

LEARNING MADE EASY

Aryaka Special Edition

# The Cloud-First WAN

for  
**dummies**<sup>®</sup>  
A Wiley Brand



Consume your  
network as a service

—  
Achieve fully  
integrated security

—  
Make your network as  
easy as your cloud

Compliments  
of  
**aryaka**

Shashi Kiran, Aryaka  
David Ginsburg, Aryaka  
Lawrence C. Miller, CISSP

# About Aryaka

Aryaka, the Cloud-First WAN company, brings agility, simplicity and flexibility to consuming the WAN-as-a-service by leveraging the cloud consumption model. Aryaka's managed SD-WAN customers include hundreds of global enterprises across more than 60 countries, with use cases spanning digital and WAN transformation, application performance optimization, global UCaaS, multi-cloud connectivity, MPLS migration, and connectivity to China. Aryaka's service is based on the company's SmartServices platform. Within the SmartServices umbrella, SmartConnect delivers regional and global connectivity via the Aryaka private core, broadband Internet, MPLS, and cellular, as well as managed Last Mile Services. SmartSecure offers multi-vendor edge and cloud security with VNF support. SmartOptimize leverages patented multi-segment network and application optimization. SmartCloud delivers pre-wired regional connectivity to IaaS, PaaS, and SaaS. SmartInsights and the MyAryaka cloud portal offer configuration and monitoring. Finally, SmartManage is the services foundation consisting of Aryaka's global service PoPs, the Aryaka Network Access Point edge appliance, dedicated L2 global connectivity, as well as orchestration and monitoring for quick service turn-up.



# The Cloud-First WAN

Aryaka Special Edition

**by Lawrence C. Miller**

**Shashi Kiran, Aryaka**  
**David Ginsburg, Aryaka**

**for**  
**dummies**<sup>®</sup>  
A Wiley Brand

# The Cloud-First WAN For Dummies®, Aryaka Special Edition

Published by

**John Wiley & Sons, Inc.**

111 River St.

Hoboken, NJ 07030-5774

[www.wiley.com](http://www.wiley.com)

Copyright © 2020 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Aryaka and the Aryaka logo are trademarks or registered trademarks of Aryaka Networks, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact [info@dummies.biz](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [Branded Rights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN 978-1-119-69972-9 (pbk); ISBN 978-1-119-69966-8 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

## Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. Some of the people who helped bring this book to market include the following:

**Project Editor:** Martin V. Minner

**Editorial Manager:** Rev Mengle

**Executive Editor:** Steve Hayes

**Business Development**

**Representative:** Karen Hattan

**Production Editor:**

Mohammed Zafar Ali

# Table of Contents

INTRODUCTION .....	1
How the World Changed in 2020 .....	2
About This Book .....	2
Foolish Assumptions.....	3
Icons Used in This Book.....	3
Beyond the Book.....	4
<b>CHAPTER 1: Digital Transformation and the Wide-Area Network.....</b>	<b>5</b>
Exploring Digital Transformation and Other Trends .....	5
Understanding the Impact on WAN Planning and Design .....	8
Recognizing Challenges in Traditional WAN Approaches.....	10
Carrier managed SD-WAN: Problems galore.....	10
SD-WAN edge overlay: A piecemeal approach.....	11
<b>CHAPTER 2: Introducing a Better Approach: The Cloud-First WAN .....</b>	<b>13</b>
Defining a Cloud-First WAN Experience-as-a-Service .....	13
Business agility .....	15
Operational simplicity .....	15
Multi-cloud readiness.....	16
Trials and tribulations with UCaaS.....	16
Seeing the Advantages over DIY SD-WAN or a Telco-Managed WAN.....	17
Adopting a Platform Approach.....	21
<b>CHAPTER 3: Discovering the Key Elements of a Cloud-First WAN .....</b>	<b>23</b>
Connectivity.....	23
Security .....	25
Multi-Cloud.....	27
Application Acceleration and Optimization.....	28
Visibility.....	31
Management and Orchestration.....	32

<b>CHAPTER 4:</b>	<b>Exploring Use Cases for the Modern Enterprise</b> .....	35
	Accelerating Application Performance .....	35
	Finding a Flexible, Simpler Alternative to MPLS .....	38
	Connecting to Any Cloud Anywhere with Ease .....	41
	Ensuring a Successful Digital Transformation .....	43
	Optimizing Unified Communications-as-a-Service.....	46
	Finding Faster Connectivity to China and Beyond .....	48
<b>CHAPTER 5:</b>	<b>Deploying a Cloud-First WAN — Ten Capabilities and Benefits</b> .....	51
	Delivering WAN-as-a-Service.....	51
	Predictable Connectivity Anywhere .....	52
	Built-In WAN Optimization .....	53
	Multi-Cloud Networking .....	53
	Security and SASE.....	54
	Automation and Orchestration .....	54
	Predictive Analytics .....	55
	Management, Visibility, and Troubleshooting .....	55
	Last-Mile Management and Monitoring .....	56
	Global PoP Architecture for Service Delivery.....	56
<b>GLOSSARY</b> .....		57

# Introduction

**D**igital transformation is perhaps the most widely used buzzword in today's business vocabulary, and with good reason. It has the power to make or break organizations. In recent years, organizations of all sizes and in every sector have embarked upon some form of digital transformation initiative.

IT leaders understand the role that technologies like cloud, big data, social, and mobile play in digital transformation. These higher-level technology stacks get most of the attention while foundational areas like enterprise wide-area network (WAN) connectivity are often ignored. It only makes sense: When people plan their vacations, they want to hear about the destination, not the plane or car that gets them there!

Digital transformation is now widely understood from both the perspective of business outcomes and high-level technology. But the importance of the underlying connectivity layer, the WAN, is often overlooked. A well-designed enterprise WAN can support successful digital transformation initiatives, whereas a poor WAN architecture for cloud, big data, and mobile has the potential to derail these initiatives.

But what should this new WAN architecture look like? It must follow cloud-first principles, echoing the flexibility, velocity, and simplicity that enterprises have embraced over the past decade with their public cloud deployments. This cloud consumption model — OpEx instead of CapEx — now has an analogy in the networking space.

This cloud-first approach is critical in supporting application modernization initiatives, and ultimately the success of digital transformation. The reverse — legacy models tied to outdated technology that results in complexity and a lack of flexibility — dooms transformation efforts to failure.

# How the World Changed in 2020

The coronavirus (COVID-19) and its impact on society have been top-of-mind for most of 2020, and millions of people have had firsthand experience with remote working as mandated by social distancing edicts in much of the world.

In many cases, corporate WANs, especially with regard to remote access and the ability to weather a lights-out operation, were unprepared. The industry learned that the WAN needed to be more flexible, scalable, simple to operate, and adaptable to change. As Heraclitus, the Greek philosopher, articulated more than 2,500 years ago, “Change is the only constant in life.”

COVID-19 taught some important lessons about how to harden business continuity practices, and at the back end of the pandemic, the industry will be stronger for the experience. SD-WAN, especially in the context of a managed service, and in an expected era of lower CapEx versus OpEx, will play a major role in this new world.

## About This Book

*The Cloud-First WAN For Dummies, Aryaka Special Edition*, consists of five chapters that explore

- »» The impact of the WAN on digital transformation and other modern trends (Chapter 1)
- »» Addressing today’s digital transformation challenges (Chapter 2)
- »» The critical elements of a cloud-first WAN (Chapter 3)
- »» Use cases for the modern enterprise (Chapter 4)
- »» Key capabilities and benefits of a cloud-first WAN (Chapter 5)

In addition, a glossary at the end of the book defines the terms you’ll encounter in your journey to the cloud-first WAN.

Each chapter is written to stand on its own, so if you see a topic that piques your interest feel free to jump ahead to that chapter. You can read this book in any order that suits you (though I don’t recommend upside down or backward).



# Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless!

Mainly, I assume that you work in an organization that is looking for a better way to design and manage your enterprise WAN in the era of the cloud. Perhaps you're an IT executive or manager such as a chief information officer (CIO) or chief technology officer (CTO) knee-deep in digital transformation. Or perhaps you're a technology decision maker responsible for infrastructure, cloud, security, or application performance.

As such, this book is written for technical readers with a general understanding of cloud and networking concepts and technologies.

If any of these assumptions describe you, then this is the book for you. If none of these assumptions describe you, keep reading anyway. It's a great book and you'll learn quite a bit about the cloud-first WAN.

## Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



REMEMBER

This icon points out important information you should commit to your nonvolatile memory, your gray matter, or your noggin — along with anniversaries and birthdays.



TECHNICAL  
STUFF

If you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon and is the stuff nerds are made of.



TIP

Tips are appreciated, never expected — and I sure hope you'll appreciate these useful nuggets of information.

# Beyond the Book

---

There's only so much I can cover in a short book, so if you find yourself at the end thinking, "Gosh, this was an amazing book, where can I learn more?" check out [www.aryaka.com](http://www.aryaka.com).

## IN THIS CHAPTER

- » Looking at digital transformation and the dawn of the cloud era
- » Rethinking the traditional enterprise WAN architecture
- » Understanding the limitations of traditional SD-WAN approaches

# Chapter 1

# Digital Transformation and the Wide-Area Network

This chapter shows you how digital transformation and the cloud have created the need for a more robust, cloud-first wide-area network (WAN).

## Exploring Digital Transformation and Other Trends

The cloud (public, private, and hybrid) is a key enabler of digital transformation in modern enterprises. Businesses now expect a cloud-like experience — better performance, more agility, operational simplicity, and greater responsiveness — in practically all facets of their operations.

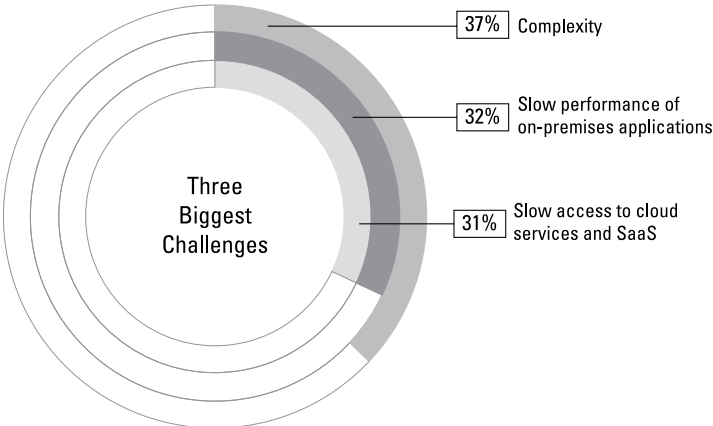


REMEMBER

The cloud isn't a destination; it's a journey and an experience.

Widespread adoption of Software-as-a-Service (SaaS) applications such as Office 365, Salesforce, 8x8, and WebEx is

putting ever-growing performance demands on legacy wide-area network (WAN) architectures. At the same time, enterprises are migrating their on-premises applications and workloads to public cloud Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) offerings and placing still greater demands on their WAN architectures. As organizations continue to migrate to the cloud, the biggest challenges with their WANs are cost, complexity, and performance (see Figure 1-1). One observation is that cost is no longer a top-three concern, speaking to the fact that the return on investment (ROI) of WAN transformation efforts is now better understood.



**FIGURE 1-1:** Enterprise WAN challenges in the cloud. (Source: Aryaka fourth annual “State of the WAN” report)

This migration has also resulted in the strong adoption of what many call the “cloud consumption model,” an OpEx versus CapEx-driven paradigm, one of *consume* versus *construct*. This model delivers the flexibility, velocity, and performance required by the modern (that is, digitally transformed) enterprise. But in following this path, they can’t leave the WAN behind.

Thus, the application of the cloud consumption model to the WAN, or what’s called the *network consumption model*. Just as the public cloud providers take responsibility for the physical infrastructure, creating a point-and-click environment for a host of services, the network consumption model does the same for the WAN. This is the essence of the cloud-first WAN.

