

BUILDING MULTI-GENERATION SCALABLE NETWORKS WITH END-TO-END MPLS

Juniper Enables Service Flexibility and Scalability
with Seamless MPLS

Table of Contents

Executive Summary	3
Introduction	3
Service Flexibility and Simplified Provisioning	3
Service Flexibility with Simplified Provisioning and Operations	3
Traffic Engineering and Deterministic End-to-End Service Restoration	3
Building Scalable Networks	4
Seamless MPLS Architecture Overview	4
Juniper's Seamless MPLS Solution	5
Decoupling Network and Service Architectures	5
Deterministic End-to-End Service Restoration	7
Failure Detection Mechanisms	8
Comprehensive OAM implementation across IP/MPLS and Ethernet Simplified end-to-end troubleshooting and performance monitoring	8
Failure Recovery and Service Repair	8
Scaling the Network with Seamless MPLS	9
Seamless MPLS and MPLS-TP Comparison	10
Conclusion	10
About Juniper Networks	11

List of Figures

Figure 1: Seamless MPLS architecture overview	4
Figure 2: Comparing traditional architecture and seamless MPLS	5
Figure 3: Connectivity blueprint for services with seamless MPLS	6
Figure 4: Decoupling network and service architecture with pseudowire headend termination (PHT) into L3 VPN	7
Figure 5: Seamless MPLS service provisioning with ACX Series and MX Series routers	7
Figure 6: End-to-end service restoration techniques with end-to-end MPLS services	9
Figure 7: Network control plane scaling with BGP LU and LDP DoD	10

Executive Summary

With the continued growth in demand for bandwidth intensive services in mobile and wireline networks, service providers are looking for ways to build an infrastructure that will not only scale to meet today's growing business needs, but also allow them to adapt and benefit from technology innovations for years to come. For both wireline and mobile networks, MPLS has become the platform of choice for packet transport and carrier Ethernet, the underlying media for service delivery. Some have even taken MPLS beyond the core to the aggregation and metro area network, and also want to deploy it in the access network to reap some of the same benefits that MPLS provides in the core. Taking MPLS to the access, and making MPLS the packet forwarding platform end-to-end across the network, requires new functionality and features and a systematic architecture that can scale to tens of thousands of nodes.

Seamless MPLS addresses the need for extending MPLS to the access network and outlines the architectural framework to build scalable and resilient MPLS networks with a flexible services delivery model. This paper describes the benefits and requirements for seamless MPLS, along with the features and functionality supported by Juniper's comprehensive MPLS portfolio to deliver a complete and flexible solution.

Introduction

There are many benefits that have led network operators to deploy MPLS in parts of the network beyond the core, deploying it in aggregation as well as access networks. However, these deployments are typically not in a single MPLS domain. In parallel, Ethernet has become the preferred cost-effective choice for building infrastructure, with increased investments being put into building Ethernet-based packet networks and services migrating away from time-division multiplexing (TDM) transport. These two trends have influenced many developments in MPLS, including ways to deliver MPLS services over Ethernet for both unicast and multicast traffic.

From a business standpoint, network operators continue to be challenged as they look for ways to make service delivery cost-effective and efficient. The need for one converged packet network to deliver all fixed and mobile services, regardless of the last mile or access technology, keeps getting stronger. These factors combined have led the industry to innovate and invest in features and functions that take MPLS to the access network and build single domain MPLS networks. Seamless MPLS is the umbrella portfolio that addresses this need, and provides the framework for taking MPLS to the access in a scalable fashion, extending the benefits of traffic engineering and guaranteed service-level agreements (SLAs) with deterministic network resiliency.

