aryaka

The Cloud Connectivity Cookbook

Table of Contents

1 Introduction

Once firmly constrained in brick and mortar, the perimeters of the enterprise network have become porous and now extend globally to a distributed pool of resources and a nomadic workforce that thrives on cloud-based communication, collaboration, and development platforms. Cloud-based solutions, being innately able to be rolled out and scaled remotely, hold a unique position to help enterprises sail through these disruptive times and keep the lights on. Swiftly and surely, even the most reluctant enterprises are now turning to the cloud. And, to an extent where most organizations on an average use approximately 1,935 cloud services.

Yes. You read that correctly. 1,935!

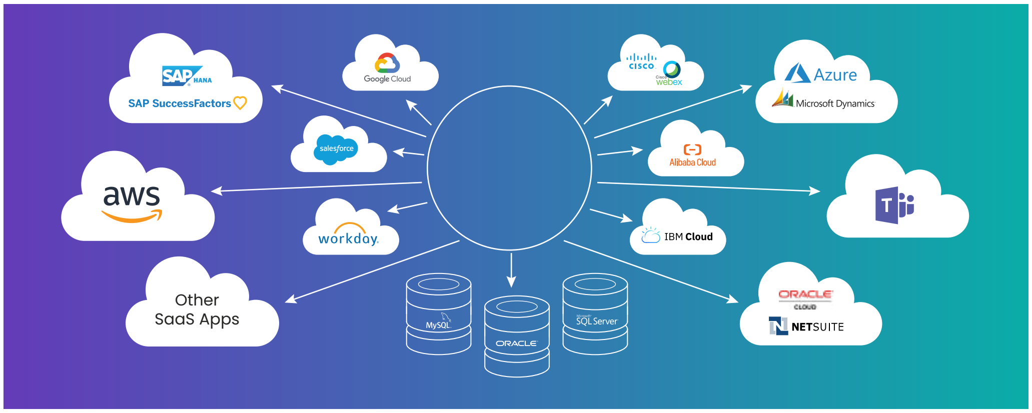

Since a multi-cloud strategy is explicitly designed to unify all the IT orchestration under a single framework, connectivity remains the holy grail of multi-cloud networking — for consumers and providers alike. And this optimized connectivity is critical as enterprises deploy a SASE architecture. In March 2023, Gartner published a report emphasizing this, specifically calling out issues arising when connectivity is not addressed at the same time as security. The figure below maps these issues to Aryaka's solution.

Leveraging Cloud Connect Network Infrastructure to Improve WAN Services at SASE Solutions for Cloud Workloads

Critical Insights

- The internet as the default WAN for user-to-cloud workloads is not reliable, and end users' self-adopting cloud connect infrastructure to enhance WAN services makes management complex.
- The poorly performing WAN interconnecting various types of SASE points of presence (POPs) greatly slows overall SASE solution performance because user traffic must go through one or multiple POPs as inspection points to fulfill network and security policies.

Cloud Connect to Improve WAN for User-to-Cloud Workloads

Impacts on Product Leaders

- Improve SASE connectivity experience for user-to-cloud workloads by preinstalling and offering choice of cloud connect infrastructure and by integrating such infrastructure to reduce customer costs and management complexity.
- Develop your WAN architecture interconnecting various SASE POPs by investigating the SASE hosting strategy and interconnection performance, bandwidth, location, and capacity requirements between various SASE solution components.

1 Internet as default WAN is NOT reliable and self-adopting cloud connect is COMPLEX

3 **Aryaka integrates Multi-Cloud Connectivity as managed services**

2 Disjointed Architecture SLOWS DOWN performance to Cloud workloads

4 **Aryaka single-pass architecture for Unified SASE**

In the following pages we'll explore the different aspects of multi-cloud networking and how Unified SASE unlocks the fastest path to value in a cloud-first world.

2 The Importance of Multi-Cloud Networking

Cloud technology and network connectivity are like hamburgers and fries. Always better when served together. However, as enterprises move away from dedicated servers in private data centers to an on-demand pool of computing capacity, latency, jitter, packet loss, and other inhibitors to a peak user experience tag along as a package deal.

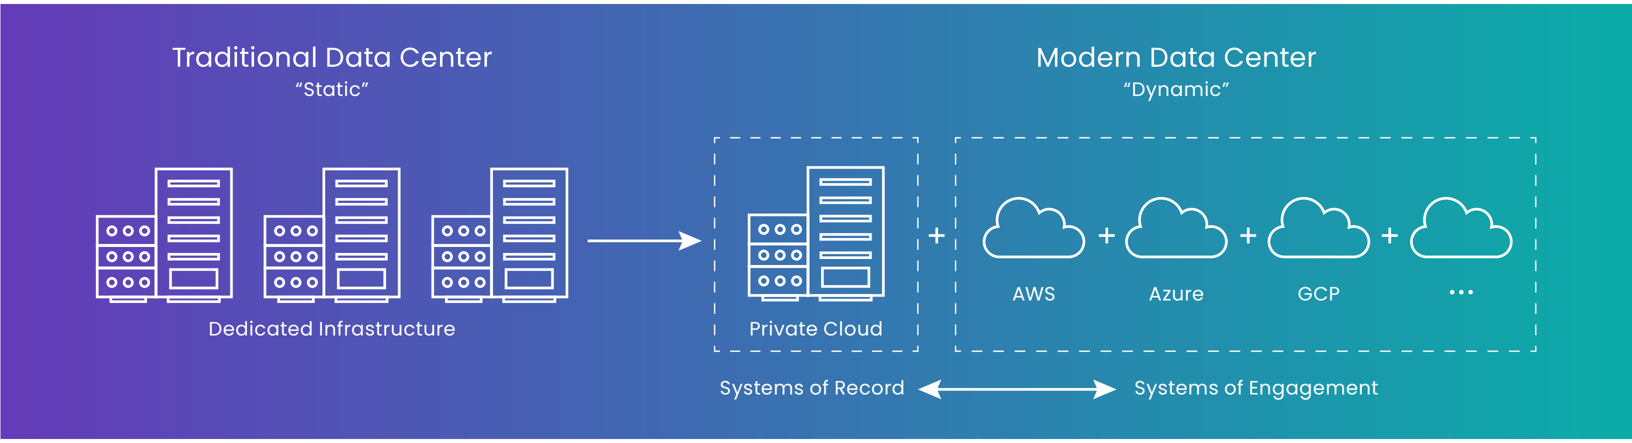

Imagine the hassle for a logistics multi-national whose ERP system goes into a never-ending loop as its consignments are unloading on a dock, or an oil & petroleum company whose operations team is fervently waiting to receive mission-critical files and blueprints on an offshore drilling rig. Your team or unit is depending on your help. But instead, because of poor network connectivity and slow application performance, you risk ruining the entire mission.

