

Brought to you by:



Secure Multicloud Networking

for
dummies[®]
A Wiley Brand

Understand
multicloud use cases



Create a secure
multicloud network fabric



Enable app-to-app
connectivity



Lawrence Miller

F5 Second Special Edition

About F5

F5, Inc. (NASDAQ: FFIV) is the global leader that delivers and secures every app. Backed by three decades of expertise, F5 has built the industry's premier platform—F5 Application Delivery and Security Platform (ADSP)—to deliver and secure every app, every API, anywhere: on-premises, in the cloud, at the edge, and across hybrid, multicloud environments. F5 is committed to innovating and partnering with the world's largest and most advanced organizations to deliver fast, available, and secure digital experiences. Together, we help each other thrive and bring a better digital world to life.



Secure Multicloud Networking

F5 Second Special Edition

by Lawrence Miller

**for
dummies[®]**
A Wiley Brand

Secure Multicloud Networking For Dummies®, F5 Second Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2026 by John Wiley & Sons, Inc., Hoboken, New Jersey. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

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, The Dummies Way, Dummies.com, Making Everything Easier, 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. 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, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-394-38022-0 (pbk); ISBN 978-1-394-38023-7 (ebk); ISBN 978-1-394-38024-4 (ePub)

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Editor: Elizabeth Kuball

Acquisitions Editor: Traci Martin

Senior Managing Editor:
Rev Mingle

Managing Editor:
Sunanda Jayakumar

Client Account Manager:
Jeremith Coward

Production Editor:
Tamilmani Varadharaj

Special Help: Karim El Jamali,
Dave Potter, Kevin Reynolds

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	2
Beyond the Book	2
CHAPTER 1: Recognizing the Need for Secure Multicloud Networking	3
What Is Multicloud?	3
Understanding Multicloud Use Cases	5
Supporting mergers and acquisitions	5
Improving resilience and disaster recovery	5
Avoiding vendor lock-in and optimizing costs	6
Choosing the best provider for your different workloads	6
Looking at Multicloud Challenges	7
Siloed environments	7
Talent/skills shortage	7
Lack of common visibility and operational tools	8
Inconsistent security enforcement	8
Competing priorities across teams	9
CHAPTER 2: Creating a Multicloud Network Fabric	11
Defining Foundational Requirements to Securely Connect Multiple Locations	11
Establishing Secure Public and Private Connectivity via a Private Backbone Network or Directly Between Locations	12
Maintaining End-to-End Segmentation	14
Deploying a Common Network Firewall for North-South and East-West Traffic across Clouds	14
Ensuring Interoperability with Third Parties	15
Leveraging Common Tools for Visibility and Troubleshooting	16
CHAPTER 3: Creating a Secure Service Network	17
Recognizing the Need for Application Networking Needs	17
Providing a Common View of Performance, Visibility, and Security Operations	22

	Ensuring Workload Interoperability	22
	Minimizing Attack Surfaces via Zero Trust	23
	Establishing Unified and Consistent Security Enforcement.....	24
CHAPTER 4:	Ten Tips for Defining Your Secure Multicloud Networking Strategy	25

Introduction

Businesses everywhere are increasingly adopting cloud-based solutions and running more of their critical workloads and business applications in the cloud — or, more often than not, in multiple clouds. Whether multicloud is a deliberate strategy or a happenstance result of various departments procuring different infrastructure-, platform-, or software-as-a-service solutions from different vendors, multicloud has become a de facto standard for enterprise IT environments today.

By using different cloud providers for different workloads and use cases, organizations can optimize their resource usage, reduce costs, and avoid vendor lock-in and other issues that may arise from vendor- or cloud-specific policies and limitations, but will need to contend with challenges around how to manage and secure each environment.

Adopting a multicloud strategy can be transformative for businesses, providing them with the flexibility and agility to thrive in today's dynamic business environment. A secure multicloud network architecture managed with a single, unified platform that spans multiple cloud platforms enables organizations to leverage the strengths of different cloud providers. Such an architecture enables organizations to optimize their workloads and operations, improve performance and agility, reduce costs and latency, and ensure availability and resilience.

About This Book

Secure Multicloud Networking For Dummies, F5 Second Special Edition, consists of four chapters that explore the following:

- » Common secure multicloud use cases and challenges (Chapter 1)
- » How to architect a secure multicloud network fabric (Chapter 2)
- » How to create a secure service network (Chapter 3)
- » Ten important secure multicloud takeaways to keep in mind (Chapter 4)

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.

Foolish Assumptions

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

Mainly, I assume that you work in a technical role. Perhaps you're a chief information officer (CIO); chief technology officer (CTO); vice president/director of IT or infrastructure; cloud architect; or NetOps, DevOps, or SecOps engineer. As such, I assume that you have a strong understanding of networking and cloud computing fundamentals.

