

LEARNING MADE EASY

Extreme Networks Special Edition

Fabric Networking

for
dummies[®]
A Wiley Brand



Scale your network
to meet demand

—
Increase availability
and fault tolerance

—
Improve security
with segmentation

Brought to you
by



Extreme[®]

Sara Perrott

About Extreme Networks

Extreme Networks delivers software-driven solutions from the enterprise edge to the cloud that are agile, adaptive, and secure to enable digital transformation. The company's 100 percent in-sourced services and support are number one in the industry. Even with 30,000 customers globally, including half of the Fortune 50 and some of the world's leading names in business, hospitality, retail, transportation, logistics, education, government, healthcare, and manufacturing, Extreme remains nimble and responsive to ensure customer and partner success. Extreme calls this Customer-Driven Networking™. Founded in 1996, Extreme is headquartered in San Jose, California. For more information, visit www.extremenetworks.com or call 1-888-257-3000.



Fabric Networking

Extreme Networks Special Edition

by Sara Perrott

for
dummies[®]
A Wiley Brand

Fabric Networking For Dummies[®], Extreme Networks Special Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

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. Extreme Networks and the Extreme logo are trademarks or registered trademarks of Extreme 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, 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 978-1-119-61937-6 (pbk); ISBN 978-1-119-61940-6 (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. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz or visit www.wiley.com/go/custompub. For details on licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

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

Project Editor: Martin V. Minner

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development Representative:

Karen Hattan

Production Editor: Tamilmani Varadharaj

Table of Contents

INTRODUCTION	1
About This Book	1
Icons Used in This Book.....	2
CHAPTER 1: Recognizing the Need for More Flexible and Automated Networking	3
Analyzing the State of Network Technology	3
Learning About the Types of Fabric Networks	5
Identifying the Benefits of Fabric Networking	6
Understanding Adjacent Networking Technologies	7
CHAPTER 2: Exploring the Use Cases for Fabric Networking	9
Understanding Data Center Use Cases for Fabric Networking	9
Checking Out Fabric Networking in the Campus	12
CHAPTER 3: Exploring Fabric Networks	15
Introducing Shortest Path Bridging.....	15
Understanding How SPB Works	17
Building the SPB fabric	17
Deploying SPB services	17
Examining the Security Attributes of SPB.....	20
Looking at Complementary Technologies.....	21
Dynamic Auto-Attach.....	21
Seamless extension of SPB.....	22
Designing multi-site networks with mobile endpoints.....	22
Seeing Where SPB Fits in the Network	22
Introducing TRILL Fabrics	23
Seeing Where TRILL Fits in the Network.....	24
Delving into IP Fabrics.....	24
Understanding IP Fabrics	24
Building the IP Fabric	24
Deploying Services	26
Layer 2 VNI.....	26
Layer 3 VNI.....	27
Looking at Complementary Features	27
Seeing Where IP Fabrics Fit in the Network	28

CHAPTER 4:	Examining Key Considerations in Fabric Technology Evaluations	29
	Enabling Automation	30
	Network automation	30
	Cross-domain automation	31
	AI/ML driven automation	32
	Enhancing Visibility.....	33
	On-box visibility.....	34
	Network analytics	34
	Integrating Ecosystem Partners	35
	Improving Security	35
	Integrating Wired/Wireless	36
CHAPTER 5:	Choosing the Right Fabric Solution for Your Business	37
	Focusing on Your Pain Points	37
	Evaluating Fabric Networking Technologies	39
CHAPTER 6:	Ten Things to Know About Fabric Solutions from Extreme Networks	41
	You Have a Choice.....	41
	Single Pane of Glass.....	42
	Visibility from the Edge to Multi-Cloud	42
	Solutions Are Truly Multi-Vendor	42
	Deployment History	42
	Innovation!	43
	Field Proven in a Demanding Environment	43
	Plug and Play.....	43
	Ecosystem Integrations	43
	Unbreakable in Hack-a-thons	44

Introduction

In today's world, businesses grow and change at a pace that can be difficult to keep up with. This is especially true if you are in charge of architecting or operating the communications networks that enable the business to serve its employees and its customers.

Networks of the past were designed for simple, static workflows like connecting an office user to the Internet or to email. Fast forward to today, and work is something you do, not a place that you go to. Today's workplace is dynamic and mobile. It requires a network that can adapt and change on demand.

Although computing technologies have matched the needs of modern organizations, networking technology has not progressed at the same pace. Many organizations still use a legacy network infrastructure that is inadequate for today's ever-changing business.

Fear not, network administrators, the next evolution in networking technology is here! Fabric networking allows you to scale on demand, create redundant paths for high availability and fault tolerance, improve your security with network segmentation, and so much more. You can move from an old legacy network to a modern, scalable, secure network infrastructure that will enable your business to succeed and grow. The network is no longer the bottleneck. It can instead be the enabler for success.

About This Book

Fabric Networking For Dummies, Extreme Networks Edition, is your introduction to the world of fabric networking and the underlying protocols and technologies that support it.

Fabric networking enables you to scale with demand and create highly available and fault-tolerant networks. This book explains what fabric networking is, including its history, use cases, current technologies available to you, and some of the key points to think about when you consider fabric networking.

Icons Used in This Book

As you read this book, notice the icons in the margins. These indicate information that may be of interest. The material that accompanies the icons can enrich your understanding of fabric networking. I highly recommend reading them!

Here's what the icons mean:



TIP

Tips provide guidance that may save you time and effort. These are typically based on real-world experience and are there to help you hit the ground running with fabric networking.



REMEMBER

The Remember icon indicates information that deserves your special attention.

- » Analyzing the state of networking technology
- » Introducing fabric networking
- » Learning about the types of fabric networks
- » Understanding adjacent networking technologies

Chapter **1**

Recognizing the Need for More Flexible and Automated Networking

Technology is advancing at a rapid pace. Robotics, Internet of Things (IoT) devices, virtual reality, and augmented reality are all making their way onto your network. Bound by complexity, traditional networks can no longer keep pace with the speed of innovation. What's needed is a simpler, automated, and secure network environment.

This chapter introduces the challenges faced today with traditional networking technology and how fabric networking can help. It shows you what fabric networking is, how it advances networking, and the significant benefits it provides.

Analyzing the State of Network Technology

As organizations rapidly adopt new technologies, applications, and delivery models, they're rendering traditional network design obsolete. From the influx of a wide variety of IoT and "bring your

own device” (BYOD) devices that require secure connectivity, to the requests for support of more real-time, bandwidth-intensive traffic such as video, to the rapid growth of hybrid cloud delivery models — all of these demands are having a significant impact on the network.

The challenges with traditional networks include:

- » **Outdated designs:** Traditional network design has remained consistent for decades, with many of the common network protocols originating almost 30 years ago. Back then networking environments were static — applications were tied to servers and employees to their desktop computers. With server virtualization, mobility, IoT, and cloud computing more common, networks are being asked to support dynamic, mobile environments, and are therefore falling short.
- » **Manual configuration:** Traditional networking is largely configured manually, switch by switch, through a command-line interface (CLI). This method was adequate in a static networking environment. However, in today’s environment where new devices and applications are moved, added, or changed frequently, manual configuration is time consuming. It also introduces the risk of an outage or even a security breach resulting from human error during a change.
- » **Too much complexity:** Traditional network design requires network administrators to understand and configure many interdependent protocols. These include Spanning Tree Protocol (STP), Open Shortest Path First (OSPF), Protocol-Independent Multicast (PIM), and Border Gateway Protocol (BGP). All this complexity can also slow network recovery because each layer is reliant on the layer below it to re-establish connectivity.
- » **Vulnerability to breaches:** Security concerns are front and center in businesses today. For example, you don’t want guest Wi-Fi devices to communicate over the same network as human resources systems. In legacy networks, this issue has traditionally required multiple virtual local area networks (VLANs) with access controls and firewalls to separate the traffic. However, if you are sharing a routing table, your IP network is flat and if someone breaches your network, they might make their way to sensitive data such as customer payment information or patient health records.

Learning About the Types of Fabric Networks

Fabric networks get their name from the diagram of their component connectivity, which resembles a piece of fabric. The network is woven together into a connectivity mesh.

