aryaka

WHITEPAPER

Aryaka Network Architecture

Unified SASE as a Service
Zero Trust WAN

Table of contents

Aryaka Network Architecture: Unified SASE as a Service PoPs

Aryaka's Zero Trust WAN infrastructure uniquely supports sophisticated network and security capabilities at the cloud-edge as part of a Unified SASE as a Service architecture. The Aryaka architecture delivers numerous advantages vs. more traditional transport layer PoPs (Point of Presence) and is the foundation for our Zero Trust approach. These advantages include:

1

Sophisticated services processing that combines routing, switching, compute, and storage, leveraging a OnePASS containerized architecture that scales to any number of enterprise customers and provides an ideal foundation for Unified SASE capabilities by avoiding service silos

2

Hardware-based multi-layer traffic management to optimize application performance across both L2 and L3 cores, while eliminating issues with disjoint overlays and underlays

3

Close handshake between the PoP and the Aryaka Network Access Point (ANAP), the Aryaka CPE, that diffuses the question, 'how many PoPs are enough?'

We describe each of these advantages in more detail below.

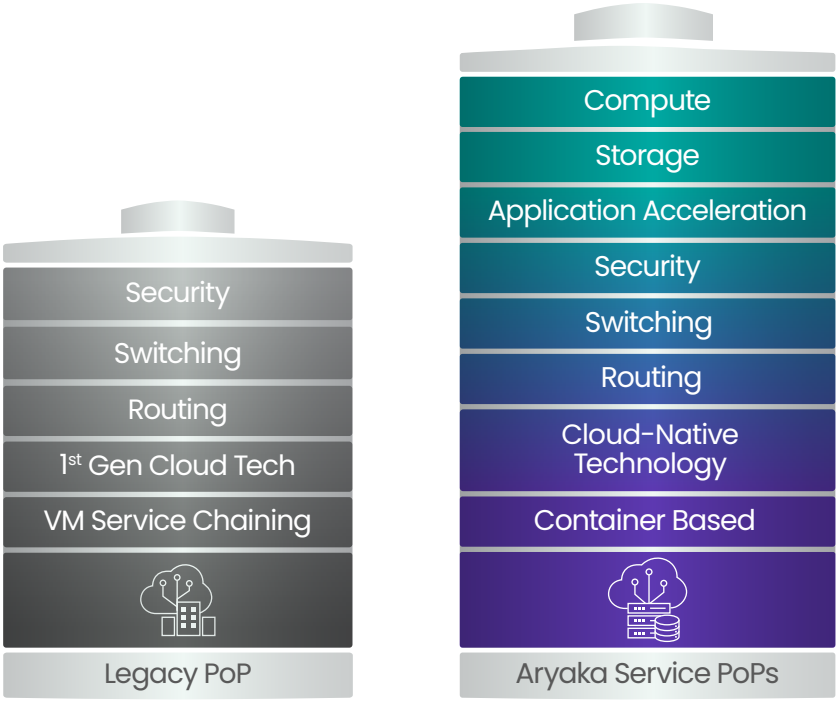

Advanced Services Processing

Aryaka Service PoPs are very different from the routing and switching transport PoPs operated by many providers. Our PoPs of course include these functions, but also integrate compute and storage for advanced services. The hardware investment is substantial, and in some locations such as Ashburn, San Jose, London, and Frankfurt, over sixty discrete elements are in play.

The diagram below depicts a typical Aryaka PoP design, with resilient core routers connecting the PoPs to one another and to ISPs, distribution switches, and any number of servers. The servers implement a single-pass containerized architecture with per-customer resource allocation and separation of tasks such as transport, encryption, or protocol acceleration. This unified hardware and software architecture permits Aryaka to tightly couple networking and security services, and one reason many vendors have problems delivering a unified architecture is that their hardware and software platforms are anything but that.

