



Keeping secrets is hard.

Is the Cloud making it harder?

A CloudGeometry Insights Research Note

EXECUTIVE SUMMARY

Effective secrets management across the modern software development life cycle requires a top-down, unified approach. But it's more than a technical problem: applying up-to date thinking about Secrets management must be adapted to the unique architectural constraints of business critical environments and the organizations that rely on them. The catch: solutions that work in one environment (e.g. a static data center VM) often break down in another (e.g. ephemeral serverless function). Among the many differences that organizations must face in managing the complexity of their business critical systems are differences in trust boundaries, identity models, automation, and ephemeral compute. Enterprises must therefore adopt both centralized best practices and environment-specific tools to keep secrets safe throughout development, deployment, and operations.

A high stakes failure mode: assuming solutions effective in one context, such as a static data center virtual machine (VM), often prove inadequate in others, like ephemeral serverless functions. IT teams, and technical leaders and business executives alike too often ignore fundamental differences in trust boundaries, identity models, automation capabilities, and the variant nature of cloud computing and hybrid environments. Read on for more detail about:

- [An up-to-date look on classic virtualization and VMware](#)
- [SaaS Architectures & Multi-Tenant Cloud Applications](#)
- [Function-as-a-Service & Microservices](#)
- [Containerized Applications \(Non-Kubernetes\)](#)
- [Kubernetes \(K8s\) Fabrics and advanced orchestration](#)
- [Cross-Environment Considerations – Why One Size Doesn't Fit All](#)

Modernizing applications and infrastructure puts secrets management in sharp relief for both business and technical leverage. Yes, systems' business value can persist long after they are developed; they enable customer retention, amortize acquisition costs, accumulate operational efficiency via well worn valuable data and processes, and more. That said, it's hard to overstate the impact of investing (or the downside of neglecting upgrades in cyber security via secrets management. Rare is the modernization that does not need to address the challenges described here.



Introduction

Secrets management – the practice of securely handling sensitive credentials (API keys, passwords, certificates, etc.) – is both difficult and critically important in distributed cloud computing environments. Secrets management is a foundational security discipline in distributed cloud environments—one that enterprise-grade architectures cannot afford to mishandle. In our real world experiences with many, many client deployments, we have repeatedly seen how gaps in secrets lifecycle management can cascade into large-scale security events, especially in dynamic, multi-cloud SDLCs.

Modern software architectures span on-premises VMs, multi-tenant SaaS applications, microservices/FaaS, containers, and Kubernetes clusters, each introducing unique complexities. In complex cloud ecosystems, numerous non-human identities (applications, services, automated processes) must authenticate and communicate across a variety of trust boundaries, dramatically expanding the attack surface. Without a robust strategy, organizations face secret sprawl, hardcoded credentials, fragmented controls, and frequent leaks.

Classic VMware Virtualized Infrastructure (On-Prem Data Center)

Classic enterprise infrastructure built on VMware vSphere and virtual machines presents the “old-school” secrets management

challenges. Here, applications often run on long-lived VMs within a corporate network, and historically, secrets (database passwords, admin keys, etc.) might be handled in ad-hoc ways – e.g. manually set in config files or stored in scripts – under the (often false) assumption that the internal network is a safe trust zone.

Challenges & Characteristics of VMware/VM Environments

*NIST Cybersecurity & Privacy
Manual Distribution & Rotation
Assumed Trust Boundaries
Scale and Heterogeneity*

In a traditional virtualized data center, secrets management difficulties stem from manual processes and implicit trust.

- **Credentials Fragmentation:** Without centralized management, credentials get scattered: hardcoded in plaintext across VM images, configuration files, or code; it's even common to find database passwords in config files or SSH keys left on VMs. These scattered secrets become hard to track and rotate, leading to “incomplete visibility” and fragmented control. Each VM, application, and admin may handle secrets differently, creating gaps that attackers can exploit.
- **Manual Distribution & Rotation:** In VMware environments, deploying a new application often meant a human or ops script manually inserting secrets (e.g.

updating a config file with a password). Without automation, secrets go stale – passwords remain unchanged for years because rotating them would require touching dozens of VMs manually. This failure to rotate credentials regularly is a classic risk. If a secret is exposed, long-lived static credentials amplify damage.

- **Assumed Trust Boundaries:** On-prem VMs typically run deep inside a trusted corporate network. This can breed complacency – teams might treat secrets with less care (storing in plain text) since “it’s all behind the firewall.” However, virtualization admin consoles themselves carry broad privileges (the vSphere admin can access any VM). A single compromised admin account or VM can lead to lateral movement and exposure of numerous hardcoded secrets. Perimeter security is no substitute for secret security.
- **Scale and Heterogeneity:** An enterprise may have hundreds or thousands of VMs across dev, test, prod. Each VM or application instance has its own set of secrets (VM credentials, application API keys, etc.) that need managing. At this scale, manually keeping track of which secret lies where – and who has access – becomes untenable. Mistakes (e.g. forgetting to remove a retired VM’s credentials) are common. The result is often “secret sprawl” – many unused or duplicate secrets lying around unmonitored.

In summary, the plethora of security risks that persist from the VMware-era is rooted in decentralization and manual practices: secrets spread across many VMs with inconsistent handling. The convenience of binding workloads to virtualized hosts in production opens a host of weak or absent upstream integration gates

and compound vulnerabilities. Developers and operators share secrets insecurely (emails, spreadsheets) during development, embed them in VM templates for deployment, and rarely revisit them in maintenance – a wide open door for breaches and exploits.

Tools & Best Practices in VMware Environments

Centralized Vaults
Enterprise Password Managers & PAM
Configuration Management With Encryption
Cloud-Based Secrets Managers For Hybrid

In recent years, enterprises have recognized these issues and begun adopting centralized secrets management even for on-premises infrastructure. Key practices and tools include:

- **CENTRALIZED VAULTS:** A core best practice is to stop storing secrets on each VM and use a central secrets vault service. Tools like *HashiCorp Vault* (self-hosted) or *CyberArk Conjur* can be deployed within the data center to act as an authoritative secrets store. Applications on VMs authenticate to the vault to retrieve passwords or keys at runtime rather than reading from local files. This centralization allows enforcing access control, audit logging, and rotation in one place. For example, Vault can hold database credentials and lease them to apps on-demand, with automatic expiration – eliminating hardcoded DB passwords on each VM. Vault’s integration with VMware is evolving (e.g. VMware’s *Tanzu* integrates External Secrets Operator with Vault), and solutions exist to **authenticate VMs to Vault** using vSphere trust info (similar to cloud instance IAM).

- **ENTERPRISE PASSWORD MANAGERS & PAM:**
Many enterprises extend their Privileged Access Management (PAM) solutions to cover infrastructure secrets. For instance, *CyberArk (Enterprise Password Vault)* or *Delinea* store not only human admin passwords but also application credentials used on VMs. These systems can rotate the credentials automatically and inject them into VMs when needed. This mitigates the risk of default or unchanged passwords on virtual servers.
- **CONFIGURATION MANAGEMENT WITH ENCRYPTION:**
Traditional config management tools (Chef, Puppet, Ansible) introduced secrets-handling features (e.g. *Chef Encrypted Data Bags*, *Ansible Vault* files) so that sensitive config values are encrypted in repos and only decrypted on the target VM at deploy time. While not as robust as a full vault service, this was a step toward automating secret distribution securely (and avoids putting plaintext secrets in scripts or VM images).
- **CLOUD-BASED SECRETS MANAGERS FOR HYBRID:**
As hybrid cloud becomes common, some organizations choose to consolidate on a cloud secrets service even for on-prem apps. For example, using AWS Secrets Manager or Azure Key Vault to manage credentials for applications running on VMware VMs, accessed via secure APIs. AWS supports mechanisms like IAM Roles Anywhere to let on-prem servers assume an IAM role and fetch secrets from AWS Secrets Manager securely. This gives a single store for secrets across on-prem and cloud, though it introduces a dependency on the cloud provider's availability.