Today, you're likely to encounter two types of fabric networks:

» **Ethernet fabrics** are based on industry standard protocols such as Shortest Path Bridging (SPB) or Transparent Interconnection of Lots of Links (TRILL) and use Ethernet-switched paths to forward traffic. They also use a link state protocol such as Intermediate System to Intermediate System (IS-IS) or OSPF as their control plane to bring more deterministic, carrier-grade functionality to Ethernet. Ethernet fabrics overcome the issues associated with traditional Ethernet network designs. They allow all links to be active with multiple equal cost paths. They eliminate hop-by-hop provisioning and they enable very large-scale Ethernet networks to be implemented, getting past the limitations of traditional VLAN and Media Access Control (MAC) scaling. Most importantly, they enable much faster reconvergence times, generally achieving sub-second network-wide recovery.

Ethernet fabrics define the type of fabric; however, they do not dictate the type of services or traffic running across it. Most support full Layer 2 and Layer 3 services within the fabric for attached devices. Certain Ethernet fabrics even offer extensions to support integrated virtual routing and forwarding (VRF) and IP Multicast routing capability. This enables them to not only be a replacement for STP, but also other protocols such as OSPF, BGP, PIM, and even Multiprotocol Label Switching (MPLS).

» **IP fabrics** are based on industry standard protocols such as BGP and Ethernet Virtual Private Network (EVPN). Like Ethernet fabrics, IP fabrics use equal cost multi-pathing to improve efficiency. They also support Layer 2 and Layer 3 services for end devices, embrace server virtualization, and provide automation for complete plug-and-play provisioning. The main differences between the Ethernet and IP fabrics are in the control and data planes used to construct the fabric. In most cases, BGP is used for the underlay (the control plane

network), and BGP/EVPN with Virtual Extensible LAN (VXLAN) tunnels are used for the overlay (the data plane).

The driver behind Layer 3 fabrics was the use of a single protocol stack and single virtualization technology. By leveraging BGP as the underlay, the network can deliver massive scalability. These fabrics all offer the benefit of full vendor interoperability for both the underlay and overlay networks, allowing for vendor-agnostic network implementations.



TIP

Although fabric technology is not new (initial deployments started as early as 2009), the capabilities have been greatly expanded upon, giving it renewed relevance when you're designing a next-generation network.



REMEMBER

Because networking standards continue to evolve, it's best to stay up to date with relevant standards-bodies such as the Institute of Electrical and Electronics Engineers (IEEE) and the Internet Engineering Task Force (IETF). In fact, the IETF is currently working on additional IP fabric technologies such as Routing in Fat Trees (RIFT) and link state vector routing (LSVR).

Identifying the Benefits of Fabric Networking

When you consider the advantages of fabric networking, these benefits stand out and make the business case for switching from traditional to fabric networking:

- » **Improved time-to-service:** Likely, you're being asked to deliver more connectivity and enable network changes at an increasing pace. Many of these changes require a maintenance window, thus introducing delays into accommodating the requests. Fabric networking can automate network changes, allowing you to make them in minutes, rather than days or weeks.
- » **No downtime from human error:** Because adding and removing network services in a fabric network only needs to be done at the edge, and is typically done from a centralized management system, mistakes are far less likely and the network is much more stable.

- » **Better resiliency:** Fabric networking provides the opportunity to simplify your network by reducing the number of protocols in use. With a simpler network, recovery times are faster. Also, more interconnectivity means the loss of a link or even a network device will not impact your applications.
- » **Enhanced security:** Fabric networks allow you to easily implement network segmentation with the creation of secure zones. This feature prevents lateral movement across the network. For example, if a system is infected by a virus, the infected system is not allowed to communicate outside of its zone or segment.
- » **Integrated wired and wireless:** With certain fabric technologies you can fully integrate your wired and wireless networks to create a unified campus network. This unification leads to greater simplicity in deploying wireless APs and providing a consistent quality of service across both wired and wireless access.
- » **Enhanced quality of experience (QoE):** Fabric technologies use algorithms to calculate the shortest path between any source and destination. This capability ensures that in any network, users and devices are connected to their applications by the shortest and most efficient path to reduce latency.

Understanding Adjacent Networking Technologies

This section examines important technologies that are synergistic to fabrics:

- » **Software-defined networking (SDN):** SDN is a concept that is loosely defined in the industry. The original approach involved an open protocol (OpenFlow) to program network switches and routers from a centralized controller. Today, the industry has largely moved away from this approach and more toward application programming interfaces (APIs) and automation.

Fabric networking can be deployed independently or as part of an SDN solution. Independently, fabrics can deliver inherent automation capabilities. Alternatively, they can be deployed with a management system or controllers for centralized control and automation.

- » **Network functions virtualization (NFV):** NFV virtualizes network services like routing, switching, load balancing, security, and wide area network (WAN) optimization so they can be deployed on commodity hardware. Because fabric networking is a feature commonly available on routing and switch OSs, it is synergistic to NFV, since the fabric OS can run within a hypervisor or a container-based framework alongside other required network functions like security or WAN optimization.
- » **Multi-Chassis Link Aggregation (MLAG):** MLAG logically aggregates two or more switches to form one logical entity. It provides link-level and device-level network resiliency and eliminates single points of failure. Originally designed to enable an STP-free core network, MLAG can be used in conjunction with fabric networking for a higher degree of resiliency, offering the ability to take nodes out of service for software upgrades and patches without any impact to availability. It can also provide enhanced resiliency for end devices such as servers, firewalls, and load balancers.
- » **Port extender technology:** Port extender technology replaces traditional fully-featured access layer switches with simple port extenders that are fully managed and controlled by an aggregation switch (typically called a *controlling bridge*). Port extender technology can be used in conjunction with fabric networking to provide a high fan-out of ports that are controlled, operated, and managed by a fabric-enabled aggregation switch.

IN THIS CHAPTER

- » Understanding data center use cases for fabric networking
- » Examining campus use cases for fabric networking

Chapter 2

Exploring the Use Cases for Fabric Networking

Digital transformation — and the influx of advanced technology that it brings onto your network — requires networking that easily adapts to changing business needs. A step in the right direction is the use of fabric networking.

In this chapter, you learn about using fabric networking in campus and data center environments.

Understanding Data Center Use Cases for Fabric Networking

Data centers and the technologies housed within them are rapidly evolving as virtualization, hyperconverged infrastructure (HCI), containers, and the cloud are becoming more prevalent. Traditional networks haven't kept pace with the changing ecosystem in the data center. With all the manual processes needed to scale up networking resources, traditional networks tend not to be easily scalable, agile, or flexible.

This section shows how fabric networking can address the challenges of the modern-day data center and enable your IT department to work at cloud speed while becoming more agile and responsive to business needs.

Identifying the need for a data center fabric network

Here are a few of the technologies that benefit greatly from the implementation of fabric networking in the data center:

- » **Server virtualization** has become common and is the best way to ensure that you are utilizing your server hardware to its fullest. One of the initial use cases for fabric networking was to solve the challenge of stretching virtual LANs (VLANs) or subnets within and between data centers so that the IP address of the virtual machine could be maintained regardless of what server it moved to. Today, fabric networking makes server and virtual machine (VM) deployments and migrations much quicker by offering any service — for example, VLAN — on any port within the data center. Automated top-of-rack switch provisioning addresses new VMs being brought up as well as moved within the network with no manual network configurations.
- » **Hyperconverged infrastructure (HCI)** is the convergence of virtualized compute, storage, and networking services onto a standard commercial off-the-shelf (COTS) server. Having a truly integrated solution provides significant value in capital expenditures as well as ongoing operations and maintenance. Fabric networking fits within HCI perfectly, providing an adaptable infrastructure that scales out in conjunction with compute and storage. Certain fabrics also inherently support advanced storage technologies to ensure optimal performance and availability.
- » **Containers** are one of the newest technologies within the data center. They have gained significant popularity because of their ease of use and quick provisioning. Containers, however, do not have long life spans. They must be spun up and shut down quickly as workloads change. The network that supports containers, as well as container hosts, must be able to adapt and respond to changing needs within seconds. Fabric networks enable a more agile infrastructure

through automation. When containers are created on a container host, fabric networks allow you to generate the resources the containers demand to support their workloads.