If any of these assumptions describes you, then this is the book for you! If none of these assumptions describes you, keep reading anyway. It's a great book and after reading it, your knowledge of multicloud networking won't be cloudy!

Icons Used in This Book

Throughout this book, I occasionally use special icons (I promise, no cutesy emojis) 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.



TECHNICAL
STUFF

This icon explains the jargon beneath the jargon and is the stuff legends — well, legendary nerds — are made of.



TIP

Tips are appreciated but 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 this short book, so if you find yourself at the end of it wondering, "Where can I learn more?," head to www.f5.com/multicloud.

- » Defining multicloud
- » Exploring common multicloud use cases
- » Addressing multicloud challenges

Chapter 1

Recognizing the Need for Secure Multicloud Networking

In today's rapidly evolving cloud landscape, a multicloud strategy can be a real game changer, offering organizations the flexibility and agility they need to compete in the modern digital era. In this chapter, you learn about multicloud: what it is, how it supports different business use cases, and what challenges it introduces.

What Is Multicloud?

Multicloud refers to a cloud computing environment in which an organization leverages various services from more than one cloud service provider (CSP).

Organizations often adopt a multicloud strategy to meet specific business needs and customer demands (we explore some of the more common use cases driving multicloud strategies later in this chapter). However, multicloud also “just happens” in many organizations, whether the result of “shadow IT” — where individual departments or teams procure cloud services from different vendors without involving (or even informing) IT — or

a rapidly growing and evolving IT portfolio of applications and services that includes, for example, Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) resources hosted in different public or private clouds.

Multicloud also increasingly includes on-premises and edge data centers. After all, the cloud isn't just "someone else's data center" — it's a highly efficient, scalable, and agile operating model. As more IT organizations strive to achieve many of the benefits of the cloud operating model in their own data centers, the "cloud" (and multicloud) becomes ever more ubiquitous.



TIP

According to the *F5 2025 State of Application Strategy* report, 94 percent of organizations deploy applications across multiple environments, including on-premises data centers and public/private clouds, and deal with an average of four different public cloud vendors. With multicloud becoming more common, many organizations are considering platform-centric approaches to manage all the complexity that comes with multicloud.

MULTICLOUD VERSUS HYBRID CLOUD

Multicloud and hybrid cloud are both cloud computing models that involve the use of multiple cloud platforms. However, there are key differences between the two models:

- **Hybrid cloud** uses a combination of on-premises or private cloud infrastructure and one or more public cloud providers. The on-premises infrastructure and the public cloud are integrated and work together seamlessly, with workloads moving between them as needed.
- **Multicloud** uses services from multiple cloud providers for different purposes, such as using one provider for application hosting and another for data storage. The workloads are distributed across these providers, and they may or may not be integrated with each other. Multicloud can also include on-premises infrastructure, private clouds, a single public cloud provider, and edge data centers.

Hybrid cloud is a subset of multicloud. All hybrid clouds architectures are multicloud, but not all multicloud architectures are hybrid cloud. Here, multicloud refers to both hybrid cloud and multicloud architectures.

Understanding Multicloud Use Cases

Multicloud enables organizations to take advantage of the strengths of multiple cloud providers and optimize their IT infrastructure and services. Organizations adopt multicloud strategies to support numerous business needs, including mergers and acquisitions (M&A), resilience, and disaster recovery.

Supporting mergers and acquisitions

M&As remain the fastest path to growth for many businesses and enterprises today. IT has always played an important role, ensuring that the various systems, networks, and applications of the different companies in a merger or acquisition are compatible and can be integrated.

In the cloud era, M&As have become more challenging in many respects. One company may still operate only on-premises data centers, while the other company is “cloud everything.” Various departments across the companies may use PaaS and IaaS resources from different cloud providers, for example, to build and host custom applications. In-house application developers write custom cloud-native applications in different programming languages and frameworks, which often dictates their choice of public cloud providers.

There is often a honeymoon period after a merger or acquisition in which the companies involved continue operating largely as they had before. However, at some point (perhaps years later), there is inevitably a drive to consolidate and assimilate. Regardless of where your organization is on the M&A cycle, there is a good chance you’ll start with, become, or morph into a multicloud environment — even after the Borg have integrated you into the collective!

Improving resilience and disaster recovery

Organizations can use multicloud to improve disaster recovery capabilities by distributing their applications and data across different cloud providers and geographic regions. In this way, multicloud can help mitigate the impact of natural or human-made disasters, cyberattacks, outages, and human errors, by

providing geographically dispersed redundancy and failover capabilities. In addition, one public cloud could be used as backup to another cloud as part of a disaster recovery strategy that ensures business continuity.

Multicloud deployments can provide resilience and redundancy by distributing workloads across multiple cloud providers. In case of a cloud provider outage, workloads can be automatically migrated and traffic redirected to other cloud providers, ensuring business continuity and minimal downtime.

