Citihub Digital

The Secret to Good Secrets Management for the Enterprise and Cloud

Achieving Maturity in Secrets Management for the Enterprise and Cloud

October 2020

Abstract

What is Secrets Management? Why is it so important to the digital enterprise to have a rationalized Secrets Management strategy for the Enterprise and Cloud? In this paper we discuss key considerations and positioning for centralized management of secrets for humans and systems.

Organizations have been managing identities and privileged access management to infrastructure and applications for decades but are now finding themselves making new investments and strategic changes to the end-to-end lifecycle management of secrets and enactment of policies.

The dynamic nature of DevOps pipelines and the ephemeral environments favored for their efficient consumption of resources require system users and orchestration engines to perform entitled activities that were traditionally handled by human users, at an order of magnitude more frequently with response times measured in milliseconds.

More than ever, service transactions are occurring outside of 'the 4 walls' of any given enterprise and the traditional safeguards and risk mitigation strategies that are inherent in a trusted perimeter are no longer applicable. While for the business 'the show must go on', in the highly regulated environments of Financial Services, Insurance, Healthcare, and Governmental organizations there is little appetite to alter the acceptable risk tolerances around access management, secure operations, confidentiality and data privacy. Simply put, failure to secure resources and data will likely get you on the front page of the news, with severe legal, financial and reputational consequences.

Managing secrets in a decentralized model, without sufficient governance and monitoring, is highly vulnerable to over-privilege abuses, cyber-attacks and operational blindness. The continual stream of high-profile data breaches and ransomware attacks have only added urgency to the need for centralizing the strategy around Secrets Management for the Enterprise and Cloud.

In order to adequately assure, monitor and secure the inevitable proliferation of secrets and the implementation of management solutions across the Enterprise it is necessary to centralize, standardize and govern use and avoid the product sprawl that would otherwise complicate management, maintenance and ability to secure.

Contents

Author

Chris Zanelli
chris.zanelli@citihub.com

Introduction

The explosion of the digital economy serves as a boon to business and technology innovators; Enterprise Security professionals - who traditionally acted as the doorman for entry, exit and lateral movement in the organization - are now faced with the rapidly growing challenge of scaling security for a huge number of B2B, B2C and S2S[1] interactions across the organization, from critical business system flows all the way down to the compute infrastructure team offering self-service compute on bare metal and all the use cases in between.

While we could write a lengthy paper on building, managing, and asserting identities, the subject of this whitepaper is to discuss Secrets Management for Enterprise and Cloud.

SECRETS

Secrets exist to establish various permutations of system-to-system, human-to-human, and human-to-system communication patterns that exist across roles and zones, trusted and untrusted networks; examples of secrets that you have likely heard of include user passwords, one time auto-generated passwords, application keys/credentials, ssh keys, database passwords, system-to-system passwords, private certificates, private encryption keys, binary license files, hardware and software generated passwords (e.g. RSA).

[1] S2S = System to System

Aside from the challenges of maintaining security posture for all of these secrets in one place, there is a non-trivial lifecycle to manage for each and every secret. This lifecycle is often neglected by anyone outside of a Security or Risk discipline. Managing secrets at scale introduces significant challenges, including:

Topic	Typical Pain Points
Segregation & Least Privilege	• Segmented management of secrets across critical systems deployed in several trusted and untrusted zones, with a variety of environments spanning development, QA, UAT and production • Designing for the principles of least privilege
Secrets Generation & Distribution	• Acknowledgement and validation of new secrets requests • Applying repeatable and sustainable standards for the generation and maintenance of secrets across all use cases (human or otherwise) • Support for on-demand secrets generation • Enacting policies for secret quality
Leasing & Revocation	• Building and enforcing granular policies for secrets revocation - whether it be age-out, invalidation/removal, or termination based on types of users and roles • Management across one-time passwords, elimination of weak secrets and maintenance of standards for secrets expiry
Security of Secrets	• Enacting appropriate access and management controls for administrators, system users, trusted 3rd party systems and human users • Governance of controls both within individual vaults and across different vaults • End-to-end encryption of secrets encrypted - in transit and at rest • Eliminating plaintext exposure of secrets • Understanding when one or more vaults has a vulnerability • Operationalizing the detection, response and recovery to cyber risks
Scaling, High Availability & Redundancy	• Supporting higher frequency requests with greater geographic distribution for system-to-system transactions • Supporting just-in-time and dynamic secrets generation at scale • Providing both good user and developer experiences
Monitoring & Reporting	• Providing reports for both producers and receivers/users of secrets, their utilization characteristics, and administrative changes performed • Monitoring, alerting and reporting on anomalies and potential threats

