HashiCorp

Anatomy of a breach:

A future post mortem of a cyberattack at your organization

Contents

Executive summary

Your organization has been hacked. It's up to your team to figure out what happened, how the bad actors were able to access your systems, and how to stop similar incidents from happening in the future.

This terrifying but all-too-common scenario forms the basis for this white paper, which walks through a hypothetical exercise held following a major cybersecurity breach at a fictional organization. The exercise leverages the well-known MITRE ATT&CK Framework — a globally accessible knowledge base of adversary tactics and techniques based on real-world observations — to understand the specific tactics used by the attackers, identifying vulnerabilities so that they can be quickly addressed to thwart future attacks.

> "This hypothetical post mortem was assembled from examples of actual security breaches."

Using the MITRE ATT&CK Framework, the post mortem identifies five key tactics used by the attackers and reviews how each one was used by the attackers (put in context of real-world breaches at actual organizations) and addresses how new processes and technologies can be implemented to close the vulnerabilities moving forward. The white paper reviews how the security vulnerabilities of the data breach can be remediated with different process and technology changes, including HashiCorp solutions purposefully designed for modern architectures. While the post mortem presented is fictional, it was assembled from examples of actual security breaches, and reviews documented tactics of how cyberattacks have been conducted against real organizations.

This approach helps illuminate the difficulties in understanding and remediating today's security vulnerabilities. As the world shifts toward cloud, multi-cloud, and hybrid architectures, traditional static and IP-based security, typically defined by the need to defend a perimeter around a datacenter, is becoming obsolete because there is no longer a clear perimeter to protect. Today's dynamic cloud-native infrastructures, characterized by an increasing reliance on remote, ephemeral access for humans, services, and devices, are increasingly publicly accessible and are likely being actively scanned by

adversaries at all times. Securing these environments requires enterprises to move beyond well-understood security techniques toward dynamic and identity-based models employing newer approaches like zero trust and least privileged access.

Even as the difficulty mounts, the stakes are higher than ever. Today's cyberattacks can affect core business operations, costing tens of millions in out-of-pocket expenses, lost revenue, damaging publicity, and even government penalties. Platform and security teams must continue to enable agile business processes and modern technology architectures, while ensuring the safety, security, and performance of systems and applications so that users and customers can still access the services they need.

Investigating the breach — a future post mortem

A breach of your organization has caused significant damage:

- The personal information and credit card data for your customers has been extracted.
- The source code for an upcoming product launch has been stolen.
- Your operational data lake has been encrypted by ransomware, shutting down all business activity.
- The impact on your business runs into millions and millions of dollars.

You've got a lot of questions to answer, starting with, "How did the attackers get in?"

This nightmare scenario is the cold, hard reality for thousands of companies. Sophos estimates that **94% of all organizations** faced a cyberattack in 2022, so you shouldn't be too surprised that it has just happened to yours.

But now you've got a lot of questions to answer, starting with, "How did the attackers get in?" You are part of a task force of platform operations, security, and engineering team members that has been assembled to quickly figure out how these attackers accessed your systems, moved around the

environment, and stole sensitive data, so that you can quickly address the security vulnerabilities. The team's goal is to understand what happened in order to implement new processes and technologies to close any identified vulnerabilities from future attacks.

The last thing you want is a repeat of what has just happened, so it's critical to understand what the attackers did and how technological and architectural changes affect your options for dealing with the vulnerabilities you uncover.

Using the MITRE ATT&CK Framework

Your team has been combing through logs and audit trails to reconstruct the different tactics the attackers used to access your systems. To organize your work, your team has adopted the widely used MITRE ATT&CK Framework to classify the different tactics and techniques that the attackers used, so that you can better understand how to address these vulnerabilities moving forward.

The MITRE ATT&CK Framework is a collected knowledge base of cyberattack tactics and techniques based on real-world observations. It is an abbreviation for Adversarial Tactics, Techniques, and Common Knowledge, and it is helpful for identifying not only ***what*** attack actions were taken, but also the ***why*** of the underlying motivation in what the attacker is trying to accomplish.

This approach helps to outline specific mitigation tactics and also to think about larger policies and infrastructure design strategies to further harden your environment moving forward.

Mitigating vulnerabilities going forward

While your team's first priority is to quickly close any vulnerabilities that are still open, your second priority is to do so in a way that will continue to hold up in the future. Your company has been growing quickly and your technology strategy has been evolving, making this crisis an opportunity to build more durable security measures.

"The post mortem has revealed your system technologies and design have changed, but that your security technologies and procedures have not kept up."