Avoiding vendor lock-in and optimizing costs

A multicloud strategy can help organizations avoid being locked into a single cloud provider, service, or solution. Choosing a single cloud provider can limit an organization's options when it comes to working with different features. For example, if an organization wanted to utilize a Prometheus open-source monitoring solution, that solution may be offered by one cloud provider but is not yet available in others. With multicloud, organizations have the flexibility to choose the best cloud provider or solution, on a case-by-case basis, for their unique business requirements, objectives, and goals.

Multicloud deployments can also help organizations optimize their cloud costs by selecting the most cost-effective cloud provider for each workload. Different cloud providers have different pricing models and cost structures. Choosing the right provider for each workload enables organizations to reduce costs and improve their return on investment (ROI).

Choosing the best provider for your different workloads

Multicloud enables organizations to choose the cloud provider that best meets their specific needs in terms of capabilities, features, performance, cost, compliance, and geographic location. For example, an organization can choose a cloud provider that offers better integrations with their existing ecosystem and technology stack or a cloud provider that offers the latest and greatest artificial intelligence (AI) models.

Different cloud providers offer different services and features, and using multiple providers allows businesses to take advantage

of the best, or the most affordable, offerings from each provider. By using different cloud providers for different workloads, organizations can optimize their resource usage and reduce costs.



REMEMBER

Multicloud deployments can help organizations address their data sovereignty and compliance requirements by selecting cloud providers that offer data residency in specific regions and comply with applicable laws and regulations.

Looking at Multicloud Challenges

Of course, multicloud isn't all rainbows and unicorns. A multicloud strategy can introduce several complex challenges that must be addressed, including managing siloed environments, talent/skills shortages, lack of common visibility and operational tools, inconsistent security enforcement, and competing priorities for different personas and uses.

Siloed environments

Siloed multicloud environments can lead to costly inefficiencies because organizations may have to deploy redundant core network services in different clouds. In addition to paying multiple cloud providers for the same services and accounting for any hidden expenses (like egress costs), you have to manage the services across all your different providers. This typically means learning multiple management consoles and a lot of manual work to ensure consistency.

Siloed multicloud challenges typically arise from unplanned or ad hoc multicloud environments. For organizations that have a well-defined multicloud strategy from the outset, siloed environments can be less problematic. Cloud and network architects can often design, select, and/or integrate various cloud services and solutions to ensure maximum compatibility, interoperability, and portability across different cloud providers.

Talent/skills shortage

Cloud computing, generative AI, machine learning, data, and cybersecurity skills are listed among the most valuable technical skills for the foreseeable future. According to *Forbes*, the global

market for cloud computing is expected to grow from \$570 billion in 2022 to close to \$2.5 trillion by 2030.

Multicloud exacerbates this talent/skills shortage. Managing and securing applications and networks across multiple clouds can be complex and often requires specialized skills and tools. It's not just about learning the different names for the various solutions: ExpressRoute in Microsoft Azure versus Direct Connect in Amazon Web Services (AWS). Many of these competing solutions are fundamentally different and have unique capabilities, features, pricing models, and management tools that can quickly overwhelm cloud architects, engineers, and admins. These differences can lead to increased operational costs, integration challenges, and potential performance issues.



TIP

If you have someone on your team who has deep knowledge and experience with *both* AWS and Azure, you have a unicorn worth more than that proverbial “pot of gold” at the end of the rainbow. Keep that invaluable asset happy with proper care and feeding!

Lack of common visibility and operational tools

Differences in controls, logs, and tooling across various cloud providers limit visibility and create complexity in multicloud environments, potentially delaying incident response and troubleshooting. Different cloud providers also use different APIs, data formats, and protocols, which can make it difficult to integrate and manage applications, workloads, and data across multiple clouds.

Inconsistent security enforcement

Multicloud environments can increase the risk of security breaches, data loss, or regulatory noncompliance, because organizations may have to manage multiple sets of security policies, access controls, and data protection measures without any kind of centralized control. Ensuring uniform, consistent, and effective security visibility and enforcement across multicloud environments is imperative to prevent threat actors from taking advantage of this complexity to exploit unknown vulnerabilities or unintended traffic flows that may leave one or more cloud environments exposed.

Competing priorities across teams

A successful secure multicloud networking strategy requires a holistic and collaborative approach among multiple stakeholders across the organization, including IT and business teams. For example, different IT teams have different priorities, including the following:

- » **Network operations (NetOps)** teams need to make sure the network is reliable and performant. Here are some of the specific challenges faced by NetOps teams:
 - **Incompatibilities and delays:** Differences between traditional enterprise networking equipment and public cloud networking constructs often cause delays. Cloud networking constructs have a complex mix of features and limitations, which can make them difficult to deploy and manage.
 - **Network performance and uptime:** NetOps teams need to ensure reliable, highly performant network connectivity across public/private cloud environments. When application performance falters, the network is often the first to be blamed — especially in distributed architectures.
 - **Complexity:** Networking between regions in a single cloud and across multiple cloud providers is complex. Each provider offers disparate toolsets that don't easily work together. Network teams are often forced to cobble together solutions from traditional network vendors that aren't designed for the cloud.
- » **Security operations (SecOps)** teams need to secure environments with consistent policies and identify and address security vulnerabilities quickly. Here are some of the specific challenges faced by SecOps teams:
 - **Increased risk:** Inconsistent policies, configurations, and operational models can result in increased security risks across multicloud environments.
 - **Limited visibility:** Native security tooling in different cloud platforms often doesn't provide end-to-end visibility. This can cause difficulty in correlating security events and responding to security incidents.

