

State of SD-WAN Connectivity

**Internet vs. Global Private Networks:
Cloud and SaaS Application Performance Comparisons**

aryaka

www.aryaka.com

Aryaka Networks, Inc. Copyright 2017
Do not reproduce without prior permission.

Contents

Executive Summary	03
Application Requirements for Today's Global Enterprises	04
Methodology	05
Sample Results	06
The Results	07
Variation of Response Time	07
Average Response Time	08
Differences between Geographies	10
What to Consider When Evaluating SD-WAN Architectures	11
A Better Way: Aryaka's Global SD-WAN	13
Core Data Analysis: Site Data Per Global Link	14



Executive Summary

Enterprises are increasingly exploring and investing in SD-WAN solutions for cloud connectivity and adding more capacity to their networks in a cost-effective manner. SD-WAN mainly relies on two types of connectivity, the public Internet and private networks.

Using the Internet as the underlying transport offers a low-cost, flexible, and rapid deployment option. While the performance over the Internet is often acceptable for users accessing Cloud/SaaS applications within the same region, the user experience deteriorates significantly with an increase in distance. Latency, packet loss and jitter are inherent to the Internet and these issues are aggravated with distance.

SD-WAN architectures built on a private network provide better security and reliable connectivity. WAN traffic traverses through a private network and is not subjected to the inconsistencies of the Internet. However, private networks such as MPLS, were not designed for cloud connectivity.

Next-generation private networks, such as Aryaka's Global SD-WAN combine the security and reliability of a private network with the flexibility, low cost, and quick deployment of the Internet to deliver a superior connectivity solution.

This report presents a quantitative analysis between SD-WAN deployments that are built on the Internet vs. a global private network such as Aryaka's Global SD-WAN. It compares and contrasts application response times between various global locations and examines key factors that impact the performance of cloud and SaaS applications. This data provides IT leaders with key insights and considerations for evaluating the best underlying transport deployment for their SD-WAN solutions to ensure enterprise-grade connectivity with reliable and consistent performance for all applications, including on-premises, SaaS, and cloud services for voice, video, and data.

4,000 Milliseconds

Application response times for one 100KB file to travel from San Jose to Shanghai over the Internet is 4,000ms. Travel times of this length lead to a loss of productivity and user abandonment of applications.

Aryaka Global SD-WAN vs. the Internet

2.5x Lower Variation in Response Time

Stable and reliable connection to data and applications improves application performance.

