A Deep-Dive into the Aembit Workload IAM Platform

This action-oriented guide covers all the essentials on securing and streamlining your workload connections with Aembit.

aembit

Gartner
COOL VENDOR 2022

Introduction: A Unified Approach to Workload Identity and Access Management

Aembit has built the industry's first Workload IAM Platform to safely allow your applications to connect to any services, applications, or APIs they need. Based on identities instead of secrets, we give you a simple, secure, and central way to provide policy-based access, instead of needing to manage low-level secrets.

Essentially, you can think of us as Okta, but for workloads. Instead of user-to-workload access, Aembit manages workload-to-workload access.

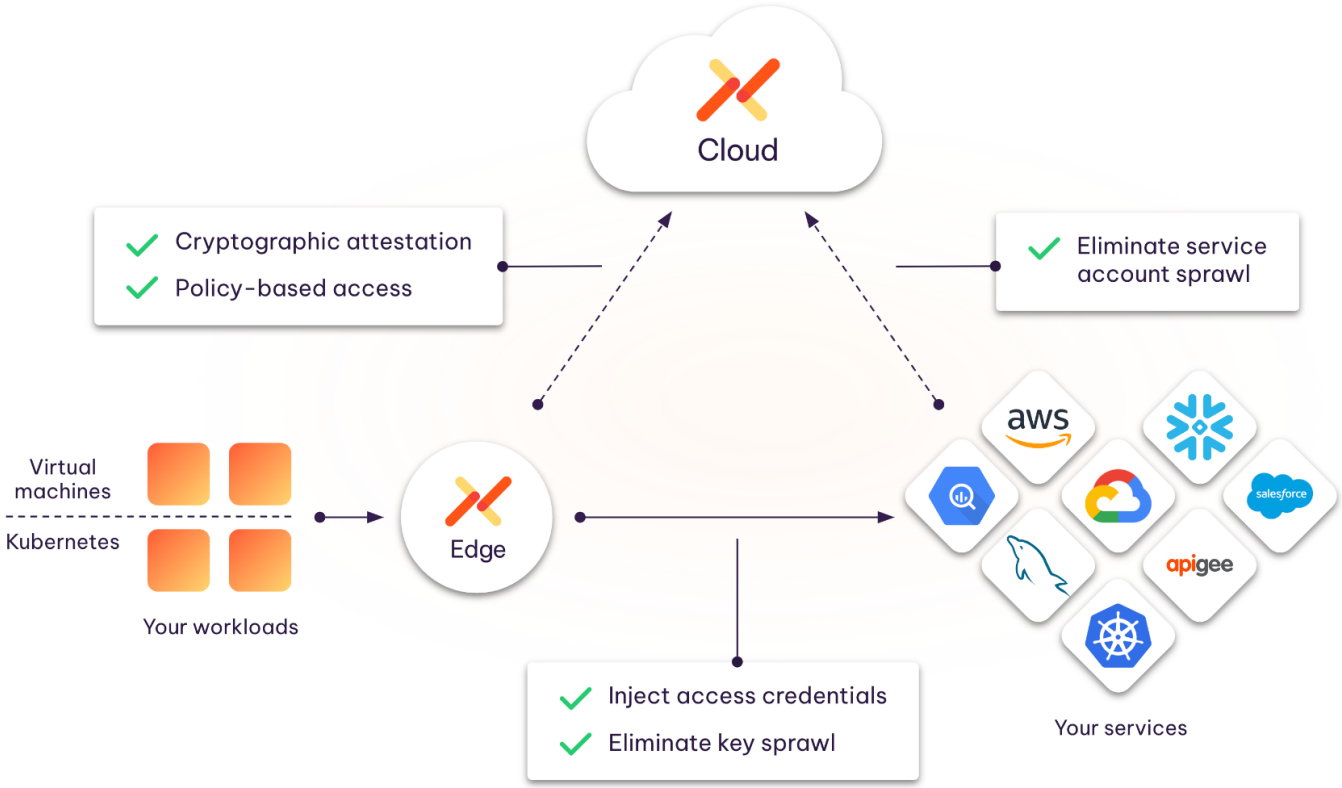

With Aembit, organizations can implement identity-first security for their workloads in the most complex of situations: where applications and services need to reach across boundaries and establish trust through identity. With Aembit, enterprises help their organizations in three key ways:

1

Stop Workload Attacks

Defend against common workload communication threats, including credential exposure, unauthorized access, lack of key rotation, service account sprawl, and weak or misconfigured authentication, to ensure robust security in your distributed application infrastructure.