- » **IP storage** connectivity requires a network that delivers the bandwidth, performance, and reliability needed in today's demanding environment. Fabric networks that support storage technologies such as Internet Small Computer Systems Interface (iSCSI), Network File System (NFS), or Non-Volatile Memory Express (NVMe), either in a dedicated IP storage network or in a hyperconverged environment, allow network administrators to architect a flexible, robust network that scales easily to align with storage expansion needs.
- » **The hybrid cloud** is becoming a common way for businesses to expand their data centers by providing additional capacity for bursting, as well as by hosting applications that can be easily housed off premises. This is often done by extending a virtualized infrastructure into a cloud provider's environment. Certain fabric technologies can extend into a cloud provider's environment to enable a robust, flexible, and scalable network that works seamlessly on and off premises — with the ability to manage both through a single pane of glass.
- » **Data Center Interconnect (DCI)** connects multiple data centers together, as shown in Figure 2-1. With multiple data centers it is critical to have active-active connectivity to ensure business continuity. Ideally, if all data centers have fabric networking implemented, and you have dark fiber interconnecting them, it's simple to extend the fabric between sites. However, in most cases, there is a wide area network (WAN) network to interconnect the data centers. In this scenario, you may have to insert a DCI solution. This allows sites to become one logical data center that has applications and services dispersed between or across locations. This can be done as an overlay to the WAN infrastructure you have in place.



TIP

When you use a DCI, the technologies in play will most likely encapsulate traffic across the WAN in Virtual Extensible Local Area Network (VXLAN), for example. This requires support of a larger maximum transmission unit (MTU) of 1600 bytes, so ensure that your WAN and WAN provider will support this before implementing your DCI.

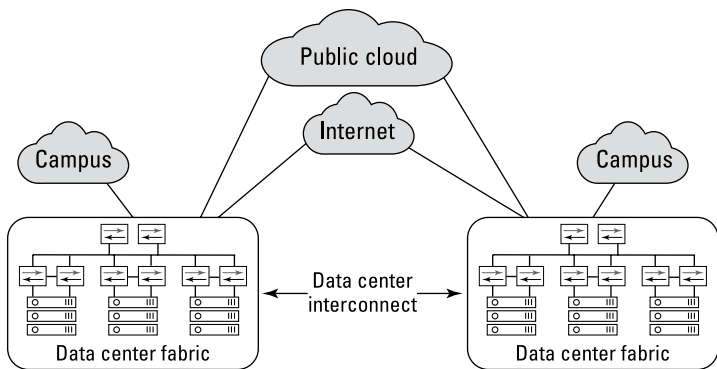


FIGURE 2-1: A data center interconnect.

Checking Out Fabric Networking in the Campus

Fabric networks represent the next evolutionary step for campus networking by offering simplicity, flexibility, and resiliency with inherent security. The right fabric technology will allow you to easily automate the secure attachment of IoT devices and enhance security by reducing the attack surface and preventing lateral movements. It will also increase agility and streamline operations with automation of network services and zero-touch provisioning.

Here are a few of the technologies that benefit greatly from the implementation of fabric networking in the campus:

- » **IoT:** One of the largest challenges you will face as a network operator is IoT. You need to onboard and extend connectivity to these devices and at the same time ensure security. Lacking in robust security features, many IoT devices rely on the network for protection.



TIP

IoT is the extension of Internet connectivity into physical devices and everyday objects. A few examples include security cameras, medical devices like magnetic resonance imaging (MRI) machines and infusion pumps, intelligent lighting systems and other smart building systems, and programmable logic controllers and sensors.

In addition to being able to easily extend connectivity for these devices with little to no manual configuration, fabric networking enables you to separate groups of IoT devices into their own secure zones or network segments, isolating them from the rest of the network far more easily than traditional VLAN, firewall, or virtual routing and forwarding (VRF) segmentation.

» **Network segmentation:** The increase in frequency and sophistication of cyber-attacks, combined with new attack vectors, such as IoT and cloud, requires reexamination of network security. In fact, companies now assume there will be a breach, and focus on minimizing the damage that can occur as a result. Network segmentation ensures that if a breach takes place, it is contained to where it occurred. Fabric networking allows you to segment the network at scale. Segments can be used to isolate IoT devices or groups of users, and segregate critical or sensitive information to assist with compliance and regulatory requirements.

Security professionals agree that network segmentation is a must for modern networks. By segmenting your network, you provide a greater degree of protection to your business's most valuable assets.



TIP

» **Integrated wired and wireless:** In the era of BYOD, you need to enable your workers to do their jobs while you also keep the network secure. Fabric networks allow you to easily segment your network and create a special segment for untrusted devices, like personal cell phones and laptops, that is completely isolated from the rest of the enterprise network. Furthermore, some fabric technologies are designed to extend to the wireless network to provide a more unified wired and wireless network. This enables the unified and dynamic attachment of users and devices to fabric services to dramatically simplify management and operations.

» **IP multicast:** Enterprise applications that rely on IP multicast can be a challenge to network managers. Applications, such as video streaming, IP television (IPTV), digital signage, software distribution, and others, rely on IP multicast to distribute traffic from a single source to multiple destinations.

The technologies required to make multicast work over a traditional network are complicated, involving protocol overlays that must be kept meticulously in sync. These protocols are difficult to configure and troubleshoot, reconvergence times can be slow, and scalability can often

be limited. Some fabric technologies excel in areas such as multicast and dramatically simplify deployment by making it easier to configure, faster to reconverge, and easier to scale.

» **IP video surveillance:** IP video surveillance is an application that is transitioning to IP multicast. The challenge is that IP multicast was designed for applications like IPTV, where a single source sends traffic to multiple destinations. IP video surveillance, on the other hand, typically involves many sources (IP cameras) sending traffic to just a few destinations. Since IP multicast wasn't specifically designed to address this scenario, implementation can be even more challenging.

The right fabric networking solution can support even this very complex form of multicast, making the network scalable, resilient, and far simpler to deploy and operate so that a surveillance deployment works seamlessly.

» **Edge computing:** Edge computing is a development on the rise due to the growth of IoT and sensors. It is the migration of compute to the edge of the network, away from centralized data centers, in order to reduce the distance that the data must travel.

Fabric networking can deliver on the network being more agile, reliable, and flexible for edge computing.

Figure 2-2 shows an example of campus fabric networking.

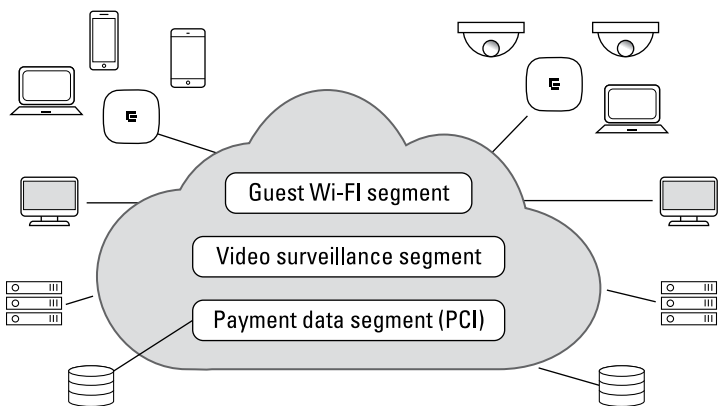


FIGURE 2-2: Campus fabric networking.

IN THIS CHAPTER

- » Understanding Shortest Path Bridging (SPB)
- » Looking at complementary technologies
- » Understanding IP fabrics

Chapter 3

Exploring Fabric Networks

Chapter 1 introduces the two main types of fabrics in the industry today: Ethernet-based fabrics, normally built with either Shortest Path Bridging (SPB) or Transparent Interconnection of Lots of Links (TRILL), and IP-based fabrics, normally built with Border Gateway Protocol (BGP) and Ethernet Virtual Private Network (EVPN).

In this chapter, you learn about the key aspects of each of these types of fabrics. After completing this chapter, you will understand how these fabrics work, what the main values of them are, and where they fit in the network.

