforcement of existing policies as resources are provisioned.
- **Session visibility:** Security engineers can monitor and manage each privileged session established with Boundary. Session logs can be exported to a wide variety of analytics tools.

Securely connect applications with HashiCorp Consul

Networking in the cloud is one of the most difficult challenges when adopting a cloud operating model. Engineers must navigate dynamic IP addresses, account for significant growth in east-west traffic in microservices implementations, and adjust to the lack of a clear network perimeter.

HashiCorp Consul enables teams to manage multi-cloud networks by discovering and securely connecting services. Consul is widely deployed to run service networks in large-scale state and local organizations.

Networking services should be based on service identity and provided centrally, allowing teams to create a centralized service registry for discovery purposes. The common registry provides a map of what services are running, where they are, and their current health. 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. Service mesh approaches can simplify the network topology, especially in multi-cloud and multi-datacenter environments.

Before Consul

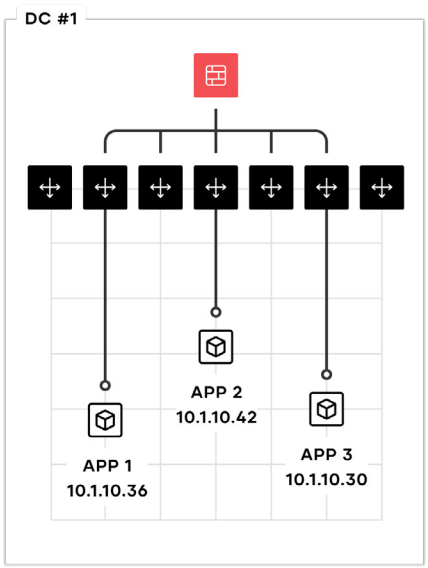

After Consul

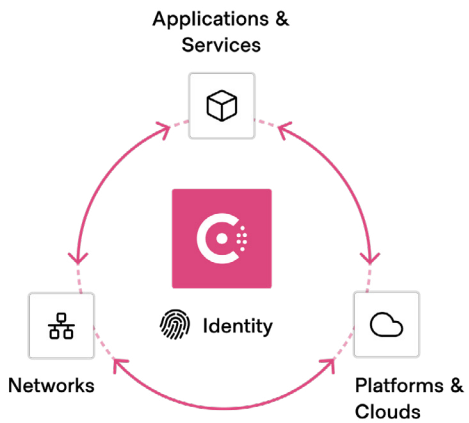

Historically, networking involved manual activities that connect hosts and static IP addresses. HashiCorp Consul gives modern cloud architectures and teams an automated, services-centric approach.

Service discovery

The starting point for networking in a cloud operating model is a common service registry, which provides a real-time directory of what services are running, where they are on the network, and their current health. Traditional approaches to networking assume a static IP for long-lived applications, using a combination of load balancers and firewalls to regulate communication. Tracking the network location of services often requires disjointed manual processes and tools such as spreadsheets, load balancer dashboards, or configuration files.

With Consul, each service is programmatically registered and DNS and API interfaces are provided to enable any service to be discovered by other services. Consul's integrated health check monitors each service instance's health status so the platform team can triage the availability of each instance and Consul can help avoid routing traffic to unhealthy service instances.

Network infrastructure automation

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

This configuration — in which Terraform executes configurations based on event changes detected by Consul — removes dependencies and obstacles for common tasks. State and local product teams can independently deploy applications while IT teams can rely on Consul to handle the downstream automation those product teams require. The benefits persist throughout the lifecycle of the service: automation can properly close firewall ports as services are retired.