How does your enterprise stack up against all these capabilities when it comes to maturity around Secrets Management for the Enterprise and Cloud? How is your organization faring with existing commercial or in-house solutions designed to tackle some of these enterprise problems? And most of all, is your enterprise's current solution stack positioned to support your business and technology demands, while improving security posture and avoiding high profile cyber incidents and attacks aimed at disrupting business services and stealing data.

Assessing Current State

Taking a cue from the infamous Deming Cycle of Continuous Improvement, **the best place to start is where you are now,** whether your enterprise is invested in commercial products - such as CyberArk, Hashicorp Vault, Beyond Trust, AWS Secrets Manager, Azure Key Vault, RedHat Vault - an in-house solution or any combination thereof.

As an Enterprise Identity and Access Management professional you have keen insight into your existing capabilities, the use cases you have adoption for and a general idea about existing strategy. You also have good awareness of the current and, to a degree, future needs of technology and business teams that will drive new secrets requirements.

If there are more questions than certainty, there has never been a better opportunity to get the attention of senior leadership on an area that's likely already being discussed at the board of directors' level.

Unless your enterprise grew up "cloud first" you probably have a mix of longstanding and short-term implementation(s) that cover Privileged Access Management (for both humans and applications), Break Glass credentials and other pockets of Secure Credentials Vaults for your core business and critical infrastructure systems. When it comes to an enterprise strategy for securing secrets and credentials, it has likely been focused largely on improving key management, rotation, complexity, and implementing MFA with 'step up' verification challenges for the most critical and entitled systems and superusers within your 4 walls. To handle identities outside of your '4 walls' you've likely set up a SAML-based system that enables a two-way channel for communicating trusted 3rd party assertations of identity.

Longstanding struggles across User Experience, Developer Experience, Operational Ownership and Risk Management have brought about stronger practices around persistent access, least privilege roles and common-sense separation of duties. The unpopular removal of super-user privileges from Development, QA and other staff is likely mostly complete, ensuring more control and standards around the restriction of Production access and reducing exposure to hijacked user accounts. To improve Developer and Operator experiences, investment in workflows and processes to handle short-lived, temporary access or monitored 'break glass' functions has helped to alleviate Developer fear of 'losing access' to 'walk the beat' on their application, giving them a reasonable option to sustain productivity while improving secure handling of credentials and entitlements.

On the system side, Application teams have created system user accounts largely as they are needed, with general re-use of credentials across application components. Within the secure enterprise network perimeter Security teams have been more relaxed about how system users are managed, provisioned and maintained. Privileged Access Management (PAM) policies for your known human users and application system users typically coexist in your Active Directory, Kerberos, Radius, OAuth, OIDC, and/or SAML systems. You use inter-process authentication for communication between backend systems and support

the creation by application teams of digitally signed certificates for web-based communications. While some teams may be using enterprise certificate management tooling to generate and store certificates, the practice of copying certificates into unencrypted and under-monitored filesystem storage is prevalent.

There is no doubt that within the last 4 years you've been having more and more conversations about public cloud computing and establishing more granular security zones within your organization, all of which have been piling on complexity as new zones of varying trust are being built and additional authentication models, roles, privileges and tokens are being constructed outside of traditional guardrails. You simply can't rely on applications being on trusted and safe networks anymore. Your hair is probably standing on end thinking about all of the AWS and Azure authentication tokens sitting inside scripts and source code repositories that have been rushed together for sprint demos. It's also likely there are several implementations of vault solutions as business and infrastructure teams stand up PoCs and trailblazing applications on Public Cloud.

