



# Building and Managing a Successful Multi-Cloud Strategy

How to overcome security, visibility, and connectivity challenges

AUTHORS

Nishant Singh, Product Marketing Manager, Aryaka

Himanshu Dhingra, Technical Marketing Engineer, Aryaka



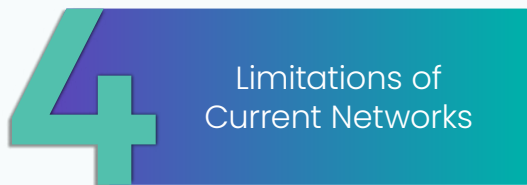
PAGE NUMBER: 3.



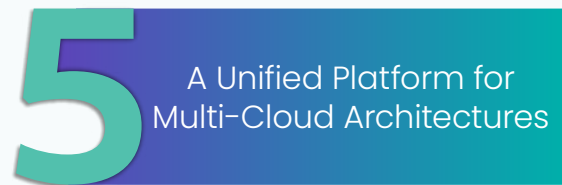
PAGE NUMBER: 4.



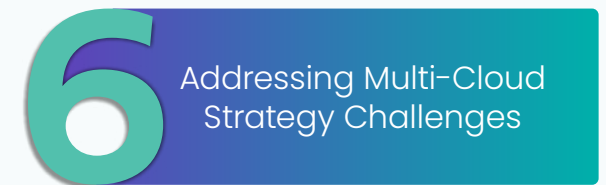
PAGE NUMBER: 5.



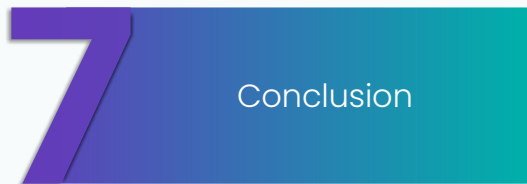
PAGE NUMBER: 9.



PAGE NUMBER: 11.



PAGE NUMBER: 12.



PAGE NUMBER: 23.

# Table of Contents

## Introduction

Of all the emerging technologies contending to grab enterprises' attention, cloud services undisputedly hold the top spot. To an extent where about **one-third** of most company's IT budget goes towards implementing cloud strategies. Heralded as the foundation of multiple technological breakthroughs, cloud-tech also emerged as IT's saving grace during the pandemic – while the brick-and-mortar branches had to pull down shutters.

Though most organizations begin their journey with a single cloud, they eventually end up running multiple workloads across different cloud platforms. For some, it is about nimbler network operations and lowering capital and operational expenses; others do so to reap the advantages of a richer feature set.

The variety of use cases is so broad that most organizations are currently using **3 ½ public and close to 4 private clouds** on average. Enterprises are going all out on multi-cloud not just in terms of service offerings, but also for resilience planning to deliver an 'always-on' business.

This whitepaper intends to demystify the dilemma faced by the frontline decision-makers charged with implementing cloud services in their organizations. By the time you're done reading this piece, expect a clear understanding of whether or not a multi-cloud strategy seems sensible for your business requirements, and should you choose to move cloudward, this whitepaper will help you make an informed choice.



## Digital Transformation and Why a Multi-Cloud Strategy Makes Sense

By definition, the term digital transformation means injecting new technologies into existing business processes, or creating new processes on top of these technologies to match the changing business and market landscape. Consider the recent COVID-19 situation, for instance, which forced businesses to dismantle their existing practices and participate in all sorts of unusual experiments to come up with new contingency plans.

The not-so-subtle push of remote work that transformed homes into virtual branch offices saw a massive shift of workload move away from on-premises environments into the cloud. But moving resources is only a job half done, and a job half done is as good as none. What's of utmost importance is to put the right workloads in the right environment for optimal performance and efficiency.

But, not all applications or workloads are created equal. Some are static, and a private cloud could be the best fit for them. Others, due to spiky processing or scaling needs, fit better in a public cloud. Making these assessments and decisions is the first step towards the transition.

Distributing workloads across multiple clouds has strong benefits...



- ✓ Maximizes availability and reliability **(63%)**
- ✓ Fulfills regulatory and compliance requirements **(47%)**
- ✓ Leverages best-of-breed services from each provider **(42%)**

Source: <https://bit.ly/32tudhX>

Similarly, while the IaaS platforms such as AWS, Azure and Google Cloud facilitate unmatched storage, compute and database services, when it comes to adding a dash of machine learning and advanced analytics, you need to be able to draw services from different providers and move data and resources between multiple clouds. It often requires a hybrid multi-cloud infrastructure environment that allows enterprises to deploy resources when they need them, where they need them.

If **predictions from IDC** and others are to be believed, 2021 is going to be the year of multi-cloud, with the vast majority of enterprises deploying combinations of on-premises, off-premises, public, and private clouds as their default environments. By 2022, over 90% of enterprises worldwide will be relying on a mix of on-premises/dedicated private clouds, multiple public clouds, and legacy platforms to meet their infrastructure needs. Which means now is the best time to get started.

## Top Three Multi-Cloud Strategy Challenges:

While connecting a single cloud data center instance to WAN may seem like a routine task to in-house IT folks, things get interesting when you need to scale that model across multiple branches. More so, if those branches run across numerous geographies and you have multiple cloud providers to connect to.

Managing multiple applications & tools that exchange a mammoth amount of data across different platforms with diverse infrastructural capabilities can give cold feet to even the most seasoned IT folks. So, which are the top three speed bumps in a multi-cloud rollout plan that keeps IT up at night?