Regardless of tool, important practices include encrypting secrets at rest (e.g. Vault stores are

encrypted, or enabling VMware VM disk encryption for any config files), role-based access controls (only the applications or admins that need a secret can retrieve it), and auditing usage. Secrets should be frequently rotated to limit exposure window – modern tools make rotation easier (automated) than the legacy manual process. Eliminating hardcoded secrets is paramount: for example, using dynamic secrets or pulling secrets from a vault at startup ensures nothing sensitive is baked into VM templates or code repositories.

By adopting a centralized vault and automating secret handling, even traditional VMware-based organizations can drastically reduce the risk of secret leaks. One government enterprise pattern, for instance, funnels all app secrets through a vault so that developers never see production credentials – the apps fetch them at runtime with an agent. These approaches bring the on-prem environment closer to cloud-grade secrets management, enforcing consistency across dev, test, and prod. The overarching goal is to standardize secrets management across the enterprise, even if multiple backends exist, to avoid the fragmented practices of the past.

SaaS Architectures (Multi-Tenant Cloud Applications)

Compliance & Regulatory Framework

Software-as-a-Service (SaaS) applications have become a predominant business and operational model in the cloud. In a SaaS model, a provider hosts a single application serving multiple customer tenants. This environment is cloud-based and highly


automated. They introduce another set of challenges via the presence of multiple external tenants (with strict data isolation needs), making secrets management critical to “getting security right.” The trust boundary now runs between tenants as well as between the SaaS provider and customers, raising the stakes for any secret compromise.

Challenges & Characteristics of SaaS Multi-Tenancy

Tenant-Specific Secrets & Isolation
Scaling and Automation
Tenant Key Management & Compliance
Integration Secrets and Third-Party APIs
Assumptions Traps & Inconsistencies

In SaaS architectures, secrets management must ensure that each tenant’s sensitive data and integrations remain isolated and secure, even though the underlying application is shared. Notable difficulties include:

- **Tenant-Specific Secrets & Isolation:** Each tenant may have unique secrets that the SaaS application needs to function on their behalf – for example, API keys to pull the tenant’s data from an external system, or encryption keys for that tenant’s stored data. The system must manage potentially hundreds or thousands of tenant-specific secrets, isolating them such that Tenant A’s keys are never accessible to Tenant B. Any flaw in segregation (e.g. a bug that returns the wrong tenant’s data) could be catastrophic. Enforcing strict access controls per tenant context is non-negotiable: a secret that’s valid for one tenant should be unreadable outside that tenant’s scope. This is a different granularity of trust boundary – not just “internal vs external,” but per-customer isolation within the app.
- **Scaling and Automation:** A successful SaaS may rapidly onboard new customers, which means new sets of secrets constantly. Manual secret setup doesn’t scale – the process of provisioning a new tenant likely needs to automatically generate or allocate secrets (for example, create a DB user and password for that tenant, or an encryption key) without human intervention. The secrets management system must integrate with the SaaS provisioning workflow to create, store, and distribute those secrets instantly. Moreover, rotating or revoking a secret (e.g. if a tenant admin updates an API token) should propagate without downtime. High automation is assumed in SaaS, so secrets management has to be API-driven and scriptable.
- **Tenant Key Management & Compliance:** Many SaaS providers choose to encrypt each tenant’s data with a tenant-specific encryption key (sometimes even giving customers control of their keys). This introduces a key management challenge – storing and protecting potentially thousands of encryption keys and ensuring the application uses the correct one for each tenant’s data. The system might use a master secret to wrap/generate all tenant keys, which itself becomes a critical secret that must be safeguarded (the classic “secret of secrets”). Compliance standards (SOC 2, GDPR, etc.) often require strong encryption and separation, putting pressure on robust secrets management (e.g. hardware security modules or cloud key management services to back the key hierarchy).
- **Integration Secrets and Third-Party APIs:** A SaaS application often integrates with third-party services (email providers, payment gateways, cloud APIs) – these require API secrets that the SaaS holds. In multi-tenant context, there might be either a



shared integration credential or distinct ones per tenant. Storing these integration secrets securely is vital, as a leak could impact all tenants or an entire business function. Additionally, when tenants bring their own credentials (for example, a SaaS data integration feature where each customer supplies a credential to connect to their system), the SaaS becomes responsible for not exposing those and for using them strictly under the right tenant's context. This demands meticulous identity context tracking – the app must only fetch a credential from the vault for the currently served tenant and not mix them up.

- **Assumption & Inconsistencies:** Compared to an on-prem app for one enterprise, a multi-tenant SaaS can't assume any user is fully trusted. Zero-trust principles apply within the app: each request is tied to a tenant identity, and the app must enforce that only that tenant's secrets are accessible. Techniques like per-tenant encryption or separate tenant data stores are used.

An assumption that was safe in one-tenant scenarios (e.g. using a single master database password) breaks in multi-tenant – instead, you may need separate credentials per tenant database for isolation. Thus, secrets management in SaaS often needs a more granular design than in single-tenant apps.

Tools & Best Practices for SaaS Secrets

Cloud Provider Key Management Services (KMS)
Secrets Namespacing And Access Control
Automated Tenant Onboarding/Offboarding
Encryption Key Management Patterns
Rotation And Monitoring

To handle these challenges, SaaS providers leverage a combination of cloud-native services and robust architecture patterns:

- **CLOUD PROVIDER KEY MANAGEMENT SERVICES (KMS):** Given most SaaS are built on public cloud infrastructure, they commonly use the cloud's managed secret and key services. For example, *AWS KMS* and *AWS Secrets Manager*, *Azure Key Vault*, or *Google Secret Manager* to store application secrets and encryption keys. These services are multi-tenant aware at the cloud level (they enforce IAM policies), and they offer features like automatic rotation and secure storage. A best practice is to use tenant-specific Key Vault keys or secrets – e.g., create a separate Key Vault secret for each tenant's API keys, each tagged or access-controlled by tenant ID.

Microsoft recommends storing tenant secrets in Azure Key Vault with controls so only the appropriate tenant's context can access them. Similarly, AWS KMS can create a CMK (customer master key) per tenant to encrypt that tenant's data. The SaaS app then must ensure it uses the correct key (often by referencing a tenant identifier mapped to a key ARN or vault entry).

- **SECRETS NAMESPACING AND ACCESS CONTROL:** Whether using a cloud vault or HashiCorp Vault, namespacing secrets per tenant can be a sound strategy. For instance, in HashiCorp Vault Enterprise, using a separate namespace for each tenant or a path scheme like `secret/<tenantID>/*` to segregate secrets. (Note: HashiCorp charges extra for these features, and you can use this best practice even without paying for the full featured Vault Enterprise suite; just using a path based scheme as a best practice in the structure of your tokens. Combined with

dynamic tokens scoped to that path, the application can be certain it can only read secrets for the tenant it's serving at the moment. Some SaaS implementations even spin up isolated vault instances or logical partitions per customer (though this can be heavy) to maximize isolation. The guiding principle is strong isolation – even within the secrets management layer – mirroring the multi-tenant isolation of the app itself.

- **AUTOMATED TENANT ONBOARDING/OFFBOARDING:** Best practices dictate that all secret provisioning is wired into the SaaS lifecycle. For example, when a new tenant is created, automation should generate any needed credentials (like creating a new database user/password for that tenant's schema) and store them in the secrets manager. No operator should need to manually create or email a password. Likewise, when a tenant is offboarded or a credential rotated, scripts or functions invoke the secret manager's API to update or delete secrets. This reduces human error and keeps secret data out of the hands of staff unnecessarily.
- **ENCRYPTION KEY MANAGEMENT PATTERNS:** For per-tenant encryption keys, SaaS providers often use a hierarchical key model – e.g., a master key (in an HSM or KMS) that encrypts individual tenant data keys. The tenant keys might be stored encrypted in a database or in a vault, and only decrypted when needed using the master key (which itself is heavily protected, ideally by hardware security modules or cloud KMS). Some SaaS allow tenants to supply their own keys ("Bring Your Own Key"), in which case the system has to interface with an external key vault – another layer of complexity. Regardless, handling this gracefully via the cloud KMS APIs (which can ensure keys never leave the service) is the preferred approach.

