# MPLS Packetized Transport for Next-Generation Networks

MPLS-based Packet Transport Ideal for Public and Private Network Resource Consolidation

## Table of Contents

## Executive Summary

Over the years, public and private network operators worldwide have made significant investments into building packet infrastructure. Networks have evolved to meet the ever-increasing traffic demands for packet services and to offer many parallel services on one network. Multiple service delivery networks have benefited from network consolidation. The general trends of increasing packet traffic versus stunted growth in TDM traffic, uncoordinated network investment at the transport and packet layers, and technology leaps in speeds and feeds in the packet network have prompted network operators to consider out-of-the-box solutions. One such solution is the migration of the underlying circuit-switched transport network to a packet-based infrastructure.

Today's network architecture with a data network overlay onto the transport network has introduced significant network inefficiency and increased OpEx for network operators. The next-generation converged network architecture must deliver higher network efficiency and must continue to deliver a network to support high availability, bandwidth guarantees, and predictable fault tolerance for mission-critical applications. In the next wave of convergence, there are multiple solutions available to network architects and operators. Any of the viable solutions must also meet the ever-increasing end-user services' scale requirements in a budget-constrained environment.

This paper discusses the new network requirements and Juniper's solution to seamlessly migrate toward the next-generation converged packet transport network.

## Introduction

Networks worldwide have evolved to support a wide range of end-user services in a fiercely competitive market. Network traffic driven by packet-based services has increased considerably (see Figure 1). Across all market segments, content-rich packet applications continue to drive network bandwidth growth while TDM- based services such as legacy voice, circuit emulation access for legacy non-packet applications, and leased-line (T1/E1) access for business services see limited growth. With cost and efficiency as key business drivers, network operators are pressed to support high-bandwidth services with virtually zero network downtime[1].

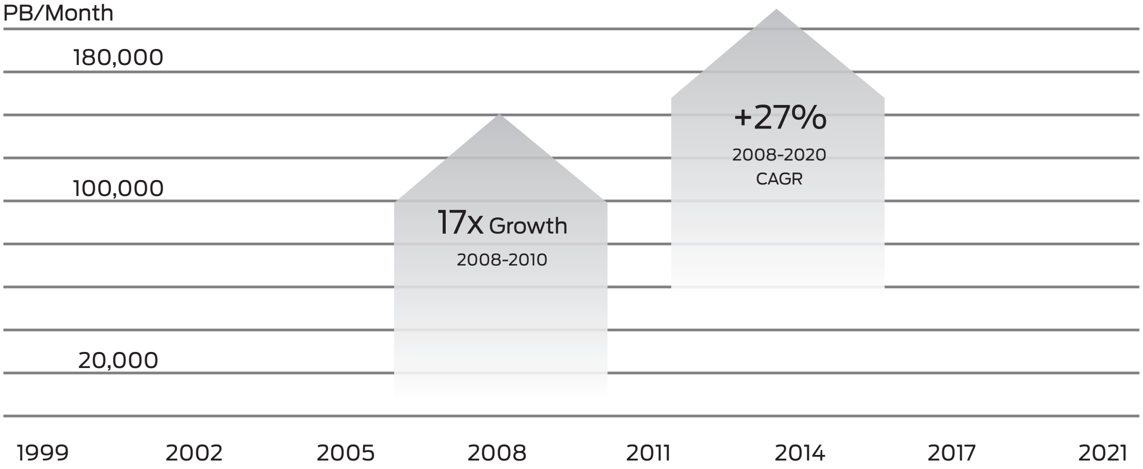

Figure 1: Exploding bandwidth growth fueled by packet applications

Advancements in information technology have transformed business processes, and IT infrastructure has become the operating backbone in a global operations environment. For instance, a segment of the federal market has a major IT transformation initiative to enhance its communication, networking, and information-sharing abilities to support the "warfighter" in a global environment. Key infrastructure and applications across the enterprise in this sector include:

- Cloud computing, virtualization schema, backup and disaster recovery
- Global connectivity for campus locations worldwide
- Streaming HD video applications
- Video surveillance, cybersecurity, forensic and big data analytics
- VoIP, IPTV, and mobile collaboration

[1] Converged Packet Transport (www.juniper.net/us/en/local/pdf/whitepapers/2000402-en.pdf)

Historically, networks first started out to deliver TDM services primarily for voice and non-packet traffic and then progressed over time to deliver data applications over Frame Relay, ATM, and then Ethernet. These networks have seen multitudes of architectural upgrades to support new world packet-based applications. However, packet networks were built over the transport infrastructure with circuit-switched SONET/SDH technology, originally designed to transport voice traffic. On the other hand, packet networks have also evolved over time to support TDM-like robust transport and deliver transport-like SLAs for bandwidth-intensive applications. As a result, SONET/SDH in the metro is now being replaced by high-speed Ethernet rings without any compromise in SLA metrics or performance. However, the convergence to an all- packet network has yet to impact the transport core in a meaningful way, thus hampering overall network and operational efficiency.