Cloud sprawl with poor network connectivity adds cost and risk. The challenge for most enterprises then is to deliver these cloud-based applications with consistency. The combination of dynamic IP addresses and exponential growth in bi-directional traffic, coupled with the rise of microservices, needs a network that can dynamically respond to changing user and traffic demands in real-time.

For a multi-cloud plan to take off, there needs to be a robust network that not only connects these cloud instances and different SaaS and IaaS providers but can also extend these resources at scale across different branch offices and global locations.

3 Getting Started: What You Need to Know

Now that you know multi-cloud is the conduit to the future of application hosting, what next? Multi-cloud adoption is a complex undertaking and should begin with a future proof assessment of your requirements. Listed below are the top five pointers you need to know before venturing out into multi-cloud networking.

3.1 Browse Through the Cloud Catalog

Begin with asking questions that include:

○ **Cloud Location** - These cloud platforms are accessed by a range of remote users and distributed branch offices. If data is transferred from a single data center, the disparate users are bound to experience performance variations. In a multi-cloud set-up, network architects need to ensure that the users connect to the SaaS / IaaS data centers that sit closest to their regions and end-users. This way, data can be served with minimum server hops. This is something that factors in deeply when trying to deliver a unified user experience while catering to data across continents and remote locations.

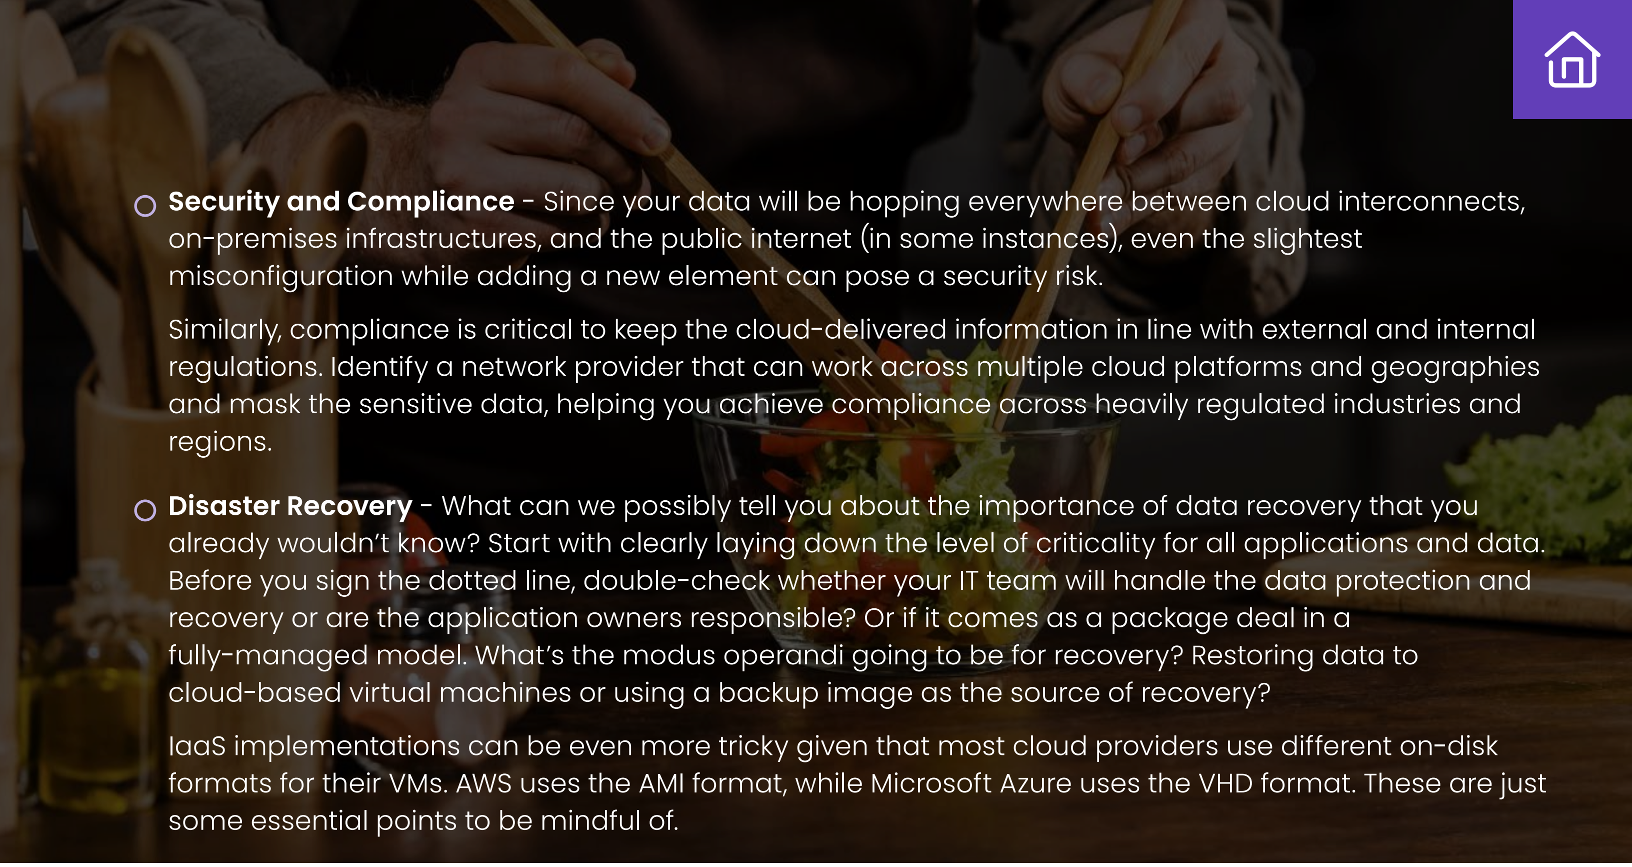

- **Security and Compliance** - Since your data will be hopping everywhere between cloud interconnects, on-premises infrastructures, and the public internet (in some instances), even the slightest misconfiguration while adding a new element can pose a security risk.

  Similarly, compliance is critical to keep the cloud-delivered information in line with external and internal regulations. Identify a network provider that can work across multiple cloud platforms and geographies and mask the sensitive data, helping you achieve compliance across heavily regulated industries and regions.

- **Disaster Recovery** - What can we possibly tell you about the importance of data recovery that you already wouldn't know? Start with clearly laying down the level of criticality for all applications and data. Before you sign the dotted line, double-check whether your IT team will handle the data protection and recovery or are the application owners responsible? Or if it comes as a package deal in a fully-managed model. What's the modus operandi going to be for recovery? Restoring data to cloud-based virtual machines or using a backup image as the source of recovery?

  IaaS implementations can be even more tricky given that most cloud providers use different on-disk formats for their VMs. AWS uses the AMI format, while Microsoft Azure uses the VHD format. These are just some essential points to be mindful of.

