White paper
Cisco public

Multiplanar Core Design Concepts in Secure Multilayer Networks

Contents

Abstract

This document discusses network design concepts currently being deployed in several of the largest carrier-class networks around the world. What is unique to this paper is, while it leverages these best practice concepts used today, it also introduces new concepts that target those unique customers that require external encryption devices for added security encryption compliance, such as public sector, defense, financial, and secure mission networks. The new approaches introduced target reducing complexity typically found in these multilayer networks, enhancing survivability, and simplifying day-to-day operations and troubleshooting, while enabling a more deterministic behavior across all layers of the backbone.

This document assumes the reader has a basic understanding of networking and Wide Area Network (WAN) design, Multiplanar Backbone (MPBB), Segment Routing (SR), and MACsec. For those not familiar with any of these concepts, links to each technology will be provided in the "References" section.

Introduction

One of the more modern core architecture designs being adopted today is the Multiplanar Backbone (MPBB). The MPBB provides a new level of service flexibility, scale, and resiliency not seen before. Adding the Segment Routing (SR) functionality to this design, the service capabilities of the backbone expand exponentially, adding another dimension of transport service capabilities and scale.

Several years ago, my colleague Michael Kowal and I co-authored "The Multiplanar Backbone," which discusses in detail the drivers, uses cases, and design options of the MPBB and, along with SR, the enhancements and capabilities it offers.

The purpose of this white paper is to leverage the MPBB and SR concepts and apply them to these highly secure networks requiring external encryption devices. The addition of these external encryptors typically adds an additional level of complexity, very analogous to the overlay and underlay concepts found in today's commercial network designs.

To illustrate, in the **Figure 1** example below, the "overlay" layer carrying the security-sensitive data runs IP over an overlay encapsulation such as Generic Routing Encapsulation (GRE), multipoint GRE (mGRE), or Virtual Extensible LAN (VXLAN). The "underlay" layer transporting the traffic between encryption (or overlay) devices is running IP, Multiprotocol Label Switching (MPLS), or SR and, to add to that complexity, each layer has a completely different topology, Interior Gateway Protocol (IGP), traffic forwarding patterns, and convergence characteristics. Many "overlay/underlay" combinations can exist, but the point remains—networks that require multiple layers increase significantly in complexity, whether the overlay is comprised of routing elements or external encryption devices, or both.