With the commercial availability of coherent 100G technology for long haul, network operators are at the crossroads of another major investment decision—an iterative and incremental bandwidth upgrade to the legacy installed base or a purpose-built converged packet core. There are multiple technology solutions available to facilitate the packet migration— for example, OTN switching, packet optical convergence, IP/MPLS, and MPLS-TP. This paper takes a close look at the networks as they exist today and how Juniper Networks' solution can fit into the transport operational model and pave the road to future-proof the core for years to come.

## A Close Look at Today's Networks

Today's networks are an overlay of data network layers stacked on top of circuit-switched transport layers with an underlying optical/physical layer (see Figure 2). The data network layer is packet based and interfaces with the back-end servers and front-end client interfaces. The optical layer provides the physical connectivity and accounts for the infrastructure needed to connect the transport and data layer equipment across the WAN. The transport network layer provides reliable and robust connectivity for the data network layer while using fixed frame channels with SONET/SDH technology. Over the years, the workflow model for transport networks has been well defined and deeply integrated along intradepartmental and interdepartmental boundaries. An integral part of the workflow is the use of an out-of-band management plane to address all circuit provisioning needs and an extensive NMS/OSS system to address network planning and design, each with clearly defined role-based access control. The workflow accounts for the complete work cycle, from receiving a customer sales order to service identification, service creation, circuit provisioning, and integrating associated OAM for the service[1].

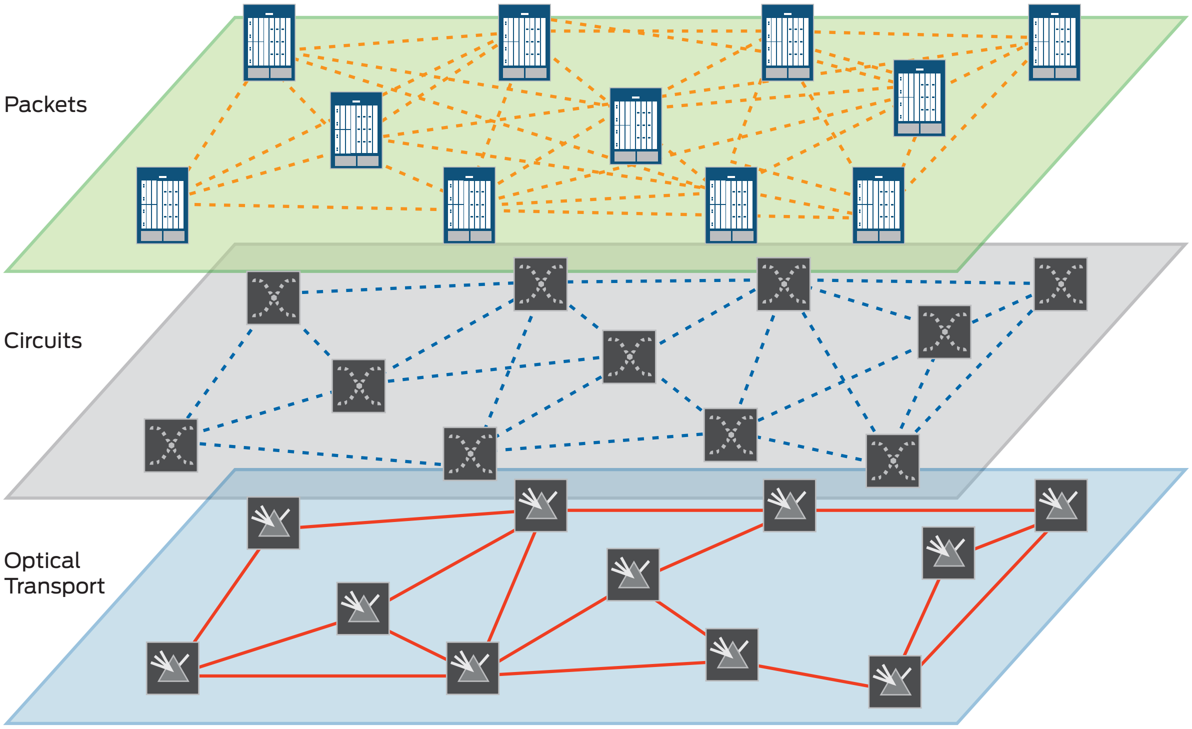

Figure 2: Current and legacy network architecture with data network and transport network layers stacked on the optical transport layer.

[2] www.juniper.net/techpubs/en_US/junos-mobility12.1/topics/concept/cac-mobility-overview.html

## Key Elements of TDM Services Provided by the Transport infrastructure

Let's take a look at key characteristics of current transport networks today. Circuit-switched transport networks were originally designed to carry TDM services such as voice traffic. While TDM services continue to see limited growth, they are still a significant source of traffic and must be seamlessly carried forward through the converged packet core migration. The services and features that the carrier-class transport infrastructure provides to its higher packet layers include the following:

**Deterministic Path with Bandwidth Guarantee**—A transport core provides a connecting "pipe" between two routers over long haul or UHL. These pipes are connection-oriented, constant bit rate circuits. And unlike data networks, there is never overprovisioning of bandwidth. Circuits are provisioned with predetermined deterministic primary and secondary paths, and network bandwidth is reserved. With bandwidth reserved, quality of service (QoS) is predetermined and predictable. L2 or L3 data flows are "mapped" to these circuits via payload encapsulation.