3.2 Direct Connection to the Cloud Provider? Think Again

Why shouldn't I do it myself? The obvious first question. If your cloud provider of choice has a nearby presence and you can get access to their on-ramp location, you should probably consider a direct interconnect. However, this is under the prerequisite that you only have one provider to connect to, and you may have requirements spread across multiple regions.

Approaching cloud interconnect via a network service provider, on the other hand, gives you more agility and the option to access multiple cloud providers and their global regions from a single point, irrespective of your location.

3.3 Dedicated Connectivity is a Must Have

The essence of multi-cloud networking lies in the actual fiber of connectivity. How well can you extend your private infrastructure to the cloud provider's network? How fluidly can data overcome the integration barriers and flow between different CSPs? How good is the bandwidth throughput for applications working with large data sets? When your data travels point-to-point, it mitigates all the insecurities around traversing the internet, and that's precisely why direct and private connectivity is an absolute must-have. And this is the architecture outlined in the Gartner report described earlier.

3.4 Look Out for Edge Intelligence

Be it HTTPS traffic, external IDs or a banana, it is all the same to a traditional router. When your mission-critical and non-critical data together hops between CSPs and branch offices, there needs to be a clear demarcation line of priority between critical business apps and lower priority traffic. Edge intelligence lets you create application-level and user-level policy management, essential for traffic classification, prioritization, and steering. Moreover, as we move away from the traditional hub and spoke architecture, placing intelligence close to where data is created makes more sense.

Note that the Aryaka SASE architecture extends security intelligence across the hybrid edge, both via our Hyperscale PoPs, as well as at our managed CPE, the ANAP.

3.5 Beef Up the Underlay

You wouldn't put a V8 engine in a food cart…right? The physical network underlay does all the heavy lifting when it comes to moving data between clouds, branches, and data centers. In most cases, the cloud-based services typically ride on the internet, not that they necessarily should. As your multi-cloud plans begin to take shape, make sure to have a robust underlay before building overlay functionalities on top of it.

Which raises an obvious question — should your data be traversing over L2 or L3? Let's discuss this in the next chapter.

4 L2 or L3: That Is the Question

So which network underlay works best for connecting to a dispersed set of cloud service providers? Considering that networks are built in layers, should it be L2 or L3? The L3 layer connectivity is a tunnel over the internet, while the L2 layer is more like wire-to-wire Ethernet connectivity providing a more robust and secure posture for passing the traffic.

Referencing the Gartner research again, to ensure application performance, the hand off between the SD-WAN/SASE connectivity and the public cloud should be at L2. However, depending upon the needs of the enterprise, the connection from the end-user, either on-premises or remote, and the first PoP, and then to the PoP connecting to the CSP, can be across a managed L2 or L3 infrastructure. By managed L3, this is an Enhanced Internet service without ISP peering and with SLAs. It is not the public internet. Note that a third option is where remote workers connect via their ISP to the cloud provider. Here, the middle mile to the cloud hand off is a function of the ISP, which may be L2, Enhanced Internet, or the public internet.

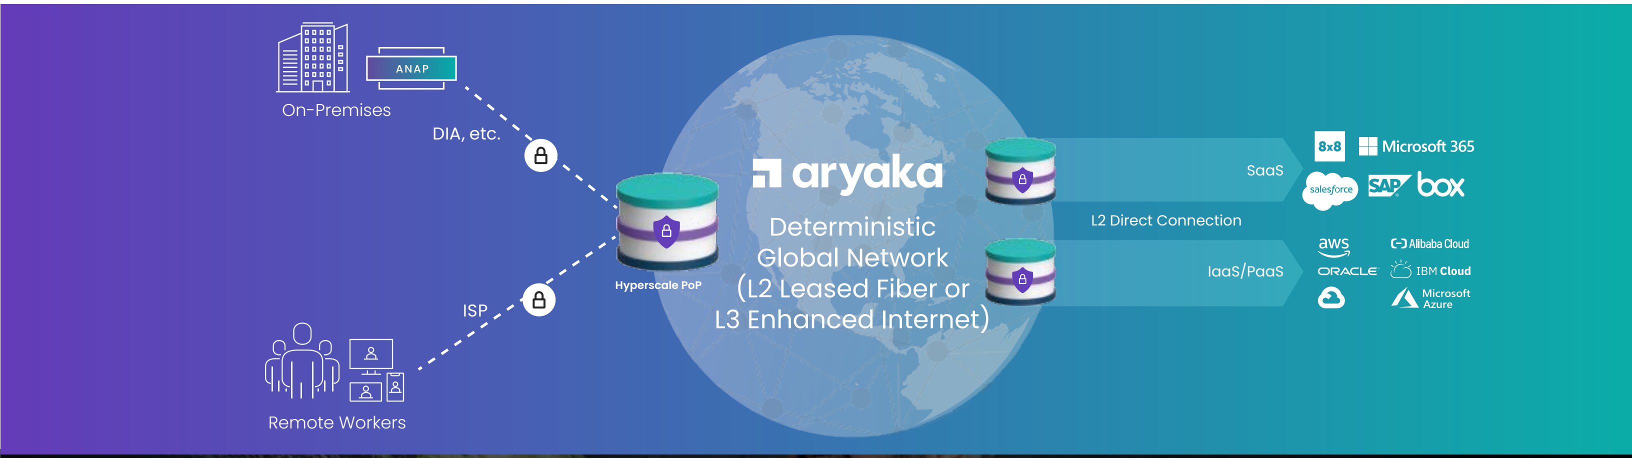

The table below summarizes the differences between the two approaches.

	L3	L2
Transport Medium	Internet as transport	Private layer 2 link as transport
Security	Secure public transport	Highly secure, SLA, private transport
Performance	Dedicated SLAs, QoS	Dedicated SLAs, QoS
Routing	Devices forward data packets based on layer 3 information (eg: IP Address)	Devices forward data packet based on layer 2 information (eg: MAC ID)
Management	Managed end-to-end by Aryaka	Managed end-to-end by Aryaka

5 Limitations Around Current Connectivity Architectures

5.1 Enterprise WAN Connecting Individually to Each Cloud and Application

- To begin with, you need detailed knowledge of terminologies such as IP address schemes, ports, access control lists, and numerous other network parameters.
- Managing a multitude of single-function devices and distributed network configurations on individual routers can be a nightmare. Policies will have to be rolled out on each device with a complete do-over in case of assigning new policies.
- Branch office and remote location set-ups will require additional hardware, successively driving the cost up. This includes manual scaling and management of bandwidth for each CSP.
- Ultimately, the solution does not meet the needs of an evolved SASE architecture due to poor network SLAs (if any), lackluster or complex security, and no traffic optimization.