## SD-WAN VERSUS WAN

Most software-defined WAN (SD-WAN) deployments today reuse the existing multiprotocol label switching networks and broadband Internet connections at enterprise locations (such as branch or remote offices). Thus, SD-WAN is a subset of a more complete WAN architecture. The WAN goes well beyond SD-WAN to include integrating security, application and network optimization, multi-cloud connectivity, and other technologies that add to complexity in a do-it-yourself (DIY) approach and make a managed, cloud-first WAN the ideal solution for enterprise WAN deployments. And, given the pace of change, will anyone still be using the term “SD-WAN” two to three years from now?

Aryaka’s fourth annual “State of the WAN” report, with more than 1,000 responses from North America, Europe, and Asia, found that the majority of surveyed enterprises operate in highly distributed and complex environments:

- » More than 50 percent have 20 or more branches around the globe.
- » More than half have 100 or more global branches.
- » Almost a third have more than 500 applications deployed.

The report is available at <https://www.aryaka.com/state-of-wan-2020/>.

For IT organizations that manage traditional WANs, solving slow application performance issues, including those in the cloud, while managing multiple network service providers and maintaining their security posture, have become the top, most time-consuming challenges:

- » Slow application performance leads to poor user experience for remote and mobile users (46 percent) and for employees in branch offices (42 percent).
- » Helping employees access and integrate cloud and SaaS application origins is time-consuming (39 percent).
- » Security breaches are still a major concern (38 percent).
- » Managing telcos or service providers is a nightmare (31 percent).

In the cloud-first era, in which critical applications and workloads are increasingly being delivered as cloud-native solutions, the traditional enterprise WAN is quickly becoming a bottleneck and a barrier to successful digital transformation.

## Understanding the Impact on WAN Planning and Design

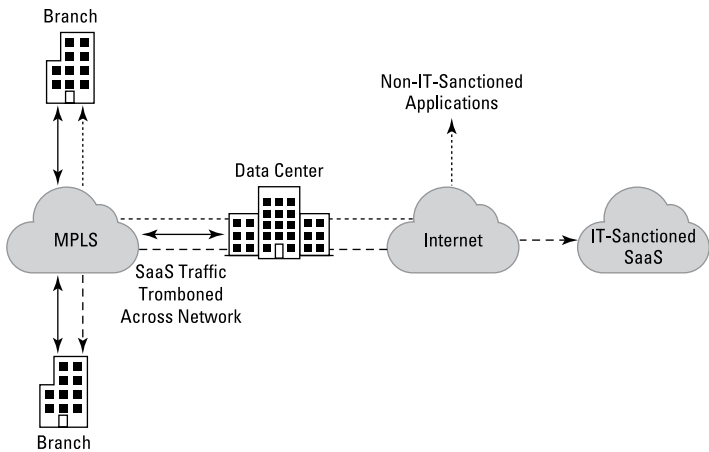
In the not too distant past, enterprise WAN traffic primarily flowed from client workstations located in headquarters and branch/remote locations to servers and applications located in an on-premises data center. Multiprotocol Label Switching (MPLS) networks — connecting numerous enterprise locations to a corporate headend or data center over high-performance, low latency private network links — were well-suited for client-server applications and became the prevalent enterprise WAN architecture. The need for Internet access was relatively limited and primarily consisted of external email and web browsing. As a result, branch and remote office Internet traffic could easily be backhauled across the enterprise WAN to the headend. This design enabled centralized management of network and security policies to ensure the WAN met the performance and security requirements of the business.

With the arrival of the cloud era, particularly the rapid adoption of SaaS applications for core business functions, the traditional enterprise MPLS WAN quickly became a bottleneck as network traffic was increasingly Internet-bound, rather than destined for the corporate data center. As the volume of network traffic bound for the Internet increased exponentially, so too did network congestion, latency, and delay on expensive MPLS links that were being used to backhaul (or “trombone”) all this traffic to the corporate headend (see Figure 1-2), ultimately causing a poor application and user experience.



REMEMBER

MPLS networks are not designed to be responsive to the dynamic needs of modern business. For an organization undergoing digital transformation, MPLS can be slow and complicated. Simply put, MPLS cannot deliver rapidly and with the sense of urgency necessary for successful digital transformation.



**FIGURE 1-2:** Traditional MPLS networks are inefficient in the modern cloud era.

As application performance issues began to degrade the user experience and reduce overall productivity, IT organizations responded by provisioning direct Internet access (DIA) links, such as broadband and/or Long-Term Evolution (LTE) wireless connections, from local Internet service providers (ISPs). Although this solution addressed some of the inefficiencies and performance issues associated with backhauling Internet traffic across the enterprise WAN, it introduced new challenges including:

- » Inconsistent application performance associated with best-effort broadband Internet access, asynchronous bandwidth (different upload and download speeds), and weak or non-existent service-level agreements (SLAs)
- » Lack of visibility and control of a majority of network traffic flowing directly from the branch or remote locations to the Internet rather than through a central headend
- » Network complexity associated with configuring route selection across multiple links, route optimization, and load balancing
- » Higher costs associated with provisioning of Internet access, procurement of networking and security equipment, and administration and maintenance

# Recognizing Challenges in Traditional WAN Approaches

To address the challenges of traditional MPLS network designs and branch/remote DIA links, software-defined wide-area networking (SD-WAN) has emerged as a possible solution. However, SD-WAN introduces new challenges as well. Two common approaches to SD-WAN today are carrier managed SD-WAN and SD-WAN edge overlay.

## Carrier managed SD-WAN: Problems galore

In the pre-cloud era, carriers were the chief providers of WAN connectivity services to enterprises. With SD-WAN establishing itself as a technology of choice for cloud connectivity, many of the same carriers are jumping onto the SD-WAN bandwagon.

Carrier managed SD-WAN networks are built using equipment sourced from multiple vendors, with each vendor providing a proprietary configuration and monitoring solution, hence making a unified view of the network hard to achieve. The situation is further complicated by the fact that carriers tend to operate within their national boundaries, thus requiring complex inter-carrier agreements for international connectivity. Some challenges associated with the carrier SD-WAN approach include:

- » Lack of unified configuration and monitoring makes providing agile, on-demand services difficult.
- » Inter-carrier agreements on international routes make them expensive and complex.
- » Many carriers provide no choice to their customers for first-mile or last-mile connectivity, often forcing them into expensive contracts.
- » With so many different types of equipment and multiple service providers, ensuring consistent SLAs and service quality can be challenging. Carriers and telcos are unable to deliver on the cloud-first experiences demanded by business today. This is reflected in middling to low Net Promotor Scores (NPS) for carrier MPLS services.



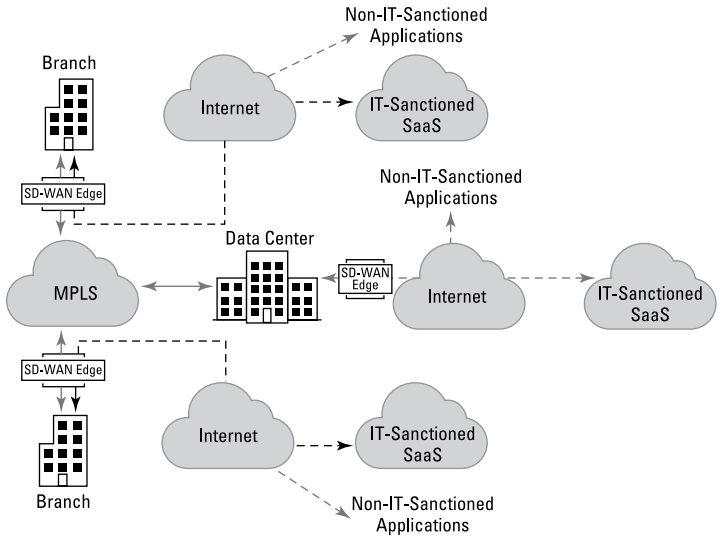
However, there are some advantages in working with a carrier or managed service provider (MSP) in that some have global reach, last-mile assets, and a history of delivering managed services. An organization that uses a given carrier for one service — voice, for example — will be more apt to use this same carrier for others.

## SD-WAN edge overlay: A piecemeal approach

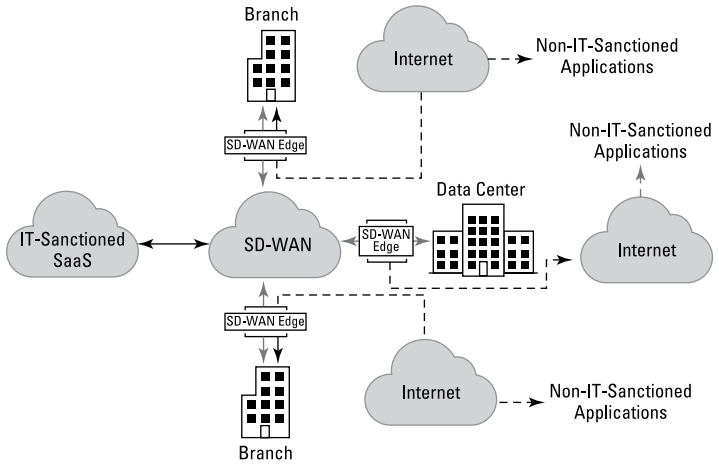
The simplest route to SD-WAN is to deploy it as an edge overlay solution. In this configuration, the overlay solution provides some benefits over the legacy MPLS network because it leverages local Internet connectivity at branch locations. The SD-WAN customer-premises equipment (CPE) provides the necessary functionality to route and distribute traffic between the MPLS network, the Internet, and any other available connectivity (see Figure 1-3). Depending on the network quality, application traffic can be routed via the MPLS network or the public Internet, neither of which is a perfect solution for application performance in the cloud era. Some challenges associated with the SD-WAN edge overlay approach include:

- » When routed over the Internet, application traffic that requires predictable performance is subject to loss and latency issues associated with the Internet. This method is not really designed for global deployments where predictable application performance is required.
- » WAN optimization is an add-on function to MPLS, making the overall solution more expensive.
- » This method relies on the underlying Layer 3 network for quality of service (QoS) convergence and thus can't offer the rock-solid SLAs of a Layer 2 network.

A cloud-first WAN (see Figure 1-4) provides a better approach to SD-WAN that supports enterprise digital transformation with flexibility, speed, and simplicity in the WAN architecture delivered “as a service.” Read Chapter 2 to learn about the different approaches to WAN and SD-WAN deployments, the pros and cons of the different models, and the advantages of a managed, cloud-first WAN experience.



**FIGURE 1-3:** The SD-WAN edge overlay approach.



**FIGURE 1-4:** The cloud-first WAN.

## IN THIS CHAPTER

- » Learning what the cloud-first WAN user experience is all about
- » Comparing the cloud-first WAN to other approaches
- » Putting it all together in a platform-based approach

# Chapter 2

## Introducing a Better Approach: The Cloud-First WAN

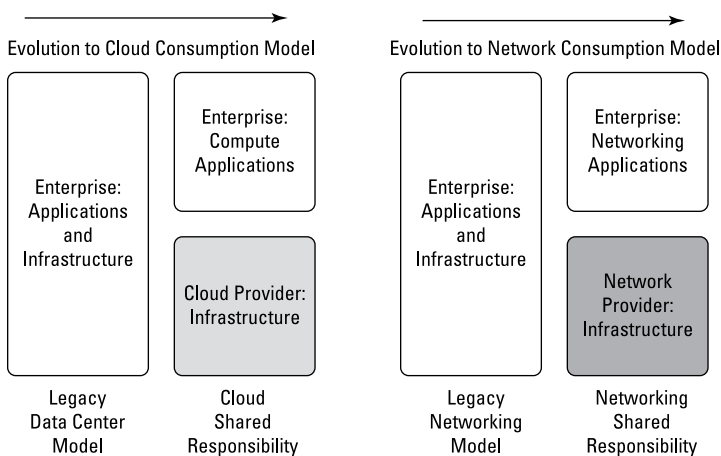
**D**o-it-yourself (DIY) software-defined wide-area networking (SD-WAN) is optimized for speed, but it isn't suitable for global deployments, doesn't give you ownership of the end-to-end user experience, and is full of complexity. Traditional telco-managed wide area networks (WANs) are optimized for connectivity but lack the agility and user experience that modern businesses need. In this chapter, you learn about the cloud-first WAN and its advantages over DIY SD-WAN and traditional telco-managed WAN approaches. You also learn about the importance of taking a holistic, platform-based approach to the cloud-first WAN.