Service Flexibility and Simplified Provisioning

There are three key categories of benefits and requirements with taking MPLS to the access and building seamless MPLS networks:

1. Service flexibility and simplified provisioning and operations
2. Network resiliency with deterministic, sub-second, end-to-end convergence for services
3. Scale to the order of 100,000 nodes network-wide without compromising any of the benefits

Service Flexibility with Simplified Provisioning and Operations

Seamless MPLS architecture suggests a systematic way of enabling MPLS end-to-end between access nodes, with all forwarding based on MPLS labels. Using this approach, packets are labeled at the access network entry point and are transported as labeled packets all through the network to the receiving end. This means that all service provisioning and operations are MPLS-based. There is a clean and homogenous separation of control plane, management plane, and data plane operations throughout the network that allows decoupling the service provisioning plane from the underlying transport technology. It also makes way for optimizing and simplifying service provisioning and operations, and it brings us to the first benefit of seamless MPLS. The number of service provisioning points can be minimized and as we will see further on in this paper, it will be possible to decouple service topology from the network topology. This architecture offers service flexibility beyond what is available with today's heterogeneous networks.

Traffic Engineering and Deterministic End-to-End Service Restoration

A second key benefit is MPLS traffic engineering capabilities and deterministic end-to-end service restoration. The ability to traffic engineer based on real-time network attributes supports strict SLAs with guaranteed service availability, as well as sub-second restoration of services with fast reroute mechanisms in the event of link or node failures. Working together, these capabilities enable providers to build highly resilient networks. The goal of seamless MPLS is to extend the same benefits end-to-end across the access network.

Building Scalable Networks

The idea of seamless MPLS elevates the requirement for scale. Can MPLS scale to the magnitude of today's largest networks with all aggregation and access nodes inclusive? Can seamless MPLS scale as the demand for network bandwidth, services, and reach continues to rise? While the WANs or core typically consist of 100 to 1,000 Layer 3 nodes, the metro access and aggregation networks may each contain the same number of nodes with up to hundreds of access and aggregation sub-networks in one single network. In effect, we are talking about a network that may scale to over 100,000 nodes. Clearly this requires some new thinking and innovative techniques to deliver this scale.

The remainder of this paper describes the seamless MPLS architecture, features, and functionality in the Juniper Networks® Junos® operating system MPLS portfolio, which enables a smooth migration from traditional networks to a scalable, end-to-end MPLS network.

Seamless MPLS Architecture Overview

Seamless MPLS architecture describes a systematic way of enabling MPLS end-to-end in a single domain. It enables complete virtualization of network services with service origination and termination at the access nodes.

The seamless MPLS architecture¹ describes the several types of “nodes” in a network, each with a different function. A physical device may combine several of these functions. Conversely, a single function may require multiple physical devices for its realization.




Figure 1: Seamless MPLS architecture overview

There are several types of nodes in the network as illustrated in Figure 1:

1. **Access Node (AN):** These are the first (and last) nodes that process customer packets at L2 or higher. Examples include digital subscriber line access multiplexors (DSLAMs), multi-tenant units, passive optical network (PON) optical line termination devices (OLTs), and cell site gateways (CSGs).
2. **Service Node (SN):** These nodes apply services to customer packets. Examples include L2 provider edge (PE) routers, L3 PE routers, Broadband Network Gateways (BNGs), peering routers, video servers, base station controllers, and media gateways.
3. **Transport Node (TN):** TNs connect ANs to SNs, and SNs to SNs. Ideally, TNs have no customer or service state.
4. **Border Node (BN):** These nodes enable inter-region packet transport (e.g., area border router and AS boundary router).
5. **Service Helper (SH):** SHs enable or scale the service control plane. SHs do not forward customer data. Examples include service route reflectors, policy and control enforcers, RADIUS and AAA devices, and session border controllers.

A physical device may, of course, play multiple roles. For example, an AN may also be an SN—or, an SN may double as a TN. SHs may be embedded in SNs. It is often useful to “virtualize” a physical device that plays multiple roles (using the notion of logical routers) so as to minimize the impact of one role on another, both from a control plane and a management point of view.

¹ Seamless MPLS www.juniper.net/us/en/local/pdf/whitepapers/2000316-en.pdf.