The post mortem so far has revealed that your security technologies and procedures have not kept up with the evolution of your system technologies and design. New technology approaches that bring new security challenges include:

- **Cloud-native architectures:** Systems design has evolved to be based on public multi-cloud containerized architectures, with more agile DevOps processes built around them.
- **Ephemeral infrastructure:** Modern infrastructure and systems are provisioned and decommissioned on demand to quickly meet technical and business needs.
- **Dynamic workloads:** Today, workloads quickly shift around different regions, availability zones, and public cloud vendors to meet different requirements.

Developing a modern security posture for cloud-native architectures

As system architectures evolve, security models need to change to keep up. In the past, operating datacenters provided a more enclosed environment, where you could use IP addresses for your endpoints and users were on LDAP, Active Directory, or similar solutions. This castle-and-moat approach provided a reasonable security design for the static and persistent needs of these environments.

Modern datacenters and cloud infrastructure require an updated security strategy. Public clouds are by definition public, with IPs that are addressable from external entities at all times. These IP ranges are being perpetually scanned, providing a rich attack surface for adversaries. In this world of multi-cloud, microservices, ephemeral remote access, and rapidly changing environments, the concept of being inside or outside the firewall is no longer relevant. Every access request must be specifically authenticated and authorized, no matter where it's coming from.

> "If you have a misconfigured Amazon S3 bucket, you have less than 16 minutes before it's been discovered and you have an attacker on the network. We have to think differently today."
>
> — Michael Wood, HashiCorp Field CTO

The move towards zero trust security

Given these system architecture changes, teams and organizations are shifting towards a zero trust security approach for managing their applications and infrastructure. Modern public cloud architectures have made network boundaries more fluid. Relying on a fixed perimeter protected by a firewall, where everything outside is banned and everything inside is trusted, doesn't align to the design of today's cloud native, ephemeral, and dynamic system architectures.

If there is no longer a clear boundary of trust, then a different approach is needed to authenticate and authorize machine and human interactions. The zero trust security model trades static and IP based security in the datacenter for dynamic identity-based security in the cloud. This identity-based approach ensures that every activity is managed regardless of its source, across multiple datacenters and different public clouds, while still ensuring a consistent standard for authentication, authorization, and access.

In our hypothetical post mortem, the security team is keeping these structural technology and security changes in mind for the future, while reconciling them with closing the security vulnerabilities of today.

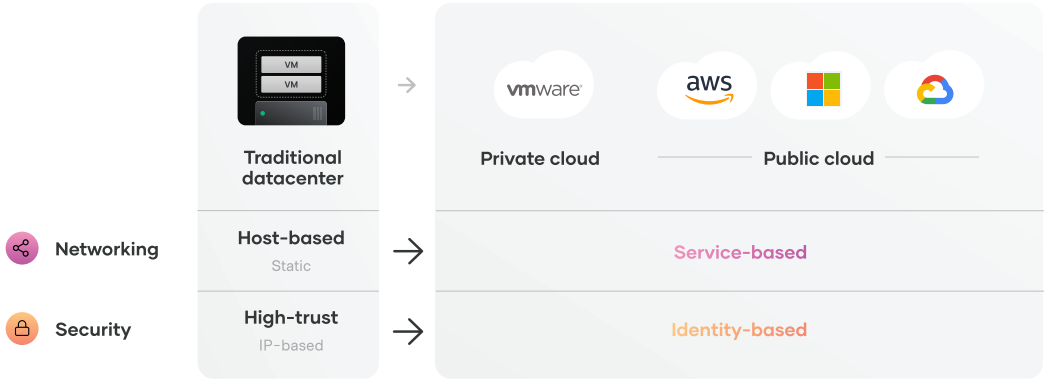

Security approaches must evolve along with system architecture.

The post mortem

Using the MITRE ATT&CK Framework, your team has outlined five key tactics that the attackers used to move through your systems, studied how to address vulnerabilities given the state of the business, and evaluated how to address those risks moving forward.

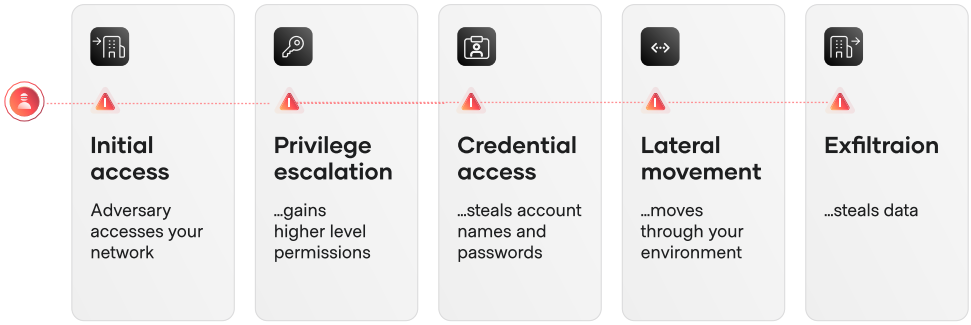

Anatomy of a cyberattack (MITRE ATT&CK Framework)

