Secure clientless access with SSE: A technical deep dive into reverse proxy based ZTNA

Delivering Zero Trust, browser-only access to internal apps without agents, VPNs, or network exposure

Clientless reverse proxy introduction

Clientless reverse proxy is a security feature of SSE (Security Service Edge) gateways that enables secure, browser-based access to internal applications without requiring installation of endpoint software such as a SASE client. It is particularly well-suited to unmanaged/BYOD devices, third-party users, and ad-hoc access where installing an agent is impractical.

(Although Versa also supports clientless reverse proxy for SaaS applications such as dropbox.com, these are outside the scope of this document).

Clientless reverse proxy architecture overview

Clientless reverse proxy acts as an intermediary between external users and internal web applications. The user interacts only with the browser; the SSE gateway terminates the external TLS session, authenticates (via the enterprise IDP) and authorizes the user, and then establishes a separate TLS session to the backend application. All traffic passes through the gateway, where security controls are enforced:

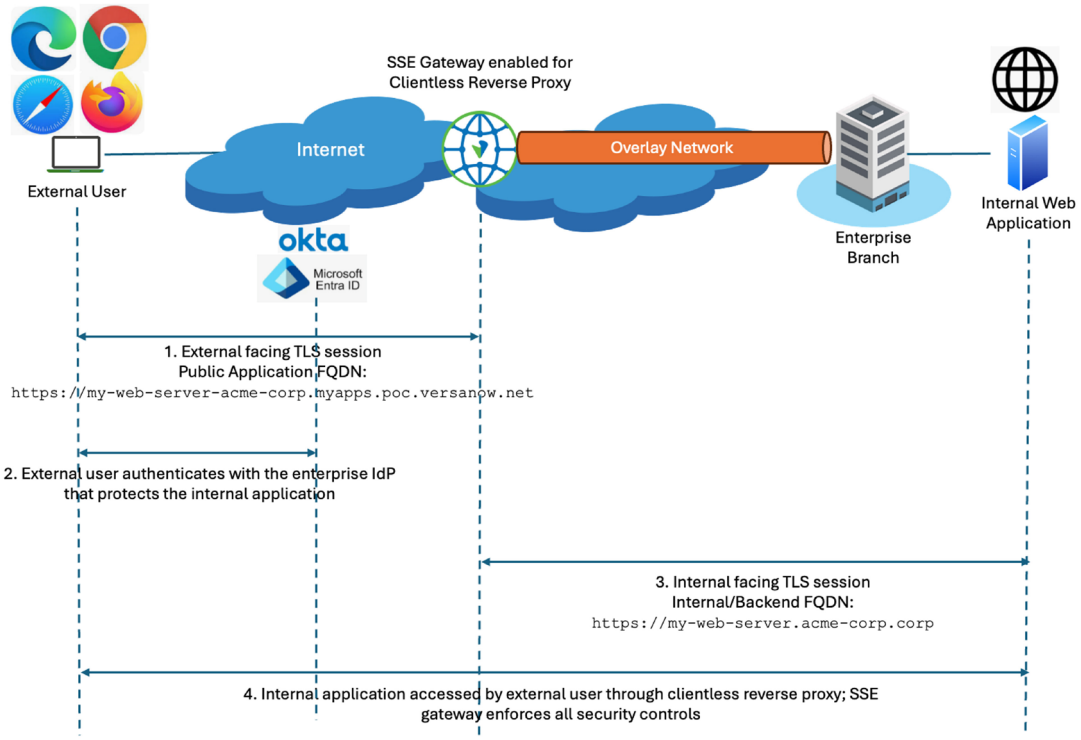

Clientless reverse proxy security model and Zero Trust controls

Clientless reverse proxy is a key enabler for Zero Trust Network Access (ZTNA):

Identity-Driven access

- User identity is evaluated per session against ZTNA policies before any connection to the origin application is established
- Conditional access policies can restrict access to browser-only or limit originating GEO locations/source IP addresses thereby reducing the blast radius compared with broad, network-level VPN access