Juniper's Seamless MPLS Solution

Juniper's solution for seamless MPLS is architected around three primary dimensions—flexible service delivery, end-to-end deterministic service restoration, and massive scale to support a network on the order of 100,000 nodes. Powered by Junos OS with over a decade of experience in carrier-class MPLS deployments, the solution is designed and developed to address these three service dimensions effectively and efficiently. Juniper's seamless MPLS portfolio is standards-based and can interoperate in a standards-based, multi-vendor solution. The solution includes various features and functionality that providers can enable incrementally to deliver end-to-end MPLS services in a scalable and resilient fashion.

Decoupling Network and Service Architectures

Traditional network deployments are built with an implicit tight coupling between the network nodes, the underlying transport technology, and the services delivered over the network. Typically, services are provisioned in multiple segments. A VLAN or an L2/L3 circuit is carried through an access/aggregation infrastructure and mapped to a VPN routing and forwarding table (VRF table) or a virtual private LAN service (VPLS) instance at the PE router in a point of presence (POP). The second service segment is then initiated at the ingress PE router. The service is carried across the core to the other end, the egress PE router, where this service segment terminates and gets mapped into the access/aggregation infrastructure on the receiving end, perhaps as the last segment of the service. This model provides limited flexibility in provisioning, as it is tightly coupled with the topological placement of the network nodes and operationally one has to deal with multiple technologies for troubleshooting and fault recovery.

With seamless MPLS, the idea is to provision the service end-to-end and minimize the number of provisioning points. The service provisioning is inline with the network architecture, maintains simplicity in the access network, and relies on increased capabilities and intelligence on the service nodes (traditional PE routers). At the same time, it also simplifies operations and makes efficient use of network resources by reducing the number of provisioning points and relying on a single MPLS-based forwarding scheme in the data plane.

Figure 2 compares service provisioning in a traditional network versus that of a seamless MPLS single domain model. As illustrated, there is a single provisioning point per connection with the notion of a single label-switched path (LSP) across the access nodes in a network-wide single MPLS domain. Minimizing the number of service provisioning points further enables decoupling of the services architecture from the underlying topology and transport.