Introducing Shortest Path Bridging

Shortest Path Bridging (SPB) is a widely adopted fabric technology. Although two variants of SPB exist (SPBV and SPBM), SPBM is where the current momentum lies and is what this book discusses. As shown in Figure 3-1, SPB is an Ethernet-based fabric where all networking services, whether Layer 2, Layer 3, IPv4,

IPv6, or multicast, are virtualized and decoupled from the physical infrastructure.

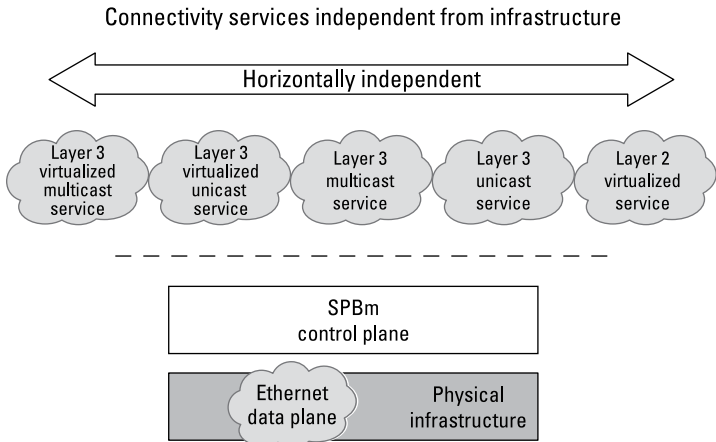


FIGURE 3-1: Shortest Path Bridging (SPB).

SPB enables businesses to implement flexible, scalable networks. Here are some of the characteristics of a fabric network that uses SPB:

- » Multiple physical topologies are supported, including mesh, partial-mesh, and rings. There are no blocked paths; thus all links are always forwarding.
- » Logical loops in the network are eliminated.
- » The control plane calculates the shortest path through the fabric and therefore removes the flooding and learning of a typical Ethernet network.
- » SPB consistently delivers sub-second failover and recovery for all network services with a single network protocol.
- » SPB supports edge-only provisioning that is done in real time with minimal configuration with dynamic auto-attach, enabling authenticated endpoints to connect seamlessly to fabric-based services.
- » SPB enables all network services with a single control plane.

Understanding How SPB Works

This section introduces some of the key concepts that explain how SPB works.

Building the SPB fabric

SPB uses Intermediate System to Intermediate System (IS-IS) as the control plane to create a “stateful” network topology. With IS-IS, each switch advertises itself to all other nodes in the fabric area so that each has a complete network topology map.

The standard defines a 24-bit service ID (I-SID). An I-SID is a unique service identifier, used within the SPB fabric, that can be extended to any service whether it is Layer 2 or Layer 3. It overcomes traditional virtual LAN (VLAN) scaling limitations and enables the creation of secure network segments.

SPB packet headers are based on Ethernet Media Access Control (MAC) layer addresses. However, when bridging traffic, it does not use typical flooding and learning mechanisms. Instead, it uses a routing protocol (IS-IS) to populate all forwarding entries. This results in a predictable and robust fabric infrastructure that is not prone to network loops.

Packets are forwarded through the fabric using the backbone MAC addresses. The user packet header is encapsulated. It is used only for forwarding, and only at the fabric boundary.



TIP

SPB is standardized by both the Institute of Electrical and Electronics Engineers (IEEE) and the Internet Engineering Task Force (IETF). In the IEEE, it was originally standardized as IEEE 802.1aq. IEEE 802.1aq was then added to IEEE 802.1Q as an amendment. The Layer 3 extensions to SPB are defined in RFC 6329.

Deploying SPB services

Typically, network operators think in terms of constructs like VLANs, IP subnets, and virtual routing and forwarding (VRF). Those connectivity configurations are treated as network services in an SPB infrastructure and are named virtual service networks (VSNs) and are comparable with virtual network identifiers (VNIs) in an IP fabric network. VLAN extensions across a network infrastructure are called L2 VSNs, and VRFs extended across the

network are called L3 VSNs. The ease of extending VSNs across an SPB network, the use of a single protocol, and the clear separation of the services or VSNs from the underlying fabric infrastructure ensure the simplicity and agility of the SPB solution.

Layer 2 virtualized network services (L2 VSNs) allow you to take Layer 2 VLANs and extend them anywhere in the network — even across geographical distances. These Layer 2 segments can either be E-Line (point-to-point) services, E-Tree (private LAN) services, or E-LAN (any-to-any) services.

At the source edge nodes of the network, VLANs are mapped to SPB I-SIDs, as shown in Figure 3-2. The I-SID provides the backbone connectivity to interconnect the L2 VSN service end points. Then at the terminating edge node or nodes, the I-SID is mapped to the VLAN.




FIGURE 3-2: SBP Layer 2 service.

Layer 3 virtualized network services (L3 VSNs) enable enterprises to separate portions of the network into logical segments, restricting communications to those members of that segment. Layer 3 virtual private networks have their own routing topology, providing a high degree of isolation. They are also scalable and easy to deploy. Layer 3 VSNs are useful in multi-tenant environments as well as environments where IoT devices, applications, user groups, or critical information must be in its own secure segment for security purposes.

At the source edge nodes of the network, Layer 3 VRFs are mapped to I-SIDs, as shown in Figure 3-3. IS-IS then advertises the service and IP routes only where the VRF IP routes are needed. These routes are installed only on nodes that contain the same I-SID. This is all done natively within the SPB fabric without requiring any additional routing protocols.




FIGURE 3-3: SPB Layer 3 service.



REMEMBER

Any SPB service is enabled simply by provisioning at the edge of the fabric, and IS-IS dynamically establishes the shortest path between the service endpoints.

IP shortcuts leverage the global routing table to forward IP packets directly over the SPB network. Conceptually they are similar to Layer 3 VSNs; however, the difference is that rather than being mapped to an I-SID, the IP encapsulation is mapped directly to an Ethernet header, as shown in Figure 3-4. A route look-up is done on the edge to determine the packet's destination. The edge node adds the Ethernet header that contains the destination MAC address. The fabric then efficiently cut-through switches to the destination node without any hop-by-hop route look-ups.




FIGURE 3-4: IP shortcuts.



REMEMBER

Layer 3 VSNs and IP shortcuts support both IPv4 and IPv6 routing.

IP multicast services bring complexity with all the configuration that goes along with setting up Protocol-Independent Multicast (PIM) or Distance Vector Multicast Routing Protocol (DVMRP). With SPB, you don't have to worry. You get high-performance, multicast functionality but without the complex configuration overhead. Multicast over SPB also offers superior performance for any type of multicast deployment whether it is one-to-many (IPTV), many-to-few (video surveillance), or many-to-many. When a multicast stream is received, it is mapped to a dedicated multicast service identifier (or I-SID), as

shown in Figure 3-5. IS-IS then advertises the I-SID to the rest of the fabric and forwards it only to nodes that register to receive it through IGMP. If a node does not request the stream, it is not forwarded, enabling far more efficient distribution of multicast traffic than in a traditional network.

IP multicast can be enabled over both L3 and L2 VSNs or over the global routing table as IP multicast shortcuts.




FIGURE 3-5: Multicast services.

Examining the Security Attributes of SPB

One of the biggest challenges with traditional IP network design is that the entire routing table is exposed to every device and every endpoint in the network. This characteristic leads to the ability for a malicious attacker to gain entry to the network, discover the network topology using common IP scanning mechanisms, then move through the network to access your valued data.

In an SPB network, this vulnerability is eliminated in many ways. First, routing is enabled only at the network edges. It doesn't reside in the network aggregation or core. Rather than forwarding based on IP look-ups, traffic is forwarded using Ethernet switched paths. Therefore, anyone running an IP scan against an SPB environment gets a list of IP subnets all showing just a single hop to the

gress of the network. Everything in between is “dark.” This inability to discover the topology of the network makes it nearly impossible for hackers to laterally move to sensitive areas of the network.

The other important piece to security is the ability to segment the routing tables through Layer 3 VSNs. These secure segments run as ships in the night without any awareness of each other and without allowing any access, in or out, unless otherwise configured. Having a highly segmented environment ensures that if the network is breached, that breach is contained to where it occurred, thus minimizing damage.