**Reliability and Availability**—Transport networks have set the precedence for carrier-class reliability and availability standards. Based on SONET/SDH framing and associated self-healing properties in the event of a failure, circuits can switch over from primary to a backup path in under 50 ms. L2 and L3 services have an independent failover mechanism based on MPLS fast reroute (FRR) protocols and divergent physical paths.

**Robust OAM (Operation, Administration, and Maintenance)**—OAM information is used to detect network faults, measure network performance, and trigger protection mechanisms. SONET/SDH support the famous "five nines carrier- class availability and reliability" with the OAM information carried in the SONET/SDH frame overhead. Link-layer management is carried inline with the SONET/SDH frame. Per-hop loss of signal (LOS) signal detection, alarm indication signal (AIS), and signal degrade information are all available inline with SONET/SDH.

## Challenges with the Hybrid Packet and Transport Network Model

While the transport network provides the robustness, availability, and deterministic bandwidth guarantee, there are inherent inefficiencies introduced in the hybrid model of packet and transport networks as they exist today.

### High CapEx and OpEx

The base transport infrastructure and the overlay packet networks (see Figure 2) have evolved in parallel over the last decade to meet increasing bandwidth and stringent latency requirements for many applications. This has resulted in significant investments by network operators across the two tracks for service-based platforms, as well as OSS for OAM and SLA management. Furthermore, since the packet core and transport core are disjoint networks and managed independently (ships in the night), there is no automated coordination or unified control between the layers of the two networks. Loss of intelligence or lack of exchange of coordinated updates in turn results in loss of efficiency and ineffective use of forwarding plane resources. Loss of signal or frame is not directly linked to bit-level errors in a unified control plane, thus blinding the operator during a failure event.

### Inefficient Network Resource Utilization

With the majority of the traffic generated by packet applications, transport over a circuit-switched core results in inherently low utilization of resources. First, IP packets are of varying packet sizes and subsequent use of a fixed-frame TDM circuit results in inefficient packing and inefficient utilization of the TDM circuit. Second, data applications are bursty in nature and to accommodate bursts, bandwidth provisioning in a circuit-switched core is designed around peak rates instead of the average rate. Since bandwidth provisioning in a circuit-based core is done around the peak, higher-order "pipes" are required. This results in overprovisioning and underutilization of network resources, significantly increasing costs.

Also, as is well known with SONET, the IP routers have to maintain logical circuits between every source and destination. This means many—perhaps hundreds, thousands, tens of thousands of—adjacencies are maintained per router, resulting in larger memory structures and complex lookup algorithms, creating an n-squared problem. This logical mesh is still required with next-generation TDM technology—that is, OTN switching, failing to reduce the control plane burden and solve the n-squared problem.

Another aspect to consider is while OTN provides the flexibility of carrying Ethernet and TDM/SONET/SDH client interfaces in the core, in a lot of cases OTN switching also ends up shifting cost from the core to the edge. Low-order channelized interfaces provide access-purposed circuits and are groomed at the edge of the network. These low-order channelized interfaces add higher cost per bit for the end-to-end solution.

## Building the New Efficient Network

Network operators are looking at ways to arrive at a converged network to achieve better network utilization and efficiency. The challenge is finding the most optimal way to accomplish this. This means maintaining network uptime and customer SLAs in the migration process, which in itself is a medium-to-long-term multiphase process. The key building blocks of today's transport networks are the operations workflow model and the centralized OSS. The workflow model has evolved over time and is tightly integrated to maintain high-level network security. Similarly, any new network element must also go through a rigorous integration with the OSS solution. The next-generation converged architecture must integrate with both of these key elements. The services driving the traffic growth are primarily packet-based, and over the years packet networks have evolved to support transport-like characteristics at much higher bandwidth (100G and beyond) to meet growing future needs. Packet-based transport architectures address the inefficiencies introduced by the overlay model and improve the network resource underutilization by eliminating the need to send variable-size packets into fixed-size frames.

### Building a Hierarchical Network

Operators who were in the business of transport services have built TDM-based networks to offer these services. Most of them are augmenting the transport services with more useful and value-oriented services like edge routing and VPNs for any-to-any connectivity. Operators with a legacy transport background are naturally inclined to build a pure transport core for interconnecting the edge routers and services nodes as well. However, building a core routing layer, which aggregates all of the edge routers in a given POP, not only allows operators to get away from provisioning and maintaining individual TDM circuits between every possible pair of edge nodes, but it supports an improvement in the utilization of underlying network infrastructure. Improvements are facilitated by statistical multiplexing between the edge routers from POP to connecting POP. In addition, statistical multiplexing reduces the control plane and networking complexity at the edge, as each edge router node is only required to maintain the protocol adjacency with its own core node, thereby reducing the total number of adjacencies. This is a clear solution to the n-squared problem seen with legacy TDM networks.