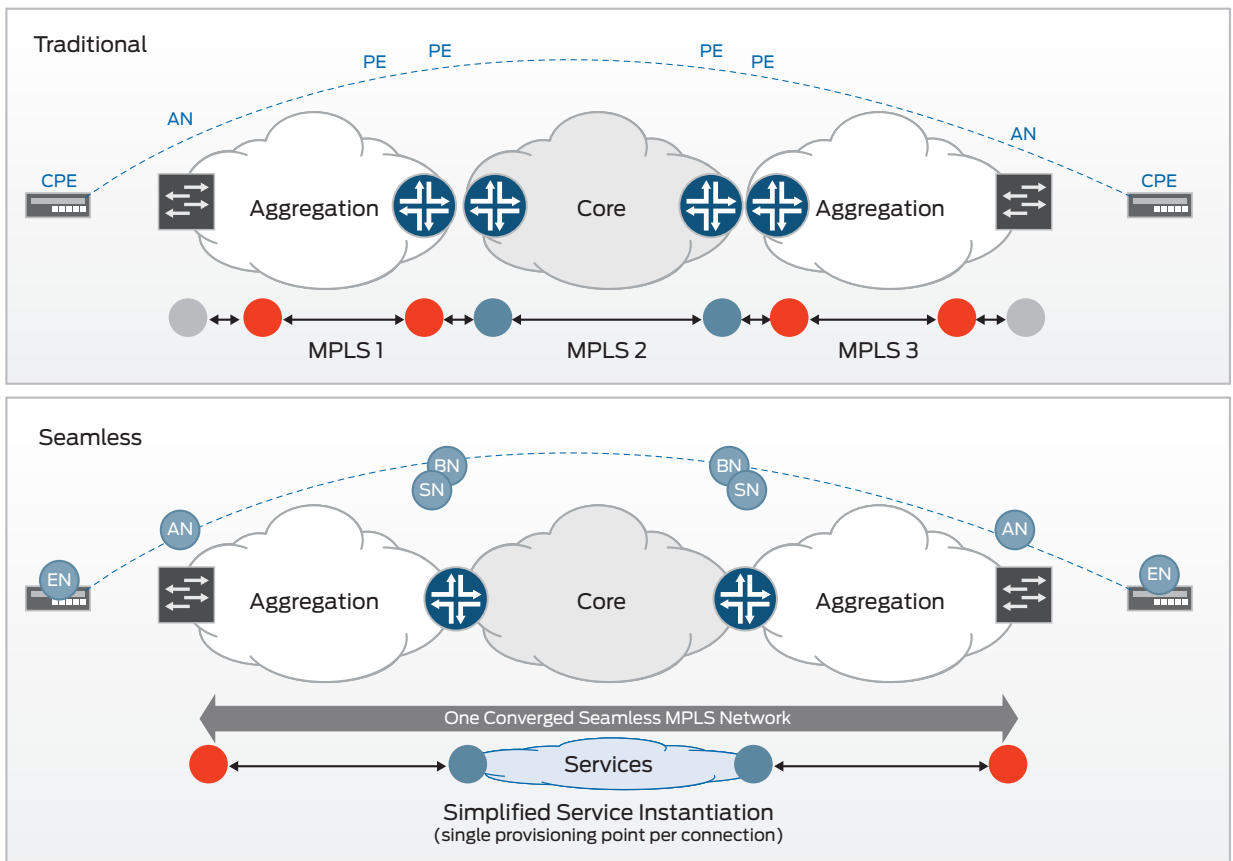


Figure 2: Comparing traditional architecture and seamless MPLS

The seamless MPLS architecture can easily be applied to various services provisioned in today's networks. Some use case examples of the application for commonly offered fixed and mobile network services are illustrated in Figure 3, a connectivity blueprint for services with seamless MPLS. The services illustrated include:

1. Basic point-to-point connectivity services: In this model, a point-to-point pseudowire is provisioned end-to-end between access nodes. The access nodes play the role of service nodes here, and there is a single LSP provisioned end-to-end.
2. L2/L3 business services in both centralized and distributed models: In this case, a pseudowire is provisioned at the AN and is then transparently mapped to a service at the SN. The service at the SN can be an L2 VPLS instance or L3 VPN. The mapping of a pseudowire to L3 VPN is enabled by a new functionality called pseudowire headend termination (PHT).
3. End-to-end MPLS service provisioning for mobile backhaul applications: For mobile backhaul applications, a pseudowire can be provisioned at the AN, and this can be carried over an MPLS RSVP or LDP tunnel in the aggregation, metro, and core. Traffic separation can be enabled at the AN, and traffic can be steered over tunnels per traffic class to ensure that the appropriate latency and bandwidth is offered to mission critical control/management and voice traffic.




Figure 3: Connectivity blueprint for services with seamless MPLS

As mentioned earlier, a key functionality that enables service flexibility in this area is the ability to provision pseudowires in the access and map these into L3 VRF instances with PHT to carry them over MPLS label transport all the way through to the receiving end. Strict SLA is maintained by classifying traffic on the pseudowire itself and managing it through the core to the other end of the network.




Figure 4: Decoupling network and service architecture with pseudowire headend termination (PHT) into L3 VPN

Since the service is initiated as an MPLS pseudowire from the origination point at the AN, any topological changes in the access can be easily made without having to completely re-provision the service layer. This can be a significant operational asset to mobile backhaul access, for example, where re-parenting of cell site routers to a different base station controller/radio network controller (BSC/RNC) is a common occurrence.