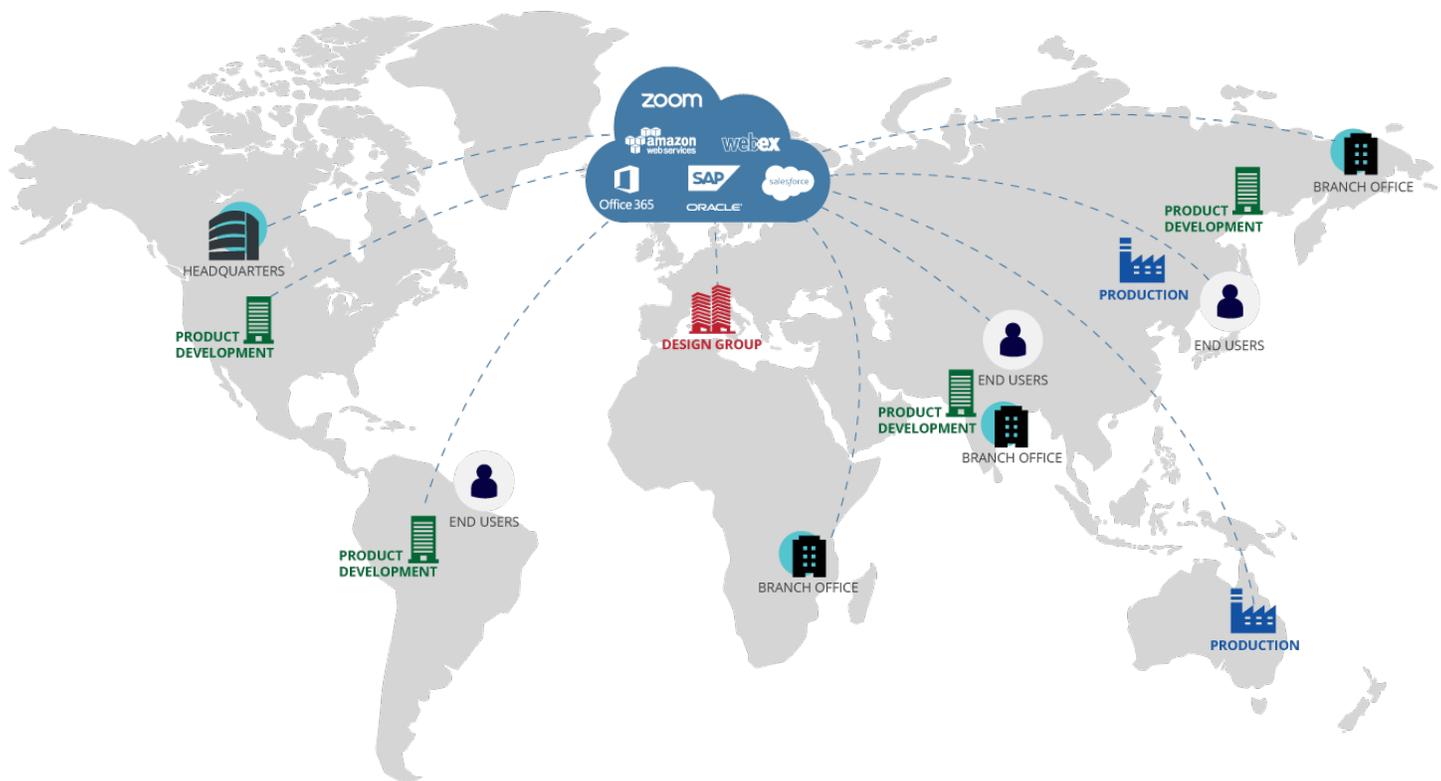
4.1x Faster Application Response Time

Improved application performance increases productivity and revenue for global enterprises.

Application Requirements for Today's Global Enterprises

Today, cloud platforms and SaaS applications are business-critical for many global enterprises. Any delay in application response time that is measured even in milliseconds can turn into hours of lost productivity per day.

While legacy WAN architectures, like MPLS and WAN Optimization, help with delivering fast performance for on-premises applications, cloud-based and SaaS applications rely on the Internet for connectivity. However, as a public medium, the Internet is prone to network congestion and packet loss. SD-WAN appliances were developed as a solution to improve these inadequacies, but only is effective when applications and end users are located within the same geographical region. When applications need to be accessed over long distances, performance is negatively impacted and cannot be improved by WAN architecture that relies on the Internet.



For many global enterprises, applications need to perform in near real time. Any latency, packet loss, or jitter over the public Internet can have a critical and negative impact on business operations and the bottom line.

Methodology

In this report, we examine the factors and parameters that impact performance for cloud and SaaS applications over long distances. To demonstrate the difference in performance that you might expect from using a private network versus the Internet, we ran a series of global application performance tests.

Results were measured and compared between the same pairs of locations, both over Aryaka's Global SD-WAN and the Internet. (The complete sample set and data is available on page 12 of the report).

The test was run as follows:

- A randomly generated 100KB file was sent from one location to another.
- The following statistics were returned:
 - HTTP result code
 - Connect time (milliseconds)
 - Transfer time (milliseconds)
- The values used for the purpose of the analysis were a summation of the transfer and connect times, when the HTTP result code was non zero.
 - Application Response Time = Connect time + Transfer Time

Two parameters were used in assessing application performance over the network

Average Response Time: The average time for an application to respond between two locations over a specified time frame.

Organizations could lose significant revenue due to just an additional second of delay in application response time. While there is no single ideal definition or industry standard for an acceptable response time, it is implicit that the lower the response time, the faster the application is responding for the user, leading to increased productivity.

Variation in Application Response Time: A measure of variation in response time relative to the average response time.

Even if the average response time is long, users can come to tolerate the wait if the response time is consistent. However, when response time varies highly, it further degrades the user experience, as users have no idea whether or not their applications will be working. In this context, maintaining a low variance in response time becomes an equally important measure of application performance.

These tests were performed over a period of a year. The data was averaged from this time period to obtain the information presented in the following sections.

Below, you'll find a graph of a sample time series of a test that was conducted between the United States and China.

Sample Results

Select Reference Site, Other Site, and Time then click 'Apply' ?