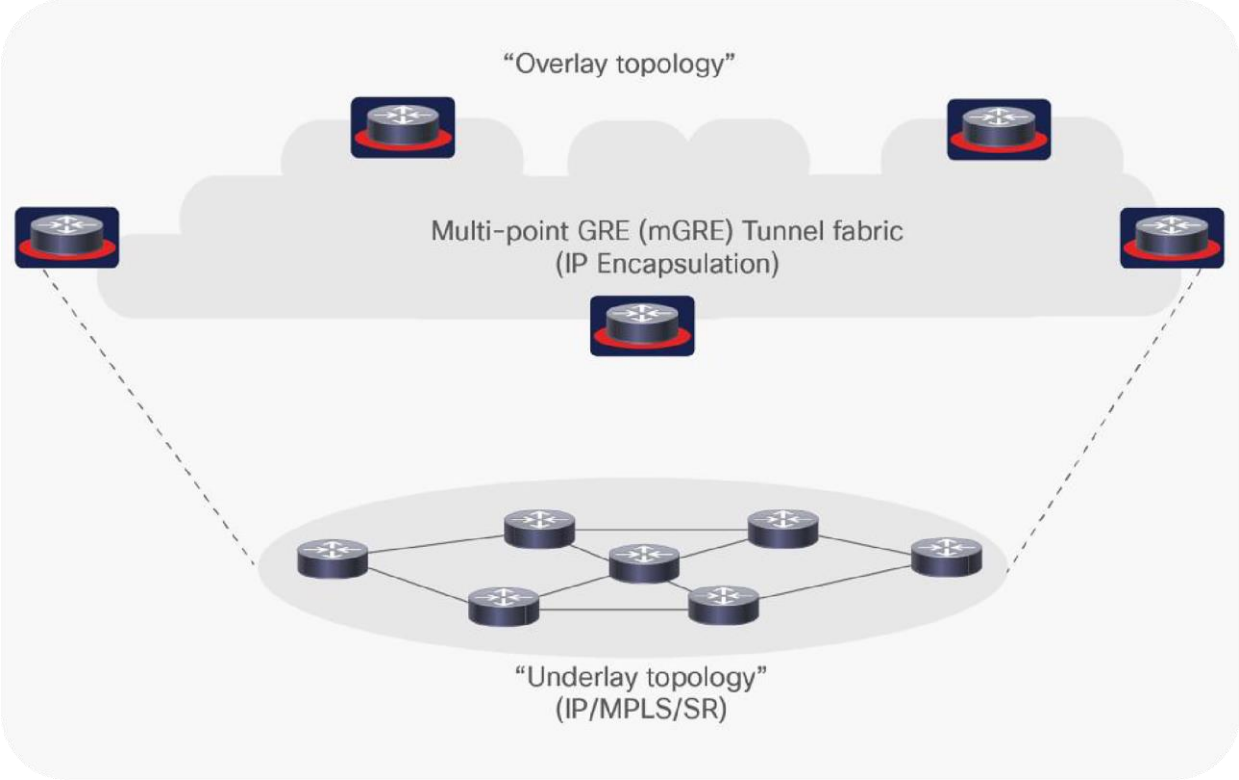

Figure 1.
Basic example of an overlay and underlay multilayer network

This paper focuses on high-speed backbone design concepts that must also integrate external encryption devices.

Complexity of two-layer architectures

To provide a quick refresher for the reader, the complexity of this two-layer architecture arises as each layer typically operates as two "ships in the night" topologies as shown in **Figure 2**. The secure IP "overlay" (shown as red) leverages its own design, protocols, and over-the-top topology. The "underlay" (shown as black) does the same, leveraging its own design, protocols, and topology.

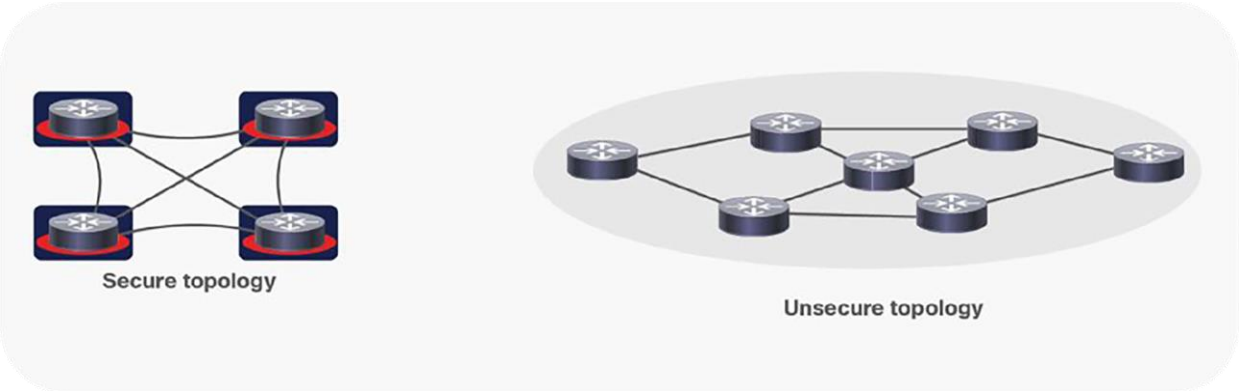

Figure 2.
Multilayer overlay/underlay "noncongruent" topology example

Both the secure and unsecure topologies together make up the overall transport network, but they are completely blind to one another, as shown in **Figure 3**, which adds an external encryption device to the example. The secured router topology (shown in red) and available paths are dictated by the unsecure (shown in black) underlay—unsecure in that it does not carry explicitly classified traffic. The challenge is that each layer is completely independent, neither has visibility of the other, and this two-layer architecture has no one-to-one relationship when a failure occurs in either layer, which poses very complex management and troubleshooting when a failure occurs.