### Defining a Cloud-First WAN Experience-as-a-Service

For network architects and engineers, the WAN experience has traditionally been defined through the lens of telco carrier relationships — and that experience, in many cases, isn't a

particularly positive one. Working with a telco carrier to design an enterprise WAN, negotiate contracts and rates, provision and install new circuits, address performance issues and outages, and manage service-level agreements (SLAs) on an ongoing basis is challenging at best. Add to this mix the complex peering relationships that telco carriers must maintain to deliver global connectivity and the fact that the term “agile” has never been used to describe the telco industry. One might even say the overall experience is a negative one.

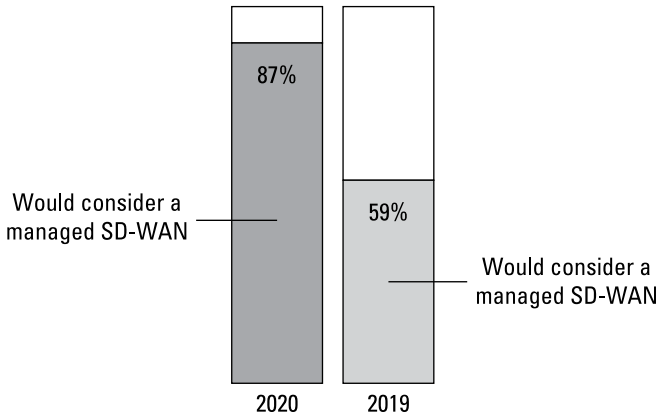
The cloud-first WAN experience changes this paradigm. It isn't just about public clouds, but rather an overall experience predicated on business agility, operational simplicity, and consistent multi-cloud deployments, all while leveraging the cloud consumption model. Think about the way computing has evolved from the enterprise owning, operating, and maintaining its applications and infrastructure in an on-premises data center to a cloud consumption model. In the same way, the cloud-first WAN experience evolves the legacy networking model to a network consumption model in which the enterprise and the network provider share responsibility for the WAN applications and infrastructure (see Figure 2-1).



**FIGURE 2-1:** The cloud-first WAN transforms the network from a legacy model to a shared responsibility model similar to cloud computing.

At the same time, the consumption model simplifies service delivery transforming it from a manual process requiring a great deal of training to an automated process that many term “intent-driven.”

Speaking of the shared responsibility model, a greater percentage of IT planners are open to considering a managed SD-WAN service, growing from 59 percent in 2019 to 87 percent in 2020 (see Figure 2-2).



**FIGURE 2-2:** Percentage of IT planners open to managed SD-WAN.



REMEMBER

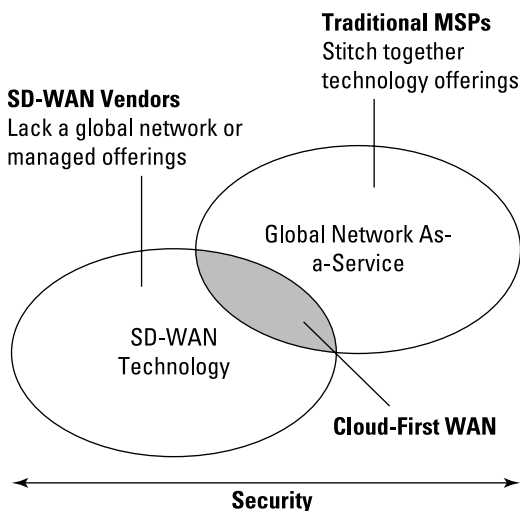
The cloud-first WAN experience delivers predictable end-to-end performance consumed “as-a-service” for an amazing user and application experience.

## Business agility

Time-to-market for businesses in the cloud era is typically defined in minutes, hours, days and, at the high end of the spectrum, weeks. Fast time-to-market in the WAN requires a cloud consumption model — an “as-a-service” OPEX-based offering with flexible billing and ease of service integration. This new network consumption model drives agility, permitting IT and infrastructure teams to rapidly adapt to the needs of the business including rapidly changing business priorities, integrated supply chains, and globalization demands.

## Operational simplicity

Operational simplicity comes from a best-of-breed managed service that permits enterprises to radically simplify complexity. The cloud-first WAN offers a unique take on the WAN consumption model by delivering both the technology (SD-WAN) and the managed service. This “best of both worlds” model offers IT the power of “and” instead of “either-or” (see Figure 2-3).



**FIGURE 2-3:** The cloud-first WAN offers a global or regional secure SD-WAN deployment on a purpose-built, multi-cloud, global L2 WAN optimized network.

## Multi-cloud readiness

A multi-cloud ready architecture offers choices to bring any application to any cloud by connecting public cloud providers, Software-as-a-Service (SaaS) providers, and partner clouds, while delivering a consistent user experience. This capability is the linchpin of a cloud-first service offering, an offering with the extensibility to connect to any Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), or SaaS provider in any region with minimal effort.



REMEMBER

IT must have the flexibility to deploy any application, anywhere, accessible by any employee in any location and at any time. In essence, you should have a local area network (LAN)-like experience extended to the WAN.

## Trials and tribulations with UCaaS

Unified Communications-as-a-Service (UCaaS), in particular, brings many WANs to their knees. IT is faced with the challenge of delivering peak performance to employees anywhere in the world, and at any time. The events of 2020 propelled these challenges to the forefront, with performance issues appearing across the first-, middle-, and last-miles. Aryaka's "State of the WAN"

report identified setup and management of the network as the top issue, surpassing symptoms visible to the end-user (see Figure 2-4).




FIGURE 2-4: Difficulties with UCaaS.

## Seeing the Advantages over DIY SD-WAN or a Telco-Managed WAN

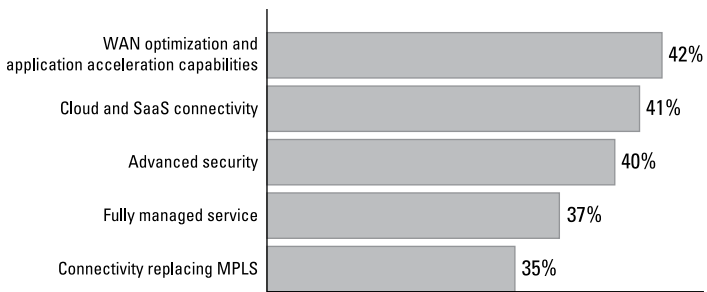
In the post-multiprotocol label switching (MPLS) world, IT planners can typically choose between building or consuming their WAN.

Enterprises that choose to build will source technology from a box vendor and add security, cloud, optimization, and orchestration components.

IT managers have high expectations of SD-WAN, including the various service components. Aryaka's "State of the WAN" report pinpointed these requirements (see Figure 2-5).

This DIY approach isn't all that simple in an era of expertise gaps, and enterprises often reach a dead end because of cost, complexity, or timing. Some challenges associated with the DIY SD-WAN approach include:

- » **Forklift upgrades:** DIY SD-WAN rollouts invariably involve hardware changes, inventory management, version control, patching issues, and more. Even in the case of SD-WAN as a software upgrade, the performance of legacy hardware deteriorates with the addition of SD-WAN features, thus necessitating an eventual hardware upgrade.



**FIGURE 2-5:** SD-WAN expectations.

- » **Lack of end-to-end security:** Because of multiple moving parts, implementing consistent security across the edge and cloud (for example, providing the ability to encrypt all traffic) is often difficult.
- » **Lack of end-to-end connectivity:** Because they don't take ownership of the last mile, it's hard to guarantee SLAs globally. Most are deployed over the Internet and work only when the quality of the underlying Internet is good.
- » **Not cloud agile:** Digital enterprises operate in an environment that requires agility. Rolling out new cloud applications, ramping down or migrating from legacy applications, and the opening and closing of branch and remote locations all require changes to the WAN.
- » **Slow rollout:** Equipment lead times, configuration, testing, and modifications in contracts with multiple last- and middle-mile service providers can delay the rollout. In the latest Aryaka "State of the WAN" report, 75 percent reported deployment times of greater than a week, with 28 percent stating that bringing up a new location took longer than a month.
- » **Sluggish applications:** Lack of direct on-ramps to cloud service providers can reduce cloud application performance. Variable latency and data loss can affect real-time, low-latency applications like UCaaS. Lack of built-in WAN optimization or application acceleration technology also degrades the user experience.



- » **Overlay issues:** An overlay deployment, such as with MPLS, results in longer quality of service (QoS) convergence times and, consequently, delayed SLAs. In addition, there is diminished end-to-end control of SLAs across a DIY global backbone because of the separate visibility of the overlay and underlay. These issues make it harder to correlate faults and provide minimal to no correlation between the overlay and underlay.
- » **Complex operations and multiple proofs of concept (POCs):** Building your own WAN requires contracts with multiple original equipment manufacturers (OEMs) and service providers, resulting in increased complexity. Problem resolution involves multiple POCs and separate contact lists for first- and middle-mile connectivity. This approach is not aligned with the cloud-like consumption model that CIOs prefer for their applications.

Another option is to consume the WAN from a service provider. The provider, in turn, sources the technology from a box vendor. This approach still doesn't engender a truly seamless experience because of the moving parts between the service provider and the technology vendor, as well as between the provider's "underlay" network and the SD-WAN technology vendor's "overlay." Challenges with this approach include:

- » **Last mile lock-in:** Carriers tend to lock customers into their last-mile solution, rather than let them choose the best available option. The last-mile service can create a poor overall user experience and defeats the agility of an SD-WAN approach.
- » **Slow rollout:** Equipment lead times, configuration, testing, and modifications in contracts with OEMs result in rollout delays.
- » **Spotty cloud co-location:** Carriers are not always co-located with cloud service providers — like Amazon Web Services (AWS), Microsoft Azure, or Google Cloud — which makes it a challenge to ensure cloud application performance and optimized regional connectivity.

- » **Not agile:** Digital enterprises operate in an environment that requires agility. Some of the events that require changes to the WAN include rolling out new cloud applications and ramping down or migrating from legacy applications, as well as opening and closing branch and remote locations. A legacy WAN is incapable of keeping pace with rapid changes.
- » **Inconsistent service-level agreements (SLAs):** Carriers typically operate within a specific service area, such as a single country or region. International connectivity is achieved through peering arrangements with multiple service providers, thus making it impossible to guarantee end-to-end SLAs. As a result, the service provided is only as good as least common denominator among the patchwork of providers.
- » **Inflexible pricing:** Carrier networks involve agreements among multiple service providers. As a result, their pricing model is designed to compensate every player in the value chain, thus making them inflexible and expensive.
- » **Low Net Promoter Score (NPS):** Multiple surveys generally rate carriers poorly overall in NPS surveys. These low scores stem from many challenges including:
  - Being consumers rather than creators of technology
  - Dependence on various OEMs
  - The need for complex inter-carrier agreements
  - The mandate to protect legacy investments in MPLS
  - The tendency to lock customers into first- and last-mile offerings
- » **Creators versus consumers:** SD-WAN is not a single-box, plug-and-play solution. A comprehensive SD-WAN solution requires interworking among various elements. As consumers rather than creators of the technology, carriers are limited in their ability to offer best-in-class service.

Table 2-1 summarizes the pros and cons of the DIY SD-WAN, traditional managed service provider, and managed cloud-first WAN approaches.