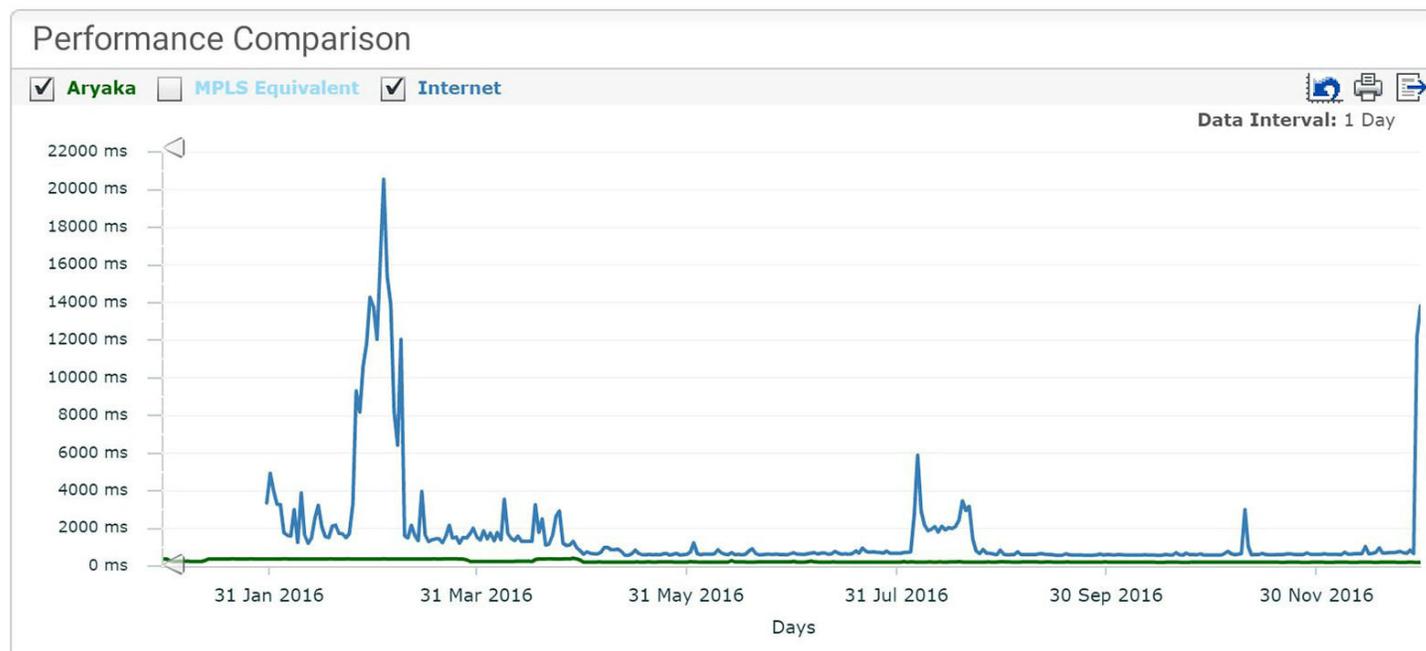
Reference Site	Other Site	Time
Bellevue, WA, USA ▼	Shanghai, China ▼	Custom ▼
		Change