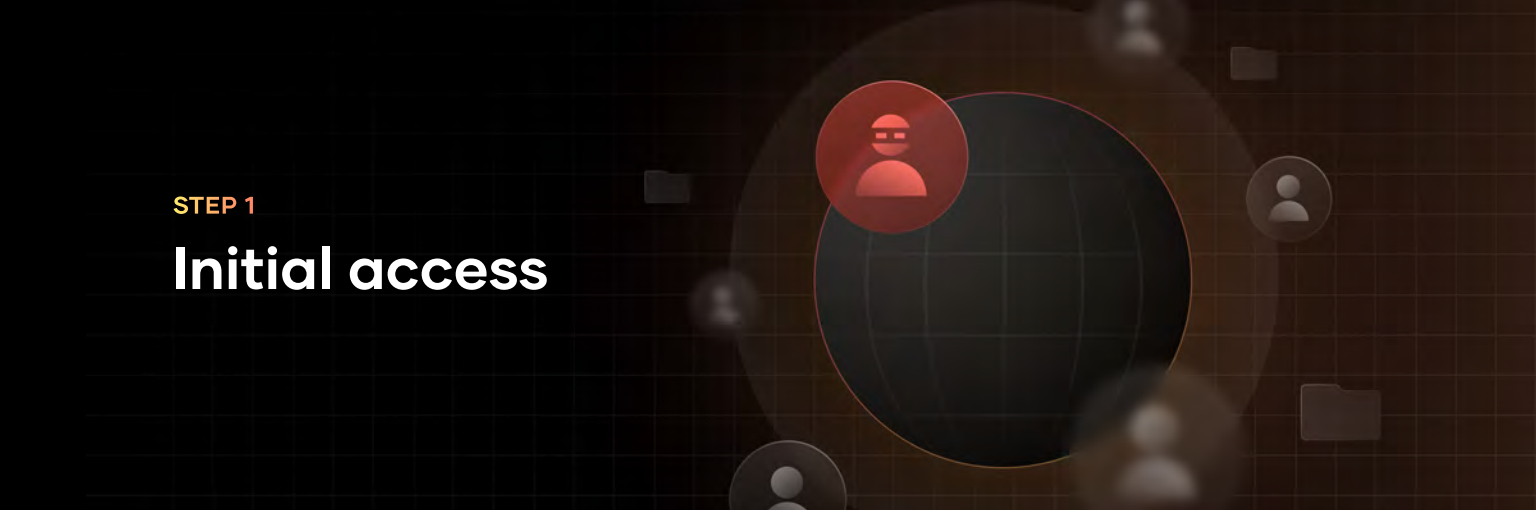

STEP 1

Initial access

The tactic

The adversary is trying to get into your network. Techniques include phishing emails, exploiting public-facing applications, and drive-by compromise.

Examples of initial access being used in attacks

At RSA, malicious phishing emails were sent to mid-level employees that backdoored employee computers.

At Sony Pictures, attackers impersonated top-level employees and used fake Apple ID verification login emails to find passwords also used on the Sony network.

How it was used against us

Keys and passwords were stolen and a set of hard-coded secrets were accidentally leaked into a public code repo as part of a training session with some contractors. Adversaries used this to gain initial entry into a sandbox system that could connect to the company network.

Understanding the problem

Our post mortem revealed that secrets sprawl is a problem at the organization, and secrets are scattered through multiple cloud systems, such as GitHub repositories and AWS Configuration Manager. Some secret storage services are used by individual teams, which creates issues with tracking the overall secret lifecycle. These services store them securely and encrypt them at rest but do not support secret versioning, regular secret rotation, or overall management. There is no programmatic way to scan for hard-coded secrets throughout our environment.

Recommendations

A more centralized approach is needed to manage the lifecycle of secrets and enable secure access:

- Scan and find secrets that may be incorrectly stored or publicly exposed and quickly revoke access.
- Provision and use just-in-time or single-use credentials to prevent or mitigate the risk of exposed credentials.
- Enable **least privileged** access by default for more granular scoping of access.

Security vulnerability remediation with HashiCorp

A new approach to managing secrets and credentials is needed to address the insufficient controls we have over secrets management. Instead of static, persistent credentialing that is more vulnerable to theft, HashiCorp Vault enables just-in-time dynamic credentialing. Credentials can be time-bound and finely scoped to the job, and can be easily changed with different jobs; an analog to the ephemeral and dynamic infrastructure of modern systems.

Additionally, **HCP Vault Radar** can scan and find secrets that have been accidentally exposed throughout our network, such as in source code, Git repositories, and collaboration tools, so that they can be rotated or stored in Vault. A data reference of managed secrets within Vault can be correlated against unmanaged secrets discovered in your scan, so you can remediate appropriately. In this case, you would immediately remove the exposed secret outside of Vault and rotate the corresponding secret being managed in Vault.