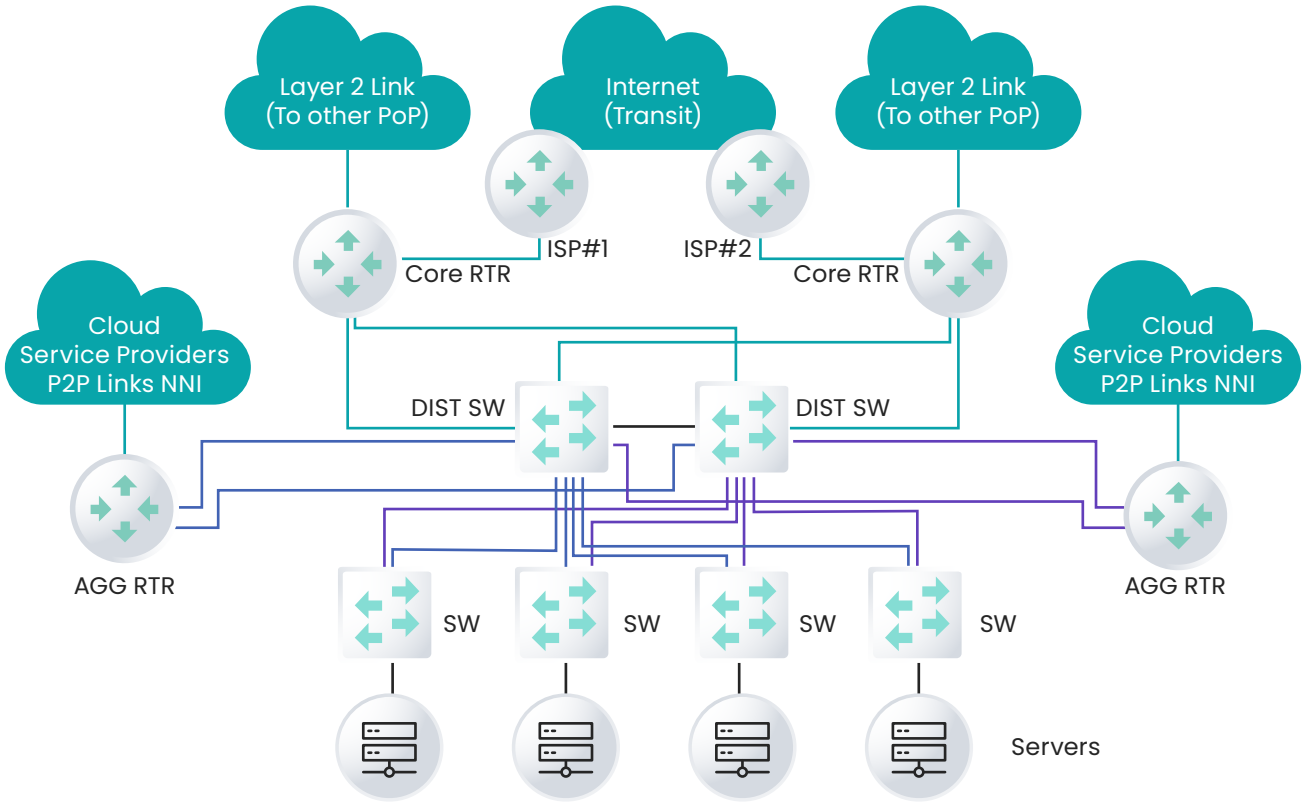

The architecture also lays the foundation for a Zero Trust WAN, by distributing networking and security functions across the fabric and in some cases to the ANAP at the edge. We know there are many views of Zero Trust, with resulting confusion. First off, Zero Trust is not a product but a set of guiding principles that may be summarized as 'never trust, always verify.' Or a change from the 'castle and moat' approach since many employees, suppliers, and applications reside outside of the moat.

We look at these principles from the lens of the WAN. Basically, how to apply them to an end-to-end networking and security architecture. By leveraging a single software codebase, the enterprise gains the observability and control required to ensure regional or global application performance and security. A way to look at this is the combination of a single management pane, unified control, and a distributed data plane supporting networking and security.

What does this three-tier architecture look like in reality? The picture below is of a typical Aryaka PoP in a cage, with most of server (compute and storage) functionality in the three racks center to right, with the switching and routing on the left.

Aryaka Dynamic PoPs

What are dynamic PoPs?

Unlike Aryaka's traditional physical PoPs, which rely on dedicated hardware deployments, Aryaka's dynamic PoPs are deployed in the cloud using AWS infrastructure. This cloud-native approach allows for a rapid, scalable, and cost-effective expansion of Aryaka's global network.

This dynamic PoP framework was intentionally designed to closely mirror that of physical PoPs, ensuring consistent support, monitoring, and troubleshooting.