2

Build Great Products, Instead of DIY Authorization

Aembit saves hundreds to thousands of developer hours by automating the heavy lifting of workload IAM, while ensuring a consistent auth implementation across your entire environment.

3

Scale DevSecOps

Save your DevSecOps teams 50% or more of the time they currently spend on workload-to-workload security. Automate and optimize activities, such as credential rotation, compliance, and audits. Aembit gives you a single place to operate, control, and analyze workload-to-workload access.

Enterprise Context and Today's Approaches

Modern applications are highly distributed, with applications or microservices leveraging many different resources to source data and delegate operations.

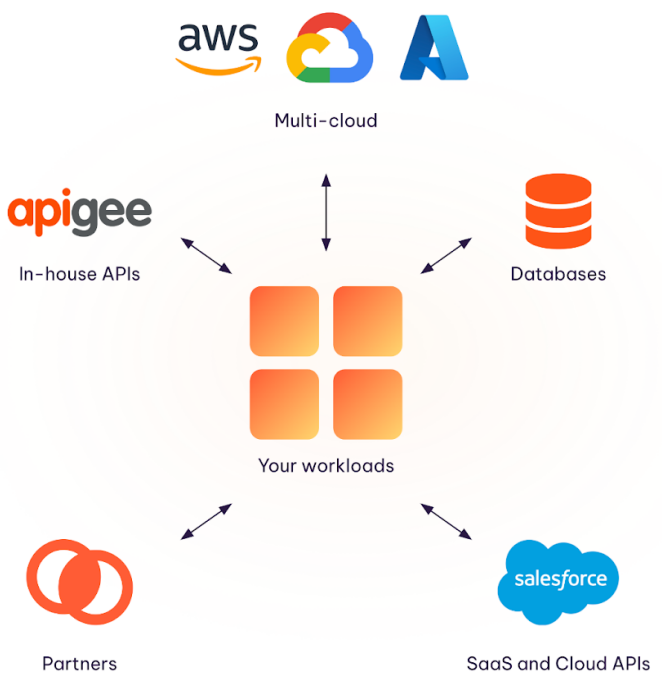

For example, you might be using services from multiple clouds, databases that are distributed in multiple locations, APIs from your own applications or other services, SaaS applications (like Salesforce or Snowflake), or even partner applications.

Against this backdrop, you can see an increasing attack surface with every new connection your application makes. Operational complexity is increasing for your teams to manage it all – conducting a regular audit or performing credential rotation, for example, can become a major project.

And secrets, the method used to connect applications today, are sprawled across applications, whether they are hard-coded or managed on a team-by-team basis.

We generally see three methods of managing access today:

Cloud IAM	This works well for controlling access to native services at a single cloud provider, but breaks down once you start crossing boundaries into other clouds, SaaS services, or even your own on-prem software.
Vaults	They are good for storing secrets, but using them implicitly assumes that you understand the identity of the client requesting access. Further, they typically require effort and maintenance from your dev team to integrate into your applications.
DIY	This is actually what we see most teams doing today. Do-it-yourself approaches generally include a set of manual configurations, best practices, and 'allow' lists that don't scale well and can easily fall apart during an incident.

The challenges for enterprises typically are seen in two ways:

1 **Providing an identity layer that works across environments is particularly cumbersome.** If you are operating in one environment (for example, wholly in AWS or using identities within a confined Kubernetes cluster), there are native options that might allow you to achieve your goals. If, however, you are trying to cross boundaries (i.e. AWS to GCP, your cloud to SaaS, your on-prem environment to a partner's cloud), access management tools break down.