This inherent security extends to Layer 2 network design. In traditional Ethernet networks, endpoint MAC addresses such as those belonging to a user’s computer or to an IoT device are visible throughout the network. In SPB, endpoint MAC addresses are only visible on the fabric edge switches.

The final way SPB enhances an organization’s security posture is by eliminating back-door entry points to the network. SPB VSNs can be elastic in nature; they can extend and retract as authenticated users and devices connect to the network and disconnect. When a user disconnects from a switch port and access to the segment is no longer required, residual configuration is automatically deleted on the edge switches. This practice not only removes the delays and risks associated with manually configured conventional networks — it also eliminates the risk of a back-door entry point to the network.

Looking at Complementary Technologies

Once an SPB core is in place, you can deploy additional features to increase the reach and the value of the deployment.

Dynamic Auto-Attach

Auto-Attach (IEEE draft P802.1Qcj) provides for automatic attachment of users, devices, and virtual machines (VMs) to connect to SPB services or I-SIDs. It uses extensions to the IEEE802.1AB Link Layer Discovery Protocol (LLDP) to automatically attach network devices to I-SIDs or VSNs in an SPB network. This auto-attach capability can be deployed on endpoints, such as IP surveillance cameras; wireless APs; and/or non-SPB-compliant access layer switches so that seamless communication with the SPB fabric is possible.

Seamless extension of SPB

Fabric Extend is a feature that enables SPB to be extended over third-party IP networks whether they are IP cores or service provider wide area networks (WANs). Any SPB service, whether it is L2, L3, or multicast, can be seamlessly extended across the WAN or across the IP core easily — without the WAN or IP core having visibility to those services.

Fabric Extend supports an IP mode, an L2 mode or an encrypted mode. It works by creating an overlay of point-to-point tunnels using either Virtual Extensible Local Area Network (VXLAN) native L2 tunnels for IP mode or encrypted IPSec tunnels.

Designing multi-site networks with mobile endpoints

Distributed Virtual Routing (DVR) is an SPB feature that improves the scale and performance of multi-site networks with mobile endpoints. These can be either VMs or wireless clients at the network edge.

DVR's main value is that it eliminates the “tromboning” problem that can occur when a user's or device's IP subnet is stretched far away from its default gateway. DVR distributes the routing function to all switches that have a presence in the IP subnet so that the default gateway is always available at the first network hop. At the same time, it keeps provisioning simple because the routing configuration is performed only on centralized controllers or spine/ aggregation nodes.

Seeing Where SPB Fits in the Network

The earliest implementations of SPB were focused on Metro Ethernet services. The technology then had uptake within the data center to solve the challenges with VM mobility. Today, most of the momentum for SPB is for campus networking and collapsed or converged data center/campus environments.

SPB's dynamic auto-attach capability, inherent network security, and integrated high-performance multicast make it ideal for these environments.



TIP

Though multi-vendor interoperability tests have been done and are publicly documented, not all implementations of SPB are the same. This chapter provides a guide as to what the technology is capable of; however, due diligence must be done to ensure that the vendor you are working with supports a comprehensive implementation of SPB.

Introducing TRILL Fabrics

If you're a Star Trek fan, I know you got excited when you saw this section. But if you were expecting a discussion on fashion trends concerning your favorite humanoid aliens with spots, you are in the wrong place. *TRILL* stands for Transparent Interconnection of Lots of Links. The TRILL algorithm was written by Radia Perlman, who also wrote the algorithm for Spanning Tree Protocol (STP). Like SPB, TRILL was designed to address the shortcomings of STP; specifically, to remove blocked links and to be topology agnostic without having logical loops in the network.

Another major advantage of TRILL technology is simplicity. This technology abstracted away the protocols and the need for intense manual configurations. These advantages helped this new way of networking gain traction quickly. Switches and routers that implement TRILL are referred to as *routing bridges* (*Rbridges*). When connected, they automatically form a fabric. As services (such as VLANs) are provisioned at the edge ports, they are automatically learned across the fabric, with no need to tag links between switches. When a new end device is attached to the fabric, the entire fabric is made aware of the device's MAC along with which switch and ports it resides on. Any other device needing to communicate with it already knows where it is attached. Thus there is no need to flood traffic looking for it.



REMEMBER

Like SPB, TRILL allows you to take advantage of the simple configuration that you get with an Ethernet network, but it uses a link-state routing protocol (control plane) to avoid traditional Ethernet-based flooding and learning. The link-state protocol determines the shortest available path to the destination and routes packets accordingly. The most common link-state protocols used with TRILL fabrics include IS-IS, Open Shortest Path First (OSPF), and Fabric Shortest Path First (FSPF).

Seeing Where TRILL Fits in the Network

Over the past several years, the concept of any service on any port, the ability to provide a reliable infrastructure for storage technologies, and ecosystem integrations with hypervisor vendors such as VMware, make TRILL a fit for the virtualized data center. This technology continues to be a viable fabric technology solution for the data center today; however, there is now an alternative solution with IP fabrics.

Delving into IP Fabrics

As data center networking continues to evolve, there is a growing need to implement fabric technology that is not only standards-based, but also fully interoperable among vendors. At the same time, providers who already had an Multiprotocol Label Switching (MPLS)-based IP-VPN solution wanted to extend this technology into the data center without the complexity of MPLS. This was the birth of the IP fabric, based on BGP as the control plane (now called the *underlay*) and BGP/EVPN as the data plane (now called the *overlay*) using VXLAN.



TIP

The underlay can use other routing protocols such as OSPF or IS-IS if desired.

Understanding IP Fabrics

This section goes into the details of how the IP fabric is implemented and how services are provisioned in and across the fabric.



REMEMBER

Although IP fabrics are based on BGP, it is only a small subset of the protocols that are necessary for the fabric.

Building the IP Fabric

The industry has moved into a common architecture for data center fabric networking by using the spine/leaf topology. This can take the form of a three-stage Clos (pronounced “clo” from its originator Frederick Clos), which is a simple spine and leaf, as

shown in Figure 3-6. Scaling this design out, the topology can move into a five-stage Clos with the addition of a super-spine that connects three-stage Clos points of delivery (PoDs). Moving to a consistent architecture makes traffic flow across the network predictable and much easier to troubleshoot.

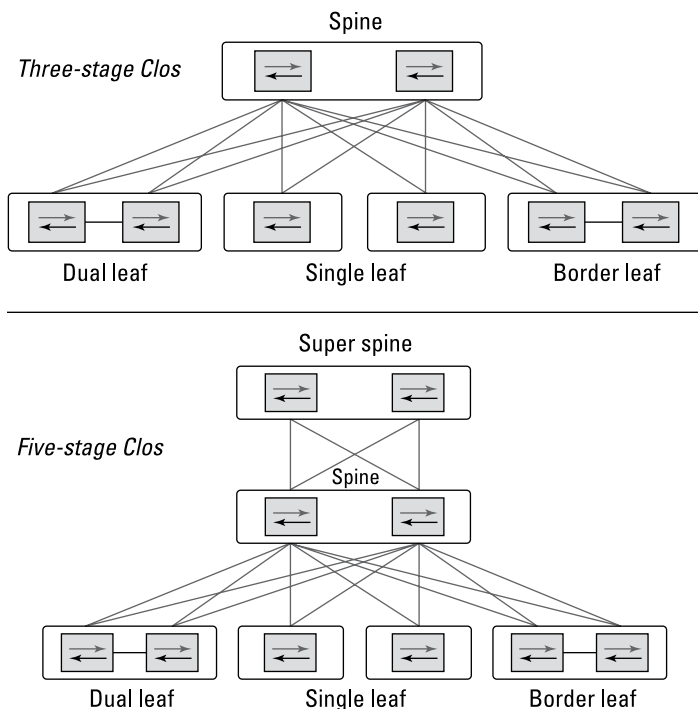


FIGURE 3-6: Three-stage and five-stage Clos designs.

One key aspect of the IP fabric is its use of BGP, which is a well-known and well-understood technology. BGP forms the underlay network for the fabric. This is how all switches and routers connect with each other and share routing and topology information. In the fabric, every leaf connects to every spine and forms a BGP neighborhood. The fabric uses equal-cost multi-path routing (ECMP) to distribute traffic across all links, which provides the bandwidth and resiliency required in the modern data center.