» **Development operations (DevOps)** teams need to create and deliver apps quickly to market. Here are some of the specific challenges faced by DevOps teams:

- **Business agility:** DevOps teams often need additional toolsets to ensure application performance across distributed cloud environments.
- **Cloud diversity and reliability:** Organizations that can't effectively manage and resource the deployment of applications across multiple clouds are beholden to a single cloud provider.



REMEMBER

A secure multicloud networking strategy can provide organizations with numerous benefits, including increased flexibility and agility, but it also requires a holistic, platform-centric approach across teams to address its security and management challenges.

IN THIS CHAPTER

- » Starting with the basic requirements
- » Looking at different options for multicloud connectivity
- » Ensuring end-to-end segmentation across clouds and networks
- » Protecting network traffic between and within clouds
- » Ensuring interoperability across different connections
- » Using common tools across teams for visibility and troubleshooting

Chapter 2

Creating a Multicloud Network Fabric

In this chapter, you learn about what it takes to create a multicloud network fabric, including foundational requirements, connectivity options, network segmentation, firewall coverage, interoperability challenges, and the need for common tools across network, security, operations, and development teams.

Defining Foundational Requirements to Securely Connect Multiple Locations

A multicloud strategy enables modern businesses to reduce their dependence on a single provider, enhance resilience, and gain greater flexibility in managing their applications and workloads.

However, organizations must first address several key requirements when connecting a multicloud environment:

- » **Core services**, such as networking, segmentation, service insertion, and traffic steering
- » **Advanced security services**, such as bot protection, distributed denial-of-service (DDoS) protection, application programming interface (API) security, and web application firewall (WAF)
- » **Centralized management**, including dashboards, reporting, and logging to assist with governance and troubleshooting across otherwise disparate and loosely connected environments

By leveraging multiple cloud providers, as well as on-premises data centers in hybrid environments, you can avoid being restricted by specific public cloud provider limitations and select the best services and capabilities for your specific workloads. It enables you to choose the best services and features to enhance digital experiences, while improving business continuity and operational stability and safety.

Establishing Secure Public and Private Connectivity via a Private Backbone Network or Directly Between Locations

Connecting an enterprise network to a cloud environment is a relatively straightforward endeavor. You can simply use one or more of the following methods to establish secure and reliable connectivity:

- » Establish a virtual private network (VPN) connection
- » Build a private connection directly to the cloud using native solutions from the different cloud providers, such as Amazon Web Services (AWS) Direct Connect, Microsoft Azure ExpressRoute, or Google Cloud Dedicated Interconnect

It isn't unusual for an organization to have multiple accounts or subscriptions within its cloud tenant. For example, organizations

commonly have separate accounts or subscriptions for their development (and possibly, test, quality assurance, and staging) and production environments. Additionally, you may have a shared services cloud environment for crucial network and IT operations — such as directory services and Domain Name System (DNS) — that need to connect to your different cloud environments without introducing unintended and potentially vulnerable communication paths between these separate environments.

You may also deploy to different cloud regions to achieve resilience, to comply with data residency requirements, to improve application performance, or for disaster recovery. Each of these environments adds more complexity and challenges to the network connectivity conundrum. Although the complexity increases significantly, the challenge is not insurmountable.

Things quickly get more interesting — and complex — in a multicloud environment. In a single-cloud environment, you can use a cloud provider's backbone network to achieve much of the connectivity required within your tenant. However, connecting two different cloud provider environments — for example, AWS and Azure — using a combination of the different connectivity methods described earlier is not an easy undertaking. Many more challenges must be addressed, such as the following:

- » **Operational scalability:** Each additional connection adds overhead, particularly in meshed environments.
- » **Internet Protocol (IP) address management:** Mergers and acquisitions (M&As) can lead to overlapping IP ranges. Application clusters operating within and across different cloud environments may also reuse existing address space, causing conflicts that can be difficult to troubleshoot and resolve.
- » **Security complexity:** Network sprawl can create blind spots for security tools and make it hard to implement effective security governance, as a lack of centralized management across management planes increases the chances for misconfigurations.

A multicloud network fabric enables organizations to solve complex connectivity challenges in multicloud and hybrid environments, including on-premises data centers, edge locations, and branch locations, by abstracting Layer 3 (Network) functionality in the various public cloud providers' native networking solutions.