2 **Enterprises would prefer to have a workload IAM platform that works for their range of environments.** Even if you cleanly separate operations in your GCP and on-premises environments, for example, you still need separate ways of managing identity and access. This is operationally burdensome for your security and DevOps teams, leading to potential data protection gaps, as well as tool fatigue and reduced responsiveness.

Targeted Benefits of Workload IAM

At a broad level, the workload IAM architecture looks somewhat similar to your user IAM strategy: a third-party, independent broker that can validate identities and evaluate access policies.

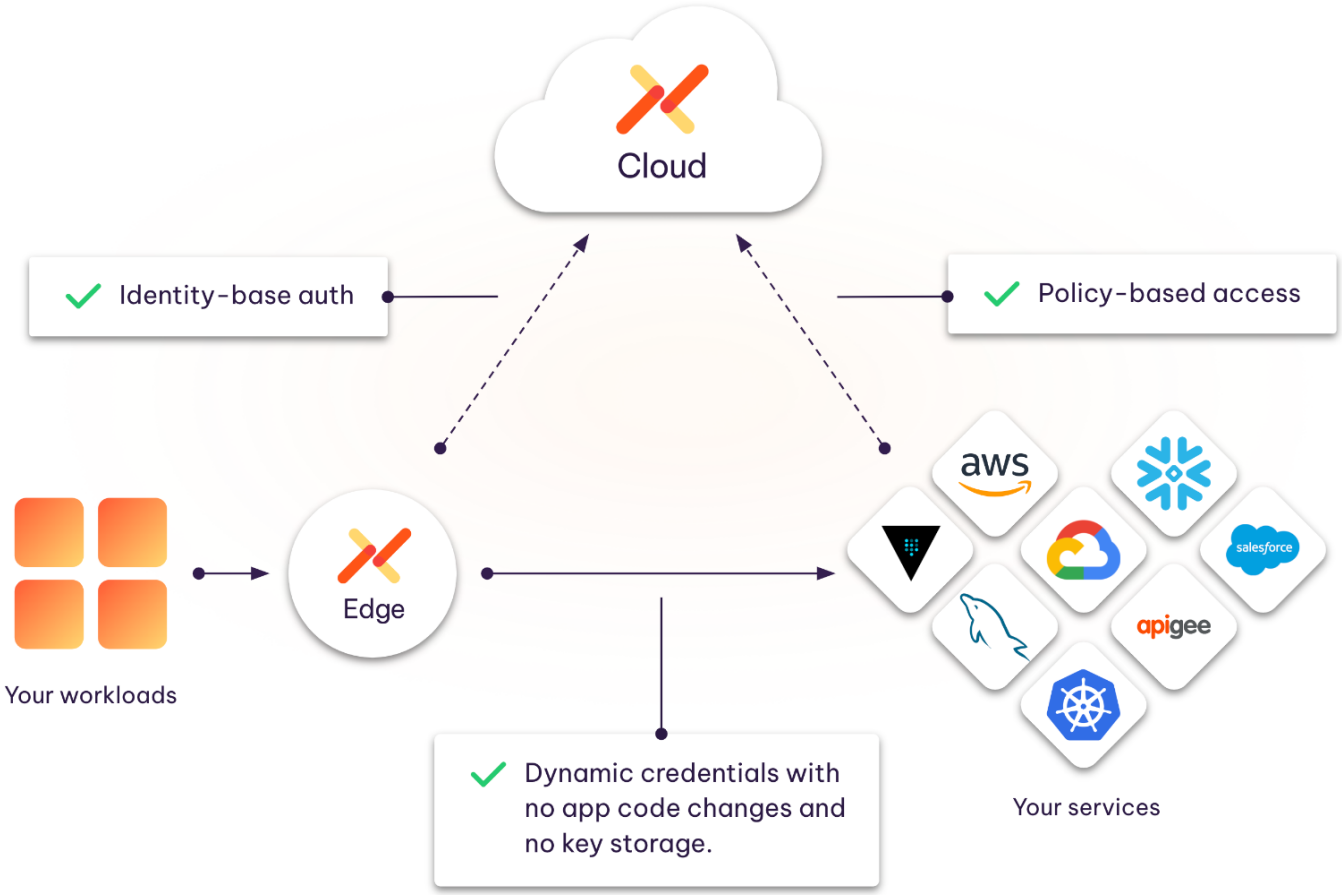

Naturally, with workload IAM the workflows need to be different to work with what machines understand and can do to validate themselves. We'll get into that in more detail.