Fragmentation in the Market

When it comes to Secrets Management you are probably hearing about new products out there, some of which are starting to pop up in the CNCF Cloud Native Landscape for Secrets Management[2] – and no doubt you are starting to hear "why don't we use XYZ?" every other day. But, as an experienced IT professional, you know that a cardinal sin of IT and a surefire way to 'miss-the-mark' is to open the pocketbook without understanding what you have, what you need to have, and assessing the cost-to-value ratio across your options. Let's say your enterprise has a considerable investment and footprint in CyberArk – does that mean you have what you need? Or do you need to widen the toolkit? Like everything in life, it depends...

[2] https://landscape.cncf.io/category=key-management&format=card-mode&grouping=category

Secrets Management Capabilities

The first question requiring clear definition is, are you looking to build a new Secrets Management system for applications, or are you going to enhance and augment an existing set of tools that you have already built to manage human users and some level of critical system secrets? It may be that this is the decision facing you as you read through this paper - so let's dig a bit deeper into the capabilities you should be thinking about in your strategy and guiding decisions on how to collectively manage human and application secrets.

Capability #1: Centralized Secrets

Centralizing Secrets Management provides a significant opportunity to take a fresh look across the infrastructures you have accumulated over time and rationalize the logical boundaries and polices that govern the interactions inside and out.

While your enterprise may have rigorous and mature practices built up around PAM of your in-house infrastructure the chances are there are still pockets of critical authentication management that are not under centralized management. These could include things like:

1. Database users, admins, and system accounts
2. Vendor and SaaS platform admins and system accounts
3. Cloud (e.g. AWS, Azure, Google, etc.) subscription and resource accounts
4. DevOps pipeline tools and orchestration accounts
5. System users and trusted delegators

Incomprehensive and decentralized management of authentication, secrets management and monitoring of these critical assets risks exposures due to mismanagement, incorrectly applied policies, unprotected secrets and a lack of visibility, monitoring and governance.

While, to non-Security and Risk personnel, this lack of centralized oversight may be benefitting 'agility' the reality is that tokens, passwords, administrative accounts and identities risk being embedded in code, configuration, scripts, database tables, bespoke key/value stores, Excel files and sticky notes. Aside from the obvious precarious security posture, it also creates a cumulative obstacle for driving the rationalization, standardization and centralization of critically important secrets for fear that 'working' implementations may break if this is changed.

But wait... isn't centralization akin to putting all your eggs in one basket? The key nuance is centralizing management, access, governance, consistency and standards, so that you can more effectively secure them. This doesn't necessarily mean you will have all your secrets in one physical storage container.

If there are gaps in your existing centralization efforts in these areas, there is good news and bad news. The good news is that most commercial software that exists today can cope well with Privileged Access Management and leaders in this space have built a robust set of plugins and/or established marketplaces

where fit-for-purpose integration approaches can be selected. The bad news is that the brunt of the integration work now falls on iterative adoption, testing and integration to a centralized solution – fighting both the cultural resistance to change and the costs of remediating this type of technical debt. While there is no adoption magic wand, there are features that enable more easier adoption or greenfield integration. But more on that in Capability #5.

Lastly, you will need to work with Engineers and Operational staff to ensure proper removal of legacy secrets that might be hiding in code repos, shared documents, databases and other key/value stores. This will involve using scanning tools, however in many cases your Vulnerability Management team may have already invested in commercial and open source tools like SonarQube, Fortify, and/or Tenable to help identify cleartext password vulnerabilities. By the way, if you don't have a good relationship with the Vulnerability Management folks in your enterprise, it's definitely time to build it – they'll be helping you harden the infrastructure that manages your secrets, in addition to budgeting, deploying, and integrating tooling that will help stop the proliferation of unmanaged secrets! But more on that in Capability #4.

Capability #2: Secrets Generation and Distribution

The guiding principle of secrets generation is that all secrets should have a unique and expected lifecycle, with a defined time-to-live (TTL) that is as short as reasonably possible. As the industry has experienced more and more credential and data breaches, there is a trend towards the elimination of static and/or shared secrets in favor of on-demand, dynamic secrets that are both generated and revoked automatically – keeping them both unique and evergreen. We'll get back to this point towards the end of this paper as we talk about the 800-lb gorilla of eliminating persistent privileged access.