Application isolation and attack surface reduction

- Internal applications are never exposed directly to the internet; only the SSE gateway's IP addresses and FQDNs are reachable
- Granular controls (URL and risk categories, DLP, threat and malware inspection) are applied at the application layer

No endpoint agent requirement

- For clientless/BYOD scenarios, all security inspection and enforcement occur in-path on the SSE gateway, avoiding the need to deploy, maintain, or support a SASE client on end-user devices

Enforced single entry and inspection point

- Because all traffic must pass through the SSE gateway, there is one controlled point for TLS termination, DLP, and threat prevention
- This prevents "backdoor" paths where traffic might bypass inspection by reaching internal addresses directly

Logging and forensics

- Because all requests traverse the SSE gateway, the enterprise gains a single, consistent source of HTTP transaction logs, DLP events, URL categories, and user identity mappings. This simplifies auditing, threat hunting, and incident response

Clientless reverse proxy end-user flow

1. A user without a native client agent opens a browser and connects over HTTPS to an external-facing FQDN that represents the protected internal application. This FQDN resolves to the SSE gateway's public reverse-proxy IP address rather than directly to the application
2. The TLS session from the user is terminated on the SSE gateway and the user is redirected to a captive portal hosted on the gateway. This portal enforces user identification and authentication:

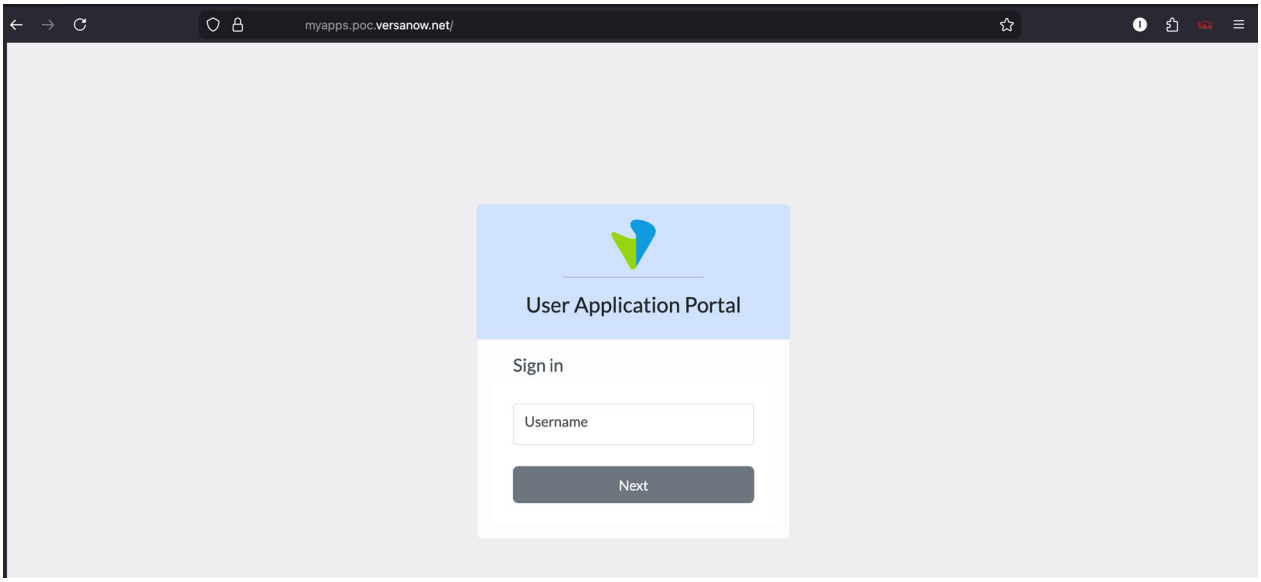

3. The username and password are authenticated using the method configured by the enterprise administrator who own the application being protected - for example, SAML via an IdP such as Okta or Microsoft Entra ID
4. Once successfully authenticated, the external user is presented with a portal listing all applications they are authorized to access via the clientless reverse proxy service:

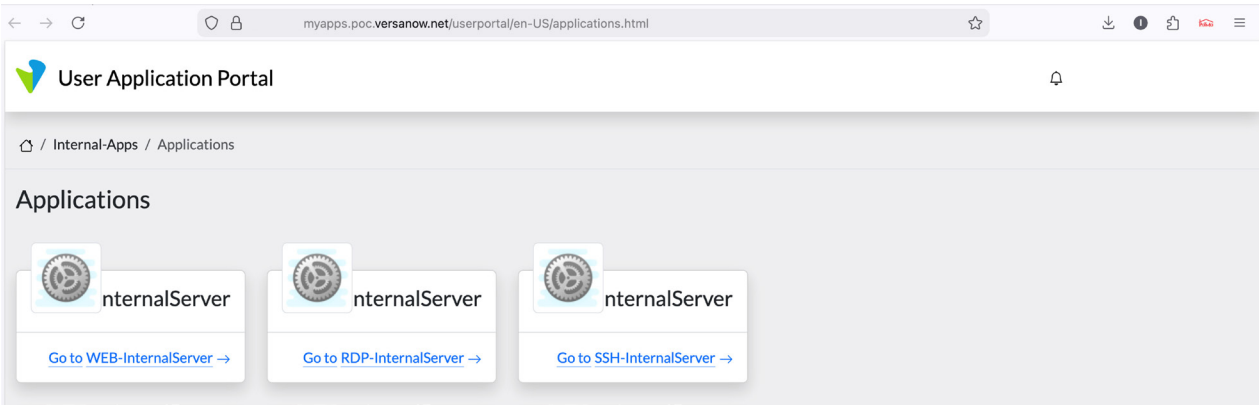

5. By clicking an application tile, the user initiates a connection to the chosen protected application – in this example, a private web server:

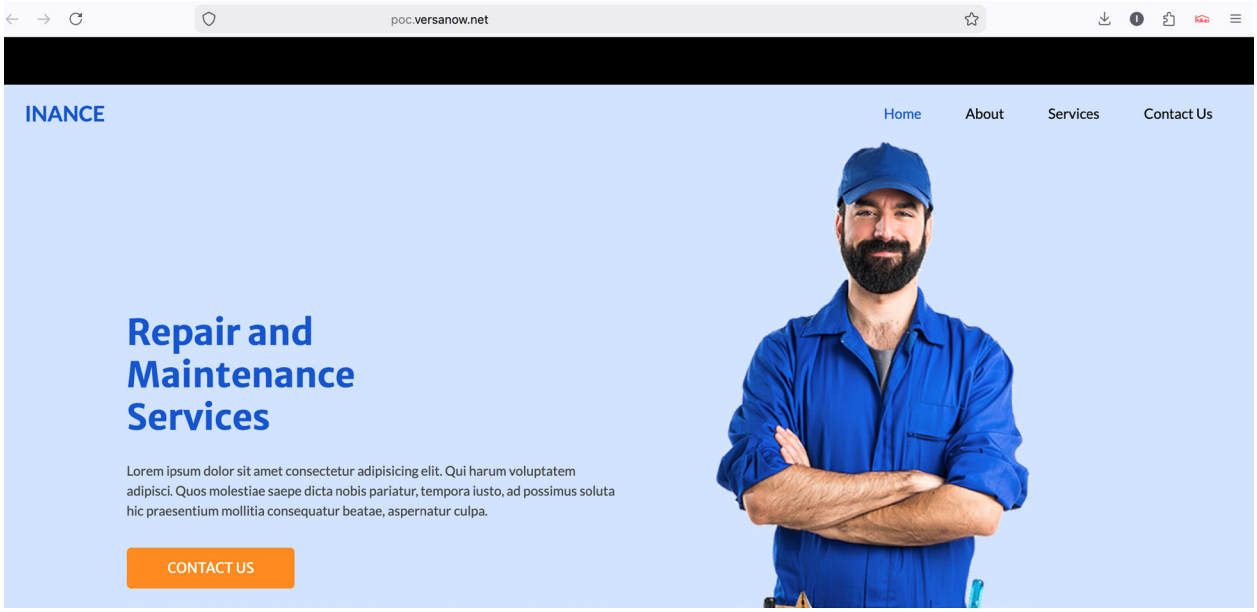

6. From the user's perspective, this appears as a normal website. However, all traffic is flowing through the SSE enforcement point, where URL filtering, DLP, malware analysis, and other security policies can be applied in real time

In this model, the SSE gateway:

- Acts as an HTTPS reverse proxy between the external user and the internal web service
- Terminates and reestablishes TLS, effectively acting as a controlled "man-in-the-middle" to enable security inspection
- Applies security profiles such as URL filtering, DLP and anti-malware scanning to all traffic
- Authenticates users before granting any access, enforcing a Zero Trust approach to application access
- Uses Real Time Private Access rules to control which internal applications each user may reach
- Ensures that traffic between external users and internal web services can traverse only the SSE gateway via the reverse-proxy feature, providing a single, centrally managed enforcement point

Clientless reverse proxy DNS and FQDN design considerations

Public application FQDN

- These are the FQDNs used by remote users and represents the portal or private application protected behind the SSE gateway
- 'Portal' Format:
  <tenant-name>.myapps.<root-domain>
- 'Private Application' Format:
  <app-name>-<tenant-name>.myapps.<root-domain>
- For Versa cloud-hosted SSE services, Versa owns the root domain and creates both the FQDN and corresponding public DNS records
- For sovereign-hosted SSE services , the sovereign provider owns the root domain and is responsible for creating the FQDN and DNS entries

Authentication FQDNs

- The enterprise IdP must be reachable from the unmanaged/BYOD network so that the user can authenticate against the enterprise identity service protecting the application
- For example, if the enterprise uses Entra ID or Okta, the associated IdP FQDNs must be accessible directly from the external user's network
- (Because the user is remote to the SSE gateway in this flow, additional Internet Protection rules on the gateway are not required to permit user-to-IdP communication)

Internal / backend FQDNs

- The SSE gateway resolves and connects to the internal application using internal (to the target enterprise) DNS or static mappings

Public and internal FQDN separation

Maintaining a clear separation between user-facing (public) and backend (internal) hostnames simplifies certificate management and allows enterprises to move or refactor applications without changing the end-user URL. The public FQDN is bound to the SSE gateway, while the internal FQDN and IP space remain private and can evolve independently.

Considerations for Sovereign SASE: Public FQDN resolution and geo-routing

For Sovereign SASE deployments, external users still connect to a Public Application FQDN that resolves in public DNS to one or more IP addresses advertised by the SSE gateway service. Sovereign SASE providers can use advanced DNS capabilities such as geolocation routing and health-aware load balancing (for example Geo-DNS or GSLB) so that queries for the Public Application FQDN resolve to the closest healthy SSE gateway point of presence. This improves latency and user experience and increases resilience, because traffic can be automatically steered away from degraded regions or sites without any change on the end-user device.

Since the root domain is under the control of the sovereign SASE provider, the same DNS platform used for FQDN resolution can also be used to implement geo-location routing and health-aware load balancing policies.