Workload IAM sits in the area between security, DevOps, and developers, and as a result provides, benefits to each group:

Security

- ✓ Enforce identity-based access policies centrally.
- ✓ Perform identity-based logging for faster incident response, audit, compliance.

DevOps

- ✓ Limit service account exposure and eliminate secret sprawl.
- ✓ Automate credential rotation.

Developers

- ✓ Avoid the need for code auth and maintain consistent auth implementation everywhere.
- ✓ Move to dynamic credentials with no app changes.

Aembit Architecture and Approach

We will begin with a high-level overview of the architecture. Following this, we will examine a detailed flowchart that illustrates the interaction between services during a request. Together, these elements provide a comprehensive understanding of how Aembit functions within your infrastructure.

Aembit is designed to be your workload identity provider (IdP). We broker between workload identity and service credentials, using policies to control access. We do this across trust boundaries, giving you a consistent implementation everywhere you operate.

Essentially, we federate with both the workload and the service environments to make this happen. This works across the array of environments and services your business operates in today, with flexibility as you adapt and change.

> **“Workload IAM sits in the area between security, DevOps, and developers, and as a result provides benefits to each group.**

Aembit Architecture

- 2-sided federation with short-lived credentials
- Centralized auth, policy-based access, and auditing
- Control-plane (not data plane) architecture
- No-code auth (proxy design)