Key features that will allow for this dynamic lifecycle of secrets will be appropriate service interfaces for use by applications in the request and retrieval of managed secrets. Because integrating applications with a centralized Secrets Management system may be non-trivial it is important to assess the current availability of plugin support for commercial off-the-shelf and open source software when considering the best solution for your enterprise. Where modern, cloud-native applications and infrastructure will find native support for managed secrets, the availability of SDK toolkits and APIs for integrating in-house or unsupported software to a Secrets Management system will further enable the transition of legacy or in-house systems to a Secrets Management solution.

If your enterprise hasn't yet provided good capabilities around management of application secrets for cloud-native developers, it's likely that tactical solutions have already been built around cloud adoption – and decisions on how and where to store those secrets are being made possibly independent of Security and Risk teams. This may also include database and storage access, administration clients, and/or Vendor SaaS software that performs privileged tasks and automation.

Capability #3: Leasing and Revocation

Secrets should have a short time-to-live, meaning as short as reasonably possible based on the use case. Our suggested way to evaluate features for the management of lease and revocation policies is to understand the types of use case your organization needs to support and build policies around. The types of features available are:

- Time-bound secrets
- One-time secrets
- Automatic onboarding / reconciliation of unmanaged accounts
- User revocation
- Session suspension
- Revocation of a tree/hierarchy (e.g. response to credential breach)
- Restoration of a tree/hierarchy (e.g. recovery from a DR or failure event)

Most IT professionals understand the concept of time-bound and one-time secrets, however more modern capabilities have enabled IAM teams to implement Just In Time (JIT) creation, sometimes referred to as Dynamic Secrets. This refers to the ability to generate secrets upon request. Supporting a feature like this enables rapid adoption and on-demand processes to take advantage of managed secrets without having to be fully setup ahead of time. For example, if an application can identify itself with a role, as long as a suitable policy exists the secret can be generated. This is not only useful for application secrets, but human users as well where you may want to build workflows around conditional approval criteria for creation of valid credentials with specific accesses for a specific lease time.

Dynamic Secrets are a key capability for moving your enterprise away from persistent privileged access completely, offering an order of magnitude improvement to security posture. Ideally, the use of any persistent privileged accounts for human or system users can be displaced in favor of Dynamic Secrets. As controversial as that sounds, this is where this space is ultimately headed, but it requires not just technical capability but significant culture change as well. The best guarded secrets are those that never existed prior to needing them, have a rigorous workflow / challenge for attaining them and then are tightly managed during their short lifespan.

Handling of Unmanaged Secrets

Some features that may exist in Privileged Access Management solutions include the ability to scan for and discover unmanaged privileged accounts such as local administrator accounts, cloud-based accounts, cloud subscriptions, and other localized infrastructure or application administration accounts. Whether you are seeking to invest in cloud-oriented Secrets Management tool or building capabilities around existing PAM solutions its important to build a strategy around simplifying the experiences around discovery, enrollment and consolidation of unmanaged secrets into the strategic platform. Coupled with Dynamic Secrets, the adoption approaches for transition to managed secrets will have higher end-user satisfaction and more likely to be executed successfully.

Cyber Resiliency

We are likely to all agree that secrets should be rotated or expire very frequently, especially when they are used for privileged access and break glass activities, but cyber resiliency concerns and the rapid,

scaled issuance of secrets underscore the need to respond & restore. Typical cyber resiliency scenarios such as hijacked credentials, man-in-the-middle attacks and infrastructure vulnerabilities may expose unauthorized access to the underlying OS of both your Privileged Access Management services and Secrets Management as well as the individual endpoints and user accounts where secrets are employed.

Having a playbook and rule-based automation for revoking users, suspending them, revoking entire namespaces or trees, and recovery across those scenarios from trusted replicated sources is essential in building robust cyber resiliency.

Capability #4: Security of Secrets

The core features of a Secrets Management system exist to provide security, monitoring, threat detection, data breach prevention and cyber recovery of secrets material. However, it is important to take into account and apply robust security across the full stack of infrastructure and data flows, not just those directly related to the secrets management system. This should include secrets-at-rest, secrets-in-motion, secrets-in-memory, least privileged access controls, security of the underlying infrastructure, audit logging and monitoring. We will cover logging and monitoring in Capability #7.