- **ROTATION AND MONITORING:** In multi-tenant setups, rotating secrets is delicate – you might have thousands of tokens to rotate without breaking service. Providers rely on staged rotation (rotate one tenant's secret at a time, or use multi-version secrets where the new version is introduced while old is still accepted, then phased out). Cloud secrets managers often support versions to assist with this. Monitoring access to secrets is equally important: the SaaS ops team should watch for any anomalous access patterns, e.g. if suddenly a service tries to access another tenant's key (which might indicate a bug or breach). Audit logs from the vault/KMS help here, as well as wrapping sensitive operations with additional checks in code (assert the tenant context matches the secret being retrieved).

In essence, secrets management in SaaS is about designing for tenant isolation and automating at scale. The use of cloud-native secret storage, coupled with rigorous identity controls, ensures that solving the problem for one tenant doesn't inadvertently expose another. Assumptions from single-tenant environments do not carry over – each piece of sensitive data must be tied to a tenant identity and secured as such. By following these practices (isolate, automate, leverage strong cloud tools), SaaS providers can maintain zero trust between tenants while still centrally managing secrets. [Table 1](#) below summarizes some tool choices across environments, including SaaS:



Secrets Management Solution	Applicable Environments	Notes
HashiCorp Vault (self-hosted)	Data centers (VMware), Kubernetes, Microservices, Multi-cloud SaaS	Flexible, enterprise-grade vault; supports static and dynamic secrets, multi-tenant namespaces, and wide integrations (apps authenticate via tokens, Kubernetes service accounts, cloud IAM, etc.). Often used as a central platform in hybrid setups.
Cloud Provider Secrets Managers (AWS Secrets Manager, Azure Key Vault, GCP Secret Manager)	Cloud-native apps, Serverless (FaaS), SaaS architectures, Hybrid (via APIs)	Managed secret storage with high availability and integration with cloud IAM. Ideal for cloud deployments – e.g. AWS Lambda can fetch from Secrets Manager natively. Can extend to on-prem via secure connectors. Support automatic rotation and fine-grained access control per secret.
Kubernetes Native Secrets	Kubernetes (single-cluster scope)	Built-in object for storing secrets in etcd (can be KMS-encrypted). Easy to use via K8s API, but limitations in encryption, rotation, and cross-env usage mean additional tools are often needed for robust security.
CyberArk Conjur / Secrets Manager	DevOps pipelines, Containers, Hybrid cloud	Enterprise PAM extension focused on non-human identities. Integrates with CI/CD and container platforms (Kubernetes, OpenShift) to provide centrally managed secrets injection with strong RBAC. Often used in regulated industries alongside human credential vaulting.
Azure AD / IAM with Managed Identities	Cloud VMs, Containers, Serverless	(Not a secrets store per se, but an identity approach.) Uses cloud identity tokens instead of static secrets – e.g. Azure Managed Identity or AWS IAM Role for a resource eliminates needing an API key at all. Simplifies secret management by relying on short-lived credentials issued by the platform, though limited to accessing that platform’s services.

Table 1: Examples of secrets management solutions and the environments they best support

Function-as-a-Service & Microservices

Modern microservices architectures – especially serverless Function-as-a-Service (FaaS) – push secrets management to its limits. Applications are composed of dozens of small services or ephemeral functions, each needing credentials to talk to databases or other services. The environment is highly dynamic (containers or functions start and stop constantly) and fully automated via DevOps pipelines. This combination makes traditional secrets-handling approaches (manual config, static keys) unworkable – instead, secrets must be delivered on-the-fly, often by the platform itself.

Challenges in Microservices & Serverless Environments

Ephemeral Instances & Short Lifecycles
Service-to-Service Authentication
CI/CD Pipeline Integration
Polyglot and Distributed State
Function Constraints

Key difficulties in microservices and FaaS include:

- **Ephemeral Instances & Short Lifecycles:** In FaaS (e.g. AWS Lambda, Azure Functions) and containerized microservices, compute instances may live only for minutes or seconds. A function is spun up to handle a request then torn down. This ephemerality means there is no “manual step” to configure a secret – secrets injection must be automatic at startup or built into the environment. It also means traditional long-lived credentials are riskier: if a short-lived container holds a secret and it somehow leaks (logs, memory), that secret


might remain valid long after the container is gone.

The push is toward short-lived (ephemeral) secrets to match ephemeral infrastructure. For example, microservices might use database credentials that expire after a few minutes and are renewed automatically, so a dumped credential from a dead container is useless. Managing rapid secret churn and distribution is a new challenge not present in static servers.

- **Service-to-Service Authentication:** A microservices application may consist of dozens of services calling each other’s APIs. Every inter-service call potentially needs a secret – e.g. a JWT signing key, an mTLS client certificate, or an API token – to authenticate the caller. In a monolithic app, internal calls didn’t require explicit auth, but in a distributed microservice design, you often implement zero-trust networking, where each service call is validated. This leads to a high volume of tokens and keys in use concurrently.

As Wiz notes, “in a microservices architecture, you usually have services accessed by several other services, and protecting those interactions requires distributing access tokens or credentials to the right clients. Ensuring each microservice gets the secrets it needs (and no more) – and that those secrets are rotated or revoked without breaking dependencies – is a delicate dance.

- **CI/CD Pipeline Integration:** Microservices and serverless encourage very frequent deployments (even continuous deployments). Secrets management must integrate with CI/CD pipelines so that new



code versions or new services get the appropriate secrets at deploy time. For instance, a build pipeline might bundle a client secret or fetch env-specific secrets to inject into a container image or function configuration. This introduces risks if not handled carefully – e.g. a misconfigured pipeline might log a secret or expose it in an artifact. The challenge is to automate secrets throughout the SDLC (from build to runtime) without exposing them. Techniques like storing secrets in pipeline-level vaults (Jenkins credentials store, GitHub Actions secrets) and using short-lived tokens for deployment jobs are used to secure this process.

- **Polyglot and Distributed State:** Microservices can be spread across multiple environments – some might run on Kubernetes, others as Lambdas, others on VMs – within the same application. This heterogeneity means there may not be one universal secrets solution at runtime. One service might use AWS Secrets Manager, another might be better served by HashiCorp Vault. Achieving consistent policy (e.g. all secrets must rotate, all access is audited) across such a diverse stack is tricky. It often requires a combination of tooling or an agreement on a common denominator (for example, using Vault as an overarching solution, or ensuring each platform's secret store is configured with equivalent policies).
- **Function Constraints:** FaaS functions, in particular, have constraints – they may have limited runtime, memory, and no local disk. This means secrets retrieval has to be efficient (you can't have a function spend 2 seconds pulling from a vault on each invocation without incurring cost and

latency). Cloud providers addressed this by enabling secrets caching in environment variables (e.g. loading an encrypted secret at startup of the function container). However, using environment variables for secrets has its own security issues (they can be read by anyone with access to the process). If developers aren't careful, they might inadvertently expose secrets by logging environment variables or not using provided encryption features.