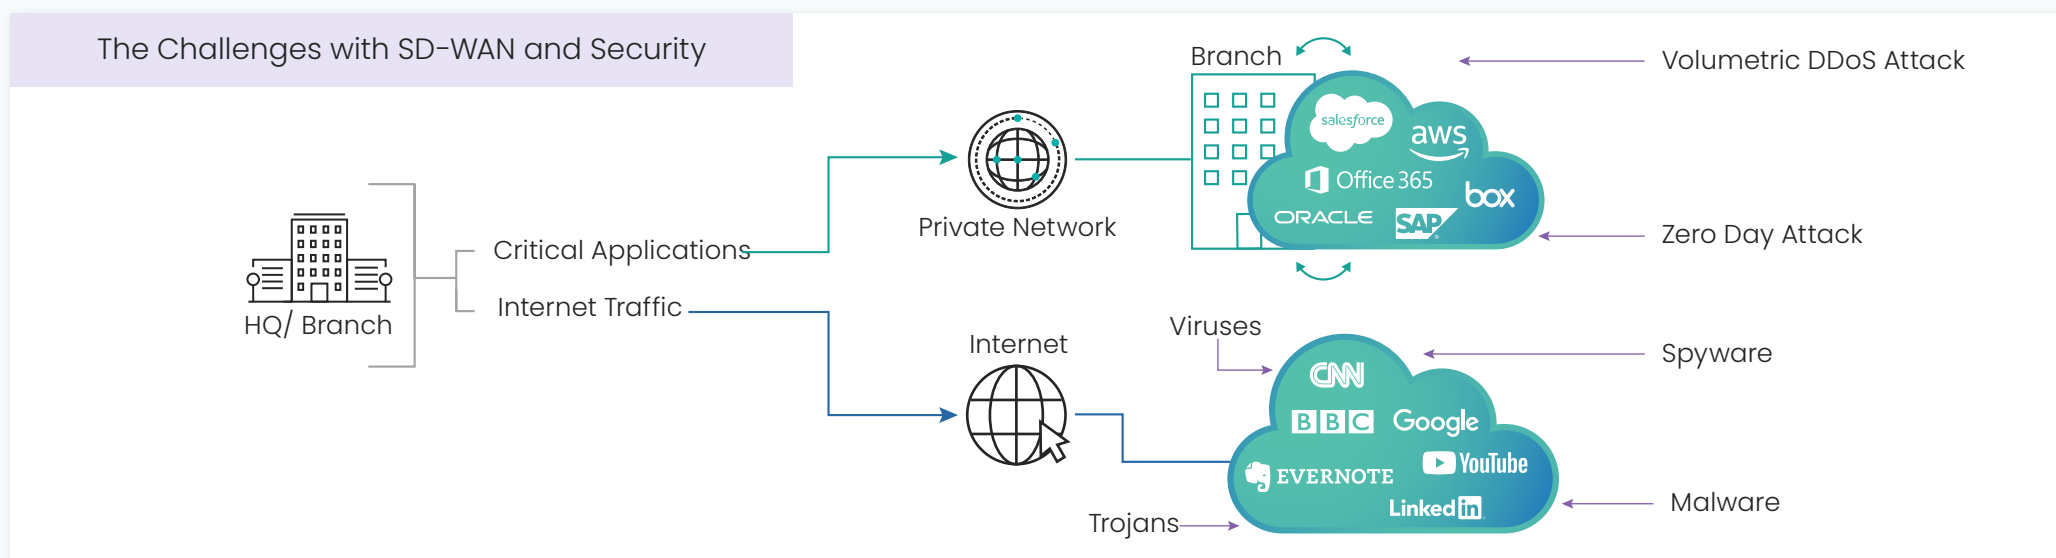




## 1. Security

Running applications through multiple cloud providers, with each one having its own set of controls and policies, gives way to numerous security challenges. More so, as synchronizing policies and settings across different providers is difficult. Though most cloud vendors spend a good chunk of their budget on security, they fall short when it comes to providing a uniform security blanket for the entire infrastructure.

SLAs, compliance, access controls, identity authentication, passwords, firewall, auditing process, app sprawl, IT documentation, patches, upgrades, and what have you? So many moving parts to be managed and done right. When it comes to cybersecurity and breaches, though malware and attacks grab all the attention, internal errors and user gaffes cause nearly the same number of mishaps.



Unlike the on-premises or single cloud model, a multi-cloud architecture means more vendors and services. Thus, a larger attack surface where the legacy perimeter-focused strategy usually fails.

Nevertheless, many enterprises still rely on old, outdated tools or one-off, isolated solutions designed to secure data in an individual environment. Transitioning to a hybrid and multi-cloud environment demands configurations and solutions built to manage new data security challenges across multiple environments.



## 2. Visibility

Poor visibility into a brilliantly crafted cloud architecture can only get you so far. Having the latest security measures in place is good, but you need someone to shine the torch on the spots where those measures need to be applied and make the most sense.

IT security and analytics tools are only as good as the data they capture. And not just security; visibility is also paramount for capacity management. So, before you sign new checks to acquire more bandwidth, you might want to ensure that the existing capacity is being used for legitimate business purposes. The point is — with enough visibility, managing a multi-cloud architecture can be easier than many believe.

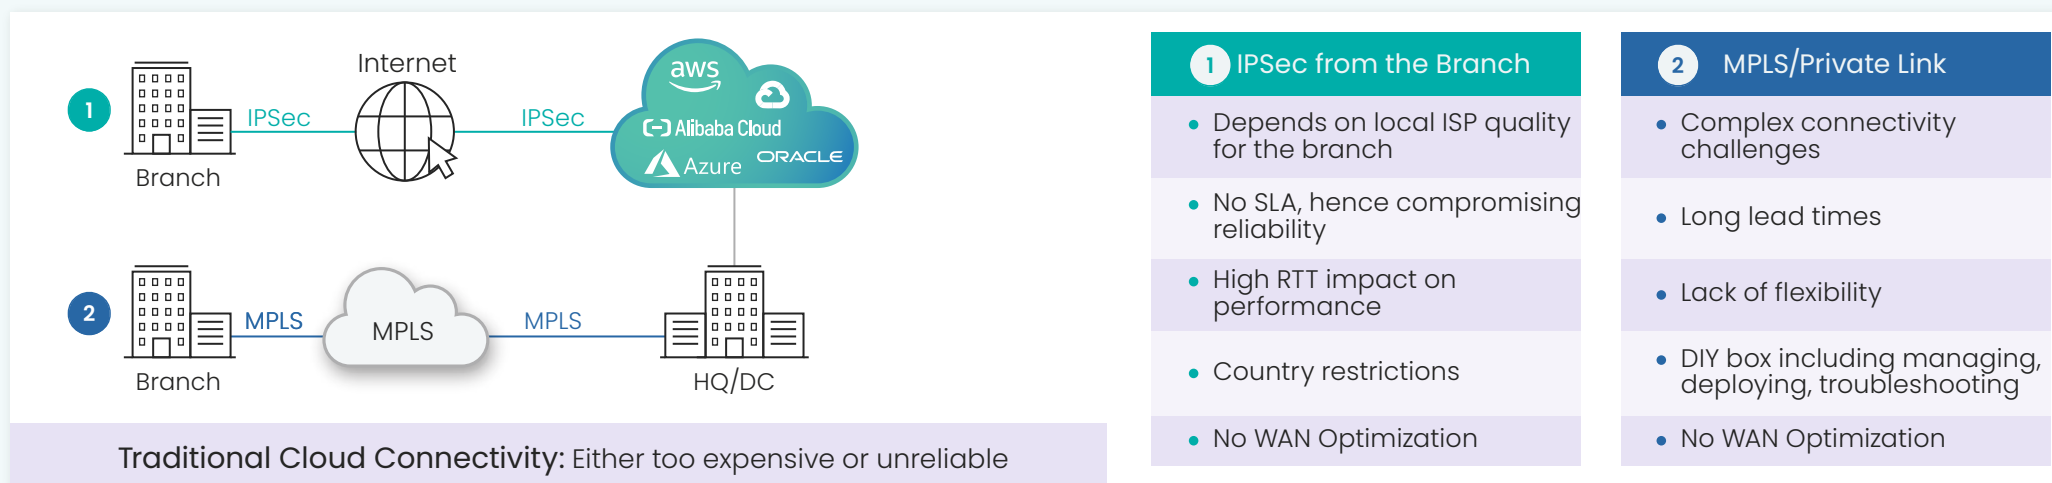


Visibility begins with in-depth access to data. You need to capture weekly, daily, and hourly fluctuations in traffic load, patterns, and behavior. The ask should be to have a single pane of glass view of pivotal information such where are the resources and data being accessed, as well as when and by whom? An eagle-eye view of the traffic flow. How much data is being consumed by different sites? And most importantly, a visual dashboard, capturing device configurations, any active or assessed attacks, non-compliance with policy, and any other associated risk.

### 3. Connectivity

Connectivity undisputedly is the holy grail of a multi-cloud architecture and cloud-tech in its entirety. All the fancy SaaS/ IaaS applications are good in theory, but eventually it all boils down to performance and the end-user experience, both of which ride on network connectivity. The efficiency of the business is directly proportional to how smoothly the information and process flows.

Unfortunately, network planning often takes a backseat. Most available network mediums still need to catch-up to support the scale of data movement and processing, and most public cloud providers are not concerned with connecting to clouds outside of their own service. The result? A cobweb of networks to work through and manage.

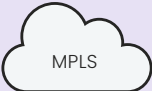


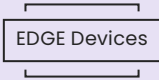


To address these diversity challenges and facilitate a centralized point for inspection and routing of data, organizations often choose to connect to cloud services through an on-premises data center. This conventional router-centric architecture not only comes with deployment and scalability challenges but was never designed keeping a cloud-first ecosystem in mind.

Contrary to the above approach, some choose to leverage the inexpensive and more flexible broadband internet services for cloud connectivity. While that may look inexpensive on the monthly bill, what that bill doesn't reflect is the risk of a downtime that can shoot up to **\$8,580 and \$74,000** per hour for small to medium businesses. When every second of downtime results in lost revenue and business opportunities, one has to be thoughtful about placing their bets on the internet for moving cloudward.

## Limitations of Current Networks

When it comes to anything 'cloud,' it is better to keep legacy networks off the table. Not only are they too rigid and complex to satisfy the growing bandwidth requirements, they provide zero optimization and are not cloud-ready to say the least.