Encryption of Secret Material

Secrets-in-motion pertains to several general paths including the transit pathways. When designing for appropriate controls and security of secrets-in-motion, one should consider encryption around both the channel of communication (e.g. TLS/SSL) and of the data (the secret) itself. In an ideal transit path, both the data payload and the transit mechanism are both encrypted at each node pathway and in memory.

Secrets-at-rest pertains to several areas where secrets may be stored in either encrypted or unencrypted forms in some type of filesystem storage.

Secrets-in-memory pertains to temporary storage at runtime in volatile memory. This could be within the application's memory space or host container's memory space – as environment variables are often used to temporarily hold runtime configuration and may include credentials.

The following diagram shows, at a high level, an example transport lifecycle of secrets, overlaid with typical data paths, in-memory storage and disk storage encryption.

Access Policy Management

A well thought out set of fine-grained policies that enforce least privileged access principles is the desired goal. It's unlikely that any mature product will lack in a core ability to create, organize and customize granular policy level definitions so the key will be to make sure you can effectively identify and correlate your needs to sets of capabilities around:

- Separation of policies into provider service groupings or 'engines' (e.g. AWS, MongoDB, SSH)
- Federation of the management/maintenance of those provider service groupings
- Ease of maintaining policies – Can they be encoded in markup language? Can they be source controlled? Is there flexible tooling or programmatic interfaces around CRUD operations? Can they be easily recovered/restored (both latest and specific version)
- Ease of maintaining identities, aliases and groups across multiple authentication methods
- Default stance of least privilege – meaning assume no access privileges
- Version control of all data & configuration objects