**TABLE 2-1 Comparing DIY SD-WAN, MSP, and Cloud-First WAN Approaches**

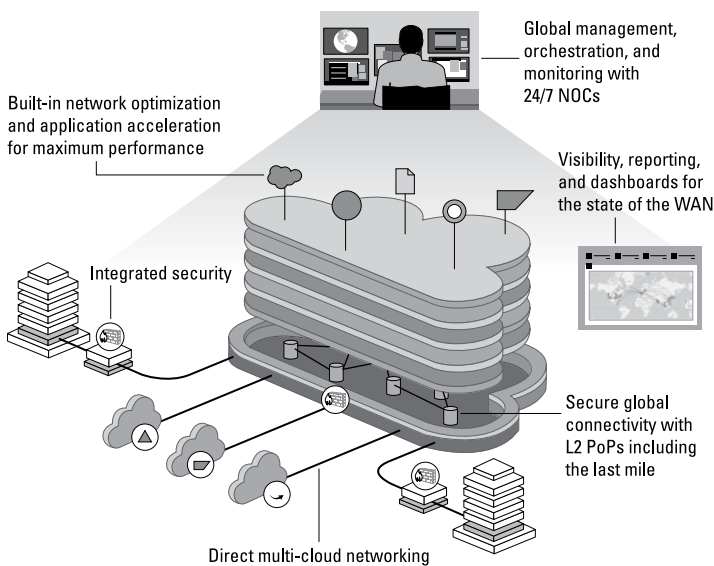
Feature	DIY SD-WAN	Cloud-First WAN	Traditional Telco MSP
Fully managed service with 24/7 support and global network operations centers (NOCs)	No	Yes	Yes
Managed last-mile connectivity with procurement and monitoring	No	Yes	Yes
Multi-cloud architecture with direct connectivity	No	Yes	No
Guaranteed application performance with built-in WAN optimization	No	Yes	No
Global L2 points of presence (PoPs) reaching 95 percent of world's knowledge workers	No	Yes	No
Consistent global SLAs with single point of contact (SPOC)	No	Yes	No
Secure transport, edge, and managed firewall-as-a-service	No	Yes	Yes

## Adopting a Platform Approach

While traditional SD-WAN vendors take a box-centric view with little accountability for end-to-end global experience, traditional service providers stitch together technology offerings from multiple vendors and consequently must compromise on delivering a seamless experience.

The path forward is to take a platform approach that leverages a unified service architecture that is extensible and reaches end-to-end. This platform must offer the service sophistication that enables the suite of connectivity, cloud, security, and optimization services — services that are deployed in a SaaS model to all customers and that are easily enhanced. As you might imagine, the sophistication of the service nodes that enable these services is an order of magnitude in capability beyond a simple PoP (see Chapter 3 to learn more about service nodes or PoPs).

This platform approach, whether deployed regionally or globally, must also leverage a sophisticated orchestration engine that offers the visibility and control into the end-to-end deployment — the first-, middle-, and last-miles. All the hardware in the world will be ineffective if the SD-WAN offering results in a piecemeal operational model (see Figure 2-6).



**FIGURE 2-6:** Taking a platform approach with the cloud-first WAN.



REMEMBER

End-to-end accountability requires a platform approach. A fully managed cloud-first WAN does this with a private global Layer 2 core and numerous points of presence (PoPs) across the globe. In Chapter 3, you learn about the key elements of a cloud-first WAN.

#### IN THIS CHAPTER

- » Delivering first-, middle-, and last-mile connectivity
- » Taking an integrated approach to network security
- » Enabling direct multi-cloud access
- » Maximizing application performance
- » Providing end-to-end visibility
- » Automating deployments with powerful management and orchestration

## Chapter 3

# Discovering the Key Elements of a Cloud-First WAN

**T**his chapter identifies the key elements to look for in a fully managed cloud-first WAN.

## Connectivity

For enterprises operating in multiple regions with traffic traversing the WAN core, a global deployment provides the required connectivity. Some traffic may optionally leverage a hybrid WAN capability for site-to-site direct Internet access (DIA) connectivity.

Enterprise connectivity requirements are driven by the needs of their applications. Some applications can be best served by increasing bandwidth. Others are latency and jitter sensitive

and need more predictable traffic engineering and optimization. Bandwidth costs also vary from one region to another as does the type of connectivity, whether it's the last mile or the core. Enterprises need options to address the different needs of their applications based on cost, application performance, and service-level agreement (SLA) considerations.

Traditional SD-WAN vendors can only provide overlay connectivity on top of the Internet. Telcos have traditionally dealt with multiprotocol label switching (MPLS) or fragmented connectivity based on global considerations, which dilutes SLAs.



REMEMBER

Your service provider needs to have global ownership of its SLAs to deliver the end-to-end application performance and user experience that your enterprise needs.

For the best performance, enterprises require flexible connectivity based on application performance, cost, and accessibility without having significant management overhead. For example, you can achieve regional or global connectivity over a guaranteed private core or a hybrid mix of Internet and private links.



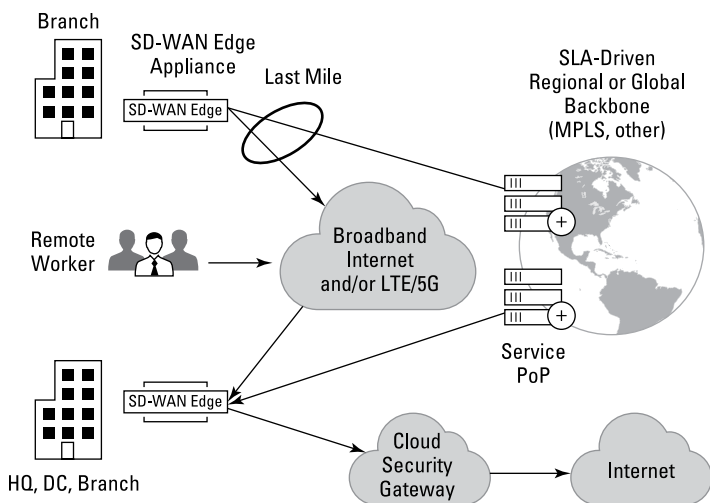
TECHNICAL  
STUFF

In a software-defined wide-area networking (SD-WAN) architecture, *hybrid WAN* typically refers to traditional multiprotocol label switching (MPLS) connectivity combined with direct Internet access (DIA), such as broadband, and/or Long-Term Evolution (LTE) wireless connectivity.

For enterprises operating primarily in a single region — defined as a regional cluster or point of presence (PoP) — with some traffic optionally carried over the network core to another region, a regional deployment is the right option. Here, most traffic will leverage a hybrid WAN capability offering site-to-site DIA connectivity.

Connectivity also extends to the last mile, a part of the network sometimes glossed over or left as an afterthought. But the last mile — procurement, provisioning, monitoring, and troubleshooting — can make or break the utility of an SD-WAN service because it is across these broadband Internet links that the SD-WAN provider has less control.

Figure 3-1 illustrates the requirements for end-to-end enterprise WAN connectivity.



**FIGURE 3-1:** End-to-end enterprise WAN connectivity.

## Security

The cloud-first WAN delivers integrated security capabilities in a platform architecture that includes features such as cloud security, micro-segmentation, secure remote access, and edge firewalls. For enterprises (or readers) asking, “Where is the most appropriate place to deploy this security functionality?” see the sidebar, “Is Cloud-First WAN ‘Sassy’ (SASE)?”

In a typical enterprise’s hybrid architecture, security may consist of various point security solutions in different physical and virtual form factors at the headquarters, the data center, the branch, and in the cloud — as well as security solutions to protect remote workers. The enterprise should have the flexibility to set their security parameters, as well as the option to use their existing security vendor as part of their SD-WAN deployment.

For example, an enterprise may have a combination of physical and virtual appliances at larger sites and select smaller sites. In other locations, the enterprise may hand off traffic to a cloud security gateway. The key is flexibility, and the understanding that the overall security posture of the enterprise is dependent on its weakest link. The addition of multiple or less capable security vendors may compromise the enterprise’s overall security posture.

## IS A CLOUD-FIRST WAN “SASSY” (SASE)?

A number of vendors have adopted the term *SASE*, or *secure access service edge*, first proposed by Gartner, to describe their converged SD-WAN and security offerings. More than a few have stated that SASE will replace SD-WAN as a way to describe this converged architecture. One reason this book is titled *The Cloud-First WAN For Dummies* instead of referring to *SD-WANs* is to stay above this fray.

Gartner defines SASE (pronounced “sassy”) as an “emerging offering combining comprehensive WAN capabilities with comprehensive network security functions (such as CASB, NGFW, DLP, ZNTA, SDP, VPN, WAF, RBI, and Sandboxing) to support the dynamic secure access needs of digital enterprises. SASE capabilities are delivered predominantly as a cloud-based service based upon the identity of the entity, real-time context, enterprise security/compliance policies and continuous assessment of risk/trust throughout the sessions.” Within SASE, a “heavy” branch with a full security stack transitions to a “thin” branch with most security functionality within a “heavy” cloud.

What you call your transformed WAN is partially a result of the lens you use, be it networking, WAN optimization, the cloud, or security. The cloud-first WAN offers the elasticity, agility, and power of choice, allowing it to be optimally tailored to the networking and security needs of any enterprise. Given that SASE proposes a cloud-first approach to security, there is alignment.

Depending upon the enterprise’s individual journey to the cloud, the organization may have a combination of edge and cloud-delivered security capabilities or a combination of heavy and thin branches, and some may never totally migrate to a pure heavy cloud architecture. The organization may also choose to adopt a single-vendor approach or may select different security vendors for different parts of its network. In any case, the cloud-first WAN concept is inclusive rather than exclusive, an architectural approach that delivers “and” instead of “or.”

Keep in mind, however, that a SASE offering cannot be delivered without a cloud-first WAN. Instead, the cloud-first WAN lays the foundation for the SASE architecture. At the time of this book’s publication, vendors were still developing their SASE architectures and none could be said to have a complete implementation.



Lastly, remote employees must have secure access to enterprise resources, as well as access to cloud-based applications, without compromising security. This is the “SD-WAN branch of one,” and the same policies and rigor implemented across the core must extend to your remote workers. Cloud-native security gateways address this requirement. Figure 3-2 illustrates the security requirements for the enterprise WAN.

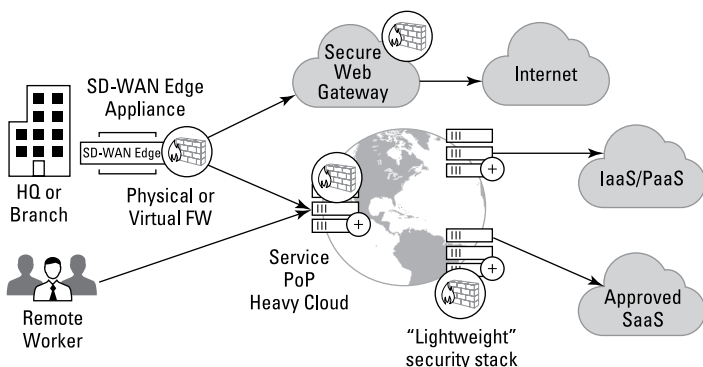


FIGURE 3-2: Security requirements for the enterprise WAN.

## Multi-Cloud


The fully managed cloud-first WAN provides plug-and-play multi-cloud and SaaS connectivity service that enables the enterprise to spin up connectivity to Infrastructure-as-a-Service (IaaS) and SaaS on-demand within a few hours. Core capabilities include:

- » **Cloud acceleration:** Enterprise users connect to their SaaS applications over a service-level agreement (SLA)-driven connection from the nearest PoP. This design overcomes the latency and packet loss issues associated with the Internet middle-mile and provides a cost-effective yet superior alternative for accelerating cloud application performance.
- » **Public cloud direct connectivity:** This feature provides an out-of-the-box on-ramp to popular IaaS providers such as Amazon Web Services (AWS), Microsoft Azure, and

Google Cloud. Regionally distributed high-speed links directly from the cloud-first WAN PoPs support Direct Connect to AWS, ExpressRoute to Azure, and Dedicated Interconnect to Google.

- » **Public cloud integration:** This feature optimizes connectivity and manageability for large public cloud users. For example, in Microsoft Azure, an organization's virtual network (VNet) becomes part of the enterprise WAN. Resources deployed in the cloud become available over the integrated WAN service.

Figure 3-3 illustrates the enterprise need for multi-cloud capabilities in a WAN solution.



**FIGURE 3-3:** The WAN architecture must provide secure and reliable connectivity to multiple public, private, and hybrid clouds.

## Application Acceleration and Optimization

Software-as-a-Service (SaaS) application performance is not just a matter of adding software-defined wide-area networking (SD-WAN) equipment into the existing network. Ensuring good application performance requires a holistic, cloud-first WAN approach that accounts for foundational aspects of technology such as:

- » **Capacity:** Optimal capacity must be provided for agility and scalability.

- » **Availability:** Superior availability is achieved through a combination of SLAs, built-in redundancy, and other redundancy options.
- » **Security:** Security, including third-party integrations, must be part-and-parcel.