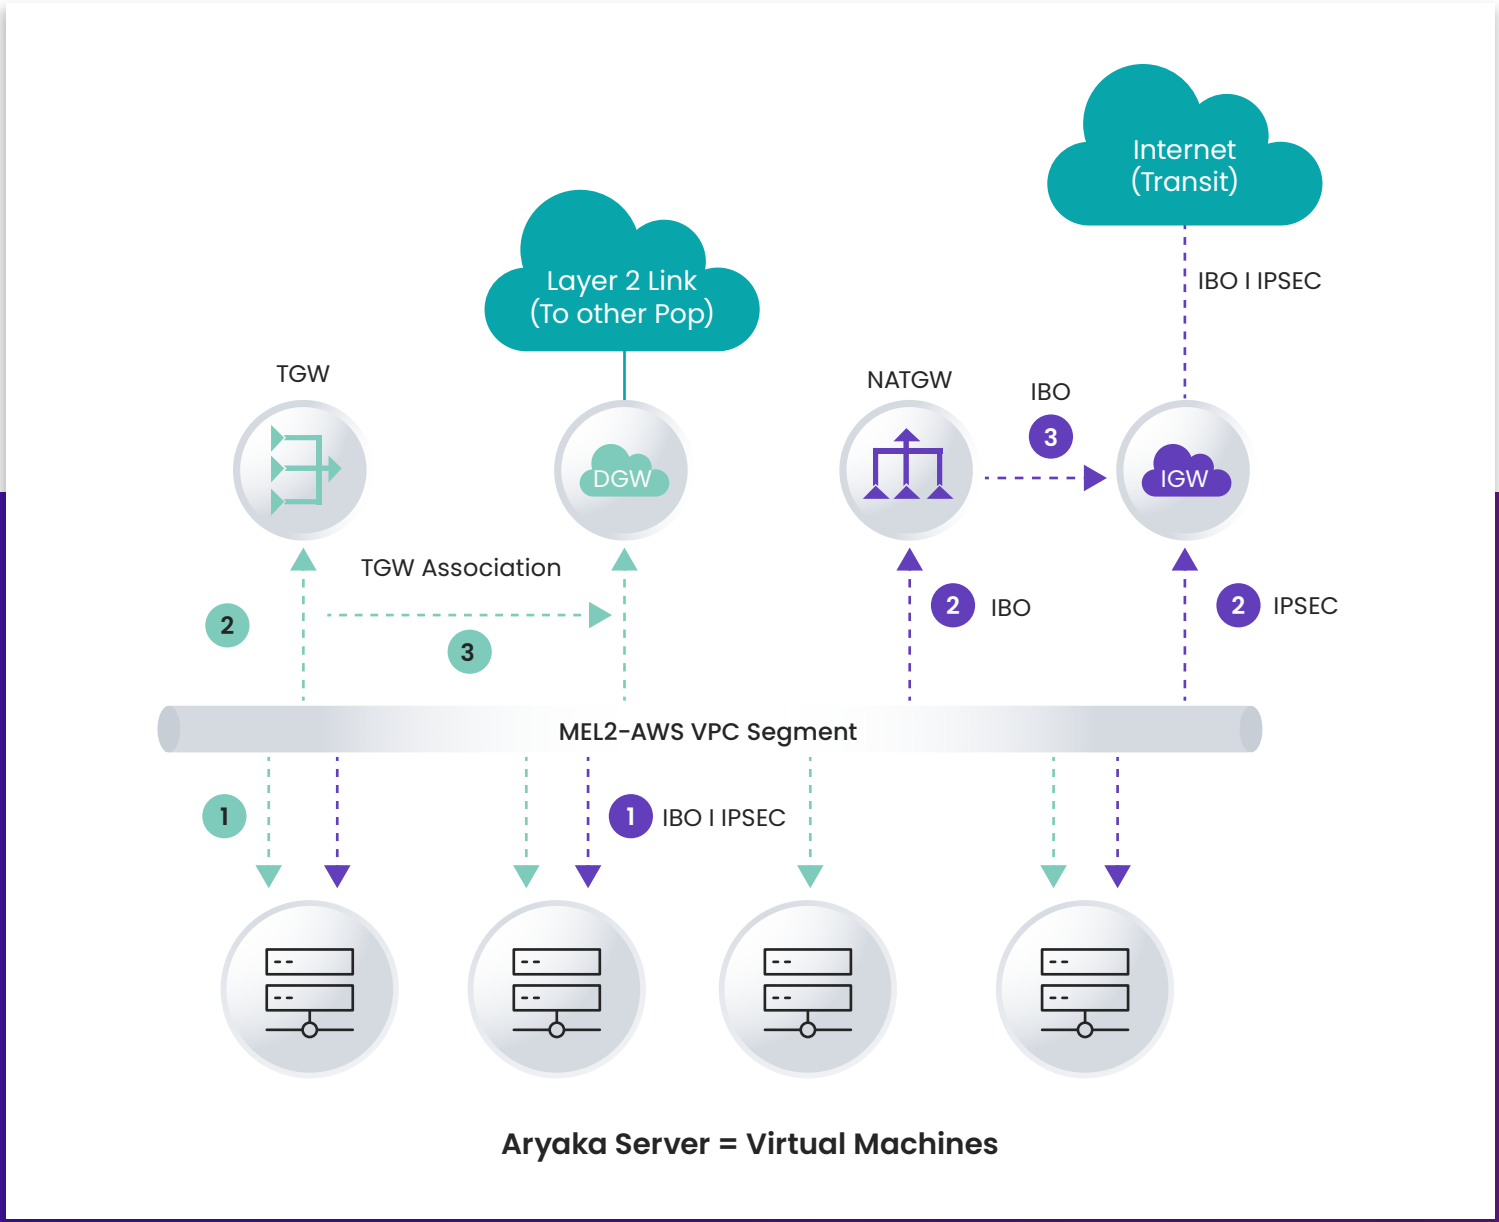

Dynamic PoP Architecture

Why Aryaka Uses Dynamic PoPs

When dealing with physical PoPs, deployment is constrained by the logistical and operational delays of configuring physical hardware. The objective of implementing dynamic PoPs is to accelerate the deployment of Aryaka's PoP infrastructure in regions where a cloud presence already exists.

Dynamic PoPs enable Aryaka to provide:

Rapid and flexible scalability:

Resource and computing capacity can be quickly scaled up or down to match demand, with minimal manual intervention required.

Global reach:

Aryaka can expand its presence across over 240 countries and territories, leveraging the distributed infrastructure of AWS.

Cost efficiency:

There is no need for upfront capital expenditure (CapEx) that is required for physical infrastructure. Dynamic PoPs can utilize an operational expenditure (OpEx) model aligned with actual usage—you can pay as you go.

Accelerated deployment:

Deployment time is significantly reduced (2-3 weeks for a dynamic PoP compared to 2-3 months for a physical PoP), allowing Aryaka to react swiftly to the needs of customers and regions.

Feature parity with physical PoPs:

All features behave consistently across physical and dynamic PoPs, ensuring service quality is upheld throughout Aryaka’s hybrid architecture.

What are the differences between physical PoPs and dynamic PoPs?

Aryaka has made an effort to minimize the differences between traditional, physical PoPs and dynamic PoPs—there are no differences in performance or feature availability when using a dynamic PoP instead of a physical PoP. However, as the dynamic PoP is a virtual AWS configuration, there are architectural differences that we must adapt to.

Dynamic PoPs exhibit the following differences compared to physical PoPs:

1. A single edge ISP link
2. No routers or switches
3. A shared IP address for internet breakout
4. Virtual machine failover

1. Single Edge ISP link

Physical PoPs: Support at least two ISPs, allowing customers with two ISPs to have ISP redundancy on PoPs.