Identity-based security controls can add an additional layer of security. Users are able to authenticate with our existing identity service provider for single sign-on access, and then automatically access any of their authorized systems with HashiCorp Boundary. Users no longer need to see, store, or copy/paste multiple keys and passwords in order to access various systems, which lowers the risk of the secret being leaked. Instead, single-use dynamic credentials can be invisibly injected into user sessions by Boundary, making for simpler, more secure, password-less access.

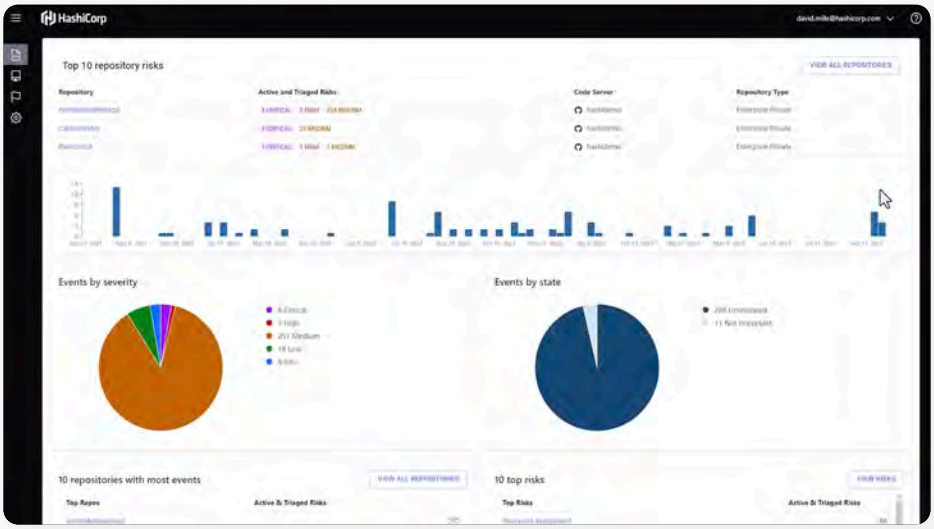

HCP Vault Radar scans for secrets in many places, including code repositories.

STEP 2

Privilege escalation

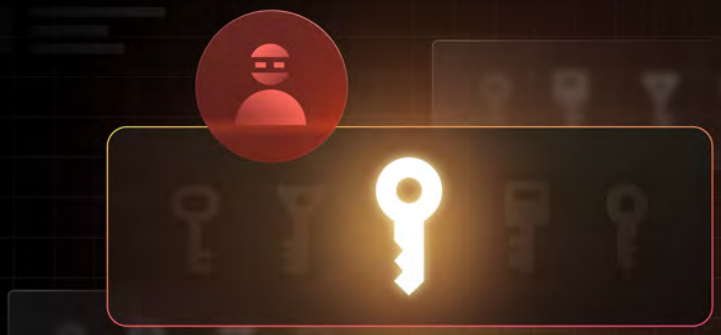

The tactic

The adversary is trying to gain higher-level permissions. Techniques include exploiting valid accounts, process injection, and access token manipulation.

Examples of privilege escalation being used in attacks

In the Solarwinds attack, compromised Solarwinds Orion software was used to harvest credentials to other Tier 0 systems like Active Directory, infrastructure systems, and even SAML systems.

The Shellshock privilege escalation vulnerability in bash on Linux compromised millions of systems and led to millions of DDoS attacks being held per day by botnets of compromised systems.

How it was used against us

From their initial access to our sandbox, attackers used this access point to dig through different credentials of contractors who were in our system. They were able to create a new credential with more privileged access rights to access deeper parts of our network.

Understanding the problem

Analysis of our security practices reveals that many of our users have greater privileges than they need, particularly across cloud infrastructure. Auditing our user accounts shows that many users have access to the company's **root of trust** when they should not. Those permissive accounts and policies mean that once an attacker is able to access that initial root of trust (such as a key management system, KMS, or hardware security module, HSM) via a compromised user account, they are able to create additional secrets with higher levels of access to network resources, databases, and other key assets beyond the initial credential, and even bypass our traditional privileged access management (PAM) system via over-escalated privileges. Managing this becomes even more complex across multiple cloud providers and system infrastructures.

Recommendations

Following the principle of least privilege access, increase permissions granularity and decrease the lifespan for credentials without impacting user productivity:

- Right-size permissions for the specific job at hand, without extending overly broad, unnecessary privileges.
- Move away from static, persistent credentials to just-in-time, short-lived issuance of credentials to shrink the attack surface area.
- Simplify and offload credentialing to a more centralized management system for more consistent control.

Security vulnerability remediation with HashiCorp

The breach reveals that we need to change our approach to how credentials are administered. We need a centralized solution to manage this process, since manually configuring access policies is too time-consuming and complex to be managed at scale, particularly across multiple clouds and infrastructures.