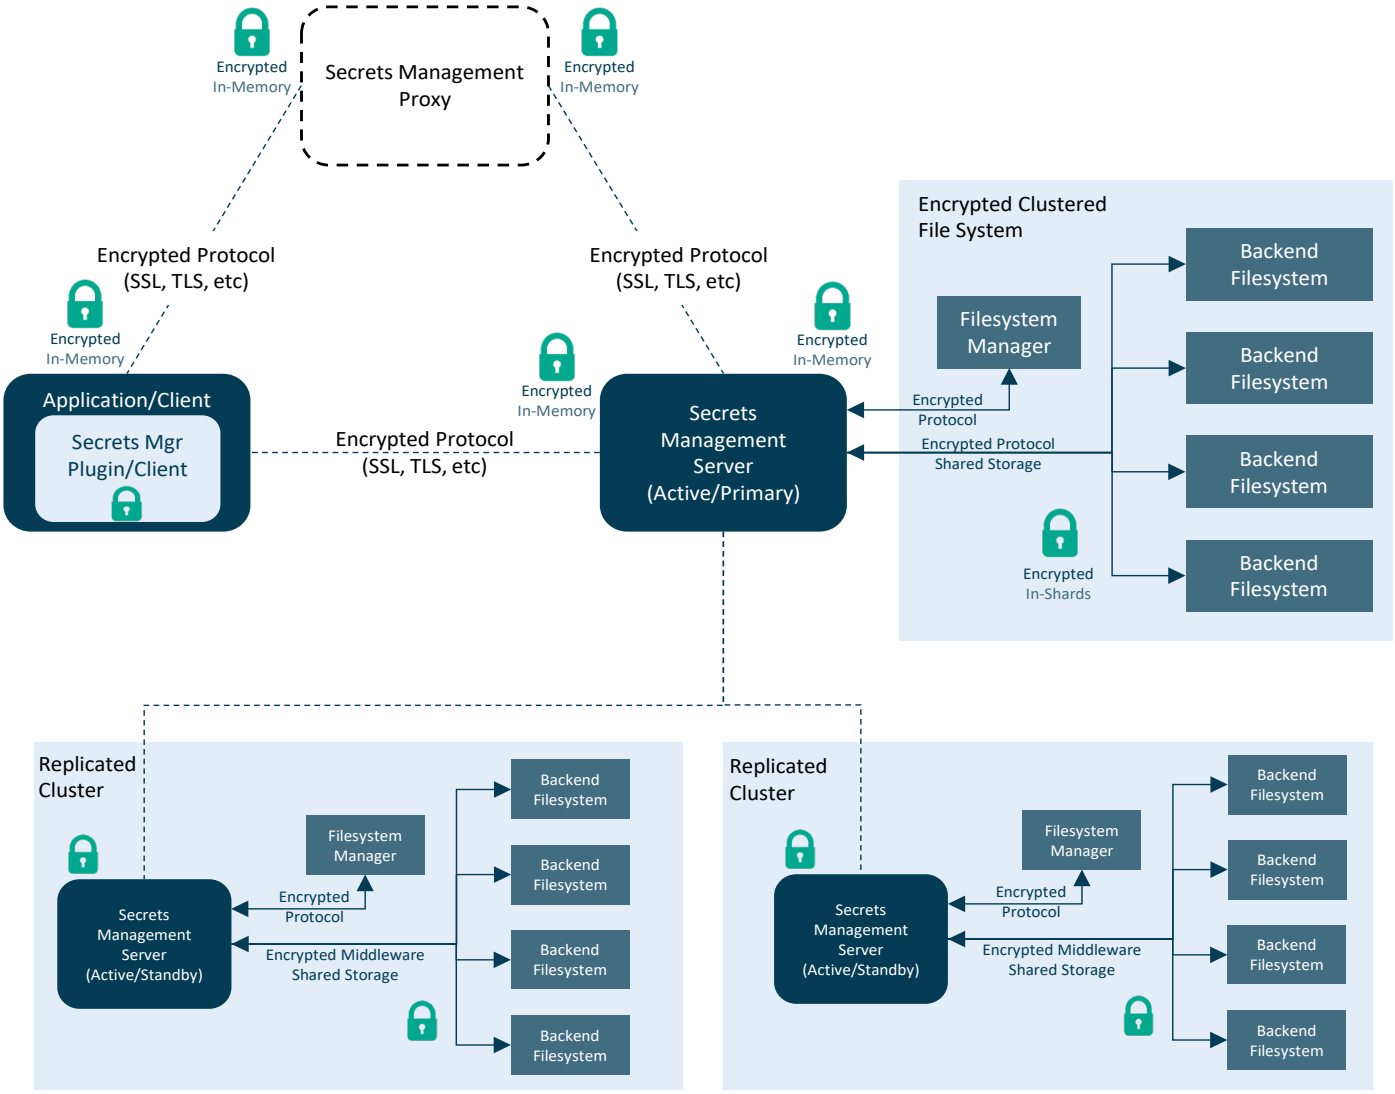

FIGURE 1: HIGH LEVEL LOGICAL DATA FLOWS TO PROTECT

Figure 1 illustrates a typical architecture of an application or client and the round-trip data flows to a centralized Secrets Management server and its backend infrastructure, including its replicated clusters.

This demonstrates an end-to-end implementation that considers handling of sensitive secrets in-transit and at-rest through the patterned use of encryption protocols and distributed backend storage.

Infrastructure Hardening

Before we can claim 'mission accomplished' on this topic of security, we also have to address operational hardening considerations around the underlying infrastructure running your Secrets Management solution, service-mesh middleware, and storage backend. Some of the top considerations in this area are:

Considerations	
Master Tokens	Does your secrets management system itself have a root token or credential? If so, can it be bootstrapped with a recovery key that can be put in a break glass system? Given the obvious threat of a hijacked root key credential, you should be asking yourself if you can delete it from existence and rely on break glass access only.
OS Hardening and Vulnerability Scanning	If you are running on RHEL you will be using SELinux and you will want to perform PCI DSS compliance scanning using OpenSCAP or SCAP Workbench to help identify and remediate exposures across your infrastructure. Similarly, you will want to seek similar PCI DSS hardening approaches for Windows Server as needed. Consider the use of FIPS encryption as well.
Software supply Chain	Use only official packages from the provider, and not from other public repos – even if they are thought to be secure repos. On linux systems enable gpgcheck on every repo, and similarly with Windows MSI ensure you are using trusted publishers.
	If you are using container images, only use the official images from the vendor. Images from public repos should be considered unverified and could potentially lead to arbitrary command execution.
	If you are using terraform or other infrastructure orchestration engines make sure you are comparing signed checksums or your organization's digital signatures so that compromised packages aren't being unknowingly deployed.
In-memory Secrets Leaks	Prevent scenarios in which memory can be dumped to the filesystem, including core files and hibernation states.
Denial of Service Protection	Set quotas, throttles, timeouts, back-offs and/or other safeguards against DDoS or unintentional resource starvation to ensure the stability of your system from excessive requests