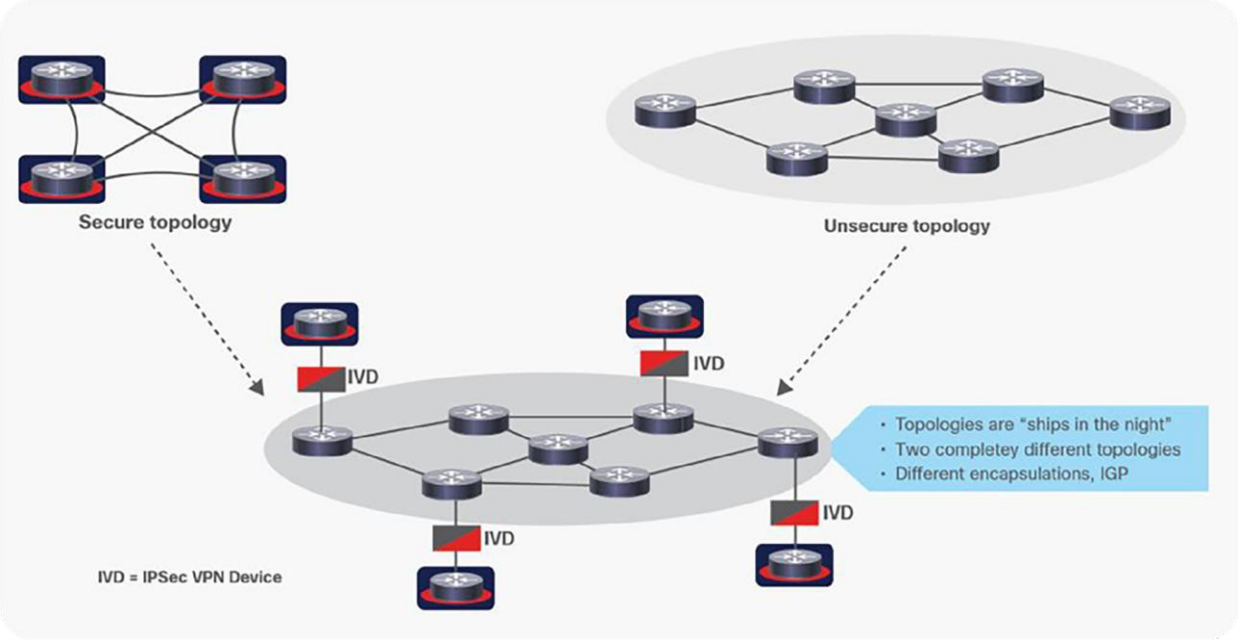

Figure 3.
Example of “ships in the night” overlay/underlay architecture

As stated earlier, these designs are analogous to commercial network topologies that leverage any overlay/underlay concepts (e.g., SD-WAN, VXLAN, GRE tunneling), specifically where there is no one-to-one alignment between the overlay and underlay. Operationally, this adds enormous complexity for troubleshooting, as any failure requires in-depth understanding of both layers, including the external encryption requirement case (IPSec VPN Device, or IVD in this example)

Introducing the secure lean core design

To overcome the specific overlay/underlay challenges that these noncongruent overlay/underlay topologies present as discussed earlier, we are introducing a new design approach, referred to as the “Secure Lean Core Design” **(Figure 4)**. The term “lean” refers to the overall simplicity of this new design that targets eliminating the complexity in the underlay, reducing the number of protocols and layers, while simplifying operations by offering pinpointed deterministic behavior for any link/node failures that occur within any layer of the topology.