### Why Converge Packet and Transport Networks

Ethernet has been the consistent data link layer for LANs for decades. The simplicity, low cost, and ease of operation of Ethernet led to the development of carrier-grade Ethernet, which is now replacing packet over SONET (POS) links. Higher-order bandwidth in the core is delivered by 100G optical transport and 100G Ethernet followed by 400GbE and 1 Tbps Ethernet, as well as coherent optics that support packet-optical integration for ultra long-haul transport. Clearly, the rate of innovation is faster for Ethernet than SONET, which is stuck at OC768 (40G), which was released to the market many years ago. Encryption technologies used by the Federal Government are also somewhat stuck, as the current Type-1 Encryptors are maxed-out at 40G to support OC768 with no near-term roadmaps to support 100GE Encryption.  After the transition to Ethernet is made, Encryption technology will likely accelerate.

Network operators have more choices in the packet technologies arena with Ethernet and MPLS-based pseudowires, which can replace TDM circuits, obviating the need for legacy SONET switching. This raises questions as to why one would want to repeat the bandwidth inefficiencies of TDM with OTN switching. The new network is Ethernet directly over the optical transport system (OTS).

In specific cases where Encryption requirements exist, an inflection point involving both aspirations for 100G-based transport will have to be balanced with access to 100G Encryption technology.  This situation may lead to a continued and necessary spend on legacy and inefficient TDM equipment.

### How to Integrate Legacy TDM Services

Now, OTN switching lets the operator switch and multiplex both legacy SONET and new Ethernet circuits over the same DWDM (dense wave division multiplexing) wavelengths. How does packet transport deal with legacy TDM?

It deals with legacy TDM in terms of circuit emulation over IP services for T-1, E-1 (which doesn't address higher-speed data rates) and serial packet transport, which does not provide an underlay for legacy TDM circuits. Service providers are choosing to build packet transport parallel to TDM transport. The idea is to move TDM circuit customers to an Ethernet pseudowire service over time, reducing use of the TDM network and allowing for eventual decommissioning.

### How to Guarantee Pseudowire Bandwidth

Another important issue is how to guarantee pseudowire bandwidth without OTN switching or TDM technology. While mapping a pseudowire to an OTUx container guarantees bandwidth, it's no better than SONET/TDM switching in terms of bandwidth efficiency. Bandwidth reservation with RSVP-TE, label-switched path (LSP) auto policing, pseudowire connection admission control (CAC)[3], and QoS profiles allows for bandwidth guarantees from the edge through the core where necessary, and bandwidth sharing where it makes sense—all coexisting on the same packet transport network.

[3] www.juniper.net/techpubs/en_US/junos-mobility12.1/topics/concept/cac-mobility-overview.html

## MPLS-TP and Native MPLS

MPLS-TP, as a subset of MPLS with certain extensions, was designed to meet transport requirements in a simplified packet-switched environment. There are applications in the access domain such as mobile backhaul that can leverage MPLS-TP due to the expected non-diversity of physical and logical pathways which would not require a dynamic signaling protocol, e.g. manual configurations are adequate and preferable.  However, in the backbone, where the data core is largely based on packet routers and is designed for substantial physical and logical path diversity, the use of native MPLS with dynamic control helps scale the end-to-end network to any required size. Juniper's vision of "one converged network" with seamless MPLS end to end is designed to support 100,000 network nodes and deliver subsecond (under 50 ms) end-to-end restoration of services.[4]

The path to migration is often a multi-stage process and requires multiple solutions to coexist. Native MPLS is a proven versatile platform for multiple cutting-edge services on one consolidated network. Juniper's native MPLS solution is standards-based and supports multivendor interoperability. This gives network operators added flexibility and choice when integrating new architecture with existing standards-based multivendor environments. The next section discusses Juniper Networks' MPLS solution that can fit into the transport workflow model and meet network operator needs to run mission-critical applications with guaranteed bandwidth, high availability, and reliability.

# MPLS in the Transport Core

MPLS was first introduced more than a decade ago and has evolved over the years to support the many needs of both businesses and end users. MPLS leverages dynamic routing information and brings connection-oriented switching via Layer 2 pseudowires (PW) to an otherwise connectionless IP network. Today, many large public and private packet networks worldwide are MPLS based and offer Layer 2 and Layer 3 MPLS-based services. MPLS is a platform for offering Layer 3 services such as MPLS VPNs over a single IP domain or multiple carrier networks. MPLS is also optimized to carry multiple protocols such as firewall, ATM, TDM, and Ethernet. Point-to-point and statically configured Layer 2 pseudowire services (very much like connection-oriented SONET circuits) can optimally transport native Ethernet or non-IP packets. One popular application of static MPLS pseudowires is in the migration of circuit-based SONET rings to MEF (Metro Ethernet Forum) services such as E-LINE, E-LAN, E-TREE.  MEF over MPLS-PW brings in the best of deterministic and secure SONET-like characteristics with the flexibility and simplicity of Ethernet. With support for QoS and traffic engineering, network operators can deliver premium SLAs and provide bandwidth guarantees. MPLS services can span multiple network domains and support unicast as well as scalable multicast services.