Maintaining End-to-End Segmentation

Network segmentation is an important element of a robust network architecture. It enables traffic within the network to be optimized and access to certain network segments to be restricted, among other things.

Most organizations typically have different environments, such as development and production, that they need to keep isolated (that is, segmented) from each other, but might want to allow both environments to access some shared services, such as directory services or DNS, which may be hosted in different environments. For example, your development and production environments may be hosted in separate AWS accounts with no connectivity between them, but both environments need to access AWS Route 53 (DNS) services in AWS, and Azure Active Directory services in Azure.

A multicloud network fabric enables organizations to segment their different network and cloud environments, as needed, to address their unique business and technical requirements.

Deploying a Common Network Firewall for North–South and East–West Traffic across Clouds

Traditional network firewalls were designed to be deployed at the perimeter between a “trusted” enterprise network and the “untrusted” internet, typically in an on-premises data center. In this perimeter-based model, all traffic behind the firewall (that is, the enterprise network) was implicitly trusted and did not pass through a firewall for inspection and policy enforcement. Any traffic that traversed the firewall — for example, between a client workstation and a web server on the internet — was considered *north–south traffic*.

Applications are also sensitive to *east–west network* latency, which is multiplied by the number of request/response pairs or *app turns*. If there are 1,000 API calls in a core loop, each millisecond of latency becomes a full second of application delay as

each loop iteration waits for the response from the remote service. To avoid these issues, it is crucial to have visibility into and control of the path of east–west connections between different regions or clouds.

Thus, while application and workload traffic traversing multicloud environments may technically be considered north–south traffic, it has the same characteristics and high-performance requirements as traditional east–west traffic. Network, security, and cloud architects must work together to design an environment in which firewalls inspect and enforce security for both types of traffic, without introducing additional latency or complexity.

Ensuring Interoperability with Third Parties

Different cloud providers offer their own native solutions to address customer needs. These solutions are designed to provide specific functionality for different use cases within each cloud environment. Customers can choose individual solutions from each cloud provider to address their unique requirements. This flexibility means customers can architect best-of-breed solutions in the environment(s) that best suit their needs.

But with great flexibility comes great . . . interoperability challenges. Cloud providers try to differentiate their cloud offerings with unique capabilities to address a myriad of customer challenges and requirements. These native cloud solutions are not usually designed for seamless interoperability with offerings from competitors. These interoperability challenges extend to on-premises environments as well, where a customer's existing investments in network equipment and tools are deployed.

A multicloud network fabric can abstract both Layer 3 (Network) and Layer 7 (Application) components to help customers achieve seamless interoperability for their applications and workloads across multicloud environments.



TIP

Ensure your multicloud network fabric leverages standard protocols to ensure seamless interoperability across multicloud and hybrid environments.

Leveraging Common Tools for Visibility and Troubleshooting

Poor application and network performance negatively impacts the customer experience. In essence, a slow application is a down application. Keeping your applications running fast and reliably is a minimum requirement for doing business in today's digital world. Customers have many options for where they deploy applications, and fast and reliable application performance is often the difference between onboarding a new customer and losing business.

Collaboration across network, security, operations, and development teams is challenging enough, given that each team may have different priorities and areas of focus. Without common tools, centralized management, or complete observability to ensure these different teams have a common view when an issue arises, everyone may as well be speaking different languages. Unfortunately, this can lead to costly delays in troubleshooting and resolving issues, as well as friction across teams pointing fingers at each other. Here again, the challenge is exacerbated in a multicloud environment in which different cloud providers offer different tools and solutions for various teams and purposes.



TIP

To successfully address these challenges, everyone needs to be using a common set of tools in a multicloud network fabric — including NetFlow information, packet captures (PCAPs), and requests initiated from gateways such as Internet Control Message Protocol (ICMP) pings, DNS, and Transmission Control Protocol (TCP) connectivity checks — to ensure full, end-to-end observability and troubleshooting capabilities across clouds and on-premises environments, including Layer 3 (Network) and Layer 7 (Application) performance and security metrics, discovery and mapping of any application running anywhere across your environments, and visibility into how and which APIs are being consumed in workflows. To manage this complexity, organizations are increasingly opting for platforms that provide all these tools integrated into a single package.

- » Focusing on velocity and agility
- » Providing a common view for DevOps, NetOps, and SecOps
- » Ensuring workload interoperability across clouds
- » Taking a Zero Trust approach
- » Enabling consistent security enforcement

Chapter 3

Creating a Secure Service Network

In this chapter, you learn how a service network enables seamless collaboration across DevOps, NetOps, and SecOps teams, while enabling each team to achieve its individual goals in support of business agility and velocity.

Recognizing the Need for Application Networking Needs

For modern enterprises, business agility is key to success in a highly competitive economy, and applications are the business. Thus, accelerating time to market for innovative new applications and application updates is key to business agility and success.