5.2 The Cloud Exchange Model (DIY or via a Telco)

The cloud exchange model was developed to provide stable latency, predictable bandwidth and better connection to meet the high throughout requirements. The idea is put your regional hubs in exchange provider's facilities, which then connects the exchange's network to a variety of cloud providers. However, this model doesn't come without some limitations:

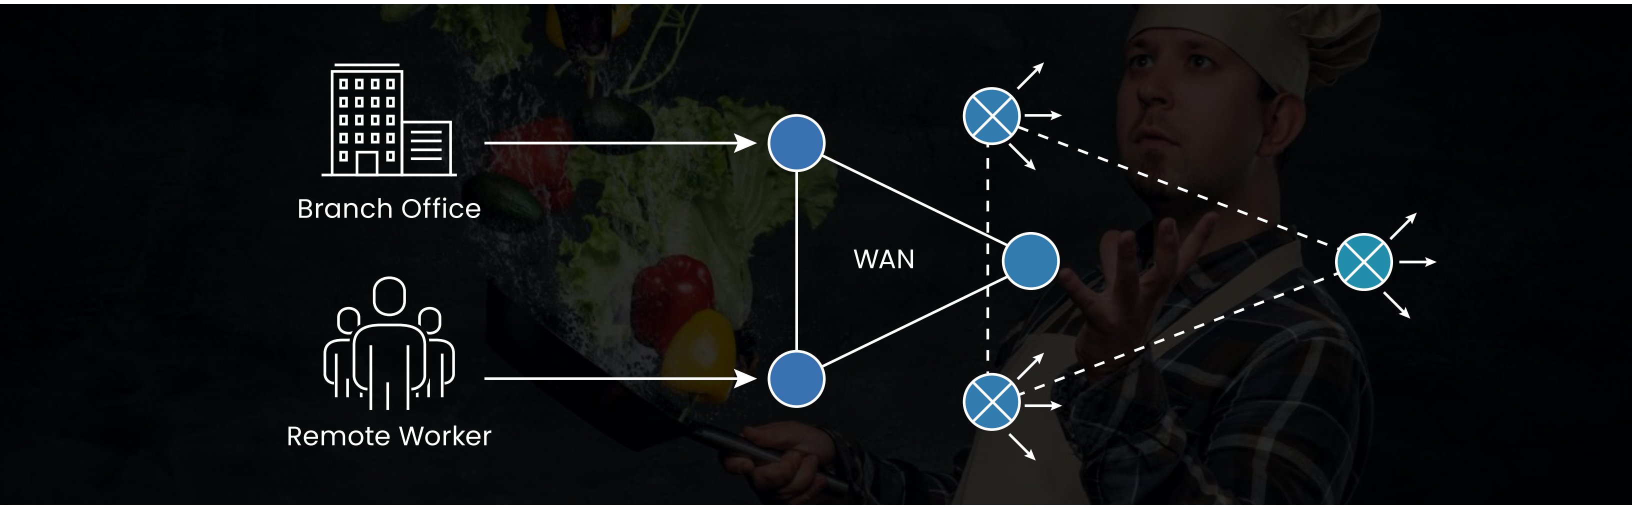

- Though the exchange model successfully connects to different CSPs, it falls short in providing regional connectivity. Furthermore, the enterprises usually connect via a separate network to the exchanges, separate from its WAN. This means more links to manage, more potential points of failure and more cost. The complication escalates further when trying to connect remote users to the exchange, who may or may not be in close proximity to an exchange center.
- Say If you want to connect from the Middle East to Europe, how will the GDPR compliance be met? Moreover, working out a redundant direct connection to the exchange PoP and setting up the support mechanism is time consuming and might not work if you have frequently moving project sites.

- Security remains an overall concern. This is due to the 'shared responsibility model,' which means that although the CSP ensures security of the infrastructure, you need to understand how to secure your data and applications. Anything less, and you are at risk.
- Given the sheer number of requests and demands placed on these exchanges, and the short span of time to process those, there is always a requirement for high throughput.
- Risk of a single point of failure, dependent on exchange redundancy.

5.2 The Unified SASE Model

A model that converges the flexibility of the enterprise WAN with the reliability of a cloud exchange, minus the hassle of management. It delivers:

- Complete global coverage, and optimized connectivity not only to CSPs but also to branch offices, remote users and on-premises deployments.
- Centralized orchestration that enables you to push policies, install features and bring up remotest site at the click of a button. A managed service completely circumvents all the deployment woes.
- A complete end-to-end security posture managed by the service provider over a highly virtualized and on-demand architecture.
- Each application receiving the attention it needs along with robust network SLAs.
- However, things do get gloomy if this technology is piggybacking on traditional networks as underlays. We are pointing at organizations who still rely on legacy connections to keep the internal and external operations afloat.

6 Connectivity 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 either one of these giants.

To match the flexibility and agility that is synonymous with cloud offerings, there are two ways of connecting to most IaaS providers. The first is a direct connection to the top IaaS platforms. We will get into this in a bit.

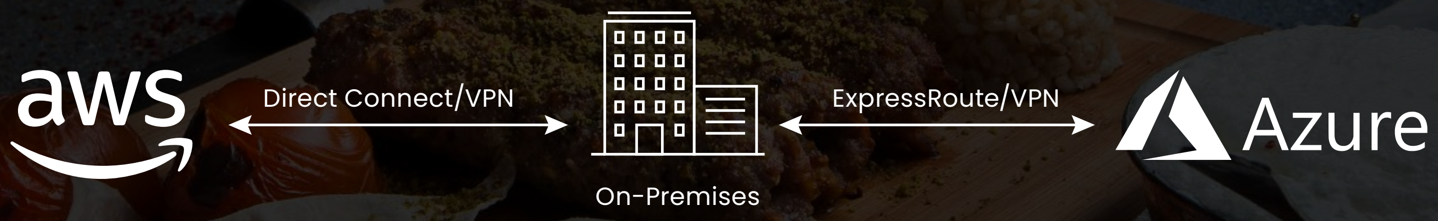

In locations where a private direct connection is not feasible, connectivity can be established through IPsec tunnels. Further, static or BGP peering is added to route traffic between the branches/DCs and the cloud provider's VPCs or equivalent.