Dynamic PoPs: AWS elastic IPs are routed through AWS's internal network stack and therefore share infrastructure, such as internet gateways and VPC routers, within an AWS Availability Zone. Adding multiple elastic IPs in the same AWS region doesn't create true ISP diversity—failures in the shared cloud network layer can affect all IPs simultaneously.

For this reason, dynamic PoPs use a single edge ISP link. Resiliency is ensured through AWS's backbone, which has extensive peering with multiple Tier 1 ISPs, redundant gateways, and built-in rerouting during failures. Customers can still maintain edge redundancy by deploying dual ISPs on their sites.

2. No Routers or Switches

Physical PoPs: Core routers handle PoP-to-PoP connectivity and terminate ISP uplinks for customer IPsec connections. Aggregation routers terminate P2P links from ANAPs (or CPE) as well as Direct Connects from IaaS platforms (for example, AWS, Azure, GCP, Oracle, and Alibaba) that use private connectivity options (for example, AWS Direct Connect, Azure ExpressRoute, or Google Dedicated Interconnect) using Equinix Fabric for high-performance cloud interconnects.

Dynamic PoPs: AWS Direct Connect gateways act as "core" routers for PoP-to-PoP connectivity, while internet gateways connect the PoP to the internet. Currently, dynamic PoPs do not include aggregation router equivalents, meaning P2P links and dedicated or private IaaS platform connections are not directly supported from the dynamic PoP itself. These are handled by Aryaka's global network of physical PoPs, which support local and remote interconnects to multiple cloud providers through Equinix ports.

However, secure VPN connectivity from dynamic PoPs to IaaS platforms is still supported. This allows customers to establish secure, encrypted tunnels from a dynamic PoP directly to their cloud-hosted workloads—providing flexibility for hybrid architectures without requiring physical PoP proximity.

3. Shared IP address for internet breakout

Physical PoPs: Each site connected to the PoP receives a dedicated, unique public IP address. Sites then use this public IP as the internet breakout IP to send internet traffic.

Dynamic PoPs: All sites connected to the PoP share the public IP addresses assigned to the PoP's NAT gateway. AWS NAT gateways can have up to 8 associated IP addresses, which are automatically load-balanced for outgoing internet traffic.