Tools & Best Practices for Microservices/FaaS Secrets

**Cloud Provider Key Management Services (KMS)
Secrets Namespacing And Access Control
Automated Tenant Onboarding/Offboarding
Encryption Key Management Patterns
Rotation And Monitoring**

For distributed microservices and serverless apps, the emphasis is on built-in platform features and automation to handle secrets, supplemented by advanced tools for more complex needs:

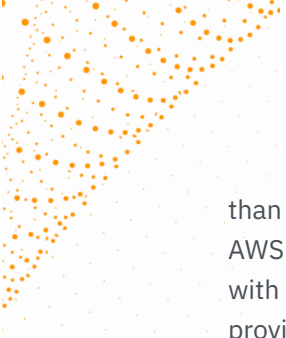
- **MANAGED SECRETS INTEGRATION IN FAAS:** All major FaaS platforms provide native ways to consume secrets. For example, AWS Lambda supports sourcing environment variables from AWS Secrets Manager or Parameter Store such that the values are decrypted and present at runtime without the developer hardcoding anything. Similarly, Azure Functions can automatically fetch secrets from Azure Key Vault via app settings references, and GCP Cloud Functions can pull from Secret Manager. Best practice is to use these platform-supported secret integrations rather than rolling your own.

Snyk's serverless security guide advises using a secure secrets store backed by the cloud provider or a trusted vault, rather than embedding secrets in code or config. The reasoning: these solutions handle encryption, access control, and versioning for you. As Snyk notes, using such a service "mitigates the risks of sensitive information stored in static files or environment variables" and greatly reduces exposure. In other words, let the platform deliver the secret at runtime, and do not bake it into the deployment artifact.

- **Ephemeral and Dynamic Secrets:** Embracing the ephemeral nature, many microservice architectures use dynamic secrets that rotate automatically. HashiCorp Vault is popular here – for example, a microservice at startup can request a database credential from Vault's database secrets engine, which generates a new user/password that expire after a configured time. The service uses it and Vault will revoke it later, forcing the service to get a new one next time. This limits the blast radius of any one credential. Similarly, Vault's approach can be used for cloud API keys, etc., but the integration complexity is non-trivial (each service must authenticate to Vault – often done via cloud IAM or Kubernetes service account auth – and handle renewal). Enterprises with high security needs adopt this to ensure no static secrets are lingering in microservices.
- **Sidecar Agents and Service Mesh:** In containerized microservices (e.g. on Kubernetes or Nomad), one pattern is running a sidecar secrets agent alongside the application container. For instance, Vault Agent Injector on Kubernetes will mount a shared memory volume into the pod with the

secret material, or Envoy sidecars in a service mesh handle mTLS certificate issuance for service-to-service encryption. This offloads secret retrieval/renewal to a separate process so the app code can just read fresh secrets from a file or memory. It's a best practice to decouple secret management logic from business logic in microservices – use the platform capabilities (sidecars, mesh) to automate secret distribution. Service Meshes (Istio, Linkerd) can manage TLS certificates for services, which is a form of secret (private keys) management, automatically rotating them. This addresses the inter-service authentication challenge by not relying on static API tokens at all but using dynamically issued identities for services.

- **Environment Variable Security:** When environment variables must be used (they are the only option for some FaaS to pass config), ensure they are encrypted at rest and not exposed in logs. AWS, for example, encrypts Lambda environment vars with KMS. Additionally, functions should retrieve secrets at invocation only if necessary and avoid lengthy storage in memory. A general rule: Do not store secrets in code (no hardcoded constants or committing secrets to source control). Instead, pull them from the environment or a secrets manager at runtime. Developer training and automated secret scanning (using tools like truffleHog or GitGuardian in CI) can catch violations early in the SDLC.
- **Cloud IAM Roles vs Secrets:** One of the best "secrets management" tactics in cloud microservices is to eliminate secrets altogether. In almost all cases, identity-based access. For instance, rather



than giving a microservice an API key for an AWS service, one can assign it an IAM Role with needed permissions – AWS then provides temporary credentials behind the scenes. This leverages the cloud’s identity model (which is short-lived and managed) instead of a static secret.

In AWS, a Lambda can run with an IAM execution role, or an EC2-based microservice can use an Instance Profile; in Azure, managed identities allow a service to access Azure resources securely without secrets. While this doesn’t solve all cases (e.g. external API keys still need to be stored), it removes a class of secrets that would otherwise need management. Many microservice architectures heavily use such features to minimize the number of secrets that must be manually managed.

- **REGULAR ROTATION AND REVOCATION:** Automation should also handle rotation in these environments. For example, if a microservice uses a database password stored in AWS Secrets Manager, using Secrets Manager’s rotation feature (possibly with an AWS Lambda function to create a new DB user and update the secret) will ensure a new secret value every X days. The microservice, on its next invocation or via a refresh mechanism, will pick up the new password. This requires that applications be built to tolerate secret changes on the fly – e.g., reading credentials at startup is not enough if they can rotate during runtime. Techniques include short connection timeouts (so new connections use updated creds) or signaling the service to reload config. While challenging, many enterprises enforce at least frequent rotation, knowing

that in a fully automated environment, the inconvenience is less than in legacy systems.

In summary, microservices and serverless demand that secrets management be integral to the platform – not an afterthought. The mantra is to avoid hardcoded or long-lived secrets by leveraging cloud-native secrets stores, ephemeral credentials, and automation. By doing so, organizations can prevent the explosion of secret leaks that could occur when you have hundreds of micro-components each with its own keys. As one security expert put it, planning for secrets management in microservices ensures tokens “are stored in a secure location and distributed only to the services that require them” reinforcing least privilege and minimizing exposure.

Containerized Applications (Non-Kubernetes)

Many organizations run containerized applications outside of Kubernetes – for example, using Docker Compose, Docker Swarm, or cloud container services like AWS ECS or Azure Container Instances. These environments share some challenges with microservices, but there are specific considerations for managing secrets at the container level (especially when Kubernetes’ more advanced tooling is not present).

Challenges in Basic Container Environments


*Image Build vs Runtime
Environment Variables & Config Files
Lack of Native Rotation Mechanisms
Secrets Persistence
Multi-Environment (Dev/Prod) Consistency*

Containers are by nature portable and ephemeral, which complicates secrets handling:

- **Image Build vs Runtime:** One challenge is avoiding secrets in container images. It's common during container build (Dockerfile) to require a secret (say, to pull a private package or repo). If not handled carefully, that secret can get baked into an image layer. Once baked, it's effectively leaked to anyone with the image. So developers must use multi-stage builds or build-time injection that doesn't persist. This is a new concern compared to VMs – the “artifact” (image) can inadvertently carry secrets if procedures aren't in place.
- **Environment Variables & Config Files:** By default, passing config to containers is often done via environment variables or mounted files. Natively using environment variables for secrets is dangerous, as noted in a CyberArk guide: they can be easily retrieved by anyone with access to the container process or host

Unlike a VM (where one might use an OS key store), containers have no innate secure enclave for secrets – everything in the container's process space could potentially be read if the container is compromised. Also, environment vars might leak through logs or error messages.

- **Lack of Native Rotation Mechanisms:** In a basic container runtime (unlike Kubernetes which has controllers, or Vault agents), there is typically no built-in way to automatically update a running container's secrets. For example, Docker Swarm secrets are immutable – if a secret value must change, you have to create a new secret and update the service to use it. This means rotation requires deploying new container instances with the updated secret. Without an orchestrator that can do rolling updates, this may cause service interruptions. Many teams therefore neglect rotation in container environments, which is a risk.
- **Secrets Persistence:** Ideally, when a container stops, any secret it had in memory or on disk should vanish. But in practice, secrets can persist on container hosts. For instance, if a secret was written to a file inside the container and the container crashes, that file might remain in a stopped container's filesystem on the host until it's removed. Similarly, environment variables might be viewable via host-level tools while the container ran. This is a challenge of “garbage collection” – ensuring secrets don't linger. Manual cleanup or using ephemeral volumes (tmpfs) for secrets can help, but not all environments enforce this. As containers scale out and back in, keeping track of where secrets might have been left (in logs, volumes, etc.) becomes complex.
- **Multiple Environments (Dev/Prod) Consistency:** Often, containerized apps are run by developers via Compose locally (with perhaps a .env file for secrets), and then in prod via a different method (maybe Swarm or mounted secrets). Inconsistencies here can lead to mistakes – e.g. a developer might



check in a Compose file with a hardcoded secret for testing, not realizing it's now in version control. Or the production team might forget to set an environment variable, and the container defaults to some insecure built-in credential. Ensuring that secrets management is consistent from dev to prod is a challenge – it might require adopting a secrets manager even for local development (which can hurt developer agility unless streamlined).