With Juniper’s vision of one network, many services, it has recently introduced Juniper Networks ACX Series Universal Access Routers, the platform that can provide a single access infrastructure for many types of services in the last mile and for both fixed and mobile access. The ACX Series brings the simplicity and ease of provisioning MPLS pseudowires from the access node, while offering the necessary operational intelligence to separate traffic and steer over specific MPLS tunnels, support strict SLAs, and provide an enhanced end user experience with its precise timing and synchronization capabilities. Figure 5 illustrates the use of pseudowires in the access infrastructure. Here, the ACX Series is the cell site router where the pseudowire originates and the Juniper Networks MX Series 3D Universal Edge Router is the service node in the mobile packet core connecting to the mobile control network. Pseudowires from the ACX Series can be mapped into Layer 3 VPNs at the service node. Junos OS makes efficient use of inline tunneling capabilities on Trio chipset-based forwarding cards on the MX Series platform to deliver this powerful capability.




Figure 5: Seamless MPLS service provisioning with ACX Series and MX Series routers

Deterministic End-to-End Service Restoration

Over the years, Juniper has pioneered a rich suite of IP/MPLS features that are used to build resilient carrier-class networks and deliver many services. Resilient networks must provide a deterministic end-to-end service restoration, and there are two broad categories of functions that help achieve this. The first set of functions includes ways to enable speedy detection of performance degradation events and location of failures. The second set of functions is comprised of the appropriate recovery actions needed to reroute and restore services.

Failure Detection Mechanisms

There are various failure detection mechanisms available within the Juniper MPLS solution. L2 failure detection relies on Ethernet Operation, Administration, and Maintenance (OAM) capabilities, as well as integration of Bidirectional Forwarding Detection (BFD) mechanisms with LSP and pseudowires (BFD-VCCV). For L3 fault detection and to test data plane consistency of pseudowires, both single hop and multi-hop BFD (RFC 5883, RFC 5884) are supported for BGP sessions and targeted LDP sessions. Table 1 provides a view of the Juniper solution's comprehensive OAM capabilities across IP/MPLS and Ethernet.

Table 1: Juniper OAM portfolio for IP/MPLS—fault detection and isolation.

Technologies	Fault Detection	Fault Verification	Fault Localization	Fault Notification	Performance Loss Ratio	Performance Fwd Delay	Performance Fwd Delay Variation
Ethernet							
802.3ah	Yes	Yes	Remote LB	(Critical) Link Events Remote Fault Indication	No	No	No
E-LMI	N/A	N/A	N/A	Yes	N/A	N/A	N/A
802.1ag	CC	LB	LT	RDI	No	No	No
Y.1731	CC	LB	LT	RDI, AIS	ETH, LM	ETH, DM	ETH, DM
IP/MPLS							
LSP	LSP-BFD	LSP-Ping	LSP-Ping/TR	LDP, RSVP	RPM	RPM	RPM
PWE3	VCCV-BFD	VCCV-Ping	VCCV-Ping/TR	BGP, tLDP, VCCV-BFD	RPM	RPM	RPM
L2VPN	BGP, VCCV-BFD	BGP, VCCV-Ping	BGP, VCCV-Ping/TR	BGP	RPM	RPM	RPM
IPVPN	BFD, IGP, MP-BGP	Ping	TR	IGP, MP-BGP	RPM	RPM	RPM

Comprehensive OAM implementation across IP/MPLS and Ethernet simplified end-to-end troubleshooting and performance monitoring.

Failure Recovery and Service Repair

The choice of recovery mechanism used to restore services is often dependent on the location and type of failure on the network. To address failure recovery effectively and efficiently, Juniper's MPLS portfolio includes an extensive suite of features and functionality. In order to achieve sub-second convergence subsequent to a network failure, the first line of defense is to use local repair techniques and precomputed paths to reroute around the failures. These are implemented on top of nonstop active routing (NSR)-enabled control plane protocols. Local repair techniques within Juniper's solution include loop-free alternate (LFA) support for ISIS, OSPF and LDP. Link and node protection can also be enabled with RVSP-TE MPLS fast reroute and provide deterministic service restoration

To guarantee strict SLA and ensure maximum network and service uptime, networks are designed with both redundancy and resiliency techniques. Juniper platforms, all powered by one Junos operating system, are built to meet carrier-class requirements. On top of the base redundancy and resiliency built into Junos OS and platform hardware, service-specific features are developed to support deterministic, end-to-end service recovery. These include ability to provision redundant pseudowires from the access node to the service node and provisioning backup logical tunnels on the forwarding cards on the service node. Figure 6 provides a snapshot of comprehensive failure recovery mechanisms available with the Junos OS MPLS portfolio.