## Benefits of Dynamic Control Plane

The choice of a static versus dynamic control plane is dependent upon the operator's preferences and the stage of transition toward a packet core. A dynamic control plane in the packet world provides reachability to all the nodes in the core, and subsequently the connecting LSP can be manually (statically) provisioned by the operator using fixed and explicit next hops. A dynamic IGP (Interior Gateway Protocol) helps to scale the number of nodes and services delivered network-wide, supporting fast restoration of services in a mesh topology with multiple entry and exit points. Core routers perform the burdensome task of providing high availability by using a dynamic IGP.

MPLS can function in a completely static model, but this creates a burden for the operator to manually configure all LSPs and manually configure reroute-based availability models.  While MPLS can function in a static model, traffic policy control is available to the network operator with a base dynamic control plane. Static LSPs for the first-mile access with single point-to-point links are adequate. However, MPLS leverages the dynamic routing intelligence of IP routers and uses signaling protocols such as LDP or RSVP to set up the LSPs. Once established, these LSPs can be dynamically or statically configured. The MPLS-PW, being connection oriented, also provides adequate means to policy-control traffic paths, enable tiered quality of service (QoS), and differentiate guaranteed bandwidth services from best-effort services in a reliable manner. This flexibility of MPLS lends its application to converged packet transport networks.

## MPLS and Transport Network Requirements

### Centralized Management Plane Provisioning

MPLS uses LSPs to transport packets from one network entity to another. The LSP can be operator controlled and statically set up with an external out-of-band management system such as Juniper Networks® Junos® Space. Junos Space is an open standards-based network management system, whereby network operators can integrate MPLS tools into existing OSS and use centralized configuration, provisioning, and OAM.

[4] Building Multi-Generation Scalable Networks With End-to-End MPLS (www.juniper.net/us/en/local/pdf/whitepapers/2000452-en.pdf)

### Bandwidth Guarantee with a Deterministic Path

Similar to allocating bandwidth for a transport circuit, network operators can allocate bandwidth to statically defined LSPs, thereby ensuring a deterministic path for traffic on an LSP. Corouted bidirectional LSPs can be set up using GMPLS extensions to RSVP to create a connection-oriented pathway. LSPs can be provisioned with connection admission control (CAC). The LSP bandwidth can be carved out of available link bandwidth when setting up the LSP. Traffic at the ingress for this LSP can be classified and allocated to a customer-defined traffic class with appropriate QoS policy on every hop along the way.

### Reliability and Availability

MPLS supports a suite of OAM tools and MPLS fast reroute to deliver carrier-class network reliability and availability. With purpose-built traffic-engineered LSPs, operators can preconfigure explicit secondary paths to switch traffic in the event of any failure on the primary path. End-to-end fast restoration for services such as pseudowires can also be achieved. With packet optical integration and ITU-T G.709 OTN framing, failure recovery can be triggered with signal degrades and forward error correction (FEC)-based mechanisms. Also, a dynamic control plane can achieve a "best-effort" pathway in the face of multiple network failures.

### OAM[5]

MPLS supports a full suite of OAM techniques to detect and trigger link-layer and network-layer fault recovery. Recent in-band OAM enhancements defined by MPLS-TP standards are also available with a native MPLS implementation. Per- LSP monitoring is supported with in-band OAM for network-layer checks, e.g. GAL (Generic Associated Channel Label), G-Ach (Generic Associated Channel Header), and BFD (Bidirectional Forwarding Detection). On-demand troubleshooting tools such as LSP ping and traceroute and Y.1731, are also available.

## Transport-Oriented Workflow Model with an MPLS Packet Core

The migration to a converged packet core to improve network efficiency must accommodate the transport network operator's workflow model. To enable this smooth transition and keep the workflow model in perspective, Juniper's leading and widely deployed MPLS solution includes Junos Space, a unified application software platform that provides the out-of-band management needed to provision, operate, and manage the packetized core. Juniper Networks Junos Space Management system consists of multiple application suites with customer-built startup profiles and is an open-standards based platform for network operators to integrate custom-built applications. Key applications available with Junos Space include Juniper Networks Junos Space Security Director, Services Activation Director, Network Director, Service Now, and Service Insight. Of these applications, Services Activation Director provides complete life-cycle management of services including workflows for service design, provisioning and timing[6].

## Services Activation Director

Junos Space Services Activation Director ensures error-free service provisioning and monitoring of legacy carrier Ethernet and MPLS services using a simple interface to design, validate, and manage these services. Services Activation Director consists of these five Junos Space applications:

- Network Activate facilitates automated provisioning and validation of MPLS and carrier Ethernet services.
- Transport Activate allows for the design and provisioning of Layer 2 paths, or LSPs, in various configurations such as P2P, P2MP or full mesh (MP2MP).
- Sync Design facilitates configuring of timing interfaces such as PTP and Synchronous Ethernet in various devices.
- QoS Design is for the design of QoS profiles for bandwidth management, traffic shaping, and network congestion control.
- OAM Insight is for service performance monitoring and SLA assurance using Y.1731, Ethernet OAM-CFM (connection fault management), LFM (link fault management), and BFD.