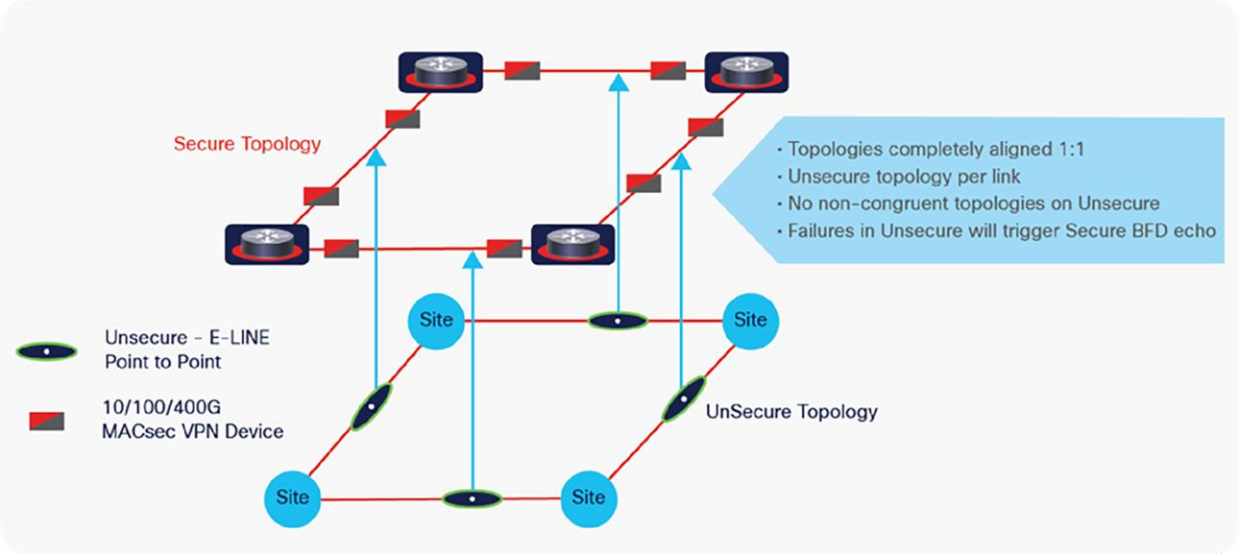

Figure 4.
Lean secure multilayer design topology

Unlike the IP-over-IP multilayer overlay/underlay design, the secure lean core design leverages Layer 2 external encryption (MVD) that offers a complete one-to-one alignment with the underlay and overlay topologies. No more noncongruent "ships in the night" behavior, this design provides operators a more simplistic and deterministic network design while leveraging their mandated external encryption needs, and doing this at 100GE+ performance.

A summary of the key components of this secure lean core design includes:

- Segment Routing (SR) as the foundational routing framework (TI-LFA, single IGP, SR Traffic Engineering (SR-TE), etc.)
- Point-to-point (E-LINE) Ethernet transport service (public/private)
- Layer 2 link-layer external encryption with the MACsec VPN device (MVD)
- Fast event detection mechanism for sub-second detection (for example, BFD)
- Use of a physical link between routing devices (no overlay encapsulations).

By leveraging the sum of these components in this new secure lean core design, the advantages are exponential in terms of providing simple deterministic behavior and providing an explicit one-to-one correlation of the secure (overlay) and unsecure (underlay) layers.

Figure 5 shows a clear depiction of how the secure lean core appears from a topological perspective, which is a very simple deterministic topology, deterministic in that with any component that fails (router, link, encryptor, COTS underlay), the resulting topology is easy to predict ahead of time without any complex calculations. Furthermore, the secure lean core design incorporates a point-to-point Ethernet/optical "underlay" transport (versus a complex noncongruent IP "underlay") that eliminates any needed complex troubleshooting of the underlay network.

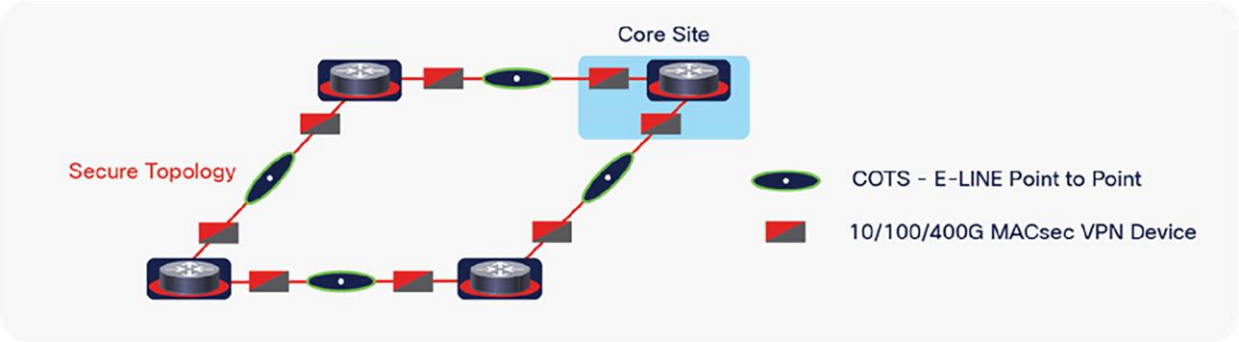

Figure 5.
Simplified secure lean core design

The advantages to this design approach cannot be overemphasized, offering consistent deterministic behavior when a failure occurs at either layer, while eliminating the need for complex troubleshooting of each network layer, independently.

The reader may ask, “What if there is a failure in the black transport and how will that event reverberate back to the red layer in order to achieve sub-second convergence?” The answer is Bidirectional Forward Detection (BFD) as shown in **Figure 6**. In the past, the event detection mechanisms in these multilayer networks were very misaligned and lacked the desired sub-second convergence goals of the designers, often causing “racing conditions” in the event detection.