Capability #5: Scaling, High Availability and Redundancy

A simple single datacenter solution design will not provide mechanisms for disaster recovery or recovery from cyber incidents. Secrets Management infrastructure will need to be replicated and made highly available in performant ways across multiple datacenters and geographies. Understanding how to scale the overall Secrets Management platform across your global enterprise entails considering not just geographic impacts to performance, but making sure the strategy includes performance of replication, backend storage and client requests against a view of current and forecast demand.

The critical end-to-end components and infrastructure as described in Capability #4 (Security) are key to ensuring each component in your design is accounted for and have known set of performance limitations in which to measure and alert by. As well as understanding the impact of the types and volumes of secret requests when considering architectural design and scale, it is imperative to understand constraints and performance impacts in allocating and scaling system resources, such as policies, groups, namespaces, secrets, and other objects that are going to have to be written to memory and/or written to backend storage.

If your enterprise is seeking to buy commercial software, it is important to get clear guidance from the vendor on the following aspects of their solution, including impact on infrastructure decisions:

- Storage requirement for a secret
- Storage requirement for policies
- Maximum size of a secret
- Maximum limits for each field in a policy
- Maximum size of a policy
- Maximum number of policies that can be assigned to a group or identity
- Maximum limits for namespaces or other hierarchal groupings
- Maximum limits for nesting of namespaces or other hierarchal groupings
- # requests / second that can be handled by a client
- # requests / second that can be handled by a server
- Maximum # of provider services or engines

Ensuring that your design takes into account these factors will reap long-term operational benefits and ensure that you can implement denial of service safeguards (as mentioned in Capability #4 - Security) and that you can apply the right telemetry and observability to ensure appropriate performance, stability and uptime across the enterprise.

From a disaster recovery and cyber resiliency perspective it will allow you to plan for impact and recovery, as well as assure your RTO/RPO when it comes to recovery of replicated secrets.

Capability #6: Monitoring and Reporting

Many of the roll-your-own Secrets Vaults stood up by Development teams for their cloud native applications are often implemented without the cooperation of Information Security teams, outside of the visibility of Cyber Security and Incident Response teams. If audit and usage logs are being generated, but not integrated into enterprise risk management tools and/or an organizational SIEM, then it is likely that user accesses are not being monitored for insider threats, brute force/probing events, or other anomalies.

From our experience, combining push and pull capabilities for humans and systems is required for complete coverage, including:

1. Standard dashboards made available to Operational, Risk, and Cyber Incident teams and reports that visualize key metrics, events, anomalies and trends pertaining to their functional role
2. Standard reports that are pushed out to key stakeholders that incorporate key metrics, events, and anomalies, and trends
3. Realtime events that are pollable (API, Database, Filesystem, etc.)
4. Realtime events that are published (e.g. SNMP traps, Application Messaging, etc.) to upstream event correlation systems (e.g. Enterprise SIEM), centralized logging platforms, and CIRT Issue Management Systems (e.g. ServiceNow, Remedy, etc.)

Events that you will want to capture will include:

- Administrative / privileged accesses and activities
- CRUD activities on any of the data and configuration objects in the Secrets Manager (Secrets, Groups, Roles, Namespaces/Partitions, etc.)
- Operational health indicators and heartbeats
- Proactive events (e.g. Trigger / Throttle based events, Resource Utilization Warnings, etc.)
- Detailed usage logging
- Anomalous behavior (excessive attempts, failed step-up challenges, use of inactive/suspended secrets, etc)
- Aging of encryption and identity keys

The granularity of data collection and ability to access/integrate data will be key in maintaining the right levels of flexibility to support your enterprise. Among the key monitoring capabilities to look for in any solution are ease of integration with downstream systems (e.g. Enterprise Logging, Monitoring, Alerting, Event, Issue Management) and the upstream systems that would need to adopt or integrate secrets (e.g. Vendor Product Authentication, Enterprise Authentication libraries, MFA solutions, and Service Orchestrators like Terraform) used in your organization.