BELLV San Jose Shanghai3 CHINA

ANAP ANAP

Apply

Start Time 01 Jan 2016 00:00 GMT End Time 31 Dec 2016 00:58 GMT

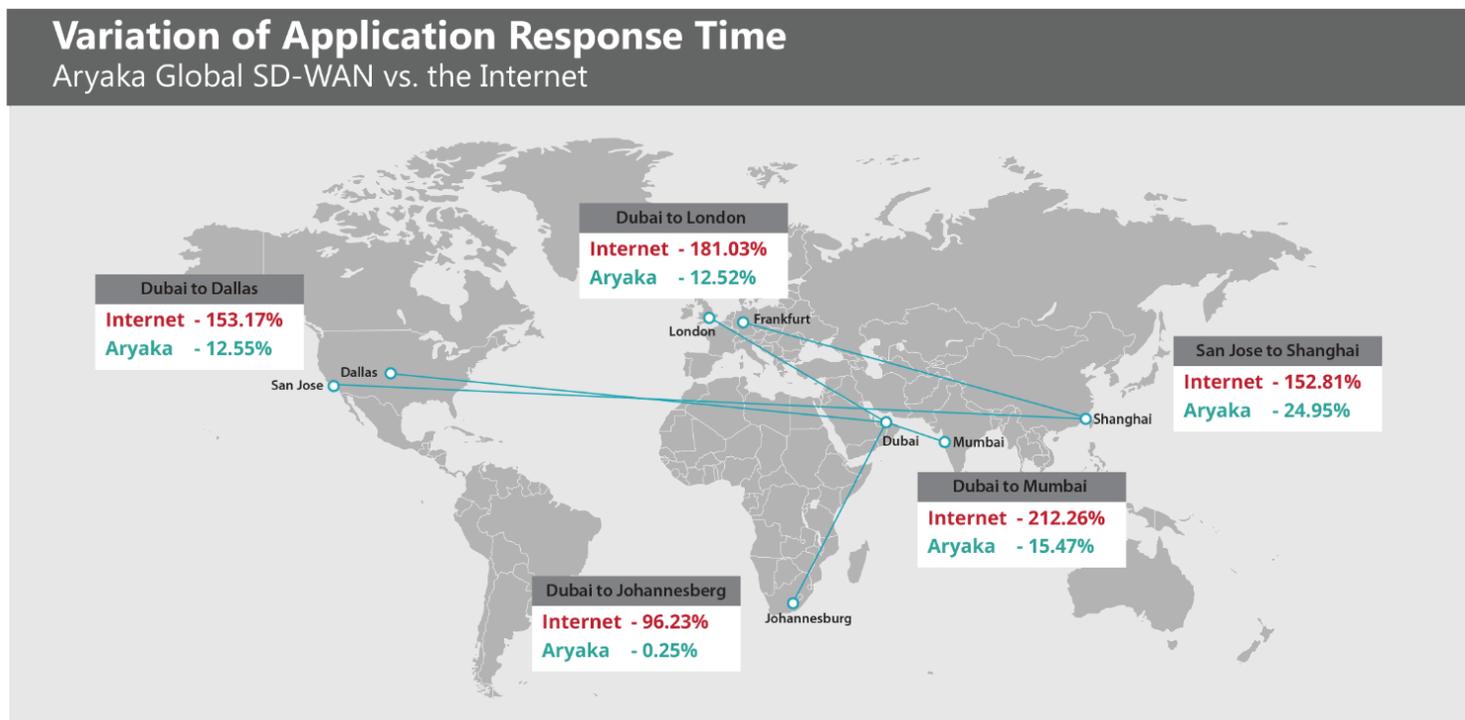


The blue line in the performance comparison graph represents the Internet, and the green line represents Aryaka.

As you can see, the response time over Aryaka's Global SD-WAN is consistent, averaging around 200ms. The response time over the Internet, however, shows significant variation, with response times often taking more than one second.

The Results

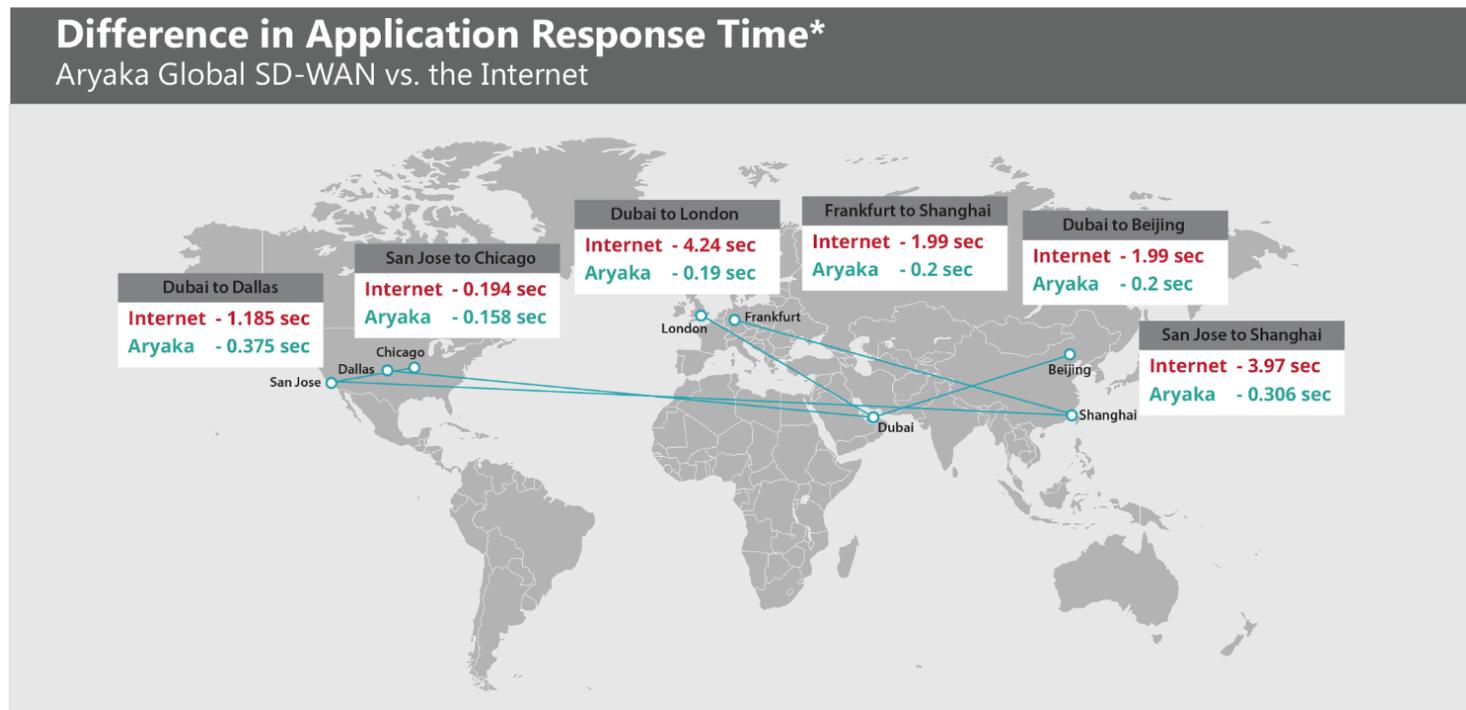
Aryaka Global SD-WAN vs. the Internet: Variation of Response Time



