

Architecture

Where the edge sits and what it hands off to. The WAF is one boundary in a layered defense, and understanding what it does and does not own makes the configuration pages make sense.

Where the WAF sits

The WAF host lives at the edge of the private network. It is conceptually the boundary between the internet and everything internal, but it is not itself outside; it sits inside the private cloud, and a cloud-level firewall in front of it exposes a single port to the world.

The request path for user traffic is a sequence of narrowing gates:

1. **Cloud perimeter firewall** allows exactly one inbound path from the internet, HTTPS to the WAF host. Nothing else public is reachable.
2. **The WAF host** terminates public TLS, inspects the request, and, if it passes, forwards it to the backend.
3. **Service firewall** allows the backends to be reached only from the WAF, on the internal side. A backend will not accept a connection from anywhere else.
4. **The backend** serves the request, having never been exposed to the internet at any point.

Every public request runs that gauntlet. There is no path from the internet to a backend that does not pass through the WAF, which is the purpose of the whole design.

What terminates here, and what begins

Two different kinds of TLS meet at the WAF, and keeping them distinct is essential:

- **Public TLS terminates at the WAF.** The certificate a browser sees is a public certificate (from a public CA such as Let's Encrypt) for the site's public name, and it lives on the WAF. Encryption between the user and the environment ends here.
- **Internal mutual TLS (mTLS) begins at the WAF.** The connection from the WAF inward to a backend is a separate, mutually-authenticated TLS connection using private certificates from the internal certificate authority. The WAF proves its identity to the

backend with a client certificate, and the backend proves its identity to the WAF with a server certificate.

So the WAF is the seam between two trust domains, the public web on one side and the internally-authenticated network on the other. It is the one component that speaks both. The details of the internal mTLS side, the private CA, the certificates, the enforcement, are the subject of the [mTLS with a Private CA](#)

A note on the single edge

Because the architecture funnels all user traffic through one WAF host, that host's availability is the availability of every public service. As stated in the concepts and design pages, this is an accepted tradeoff for a small ecosystem and would be replaced by a high-availability tier in any larger environment. Architecturally, the thing to understand is that the WAF is not just a security boundary but a throughput and availability chokepoint, and every property of the environment that depends on "the edge is up" depends on this single host in this design.

Revision #3

Created 2026-07-27 17:10:42 UTC by Chris Landis

Updated 2026-07-27 17:47:30 UTC by Chris Landis