

13 Must-Have Features to Adopting a Zero-Trust Model

The best way to approach service mesh security is to adopt a zero-trust model, which means that every connection, no matter its origin, must be validated and secured. To help you evaluate service mesh technologies and implement zero-trust security, we present our **13 must-have features** to keep your application connections safe. Here's the checklist:



1 Transport Layer Security (TLS & mTLS)

Provides end-to-end encryption to protect data in motion between any pair of end points. It's perhaps the most fundamental component, but surprisingly not all service meshes fully support mutual TLS.



2 Authorization

For example with Open Policy Agent (OPA), which defines service API policies as code. Authorization is the flip side of authentication and controls who has access to what resources once you've verified their identity.

3 Certificate Management

Controls and executes SSL certificates from a centralized platform to authenticate connections. Certificate rotation can be a painful administrative step that should be accounted for gracefully. This should be extensible to support external authorities, meaning it'll work with the enterprise identity and access management solutions you already use.

4 Federated Role-based Access Control (RBAC) and Delegation

Grants permissions to users appropriate to their responsibility and, again, applies this consistently everywhere. These controls can be applied at different levels for operators managing the service mesh and also for developers building applications that run in the mesh.

5 Federated Trust Domains

Safely authenticate users and applications across environments, extending the authentication policies consistently everywhere. Without this, you'll spend a lot of effort trying to keep various roles updated and in sync — and probably make some mistakes.

6 Federal Information Processing Standards (FIPS) 140-2

Means your service mesh technology has been validated to meet specific strict security standards as set out by the United States government. There are many government regulations and industry best practices, but FIPS is one common way to baseline your security.



7 Multitenancy and Isolation

Lets users and applications in service meshes share resources securely. Once you have RBAC, you can safely define who can touch what and effectively create isolated workspaces for different roles. Istio's Authorization Policy can also be used to prevent unwanted traffic from reaching your application.

8 Built-In Web Application Firewalls (WAF)

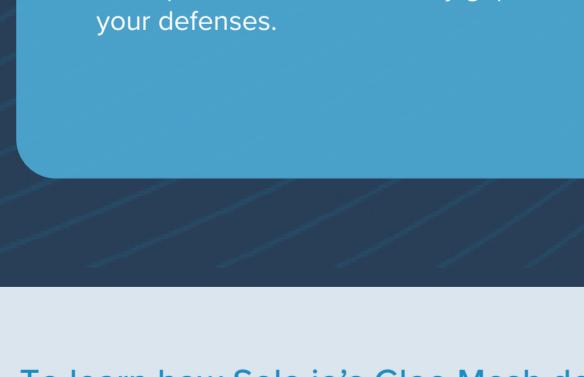
Screen inbound traffic for threats and stop attacks from penetrating your perimeter. It's essential for any edge gateways that are exposed to the internet for incoming user and application connection requests.

9 Data Loss Prevention (DLP)

Monitors for data breaches or exfiltration to prevent data loss and data leaks. If your applications are somehow compromised, you don't want data escaping your perimeter.

10 Integration With Secrets Management

For Kubernetes, which manages sensitive credentials like passwords, security tokens and encryption keys. You'd be alarmed at how often this information is still hard-coded into applications or stored in plain text.



11 Multicloud Access Observability

Provides complete log aggregation and auditability of all activity across the system. This can be useful both for real-time monitoring and forensics after an incident. With distributed applications, it's necessary to get a holistic view. Many use open source tools like Prometheus and Grafana for observability.

12 Vulnerability Scanning and Publications

Find, address and alert on any weaknesses in the system. Security is only as good as its weakest link, so it's important to check for any gaps in your defenses.

13 A Secure Pull Model for Cluster Relay

Safely shares configurations across the system. This is pretty nuanced, but you want to make sure that any configuration changes are distributed to the edge when requested, and only when requested.