Best Practices & Tools for Container Secrets

Use Orchestrator Secrets Features
External Secrets Store With Injection
Encrypt At Rest On Hosts
Twelve-Factor App & Security
Developer Training & Tooling

Here is set of practices improve secrets security For container environments that are totally non-orchestrated or even lightly-orchestrated:

- **USE ORCHESTRATOR SECRETS FEATURES:** If using Docker Swarm or a similar system, take advantage of its secrets management. Swarm, for instance, allows you to create a secret and then attach it to services. The secret is stored encrypted on manager nodes and mounted in containers as a file, not an env var This approach means the secret never appears in plaintext in images or in the docker inspect output, etc. It also can be detached when not needed – e.g. Swarm allows a service to remove a secret after initialization, reducing exposure during runtime.

While Swarm lacks one-command rotation, the practice of regularly updating secrets via

service updates (as described in Docker docs) can be followed. The key: avoid baking secrets into images or passing via command-line, and instead use the orchestrator's mechanism which was designed for this purpose.

- **EXTERNAL SECRETS STORE WITH INJECTION:** In environments like *AWS ECS* (Elastic Container Service) or *HashiCorp Nomad*, there are integrations to pull secrets from external stores. For example, ECS allows specifying secrets from AWS Secrets Manager to be injected into containers as env vars (encrypted by AWS and delivered to the container at start). Similarly, HashiCorp Vault has a Nomad integration and also a tool called Vault Agent that can render secrets to a file inside a container. Using these, one can keep the source of truth in a robust secrets manager and let the scheduler place the secret into the container at runtime. This avoids managing a separate secrets distribution mechanism by hand.
- **ENCRYPT AT REST ON HOSTS:** Ensure that wherever secrets do touch the host (like the Docker Swarm manager storing secrets, or volumes containing secrets), encryption is enabled. Docker Swarm, for instance, can encrypt the secrets at rest on disk. If using file-based secrets (like **mounting** a volume with a config file of secrets), consider using disk encryption on that volume. This way, if an attacker gets a copy of the host filesystem or a backup, they can't read the secrets.
- **TWELVE-FACTOR APP VS SECURITY:** Containerized apps often follow the Twelve-Factor methodology which encourages storing config in environment variables. While great for stateless config, this advice is risky for secrets. Best practice

is to use env vars for non-sensitive config but use secret management for anything sensitive. For example, an app might have an environment variable pointing to `DB_PASSWORD_FILE=/run/secrets/db_pw` – the actual password file such provided by the orchestrator. This way, the secret isn't directly in the environment, but the app can read it at runtime. Adopting this pattern keeps the separation between config and secret while preserving the convenience of the twelve-factor approach.

- **DEVELOPER TRAINING & TOOLING:** Since containers are developer-driven, make sure engineers understand these best practices. Integrate container security scanning that checks images for hardcoded secrets or keys (tools can scan Docker images for AWS keys, etc. that might have slipped in). Use linters or commit hooks to catch if someone tries to put secrets in Dockerfiles or Compose files. Essentially, enforce that no secret ever enters the source code or image – only inject at runtime from secure storage.

By following these practices, even organizations not using full Kubernetes can achieve a decent security posture for container secrets. The principle remains to centralize and automate: use a secrets manager or orchestrator feature to handle the heavy lifting, rather than custom ad-hoc solutions. As CyberArk emphasizes, understanding where authentication secrets are needed in your container setup and controlling their configuration and storage makes the environment much more secure

Containers add speed and agility – but without proper secrets management, they could also rapidly multiply your secret leakage points. Thus, getting this right is essential before scaling out container deployments.

Kubernetes (K8s) Fabrics

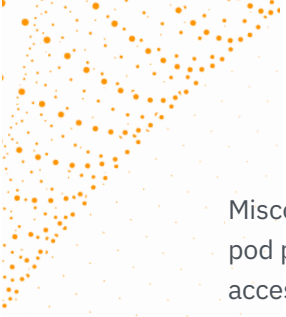
Kubernetes has become the de facto orchestration platform for containers in enterprises. It provides some native constructs for secrets, but also brings new challenges due to its dynamic, distributed nature. Kubernetes clusters often run many microservices, possibly even from different teams or tenants, on a shared fabric. Secrets management in K8s is both critical and notoriously tricky – in fact, “keeping secrets secret is tough because they need to be shared between so many different resources (in a K8s cluster).

Challenges Unique to Kubernetes

*K8s Native Secrets Are Basic
High Secret Volume & Distribution
Trust Boundaries and Multi-Tenancy
Ephemeral and Automated Nature
Interfacing External and Legacy Systems
Human Error and Exposure*

Kubernetes (K8s) introduces a mix of issues – some inherited from containerization and microservices, and others stemming from Kubernetes' design and multi-tenancy:

- **K8s Native Secrets Are Basic:** Kubernetes has a built-in object type called Secret, which is used to store small bits of sensitive data (like passwords, tokens). However, by default, K8s stores these secrets in plaintext (Base64-encoded) in its `etcd` database – without encryption at rest unless extra config is enabled. This means that anyone with access to the `etcd` data or a backup can read all secrets. Moreover, K8s secrets are namespaced objects accessible to any pod in the namespace by default if referenced.



Misconfigurations in RBAC or overly broad pod permissions can lead to unintended access. The platform won't prevent a pod from reading a secret it's been given, even if that secret was meant for another app – it's on administrators to set correct policies. In summary, the default state of K8s secrets is not as secure as many assume, requiring additional measures (encryption, RBAC) that are sometimes overlooked.

- **High Secret Volume & Distribution:** In a K8s cluster running dozens of apps, the number of secrets can be large. Each application might have several (DB creds, API keys, TLS certs, etc.), and these are all stored through the same K8s API. There is a scalability and manageability challenge – secrets must be created and updated through YAML manifests or API calls, often manually.

Kubernetes provides no built-in mechanism to automatically rotate or prune secrets. If a secret value changes, an admin has to issue `kubectl apply` with a new secrets manifest and trigger pods to reload it (usually by restarting them). Note that in a traditional Kubernetes setup, updating a Secret or ConfigMap does not automatically restart or redeploy your workloads; i.e., Kubernetes secrets management “lacks scalability” for large numbers of secrets without additional tooling. This can lead to stale configurations running in production, especially when dealing with dynamic values like credentials, feature flags, or environment configs.

Similarly, when secrets become unused (e.g. an app removed), they may linger in `etcd` unless cleaned up. This manual overhead at scale often results in stale secrets that never

get rotated or removed, simply because tracking and updating hundreds of secret objects is burdensome. A more effective approach: use [Reloader](#), an open source K8S controller that automatically triggers rollouts of workloads (like Deployments, StatefulSets, and more) whenever referenced Secrets or ConfigMaps are updated.

- **Trust Boundaries and Multi-Tenancy:** Many clusters are multi-tenant (either multiple teams or even multiple applications of different sensitivity sharing a cluster). The K8s access model is complex – service accounts, roles, and role bindings determine which pod or user can access which secret. Mistakes in RBAC configuration can accidentally expose secrets cluster-wide. For example, giving a service account access to list all secrets in a namespace could be dangerous if that account is used by many pods. There have been real incidents where poorly configured RBAC allowed an attacker in one compromised pod to retrieve all secrets in the namespace.