## Services Activation Director Deployment Example

Let's consider a service deployment and see how Junos Space fits into the workflow mode. In the following example a network operator provides intersite connectivity to its customers.

a. To request a service, a customer fills out a service request template provided by the network operator. This network request is then received by a service team and translated into a service order by identifying the geography of service location, the data rate requested, type of connectivity, QoS, security, and SLA requirements.

[5] www.juniper.net/techpubs/en_US/junos/topics/topic-map/mpls-tp-oam-configuration.html

[6] www.juniper.net/techpubs/en_US/junos-space12.2/information-products/topic-collections/junos-space-transport-activate-pwp/junos-space-transport-activate-pwp.pdf

b. A separate network design team then maps out the service points based on the network topology, categorizes the service for its QoS and OAM requirements, and assigns network resources based on available bandwidth on each hop along the way from the service origin point to the service destination. Role-based access control can also be attached to this service order. This can be done via a template generated in the Junos Space Transport Activate module of Junos Space Services Activation Director.

c. Based on the role-based access specified in the service order and the associated service attributes, a separate network deployment team then uses Transport Activate to provision the circuit (LSP). The LSP can be provisioned as a multi-hop static path or with RSVP signaling (see Figure 3 and Figure 4). With RSVP-TE signaling, the operator can control the pathway by choosing an explicit path LSP type. The operator can control the pathway the LSP takes by defining each intermediate link or next hop from the ingress node to the egress node. This is similar to a static LSP but needs much less configuration. With RSVP-TE as the signaling protocol, the user can also allocate bandwidth for the LSP and RSVP checks for network resource availability along every hop of the LSP path. The user can also define a protection path for this LSP (see Figure 5) that is used in the event of any link or node failure along the primary path.