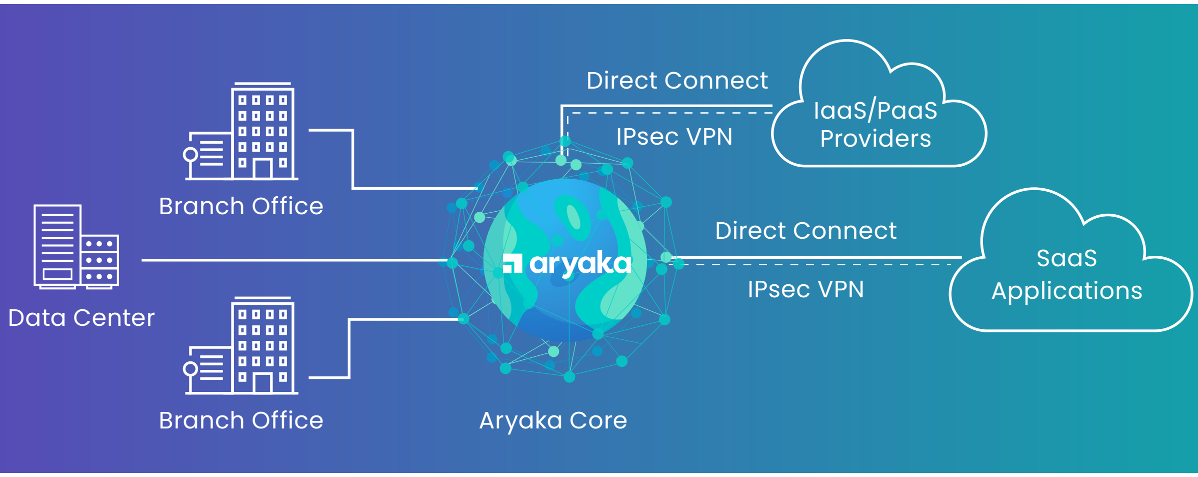

Aryaka's network connects branches using a single IPSec tunnel going from the PoP to the cloud provider's VPC. This architecture is far superior to the traditional SD-WAN architecture, where a single VPC has a number of IPSec tunnels to connect the branches. Advantages of this solution include better reliability and a smaller Round-trip time (RTT).

Note that most of the CSPs don't only leverage their own middle-mile, but also now offer it to other vendors for connectivity. Although this is a viable solution, if an enterprise or an MSP has a multi-cloud strategy, this introduces limitations. This is especially true for SASE SSE vendors from the security space that do not control their own middle-mile, partnering with a single CSP as an alternative.

Coming to the direct connect model, let's look at the different private direct connections for the top IaaS Providers. To avoid any confusion around the terminologies used by these CSPs, that may sound different but are similar, here is a list:

	aws	Azure	Google Cloud	Alibaba Cloud	ORACLE	IBM Cloud
Compute	EC2 Instance	Virtual Machine (VM)	Compute Engine VMS	ECS	Cloud Infrastructure Compute	Virtual Servers for VPC
Object Storage	S3	Blob Storage	Cloud Storage	Object Storage Service	Cloud Platform	Object Storage
Logical Data Center	VPC	VNet	VPC	VPC	VCN	VPC
Private Connectivity (L2)	Direct Connect	ExpressRoute	Interconnect	Express Connect	FastConnect	DirectLink
Gateways	TGW,VGW, DGW	VNet Gateways	Cloud Router	Virtual Border Router	Cloud Infrastructure VCN	Virtual Router Appliance

7 SaaS and UCaaS Connectivity

Aryaka uses direct Layer 2 peering and the Virtual Office (VO) methodology for faster connectivity and enhanced application performance of SaaS and UCaaS applications accessed over the internet.

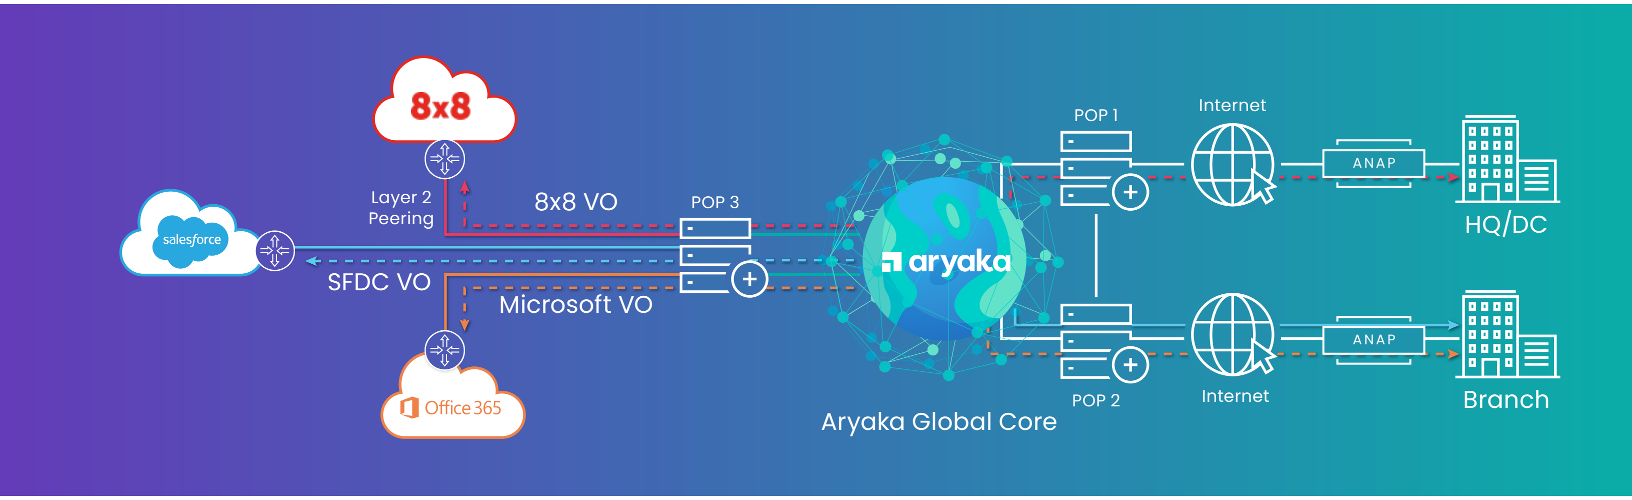

Here, Aryaka recognizes and marks traffic coming from UCaaS applications such as 8x8. The 8x8 client then leverages our QoS-driven and resilient SmartLink capabilities to connect to the closest Aryaka PoP. Traffic then traverses Aryaka's SLA-driven backbone to the nearest 8x8 peering point, bypassing the jitter, latency, and packet loss endemic to the public internet while also offering a highly available alternative to MPLS.

Note how Aryaka's optimized core provides stable core latency across both L2 and L3.