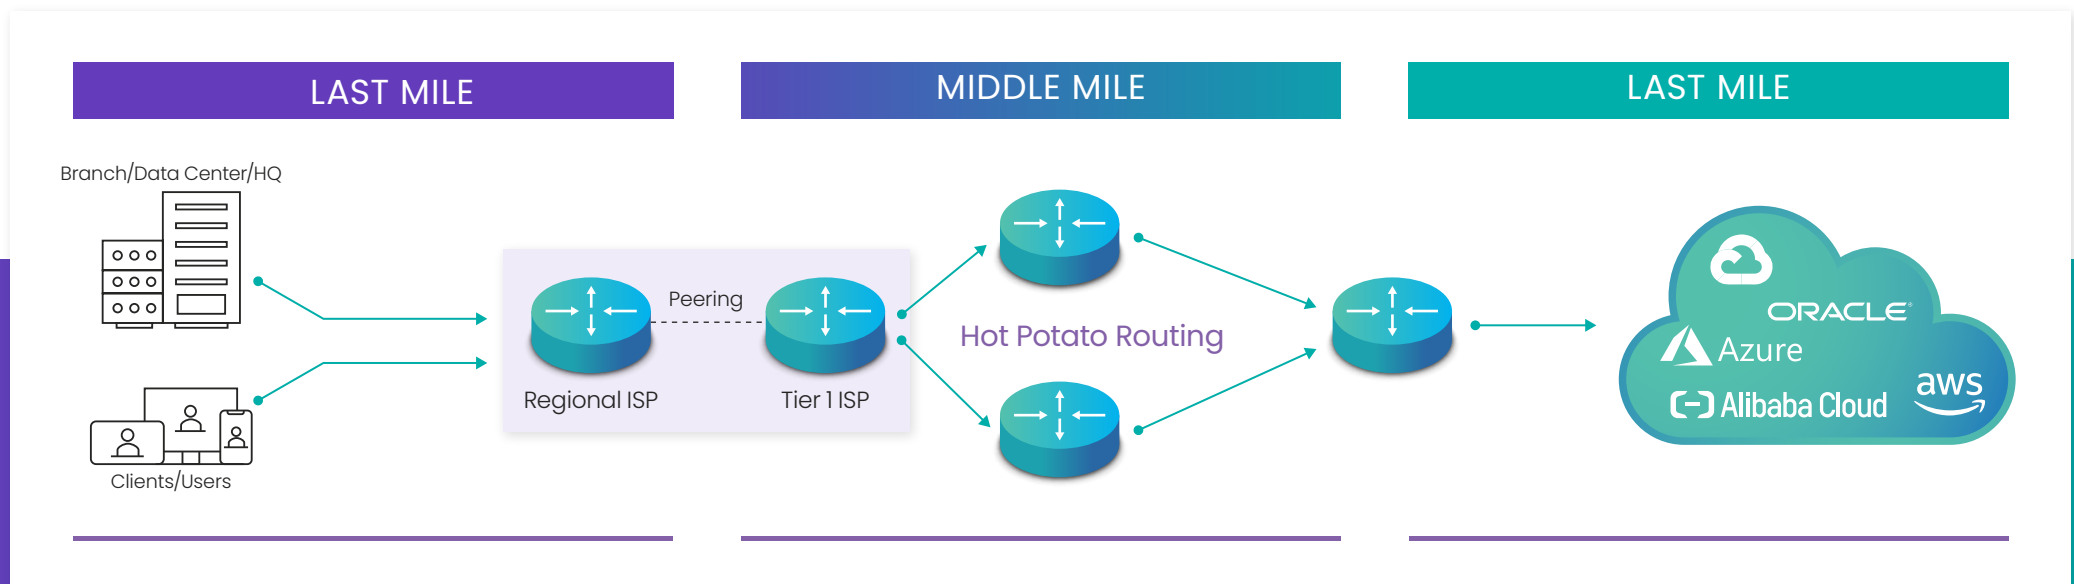
| Current ways enterprises connect to networks  | While these solutions provide certain benefits   | None of them fully meets the needs of dynamic enterprises  |
|---|--|--|
|    | <ul style="list-style-type: none"> <li>✓ Guaranteed throughput</li> <li>✓ Low latency</li> <li>✓ No CAPEX</li> </ul> | <ul style="list-style-type: none"> <li>✗ Expensive</li> <li>✗ Slow deployment</li> <li>✗ Inflexible</li> </ul>   |
|    | <ul style="list-style-type: none"> <li>✓ Rapid deployment</li> <li>✓ Flexibility</li> <li>✓ No CAPEX</li> </ul>      | <ul style="list-style-type: none"> <li>✗ Latency, Throughput, Jitter, Reliability and Security</li> <li>✗ Best effort</li> </ul>                           |
|   | <ul style="list-style-type: none"> <li>✓ Optimized throughput</li> <li>✓ Real-time visibility</li> </ul>             | <ul style="list-style-type: none"> <li>✗ Slow deployment</li> <li>✗ Inflexible</li> <li>✗ Not application agnostic</li> <li>✗ Significant CAPEX</li> </ul> |
|  | <ul style="list-style-type: none"> <li>✓ Rapid deployment</li> <li>✓ Flexibility</li> </ul>                          | <ul style="list-style-type: none"> <li>✗ Software overlay - No WAN</li> <li>✗ Does not address middle mile</li> <li>✗ Requires CAPEX</li> </ul>            |

Why can't I have all of this, in one solution?

Then there is the internet. The proliferation of cloud applications, when mixed with the sheer volume and variety of data traversing through the internet, creates congestion, killing application performance.

UCaaS applications, as critical they are to run globally distributed businesses, are notoriously susceptible to jitter and packet loss. Cloud services from Amazon Web Services (AWS), Microsoft Azure, and Google Cloud are equally susceptible to slowdowns. When any of the routers and softwares they run upon are impacted, it has a cascading effect on the overall service delivery and the entire experience takes a hit.

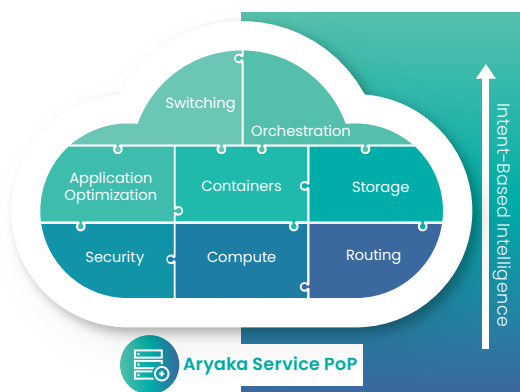
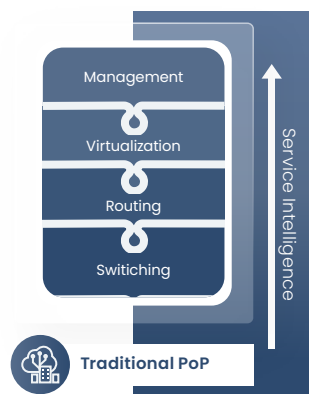
A good chunk of these speed bumps occurs at the nebulous middle-mile, that spans thousands of miles and does not belong to a single owner but is managed by multiple entities. The data has to traverse through various congested peering points and get laden with latency bottlenecks and throughput constraints while doing so. This eventually negatively impacts the entire user experience. The traditional hot-potato routing practices only add to the problem.



Many might claim that the last mile is an equally vital culprit, but the last mile is relatively short and mostly robust. Thus the impact of latency is minimized to a great extent. In fact, availability and lack of redundancy might be a bigger issue for the last mile than latency or packet loss.