If clusters are truly multi-tenant (different clients), it's even more challenging – one might resort to strict namespace isolation or even separate clusters, because K8s Secrets are not designed for cross-tenant isolation beyond namespaces. Essentially, the trust boundary in Kubernetes is often at the cluster or namespace level, which might be coarser than the application's actual trust needs. This mismatch means operators must be very cautious in how they distribute and permit access to secrets.

- **Ephemeral and Automated Nature:** Similar to earlier points, K8s is highly dynamic – pods

come and go based on scaling or updates. This dynamic nature demands that secrets injection be seamless. Kubernetes will mount secrets into pods (as env vars or volumes) at pod startup, but if a secret is updated, pods won't automatically get the new value unless they are restarted. This can make secrets rotations that impact a running workload tricky – how to ensure all pods pick up the change?

Typically, a rollout (restart) is needed, causing a brief disruption or at least requiring coordination. If not handled, some pods might still hold old credentials while others have new ones, potentially causing authentication errors. This orchestration of change is something teams must plan (often by treating secret updates as a deployment). Again, the Reloader open source K8S controller automatically triggers rollouts of workloads whenever referenced Secrets or ConfigMaps are updated.

- **Interfacing External and Legacy Systems:** Many K8s apps still need to consume secrets from or share secrets with systems outside the cluster (e.g. a legacy database on a VM, or a SaaS API key that is also used by non-K8s services). But K8s Secrets are not easily accessible outside the cluster – they are meant for internal consumption. This creates a need to synchronize secrets between Kubernetes and external vaults or vice versa.
- Some admins in the past chose to use duplicate secrets management (store X in Vault for legacy systems and also put X in K8s Secret for cloud-native apps). We recommend that you not do this; if you see this practice in older K8S environments, make plans to get away from it. This sort of

duplication can lead to inconsistencies (one gets updated, the other forgotten). What's more, you will find it quite a challenge to unify secrets management across cluster and non-cluster environments. For this reason, many opt to use external secrets managers even with K8s (using K8s just to distribute the secret locally).


- **Human Error and Exposure:** Kubernetes manifests (YAML files) often end up in Git repositories (for GitOps or just config). If secrets are stored directly in these manifests (even Base64), there's a risk of leakage via source control. It's a common mistake to accidentally commit a secret. K8s doesn't solve that – instead, solutions like Sealed Secrets or SOPS (age or PGP encrypted files) are used to keep secrets encrypted in Git. This is an additional complexity in the SDLC: ensuring developers do not leak secrets when managing K8s configs. It requires processes and tools on top of Kubernetes.

Kubernetes Secrets – Tools & Best Practices

**Use Orchestrator Secrets Features
External Secrets Store With Injection
Encrypt At Rest On Hosts
Twelve-Factor App & Security
Developer Training & Tooling**

To handle Kubernetes secrets safely, open source the community has developed a robust ecosystem of tools and guidelines:

- **ENCRYPT SECRETS AT REST:** The first step is always to enable **etcd** encryption for secrets on the cluster. Kubernetes allows you to configure an encryption provider (often using a KMS plugin or a locally managed key) so that all Secret objects are stored encrypted



in `etcd`. This ensures that if someone gains read access to `etcd` data or a backup, they cannot retrieve plaintext secrets. It closes the glaring hole of the default plaintext storage. Many cloud-managed K8s (like AKS, EKS) offer an option to enable such encryption with a customer-managed key (e.g. AWS KMS key). This is a must in enterprise environments handling prod secrets.

- **TIGHT RBAC AND NAMESPACE ISOLATION:**

Implement least privilege for secret access. Define Kubernetes RBAC roles such that only the specific service account running a pod has access to the Secret objects it needs. For example, if Namespace “payment” has a secret “db-password”, only the deployment in the payment namespace for the DB connector app should be able to get that secret. You can bind a role to that service account allowing you to get on that one secret **only**. Avoid giving broad access like reading all secrets in a namespace unless absolutely necessary. In multi-tenant clusters, use separate namespaces per tenant or team and consider even multi-cluster if hard isolation is needed. Essentially, treat Kubernetes secrets with the same zero-trust mindset: by default no pod should be able to access any secret until explicitly granted. Kubernetes will not do this automatically – it’s on administrators to configure. Regular audits of RBAC policies help catch misconfigurations that could expose secrets.

- **EXTERNAL SECRETS MANAGEMENT:** One of the most popular trends is to integrate external secret managers with Kubernetes. Rather than storing sensitive values directly in K8s Secrets, many use controllers/operators that

fetch secrets from an external vault and inject into Kubernetes on the fly. For example, Kubernetes External Secrets (now CNCF project External Secrets Operator) can sync secrets from AWS Secrets Manager, HashiCorp Vault, Azure Key Vault, etc., into K8s Secrets resources. This means your source of truth remains the external vault, but your pods still get a native K8s Secret to consume. It bridges that gap and prevents duplication.

HashiCorp Vault has a dedicated Kubernetes auth method and injector that can directly render secrets into pods without them ever being stored persistently in `etcd` (the Vault agent runs in-memory).

Tools like *Sealed Secrets* by Bitnami allow you to encrypt a secret with the cluster’s public key and commit the ciphertext to Git; the cluster then decrypts it into a Secret at runtime.

All these solutions aim to augment Kubernetes’ basic secrets with enterprise-grade security and automation. For instance, by using Vault, you can get features like automatic secret rotation and leasing inside Kubernetes which vanilla K8s lacks.

Enterprises often choose a solution like Vault or *Akeyless* or CyberArk Conjur and use their official integrations with Kubernetes to ensure a pod can seamlessly retrieve secrets at startup using its service account identity. This also solves the multi-platform issue: the same vault can serve both K8s and legacy apps, ensuring consistency.

- **SECRETS ROTATION AND RELOAD:** Since Kubernetes won’t rotate secrets by itself, organizations implement patterns to handle this. One approach is using short-lived

secrets via external managers (as discussed, Vault dynamic secrets for K8s apps – e.g. every 30 minutes the pod gets a new DB password). Another is using operators that watch for secret changes and trigger pods to reload. Kubernetes lacks a native secret update propagation, but tools like Reloader or built-in features like mounting secrets as volumes (with subPath) can trigger restarts.

A simpler approach is versioned secrets – e.g., create a new secret with a version number in the name and update deployments to use that, which triggers a rollout. While these require operational discipline, they are necessary to avoid the “secret value never changes” syndrome. At minimum, have a process (and possibly a Jenkins job or Argo CD workflow) to periodically update secrets and redeploy dependent pods.

- **AUDIT AND MONITOR:** Turn on Kubernetes Audit Logs for secret access events. This can log any get or list of Secret objects, providing a trail of which service accounts or users accessed them. Monitor these logs for suspicious access (e.g. a pod that normally doesn’t read certain secrets suddenly does). Kubernetes by itself won’t flag this, but a SIEM or monitoring tool can. Also, keep an eye on network egress – if a compromised pod exfiltrates a secret, network policies might help contain that. While these go beyond pure secrets management, they are part of an overall strategy in K8s to assume a secret might get out and to have layers of defense and detection around it.
- **KUBERNETES SECRET STORE CSI/CSI DRIVERS:** The cloud providers and community have created CSI drivers (Container Storage Interface drivers) that allow mounting

external secrets as files in pods without ever creating a K8s Secret object. For instance, Azure has a Key Vault CSI driver; AWS has Secrets Store CSI. These essentially let the kubelet directly fetch a secret from the external store and present it in a volume. This is a great approach to skip the Kubernetes etcd storage altogether for highly sensitive material. It does, however, require using those specific drivers and might not cover all use cases (some apps expect env vars, not files).