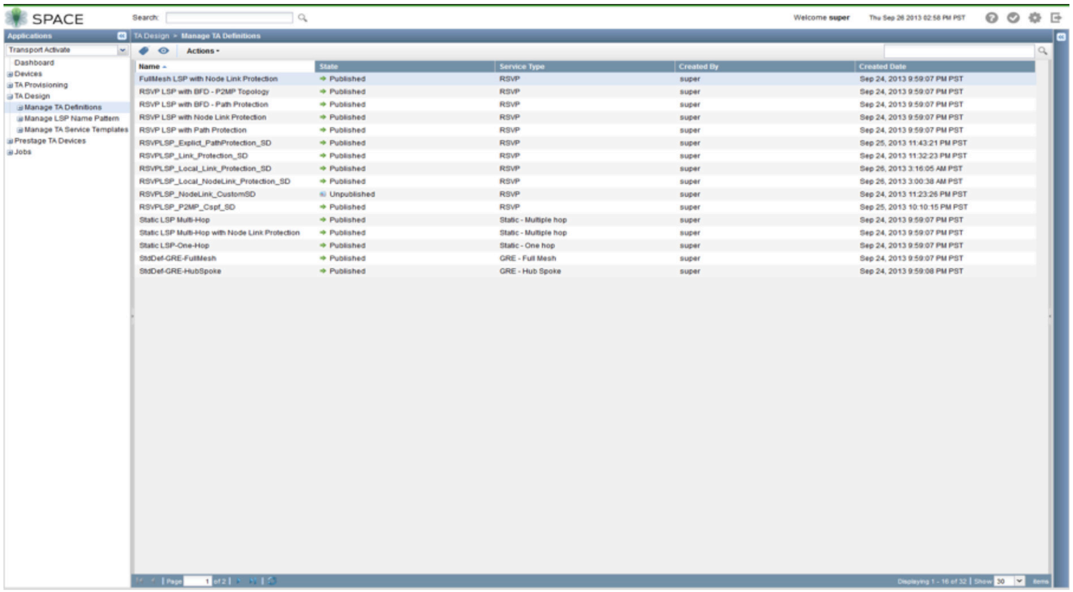

| Name | State | Service Type | Created By | Created Date |
|---|---|---|---|---|
| FullMesh LSP with Node Link Protection | Published | RSVP | super | Sep 24, 2013 9:59:07 PM PST |
| RSVP LSP with BFD - P2MP Topology | Published | RSVP | super | Sep 24, 2013 9:59:07 PM PST |
| RSVP LSP with BFD - Path Protection | Published | RSVP | super | Sep 24, 2013 9:59:07 PM PST |
| RSVP LSP with Node Link Protection | Published | RSVP | super | Sep 24, 2013 9:59:07 PM PST |
| RSVP LSP with Path Protection | Published | RSVP | super | Sep 24, 2013 9:59:07 PM PST |
| RSVPLSP_Explict_PathProtection_SD | Published | RSVP | super | Sep 25, 2013 11:43:21 PM PST |
| RSVPLSP_Link_Protection_SD | Published | RSVP | super | Sep 24, 2013 11:32:23 PM PST |
| RSVPLSP_Local_Link_Protection_SD | Published | RSVP | super | Sep 26, 2013 3:16:05 AM PST |
| RSVPLSP_Local_NodeLink_Protection_SD | Published | RSVP | super | Sep 26, 2013 3:00:38 AM PST |
| RSVPLSP_NodeLink_CustomSD | Unpublished | RSVP | super | Sep 24, 2013 11:23:26 PM PST |
| RSVPLSP_P2MP_Cspf_SD | Published | RSVP | super | Sep 25, 2013 10:10:15 PM PST |
| Static LSP Multi-Hop | Published | Static - Multiple hop | super | Sep 24, 2013 9:59:07 PM PST |
| Static LSP Multi-Hop with Node Link Protection | Published | Static - Multiple hop | super | Sep 24, 2013 9:59:07 PM PST |
| Static LSP-One-Hop | Published | Static - One hop | super | Sep 24, 2013 9:59:07 PM PST |
| StdDef-GRE-FullMesh | Published | GRE - Full Mesh | super | Sep 24, 2013 9:59:07 PM PST |
| StdDef-GRE-HubSpoke | Published | GRE - Hub Spoke | super | Sep 24, 2013 9:59:08 PM PST |

Figure 3: Managing LSP definitions with Junos Space Transport Activate.

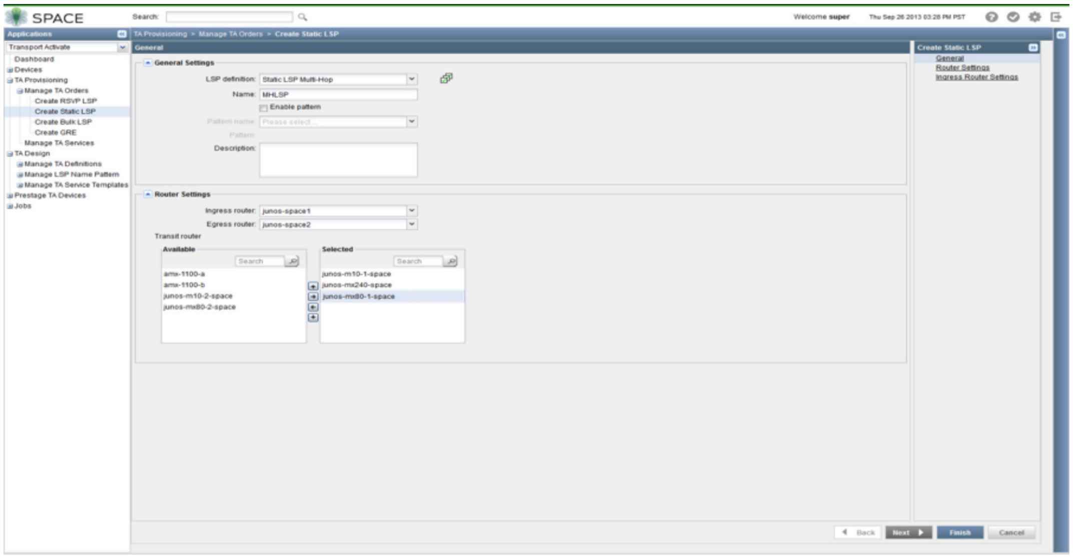

General Settings

LSP definition: Static LSP Multi-Hop

Name: MHLSP

□ Enable pattern

Pattern name: Please select ...

Pattern:

Description:

Router Settings

Ingress router: junos-space1

Egress router: junos-space2

Transit router

| Available | Selected |
|---|---|
| amx-1100-a | junos-m10-1-space |
| amx-1100-b | junos-mx240-space |
| junos-m10-2-space | junos-mx80-1-space |
| junos-mx80-2-space | |

Figure 4: Configuring a static mult-hop LSP with Junos Space

d. With Junos Space Transport Activate, the network deployment team can then define the appropriate QoS policy for traffic on the provisioned LSP. The QoS policy definition includes defining the class of service for the traffic on the LSP and the metric for the link to the next-hop destination address (see Figure 6).