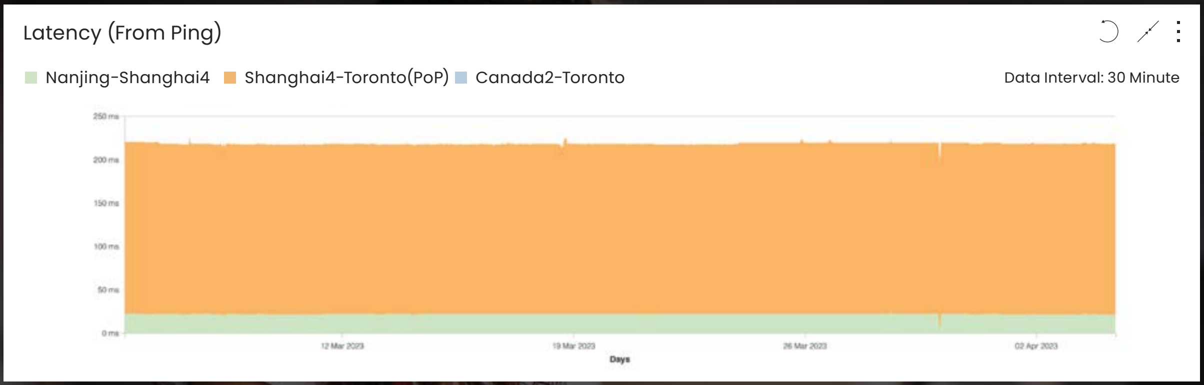

Aryaka L3 (EZ) performance from China to Canada

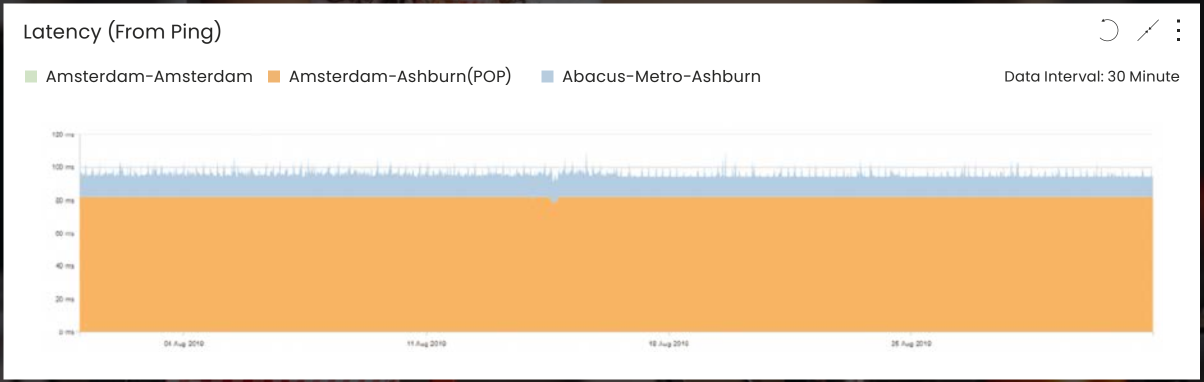

Aryaka L2 (Pro) performance from London to the US

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

As part of the Aryaka solution, VOs are established in optimal peering areas for cloud-based collaboration. This means your business-critical collaboration traffic is always taking the optimal route through Aryaka SmartConnect and SmartCloud to the best possible peering point.

So how does a virtual office work? To begin with, it connects a private branch to a SaaS or UCaaS application. 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 site 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 instance. The VO mechanism then replicates a virtually hosted public site on the PoP itself and performs SNAT (Source Network Address Translation), allowing the private network to go over the internet for the traffic going out to the Salesforce server. SNAT provides a public IP that the public Salesforce server can access, enabling it to talk to the private network where the Bangalore site sits.

For guaranteed application performance over the first and last mile, Aryaka provisions multiple features. Take path replication, for instance, that solves the packet loss issues over the internet's first-mile between the user and the Aryaka PoP. The Aryaka VO at the egress PoP applies application optimization techniques to further accelerate the applications traffic towards the application gateways.

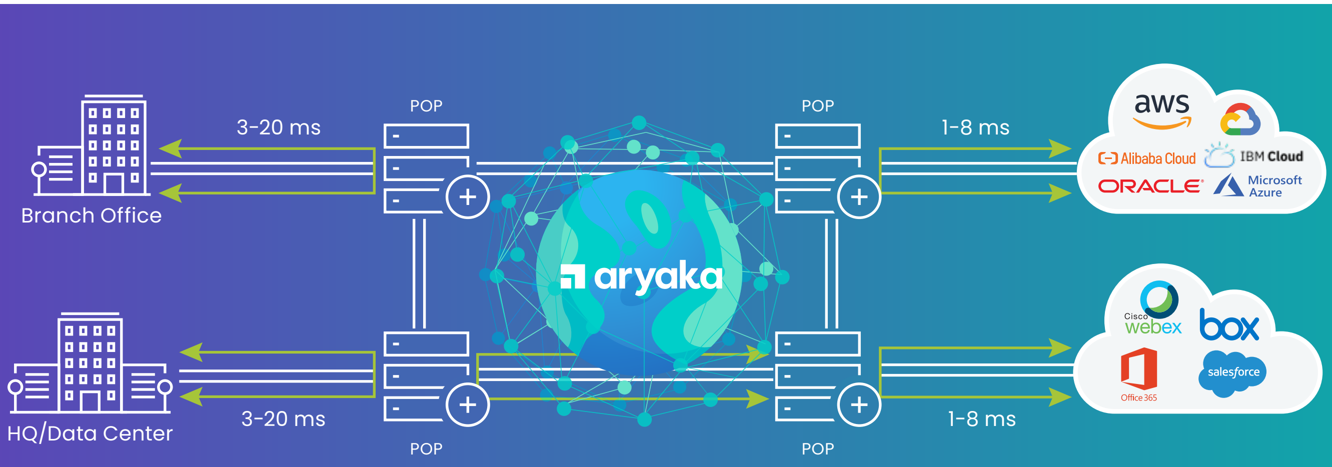

This entire process mitigates middle-mile limitations and optimizes the traffic that traverses the Aryaka core network, boosting SaaS application performance up to 20x with as low as 0% packet loss and no performance degradation.

8 Best Practices for Selecting the Right CSP

Now that you have figured out the connectivity part of the equation, how do you plan on choosing your Cloud Service Provider? When selecting the right CSP, certain aspects are simpler, while others need long and hard planning & consideration. The list of CSPs is vast and filled with choices — from AWS and Microsoft to niche players offering bespoke services. Here are the top 5 best practices that should help you work through the selection process.

8.1 Assess Their Regional Footprint

How many of the provider's regions and availability zones lie in the radius of your operational geography? The closer they are, the lesser latency and other performance degrading factors will come into play. For instance, if the user has employed more than one data center in a given geography, the chances of disruption in case of failover is almost zero.

Higher regional density translates directly into resilient delivery, better replication and redundancy, facilitating a higher disaster tolerance. The idea is to keep your computation local and engage in as little cross-regional operations as possible, to liberate your system from the possibilities of hardware and infrastructure failures.