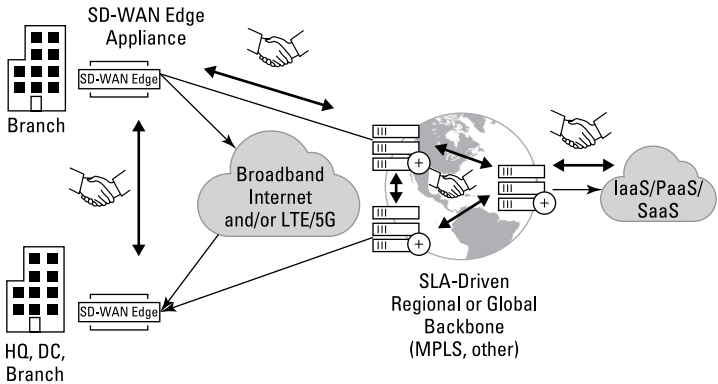
Though important, these foundational aspects alone are not enough. Building on this foundation, an effective cloud-first WAN solution must also address the following:

- » **Quality of service (QoS):** Customers should be able to easily flag and prioritize their applications and traffic on the network with intuitive classifications like transactional, real-time, productivity, critical, and best effort.
- » **Topology:** Users should connect to SaaS applications in a full mesh architecture regardless of where the applications reside, rather than backhauling traffic through multiprotocol label switching (MPLS) headend locations and data centers, which further increases latency and unpredictability.
- » **Application routing:** Connectivity to SaaS applications like Office 365, Salesforce, or WebEx is a challenge. Traditional connectivity solutions for access SaaS applications depend on the public Internet, which can be slow and unreliable in places.
- » **Application acceleration and optimization:** Data deduplication, compression, bandwidth management (QoS, prioritization), Secure Sockets Layer (SSL) acceleration, and other innovations accelerate and optimize application performance in the cloud-first WAN.

Another important area to consider is the deployment model (discussed in Chapter 2), namely do-it-yourself (DIY) versus a managed service. Organizations need to decide whether it is more cost effective to constantly recruit, train, and upskill employees or to leave the complexity to specialty players and simply consume connectivity as a service.

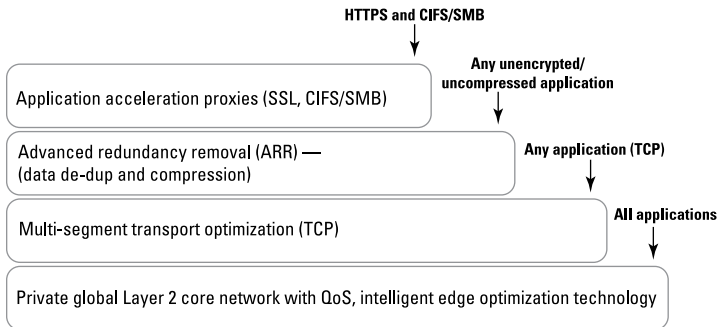
Bringing the cloud-first WAN solution all together is the process, which should be simple but still allow technology to move at the pace of the business.

Figure 3-4 illustrates the need for the enterprise WAN to optimize multiple links and connections.




**FIGURE 3-4:** Providing WAN optimization capabilities across different links and connections.

Figure 3-5 depicts the different layers of network and application optimization possible across the SD-WAN first and middle-mile. One outcome of this is the ability to reduce the perceived connection setup time, where application throughput increases. Optimization also improves the performance of collaboration applications, as evidenced by the mean opinion score (MOS), as shown in Figure 3-6.



**FIGURE 3-5:** Layers of network and application optimization.



Even though the physical latency is ~200–210 ms, the latency perceived by applications is ~5–30 ms. Applications begin behaving for this latency.

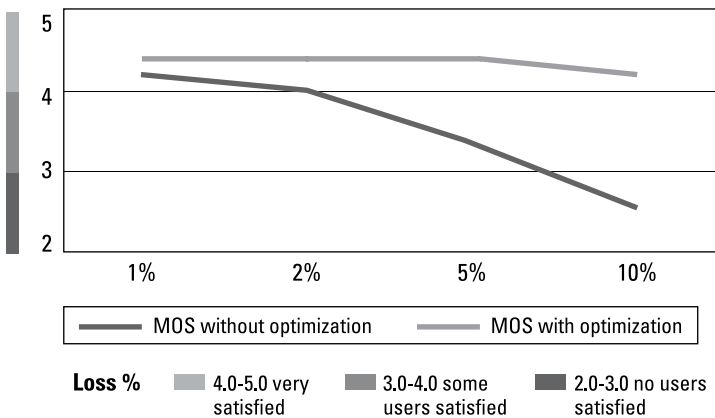


FIGURE 3-6: Improving the performance of collaboration applications.

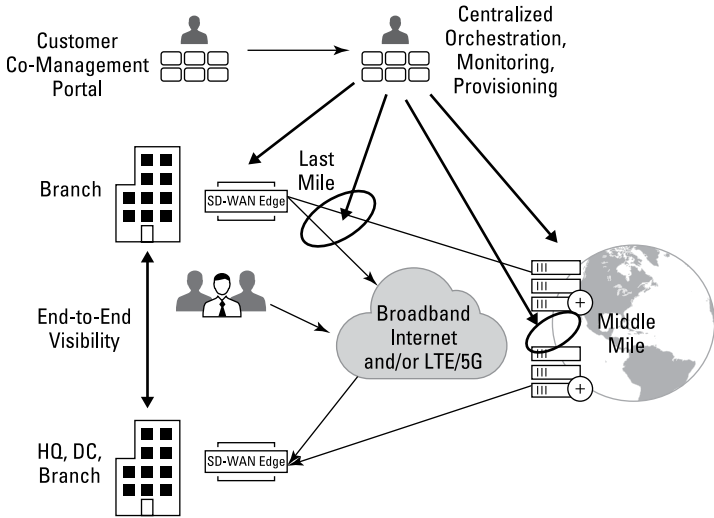
## Visibility

End-to-end visibility of the entire WAN is a critical component of the cloud-first WAN. Traditional MPLS WAN architectures augmented with piecemeal DIA connections from a multitude of local Internet service providers (ISPs) and telco carriers, as well as various service bolt-on components, are unable to provide this “single pane of glass” visibility. This may result in individual visibility silos and blind spots. Without complete visibility, network teams cannot effectively manage performance and bandwidth utilization, troubleshoot network issues, and secure the network. It’s like trying to run a network with one arm tied behind your back and one eye closed!



TIP

A fully managed cloud-first WAN enables end-to-end visibility through an intuitive web-based portal that provides real-time contextual insight into your network (the “state of the WAN”) and applications and a wide range of functions to speed up service delivery (see Figure 3-7).



**FIGURE 3-7:** An intuitive customer co-management portal and centralized orchestration, monitoring, and provisioning are key enterprise WAN capabilities.

## Management and Orchestration

The essence of a “software-defined” WAN is centralized orchestration. Not every forwarding and policy decision must be made centrally. On the contrary, orchestration maintains visibility over the end-to-end deployment, from the first- and last-mile to the service nodes and the SD-WAN edge appliances. Changes to an enterprise’s topology, including the addition of new sites, are seemingly instantaneous, requiring hours or days instead of weeks and months.

Besides provisioning, any anomalies across the enterprise’s SD-WAN may be immediately identified and corrected. Looking to the platform concept (introduced in Chapter 2), to be truly effective, the orchestration must have visibility into not only the WAN infrastructure itself, but also the additional security, multi-cloud, and optimization services consumed by the enterprise as part of the managed SD-WAN.



REMEMBER

Management and orchestration in the cloud-first WAN enables automation at scale with predictive analytics to deliver proactive application performance optimization, enhanced security, and more.

It's also important to know “who you gonna call” when something goes wrong on the WAN. Particularly in the case of DIY SD-WAN deployments, enterprises sometimes take on more than they can handle. It's important to have a managed, cloud-first WAN provider that offers global 24/7 service to address service delivery, troubleshooting, and support. This human aspect of WAN management is also critical to delivering the same kind of exceptional experience as the technology and deployment they support.

## THIS ISN'T YOUR POP'S POP — THE IMPORTANCE OF SERVICE DELIVERY

There is a great deal of market confusion regarding how best to architect an SD-WAN, and more specifically, the points of presence (PoPs) through which the enterprise traffic flows. In the case of an SD-WAN overlay, the PoP integrates routing and most likely switching, forwarding the MPLS and/or IP traffic from one locale to another. This PoP also has a complement at the transmission layer — the underlay. And, at a minimum, what an SD-WAN vendor terms a PoP may be only a light-weight virtual machine spun up in the public cloud. Are either of these approaches sufficient?

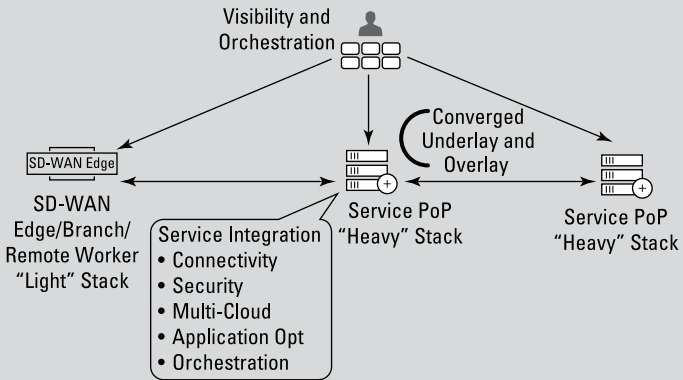
Considering the service sophistication demanded by IT as part of their WAN transformation and some of the complaints leveled against current approaches, the security, optimization, multi-cloud, and orchestration capabilities require much more than a mom-and-pop (or your pop's) PoP.

The solution is what Aryaka refers to as an *SD-WAN service PoP*. It not only integrates routing and switching, but also includes compute and storage. This richness, in combination with the SD-WAN edge, supports the set of services that enable a truly functional SD-WAN that meets enterprise business objectives.

(continued)

(continued)

The figure highlights the rich handshake between the PoP and the SD-WAN edge appliance, with the PoP handling the heavy services lifting and the edge, with a lighter weight stack, as the services endpoint. In terms of the SASE description earlier in this chapter, this architecture is in full alignment with the concept of a “thin branch” and “heavy cloud.”



A related discussion is whether the node operates at Layer 2 or Layer 3. A Layer 3, routing-only PoP will not have full visibility into the underlying transmission architecture, and performance is dependent upon peering between ISPs, which is never an exact science when speaking about QoS. Sure, there are ways to gain partial visibility, but guaranteeing an end-to-end SLA isn't simple, especially as part of a global deployment.

A more effective architecture is a Layer 2 PoP, which by definition also includes routing. Here, the PoP has full visibility into the underlying transmission infrastructure and the provider operating the PoP controls the direct Layer 2 connectivity from one PoP to another. The QoS across this connection is therefore very deterministic, and end-to-end guaranteed SLAs are a reality assuming a global footprint. A Layer 2 PoP is never more than 25 milliseconds from the nearest knowledge worker and is built with a multi-segment architecture. Visibility offers a converged view spanning both Layer 3 (and above) and Layer 2.

Finally, Layer 2 PoPs can be easily extended to public, private, or partner clouds, as well as SaaS providers, allowing for seamless managed networking and multi-cloud connectivity with consistent treatment for applications across clouds.



## IN THIS CHAPTER

- » Improving application performance
- » Migrating from multiprotocol label switching (MPLS) networks
- » Supporting multi-cloud connectivity
- » Accelerating digital transformation
- » Extending your global network into China
- » Delivering reliable network performance for voice and video

# Chapter 4

## Exploring Use Cases for the Modern Enterprise

This chapter introduces real-world cloud-first wide-area network (WAN) use cases and shows you how Aryaka helps its customers address their networking and digital transformation challenges.

### Accelerating Application Performance

Organizations are increasingly migrating their on-premises applications to Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) cloud platforms — such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud — and adopting Software-as-a-Service (SaaS) applications — such as Office 365, Salesforce, and WebEx or 8x8. However, many of these organizations find that their user experience has deteriorated as they experience new application performance challenges including:

- » Many applications are slow and sluggish despite deploying multiprotocol label switching (MPLS) or even software-defined wide-area networking (SD-WAN).