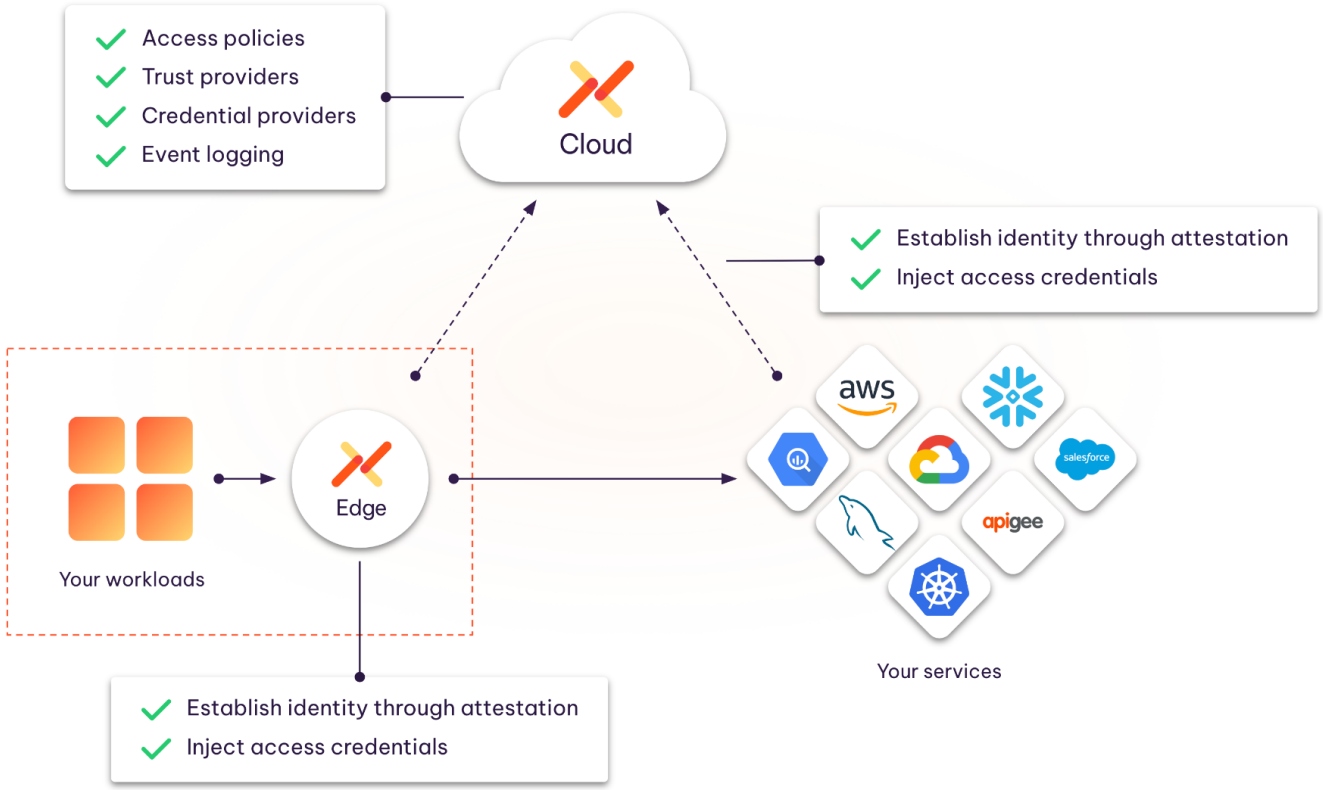

Imagine your workloads on the left – custom code that you write, or packaged applications that you run. They can be in the cloud or on-prem, running in Kubernetes or in VMs for example. On the right are services you access – APIs, SaaS applications, apps run by your business partners, or additional cloud services.

Up top, you see Aembit Cloud. Aembit Cloud is our SaaS service that provides a centralized approach to workload identity management, workload access policies, visibility, and logging. It is operated by Aembit as a SaaS platform. We built this with a control-plane architecture in mind: Aembit manages access, but doesn't see your sensitive application data.

On the left, Aembit Cloud federates with your workload environment through trust providers, which allow us to cryptographically validate the workload's identity.

We do this by leveraging the Aembit Edge. Aembit Edge is a transparent proxy that you deploy alongside each of your applications for which you'd like to manage workload identity. In Kubernetes environments, Aembit Edge is deployed as a sidecar, whereas in VM setups, it functions as an agent. Aembit Edge both works to validate the identity of the workload and inject credentials on behalf of it.

Aembit Edge is what makes no-code auth possible. Instead of requiring your developers to modify code in their application, Aembit Edge intercepts authorization requests and injects credentials into validated requests. That means we'll work well for greenfield but also existing apps. We just slot in and can manage credentials without app changes.

The other big benefit here is that you reduce key sprawl. Workloads don't even need to store keys as Aembit Edge takes care of this dynamically for the app.

On the service side, Aembit Cloud federates with services through credential providers, which allow us to be trusted by the services you rely on to issue credentials that the client can use to access the service.

Finally, Aembit Cloud provides a structured logging approach that allows you to see and analyze access based on identity. This can be viewed in our console or sent to your SIEM for further correlation or alerting.

This architecture gives you a single workload IAM platform across your environment today – and in the future as you add clouds, services, or workloads.

> **Aembit Edge makes no-code auth possible.**

Detailed Flow of a Workload Access Request

Let's walk through an example of a workload making a request.