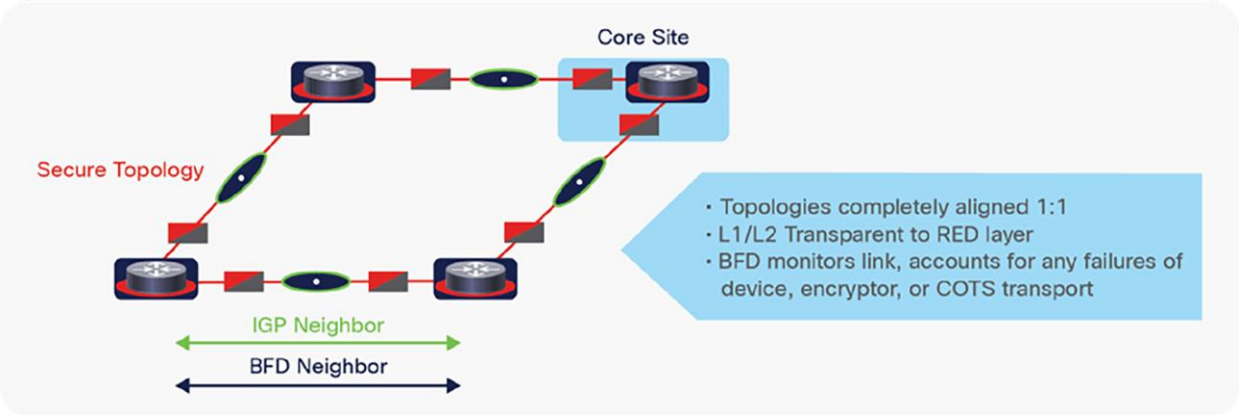

Figure 6.
Secure lean core design using bfd for sub-second event detection

This is not the case for the secure lean core design utilizing BFD, as BFD is only required on the Secure Topology links and any failures within any element of the underlay (E-LINE or Ethernet MVD encryptor) will trigger BFD on the secure overlay to declare the link unusable. When using SR in the secure lean core, fast convergence is incorporated as part of the SR framework using Topology Independent-Loop-Free Alternative (TI-LFA), which preprograms the Forwarding Information Base (FIB) before the failure happens, rapidly steering traffic to the preprogrammed backup path, and avoiding the flood-and-learn latency imposed to recalculate the shortest path.

Putting it all together

The key advantages for network services and operators of this secure lean core design cannot be overemphasized, offering a high-speed, secure, one-to-one alignment of Layer 3 to Layer 2 to Layer 1 and eliminating the overlay/underlay inconsistencies that have existed for decades in these types of secure external encryptor designs.

To summarize the secure lean core design capabilities:

- Core/edge design that delivers the feel of a single layer, while still accommodating the multiple layers required in external encryption designs
- High-speed, extremely resilient, deterministic routing core using Segment Routing
- Simplified and fewer protocols in the core, with ability to leverage a single IGP, eliminating the complex IGP-Label Distribution Protocol (LDP) sync functions found in today's MPLS backbones
- Inherently native sub-second convergence
- Consistent one-to-one mapping of the secure and unsecure transport layers
- Elimination of complex troubleshooting necessary in today's noncongruent overlay/underlay designs
- Option to link the cipher-text side of the Ethernet MVD to an optical/Optical Transport Network (OTN) multiplexer, if so desired.

Putting this all together, **Figure 7** depicts an example of the secure lean core leveraging a multiplanar topology with MVDs securing each link. A typical use case that can be leveraged in this topology with SR is the ability to create virtual topologies that maintain and enforce certain topological constraints. For example, the **BLUE** plane could provide a "low latency" transport, while the **GREEN** plane could provide a "lowest routing cost" transport. The edge router will leverage a policy and forward the traffic to the appropriate plane that meets the specific level of service. So, for example, mission video traffic would be forwarded to the "low latency" plane on the blue path.