- **The Internet is subject to inconsistencies and high variation in application response time, which create erratic application behavior and affect the productivity of end users accessing these applications.**
 - The response time fluctuates between 750ms and 2 seconds on the link between Dubai and Dallas.
 - 750ms by itself is an unusually high application response time, but the variation further degrades the user experience since the user cannot predict what kind of performance to expect.
 - Another example of variation is the link between Dubai and Mumbai, the application response time varies between 127ms and 4 seconds. The frustration of a user who has experienced a response time of 127ms, who then must wait for 4 seconds for an application to respond, refresh or reload is compounded in comparison to a user who has simply come to expect a high response time.
- **Aryaka’s Global SD-WAN reduces the variation in application response time on an average, by a factor of 4.**
 - The values of the variation in response time for the Dubai-Dallas and Dubai-Mumbai over the Aryaka global private network is 12.5% and 15.47% respectively.

A lower variation in response time helps deliver a more consistent user experience, especially for voice and video applications.

Aryaka Global SD-WAN vs. the Internet: Average Response Time



*This graph compares the average response time across both networks between the same pairs of locations.

- The Internet cannot be your corporate WAN
 - Enterprises that utilize the Internet for their mission-critical applications will always be prone to poor application performance due to the distance data must travel across oceans and continents.
- A global private network addresses the issues associated with the Internet over long distances.
 - The data indicates it takes ~200ms for a 100KB file to travel from San Jose to Chicago over the Internet. The same file takes nearly 4 seconds or 4000ms to travel from San Jose to Shanghai.
 - When the file is transferred over Aryaka’s Global SD-WAN, the average transfer time between San Jose and Chicago is ~150ms, and the time between San Jose and Shanghai is ~300ms.

Global private networks shrink the perceived distance between locations to deliver an application performance experience that is nearly identical to those where applications and users are located in the same geographical region.

- The greater distance between locations, the greater the performance benefits of a private network.
 - Application response time is reduced by a factor of ~13 on the route between US and China (San Jose – Shanghai) and the UK-Middle east path (London-Dubai) sees an improvement by a factor of nearly 22.

Applications respond much faster for users and provide a connectivity experience similar to accessing an on-premises application when using Aryaka's Global SD-WAN.

- In the case of regional data transfer, where the distance between locations is smaller, the performance difference between the Internet and a private network is not as stark, but improvements remain visible.
 - Examples of "regional" links include the San Jose – Chicago and Frankfurt – London links.

A typical distributed enterprise as depicted in the figure on page 4 of the report can thus benefit from the enhanced user experience of the private network irrespective of whether the applications are being accessed locally or over long distances. This translates to enhanced productivity and time-to-revenue.

Differences between Geographies

- **More than just distance affects the average and variation of application response times: the maturity of infrastructure in a given geographical region plays a part.**
 - For example, data transfer in the Americas has a higher consistency than data transfer in other regions.
 - This can be attributed to the maturity of infrastructure in the Americas in comparison to other regions.
 - Yet the Americas feel the effects of latency and larger degrees of variation when connecting to end users outside of the Americas and particularly in locations that normally exhibit higher variation and response times.
- **The Asia-Pacific (APAC) links exhibit the highest unreliability. Dubai, in the EMEA region, exhibits the largest variation and highest response times.**
 - Depending on the connection point, variation for Dubai starts at a minimum of 96% (Dubai-Johannesburg) and balloons to 261% in the highest instance (Dubai-Singapore).
 - Latency for India falls into a reasonable norm, though variance is problematic.
 - Variation for China is not as extreme as for India, though response times are still problematic, ranging from 12 and 23 seconds.
 - When connecting from the Americas or Europe to countries in the APAC region over the Internet, expect slower overall performance.