As organizations increasingly adopt multicloud strategies, they need solutions that allow them to rapidly and easily interconnect applications and services hosted in different clouds (not just cross-connect virtual networks) and distributed microservice clusters. They need to connect applications to specialized

services, third-party identity and access management (IAM) providers, and in-house services hosted in different clouds.

Traditional Layer 3 networks cannot adequately scale to meet the connectivity needs of modern microservices-based application architectures that are distributed across multicloud and edge environments with internal app-to-app communication requirements.



REMEMBER

According to the *F5 2025 State of Application Strategy* report, only 6 percent of organizations deploy their applications in a single environment. Some 94 percent deploy applications across multiple environments, including on-premises data centers and public/private clouds, and deal with an average of four different public cloud vendors.

Unfortunately, network and security teams often become the bottleneck for traditional organizations when it comes to deploying applications to multicloud environments, resulting in a drastic slow-down of application velocity — even though the applications themselves may be ready to go. This ultimately brings business agility to a crawl and puts downward pressure on the business (see Figure 3-1).




FIGURE 3-1: Traditional organizations are subject to what the network can support, slowing application and business velocity.

Let's consider the deployment timeline for a typical application, such as a modern (that is, cloud-native) application running in a public cloud, such as Amazon Web Services (AWS). Various clients, including other applications and users, connect to your application from different environments, including on premises, other public clouds, and the internet (see Figure 3-2).




FIGURE 3-2: Traditional application deployment can take weeks or months.



The Open Systems Interconnection (OSI) reference model defines seven layers that facilitate communication between systems and components: Application (Layer 7), Presentation (Layer 6), Session (Layer 5), Transport (Layer 4), Network (Layer 3), Data Link (Layer 2), and Physical (Layer 1).

After the application has been deployed in AWS, DevOps, NetOps, and SecOps teams work together to bring the application online:

- » **DevOps** will need to configure the Layer 7 (Application) proxy and ingress rules to route the appropriate application programming interface (API) traffic to the application.
- » **NetOps** will need to configure a frontend load balancer and configure network routing to the frontend load balancer. Because the various clients will be connecting from other environments, the NetOps team will also need to configure an appropriate network address translation (NAT) policy, create any required rules on the perimeter firewall, and configure virtual private network (VPN) and inter-environment connectivity.
- » **SecOps** will need to configure web application and API protection (WAAP) policies to protect against application security risks from vulnerability exploits, bots, automated attacks, denial of service, fraud and abuse, and insecure third-party API integrations.



WAAP refers to an integrated set of security services that work together to mitigate security risks from APIs and web applications.

Each of these tasks and functions may be managed by various devices from different vendors, each with its own management console, making collaboration across teams extremely challenging — and extending application deployment timelines by weeks or months.

DevOps, NetOps, and SecOps teams often have different priorities, which creates conflict, friction, and siloes within the organization. Despite being aligned with the business objectives (that is, velocity and agility), they often find themselves at odds with other teams (see Figure 3-3).

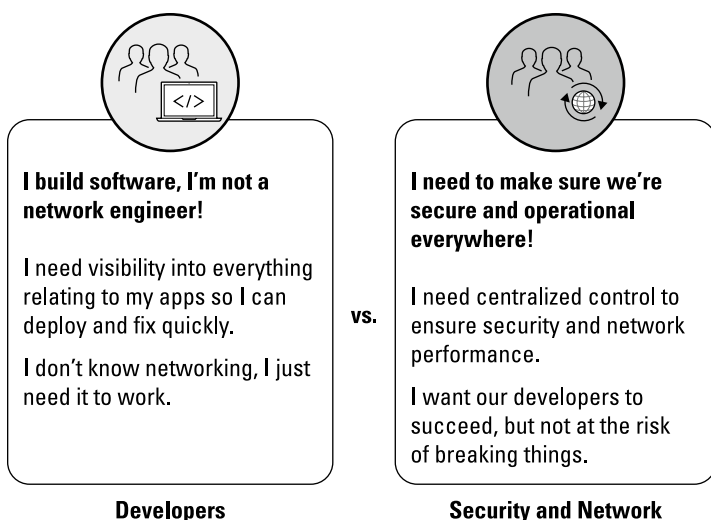


FIGURE 3-3: DevOps, NetOps, and SecOps teams often have different goals, which creates conflict across the different teams.

This same process happens, more or less, for every application deployed in an environment. If you don't have the in-house resources and expertise to do all of this yourself, then configurations may be missed, security policies may not be deployed or enforced properly, and cyber risk will increase for every application in each of your environments.

To increase agility, DevOps teams need self-service capabilities that allow them to deliver applications without worrying about underlying network complexities. NetOps teams need the ability to seamlessly configure and manage a performant, reliable network across all environments. Finally, SecOps teams need the ability to consistently deploy, manage, and enforce security policies for all applications across the environment (see Figure 3-4).