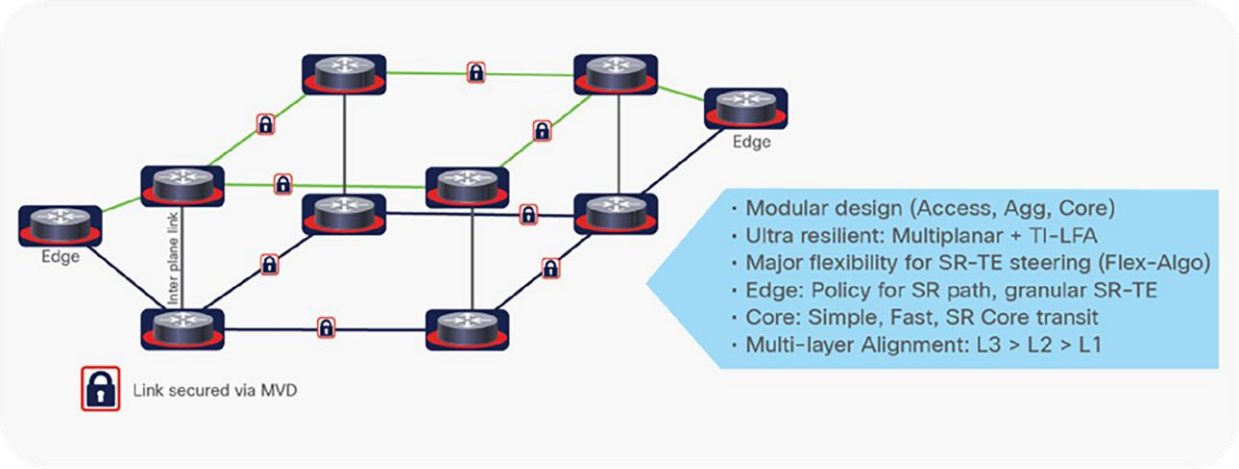

Figure 7.
Secure lean core design with MVD leveraging a multiplanar topology

Commonly, each edge device in this design (assuming redundant edge devices for high availability) has a link to both planes and, using this with the concepts in the MPBB design, extends the resilience, convergence, capacity, and traffic engineering capabilities of the secure lean core design. Regardless of which edge routing device the packet arrives on, the "match/action" of the policy can be applied and the packet forwarded to the proper plane based on the policy intent. The edge/core node combination can reside in any location, such as a private facility, private data center, or co-location center.

The scale-out options of the edge/core devices also offer various options that are outside the scope of this white paper, but can be found in detail at the Core Fabric Design overview authored by Shelly Cadora at the @XRDOCS website: https://xrdocs.io/design/blogs/latest-core-fabric-hld

Summary

For years, those customers responsible for designing and maintaining multilayer networks to support external encryption capabilities have struggled with the complexities that multilayer imposes, specifically the independent "ships in the night" nature in which the overlay and underlay operate.

Combining the innovation of the multiplanar backbone concepts, Segment Routing, and Layer 2 MACsec-based external encryption with a simple high-speed underlay all adds up to what is the secure lean core design.

By using this new secure lean core design, the advantages for operators are endless in terms of providing simple deterministic behavior in the backbone with more services, high-speed convergence, and the one-to-one correlation of the secure (overlay) and public (underlay) transport layers, while meeting the need for supporting external encryption in a carrier-class backbone design.

References

Hill, C., and Kowal, M. (2019). The Multiplanar Backbone (MPBB) White Paper.https://www.cisco.com/c/dam/m/digital/elq-cmcglobal/OCA/Assets/Federal/The_Multiplanar_Backbone_MPBB.pdf

Hill, C., and Orr, S. (2016). Innovations in Ethernet Encryption (802.1AE - MACsec) for Securing High Speed (1-100GE) WAN Deployments White Paper.https://www.cisco.com/c/dam/en/us/td/docs/solutions/Enterprise/Security/MACsec/WP-High-Speed-WAN-Encrypt-MACsec.pdf

Hill, C., and Mosher, S. (2020). Cisco Software-Defined WAN for Secure Networks - Redefining WAN Delivery in the Cloud Era White Paper. https://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/sd-wan/white-paper-c11-741640.html

Katz, D., and Ward, D. (2010). Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop) IETF RFC 5881. https://datatracker.ietf.org/doc/html/rfc5881

About the author

Craig Hill is a Distinguished Architect in the U.S. Public Sector CTO Office working at Cisco for over 27 years. His recent focus is in designing large-scale, carrier-class service delivery architectures in DoD/Intelligence and large enterprise networks that include Segment Routing, Zero Trust Architectures, SD-WAN and SASE transitions, high-speed encryption, cloud network transition architectures, and incorporating NetDevOps tools and concepts into network operations. Craig is a 26-year CCIE (#1628) in Routing and Switching, a 12-year speaker at Cisco Live, and is based out of the Cisco office in Herndon, VA.

Acknowledgements

Craig Hill would like to thank Tim Thomas, Steve Blasiol, Marc Moffett, and Matt Breneisen for their suggestions, review, and input to this paper.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-3160423-00 10/22