Deploying Services

Once the fabric infrastructure is in place, the next step is to enable services on the fabric to support connections from end devices such as servers, storage, and appliances. The overlay is used to extend services across the fabric. BGP/EVPN using VXLAN makes this extension simple and easy, as shown in Figure 3-7. The default gateway for every end device is at the leaf (top of rack) using a static anycast gateway. When VLANs or VRFs need to extend beyond the leaf, a VXLAN tunnel is used. Each leaf or leaf pair becomes a VXLAN tunnel endpoint (VTEP). This arrangement allows the leaf to encapsulate VLANs to VXLAN tunnels and do the reverse, taking VXLAN tunnels and breaking out the VLANs. As per the standard, the VTEPs use an auto-discovery mechanism to create tunnels across the fabric, thus eliminating manual configuration.




FIGURE 3-7: BGP-EVPN (VXLAN) design.

The creation of a Layer 2 or Layer 3 VNI facilitates the transport of traffic as described in the next two sections. Both of these concepts use Integrated Routing and Bridging (IRB). *Asymmetric IRB* is the case where the same VLANs exist on both ends of the VXLAN tunnel. Routing is done on the ingress leaf and then traffic is sent across an L2 VNI to the destination leaf (route first, bridge second). *Symmetric IRB* is used when the same VLANs don't exist on both ends of the VXLAN tunnel. On ingress, the VLAN routes to its destination via the L3 VNI to communicate to a different VLAN across the VXLAN tunnel.

Layer 2 VNI

VLANs entering the leaf from the edge device are mapped into a VXLAN tunnel using a VNI. In most cases, it's easy to map the VLAN ID to the VNI ID without creating a complicated mapping

structure. The VLANs are placed into an EVPN instance on the leaf and VXLAN tunnels are automatically created to all other leaf switches in the fabric where that VLAN exists. You don't need an external controller or manual mapping of tunnels; it's all automatic.

Layer 3 VNI

When there is a need to extend VRFs between leaf switches, a Layer 3 VNI is created. This is a dedicated VLAN with routing enabled on it that is used as the “connector” between leaf switches for that VRF. This allows VLANs within VRFs to communicate across the fabric and eliminates the need for any additional routing protocol for this purpose.

Looking at Complementary Features

The following features and functions are typical in IP fabric implementations:

- » **EVPN Instance (EVI)** is a construct used for extending VLANs or VRFs across the IP fabric. When a VLAN is added to the EVI, this triggers a routing update to be sent, and VXLAN tunnels are automatically created to all other leaf switches where the VLAN exists. When the VLAN is deleted, the tunnel is taken down automatically.
- » **Automatic VTEP discovery** happens during the EVI process. There is no need to configure any static tunnels because the creation and deletion are automatic. The VTEPs talk to each other over the underlay network that is providing IP connectivity across the fabric.
- » **VXLAN traffic distribution** is achieved by changing the source User Datagram Protocol (UDP) port in the VXLAN packet. This is part of the VXLAN standard and happens automatically. By changing the source UDP port, the hashing algorithm sends traffic across different ECMP links in the fabric. This ensures that one link will not get saturated while other links remain idle.
- » **Static Anycast Gateway** is the default gateway for the end devices. This gateway will exist on every leaf switch where the VLAN is present. No control protocols are needed because this is a static configuration across the fabric.

- » **Address Resolution Protocol (ARP) suppression** is accomplished by keeping ARP information in a suppression cache on each leaf, thus reducing the amount of traffic across the fabric for end station discovery.
- » **MAC/IP Learning** is done via BGP/EVPN, which eliminates nearly all flooding of traffic to find end stations. As soon as a MAC is learned on the leaf, a routing update is sent to all leaf switches telling each of them the MAC just learned and the VTEP IP where that MAC is attached.

Seeing Where IP Fabrics Fit in the Network

IP fabrics have gained significant momentum in the data center. This success can be attributed to the simple scaling that the fabric provides, its virtualization capabilities to have any VLAN on any port in any data center, and the inherent automation. Because deterministic behavior is so critical in the data center, the IP fabric and its spine/leaf topology are the perfect fit for this.

The IP fabric is also seeing traction outside of the data center in uses for data center interconnects as well as campus implementations. The attraction of having a standards-based and fully interoperable fabric that uses a simple underlay/overlay concept is piquing the interest of network administrators looking for ways to streamline operations with automation and provide services quickly and easily across an infrastructure that is adaptable through virtualization.

Every use case and need is different; therefore, picking the right fabric for the right use case will vary. As IP fabrics continue to gain momentum, their use will continue to diversify.

IN THIS CHAPTER

- » Enabling automation
- » Gaining operational efficiency through artificial intelligence and machine learning
- » Enhancing visibility
- » Integrating ecosystem partners
- » Improving security
- » Integrating wired and wireless

Chapter 4

Examining Key Considerations in Fabric Technology Evaluations

If you think that fabric networking sounds like the coolest thing in the world . . . you're right. Your next step is to choose a solution that will meet your needs. As with any solution, you should take certain key considerations into account. New infrastructure is an investment after all, so you want to choose a solution that meets your needs now but can grow with you as your business needs evolve.

In this chapter, you learn about the key considerations that you need to review with any fabric networking solution. These traits far outweigh price and product feeds and speeds in the evaluation process.

Enabling Automation

Automation is the new buzzword in information technology circles. With all the manual processes that are done daily, it's easy to see why automation is so critical. It allows manual operational tasks to be offloaded from the network administration team, leaving network administrators to focus on strategic projects.

There are different stages of automation, so think of it as a journey. Where you start and where you stop is your choice and based on the needs of the business.

Network automation

Automating network tasks falls into two areas:

- » **Infrastructure provisioning** entails all the configuration tasks to instantiate the fabric. These extend from the point where switches and routers are powered on and cabled to the point where the fabric is ready to have end devices attached to it.
- » **Tenant/services provisioning** includes creating all the necessary configuration for end devices to attach and use the fabric. Over the lifecycle of a network, this is where most configuration tasks occur.

With regard to infrastructure provisioning, once the devices are powered on and cabled, it should be simple plug and play to build the underlying fabric topology.

Once the fabric network has been built, it is time to build the network services. Being able to leverage the power of automation is where you can recognize significant time savings. Some fabric technologies have embedded automation features and can establish network connectivity services on demand without relying on any external controllers or management systems. An example is the dynamic establishment of a guest Wi-Fi service as a personal iPad is detected through Institute of Electrical and Electronics Engineers (IEEE) auto-attach features and on-boarded onto the network.

Some fabric technologies work synergistically with centralized management or controller-based systems that can deliver automated provisioning of the network based on an external

event or trigger. An example is the dynamic provisioning of top-of-rack switch ports as a new virtual machine (VM) is turned up, through integrations between the hypervisor and the network environment.



REMEMBER

Differing approaches to automation exist. One is for the protocol to deliver inherent automation capabilities so that services can be dynamically established. The other is for the fabric infrastructure to be programmable, meaning that it supports application programming interfaces (APIs), or it supports open configuration through programming languages such as NetConf, RESTConf, Ansible, or Python. An external controller or management system would then be used to automate the network infrastructure.



TIP

What's key is not the underlying technology used to automate — but the result! That means what's important is that the technology enables you to achieve better agility. This, in turn, enables you able to respond quickly to business requirements while eliminating time-consuming and risky manual provisioning.

Cross-domain automation

Cross-domain automation is a step forward from network automation because you are automating components from other IT domains to complete the entire service chain. This process automation can include the provisioning of servers, storage, and networking, as well as network functions such as firewalls and load balancers, as shown in Figure 4-1.

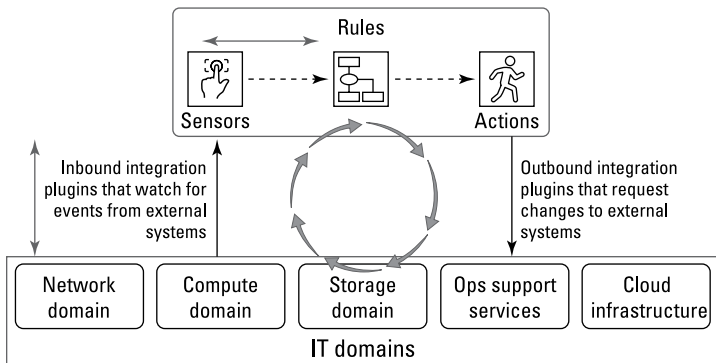


FIGURE 4-1: An overview of cross-domain automation.