HashiCorp Vault provides ephemeral, on-demand dynamic secrets creation and revocation to safeguard against credential leakage. It functions as a centralized identity broker, secrets management platform, and encryption service across an enterprise. For secure user access, HashiCorp Boundary evaluates permissions every time a user logs into the system and helps right-size privileged access at scale.

With Boundary in place, privileged access becomes more secure and scalable than a VPN or bastion host approach. It reduces the attack surface by using workers as the entry point into a network. And Vault limits secret sprawl by keeping secrets in a centralized location.

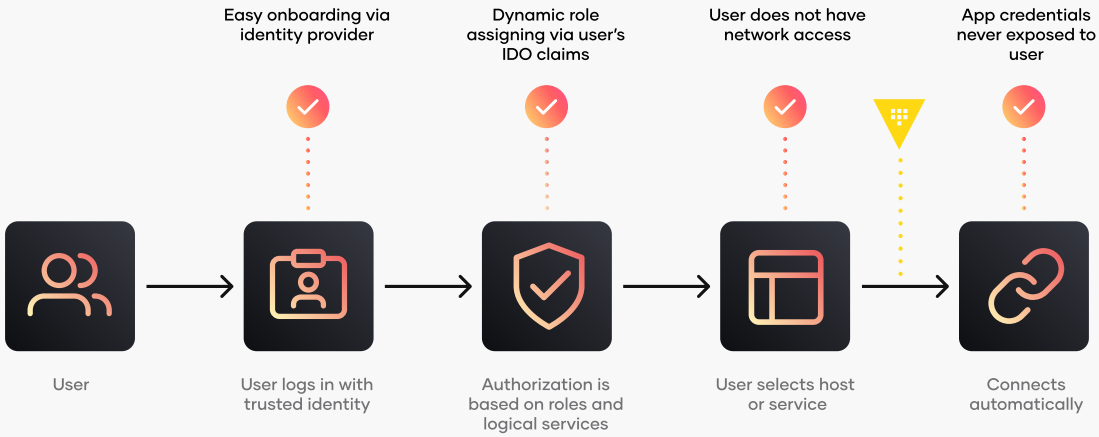

Simplifying authentication and authorization with Hashicorp Boundary and Vault while enhancing security

STEP 3

Credential access

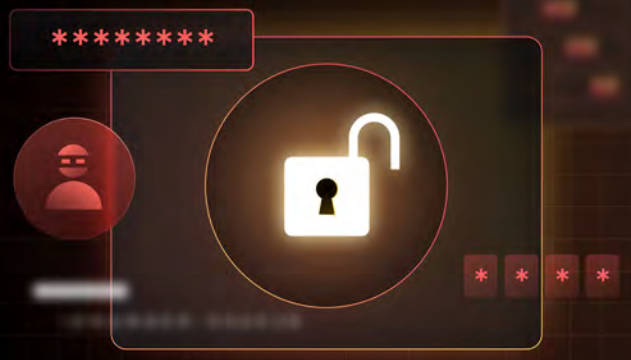

The tactic

The adversary is trying to steal account names and passwords. Techniques include harvesting credentials in files, OS credential dumping, and keylogging.

Examples of credential access being used in attacks

At CircleCI, malware installed on an employee laptop was used to steal a valid 2FA-backed SSO session, from which the employee's production access was used to extract production encryption keys from a running process.

The worldwide WannaCry ransomware attack leveraged credential dumping through the commonly used Mimikatz tool in order to harvest credentials from compromised machines and spread further into networks.

How it was used against us

With a set of credentials that allowed attackers deeper access into our systems, they found credentials on a shared network drive, which they used to create another set of credentials with admin access and gained root access.

Understanding the problem

The increasing complexity and scale of systems leads to significant sprawl of credential sets, which increases our attack surface. Credentials are stored in too many systems, further increasing complexity and friction. Given business pressure to move fast to keep up with market demands, users respond to these forces by storing credentials insecurely for ease of access, leaving the organization vulnerable.

Recommendations:

Disincentivize users from taking high-risk security shortcuts such as leaving credentials exposed, by making the right behavior easier:

- Automate authentication and authorization into systems for users through our existing identity provider (IdP).
- Shift towards ephemeral and scoped credentials, replacing static and broad credentials to shrink the attack window.
- Deploy a centralized and standardized way to manage keys across multiple cloud providers and systems to promote consistency and prevent secret leakage.

Security vulnerability remediation with HashiCorp:

Users can authenticate into our existing identity providers (such as Okta, Kubernetes, and Active Directory) to gain access to HashiCorp Vault and access the systems and secrets they need from there. This integrates well into existing IdP workflows with which users are already familiar.