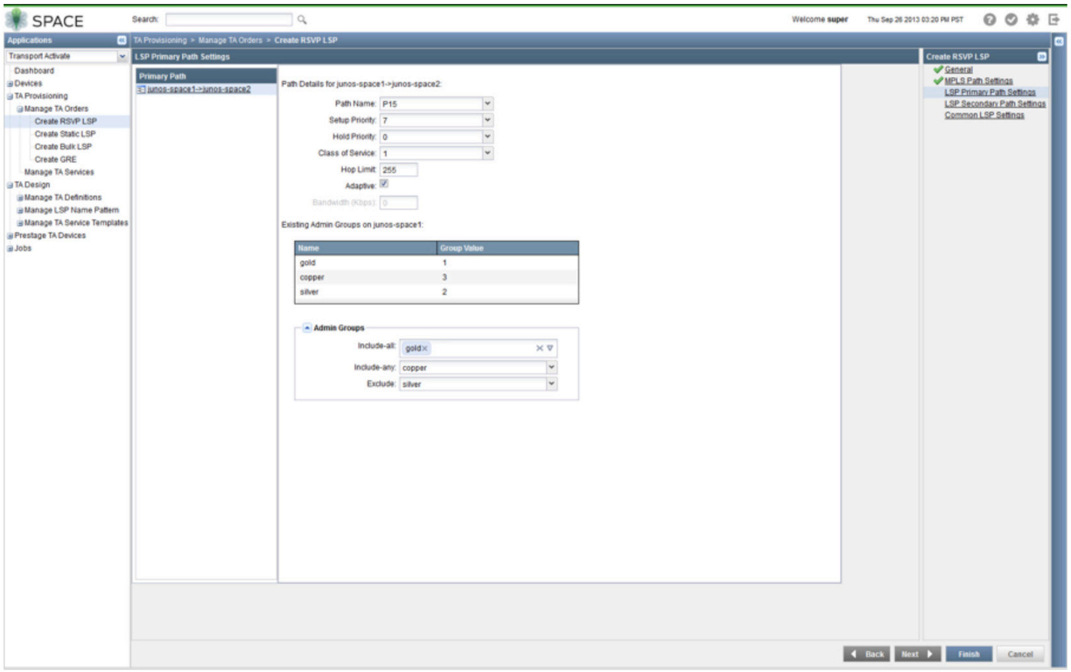

Path Details for junos-space1->junos-space2:

Path Name: P15
Setup Priority: 7
Hold Priority: 0
Class of Service: 1
Hop Limit: 255
Adaptive: ☑
Bandwidth (Kbps): 0

Existing Admin Groups on junos-space1:

| Name | Group Value |
|---|---|
| gold | 1 |
| copper | 3 |
| silver | 2 |

Admin Groups

Include-all: gold
Include-any: copper
Exclude: silver

Figure 5: Configuring LSP path protection with Junos Space.

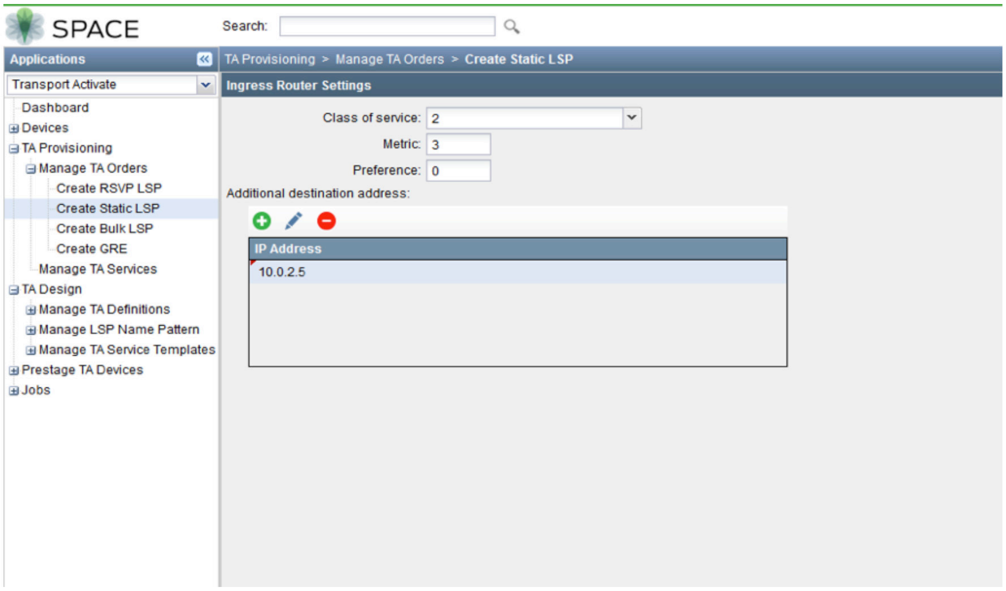

Class of service: 2
Metric: 3
Preference: 0
Additional destination address:

| IP Address |
|---|
| 10.0.2.5 |

Figure 6: Configuring quality of service for LSP with Junos Space with metrics for LSP path protection.

e. Finally, a separate team focused on SLA enforcement can specify OAM requirements for the LSP. For example, BFD timers can be set up for In-band management, and demand-driven OAM can be specified as well. Link fault management (LFM) and connectivity fault management (CFM) are also available to monitor the connections.

## Conclusion

Packetizing the transport network is a means of gaining network efficiency and optimal utilization of network resources. There are multiple solutions available to transition the transport network from circuit-switched technology to a packetized core. The choice of MPLS-TP versus MPLS varies from one network operator to the other. Juniper's converged MPLS solution with a centralized network management system—Junos Space—provides a standards-based simplified solution and supports multivendor interoperability, thereby allowing integration with open OSS and paving a way to future- proof the network for years to come.

## 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.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701

2000543-001-EN Sept 2015