- » Existing WAN infrastructures based on MPLS architectures can't deliver the required agility.
- » Employee productivity and corporate profitability suffer as the WAN becomes a barrier to successful digital transformation.
- » Once a WAN connectivity solution is operational, challenges shift to centralized configuration and monitoring.
- » Monitoring WAN performance at an aggregate link level is insufficient for cloud-based SaaS applications.
- » Administrators need to be aware of application-specific data flowing through their network and be able to detect and fix any performance degradation in end-user applications.



TIP

A fully managed cloud-first WAN includes features such as built-in WAN optimization and direct connectivity to leading IaaS, PaaS, SaaS, Unified Communications-as-a-Service (UCaaS), and other “XaaS” service providers to ensure optimal application performance. Capabilities and benefits of a fully managed cloud-first WAN include:

- » **Low latency and jitter:** Delivers SaaS acceleration through a private, software-defined Layer 2 network, with points of presence (PoPs) located within 25–30 milliseconds from leading SaaS and IaaS providers
- » **Compression:** Reduces the file size of data that is transmitted over the network, optimizing use of expensive regional and global bandwidth
- » **Direct connectivity (IaaS and PaaS):** Provides out-of-the-box connectivity to leading IaaS and PaaS providers with pre-wired regional links such as Direct Connect to AWS or ExpressRoute to Microsoft Azure
- » **Multi-segment optimization:** Achieves optimal application performance with independent first-mile, middle-mile, and last-mile proxies, optimizing the data flow by reducing the time taken for the first-byte transfer, using bigger payloads sizes per packet, and providing recovery from up to 5 percent packet loss
- » **WAN optimization:** Includes data deduplication, compression, bandwidth management (quality of service [QoS], prioritization), and Secure Sockets Layer (SSL) acceleration

# MANUFACTURER IMPROVES PRODUCTION TIME BY 20X

A global leader in metal-cutting and manufacturing for more than 80 years has locations in Europe, North America, Japan, and other regions in Asia. As part of its digital transformation, the company has been building a global network to meet its business expansion requirements and expedite data exchange between distributed research and development (R&D) departments and tech centers.

## Challenge

The company synchronizes massive amounts of data for machine specs and schematics from its headquarters in Tokyo to its tech center in Ohio. File synchronization took roughly six to seven hours daily and had to be done overnight to minimize adverse effects on performance. Sometimes the process was not completed by the time U.S. employees started their workday, creating a drain on operational efficiency. In addition, the company had recently acquired several other companies and knew that legacy network solutions such as MPLS and WAN optimization hardware would not be able to keep pace with the number of sites the company needed to onboard quickly.

## Solution

WAN optimization hardware had been implemented in the past to improve data and application delivery, but to upgrade at every existing site would have been cost prohibitive. The company also considered deploying Internet-based SD-WAN, though that wouldn't have solved the latency issues between the sites in Asia and the United States. It also wouldn't have provided the stable connection that the company needed for data synchronization.

Instead, the company went with Aryaka SmartServices. After deploying the solution, the company noticed a dramatic improvement in performance and data transfer times immediately. The file synchronization that had taken 6 to 7 hours now took only 22 minutes, which allowed the company to become more responsive. This improvement opened new possibilities for the business, giving the company a massive competitive advantage. As additional sites were needed, setup took two or three days compared to the weeks or months it would take for an MPLS deployment, allowing the company to ramp up its business rather than waiting to bring sites online.

*(continued)*

(continued)

The company is now embarking on its cloud strategy with a goal of having 90 percent of its data and applications in the cloud. Because Aryaka provides accelerated access to any application, on-premises or in the cloud, the company now has an infrastructure in place to handle the upcoming migration and can deliver data and applications to every end-user as if it lived in the local data center.

### Results

- Data replication times have been reduced from 6–7 hours to 22 minutes.
- Application performance has improved 20x.
- As much as 99 megabits/second (Mbps) of peak bandwidth has been saved.
- The company has experienced a 97 percent data reduction across applications.
- Deployment now takes days, instead of weeks or months, for MPLS.

## Finding a Flexible, Simpler Alternative to MPLS

MPLS still represents a foundational technology in most global enterprise WANs. MPLS can deliver high availability and deterministic QoS within a service provider's domain, but it is also costly and slow to deploy, and its traditional hub-and-spoke architecture from the branch to headquarters and/or to the data center does not support the overwhelming need to optimally support cloud deployments. Do-it-yourself (DIY) SD-WAN solutions allow traffic to be routed directly to the Internet at the branch. However, they often still rely on costly MPLS links to support traffic that is perceived as business critical.



TIP

A fully managed cloud-first WAN can deliver the predictability of MPLS with the agility of SD-WAN, offering capabilities and benefits that include:

- » **Deployment speed:** Branch connectivity is available anywhere within 48 hours, with “Day-1” service-level agreements (SLAs).

- » **Lower total cost of ownership (TCO):** The cost is lower compared to MPLS, with optimization, security, and multi-cloud connectivity.
- » **Low latency:** Less than 30 milliseconds (ms) of first-mile latency is available to 95 percent of knowledge workers globally.
- » **Simplicity:** Deployment is simple, with a fully managed model based on business intent.
- » **Superior performance and connectivity:** Deterministic traffic behavior delivers performance equal or superior to MPLS, with hardened last-mile Internet connectivity leveraging the best local Internet service providers (ISPs) and technologies to eliminate packet loss while minimizing latency and jitter.
- » **Operational excellence:** Network managers can deliver on a global network infrastructure optimally suited to the needs of digital business.

## CHEMICALS COMPANY REPLACES MPLS NETWORK WITH ARYAKA SERVICES

Element Solutions Inc. (ESI) is a global, diversified manufacturer of high-tech specialty chemicals and electronics products. ESI selected Aryaka to replace a global MPLS and ad-hoc tunnel network with SmartServices.

### Challenge

ESI is a blend of multiple distinct business units including a series of acquisitions that played a major role in the company's evolution in recent years. The company developed a strategy of aligning shared services with business operations. This strategy created several global business challenges, which required the IT organization to quickly and effectively solve the unification of network communications among multiple businesses.

The challenge was to consolidate a wide-area network comprised of variable routing methods. Connecting all global users to a set of combined and common business critical applications proved to be difficult over a disparate network.

*(continued)*

(continued)

While ESI contemplated a global, traditional MPLS deployment, the estimated 18–24-month time to convert more than 200 locations in dozens of countries and the associated high costs led to an exploration of other options. ESI simultaneously adopted a “cloud-first-and-only” consolidation approach for production workloads, with a deliberate intention of eliminating on-site data centers, servers, and other network services. After considering each of the major SD-WAN providers, ESI determined that a combination of network optimization, automated routing, and fully managed middle-mile solution was the best catalyst toward this overall strategy.

### **Solution**


After deploying Aryaka services in 200 sites, including 5 cloud data centers, the company consolidated workloads, carving several months from the original project plan and millions in projected operating costs. ESI and Aryaka teamed together to deploy most of the network in less than six months. The scope of the deployment included major efforts in China with no delays or downtime. The results were significant cost savings and accelerated application performance to end-users.

Not long after the Aryaka SD-WAN deployment, ESI announced the sale and divestiture of the Arysta business unit. ESI IT was tasked with the orderly separation of about 40 percent of the global sites. ESI again turned to Aryaka to build a strategy to split the network in preparation for the separation. The effort was completed successfully as expected with zero downtime. A similar transition might have required months with traditional vendors, but Aryaka and ESI managed to split the network in less than two months, saving ESI countless internal IT hours and aggravation.

### **Results**

- ESI achieved \$2 million to \$3 million in cost savings over MPLS.
- Performance for file transfers increased 20X. The figure, “Bandwidth Optimization,” depicts one factor in increasing performance. Connection setup time minimization and latency management are two additional factors.

- Integration with cloud services is seamless.
- ESI gained agility to meet large-scale strategic business objectives.



## Connecting to Any Cloud Anywhere with Ease

For most organizations, a multi-cloud architecture is a practical reality. According to the Aryaka 2020 “State of the WAN” report, enterprises leverage two or more public cloud platforms and dozens of SaaS applications. In some cases, multi-cloud may be a conscious risk mitigation strategy to avert commercial, technical, or operational reliance on a single service provider. In other cases, it may be a transient migration strategy as the company moves from one cloud service provider to another. In either case, multi-cloud connectivity is a necessity and enterprises must have a consistent approach that normalizes application performance, SLAs, and operations in this heterogeneous environment.

Current approaches for multi-cloud connectivity are inefficient because they require traffic to be backhauled across the entire enterprise network. Legacy WAN architectures weren’t designed for the cloud and place the burden on the enterprise to provision and maintain multi-cloud connectivity. These legacy WAN architectures:

- » Don’t offer seamless connectivity for IaaS, PaaS, and SaaS.
- » Don’t offer adequate application SLAs.
- » Are often inadequate in dealing with the volume and variety of traffic that traverses modern enterprise networks.

Worse, a badly designed WAN can degrade the performance of cloud-based applications and adversely affect end-user experience and productivity.



TIP

A fully managed solution removes the complexities of a multi-cloud operational model to deliver a consistent operational experience to the client, as well as application performance to the end-users.

## FUEL LOGISTICS COMPANY ENABLES MULTI-CLOUD CONNECTIVITY

A leading global fuel logistics company selected Aryaka to assist the company with its digital transformation, network simplification, and managed services. The company markets, sells, and distributes aviation, marine, and land fuel, as well as related products and services, to its clients at more than 8,000 locations in more than 200 countries and territories worldwide.

### Challenge

The company was seeking a unified network architecture across its business locations to deliver traditional and cloud-based services such as Office 365, Box, Slack, and Zoom. Business needs dictated a cost-effective, cloud-based architecture to deploy applications across a secure global network.

“When evaluating vendors for this project, we wanted to see cost efficiencies, last mile management, low latency access to multiple cloud services, and the ability to accommodate all of our diverse global locations,” said the company’s Vice President of Global Infrastructure.

### Solution

The project was completed together with Aryaka’s U.S. partner Pluto Cloud Services, a worldwide leader in emerging technologies specializing in WAN acceleration, application delivery, WAN optimization, telco carrier MPLS, and IP-VPN services. Pluto Cloud Services introduced Aryaka’s solution to the company and managed the communication between the two companies during implementation.



“Being able to provide . . . a global infrastructure that would match [their] needs was our goal,” said Larry Chaffin, CEO at Aryaka partner Pluto Cloud Services. “But it’s more than that as we have provided global telecom procurement as well. It’s all about the relationship with our customers and how we can provide them the best solutions and services around the world.”

#### Results

- One unified global network architecture
- SaaS application support through a cloud-first approach
- Quick site turn-up

## Ensuring a Successful Digital Transformation

Digital transformation initiatives, as well as associated WAN transformation, introduce both challenges and new opportunities. The cloud, big data, social, mobile, and the Internet of Things (IoT) play a central role in this transformation, and it is critical that enterprises get the WAN architecture right. Unfortunately, traditional WAN options are:

- » Inflexible
- » Expensive
- » Incapable of handling the demands of a modern digital enterprise

Direct cloud connectivity, application acceleration, end-to-end security, and global SLAs are features that either are not available in a traditional WAN offering or require expensive new hardware.

The network transformation journey is not an easy one, given the critical nature of connectivity and the consequences of a network outage. Unlike the static enterprise WAN of yesteryear, the enterprise WAN network of today is a living, breathing, dynamic entity that must continually evolve with the needs of the business. As such, both the choice of technology (such as SD-WAN) and the mode of operation matter.

In such a dynamic environment, the DIY approach to SD-WAN has many challenges, including the need for expensive resources and purchasing contracts with original equipment manufacturers (OEMs). Carrier-provided SD-WAN may be an option, but it comes with conditions attached, like last- and first-mile lock-in, long rollout times, no single SLA, and multiple contracts for end-to-end connectivity.