Difference in Application Response Time

Aryaka Global SD-WAN vs. the Internet

Region	Aryaka	Internet	Ratio
Americas	0.27 ms	1.04 ms	3.80
EMEA	0.23 ms	1.31 ms	5.69
APAC	0.25 ms	1.44 ms	5.69

Table 1

Table 1 indicates the difference in average response time between the Internet and Aryaka's Global SD-WAN in each geography. The ratio represents the average response time over the Internet divided by the average response time over Aryaka. For example, in the APAC region, the Aryaka response times are lower by a factor of 5.7.

Variation of Application Response Time

Aryaka Global SD-WAN vs. the Internet

Region	Aryaka	Internet	Ratio
Americas	36.05%	105.45%	2.93
EMEA	32.33%	67.32%	2.08
APAC	16.85%	53.72%	3.19

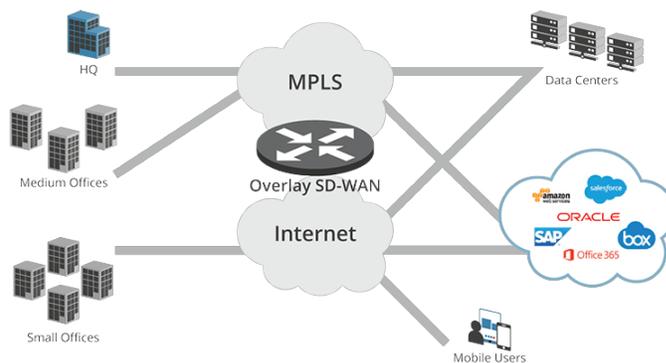
Table 2

Table 2 represents the variation in response time between the Internet and Aryaka's Global SD-WAN in each geography. The ratio represents the variation in response time over the Internet divided by the variation in response time over Aryaka's Global SD-WAN. The data indicates that this variation is almost 3.5 times higher in APAC when applications are accessed over the Internet.

What to Consider When Evaluating SD-WAN Architectures

As new solutions come onto the market purporting to offer ways for IT teams to cope with the shortcomings of legacy solutions like MPLS, IT leaders must remain critical of solutions that simply layer new problems on top of the old ones.

Many network planners try to design a hybrid WAN architecture that incorporates both the Internet (via SD-WAN) and MPLS. However, this leaves them with two different transport networks, both riddled with inherent challenges. Enterprises are still left with the need to establish connectivity between these networks in order to provide seamless application delivery to end users.

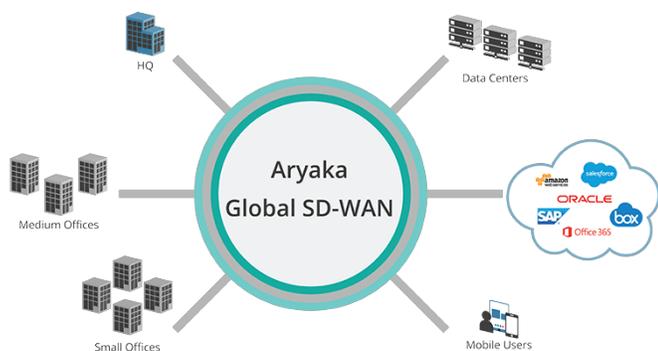


Typical SD-WAN Architecture

One common approach is to connect MPLS and Internet at the enterprise data centers. This method is suitable only for enterprises with a limited geographic distribution, as it requires a backhauling of traffic, which causes poor performance, especially for distributed cloud services.

Another approach is to purchase direct cloud connectivity from individual cloud providers. However, carriers currently only offer direct connectivity to leading cloud services in select locations, and connectivity is not likely to be scalable.

In addition, when multiple transport services are in the picture, enterprises then need to overlay SD-WAN services that will enable controlling, defining, and configuring end-to-end application flows across all of the elements of the hybrid WAN.



Aryaka's Global SD-WAN

Even with the seeming control granted in the process of designing an elaborate WAN, network administrators still do not have complete visibility into their traffic flows, which could lead to under-utilization of the network and thereby poor application performance. Rolling out new applications becomes a complicated process, while ensuring security could become an issue too. In addition, costs can inflate wildly, as more and more hardware, software, and direct connections need to be purchased and maintained to keep the network running.