In practice, a common enterprise setup is: Kubernetes for orchestration, but HashiCorp Vault (or cloud secrets manager) as the source of secrets, using a syncing mechanism. This gives the best of both – Kubernetes automates deployment and scheduling, and the external vault ensures robust secret control (encryption, rotation, audit). By addressing Kubernetes’ built-in limitations – lack of encryption by default, no rotation, limited audit, etc. – with these tools, teams can safely manage secrets even in large multi-team clusters.

The critical insight is that solving secrets in Kubernetes requires a holistic approach: relying solely on the native Secret objects without additional guardrails is usually insufficient for enterprise needs. Kubernetes is a powerful fabric, but it needs to be paired with equally powerful secrets management practices to ensure that the convenience of central orchestration doesn’t become a single point of failure for security.



Cross-Environment Considerations – Why One Size Doesn't Fit All

*Trust Boundaries Vary
Identity Models Differ
Ephemeral vs. Static
Mixed Ecosystems, Mixed Support:*

As shown in previous sections, each environment (VMs, SaaS, microservices, containers, K8s) presents distinct conditions for secrets management. It is dangerous to assume that a solution or practice from one domain will directly apply to another. Let's take a closer look at what's behind these differences and why a tailored approach is necessary for each:

- **Trust Boundaries Vary:** In a traditional on-prem VM environment, the primary trust boundary is often the corporate network (inside vs. outside). In contrast, a SaaS multi-tenant app has trust boundaries between each tenant, and a microservices app treats every service-to-service call as crossing a trust boundary (zero trust internally). Thus, the granularity of secrets isolation needed is different. A method that suffices for one big trusted zone (e.g. one vault for the whole data center) might not granularly isolate secrets for many tenants.

You might need per-tenant compartments or multiple vault instances to achieve the same level of isolation SaaS requires. Similarly, Kubernetes introduces trust boundaries at namespace/cluster levels that might not map one-to-one to organizational boundaries, requiring extra care. Bottom line: The scope at which you must enforce “who should see this secret” changes in each context, so your

solution must accommodate the smallest required trust unit (be it a VM, a container, a tenant, or a service).

- **Identity Models Differ:** Different environments authenticate machines and applications in different ways, which affects how they retrieve secrets. On VMs, one might use machine credentials (like Kerberos domain accounts or instance certificates). In cloud VMs or K8s, you have cloud IAM roles or service accounts. In serverless, the platform's identity (Lambda execution role) might be the only identity. A secrets management solution needs to integrate with these identities. For example, HashiCorp Vault can use AWS IAM to authenticate an EC2 instance – great for cloud, useless on pure VMware without AWS.

Conversely, a Windows-oriented solution that relies on Active Directory might not work in a stateless container environment. No single auth method spans all environments seamlessly. This is why enterprises often run multiple secrets backends or at least multiple auth pathways – one size doesn't fit all. An illustration from OWASP: cloud-native teams might use the cloud provider's vault, while private cloud teams use a third-party vault, and both need to be accommodated. The key is to standardize the interfaces as much as possible (so developers have a consistent way to request secrets) even if behind the scenes there are multiple systems bridging different identity models.

- **Ephemeral vs. Static:** Traditional infrastructure assumed servers are relatively long-lived, so manual secret updates were feasible albeit risky. Modern cloud-native infrastructure is highly ephemeral –

instances might last hours or less – so manual secret provisioning is practically impossible. Solutions that require human intervention (e.g. an admin approving access or a manual copy-paste of a secret) simply do not work at scale in ephemeral environments. Instead, ephemeral environments demand dynamic, automated secrets distribution and often prefer short-lived secrets to match (e.g. issuing 15-minute credentials for a function). Conversely, in a static environment, pushing constantly rotating ephemeral creds might be overkill or incompatible with legacy software. Thus the cadence and style of secrets (static vs dynamic, long-term vs short-term) need to match the environment. Many companies discover that the secret rotation policies they can enforce on Kubernetes or cloud (maybe daily rotations) cannot be imposed on a legacy on-prem app that can't reload credentials without a restart. This leads to different policies in different environments – not ideal, but sometimes necessary while transitioning to more dynamic-friendly systems.


- **Automation and Tooling:** The level of automation in deployment also differs. Cloud and containerized deployments are heavily automated via CI/CD – which allows embedding secret management into the pipeline (e.g. pulling secrets at build or deploy). In older environments, deployments might be manual or via older tools, making it harder to insert secrets retrieval steps – thus people might hardcode values for simplicity. Moreover, modern environments have infrastructure-as-code, allowing integration of secrets (like Terraform pulling from Vault). If one tries to use a very manual secret process in a fast-paced DevOps environment,

it will be bypassed or cause failures. Conversely, implementing a full Vault+CI integration for a small static environment could be over-engineering. The solution must align with the operational practices of the environment. In practice, this means using cloud-native solutions for cloud-native apps, and perhaps more human-centric or ITIL-like solutions for legacy – or investing in refactoring legacy processes to adopt automation.

- **Mixed Ecosystems, Mixed Support:** A common security gap emerges when trying to reconcile varied environments each with a different ecosystem of supported tools.
 - Kubernetes, for instance, has dozens of open-source operators for secrets; using those might be obvious in K8s but meaningless if your app is on bare EC2 instances.
 - SaaS providers on AWS might lean on AWS KMS, while an on-prem team might use an HSM appliance.

Organizations must recognize that no single vendor currently provides a one-stop secret solution that natively plugs into everything – there will be a mix. What's important is to set overarching governance (policies for how secrets are handled, regardless of tool) and ensure each team's chosen tool meets those.

For example, your policy might state: all secrets must be encrypted in transit and at rest, must be rotated at least every 90 days, and access must be auditable. How it's implemented can differ: one team uses Azure Key Vault with those settings, another uses Vault with policies, another uses



CyberArk – but can still all adhere to crucial high-level requirements.

In short, assumptions need to be re-validated for each new platform. One company that moved from on-prem to cloud found that where they once relied on network isolation, they now had to encrypt everything and manage IAM roles. Another that adopted Kubernetes found developers assuming K8s Secrets were “secure enough” had to be educated on additional controls.

The lesson is that secrets management must be context-aware: architecture and security teams should explicitly design a secrets strategy for each environment and then integrate them, rather than trying to retrofit one into another. This often results in a hybrid secrets management landscape – e.g., a central Vault that syncs to cloud-specific services, or a federation of secrets managers – which if managed well can still provide a unified security posture. As the OWASP guidance suggests, even if multiple solutions exist, standardizing interactions and having a clear picture of “what is stored where” is crucial to maintain control.

Key Recommendations

Secrets management in distributed cloud computing is a complex, critical discipline that touches all phases of the software lifecycle. From a developer writing code (who must avoid hardcoding secrets) to a CI/CD pipeline deploying an app (which must inject secrets securely) to runtime in production (where systems must retrieve and rotate secrets safely), a holistic strategy is needed. The Minto Pyramid analysis above has illustrated that while the core principles of secrets management – least privilege, secure storage, encryption, rotation, and audit – remain

constant, their implementation varies widely by environment.

For enterprise architects and security leaders, the following key recommendations emerge:

(1) ADOPT A CENTRALIZED (OR FEDERATED) SECRETS MANAGEMENT PROGRAM: Inventory your secrets and avoid ad-hoc approaches. Use enterprise-grade secret managers (Vault, AWS/Azure/GCP managers, etc.) as appropriate, and consolidate where possible. Centralization doesn’t mean a single tool for everything. Your goal is to engineer a single policy framework with possibly multiple integrated tools. Aim for visibility – you should know where every critical secret lives and who can access it, even if it’s across different platforms.