Vault enables us to securely store and manage credentials, keys, and other secrets without having to employ fragmented and complex key-management practices on additional infrastructure. Vault allows us to automate generating, rotating, and expiring secrets to limit the long-term exposure of static secrets.

Finally, **Vault secrets sync** provides an API and standardized workflow for the distribution and lifecycle management of keys in our multi-cloud environment across various KMS providers. Instead of having to manage keys in each cloud provider's KMS service, Vault secrets sync provides a unified control point for all our encryption keys across all our clouds. They can be centrally generated on demand and easily rotated, while also leveraging the unique cryptographic capabilities native to different KMS

providers, including AWS KMS, Microsoft Azure Key Vault, and Google Cloud KMS. Vault also provides a secure, encrypted way to natively retrieve and sync secrets in our containerized Kubernetes systems, allowing each service to uniquely authenticate and request their own credentials.

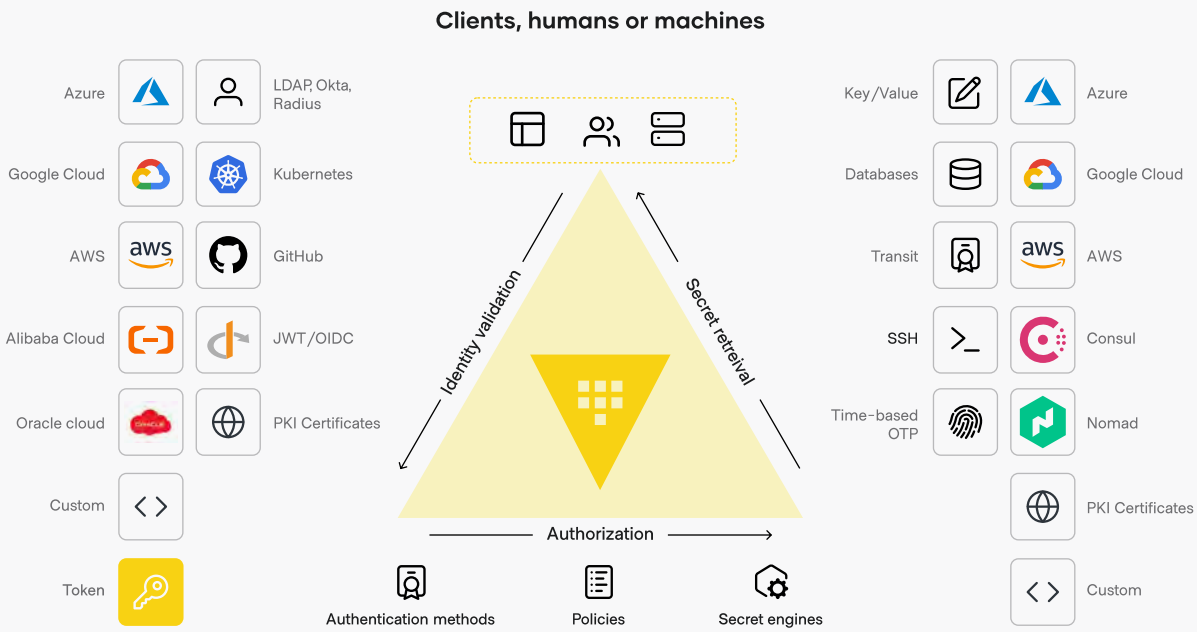

HashiCorp Vault centrally manages secrets across various systems and clouds.

Additionally, HashiCorp Boundary can be leveraged to implement password-less access to systems by injecting single-use, dynamic credentials into user sessions without exposing the credentials to end users, reducing the risk that they might be lost or stolen.

STEP 4

Lateral movement

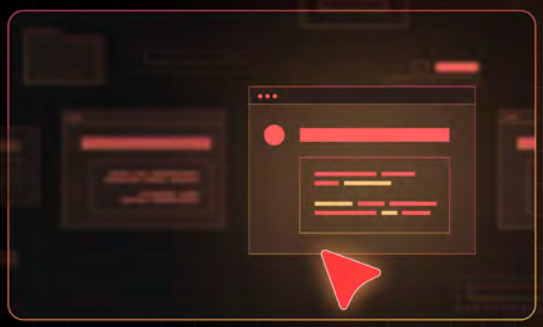

The tactic

The adversary is trying to move through your environment. Techniques include pass the hash, remote desktop protocol, and Windows admin shares.

Examples of lateral movement being used in attacks

At Target, attackers used HVAC contractor credentials to move from an external portal system onto the company network.

Marriott announced hundreds of millions of customer records like credit cards and passports were exfiltrated after a multi-year breach at Starwood, even before Marriott's acquisition of the Starwood business.

How it was used against us

Once they had further entry into our network, the attackers began to move around the environment. They scanned the network, likely looking for high-value systems to exploit. They also accessed more systems, employing techniques like remote desktop connections to infiltrate more servers and to install ransomware.