Whether or not your enterprise is investing in a 3rd Party solution that provides standard out-of-the-box reporting, event correlation, threat detection, or other more advanced monitoring capabilities, the implementor will need to work with stakeholders in Cyber Analytics, Issue Management, and IAM to align and integrate to standard enterprise tools, further emphasizing the importance of data access and portability.

Balancing Build vs Buy

Part of assessing your current capabilities is understanding your workforce and strategy for Identity and Access Management. Will you maintain a staff of engineers and/or developers to build, maintain, and support integrations to vendor, in-house, and open source tools? Or is your organization looking to shrink its proprietary integration footprint and leverage more customizable, but out-of-the box software?

If you are leaning towards building, is your organization prepared to handle the investment needed to build product plugins and engineer product integrations? What about version control and change management of policies, groups, roles, and other objects? Is there a defined team that will analyze patterns of usage, anomalies and threats? Is the Cyber Security team onboard to take on the responsibility of building compliance and monitoring in the SIEM?

5 years ago, if you had asked an IAM Professional in a large financial enterprise if CyberArk – the most widely used solution for managing human privileged access in financial services – was still the right solution for Cloud adoption, you probably would hear about interoperability and integration challenges, difficult plug-in development for in-house and vendor software, and relatively slow vendor market response. However there is now a fairly robust marketplace for plugin support available in CyberArk's Marketplace[3].

If you contrast this against a product like Hashicorp Vault, which offers flexibility for designing the right deployment model across regions, lanes, and clusters encompassing your Development, QA, UAT and Production Application Environments, you may find more difficulty in building up the in-house expertise needed and instead source an experienced integrator. You will also need to develop a strategy that considers management of an application secrets vault in addition to traditional PAM done for human users, with rationalization opportunities further down the roadmap.

The flexibility and time-to-market of a separate solution for application secrets could be a significant benefit - however, only to the extent that you are willing to staff, source or re-assign resources to engineer Secrets Management Architecture, Role Management & Policy Design, plugin-development, operational support expertise and compliance monitoring. Given the necessary investments that need to be made in traditional PAM and Secrets Management toolsets and the recent disruption in this space, you will likely find that you will need significant investment in CAPEX and OPEX to build or buy modern solutions for human and application secrets to complement existing services and capabilities in your organization, regardless of which journey you are on.

[3] https://cyberark-customers.force.com/mplace/s/#---Privilege_Cloud_Support__c-ALL

How Citihub can Help

Citihub have experience of implementing both cloud-oriented Secrets Management products and traditional PAM solutions. Leveraging our industry expertise and history of both application and risk-oriented assessments we can play a role in helping your organization rationalize the most optimal strategic product mix of IAM capabilities, but also helping on design, implementation, scaling, hardening, and adoption that is required – all of which are crucial to ensuring centralized management, oversight and security across encoded secrets.

We work with the majority of global financial services providers across IT Delivery, Operations, Risk and Regulatory disciplines and are familiar with leading PAM and Secrets Management products and critical integration of business and technology applications in public, private and hybrid operating environments. Our consultants have extensive pedigree in financial services, capital markets, and heavily regulated industries and understand how to apply and scale solutions to the highest level of security and quality.

About Citihub Digital

Recoding the Digital DNA of Financial Services

Citihub specializes in digital transformation for financial services. We are passionate about building more agile, dynamic companies by harnessing the full power of cloud technologies. Together with our clients, we solve some of the industry's toughest challenges.

Digital Natives, Born in Financial Services

We thrive in the complex, highly-regulated environment of financial services. We're intimate with the industry's application ecosystems, from high frequency trading to retail payments services, and have deep experience resolving security, risk and compliance challenges. We bring domain-specific methods and IP to every project.

Going Beyond Technology

The work we do reaches far beyond technology; it requires fundamental changes to the operating model, processes and culture of an organization. Our blended teams of technologists and industry specialists span diverse functions and disciplines, acting as catalysts to bridge silos and maximize the value of your own team.

Building Lasting Relationships

We work with 9/10 of the biggest banks in North America and Europe. Our clients trust us to deliver and they stay with us because we do. We've had year-on-year relationships with our top 15 clients for an average of 8 years.