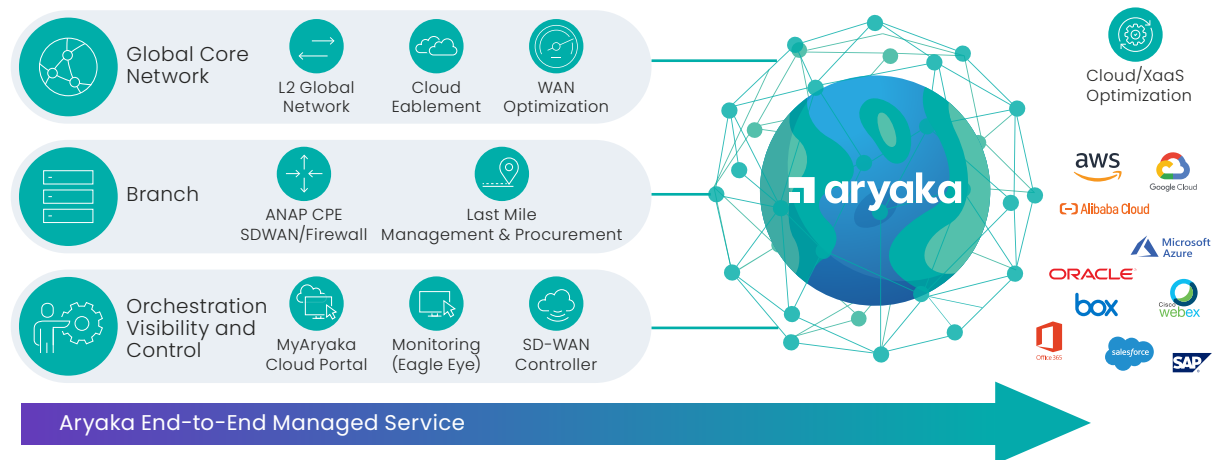
## A Unified Platform for Multi-Cloud Architectures

Imagine having a fully meshed L2 network with 40+ services points-of-presence spread out across the globe. Each PoP providing services that go well beyond internet access and include advanced first/last mile optimization, protocol acceleration, application optimization, optimal peering to cloud services and orchestration capabilities that enable true intent-based networking.



Traditional PoP vs Aryaka Service PoP

These global PoPs can easily be peered with IaaS, PaaS, SaaS and UCaaS providers and help to deliver on an optimal cloud-friendly network topology. This distributed peering model with leading cloud services providers delivers superior performance and user experience, since the Layer 2 infrastructure will always perform better compared to a direct break-out to the public internet in the branch in order to reach most business-relevant cloud-based services.



Aryaka Cloud-First WAN

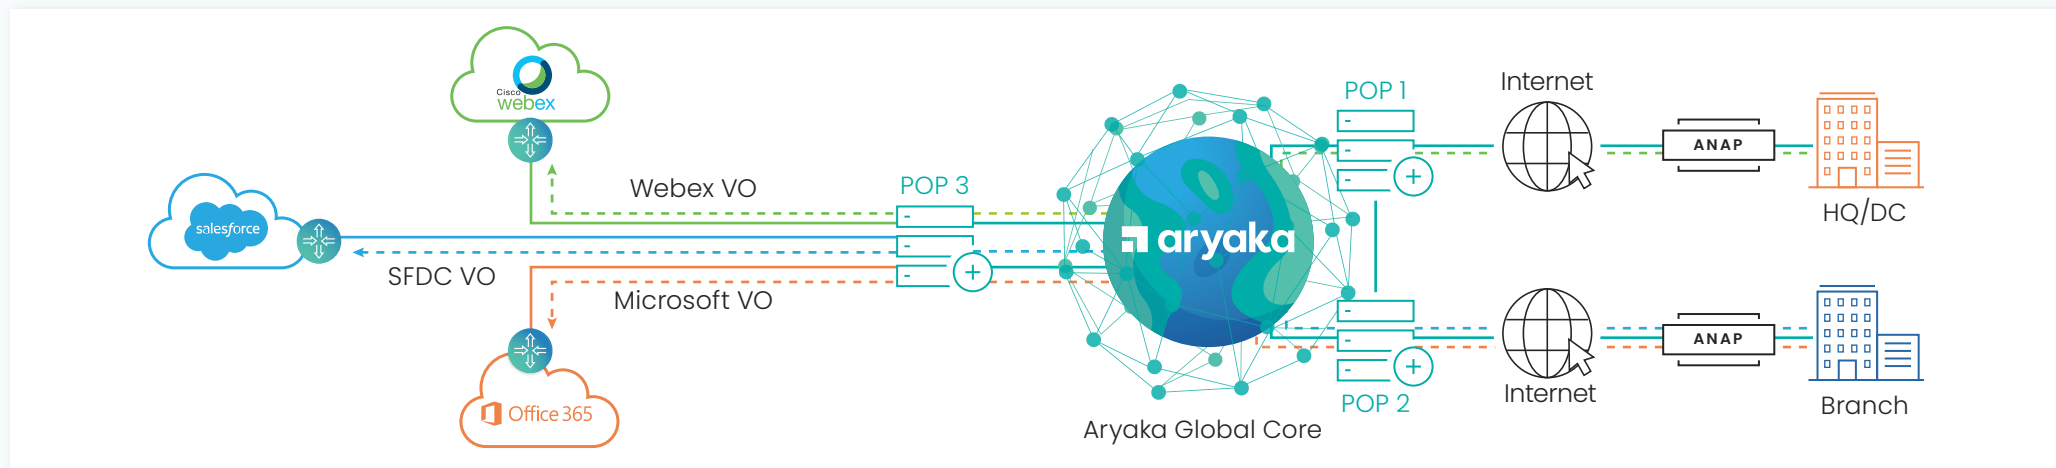
Let's see how this network infrastructure addresses the top challenges for multi-cloud deployments.

## 1. Connectivity

### The VO Model for SaaS Connectivity

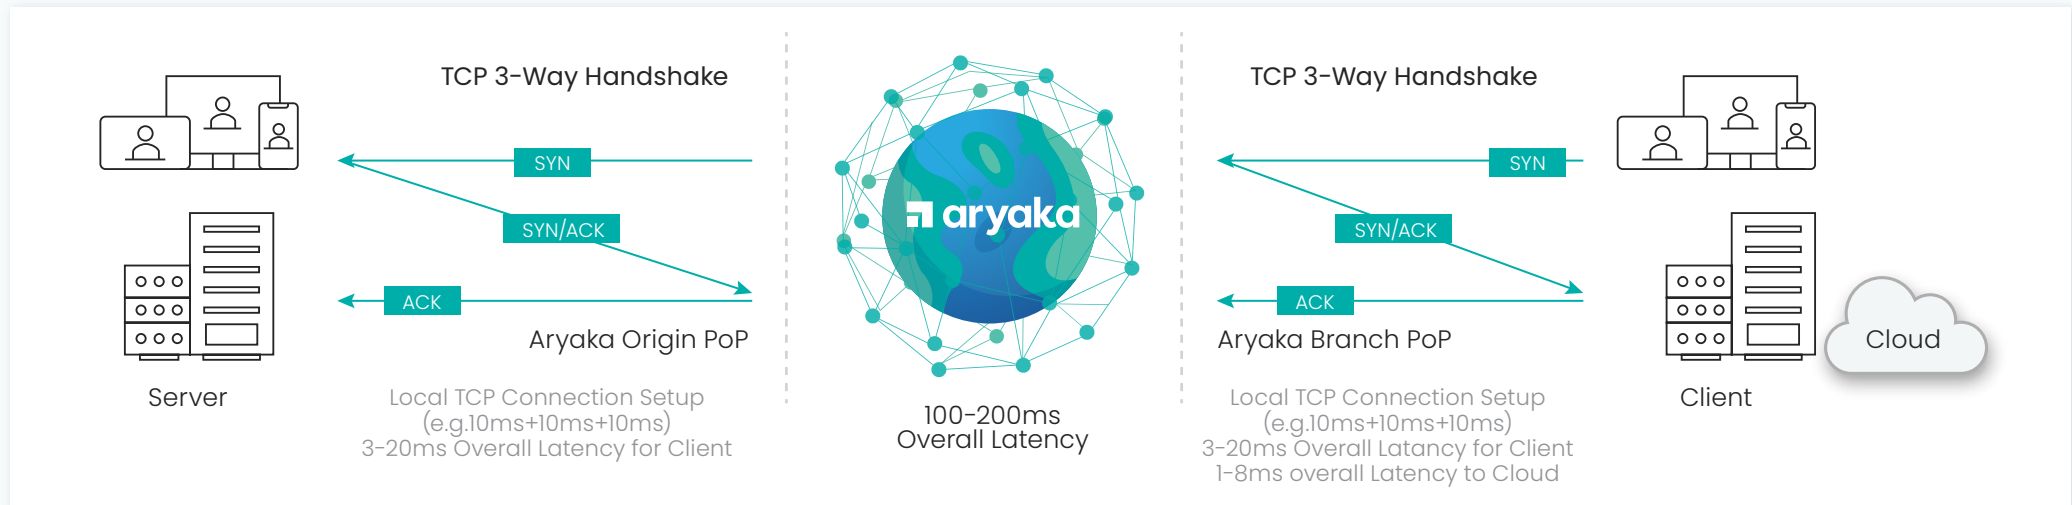
The Virtual Office (VO) model facilitates faster connectivity and improved performance for SaaS applications accessed over the internet. A VO is an Aryaka virtual router with a public IP address and Layer 4 stateful firewall capabilities that provides an optimization container and multi-segment TCP architecture to reduce the RTT.