This type of automation is important because most organizations have silos of teams that need to work together in order to bring a new application or service online. There is inherent delay because of manual hand-offs between teams and the back-and-forth to clarify what's needed and when. Cross-domain automation simplifies this process and ensures that teams are working in the most efficient way possible.



REMEMBER

This stage of automation is not solely about provisioning; it also includes full lifecycle management. It allows you to automate any manual process that is repeated on a regular basis, including processes related to validation, troubleshooting, and remediation. An example is the creation of a workflow for your help desk that automates the gathering of information when a user calls in, improving the speed and efficiency of that process.

A key requirement of a cross-domain automation tool is that it is customizable so that it can fit into your existing IT processes and workflows. Generally, these tools are DevOps inspired and can be programmed easily through common tools such as Python. Another consideration is to ensure that the tool has ties to an open source community, such as Github. When you are looking to automate a particular process, you can see what others have done and simply modify existing workflows to fit your environment.



TIP

Take the micro-services approach to process automation by creating several small workflows and tying them together into a larger workflow. This approach makes the creation of workflows easier and enables you to reuse the small workflows many times in different larger workflows, thus getting the benefits of reuse.



REMEMBER

Cross-domain automation is a big step forward from network automation, but you can ease into it by creating small workflows and building from there. This is a great way to dip your toe in the water.

AI/ML driven automation

Wouldn't it be great if your fabric networking solution could use artificial intelligence or machine learning? I know, it sounds like science fiction, but it's worth your consideration. AI/ML technologies hold the promise of achieving true autonomous networking — where the network becomes self-learning, self-healing, self-securing, and self-optimizing.

What's the difference between automation and true autonomous networking? Automation streamlines manual processes, but it isn't yet intelligent. When you combine automation with AI/ML within a reliable network infrastructure (like a fabric), you have a truly autonomous network.

Although truly autonomous networking will develop gradually over time, AI/ML technology is here today and can make your life easier. For example, you can troubleshoot performance problems by extracting meaning from large amounts of data in network logs and databases. Another example is 24/7 monitoring of IoT devices, where normal behavior is learned through AI/ML, and anything that deviates from that behavior is remediated against quickly and efficiently.



REMEMBER

In general, AI and ML close the feedback loop between large amounts of data/analytics and automation tools. This form of intelligence can help reduce costs while enhancing reliability, availability, overall performance, and the productivity of network-operations teams and end-users alike.



TIP

When evaluating fabric networking solutions, you should consider the vendor's strategy regarding automation and AI/ML. Because AI/ML is poised to be a significant part of next-generation networking, it is recommended that you start evaluating AI/ML technologies for targeted use cases.

Enhancing Visibility

The fabric solution you choose should allow for an in-depth view of the network, its services, its applications, and its connected users and devices.

Having the right visibility can help avoid business disruptions through the proactive monitoring of application performance and the network. It facilitates troubleshooting by giving a deep, real-time view into the network and its traffic, and it can improve network security through extensive visibility into unapproved applications, unusual traffic, and shadow IT.

On-box visibility

On-box visibility is the ability to view traffic that traverses a particular fabric-enabled switch or router. Previously through features such as Switched Port Analyzer (SPAN) or Remote SPAN (RSPAN), a traffic sniffer or protocol analyzer would be plugged into the switch or router, either locally in the case of SPAN, or remotely in the case of RSPAN, for traffic on multiple virtual LANs (VLANs) or ports to be analyzed.

Now, with many switches using x86 hardware, it is possible for the traffic analyzer function to be run directly on the switch within a guest VM. This approach reduces the mean time to repair. Another big advantage is that multiple types of analysis tools can be run on the VM. Examples include Wireshark, TCPDump, Splunk, and PerfSONAR.

Network analytics

In addition to on-box visibility, having an end-to-end view of the network, its services, its applications, and its connected users and devices is essential. Seeing application and network response time, top talkers, and top applications provide actionable insights into the overall IT infrastructure. Having the right analytics data can give you a better understanding of user behavior on the network. It can also help with enhancing security by being able to pinpoint unusual traffic and unauthorized applications.

Some of the key criteria in an analytics platform include having full Layer 4-7 visibility end-to-end across the network — from the user or end device all the way to the data center and out to the cloud. It's important to consider a tool that extends into the hypervisor environment in the data center for visibility into traffic between VMs that never touch a physical switch.

Another criterion is how easily the data is collected from the network devices. With some solutions, the data can be collected directly from the network switches and routers in an efficient and cost-effective way through embedded telemetry features. In other scenarios, external probes or sensors might be required. These add significant cost and complexity when altering the physical cabling environment.

You should deploy a solution that has a robust catalog of application fingerprints for detecting and measuring the performance of

the applications that your business uses today. Having the ability to easily customize the solution to include any custom applications is also a huge benefit to ensure full coverage of applications across the infrastructure.

Integrating Ecosystem Partners

A fabric networking solution should not create vendor lock-in. Instead, it should provide automated, simplified network connectivity throughout your overall IT ecosystem.

With so many solutions available across the campus and data center, it is virtually impossible for a fabric solution to support every ecosystem partner. However, there are high runners that all should and likely do support. Examples of these include hypervisors such as VMware, Microsoft, and KVM. These might also include a multitude of storage options that utilize Internet Small Computer Systems (iSCSI), Network File System (NFS), or Non-Volatile Memory Express (NVMe); either as stand-alone storage arrays or hyperconverged infrastructure (HCI).

In the campus, many IoT and security solutions are available that should provide some level of integration with the fabric or its associated management platform. These include IoT devices like IP cameras, third-party switches, and security products like firewalls and access control systems. This capability enables the infrastructure to be managed holistically, allowing the network to react in real time to alerts received by third-party solutions. It is also important for fabric-based switches to support traditional protocols so that interoperability with existing networks is not an issue.



REMEMBER

With your fabric solution, APIs should be available to ensure smart integrations between your networking environment and the rest of your IT infrastructure.

Improving Security

A fabric networking solution should make securing your network simpler. Many methods of security are available including macro-segmentation, micro-segmentation, hyper-segmentation, and virtual routing and forwarding (VRF). All of these accomplish

the segmentation of traffic across the network to add levels of security.

However, having a network that can easily be segmented at scale allows you to improve your overall security posture by dramatically reducing the attack surface and preventing lateral movement to more sensitive areas of the network. A stealth network prevents malicious actors from discovering the network topology. In addition, service elasticity removes potential back-door entry points to the network by removing residual configuration.

Are you in a network that has strong regulatory requirements? Through fabric security features like segmentation and stealth, compliance with the European Union's General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry Data Security Standard (PCI-DSS) can be facilitated.

Integrating Wired/Wireless

You need a fabric solution that can work across wired or wireless networks. For maximum simplicity, the solution should offer the ability to converge the wired and wireless networks. Having these tightly integrated provides a seamless user experience no matter what device or devices your users attach to. It also significantly eases the administration of the network by eliminating the need to manage two disparate network environments.



TIP

Being able to configure network services from a single pane of glass — regardless of whether they are wired or wireless — is a huge win for your business as well as for your network administration team. It is also important that common policies be applied to users and devices across wired and wireless, and that common telemetry information be collected from the network switches as well as the APs so that you can gain a clear view of network services and application performance to simplify troubleshooting.

- » **Focusing on your pain points**
- » **Evaluating fabric networking technologies**

Chapter **5**

Choosing the Right Fabric Solution for Your Business

Chapter 3 introduces the different fabric technologies and the relative fit of each technology in the network. As that chapter explains, Shortest Path Bridging (SPB) is well positioned for the campus or metropolitan area network (MAN) as well as collapsed data center/campus cores, while Transparent Interconnection of Lots of Links (TRILL) and now more significantly IP fabrics are well positioned for the data center with increased traction in the campus.

With the lines blurred between the different fabrics, where do you start with the evaluation process?