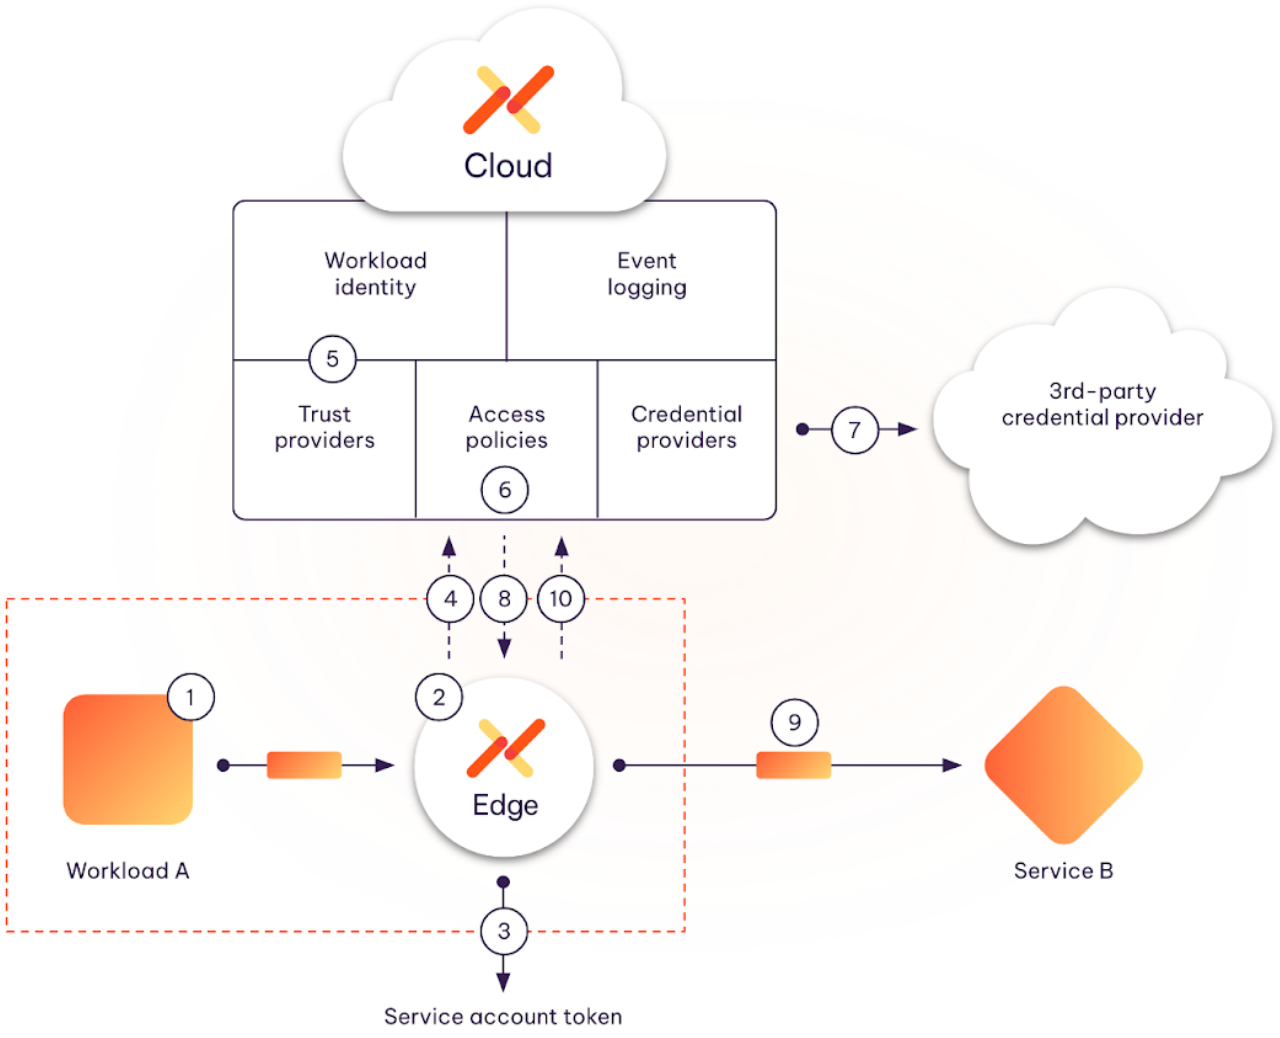

1. Client workload makes a request to service.

Custom or packaged software that you operate makes a request to a SaaS service, API, database, or another workload. It makes the request as if it doesn't exist – the workload doesn't need to know anything about Aembit.

2. Aembit Edge intercepts client request.

The Aembit Edge is a transparent forward prLS, ensuring that you can continue best practices and will work for the range of your applications.

As a reminder, Aembit Edge is a multi-protocol proxy, deployed as a sidecar or agent, whichinjects service credentials without workload code changes. It performs TLS decryption for HTTPS and more, so it will work for much more than HTTP-based apps.

3. Aembit Edge retrieves service account token for secretless auth.

Aembit Edge extracts signed metadata directly from the workload's environment, serving as a robust attestation to the workload's identity. This approach eliminates the requirement for an identity secret by substituting it with secure attestation.