Teams can also leverage AI-powered tools to intelligently configure interconnections, optimize application routing, and predict potential connectivity bottlenecks across multicloud environments, significantly reducing manual effort and deployment timelines.



REMEMBER

The application, and all of its various components, is the center of the universe in this new reality. Service networking or application-level networking provides resiliency, security, and deep-level monitoring and visibility for the application and, in turn, abstracts the networking fabric focus.

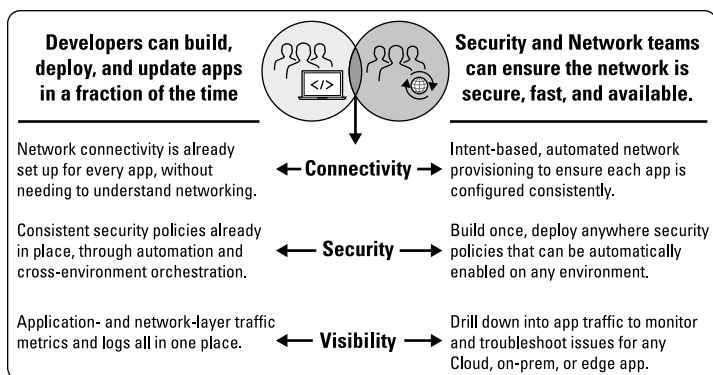


FIGURE 3-4: Speed and agility are paramount for developers, while network and security teams prioritize performance, reliability, and protection.

Providing a Common View of Performance, Visibility, and Security Operations

A secure service network addresses management challenges with a centralized management and analytics console that DevOps, NetOps, and SecOps teams all can use to share a common view and seamlessly collaborate with each other.

They can use the platform to plug in their automation frameworks, continuous integration/continuous delivery (CI/CD) pipelines, alerting, and security information and event management (SIEM) systems.

Ensuring Workload Interoperability

Ensuring workload interoperability across multicloud environments can be a challenge in and of itself. Various application components and services may be hosted in different clouds and on different technology stacks, which increases complexity and exacerbates this challenge.

Regardless of the locations or form factors on which application workloads and components are hosted — whether a virtual

machine (VM), Kubernetes service, or serverless function — a service fabric solution should provide seamless interoperability, deep-level visibility, and end-to-end monitoring capabilities.

AI can also act as a unifying layer, intelligently mapping dependencies between components, optimizing resource allocation across environments, and ensuring seamless interoperability despite differing technology stacks and architectures.

Minimizing Attack Surfaces via Zero Trust

A zero trust architecture replaces the traditional perimeter-based network model — in which a network firewall sits between the “untrusted” (that is, internet) and “trusted” (that is, enterprise) networks — with a “Never trust, always verify” security model.



TIP

The Zero Trust approach is primarily focused on protecting data and services, but it can and should be expanded to include all enterprise assets (devices, infrastructure components, applications, virtual and cloud components) and subjects (end users, applications, and other nonhuman entities that request information from resources).

AI can enhance Zero Trust architectures by leveraging advanced machine learning algorithms to monitor real-time data, detect unusual activity, and predict coordinated threat campaigns across vectors such as APIs, bots, and distributed denial-of-service (DDoS) attacks. This level of proactive intelligence helps organizations stay ahead of evolving security risks.

A secure service network enforces Zero Trust by leveraging a WAAP platform with integrated security controls, which provides organizations with deeper visibility and actionable insights on their security posture that can help stop specific attacks — such as bots and distributed denial-of-service (DDoS) — and identify coordinated threat campaigns that span multiple threat vectors, including API security.

Establishing Unified and Consistent Security Enforcement

When organizations have applications running in multicloud and hybrid environments, they often have to manage different tools, such as DDoS protection and web application firewalls (WAFs), for each environment — each with its own unique capabilities, features, and management consoles. For example, an organization with a relatively simple environment consisting of DDoS protection and WAFs in AWS, Azure, and on-premises data centers may have to manage six different consoles, quickly creating a complex burgeoning management nightmare challenge.

Managing multicloud and hybrid environments with a single, unified platform enables granular security policies to be centrally defined and consistently enforced across different clouds and technology stacks. Without a secure service network with centralized management, security teams must learn and configure the native security tools offered by each cloud provider. Here again, the cloud providers' primary focus is functionality, not interoperability with other clouds. Even when organizations procure third-party tools offered within a cloud provider's environment, these tools often function differently across different clouds.

AI-powered solutions simplify the enforcement of security policies by dynamically adapting rules based on real-time threat intelligence, enabling consistent protection across multicloud and hybrid environments without manual intervention.

Finally, a unified monitoring capability, at a per-application or per-service level, enables security teams to visualize any type of security event or incident, and network teams to view performance information, such as how each connection request gets processed in terms of latency, errors, and so on. Doing all of this through a single platform is now the de facto choice for many organizations.