The VO methodology creates a virtual site, instead of a real physical one, to hand off the traffic from the Aryaka PoP to the nearest SaaS/ UCaaS entry point.



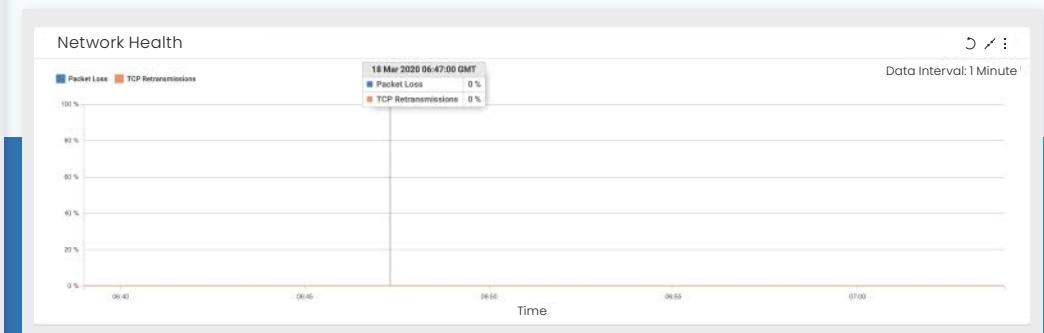
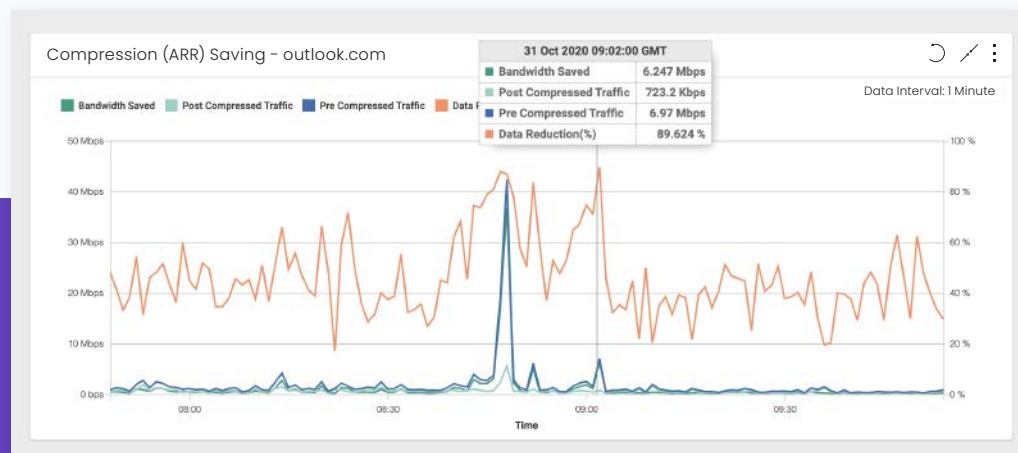
So how does a virtual office work? To begin with, it connects a private branch to a SaaS or UCaaS provider. Take Salesforce, for instance, which runs over a public server. We are well versed with the internet machinery if traffic were to be sent from a private server in Bangalore, India to the Salesforce server in the US.

With the VO model, the traffic rides over the private, optimized middle-mile and lands on the closest PoP to the Salesforce server. The VO mechanism then replicates a virtually hosted public site on the PoP itself and performs SNAT (Source Network Address Translation) for the traffic going out to the Salesforce server. SNAT NATs a private IP to Public IP so that the clients can talk to the public Salesforce server.

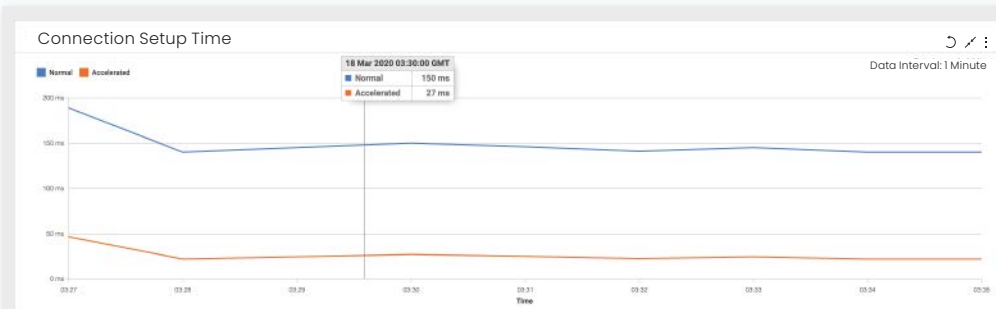


This entire process mitigates the middle-mile limitations and optimizes the traffic that traverses the Aryaka core network, boosting SaaS application performance up to 20X with as low as 0% packet loss and performance degradation. The immediate benefits look somewhat like this.

The non-optimized Outlook traffic that hogged 7 Mbps of bandwidth shrinks down to 723 Kbps post-optimization over the Aryaka core with as low as zero percent packet loss.



Upto 6x faster TCP Connection set-up time and a stable core latency that facilitates consistent application performance.



Up to 6x Faster TCP Connection Set-Up Time



Stable Latency over the Aryaka Core

## Fuze UCaaS Service

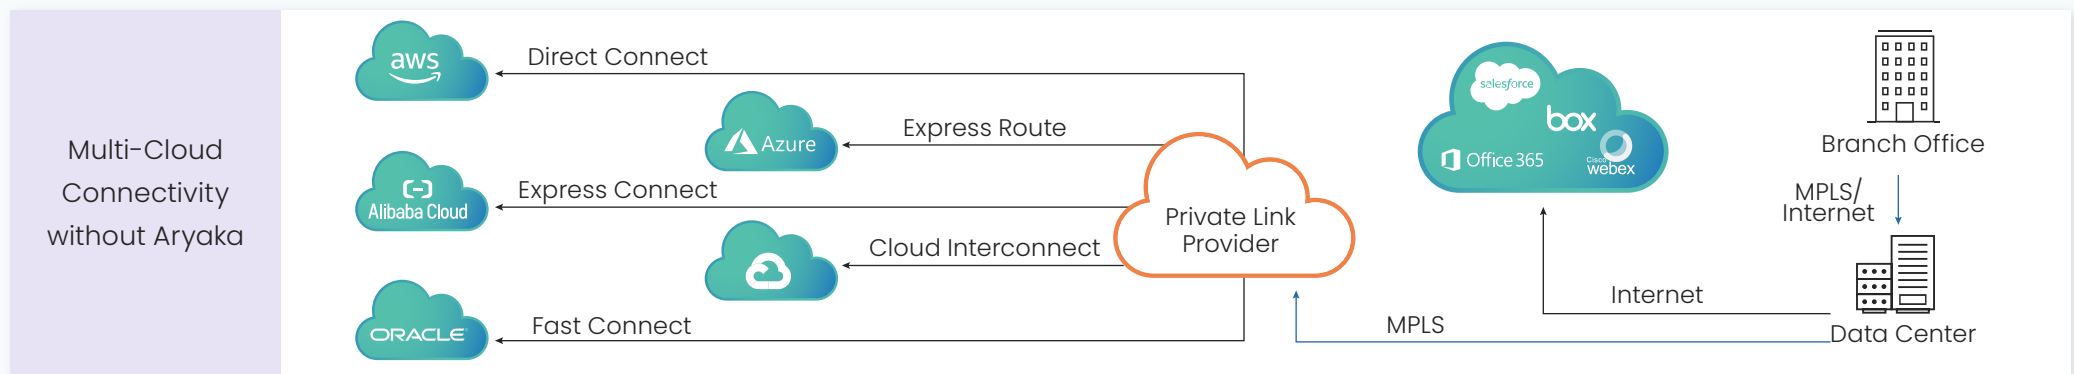
Aryaka implements **connectivity to Fuze's platform** using its Virtual Office (VO) implementation as well as direct Layer 2 peering. The UCaaS traffic is marked and injected with all Aryaka LinkAssure intelligence before being optimally routed to the closest Aryaka PoP. Traffic then traverses Aryaka's SLA-driven backbone to the nearest Fuze peering point, bypassing the jitter, latency, and packet loss endemic to the public internet while offering several features to improve the performance and resiliency of the last mile. The intelligence to best optimize the traffic can be tailored based on the customer requirements and underlying circuits.