The fast-track to cloud: People, processes, and tools

For educational organizations, the fastest path to value in modern cloud environments isn't just a matter of installing the right technology. It requires unleashing the power of a platform approach across three dimensions — people, processes, and tools. Each dimension requires its own shifts for optimal results:

People: Shift to site reliability engineering (SRE) practices

- Reuse expertise from internal datacenter management and single cloud vendors and apply them consistently in any environment
- Embrace DevSecOps and other agile practices to continuously deliver increasingly ephemeral and distributed systems
- Enable IT staff to automate with code and treat operations as a software problem

Processes: Shift to self-service

- Set up tools as an enabling shared service focused on application delivery velocity — shipping software ever more rapidly with minimal risk
- Establish centers of excellence for infrastructure, security, networking, and other functional areas for self-service delivery of capabilities
- Standardize on a service catalog with a set of common services for development teams to use and add new capabilities according to feedback

Tools: Shift to dynamic environments

- Use tools that support the increasing ephemerality and distribution of infrastructure and applications and that support critical workflows rather than being tied to specific technologies
- Provide policy and governance tooling to match the speed of delivery with compliance to manage risk in a self-service environment
- Embed security and compliance requirements within the platform itself to accelerate production deployments

Unlocking a cloud operating model

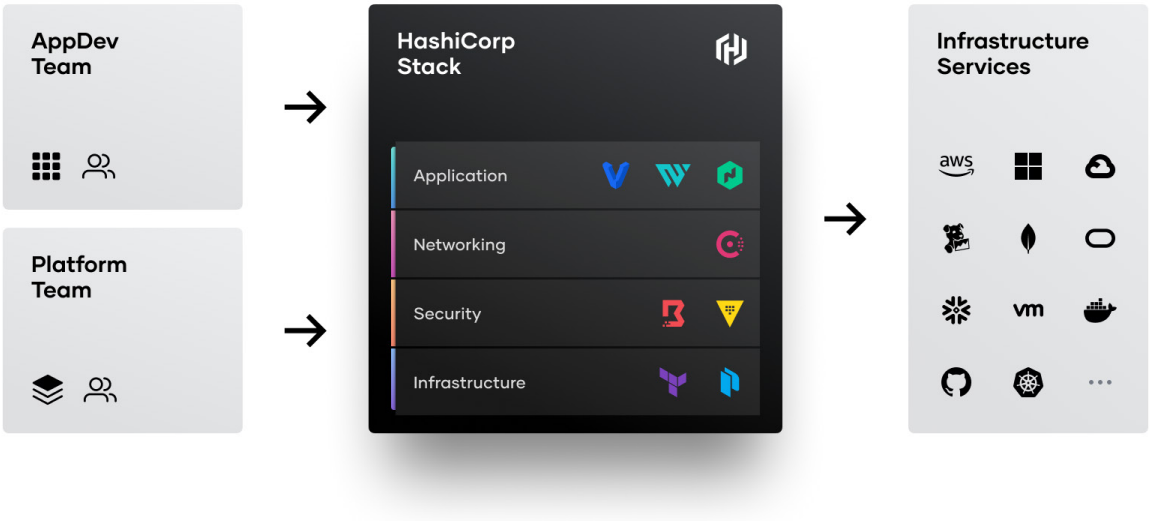

Put it all together, and it's clear that adopting a common cloud operating model is an inevitable shift for educational organizations aiming to maximize their digital transformation. IT is evolving away from ITIL-based control points — focused on cost optimization — toward becoming self-service enablers focused on speed. HashiCorp's suite of products provides solutions for each layer of the cloud infrastructure to enable educational institutions IT teams to successfully lead this shift to a cloud operating model.

By doing so, IT teams will be able to propel their digital transformation, shifting key infrastructure, applications, and services to the cloud and driving greater constituent experiences. In addition, by applying a platform mentality to this cloud operating model, they will create tools and processes that will optimize cloud spend and standardize organization cloud development, freeing up those resources to better serve the needs of their organizations.

About HashiCorp

HashiCorp is a leader in multi-cloud infrastructure automation software. The HashiCorp software suite enables organizations to adopt consistent workflows and create a system of record for automating the cloud: infrastructure provisioning, security, networking, and application deployment. HashiCorp's portfolio of products includes Vagrant™, Packer™, Terraform®, Vault™, Consul, Nomad™, Boundary, and Waypoint™. HashiCorp offers products as open source, enterprise, and as managed cloud services. The company is headquartered in San Francisco, though most of HashiCorp employees work remotely, strategically distributed around the globe. For more information, visit hashicorp.com or follow HashiCorp on Twitter @HashiCorp.

HashiCorp