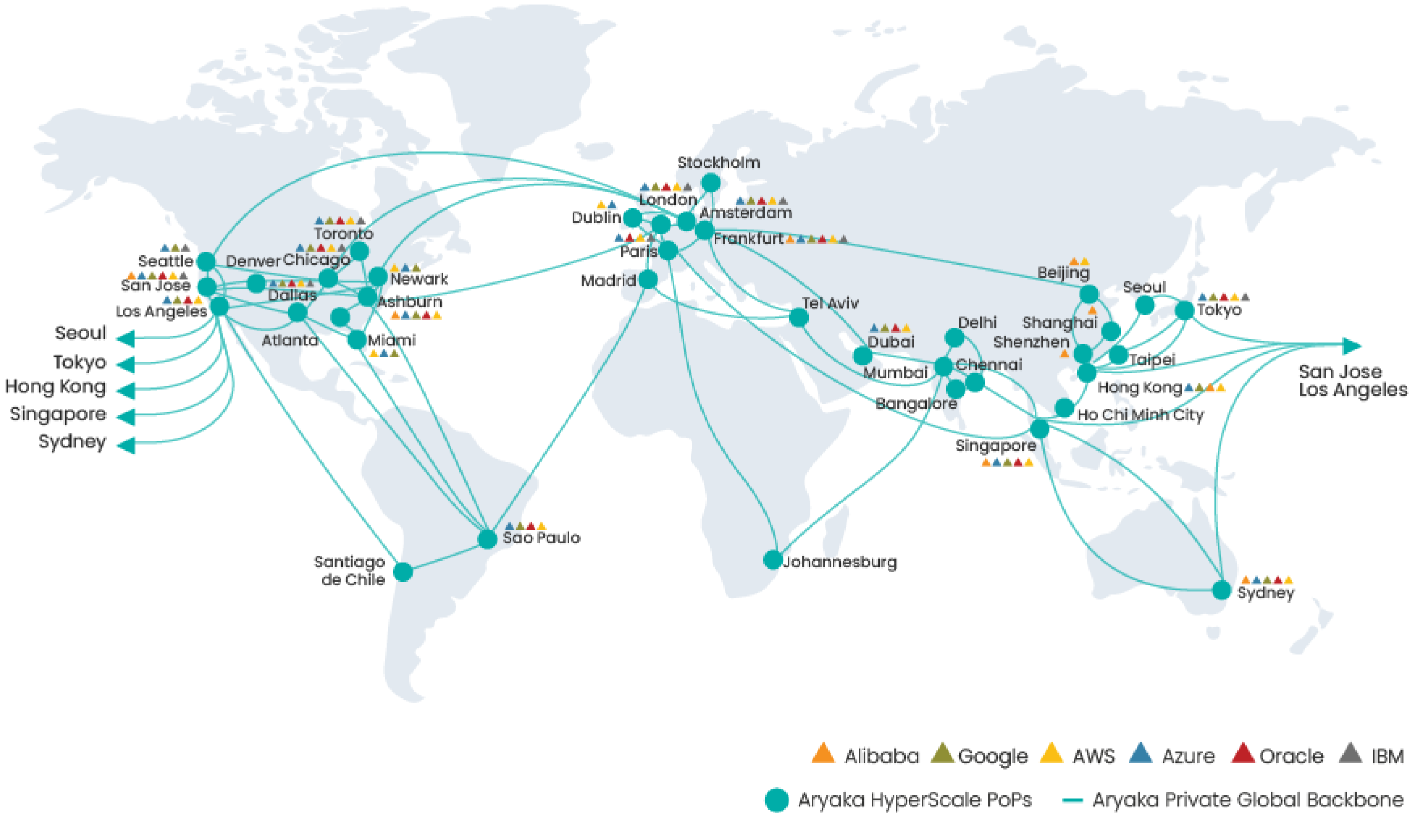

8.2 How Is the Availability In the Co-Location Facility?

Co-location can be a shared pool of benefits for both providers and consumers. It helps the providers to expand their footprints and consumers to have direct connectivity for accessing public cloud services, while housing critical workloads in their own or a colocation data center.

CSPs along with their colocation and carrier partners offer dedicated links to CSP networks. Though businesses can set-up a direct link with the cloud provider, they often prefer going via a third-party network service provider. Choosing the right provider often depends on the location of the enterprise WAN with respect to the cloud provider's data centers or zones.

When your routers are located in the same colocation facility as the CSP, you can often work directly with the CSP to set-up the direct connections between the networks.

Having a dense presence in colocation facilities enables the customer to peer through regional boundaries and instantly connect to the IaaS provider from different locations. Having the CSPs DC in the same location as the network provider can be a good bounty.

8.3 Decide Based on Your Budget and Operational Requirements

This is probably the top deciding factor when choosing a CSP. However, direct cost comparison is not as transparent as it may seem. While service providers such as Google charge separately for each service characteristic, giving users the freedom of configuration, in contrast, AWS and Microsoft lean towards pre-defined bundles.

Furthermore, different cloud providers have their forte in different domains. For example: Azure leads the way in hosting SQL servers, while AWS is the king of computing power. You need to take a call based on your predicted and actual usage pattern vis-a-vis the CSP's plans after checking what works best for your business model.

8.4 Decide Whether You Want to Use L2 Connectivity or VPN IPSec

There is a reason we dedicated an entire section to this topic and we still cannot stress enough on how important it is. Once the aggregation sites are shortlisted, they need to be connected. When your team is collaborating in that IaaS instance from different corners of the world, architectural considerations around bandwidth, throughput, traffic flow and latency etc. becomes very important.

There is a lot of private and mission-critical data flowing between continents which cannot be left to the laws of luck and probability. Thus, the network carrying such data should have the robustness of point-to-point connectivity with stable core latency.

8.5 Vendor Lock In

Almost any article that you will read on multi-cloud will have a point mentioning how distributing your workloads avoids across multiple clouds avoids vendor lock-in. However, more often than not, these CSPs run on applications that are configured to work on their and their platform only. Which means if you've built up your cloud environment using one of these applications, it will be difficult to carry them along should you decide to switch your provider. Something to be mindful of when you start looking.