Note that Aryaka also has a **global technology partnership with 8X8**, Inc. that helps its users with blazing fast 8X8 performance, and of course supports all leading cloud collaboration platforms.

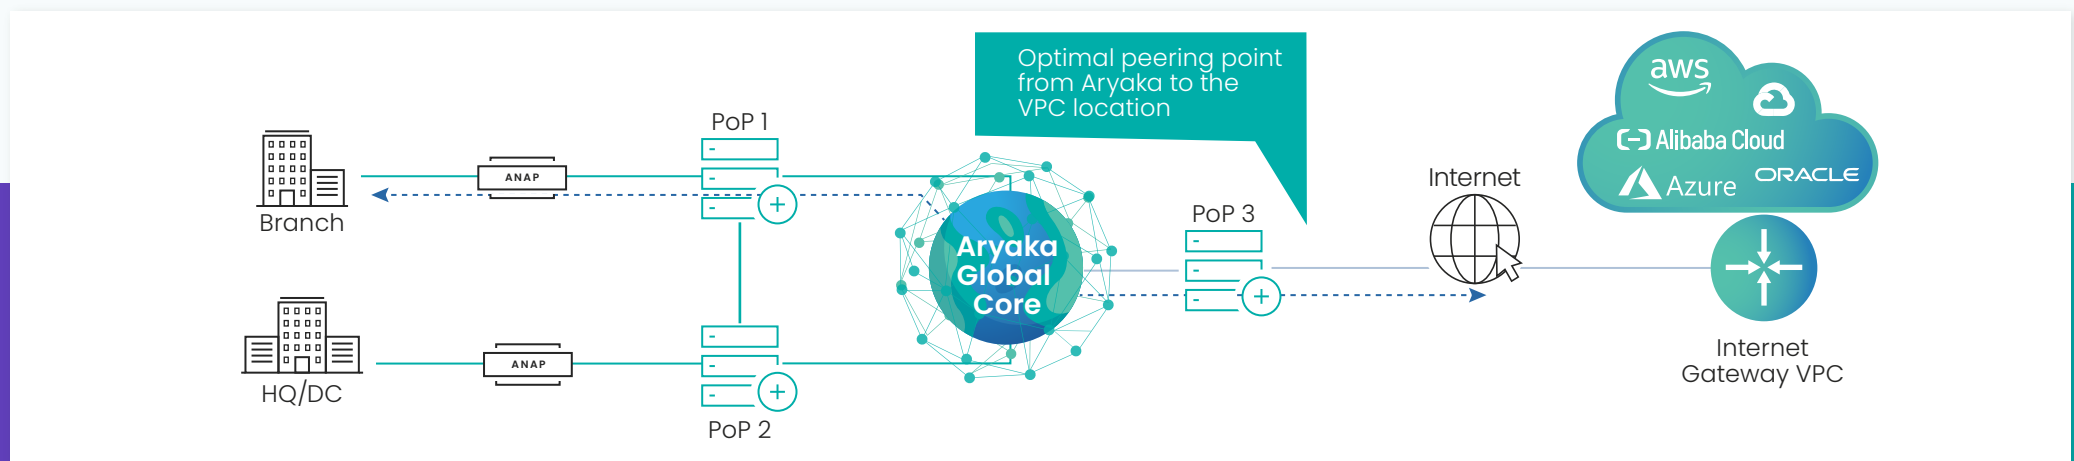
## Private Connections to IaaS Providers

According to Gartner, AWS, Azure, Google Cloud and Alibaba Cloud together control more than two-thirds of the IaaS Market share — which is indicative of the fact that most IaaS activities are powered by one of these giants.

To match the flexibility and agility that is synonymous with cloud offerings, Aryaka provides two different ways to connect to IaaS providers. The first is a direct connection to the top brass IaaS Platforms.



The second is an IPsec tunnel from the nearest PoP router. The PoP closest to the IaaS location acts as the transit point for the inter-cloud traffic, thereby providing an efficient alternative to the traditional solutions.



## AWS

Most Aryaka PoPs are located in the same AWS IaaS data centers, giving users the choice to connect to their VPCs using either Direct Connect or an IPsec VPN (IKEv1/IKEv2). Direct Connect links your internal network to an AWS Direct Connect location at Layer 2 over a standard Ethernet fiber-optic cable, where one end of the cable is connected to your router, and the other to an AWS router.

This connection enables you to create virtual interfaces directly to public AWS services (for example, to Amazon S3) or to an Amazon VPC, bypassing internet service providers in your network path.

## Azure

Similar to AWS, most Aryaka PoPs are co-located in the same Azure IaaS data centers. Users can connect to their VNETs using ExpressRoute or an IPsec VPN (IKEv1/ or IKEv2). ExpressRoute offers Layer 2 connectivity between your on-prem network and the Microsoft Azure Cloud through a connectivity provider.

Connectivity can be from an any-to-any (IP VPN) network, a point-to-point Ethernet network, or a virtual cross-connection through a connectivity provider at a colocation facility.

ExpressRoute connections don't go over the public internet. This allows ExpressRoute connections to offer more reliability, faster speeds, consistent latencies, and higher security than typical connections over the internet.

## Google Cloud

Aryaka's partnership with Google Cloud enables us to provide the most complete multi-cloud footprint of any SD-WAN provider, with direct connectivity to the Google Cloud.

Just like Azure and AWS, Direct Connect links the internal network to a Google Cloud router at Layer 2 over a Ethernet fiber-optic cable, where one end of the cable is connected to your router, and the other to a Google router, forming up a GRE-BGP connection between the routers. BGP helps in learning and advertising the route information between ASs (Autonomous Systems). This enables organizations to eliminate the need for ISPs and get a privileged lane to Google services.