While users do not receive a unique breakout IP for each site, this shared IP model is optimized for the most common requirement—secure, reliable internet access for branch and remote users, using PAT overload for efficiency.

4. Virtual machine failover

Physical PoPs: If a physical server or service fails, failover to healthy hardware requires a manual process.

Dynamic PoPs: AWS auto-recovery ensures that if the underlying hypervisor fails, the virtual machine is automatically restarted or moved to healthy hardware—without manual intervention.

This results in minimized downtime and accelerated recovery in the event of infrastructure issues, allowing users and enterprises to maximize efficiency and business continuity.

OnePASS Architecture

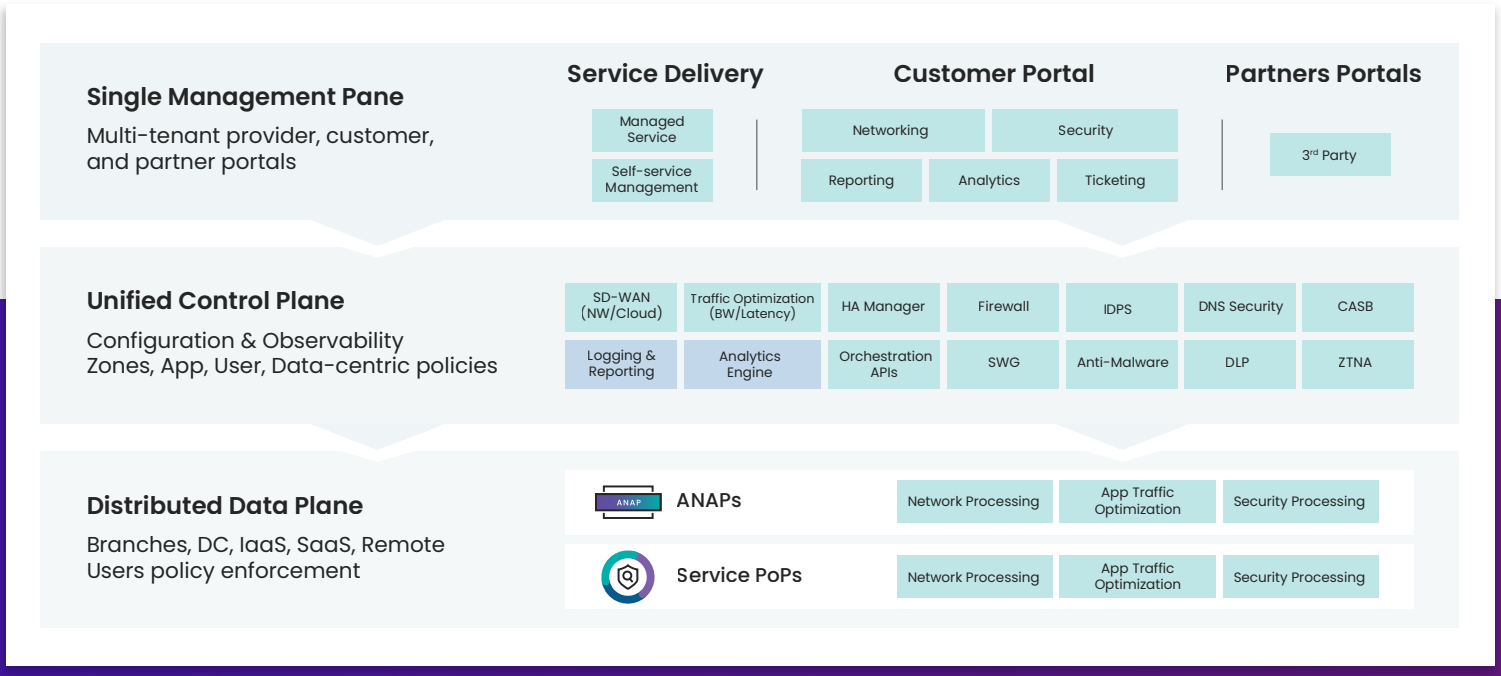

Another key advantage of Aryaka's architecture is its ability to automatically bring online additional resources when utilization nears a defined threshold. This reserve capacity ensures uninterrupted performance and prevents any service degradation for enterprises.

Aryaka's hardware and software architecture offers a distinct edge over approaches like Network Function Virtualization (NFV), often promoted by telcos. While NFV has its applications, such as supporting Virtual Network Firewalls (VNGFWS) on Aryaka's ANAPs, it can face performance issues at the PoP level due to service chaining and resource allocation challenges. Additionally, VNFs originating from different code bases or acquisitions may not play nice due to a lack of seamless interoperability.