(2) TAILOR SOLUTIONS TO ENVIRONMENT, BUT ENFORCE BASELINE CONTROLS: A classic VM app might stick with a static vault, whereas a cloud-native app uses cloud secrets services – this is fine as long as each meets your baseline: e.g., secrets are encrypted at rest, never stored in code, and rotated regularly. Do not try to force-fit one environment’s tools into another if it creates friction; instead, bridge them. For instance, if developers love AWS Parameter Store for serverless, allow it, but maybe sync those secrets to your on-prem vault for auditing, or vice versa. Integration beats uniformity in many cases, as long as security controls are equivalent.

(3) AUTOMATE EVERYTHING (CI/CD AND RUNTIME): Human error is a top cause of secret leaks. Strive to remove humans from the loop of managing secrets in day-to-day operations. Use pipeline integrations so that when deploying to Kubernetes or VMs, the pipeline fetches the needed secrets from a secure store (never exposing them in logs). Implement

auto-rotation for databases, API tokens, certificates – machines handle it, humans just get notified on failures. This also ensures that across the SDLC, from dev (who might get an ephemeral secret for testing) to prod (which has long-lived but rotated secrets), the process is consistent and reliable. Tools like GitHub Actions secrets, Jenkins credentials, etc., should be linked back to your main secrets store so that there aren't disparate pockets of hidden credentials.

(4) EDUCATE AND ENFORCE POLICIES: Make secrets management a part of the engineering culture. Developers should understand why checking in a password is dangerous, ops engineers should know how to use vault tools, and product managers should recognize the need to invest in this infrastructure. Use commit hooks, code reviews, and secret scanning tools

to prevent mistakes (e.g., prevent that AWS key from ever entering Git). Regularly review access logs and perform secret rotation drills (simulate a leaked secret and rotate it – are all systems still working?). Foster a mindset that credentials are as important as code – they must be handled with care at every step.

(5) PREPARE FOR CROSS-ENVIRONMENT SECRETS

INCIDENTS: Assume that at some point, a secret from one environment will leak (it happens – an intern posts a code snippet with a key, or an extracted container image has a leftover credential). Your response should be swift: knowing where else that secret is used, being able to rotate it everywhere, and updating dependent systems. This is much easier if you have centralized tracking and unified processes. If each environment is a silo, a single secret leak can turn into a protracted scavenger hunt.

Conclusion

Effective secrets management across the SDLC requires a top-down, unified approach, yet one that is tailored to each architectural environment's nuances. Solutions that work in one environment (e.g. a static data center VM) often break down in another (e.g. ephemeral serverless function) due to differences in trust boundaries, identity models, automation, and ephemeral compute. Enterprises must therefore adopt both centralized best practices and environment-specific tools to keep secrets safe throughout development, deployment, and operations.

Secrets management is difficult precisely because it underpins all layers of modern infrastructure – but it is an essential foundation for cloud security. Breaches like source code leaks with hardcoded keys, or public GitHub repos exposing API tokens, have taught the industry hard lessons. The path forward is clear: treat secrets as first-class assets, use robust tools fit for each environment, and architect your systems (from development practices to runtime platforms) to minimize secret exposure. Organizations that do this successfully gain not only improved security but also operational benefits – easy key rotations, flexible integrations, and confidence to adopt new cloud technologies without introducing new secret-related risks.

In a distributed cloud era, secrets management done right is a crucial business enabler, allowing teams to move fast and stay secure. The challenge is non-trivial, but with the strategies outlined above, enterprises can master it across the entire software development lifecycle and cloud stack.

SOURCES

1. [What Is Secrets Management? – Palo Alto Networks](#)
2. [OWASP Secrets Management Cheat Sheet](#)
3. [Unified Secret Management Across Cloud Platforms: A Strategy for Secure Credential Storage and Access – ResearchGate](#)
4. [Secrets Management in Cloud – Medium \(K. Bangalore\)](#)
5. [OWASP Secrets Management Cheat Sheet – Multi-Cloud Environment Challenges](#)
6. [AWS Secrets Manager – Service Overview](#)
7. [Azure Key Vault – Microsoft Azure Documentation](#)
8. [Secret Manager Overview – Google Cloud Docs](#)
9. [Comparative Analysis of Native Secrets Management Services – Online Scientific Research](#)
10. [Unlocking Multi-Cloud with HashiCorp Vault Secrets Management – CloudOps \(Medium\)](#)
11. [Conjur Open Source – Secrets Management \(CyberArk\)](#)
12. [Secrets Management Best Practices – Palo Alto Network](#)

About the Author



Serg Shalavin, Lead Cloud DevOps Architect at CloudGeometry, has over 20 years of experience in hard core and on systems engineering, infrastructure, AWS DevOps, cloud computing, and release management. A certified Kubestronaut, he holds credentials including CKA, CKS, KCSA, and 3x AWS Professional certifications (DevOps Engineer, Networking and Solutions Architect). Serg is known for his disciplined approach to operational processes, customer-focused mindset, and ability to bridge the gap between developers and business needs. At CloudGeometry, he leads enterprise-grade Kubernetes adoption, secure cloud modernization, and AI/ML infrastructure integration across AWS, GCP, and Azure for high-growth tech organizations.

CloudGeometry Insights is our applied research initiative where architecture decisions, security tradeoffs, and emerging tech are dissected by our engineers who *live* in production. This isn't editorial. It's strategic intelligence from practitioners shipping real systems at scale – the kind of thinking we bring to client roadmaps: shared openly, challenge-ready, and rooted in field experience. If you're a CTO, CIO or founder trying to separate signal from noise, this is where we publish what's worth your time.

Want to hear from our technical experts on a topic of particular interest? Contact us via insights@cloudgeometry, and let us know what you want to learn more about. If we've already written on the subject, we will send it on; if not, we will let you know and add it to our research agenda for future publication.

ABOUT CLOUDGEOMETRY

CloudGeometry delivers expert technical services, helping our clients unlock the full potential of cloud-native open-source tooling and commercial platform technologies. As an **AWS Advanced Consulting** partner, our certified solution architects and platform engineers help address the range of challenges facing enterprise innovators and venture-funded startups alike. The Cloud Native Computing Foundation has accredited us as a **Kubernetes Certified Service Provider**.

- [Foundation Services](#) solidify, simplify, and modernize your tech foundations: healthier systems, streamlined migrations, and robust uptime for future-ready growth.
- Advanced Services provides [software development teams](#) with proven strategies and tactics for better technical & commercial outcomes for a clear path to cloud-native velocity
- AI/ML & Data Services help you apply the latest in [AI and data analytics technology to your crucial business processes](#): MLOps, Generative AI, AI/ML Development. As a **Databricks Systems Integrator**, our data engineering experts ensure your data is accurate, accessible, and ready to fuel decision-making.

For the last decade, we've built and deployed hundreds of big, fast apps backed by scale-out infrastructure across a range of industries. Working across [a wide variety of applications](#) – Financial Services, Industrial Automation, HIPAA-compliant Healthcare, AdTech, Consumer-grade Mobile, IoT and smart devices, among others – has given us valuable insights into essential full-stack application patterns.

Learning from these similarities lets us develop adaptive, flexible, technology-driven solutions even faster. It's why we created **CGDevX**, our free, open-source hybrid-cloud application delivery platform. By automating many of the routine tasks associated with software development, CGDevX gives you a shorter, more effective path for engineering teams to focus on creating value-added features and innovations.

Count on us to accelerate application modernization, Kubernetes adoption, developer enablement, secure multi-tenancy, DevOps automation and more. With roots in Silicon Valley, we've seen firsthand what works (and what doesn't). We've completed hundreds of migrations and application modernization projects, so we're well-versed in solving for rising costs and evolving technology and helping clients choose solutions that prevent needless vendor lock-in.

We'll help you find the right technology for your business goals and budget, and we'll integrate the right tools with your Software Development and IT processes, all backed by a complete portfolio of ongoing application development and full-stack support services. From enterprise upgrades to cloud-native scale-out to practical applications of AI, CloudGeometry helps plot the shortest path across all dimensions of modern cloud software and data engineering.