In addition to these IT considerations, organizations must contend with the change management challenges inherent to large digital transformation initiatives that span departments. A common roadblock when rolling out anything new is ensuring adoption so the benefits can be realized. Organizations undergoing a digital transformation also need a reliable network that can remain flexible and adapt to the business's evolving needs.



TIP

A fully managed cloud-first WAN is tailored to the needs of digital and WAN transformations and delivers capabilities and benefits that include:

- » **Business outcomes:** Digital transformation is where SD-WAN crosses over most clearly into the C-suite, with objectives that include organizational speed, competitiveness, productivity, and time-to-market.
- » **Flexibility:** The cloud-first WAN is the most flexible networking solution available. As organizations' priorities or business needs evolve, the cloud-first WAN supports and adapts to those needs. The network is a flexible global solution that permits applications and data to reside anywhere in the world and to be accessible by employees in any region, at any time.
- » **Speed of deployment and ease of use:** A fully managed solution enables IT organizations to implement the network quickly and easily, so they can focus on other priorities.
- » **Security:** Increased traffic and Internet access from the branch increases flexibility and performance. However, it also creates security risks that must be addressed. Installing security hardware in every branch is an expensive proposition. A fully managed cloud-first WAN offers end-to-end security from the physical layer to the data link, network, and application layers.

# TRANSPORTATION COMPANY ACCELERATES TRANSFORMATION

A transportation and logistics company headquartered in Pennsylvania is the largest privately held U.S. freight forwarder with more than 75 offices throughout North America and a global network across Europe and Asia.

## Challenge

When the nearly 50-year old company embarked on a digital transformation initiative with a vision to move to a cloud-based infrastructure, the first step was modernizing the network. The company's core on-premises architecture limited its ability to adopt cloud-based and SaaS applications. The company also lacked full end-to-end visibility into its network.

## Solution

By deploying Aryaka's managed SD-WAN-as-a-service solution, the company has completely transformed its WAN. The company has been able to adopt a cloud-first approach to its infrastructure by leveraging Azure ExpressRoute, adopting Unified Communications-as-a-Service (UCaaS) applications like 8x8 that improve productivity and enable remote workers, and integrating Zscaler to ensure security.

With Aryaka's managed services, the company has also been able to solve the key issue of servicing remote sites that previously had limited IT support and gain visibility into those sites. Free from managing tedious installations and configurations, the company's lean IT team has been able to add more value to the business by focusing on other priorities.

The company's infrastructure soon will be entirely cloud-based. What was originally planned as a three-year transformation has been cut in half.

Looking ahead, Aryaka will continue to be a key partner in enabling the company's growth through global expansion. As the company makes acquisitions across the globe, Aryaka will provide the flexibility to add network connectivity anywhere in the world.

## Results

- Ninety sites were deployed in less than 120 days.
- The timeline for digital transformation was cut in half from three years to one and a half years.

# Optimizing Unified Communications-as-a-Service

Unified Communications-as-a-Service (UCaaS) is expected to reach \$167.1 billion in market size by 2025. As a single integrated solution for collaboration — from email and video conferencing to file sharing and directory services — UCaaS has become an essential tool for global enterprises. However, UCaaS is only as good as the network it's delivered upon. A well-performing UCaaS service can mean better collaboration and improved productivity for a company, while dropped calls, distorted video, and slow-to-send files can result in the opposite. UCaaS solutions are especially prone to packet loss, latency, and jitter — the public Internet can't support UCaaS performance and purchasing private circuits can be prohibitively expensive.

As enterprises roll out UCaaS, availability, user experience, and security are top priorities. However, the lack of deterministic behavior in the underlying Internet network, as well as legacy network architectures that fail to take optimal cloud application support into account, often stand in the way of delivering on these top priorities.



TIP

A fully managed cloud-first SD-WAN can mark UCaaS traffic, steer it optimally and dynamically across Internet access links and through the core infrastructure, minimize packet loss and latency, and deliver an optimized user experience.

## TRANSPORTATION PROVIDER ACCELERATES 8X8 UCAAS

A leading U.S.-based transportation and logistics company has offices throughout the world.

### Challenge

When the company initiated a company-wide digital transformation project, implementing improved performance with its 8x8 UCaaS solution was key to the first phase of the plan. In order to address

performance issues, as well as future-proof its network for future cloud-based and SaaS application deployments, the company needed to transform its core on-premises network.

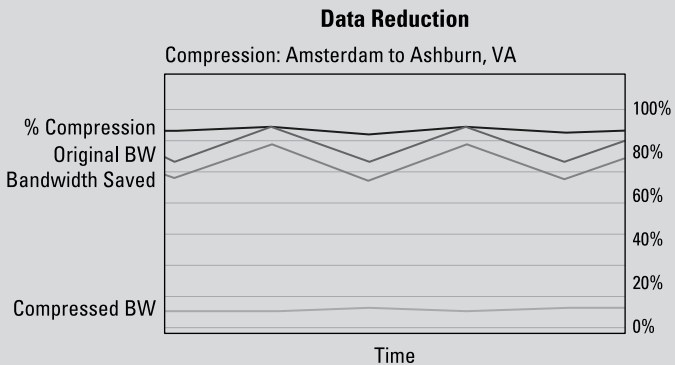
### Solution

Aryaka implements connectivity to 8x8 data centers using its Virtual Office (VO) implementation. It creates five VOs on its backend. Each customer site belonging to Aryaka is mapped to a region, connects to a VO, and is used to access 8x8 for that region. The customer site is mapped based on the PoP to which it connects.

8x8 provides the public IP subnets to Aryaka, which is hosted in Aryaka's data centers. The traffic to the 8x8 global traffic manager (GTM)/domain name system (DNS) server IP address is routed over the regional VO. 8x8 maintains a mapping of the Aryaka public subnets to that region, so that when endpoints register via Aryaka, they are redirected to the services hosted in that region.

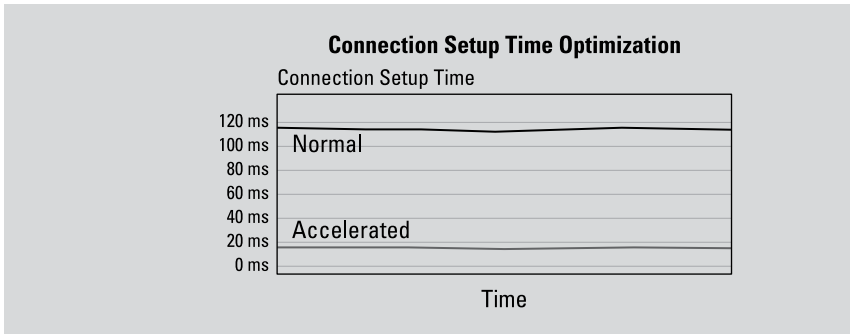
### Results

- Data was reduced by 93 percent and UCaaS performance increased by up to 20x (see the figure, "Data Reduction").
- Packet loss decreased to almost zero.
- Transmission Control Protocol (TCP) connection setup time is now 8x faster over the Aryaka core network (see the figure, "Connection Setup Time Optimization").



(continued)

(continued)



## Finding Faster Connectivity to China and Beyond

Globalization is a core driver for digital transformation in the enterprise. In that context, the strategic importance of China and other Asia Pacific (APAC) markets is significant. For many enterprises, the region is a fast-growing target market, a strategic hub for partners and suppliers, and a key corporate location leveraging the skilled local talent pool. In the case of China, enterprise network connectivity presents local challenges, ranging from availability and quality of Internet connectivity to providing proof of compliance with local regulations. The Internet infrastructure in China is characterized by several potential choke points that routinely lead to high latency and packet loss.



TIP

A fully managed cloud-first WAN can provide multi-cloud connectivity for optimal coverage in China and throughout the region. Look for the following capabilities and benefits in a fully managed cloud-first WAN:

- » **Broad coverage:** Optimal coverage of key locations including state-of-the-art PoPs and a constantly expanding scale and presence
- » **Low latency:** Less than 30ms of first-mile latency to customers, partners, suppliers, and employees in the China region

- » **Optimized connectivity:** Last-mile services that include contracting and monitoring for optimal last-mile connectivity in the China region
- » **Simplified compliance:** Experience balancing complex compliance and local regulatory requirements without degrading performance
- » **Remote worker access:** Optimized onboarding to the SD-WAN backbone for performance and consistent security and the ability to leverage out-of-region capabilities in times of network congestion

## LOGISTICS PROVIDER SOLVES CONNECTIVITY ISSUES IN CHINA

A global logistics provider headquartered in Lisbon, Portugal with offices in 23 countries across the world has grown from Europe into Africa, Latin America, North America, and Asia since its founding nearly 20 years ago.

### Challenge

The company frequently works with many exporters in China. For every export process, its Shanghai-based team opens a new order in an enterprise resource planning (ERP) system located in a data center in Portugal, to be sent to one of many destination offices around the globe. For each of these orders, the Shanghai team often needs to attach 20 to 30 pieces of critical documentation.

Because of the Great Firewall of China (GFW), however, the company's Shanghai team often lost connectivity. During downtime, which totaled around 100 days per year, the team wouldn't be able to access critical applications including the ERP system, Microsoft Exchange, Microsoft SharePoint, and Microsoft Dynamics 365. This held up business to the tune of an estimated €13,000 per month in lost productivity.

These connectivity issues were compounded by the challenge of having a lean IT team of three based in headquarters in a different time zone.

*(continued)*

(continued)

### Solution


After deploying Aryaka Secure Remote Access through its key IT business partner Cloud365 (cloud365.pt), the company's China connectivity issues were resolved immediately. The Shanghai office now experiences 100 percent availability. No downtime, as depicted in the figures "Latency and Downtime Management" and "Link Management for Zero Loss" means no missed revenue opportunities.

Additionally, the company's IT team is able to leverage the 24/7 support team as part of Aryaka's managed services to solve any potential challenges before they become networking issues.


### Results

- The company's offices in China now experience 100 percent network availability.
- The company has saved an estimated €156,000 annually.

#### Latency and Downtime Management



#### Link Management for Zero Loss





## IN THIS CHAPTER

- » Leveraging WAN-as-a-Service and providing predictable connectivity
- » Optimizing WAN performance and enabling multi-cloud networking
- » Enhancing security and taking advantage of automation and orchestration
- » Being proactive with predictive analytics
- » Simplifying management, visibility, and troubleshooting
- » Managing the last mile and accessing a global point-of-presence (PoP) architecture

# Chapter 5

## Deploying a Cloud-First WAN — Ten Capabilities and Benefits

**H**ere are ten important capabilities and benefits to look for in a fully managed cloud-first WAN.

### Delivering WAN-as-a-Service

The foundation of the cloud-first WAN is that it is a service, consumed rather than constructed, OpEx instead of CapEx, and aligned with the cloud consumption model. A well-integrated service will deliver most, if not all, of the capabilities and advantages outlined here.

WAN-as-a-Service cuts through the complexity of understanding the connectivity options such as multiprotocol label switching (MPLS) and direct Internet access (DIA), as well as the task of provisioning and troubleshooting local Internet service provider (ISP) connections. It also hides the mix of optimization, security, and cloud connectivity components, but still lends itself to the visibility and transparency required by IT.



TIP

The benefits of WAN-as-a-Service include:

- » Flexibility enabled by moving beyond legacy architectures and artificial service barriers
- » Simplicity by leveraging a consumption-based approach
- » Velocity with the ability to adapt to fast-changing business needs

## Predictable Connectivity Anywhere

Enterprises expect predictable end-to-end connectivity, so service-level agreements (SLAs) can't stop at a regional border. Applications like Office 365 and Salesforce are more challenging because they utilize traditional connectivity methods for accessing Software-as-a-Service (SaaS) applications and depend on the public Internet, which can be unreliable and slow.

Although WAN optimization techniques can mitigate some public Internet deficiencies, they can't overcome the inherent limitations of ISP peering and congestion. The path forward is a dedicated global and regional backbone, supplemented by reliable last-mile ISP links connecting software-defined wide-area networking (SD-WAN) edge appliances to the nearest points of presence (PoPs).