Understanding the problem

Our castle-and-moat security posture is permissive once entry into our network is made. Any given system or user inside the network can freely see all the other systems within the network. Even a single compromised system can give attackers a significant toehold to access and compromise other systems.

Recommendations

Adopt zero trust security principles to increase access granularity and scoping for systems. This should help prevent movement across systems:

- Establish network rules that scope services and applications to be able to access only the specific, defined endpoints and not the whole network, preventing compromised services from affecting other services.
- Manage employee access so that they can't see (or access) everything on the network. Attackers that establish a root of trust outside the perimeter can breach the network, allowing a single compromised employee account to expose the whole system.
- Automating this system is necessary to manage at scale. The increasing complexity of modern environments means there are too many access points to track and secure manually.

Security vulnerability remediation with HashiCorp

Adopting a zero trust security posture for our infrastructure can significantly limit the blast radius of an attack. We can keep the network more secure by keeping users off the network altogether with HashiCorp Boundary. Users would connect to a HashiCorp Boundary worker inside the firewall, which then brokers a secure direct connection to the target system and can dynamically inject necessary secrets for the job. Users no longer need to access the network with a VPN, and no longer need to manage long-lived secrets that can potentially leak to bad actors. Boundary also limits the scope of systems a user is allowed to access.

HashiCorp Vault can also be deployed as a Certificate Authority (CA) to verify the identity of each service with with the appropriate authentication method prior to issuing TLS certificates. The certificate creation, signing, and renewal process can also be automated for better scaling and security. Vault's

central management allows credentials to be narrowly scoped on a given target system. For example, it can limit user access to one application process on the Linux system that they are using, and not allow access to other Linux processes on the host. This fine granularity of control helps keep compromised users or machines from moving laterally within a system.

Network services can be further hardened for least privilege access by denying all service-to-service communication unless explicitly permitted by HashiCorp Consul. Each specific application endpoint is uniquely identified, authenticated, and authorized, and the channels between them are encrypted with TLS certificates. Even if an attacker is able to access a weakness in a given application, they are limited in their ability to move around laterally, and the encryption prevents intercepted network communication from being viewed.

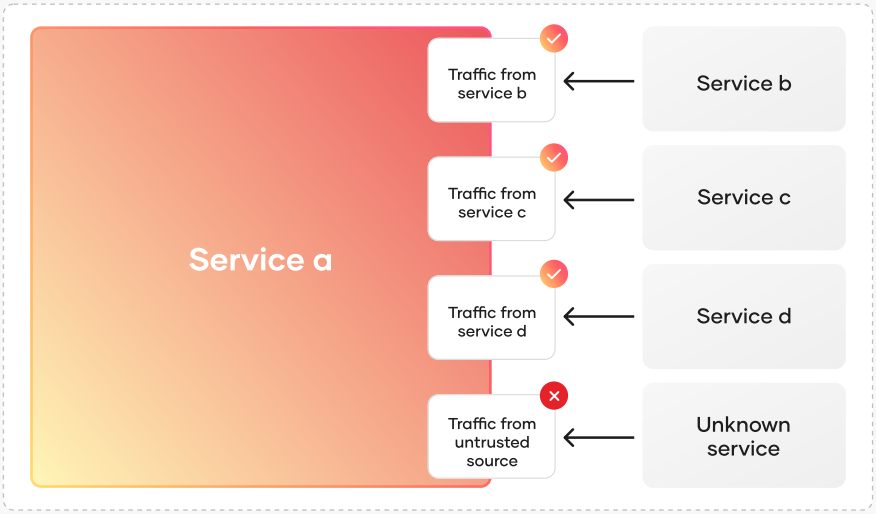

Service-to-service communication is denied unless explicitly permitted by HashiCorp Consul.

STEP 5

Exfiltration

The tactic

The adversary is trying to steal data. Techniques include moving data over existing command control channels, via alternate protocols, or through automated exfiltration.

Examples of exfiltration being used in attacks

During the Equifax data breach, the personal information of more than 145 million people — including full names, Social Security numbers, birth dates, and addresses — was stolen in small increments using standard encrypted web protocols to disguise the exchanges as normal network traffic.

In the Shields Healthcare Group data breach, impacted data from this possibly included name, diagnosis information, treatment formation, provider information, and insurance information.

How it was used against us

The attackers used their root access to create an account that was used to call out to our customer payment database and pull data which was then routed to an external system in small increments to evade detection.

Understanding the problem

Our security posture invests resources into encryption at rest as well as encryption in transit, including using an HSM for sophisticated encryption approaches that make it difficult to break encrypted data. However, the data still remains vulnerable if a trusted IP puts in a valid request, from which the HSM decrypts all the data and delivers it in plain text, which can then be easily stolen.

