

# PTX5000 100G Packet Optical Solution

---

Field trial with Third Party DWDM system

## Table of Contents

Executive Summary .....	3
Introduction.....	3
Integrated Optics in Packet Switching .....	3
Colored Optics in Production DWDM Networks .....	6
Trial Overview.....	6
ESnet BayExpress Network .....	6
Test Scenarios.....	8
Preparatory Lab Tests.....	9
Test Case A.....	11
Test Case B.....	11
Test Case C.....	12
Results.....	12
Measured Optical Parameters – PTX5000 Wavelength.....	12
Measured Optical Parameters – Native Wavelengths.....	13
Measured Tx Power Levels – Case C.....	13
Conclusions.....	15
About ESnet.....	15
About Juniper Networks.....	16

## Executive Summary

Strong traffic growth in combination with stagnating or decreasing revenue force commercial and government-funded network operators to reduce network costs and simplify network layers and operation. Meanwhile, advances in optical technology have simplified the process of dense wavelength-division multiplexing (DWDM) network design and allowed the integration of colored optics into routers. In this document, we are describing a field trial implemented by ESnet and Juniper Networks to prove the ability of transport packet platforms with coherent optical interfaces to operate over a third-party DWDM system from a technical and operational point of view.

## Introduction

With traffic rising and the average revenue per user (ARPU) floundering or falling, service providers are looking for ways to reduce network costs (both CapEx and OpEx) while providing competitive speed, reliability, and resiliency. Eliminating unnecessary equipment and increasing integration between the network layers are key steps on this path forward.

Enhancements in optical technology such as coherent transmission with digital signal processing (DSP) to compensate for chromatic dispersion and polarization mode dispersion (CD & PMD) have dramatically improved transponder design and allowed the integration of colored optics into routers. Such integration eliminates the need for short-reach “patches”, between the gray optical interface on the router and the gray optical interface on the transponder, allowing for direct interaction between packet and long-haul optical layers.

In Autumn 2013, ESnet and Juniper went on to a joint effort to demonstrate the operation of routers with integrated optical interfaces over a production third-party DWDM system. This field trial included the following sections:

1. Proof that the integrated router optics can provide error-free transmission over extended distances
2. Proof that the alien wavelength provisioned from the router port does not interfere with performance or color plan of previously installed transponders in operational DWDM plant
3. Proof that the alien wavelengths can be added and removed without power budget adjustments of existing transponders and amplifiers

This trial resulted in valuable data, which can be used to de-risk the prospective deployments of integrated DWDM optics alongside with transponder shelves. This test can be seen as an important step towards moving this advanced technology into the mainstream networks.

## Integrated Optics in Packet Switching

Since 100G DWDM optics use dual-polarization quadrature phase shift keying (DP-QPSK) modulation and coherent transmission with DSP for compensation of chromatic and polarization mode dispersion, DWDM networks are now becoming significantly easier to operate. This has created a big momentum for the integration of DWDM technology directly into router linecards (Fig 1).




Figure 1: Integrated optics

Integrated optics may provide significant CapEx savings due to the fact that a traditional transponder-based approach requires optics for the DWDM long-reach (line) side plus two grey (short reach, non-DWDM) couplers: one on the router and another one on the client side of the transponder. In addition, the external transponder placement needs shelf slots with the associated power supplies, controllers, fans and management logic. Incorporation of the DWDM photonics into the router reduces the number of optical interfaces per link from six to two. In addition, thanks to recent advances in form factor of tunable optics, colored interfaces can now provide the same (or similar) port density as grey optical interfaces (as for example with 100GbE DWDM PICs in Juniper Networks® PTX5000 Packet Transport Switch and PTX3000 Packet Transport Switch).

Besides lower capital cost, the reduction in hardware with optical integration enables facility savings with smaller footprint (no external transponder shelves) and curtailed energy consumption. Additional monetary effect can be seen from reduction of the operator’s cost of provisioning: the elimination of intermediate layers means fewer configuration touchpoints and simpler deployment procedures.

Finally, integrated packet-optical network management with interoperable control plane is preferable to the “ships-in-the-night” layered model of network planning. Joint optimization of router and transponder interfaces against the matrix of planned demands may help to eliminate some characteristic multi-layer inefficiencies (for example, where small demands are riding high-speed pipes because the possibility of intermediate-node packet multiplexing was not taken into account).

On the operational side, there are a number of advantages for the router to have direct access to the optical layer at the level of network performance as well. For example, an MPLS fast re-route (FRR) can be triggered by pre-forward error correction (pre-FEC) bit error rate (BER). Typically, a fiber gets bent just a moment before it breaks, which results in rapidly increasing BER. The observation of this fact allows the router to perform the switchover before the fiber break occurs. To set a balance between the reach and desired performance, the operator can set a BER threshold for the switchover (and switchback after repair).