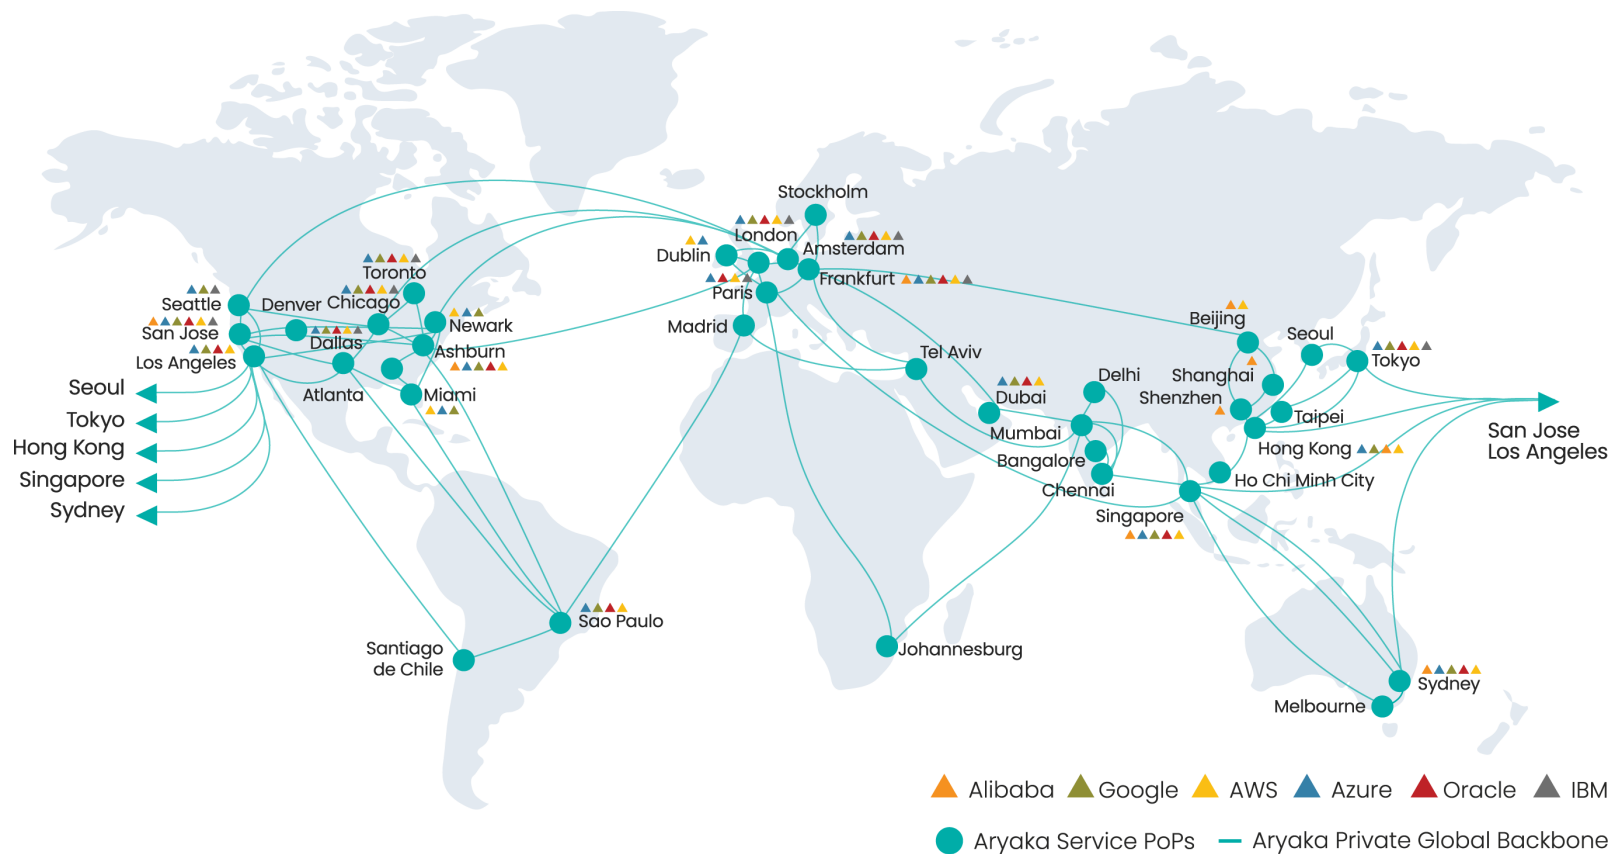
## Alibaba Cloud

With 63 availability zones in 21 regions, Alibaba Cloud is a giant that covers the largest IaaS market share in the Asia Pacific region. The partnership between Alibaba and Aryaka lets enterprises outside of the Alibaba Cloud footprint leverage Aryaka's end-to-end global WAN for access into China and elsewhere.

## Oracle

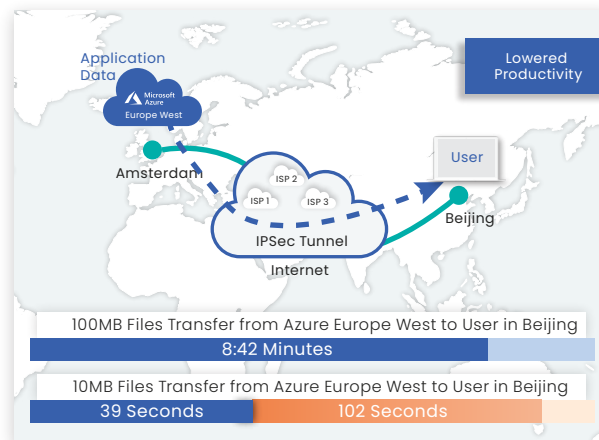
Aryaka supports connectivity to Oracle via both FastConnect and IPsec VPNs, delivering up to 8 times faster application performance compared to the existing network options, providing access to business-critical applications in 30 milliseconds or less.

### Seamless Multi-Cloud Connectivity: IaaS, PaaS, SaaS



The End Result? Much faster application performance with reliable and optimized connectivity.

## Before Aryaka:

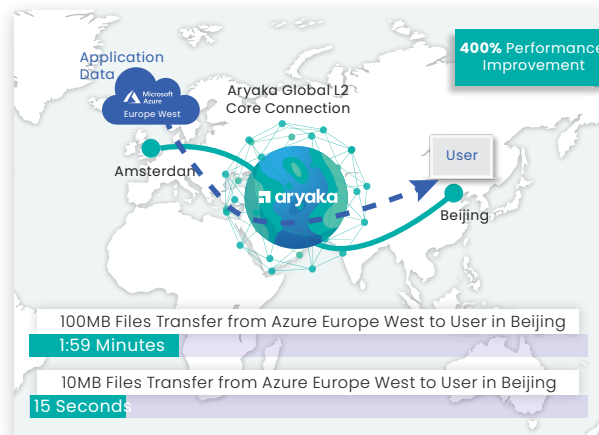


Remote users often complain about problematic and/or unpredictable performance when they access cloud-based applications and data via IPsec tunnels.

The cause: global distances with unpredictable latency, jitter, and packet loss that severely impacts TCP performance.

File transfer measurements shows a real-world test scenario between a user in Beijing and application data residing in Azure Region Europe West.

## With Aryaka:



The Aryaka solution provides a 400% measured improvement over a regular IPsec tunnel.

The Aryaka Global L2 Network ensures deterministic delay, jitter and eliminates packet loss. Furthermore, it implements traffic optimization.

Global companies choose the Aryaka solution to accelerate cloud application performance, both for remote branches as well as mobile, remote users.

## The Equinix ECX Advantage

Equinix is the global leader in co-location and data center services. Today, Aryaka runs 16 of its PoPs in Equinix data centers. At 15 of these 16 PoPs, Aryaka has purchased Equinix Cloud Exchange (ECX) Ports to set up private connection on-ramps to IaaS providers including Azure, AWS, Oracle Cloud, Google Cloud, and Alibaba Cloud. In addition to the Equinix DCs where Aryaka is co-located, Aryaka can set up remote virtual circuits to 29 other Equinix DCs that are on the

ECX Fabric, allowing Aryaka to extend its cloud reach. Thus, Aryaka can scale its IaaS on-ramps significantly, i.e. today up to 4x worldwide. Depending upon region and in which locations Equinix has enabled ECX Fabric. The partnership enables Aryaka to establish local or remote virtual circuits over the ECX Fabric with the IaaS providers in two ways:

#### a. Local Virtual Circuit:

This circuit is established if the IaaS provider is in the same ECX-enabled metro as the Aryaka PoP.

#### b. Remote Virtual Circuit:

This circuit is established if the IaaS provider is in a different ECX-enabled metro as the Aryaka PoP, i.e. in an ECX-enabled metro where Aryaka does not have a PoP.



How does the new partnership benefit customers?

- The new partnership allows Aryaka customers to gain an expanded WAN from Aryaka to connect to IaaS providers.
- The partnership also allows Equinix ECX customers to set up private connections with Aryaka's WAN.