Based on our data and analysis, we recommend that enterprises consider these key points when evaluating SD-WAN architectures:

1. **Application Performance.** Although it depends on the application and each individual user, delays in response time can be detrimental—from lost productivity to lost conversions.* From a regional perspective, our data indicates that the Internet experiences the biggest application performance problems in EMEA and APAC. In some cases, these problems are caused by the latency introduced by longer distances, and in others, the underlying network infrastructure is a major contributing factor.
2. **Variability.** Variation in network performance and low reliability is a significant issue with the Internet. The inability to predict application performance and reliability for each geographical location is problematic for businesses, as it affects end user productivity. The Internet has the strongest variability challenges when it comes to the APAC region, where response time variance shoots above 100% (as compared to 67% and 53% for EMEA and the Americas respectively).
3. **Complexity.** Building a complex WAN infrastructure by combining a private WAN or MPLS lines with lower cost Internet connectivity and overlay SD-WAN is not the most cost-effective approach an enterprise can take. In addition to the cost, such a solution requires both IT and network expertise, as well as a significant outlay of time for maintenance and support.
4. **Flexibility.** Enterprises must consider how agile and flexible their WAN needs to be to support distributed cloud deployments. Network resources must be allocated dynamically, quickly, and efficiently across multiple locations and different types of cloud services and providers. Depending on the application requirements, network administrators must be able to flexibly increase and decrease bandwidth as needed, without having to wait months or days.
5. **Security.** Another key consideration when building a WAN for the cloud is security. Sensitive and business-critical data flowing across the network and between different cloud types must remain protected and secure. The purpose of using legacy technologies like MPLS in the first place was to ensure not just the reliability of WAN traffic, but also its security. In moving to the cloud, enterprises shouldn't have to compromise on those expectations. Therefore, enterprises should adopt a network solution that ensures data confidentiality, integrity, and availability. In many instances, this requires network-based encryption and, in some cases, encryption at multiple layers.

Evaluation Suggestions:

1. Global connectivity. Look for an SD-WAN solution that delivers global connectivity and not just regional connectivity.

2. Dependency on the Internet. Knowing that the Internet is not your corporate WAN, find an SD-WAN solution that removes Internet dependency, not one that simply finds an alternate path.

3. Application performance requirements. Know your users, applications, and associated performance requirements. An ideal SD-WAN solution must map to these needs.

4. Low variability. Understand differences in regional performance issues and variability. Look for solutions that flatten variability across each of them.

5. Cost-effectiveness and agility. Look for next-generation SD-WAN solutions that are cost-effective and provide agility (the ability to extend connectivity to new offices within hours vs. weeks or months) as compared to legacy MPLS and WAN Optimization solutions.

Based on these recommendations, this report clearly indicates that a global enterprise that has users spread around the world and accesses applications in the cloud simply cannot rely on the Internet or Internet-dependent SD-WAN solutions to deliver a reliable consistent application delivery experience.

* Scott Barber, "Acceptable Application Response Times vs. Industry Standard," Search Software Quality, March 2007.

A Better Way: Aryaka's Global SD-WAN

Aryaka takes a unique approach to solving global application performance. Aryaka's Global SD-WAN reduces the perceived latency between locations by providing an optimized network. Aryaka has 28 globally distributed POPs, connected by private leased lines that are located within 30 milliseconds of 95% of the world's business users. This removes barriers for global data transfer and lowers latency within the middle mile.

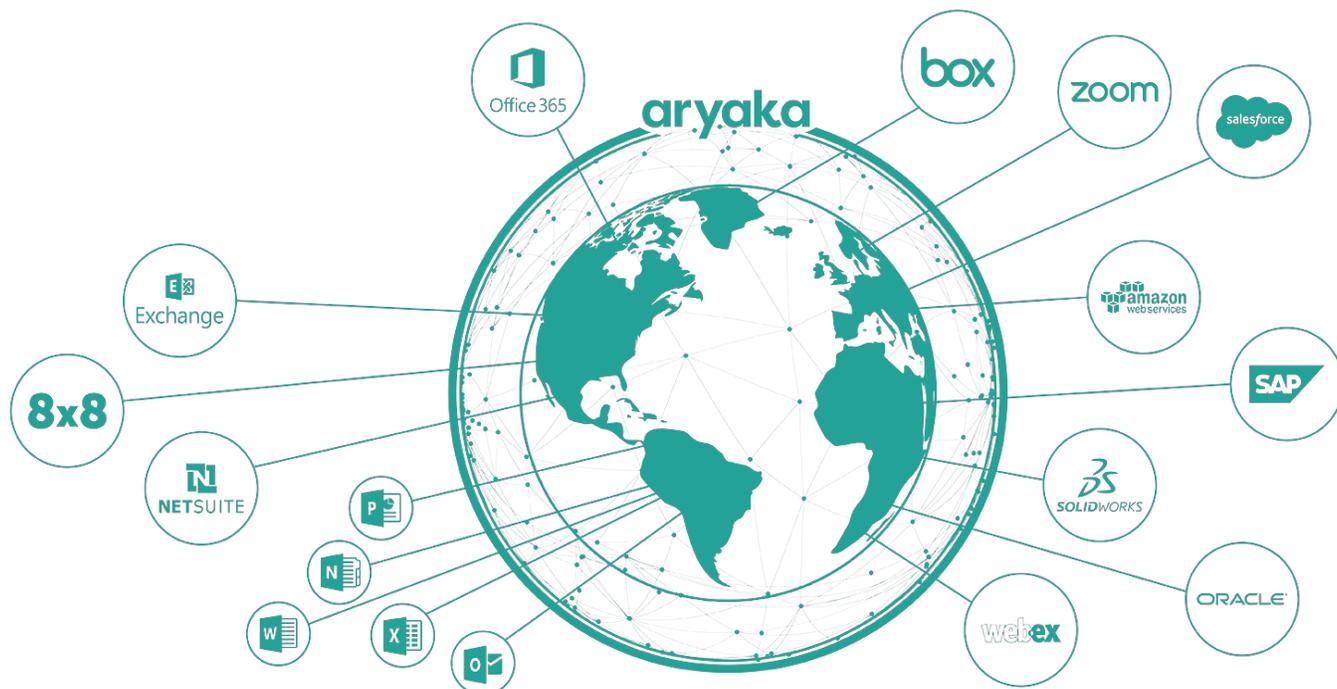
Aryaka's Global SD-WAN includes built-in WAN Optimization and acceleration, reducing packet loss down to 0–2%, while accelerating application performance up to 40 times. This enables IT organizations to avoid investing in additional hardware to optimize throughput, as well as the resources to manage those assets. The latter is particularly important to IT departments under increasing budgetary constraints.