The required hypervisors to support NFV are also more of a 'heavyweight' approach vs containers. In fact, we only support a single VNF on our ANAPs to ensure line-rate performance. Yet another PoP architecture is to just run routing and services in a VM within a cloud provider, such as as an AWS EC2 instance. These nodes are very quick to spin-up but are limited in performance.

Multi-Layer Traffic Management

The next advantage of the Aryaka architecture is traffic engineering, leveraging the hardware and software architecture described above. Each Services PoP connects to both a Layer 2 and Layer 3 core, the former based on leased fiber, and the latter based on point-to-point non-peered IP connections. Unlike many solutions that only offer Layer 3 (internet) handoffs to the cloud providers, Aryaka's dedicated connections are based on AWS's Direct Connect, Azure ExpressRoute, and the equivalent offerings from the other CSPs, guaranteeing end-to-end application performance and preserving SLAs.

These performance SLAs are critical, with packet loss, latency, and jitter all controlled, especially important for real-time applications such as unified communications or production monitoring. The elimination of the traditional underlay-overlay boundary results in additional control and observability that an 'underlay' provider partnering with an 'overlay' SD-WAN vendor can't replicate.

Another way to look at the Aryaka's architecture is how the middle mile is treated. Aryaka's L2 middle-mile relies on dedicated fiber interconnects between its Service PoPs, ensuring a single provider is responsible for the link, without relying on peering. Similarly, Aryaka's L3 service uses point-to-point IP Layer links sourced exclusively from Tier 1 providers.

In contrast, other approaches in the industry often depend on Layer 3 and involve multiple ISPs, which can introduce ISP peering. Additionally, some solutions rely on CSP peering, where performance guarantees may be limited. While such configurations can be effective for content acceleration, they are often less suited for bi-directional interactions that require consistent and predictable performance.

Key metrics for ensuring guaranteed application performance include packet loss, latency, and jitter. Packet loss leads to data retransmissions, while jitter represents variations in latency. For example, the diagram below illustrates the latency across the Aryaka network from London to a company office in Melbourne, Australia. The total latency is approximately 265ms—the lowest possible due to the speed of light. The connection from the London office to the PoP adds negligible latency, while the segment from the Sydney PoP to Melbourne contributes about 15ms. Beyond low latency, applications demand predictability, which highlights the importance of controlling jitter. The graph clearly demonstrates this consistency, showcasing the performance Aryaka customers can expect.

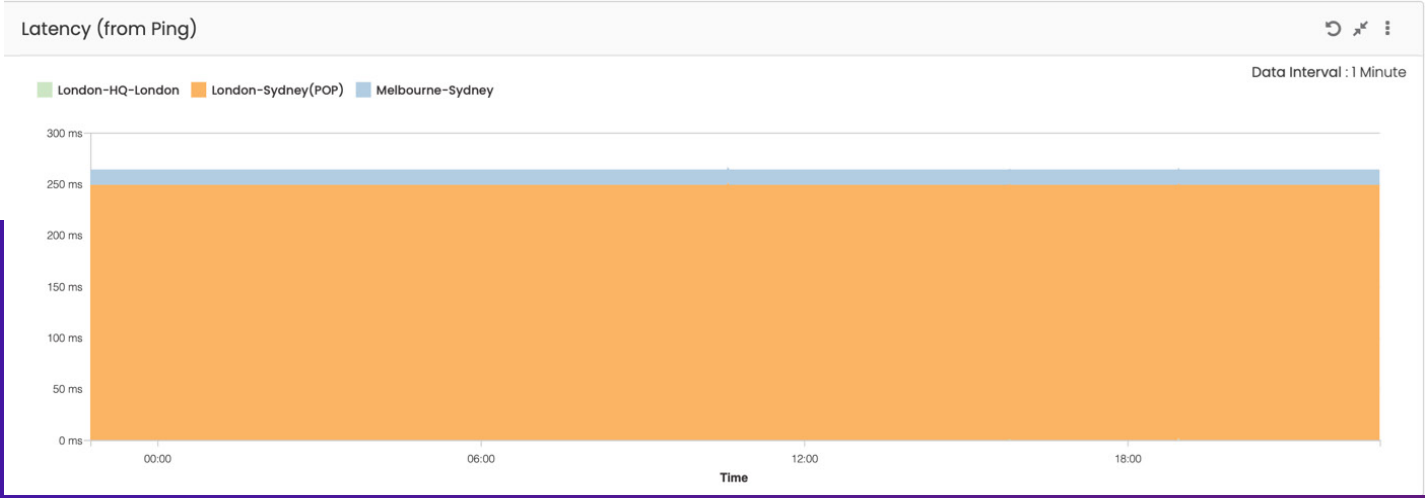

The minimization of loss and control of latency and jitter results in better performing applications, both at L2 and L3, as depicted below. The first diagram visualizes not only the advantages of the Aryaka L2 middle-mile, but also Aryaka's managed L3. This is very evident in the second diagram, where control of L3 eliminates loss.