Clientless reverse proxy TLS certificate design considerations

Certificates presented to external users

- The SSE gateway presents valid, public CA-signed certificates for every user-facing FQDN so unmanaged devices can trust the connection without installing a corporate root CA

- For Versa cloud-hosted or sovereign SSE services, the provider is responsible for obtaining and managing these public certificates because it owns the user-facing root domain. The enterprise does not supply these certificates
- For sovereign SSE services, it is recommended to use a wildcard SAN pattern that can accommodate multiple tenants/applications exposed via the clientless reverse proxy, avoiding frequent certificate changes as new apps are published
- For example, a wildcard SAN of *.myapps.<root domain> will support all combinations of application and tenant that could be provisioned on the SSE Gateway

Certificates toward the origin application

- During application onboarding via Versa Concerto, the enterprise administrator uploads the application certificate used between the SSE gateway and the origin application
- This certificate must include the SAN matching the external FQDN used by clientless users (for example, <app-name>-<tenant-name>.myapps.<root-domain>) in addition to the application's original Common Name/SAN. (Without this additional SAN entry, the TLS connection will fail validation during the TLS handshake)

General certificate behaviour

The external user's browser sees only the certificate chain presented by the SSE gateway's reverse-proxy listener, not the backend application's certificate. No certificate or key material from internal applications is ever exposed directly to the public internet.

- The browser establishes TLS directly with the SSE gateway, which presents its own server certificate and any required intermediate CA certificates. Because this certificate is signed by a public CA, the browser can validate the connection without additional trust configuration
- The reverse proxy then establishes a separate TLS session to the backend server. That second certificate exchange is invisible to the external client
- In the browser's certificate viewer, the subject, issuer, and chain will always correspond to the reverse proxy's certificate rather than the origin web server

Clientless reverse proxy IP addressing design considerations

External user IP addresses

- The clientless reverse proxy feature **obfuscates external user IP addresses** by performing Source Network Address Translation (SNAT) on traffic flowing through the SSE Gateway
- During tenant onboarding, each SSE Gateway is assigned a dedicated pool of IP addresses by the network administrator. This pool is typically already used for agent-based users connecting with the SASE client to the SSE service
- Commonly, each SSE Gateway is provisioned by the network administrator with a **unique RFC1918 address block** that is globally unique within the enterprise network, ensuring unambiguous routing and policy enforcement
- The clientless reverse proxy reuses this same address pool for agentless (browser-based) sessions, so no additional routing configuration is required between the internal web application and the user; the path is already established for agent-based traffic
- SNAT also ensures that all return traffic destined for external users is steered back through the SSE Gateway as the security enforcement point, preventing bypass via alternative routes inside the enterprise's network

Internal Web application IP addresses

- From the external user's perspective, the clientless reverse proxy hides the internal web application IP address via Destination NAT (DNAT)
- DNAT rules are automatically created during clientless reverse proxy configuration – therefore no manual configuration is required when onboarding the service

Clientless reverse proxy summary

Clientless reverse proxy on SSE gateways delivers secure, browser-based access to internal applications without requiring an endpoint agent, making it ideal for BYOD, third-party, and ad-hoc users. It enforces Zero Trust principles by authenticating identity and applying granular, Layer-7 controls before any connection to the protected application is established.

By terminating external TLS sessions and creating separate, internal TLS connections, the SSE gateway keeps internal apps and IP space hidden while still allowing full inspection for URL filtering, DLP, and malware protection. Split public and backend FQDNs simplify DNS and certificate management, and sovereign deployments can add geo-aware DNS routing to optimize latency and resiliency.

Because all traffic and authentication events traverse a single enforcement point, enterprises gain centralized logging, forensics, and policy control, reducing operational complexity compared with traditional VPN models.

https://versa-networks.com/documents/solution-briefs/infographic-sovereign-sase.pdf

Versa Networks, Inc
2550 Great America Way, Suite 350
Santa Clara, CA 95054
Tel: +1 408.385.7660
Email: info@versa-networks.com
www.versa-networks.com

About Versa
Versa, a global leader in SASE, enables organizations to create self-protecting networks that radically simplify and automate their network and security infrastructure. Powered by AI, the VersaONE Universal SASE platform delivers converged SSE, SD-WAN, and SD-LAN solutions that protect data and defend against cyberthreats while delivering a superior digital experience. Thousands of customers globally, with hundreds of thousands of sites and millions of users trust Versa with their mission critical networks and security.

Part# TB_CLIENTLESS-SSE-01.0