Recommendations

Align with encryption best practices, and also implement additional layers of security around how data can be accessed and decrypted:

- Implement additional protections around access authorization into a database, so that a trusted system that is compromised will not be able to directly extract sensitive data.
- Implement additional partition and scoping of data so that even internal applications or employees are not able to access the entirety of larger data stores.

Security vulnerability remediation with HashiCorp

Improvement to our encryption strategy is needed not only from a technology perspective, but also in terms of process. Leveraging HashiCorp Vault as an encryption services engine lets us encrypt data before it lands in the database. This means that having the user ID password for the database is not sufficient, users will need to be authenticated with Vault, and a Vault token needed in order to decrypt the data, making it more difficult for an attacker to exfiltrate data, as they must defeat two systems.

Vault provides additional granularity for the scoping and partitioning of data access. Large data lakes or databases may be shared by many different applications, and Vault can be used to centrally define different keys for different applications. A column or row of data can be encrypted differently with different scoping access. This granularity of control helps protect from attacks within the network or even by internal employees, and is critical in highly regulated industries such as finance, healthcare, government, and others.

Managing keys application-by-application can quickly become cumbersome, but this can be offloaded to Vault for centralized management, control, and auditing. HashiCorp Vault and Consul use encryption as a service (AES-GCM) with 256-bit AES to secure data during transit and at rest.

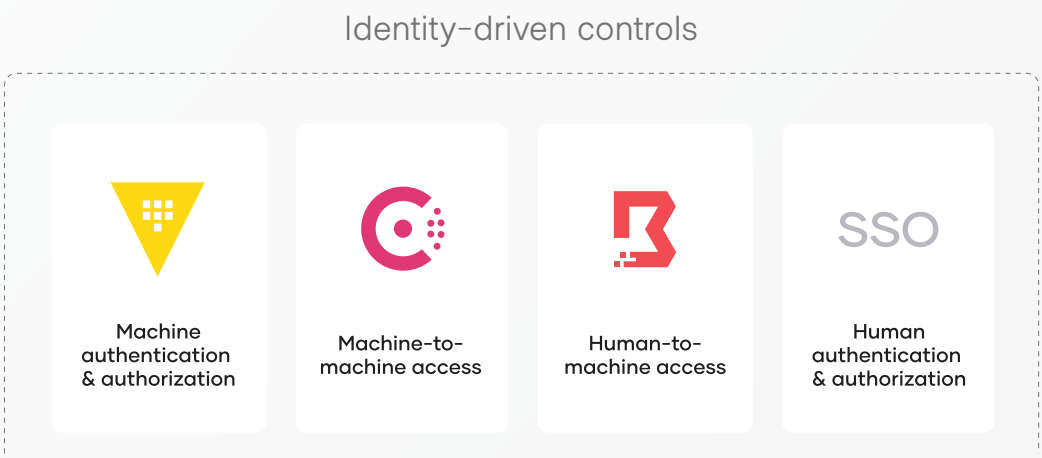

Identity-driven controls across machines as well as humans with HashiCorp help maintain data security.

Finally, Boundary session recording allows security admins to play back sessions, which is especially important after a breach has occurred. In addition to helping meet compliance requirements, session recordings also help expedite remediation by letting admins see exactly what commands and actions an adversary performed. Notably, session recordings can be useful in addressing all five steps of an attack.

The stakes have never been higher

The ramifications of a security breach are higher than ever. Major data breaches regularly cost victims millions of dollars to remediate and can significantly erode customer trust. And as organizations become increasingly reliant on their systems and infrastructure for core business operations, service interruptions caused by cyberattacks can be devastating, even if no data is lost. It can take months of work to remediate systems and deal with the effects of a data breach, especially while balancing the security of the organization against all of the other business critical projects and initiatives in the works.

> "As organizations become increasingly reliant on their infrastructure for core business operations, service interruptions caused by cyberattacks can be devastating even if no data is lost."

Putting yourselves in the shoes of a team reviewing a security breach can help bring home the impact of such incidents, and clarify where you need to improve your security posture. And as this exercise has shown, deploying tooling using identity-based authentication and authorization is the best way to reduce your exposure to these costly attacks.

New technologies and platform architectures are opening new possibilities, but also demanding new approaches to your security design and principles to match these new capabilities. Keeping teams, systems, and processes agile is essential for future growth and success while ensuring operations remain safe and secure.

HashiCorp can provide critical zero trust security and access solutions using identity-based controls to protect, inspect, and connect without slowing down your organization.

To learn more and get a demo of HashiCorp cybersecurity solutions, please contact the HashiCorp sales team: sales@hashicorp.com.

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 free community source-available products, enterprise products, and managed cloud services. The company is headquartered in San Francisco, though most HashiCorp employees work remotely, strategically distributed around the globe.
For more information visit hashicorp.com