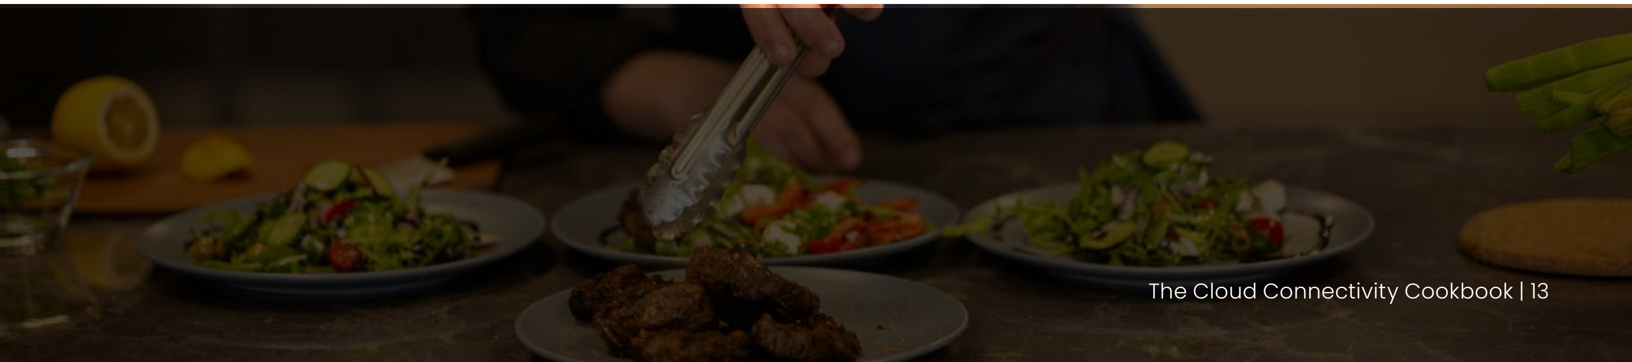

9 Conclusion - Aryaka Unified SASE

- When constructing a network architecture to support a multi-cloud rollout, there are a million moving parts that need to be taken care of.
- While most SD-WAN / SASE providers label their solution as-a-service, they're far from it. Our notion of as-a-service is where the consumer doesn't have to move a finger to do anything but consume.
- Our Unified SASE was built from the ground up on cloud-first principles, much like the cloud service providers. Our PoP and services edge architecture tied together via an automated multi-cloud services fabric is much closer in architecture to a CSP than a traditional provider and can be scaled up and down just like SaaS services.
- The PoPs connect to the service edge, and interconnect via dedicated leased fiber or Enhanced Internet links. When combined with our global orchestration, they operate as a well-integrated machine that ties together connectivity, security, optimization, and clouds with no blind spots. This is the essence of a high-performance SASE architecture with native multi-cloud connectivity.

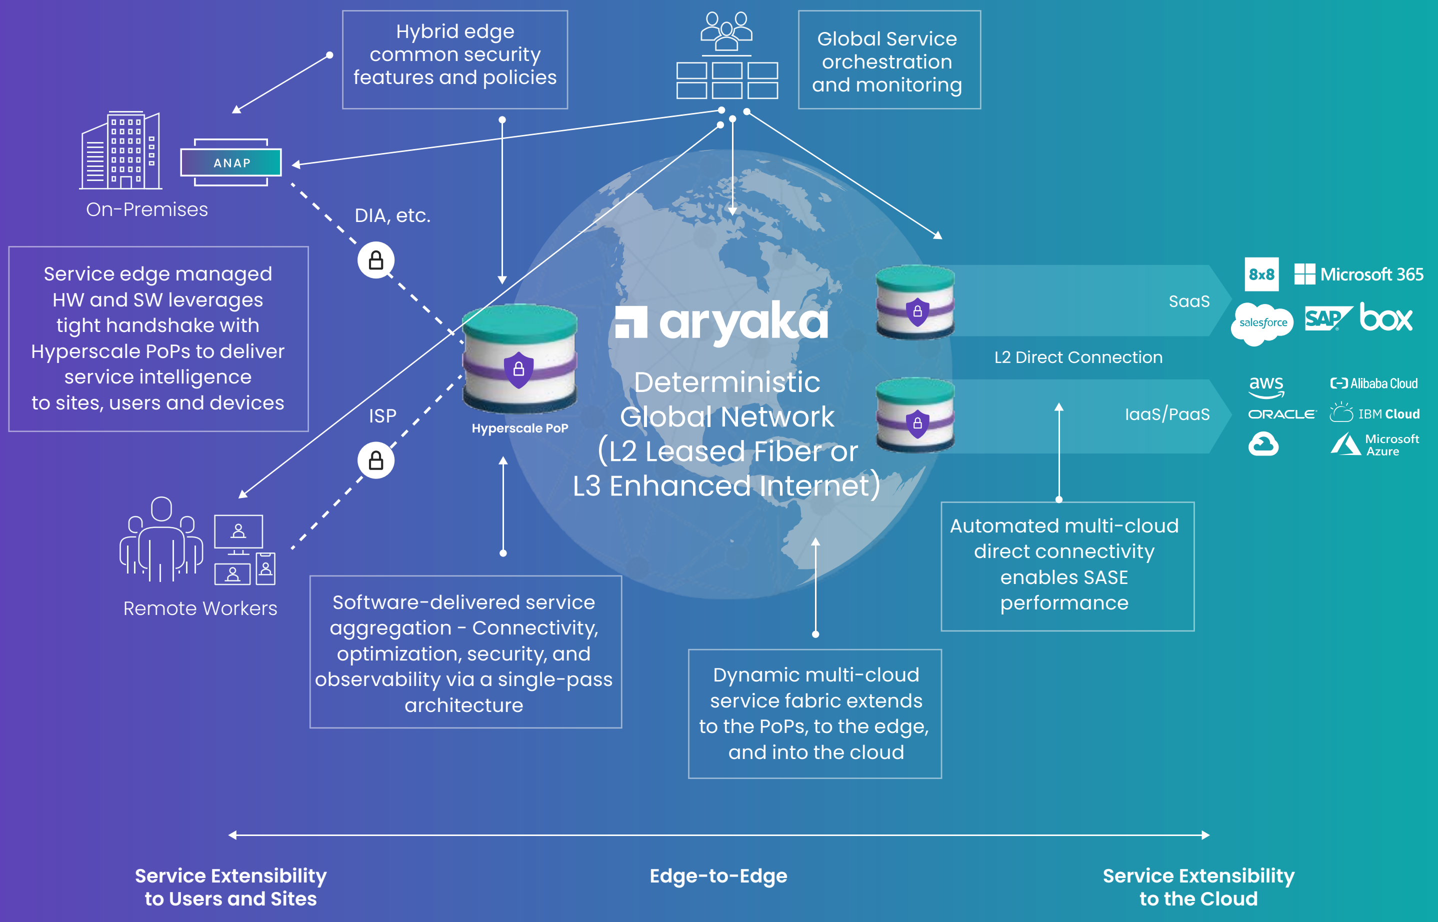

Hybrid edge common security features and policies
Global Service orchestration and monitoring
ANAP
On-Premises
DIA, etc.
Service edge managed HW and SW leverages tight handshake with Hyperscale PoPs to deliver service intelligence to sites, users and devices
ISP
Hyperscale PoP
aryaka
Deterministic Global Network (L2 Leased Fiber or L3 Enhanced Internet)
SaaS
L2 Direct Connection
IaaS/PaaS
8x8
Microsoft 365
salesforce
SAP
box
aws
Alibaba Cloud
ORACLE
IBM Cloud
Microsoft Azure
Remote Workers
Software-delivered service aggregation - Connectivity, optimization, security, and observability via a single-pass architecture
Dynamic multi-cloud service fabric extends to the PoPs, to the edge, and into the cloud
Automated multi-cloud direct connectivity enables SASE performance
Service Extensibility to Users and Sites
Edge-to-Edge
Service Extensibility to the Cloud

About Aryaka

Aryaka, the Cloud-First WAN and SASE company, and a Gartner “Voice of the Customer” leader, makes it easy for enterprises to consume network and network security solutions delivered as-a-service for a variety of modern deployments. Aryaka uniquely combines innovative SD-WAN and security technology with a global network and a managed service approach to offer the industry’s best customer and application experience. The company’s customers include hundreds of global enterprises including several in the Fortune 100.

1850 Gateway Drive, Suite 500, San Mateo, CA 94404

Follow us on :