With Aryaka, businesses have the flexibility to bring new branch offices and locations online in days (between eight and 48 hours), compared to the months required to deploy MPLS solutions. In a digital age where time means everything, this agility can translate into substantial benefits.

Benefits of Aryaka's Global SD-WAN

- Up to 40x better application performance over the internet
- Application response times 4.1x faster
- Up to 56% cost savings over MPLS
- Deployment in days vs. weeks or months with MPLS

Aryaka's Global SD-WAN is the only solution that works for any application, anywhere in the world: on-premises, cloud, and SaaS.



Core Data Analysis: Site Data Per Global Link

All data in this report is based on tests between the following pairs of locations.

Link	Application Response Time (in sec)	
	Internet	Aryaka
Dubai - Dallas	1.185	0.373
Mumbai - Dallas	0.791	0.355
San Jose - Singapore	0.661	0.285
San Jose - Shanghai	3.97	0.306
San Jose - Chicago	0.194	0.158
Frankfurt - Los Angeles	0.486	0.219
Amsterdam - Chicago	0.629	0.25
Frankfurt - Chennai	1.098	0.363
Dubai - Singapore	0.708	0.162
Dubai - Mumbai	0.4616	0.09
Dubai - Beijing	1.99	0.2
Frankfurt - Shanghai	1.99	0.2
Dubai - London	4.24	0.19
Dubai - Johannesburg	1.13	0.282
Delhi - Chicago	0.932	0.338
Frankfurt - London	0.258	0.18

Link	Variation Response Time (in %)	
	Internet	Aryaka
Dubai - Dallas	153.17	12.55
Mumbai - Dallas	15.67	10.02
San Jose - Singapore	66.39	22.4
San Jose - Shanghai	152.81	24.95
San Jose - Chicago	8.06	4.15
Frankfurt - Los Angeles	8.25	5.05
Amsterdam - Chicago	29.28	25.3
Frankfurt - Chennai	32.11	9.81
Dubai - Singapore	261.79	102.33
Dubai - Mumbai	212.26	15.47
Dubai - London	181.03	12.52
Dubai - Johannesburg	96.23	0.25
Delhi - Chicago	11.80	9.25
Frankfurt - London	70.71	31.57

Next Steps

Find out how your enterprise can benefit from using
Aryaka Global SD-WAN compared to the Internet.

Call us at 1-877-727-9252 or email us at sales@aryaka.com to learn more.

For a more in-depth analysis of WAN solutions today, download a copy of Aryaka's "2017 State of the WAN Report"

About Aryaka Networks

Aryaka Global SD-WAN is the #1 MPLS replacement for global enterprises. Aryaka delivers a global enterprise connectivity solution for data centers, branch offices and remote/mobile employees with significantly better performance for on-premises and cloud applications – voice, video and data. It is delivered as a service, eliminating the need for complicated, costly, and time-consuming WAN construction and maintenance. Unlike MPLS, Aryaka Global SD-WAN can be deployed anywhere in the world within days and without requiring additional hardware or resources.