IN THIS CHAPTER

- » Maximizing flexibility and reducing complexity
- » Ensuring security and leveraging ecosystem integrations
- » Adopting new and emerging technologies like AI
- » Enabling single-pane-of-glass management for multicloud

Chapter 4

Ten Tips for Defining Your Secure Multicloud Networking Strategy

Here are ten tips and best practices to help you define (or refine) your secure multicloud networking strategy:

» **Select the right cloud providers for your business needs.**

Choose cloud providers that can meet your business goals and requirements, and consider their services, features, pricing, support, reputation, and ecosystem. Also, ensure interoperability and compatibility between your different cloud providers.

» **Build “best of breed” solutions and increase resilience.**

Multicloud enables organizations to leverage best-in-class services from each cloud provider while mitigating the risk of downtime from a catastrophic outage at any single cloud provider.

- » **Take control of your multicloud strategy.** Whether your multicloud deployment is the result of a carefully planned multicloud strategy or a more ad hoc shadow IT environment, you need to ensure that it meets your business needs, including agility, compliance flexibility, governance, performance, and application security requirements.
- » **Define strategies to manage and secure applications and workloads across multiple clouds.** Leverage capabilities and solutions such as Web App and API Protection, Layer 7 (application) gateways, backup and recovery, container orchestration frameworks, geo-proximity, and load balancers as part of a holistic platform to manage multicloud applications and workloads. Enhance these capabilities with AI-powered tools, which can dynamically optimize policies to respond to emerging threats.
- » **Enable workload portability with containerization and virtualization technologies.** Workload portability helps organizations easily move workloads across cloud providers or bring workloads back to on-premises data centers as needed. This portability is often achieved through containerization or virtualization technologies, which can enable workloads to run consistently across different cloud platforms. This gives organizations more flexibility in terms of where they deploy their applications and workloads.
- » **Simplify multicloud network transit.** With multicloud network transit, data centers route through the same routers as your cloud application flows, allowing you to see each cloud provider as a metered resource for app consumption. No need to worry about addressing, Domain Name System (DNS) resolution, or routing for each environment.
- » **Ensure consistent application security and compliance.** Use a centralized security and compliance framework and AI-powered tools such as identity and access management (IAM), encryption, threat detection and response, and compliance monitoring to ensure consistent security and compliance across multiple clouds.
- » **Don't forget ecosystem integrations.** By leveraging advanced application delivery and security capabilities, Layer 7 gateways, and automation tools for software development and deployment, businesses can optimize

their multicloud strategy for cost efficiency, improved resilience, and greater agility. Consider platforms that offer these integrations alongside AI-powered tools to make managing your multicloud environment easier.

- » **Leverage new and emerging trends.** Businesses can leverage serverless architectures across multiple cloud providers, taking advantage of the best services and capabilities from each provider without managing infrastructure. They can also deploy edge computing solutions across multiple cloud providers in different geographic locations, allowing them to optimize performance and reduce latency in their applications. AI-driven solutions can further optimize serverless architectures by dynamically allocating resources and predicting scaling needs. In edge computing, AI can analyze data locally to improve decision-making, enhance performance, and reduce latency without relying on centralized cloud processing.
- » **Use a cloud management platform.** Use a cloud management platform that provides a single, unified view of your multicloud environment. This platform should support provisioning, monitoring, billing, and governance across multiple clouds, as well as provide end-to-end observability of your security posture, and provide AI-powered tools to enable stronger security.



Simplify hybrid multicloud networking

Connect and protect apps across public cloud, edge, and on-premises environments with the F5® Application Delivery and Security Platform, a SaaS-based solution that delivers full-stack multicloud networking to help you:

- **Reduce the number of tools** needed to connect, extend, and secure networking between clouds and data centers
- **Centralize observability** to monitor health and speed troubleshooting across environments
- **Accelerate time-to-service** for new app deployments with unified management
- **Enforce consistent security policies** to minimize downtime and defend against threats

Hybrid multicloud networking can be complex. Managing it doesn't have to be.

f5.com/multicloud



Take back control of your multicloud environment!

Deciding which environment to host your applications is a choice often driven by what's best for the application, even when different applications require different cloud services. A multicloud strategy can help manage and accommodate different applications and workloads, while also increasing business agility, reducing costs, and optimizing IT operations. But multicloud also increases complexity and introduces new challenges. This guide helps you discover how to you can take back control of your multicloud environment with a platform approach to building a multicloud network fabric and service network.

Inside...

- Recognize multicloud benefits
- Address multicloud challenges
- Enable end-to-end observability across clouds
- Secure every environment with Zero Trust architecture
- Leverage common tools
- Ensure workload interoperability



Lawrence Miller served as a Chief Petty Officer in the U.S. Navy and has worked in information technology in various industries for more than 25 years. He is the coauthor of *CISSP For Dummies* and has written more than 250 *For Dummies* books on numerous technology and security topics.

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

ISBN: 978-1-394-38022-0

Not For Resale

for
dummies®
A Wiley Brand



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.