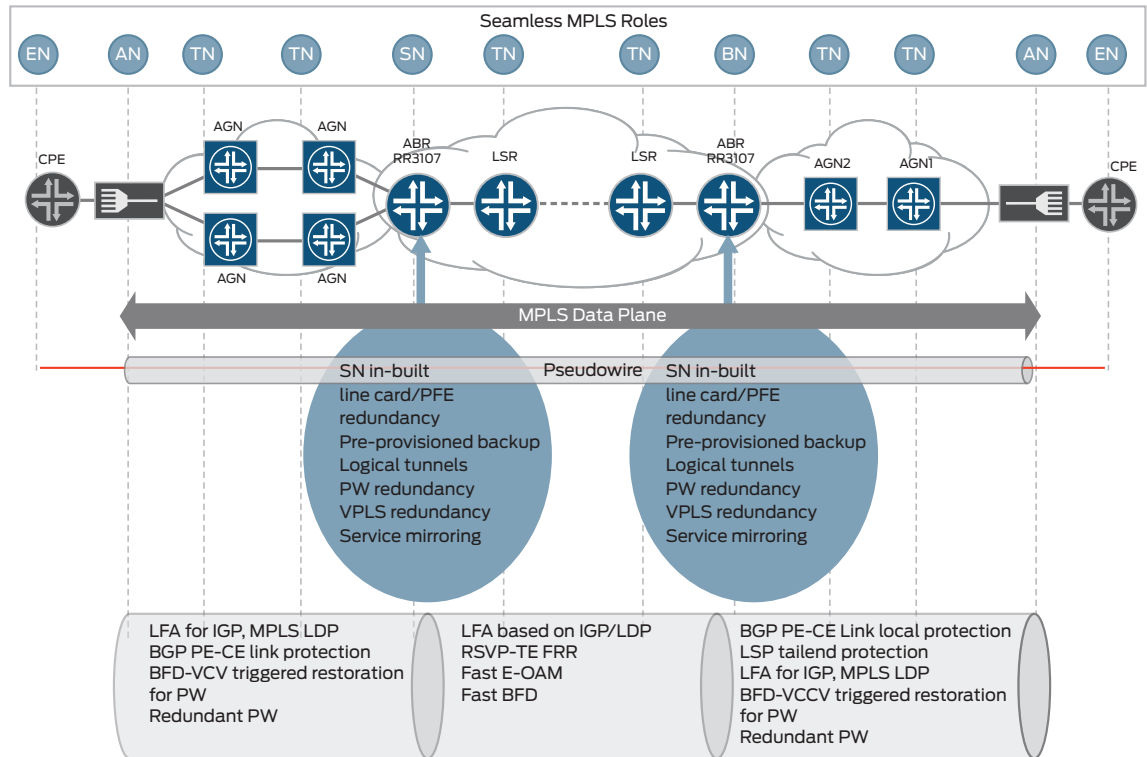


Figure 6: End-to-end service restoration techniques with end-to-end MPLS services

Scaling the Network with Seamless MPLS

As discussed early in this paper, the seamless MPLS architecture must scale to support end to-end MPLS in a network on the order of 100,000 nodes. We build upon the merits of existing control plane functionality and deliver a scalable and operationally intelligent network. The key to scaling is to create hierarchical end-to-end LSPs by distributing the right amount of routing intelligence with BGP Labeled Unicast (BGP-LU). Hierarchy is created by segmenting the network into regions, running closed interior gateway protocol (IGP) within the regions, and restricting inter-region IGP communication. All inter-region control plane information is shared via BGP-LU¹.