Focusing on Your Pain Points

A logical way to figure out what fabric technology is best suited for your business is to determine what your biggest networking pain points are and what technology might be the best fit to solve those.

Here's an overview of the top values for each technology:

»» **SPB**

- Faster time-to-service by eliminating core and aggregation node reconfiguration (provisioning at the edges only)
- Dynamic auto-attach and segmentation for IoT devices, users, and virtual machines
- Simplification of the network with the ability to use just a single control plane for all Layer 2 and Layer 3 services
- Unification of wired and wireless networking
- Inherent security through network segmentation at scale, stealth topology, and service elasticity
- Less complex administration of IP multicast
- High-performance and scaling for IP video surveillance
- Topologies supporting any service, any port, any place
- Ecosystem integrations, virtual machines (VMs), Internet of Things (IoT), security

»» **TRILL**

- Simple plug-and-play data center networking
- Built to support server and network virtualization
- Embedded automation to simplify provisioning and operations
- Fully programmable with application programming interfaces (APIs)
- Topology agnostic supporting any virtual LAN (VLAN), any port, any place
- Simplified convergence of IP storage
- Ecosystem integrations, VMs, storage, security

»» **IP fabrics**

- Vendor-agnostic deployments with full interoperability
- Scale from the very small to the very large
- Simple plug-and-play data center networking
- Built to support server and network virtualization
- Embedded automation to simplify provisioning
- Fully programmable with APIs

- Topologies supporting any VLAN, any port, any place
- Simplified convergence of IP storage
- Ecosystem integrations, VMs, storage, security

Evaluating Fabric Networking Technologies

To help you navigate the confusing process of evaluating fabric solutions, Chapter 4 provides guidance on what to look for in your fabric technology. This section builds on that discussion, providing simple checklists that you can use and customize. These checklists identify the important features to look for to help you modernize campus and data center networks:

» Data center networking checklist

- Must be easy to deploy and operate
- Vendor-agnostic, industry standard with full interoperability
- Embedded automation to simplify provisioning
- Scales from the very small to the very large
- Topologies supporting any VLAN, any port, any place
- Supports the ability to add new infrastructure (links and nodes) with no service interruption
- Designed from the ground up to support virtualization
- Fully programmable with APIs
- Decouples connectivity services provisioning from infrastructure topology configuration
- Provides rapid time-to-service for provisioning in real time, and must not require device-by-device, hop-by-hop configuration
- Simplified convergence of IP storage
- Ecosystem integrations, VMs, storage, security
- Enables the scalable deployment of network segments to isolate and protect critical information

- Must provide on-box, application telemetry, and analytics capabilities to provide visibility into the network, users, traffic, and devices.
- Simple and resilient interconnect capabilities

» **Campus networking checklist**

- Must be easy to deploy and operate
- Supports small to very large campus environments
- Provides design flexibility allowing for hierarchical, full-mesh, partial-mesh, and ring-based topologies
- Decouples connectivity services provisioning from the fabric infrastructure
- Supports the ability to add new links and nodes without network downtime
- Provides rapid time-to-service for provisioning without requiring device-by-device, hop-by-hop configuration
- Provides virtualized IP multicast at scale without requiring the use of Protocol-Independent Multicast (PIM)
- Provides a secure and robust IP-based video surveillance service infrastructure that scales to thousands of sources seamlessly
- Provides a dynamic auto-attach feature to facilitate the connectivity of users and devices
- Enables the scalable deployment of network segments to isolate and protect critical information, groups of users, and IoT devices
- Limits lateral movement by making it difficult for hackers to discover the network topology using typical IP scanning techniques
- Eliminates residual configuration on edge switch ports as users and devices disconnect from the network dynamically
- Is built on industry standards
- Provides on-box, application telemetry and analytics capabilities to provide visibility into the network, users, traffic, and devices
- Delivers wired and wireless integration
- Delivers simple and resilient interconnect capabilities

IN THIS CHAPTER

- » Exploring the breadth of solutions
- » Managing through a single pane of glass
- » Supporting multi-vendor solutions

Chapter 6

Ten Things to Know About Fabric Solutions from Extreme Networks

Extrême Networks delivers world-class fabric networking solutions to organizations needing scalability, flexibility, and ease of administration. Industry leading and differentiating technology will help you transform the data center and campus network into an infrastructure that operates at cloud speed with efficiency and operational simplicity.

You Have a Choice

With Extreme, you can choose the fabric solution that best suits your use case and network requirements. If you are looking for campus or data center transformation, Extreme can solve your networking pain points with a flexible portfolio of solutions that includes Ethernet and IP fabrics. Extreme can work with you to ensure that you have the right technology for your environment.

Single Pane of Glass

Extreme's single operational model extends from the wired/wireless edge to the data center and multi-cloud, providing comprehensive management, policy, and analytics for fabric networking solutions. All Extreme fabric solutions for the campus and the data center use the same operational model with visibility and policies consistent end to end.

Visibility from the Edge to Multi-Cloud

Extreme provides a unique end-to-end view of application flows that extend from the campus edge to the data center, the virtualized environment, and even into multi-cloud environments. Regardless of the chosen fabric solution, you have the option of complementing agile, automated networking with deep insights into the network, applications, and users through application telemetry and analytics capabilities.

Solutions Are Truly Multi-Vendor

The solutions from Extreme Networks are industry aligned and based on industry standard technology. Extreme's solutions are designed to work seamlessly with whatever you have installed in your network today. This capability allows you to continue to get value from devices you have already purchased while gaining flexibility, agility, and scalability from the fabric networking solutions.

Deployment History

With thousands of fabric deployments worldwide, Extreme Networks has the knowledge and experience to provide a solution that will meet your needs today and in the future. Extreme has deployed solutions that span sizes and segments of both service provider and enterprise markets and have deployed fabric solutions in the most demanding, highly sensitive, and highly secure environments.

Innovation!

Extreme Networks holds more than 500 patents related to fabric networking. In addition to developing many of the drafts related to Shortest Path Bridging (SPB), the company continues to invest in research and development to continually improve its products and increase the functionality of its solutions. You can view Extreme Networks' patents at www.extremenetworks.com/company/legal/patents/.

Field Proven in a Demanding Environment

Extreme's fabric solutions are deployed worldwide and are at the core of most Internet exchange providers, along with deployments among cloud service providers, content delivery networks, and enterprises across all verticals and all sizes. This includes mission critical environments like hospitals, emergency response, and national defense as well as high-profile events such as the Sochi Olympic Winter Games.

Plug and Play

Fabric solutions from Extreme Networks can be deployed and fully operational within minutes, rather than days or weeks. They can self-provision rapidly, ensuring that you are up and running quickly with your next-generation networking solution. In addition, through network service automation and even cross-domain automation, you can reduce operational expenses and do more with less.

Ecosystem Integrations

Extreme Networks provides the networking foundation for your business. This is more than just plumbing; it's the enabler for services and applications that the business runs on. Having the

fabrics ready and able to integrate ecosystem partners is critical to success. By providing seamless integration with virtual machine (VM) providers, storage solutions, security partners, and IoT vendors, Extreme provides you with an integrated network that can be managed holistically.

Unbreakable in Hack-a-thons

The secure fabric solutions from Extreme Networks have been part of multiple private and public hack-a-thons. To date, not one has been breached. This is because of the inherent security of the solution with features such as hyper-segmentation, which isolates your virtualized networks; stealth networking, which hides your network's topology from prying eyes; and service elasticity, which allows your segmented network to extend and retract when devices connect or disconnect.



Extreme Elements™

Real Tech. Really Human.



The future of networking technology is here!

Fabric networking allows you to scale with demand, create redundant paths for high availability and fault tolerance, improve security with network segmentation, and much more. You can move from an old legacy network to a modern, scalable, secure network infrastructure that enables your business to succeed and grow. This book is your introduction to the world of fabric networking and the underlying protocols and technologies that support it.

Inside...

- Reduce time to service
- Improve network stability
- Eliminate downtime from human error
- Achieve better resiliency
- Enhance security



Extreme

Sara Perrott has an MS in cybersecurity and information assurance and holds several industry certifications such as CISSP and GCIH. She works full-time as a security engineer and teaches part-time.

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

ISBN: 978-1-119-61937-6

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.