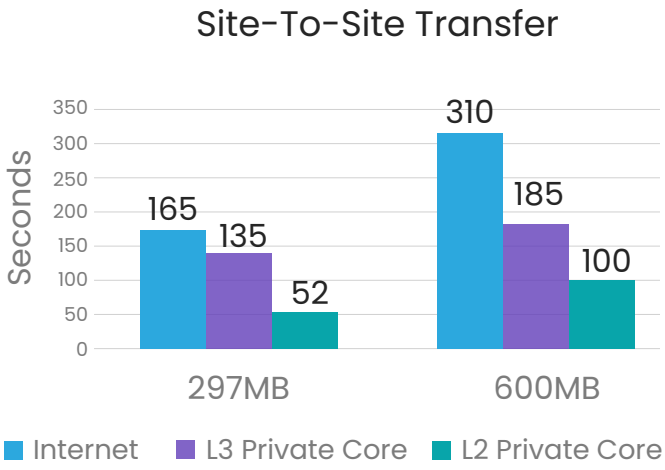

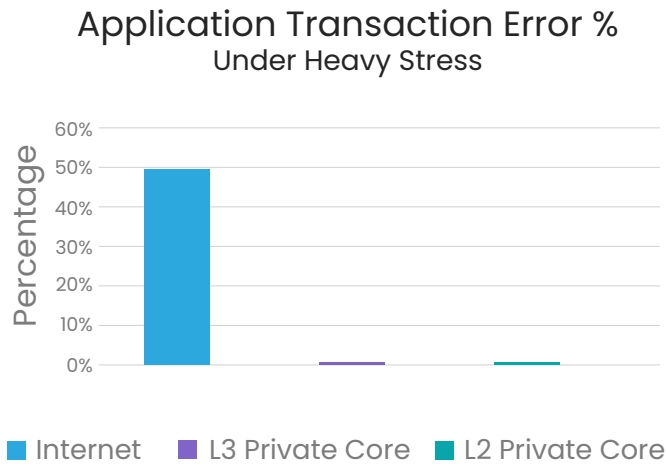

Aryaka PoP to ANAP Handshake

The last advantage is the tight interworking between the Aryaka PoP and the ANAP. This influences both the user experience, as well as how many PoPs are required. Looking back at the latency discussion, there is a great deal of focus (and confusion) regarding latency across the last mile, the link between the ANAP and the Services PoP. But the real issue is not the absolute latency, but how well the service handles it.

Aryaka has developed a best-in-class multi-segment optimization architecture, and between the ANAP and the PoP, capabilities include strict QoS, encryption, redundancy, application acceleration, data deduplication, and compression. Our target is that any enterprise should be within 30ms of one of our PoPs to be able to take best advantage of these capabilities, and this is true for what we estimate to be 95% of the knowledge workers or enterprise sites in the world, including remote workers. In reality, many of our PoPs are much closer to the user.