In addition to the use of BGP-LU to create LSP hierarchy, to keep the access nodes cost-effective and functionally simple but still operationally intelligent, LDP DoD (Downstream on Demand) (RFC5036) with static routes can be used. Given that the simplest access router infrastructure may only use static routes, LDP DoD ensures that a label table is created only for the specific routes installed in the RIB table. The general MPLS use of LDP DU (Downstream Unsolicited) advertises labels for all routes in the routing table. This would be unnecessary for the simple and cost sensitive access router. LDP DoD specification was originally written for ATM and Frame Relay-based MPLS label-switching routers (LSRs), and now this specification is being used for Ethernet-based MPLS access. For some access topologies, an IGP may be used, but both principles of effective use of BGP-LU and LDP DoD still apply and keep the access node functionality simple and cost effective. Figure 7 illustrates end-to-end control plane and service plane views with BGP-LU and LDP-DoD to scale a seamless MPLS network.

¹ Refer to *Network Scaling with BGP Labeled Unicast Design and Configuration Guide* at <http://www.juniper.net/us/en/local/pdf/design-guides/8020013-en.pdf>.




Figure 7: Network control plane scaling with BGP LU and LDP DoD

Seamless MPLS and MPLS-TP Comparison

There are often questions around the interaction and applications of MPLS Transport Profile (MPLS-TP) and seamless MPLS. Seamless MPLS is an architectural enabler for scale, resiliency, and service flexibility. Seamless MPLS allows multiple choices for the underlying transport—dynamic signaling, static signaling, and static provisioning. One of the choices for transport is MPLS-TP. There is a misconception that MPLS-TP primarily translates to statically provisioned MPLS LSP. On the contrary, dynamic signaled LSPs are also part of the MPLS-TP profile. Static provisioning creates a dependency on a network management system, which is usually proprietary and can tie the operator into a single vendor environment. Dynamic signaling, on the other hand, uses standards-based protocols that allow multi-vendor interoperability. The dependency on a proprietary network management system is reduced. A network management system can be used to enhance the overall provisioning experience, even with dynamically signaled network architecture. The main characteristics of MPLS-TP include a rich OAM feature set for both fault and performance management. The OAM extensions introduced via MPLS-TP can be effectively used in a generic seamless MPLS context.

Conclusion

The use of MPLS in the access infrastructure can help network operators achieve their operational efficiency goals and successfully build one converged network for many services. Juniper’s seamless MPLS solution is designed to provide service flexibility and deterministic end-to-end service restoration at large scale, and it can be incrementally enabled in current networks to support multi-generation service needs.

Juniper’s seamless MPLS portfolio is standards-based and can interoperate in a standards-based, multi-vendor solution. With Juniper’s vision of one network, many services, Juniper Networks ACX Series Universal Access Routers bring simplicity and ease to the provisioning of MPLS services from the access network, while offering the necessary operational intelligence to traffic engineer over specific MPLS tunnels. The migration to MPLS represents a key trend in the access market; the need for more intelligence and higher intelligence in the access layer plays into the strengths of MPLS and the rise of converged networks.

For More Information

For more information, please refer to the following documents:

Seamless MPLS Whitepaper

www.juniper.net/us/en/local/pdf/whitepapers/2000316-en.pdf

Network Scaling with BGP Labeled Unicast Design and Configuration Guide

<http://www.juniper.net/us/en/local/pdf/design-guides/8020013-en.pdf>

LDP Downstream-on-Demand in Seamless MPLS

<http://tools.ietf.org/html/draft-beckhaus-ldp-dod-00>

Juniper Networks ACX Series Universal Access Routers Datasheet

www.juniper.net/us/en/local/pdf/datasheets/1000397-en.pdf

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803


EMEA Headquarters

Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2012 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

2000452-001-EN Jan 2012

 Printed on recycled paper