4. Aembit Edge requests access a credential on behalf of the client.

Aembit Edge sends the request and the attestation metadata to Aembit Cloud.

5. Aembit Cloud authenticates client using attestation.

Using a trust provider that was configured to trust the workload's environment, Aembit Cloud cryptographically validates the workload's identity.

6. Aembit Cloud checks access policy.

Policies are defined in the cloud by your team to manage access between workloads. As described in the next section, this is also where conditional access policies are evaluated.

7. Aembit Cloud requests access credential from Credential Provider

A credential provider may be the service itself (i.e., we communicate directly with a SaaS service to provide a credential) or it may be a trusted provider who provides a credential on behalf of the application (e.g., vault, or an IAM system).

Depending on the application and use case, Aembit can act as the credential provider as well, further simplifying the environment.

Aembit uses credential providers to generate credentials (short or long lived) for many different services.

8. Aembit Cloud responds with policy and access credential.

Aembit passes the shortest-lived credentials the service supports back to Aembit Edge.

This is particularly valuable because you can move the security posture of the communication to a shorter-lived token without asking developers to make any changes to apps.

If you are using Aembit with a pre-packaged application, you may not have the option to change the secret type within the application. Aembit manages this transition without changes to the application.

> **Injecting credentials via Aembit Edge reduces work related to credential rotation up to 95%.**

9. Aembit Edge injects credentials into client request and forwards it to the service.

The workload itself never sees or stores the credential. This significantly reduces key sprawl, as workloads no longer need to know about the credential. Additionally, credential rotation is dramatically simplified. While any service that depends on long-lived credentials still needs to be manually rotated, no downstream activity must happen in workloads requesting access.

This reduces work related to credential rotation up to 95%.

After this, communications continue between workload and service as they normally would.

10. Aembit Edge sends access event log to Aembit Cloud.

The access request and access metadata are logged for analytics, auditing, and compliance. You can easily send these into a data warehouse or SIEM.

Extending Aembit to Provide Zero Trust for Workloads

Enterprises are increasingly looking to implement a Zero Trust architecture for workload-to-workload access, much like they have already moved or are moving to Zero Trust Network Access (ZTNA) for user access to applications.

Aembit provides the most foundational elements for Zero Trust for workloads:

Identity + policy-based access to services. In addition, Aembit enables conditional access based on workload posture.

Not only can you use Aembit to validate the identity of an application, you can also use contextual information to determine the current state of the machine and provide conditional access based on a range of conditions.

These conditions are easily defined within an access policy.

The diagram below shows how Aembit implements Zero Trust with CrowdStrike, as an example.