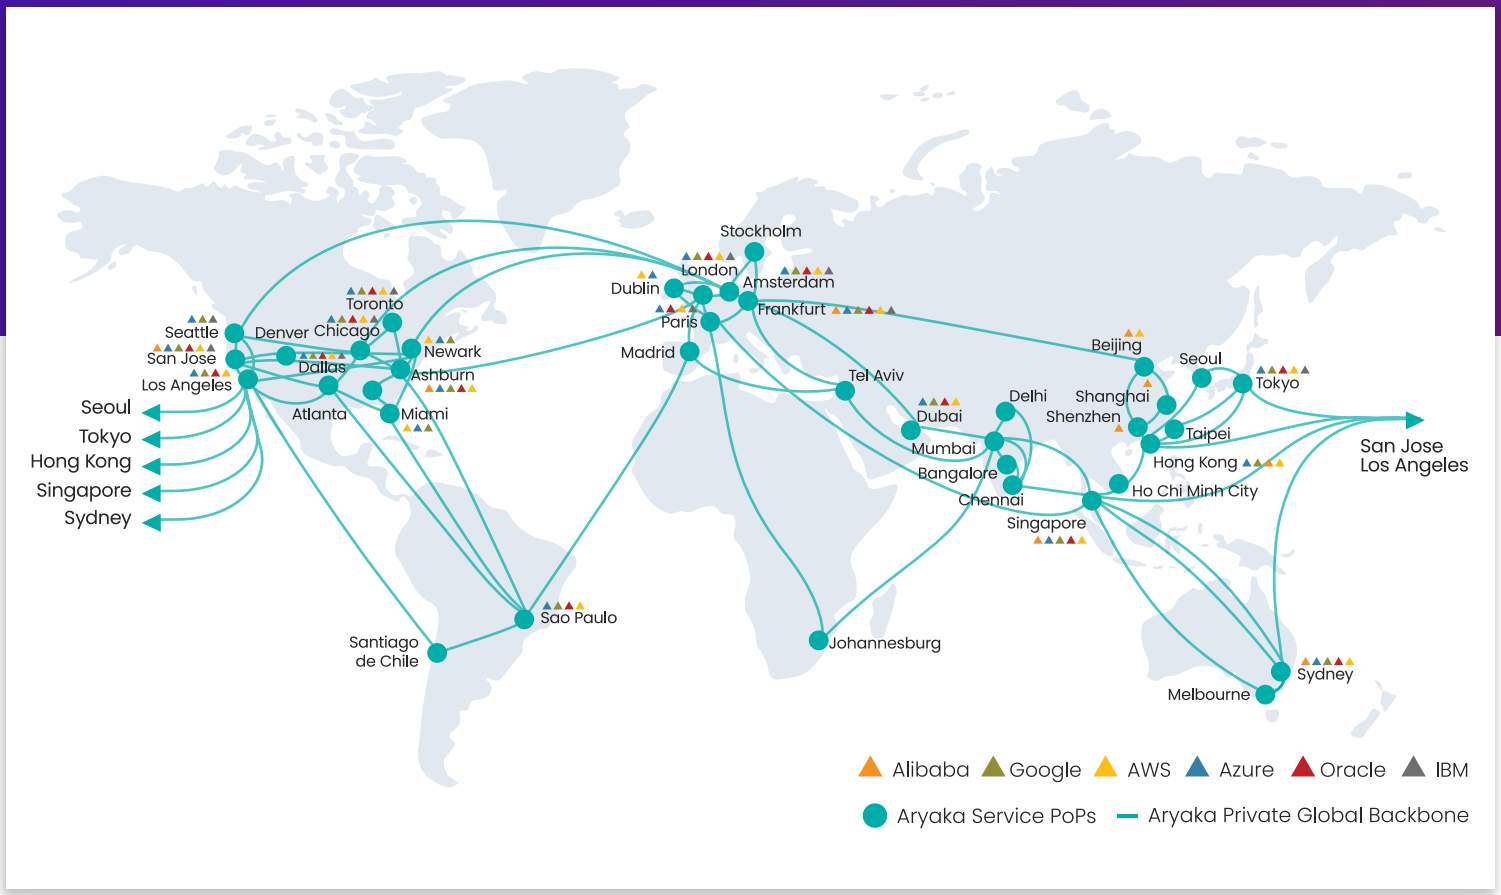

The number of PoPs required often becomes a marketing point, with many vendors claiming "more is better." But is that true? Aryaka's scale-out architecture shows otherwise, with no limit on the number of customers a single PoP can support. In fact, regions like Ashburn or Frankfurt function as "PoP clusters," equivalent to 10 or even 50 PoPs by other vendors' definitions. Aryaka's existing PoP footprint and single-pass architecture with a distributed data plane are sufficient for reliable and high-performant service delivery, including Unified SASE. In the middle mile, where regions and continents are interconnected, fewer PoPs often deliver better performance. For instance, a direct Sydney-to-San Jose link performs better than routing through Tokyo or Seoul.

Summary

The Aryaka Service PoP architecture and global footprint deliver unparalleled connectivity to enterprises wherever their employees and applications reside. It is an ideal foundation for deploying advanced network, security, and observability functionality to future-proof the modern workplace and distributed enterprise.

About Aryaka

Aryaka is the leader in delivering Unified SASE as a Service, a fully integrated solution combining networking, security, and observability. Built for the demands of Generative AI as well as today's multi-cloud hybrid world, Aryaka enables enterprises to transform their secure networking to deliver uncompromised performance, agility, simplicity, and security. Aryaka's flexible delivery options empower businesses to choose their preferred approach for implementation and management. Hundreds of global enterprises, including several in the Fortune 100, depend on Aryaka for their secure networking solutions. For more on Aryaka, please visit www.aryaka.com.

Schedule a Free Network Consultation with an Aryaka Expert

See How It Works Live →

Experience Aryaka's Unified SASE as a Service

View Interactive Tour →

Contact Us

+1.888.692.7925 info@aryaka.com

aryaka