TIP

The benefits of predictable connective anywhere include:

- » Predictable application performance leading to greater employee productivity
- » Reduced troubleshooting because connectivity is no longer a variable
- » No indirection and lack of visibility between an underlay and overlay

# Built-In WAN Optimization

WAN optimization is a critical feature in a fully managed cloud-first WAN to ensure application performance. WAN optimization techniques should include innovative techniques such as multi-segment optimization and data deduplication, along with other standard techniques like compression, bandwidth management (such as quality of service [QoS] and prioritization), and Secure Sockets Layer (SSL) acceleration.



TIP

The benefits of WAN optimization include:

- » Optimized network and application performance
- » Most efficient use of WAN bandwidth: for example, using compression and deduplication to reduce total cost of ownership (TCO)
- » Helps ensure SLAs are met

# Multi-Cloud Networking

Multi-cloud connectivity is never an afterthought for the cloud-first WAN. It offers direct, regionally based connectivity to the most popular public cloud platforms, leveraging high-speed access technologies such as Direct Connect for Amazon Web Services (AWS) and ExpressRoute for Microsoft Azure.

For SaaS applications like Office 365 and Salesforce, application traffic should be transported over a private network core to the PoP that is closest to the SaaS provider.



TIP

The benefits of multi-cloud networking include:

- » Optimal application performance.
- » Simplified multi-cloud deployment removes the operational burdens from the enterprise.

# Security and SASE

As the threat perimeter grows with SD-WAN and multiple cloud deployments, security can't be an afterthought. The cloud-first WAN embraces flexibility with the choice of security vendor(s) as well as where to deploy — whether it's at the edge, in the cloud, or both. This also includes securing remote access, a superset of secure access service edge (SASE) functionality.

In all but the simplest of deployments, the enterprise will be working with a trusted security vendor. The SD-WAN service must interwork with this vendor and, if desired by the enterprise, should also provide for management of both physical and virtual security appliances.



TIP

The benefits of security and SASE include:

- » Supports the best security solution at all points in the network to minimize the threat of breach
- » No lift-and-shift of existing security vendors

# Automation and Orchestration

In the modern cloud era in which compute, storage, and other cloud resources can be provisioned on demand in minutes, it still takes weeks or months to provision new WAN circuits from a telco carrier or service provider.

A fully managed cloud-first WAN offers cloud-based network provisioning on par with other cloud services. This feature enables the enterprise WAN to keep up with the speed of business rather than becoming a bottleneck to innovation and digital transformation.



TIP

The benefits of automation and orchestration include:

- » Quick turn-up of new locations and services, offering flexibility to the enterprise
- » Optimized allocation of regional and global SD-WAN resources
- » Facilitates troubleshooting and enterprise visibility

# Predictive Analytics

In the same way that navigation applications warn you of upcoming traffic on the highway, the cloud-first WAN integrates predictive analytics to help IT navigate around potential outages. The central orchestration and monitoring system maintains a real-time view into the WAN, leveraging data analytics to predict and determine any issues.



TIP

The benefits of predictive analytics include:

- » Less troubleshooting for enterprise IT troubleshooting
- » Lower TCO
- » Peak performance maintained at all times

# Management, Visibility, and Troubleshooting

Your fully managed cloud-first WAN should offer a powerful, intuitive, web-based management and analytics portal that provides real-time, contextual insight into your network and applications. It should also enable you to perform complete configurations in real time across edge access network locations as well as in the core private network.

The provisioning model for the cloud-first WAN is just like the public cloud with a simplified point-and-click interface that hides the underlying service. Provisioning is the responsibility of the WAN provider.



TIP

The benefits of management, visibility, and troubleshooting in a fully managed cloud-first WAN include:

- » SLA verification and state of the WAN real-time views
- » Control over application performance as if you owned the WAN

# Last-Mile Management and Monitoring

With a fully managed cloud-first WAN, you have no first- or last-mile lock-in, and you no longer have to deal with the complexity of having to understand the capabilities of the ISPs in every country in which you operate. You are free to opt for your preferred service provider rather than the one mandated by a carrier.

The cloud-first WAN provider can handle link subscriptions, monitoring, and troubleshooting, relieving your highly skilled IT team from these more mundane tasks.



TIP

The benefits of last-mile management and monitoring include:

- » It removes the last bit of friction in adopting an end-to-end managed service.
- » You don't need to build ISP expertise within your organization.

## Global PoP Architecture for Service Delivery

Last, but certainly not least, the cloud-first WAN's service richness depends upon an architecture based on distributed service delivery nodes. These nodes, in contrast to traditional SD-WAN PoPs that only support data forwarding, integrate routing, switching, compute, and storage. They are the essential middle-mile component, countering a "hollowed-out" SD-WAN offering that offers no core intelligence.

A key part of the architecture is the handshake between the SD-WAN edge appliance and the node, enabling a host of advanced capabilities across the last-mile. All hardware and services are centrally orchestrated, permitting timely and consistent service enhancements.

The benefits of a global PoP architecture for service delivery include:

- » Service-rich end-to-end architecture, with capabilities propagated from the core outward
- » A consistent edge and node codebase for quick service enhancements

# Glossary

**5G:** 5G stands for 5th Generation Wireless Technology as standardized by ITU in IMT-2020. Large-scale adoption began in 2019. It promises to deliver speeds of 1–2 GB/s as well as optimal support for Internet of Things (IoT) applications.

**AWS:** Amazon Web Services is a subsidiary of Amazon that provides on-demand cloud services and delivers an abstracted technical infrastructure and distributed computing building blocks and tools.

**Alibaba Cloud:** Alibaba Cloud Intelligence is a subsidiary of Alibaba Group and the leading provider of cloud computing services in China.

**CapEx:** Capital expenditure or capital expense (also capex or CAPEX) is the money an organization or corporate entity spends to buy, maintain, or improve its fixed assets, such as CPE for network hardware.

**CASB:** A cloud access security broker sits between cloud service users and cloud applications, monitors all activity, and enforces security policies.

**CPE:** In telecommunications, a customer-premises equipment or customer-provided equipment (CPE) is any terminal and associated equipment located at a subscriber's premises.

**DIA:** Direct (or Dedicated) Internet Access typically represents a broadband service directed to business customers that entails faster response time for support issues, yet still ultimately provides best effort unless other traffic optimization technologies are deployed.

**DIY:** In the context of wide-area networking, do-it-yourself represents an approach where enterprise IT staff take on full ownership for the planning, design, implementation, and ongoing operation of a network. This typically in a focus on day-to-day operations at the cost of neglecting strategic business initiatives.

**FWaaS:** Firewall-as-a-Service is an emerging concept as part of SASE to host more advanced firewall rules in the cloud.

**Google Cloud:** The Google Cloud Platform is a suite of cloud computing services that runs on the same infrastructure that Google uses internally for its end-user products.

**IaaS:** Infrastructure-as-a-Service consists of online services that simplify the consumption of underlying infrastructure like physical computing resources and storage. They support vast numbers of virtual resources along with the ability to scale services up and down according to customers' requirements.

**ISP:** An Internet service provider (ISP) is an organization that provides services for accessing, using, or participating in the Internet.

**L2:** Layer 2 represents the data link layer in the seven-layer OSI model of computer networking. This layer is the protocol layer that directly transfers data between adjacent network nodes on a dedicated physical medium. The most prevalent Layer 2 technology is Ethernet.

**L3:** Layer 3 represents the network layer in the seven-layer OSI model of computer networking. The network layer is responsible for packet forwarding including routing through an intermediate router. The network layer provides the means of transferring variable-length network packets from a source to a destination host via one or more networks, providing an abstraction layer at the cost of potentially increased processing time and deprecated QoS guarantees.

**LTE:** Long-Term Evolution (LTE) is the 4G standard for wireless broadband that precedes 5G. It is sometimes used as an alternative backup connectivity link in case MPLS or DIA links fail.

**Microsoft Azure:** Microsoft is a cloud computing service created by Microsoft and supports many programming languages, tools, and frameworks.

**MPLS:** Multiprotocol label switching is a routing and forwarding technique in telecommunications networks that directs data from one node to the next based on short path labels rather than long IP network addresses. It is sometimes referred to as a Layer 2.5 technology because it often co-exists with IP in carrier class IP core routing infrastructures.

**Multi-cloud:** Multi-cloud is the use of multiple cloud computing and storage services in a single architecture, typically used to reduce dependency on any single public cloud provider.



**NPS:** Net Promoter Score is a management tool that can be used to gauge the loyalty of a firm's customer relationships.

**OpEx:** Operating expense (or expenditure) is an ongoing cost for running a product, business, or system. In technology consumption, it has become increasingly preferable because of elasticity and business agility considerations associated with XaaS (anything-as-a-service) models.

**Oracle Cloud:** Oracle Cloud is a cloud computing service offered by Oracle Corporation providing servers, storage, network, applications and services through a global network of Oracle Corporation-managed data centers.

**PaaS:** Platform-as-a-Service (PaaS) provides a platform allowing customers to compose applications via API calls from the platform.

**QoS:** Quality of service (QoS) in networking refers to several technologies (prioritization, queuing, marking and policing) implemented to guarantee overall latency, jitter, and packet loss.

**SaaS:** Software-as-a-Service is a software licensing and delivery model in which software is licensed on a subscription basis and is centrally hosted in the cloud. SaaS has become a common delivery model for many business applications.

**SASE:** The secure access service edge is an emerging offering combining comprehensive WAN capabilities with comprehensive network security functions. SASE capabilities are delivered predominantly as a cloud-based service.

**SLA:** A service-level agreement (SLA) is a commitment between a service provider and a client to provide a service with deterministic, measurable attributes like — among others — availability, latency, jitter, and packet loss.

**SSL:** Now often referred to as Transport Layer Security (TLS), Secure Sockets Layer (SSL) consists of cryptographic protocols designed to provide secure, encrypted communications security over a computer network.

**SWG:** A secure web gateway is software that restricts or controls the content an Internet user is capable of accessing.

**TCO:** Total cost of ownership is a financial modeling tool intended to help buyers and owners determine the direct and indirect costs of a product or system over time.

**UCaaS:** Unified Communications-as-a-Service provides enterprise communications (voice, video, messaging, conferencing, and so on) as a service. User experience is highly dependent on QoS guarantees over the underlying infrastructure.

**ZTNA:** Zero Trust is an information security framework stating that organizations should not trust any entity inside or outside of their perimeter at any time, effectively discontinuing the premise of universal connectivity the Internet initially enabled.

# aryaka

The Cloud-First WAN Company

## #1 End-to-End Only Managed SD-WAN Provider for the Cloud-First Enterprise

### We Help With



### Why Aryaka?

- ✓ **Leading ROI - OpEx Only Network Consumption Model**
- ✓ **Regional and Global**
- ✓ **Private Core, Internet, MPLS**
- ✓ **The Best of Both Worlds - Leading Technology and Top-Rated Managed Service**
- ✓ **Flexible Edge and Cloud Security - Tier-1 Partners**

### Aryaka SmartServices Platform



Global Management, Orchestration, & Monitoring with 24x7 NOCs


Built-in Network Optimization & Application Acceleration for Maximum Performance

Visibility, Reporting, Co-Management, and Dashboards for the State of the WAN

Integrated Edge and Cloud Security

Direct Multi-Cloud Networking

Secure Global Connectivity with L2 Service PoPs and Last Mile Management



# Transform your enterprise with a cloud-first WAN

A well-designed enterprise wide-area network (WAN) can support successful digital transformation initiatives by following cloud-first principles, thus echoing the flexibility, velocity, and simplicity that enterprises have embraced over the past decade with their public cloud deployments. The cloud consumption model, OpEx instead of CapEx, now has an analogy in the networking space. This book is your introduction to the new WAN technology.

## Inside...

- Rethink traditional WAN architecture
- Compare current WAN approaches
- Imagine a cloud-first WAN — as a service
- Take an integrated approach to network security
- Provide end-to-end visibility
- Automate deployments
- Explore cloud-first WAN use cases

Go to **Dummies.com**<sup>™</sup>  
for videos, step-by-step photos,  
how-to articles, or to shop!

## aryaka

**Lawrence C. Miller** has worked in information technology for more than 25 years. He has written almost 200 For Dummies books.

ISBN: 978-1-119-69972-9

Not For Resale

for  
**dummies**<sup>®</sup>  
A Wiley Brand



# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.