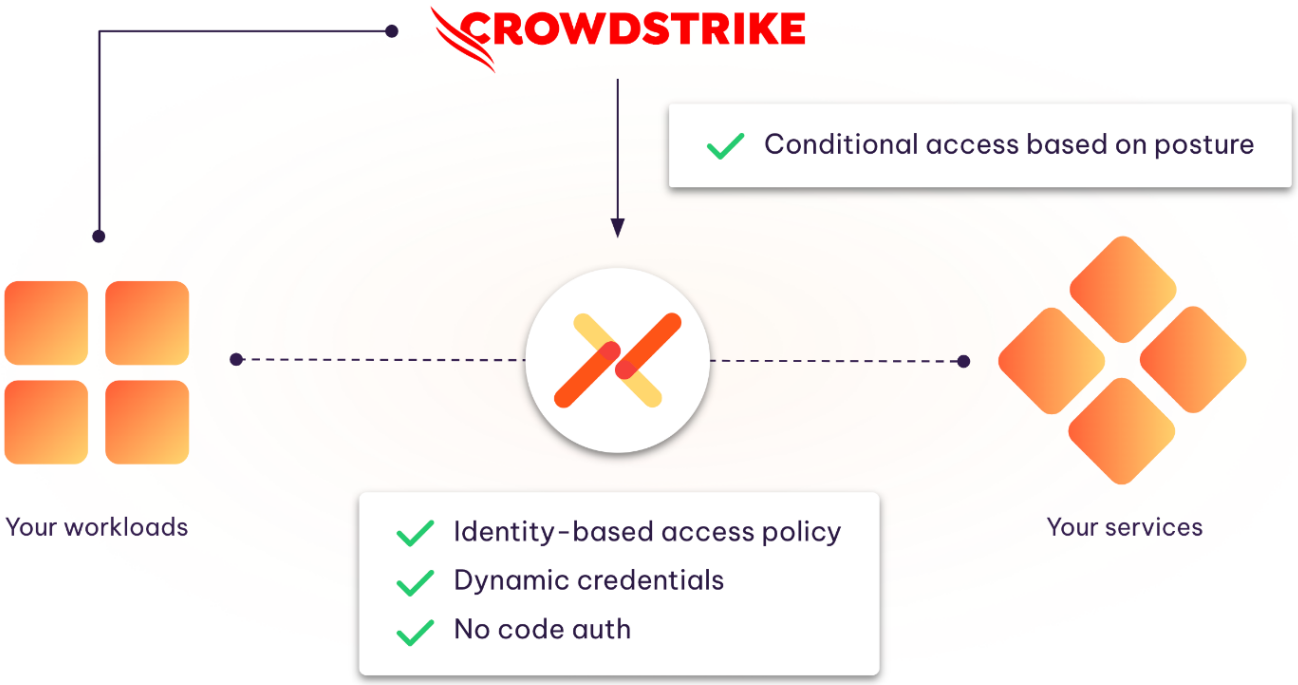

Most Common Use Cases

Enterprises generally are taking a step-wise approach to implementing workload IAM, focusing on particular environments where they believe the need for the most amount of control immediately – and then expanding to include more resources to create a unified management environment.

We most frequently see these drivers:

SaaS & API Access and Governance

Control access to third-party SaaS services in a policy-based manner. (e.g., Salesforce, MSFT Graph, Stripe and more)

Secure Data Lake and Database Access

Policy-driven access based on workload identity, consistently applied across legacy and modern environments. Applies across multiple clouds.

Protect Critical Infrastructure

Secure access to critical components like CI/CD, secrets vaults, and message buses or queues.

Zero Trust for Workloads

Allow for conditional access by checking workloads for posture conditions before access. (i.e., CrowdStrike installed with checks passing)

Typical Deployment, Packaging and Benefits

Aembit is designed as a product that you could easily self-deploy (**sign up on our website and begin yourself**).

You or your colleagues need to be able to deploy Aembit Edge within your infrastructure or lab and then set up policies in Aembit via UI or API. Typically we see customers are able to deploy initial scenarios within a day.

As a way to begin rollout, we also encourage customers to focus on a particular application or set of applications that have sensitive data.

Typical examples include:

- Sensitive data such as a database or data warehouse.
- A particular SaaS service you'd like to begin to control, such as Salesforce.
- Strategic infrastructure, such as HashiCorp Vault or a message bus.

Our model is designed to allow you fast, cost efficient deployment

- Make it frictionless to start. Our free tier offers up to 10 workloads at no cost, with production-class performance.
- The price is based on workloads so that costs are visible and predictable without complex per-request charges.

Based on this, we are targeting a six-month return on investment for paying customers. We know it will vary depending on your environment and use cases.

We see our customers deriving benefits in three areas

1. More robust security.

2. Simpler operations, which include less manual labor and reduced toolset.

3. Reduced operational load on your dev team, freeing them up to focus more on shipping your revenue-generating products than worrying about your security infrastructure.

Summary

Workload IAM is a powerful way to secure your application-to-application communications, especially where applications need to communicate across trust domains to accomplish their tasks.

Leveraging the Aembit Workload IAM Platform can simultaneously help you prevent credential loss, accelerate your development cycles, and eliminate the manual burdens on DevSecOps teams, including credential rotation, credential tracking, audits, and compliance.

We encourage you to **start today with our forever-free tier**, enabling you to secure up to 10 workloads at any scale.

aembit

Gartner COOL VENDOR 2022