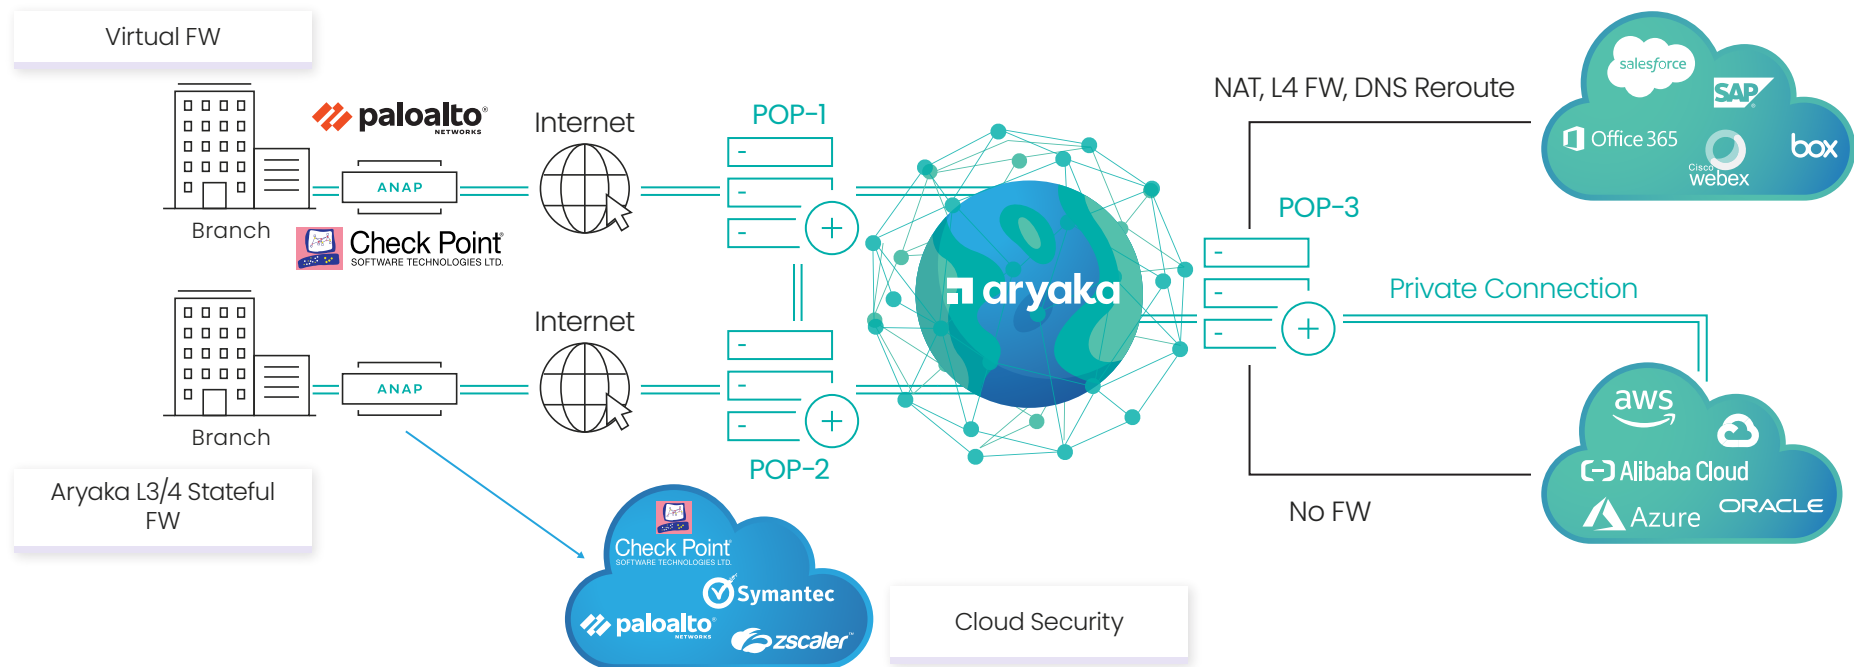


## 2. Security

A single security solution can barely suffice all enterprise's requirements, for reasons spanning between architectural concerns down to regulatory requirements. Most prominent industry analysts also swear by the need for a multi-layered approach to security.

So how do we do it? First, there is the foundational aspect of SD-WAN security. In the branch, you need a basic stateful firewall and policy-based segmentation. Aryaka integrates both functions into CPE, the ANAP (Aryaka Network Access Point).

Second, Aryaka partners with cloud security leaders like Check Point, Zscaler, Palo Alto Networks, Symantec, and others to deliver choice and easy integration. In a nutshell, Aryaka's open approach to multi-layered SD-WAN security delivers on the option that enterprises prefer and which emerging Zero Trust security postures require as a foundational enabler.



Aryaka Security Architecture: A World of Options

### Cloud Security

Secure, local Internet breakout is enabled through Check Point CloudGuard Connect, Palo Alto's Prisma Cloud Security Suite, Symantec's Web Security Service, and Zscaler's Cloud Security, protecting all ports and protocols without appliances.

### Micro-Segmentation

Micro-segmentation adds to the ANAP's zones functionality. On-site, VLANs provide the local segmentation for internal and DMZ zones. Micro-segmentation extends this across the Aryaka core network via a BGP-driven VRF 'lite' functionality.

### Virtual Firewall

Aryaka's ANAP supports NFV functionality for additional SDN-delivered services. We are partnering with multiple Tier-1 security vendors such as Palo Alto Networks and Check Point, enabling choice. Management of physical firewalls part of the SD-WAN deployment is also supported.

### Secure Remote Access

Aryaka's Secure Remote Access is the first clientless SD-WAN for software-defined remote access. It significantly enhances both on-premises and cloud/SaaS application performance for the remote and mobile workforce without requiring additional hardware or software clients.

### Aryaka Core Protection

In parallel, the Aryaka private core delivers partitioned connectivity to all enterprises, encrypting the data and protecting against DDoS attacks. Within the branch, enterprises have access to Syslog and Netflow logging, and at the network level, the MyAryaka cloud portal provides a single pane of glass for service configuration, monitoring and health.

### Stateful Firewall

The ANAP includes a virtual stateful firewall that delivers north-south access protection as well as a 'zones' capability offering site-segmentation to secure east-west traffic within the branch. The ANAP aligns with the evolving Secure Access Service Edge (SASE), offering a choice of edge and cloud security.



### 3. Visibility

Aryaka bridges this gap between DevOps and NetOps with centralized monitoring and orchestration. The ability to provision and make changes to the network from a centralized application dramatically speeds up the network rollout time.

Our **SmartInsights** provide deep insights into the state of a customer's WAN, accessible via MyAryaka, a cloud-based orchestration, and **visibility portal** that gives customers the ability to configure, control, and manage all Aryaka's **Managed SD-Services**.

The MyAryaka portal also offers deep, end-to-end network and application visibility for your business and provides APIs for graph data as well as an option to embed URLs that can be used to integrate to a customer's existing visibility tools.



#### Monitor

- Application Insights • Traffic insights • Optimization Benefits • Latency • Network Health • Feature related visibility • CIFS/SMB • Cloud Security • All Sites & Links • QoS • Domains



#### Configure

- Sites • Links • QoS • NAT/Firewall • Domains • Route Controller • SSL Certificates • Feature Configuration – such as Cloud Security • Network Object Groups • Manage Users • Manage Orders



#### Status

- Edge, IPSec Status • Links Status • ANAP related Status information • Route Tables

## Conclusion

Multi-cloud strategies are witnessing a massive surge in uptake, partly because the coronavirus situation has accelerated enterprise interest and adoption of cloud. As the majority of the workforce goes remote and on-the-move, they still need access to their applications and resources.

This requirement is pushing organizations towards a cloud-based model that lets them expand their infrastructural capabilities without wasting precious time in planning. The idea is to shift the processing power of applications from localized computers to a global network of connected facilities, creating a platform that can deliver any type of service, anywhere in the world.

As these workloads continue to spread across multiple cloud platforms, it also means more data and more real-time performance flowing through the system and, thus more information to capture and work upon. You're looking at multiple cloud instances and hardware stacks managed by multiple vendors running different OS, applications, and mechanisms to provide the security, visibility, and connectivity you need to keep running the marathon. Needless to say, when you get down to the nitty-gritty of management, this mish-mash of vendors can be a nightmare.

Which begs the obvious question — what if it could all be taken care of with one single vendor? One hand to shake across all your cloud connectivity, security, and visibility requirements.

## A Cloud-First WAN for Cloud-First Enterprises

[Click here](#)



to see how we can take your multi-cloud planning from 0 to 100 within days.

## 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](http://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 Demo →](#)



PO Box 610307 San Jose, CA 95161

Follow us on :