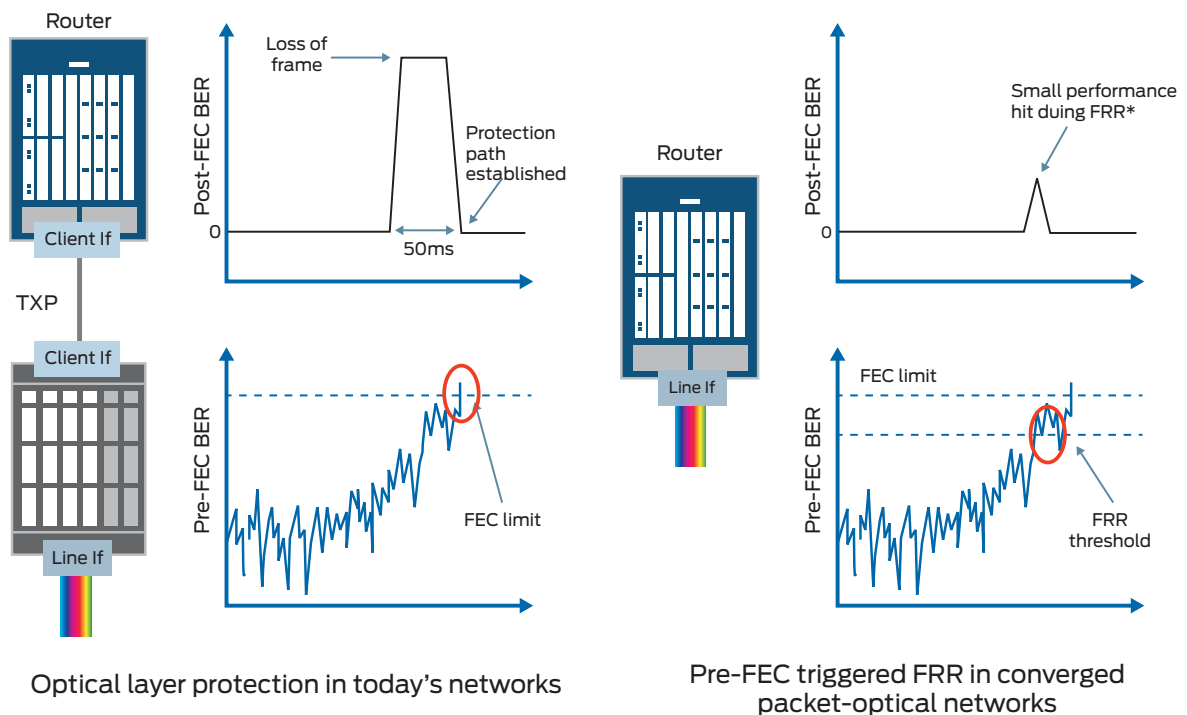


Figure 2: FRR triggered by BER increase

A key enabler and operational requirement in successfully integrating DWDM optics directly into router interfaces is the amalgamation of those optics into existing DWDM network management systems. For alien wavelength application, the concept of “virtual transponders” (VXP) allows the direct incorporation of integrated optics into 3rd party DWDM infrastructure<sup>1</sup>. This way, the DWDM management system can keep control over the optical parameters of the integrated optics as if it is a standalone transponder (Fig 3).

<sup>1</sup> The opposite scenario can be seen if a router needs to control the external transponder. Juniper Fusion technology allows for configuration touchpoint reduction in the latter case, and can be mixed and matched with VXP in one network

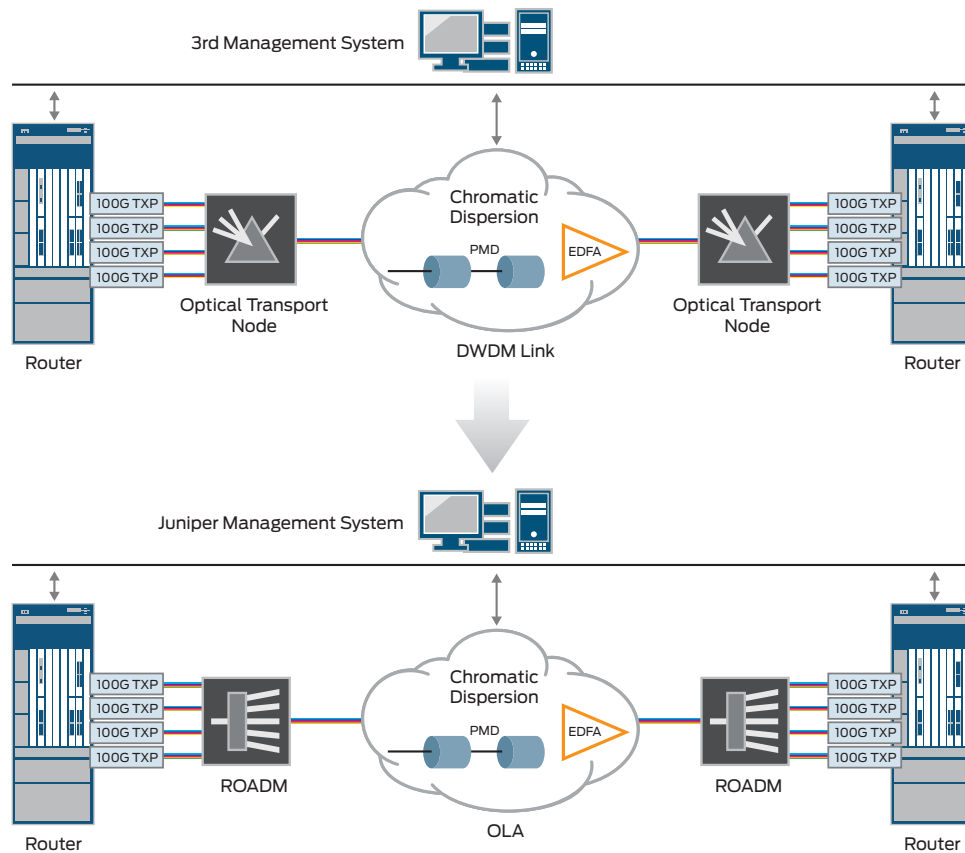


Figure 3: Virtual transponder and network management

Note, that virtual transponder functionality can work in addition to interoperability of MPLS and optical control plane via Generalized MPLS (GMPLS) messaging. Interoperable control planes allow for automated set-up of wavelengths between routers—triggered, for example, by a provisioning command on the common network management system. This tight integration prevents situations where double failures may occur: for instance, the router can communicate with the DWDM network to reroute an optical channel so that there is a new independent protection path is established<sup>2</sup>.

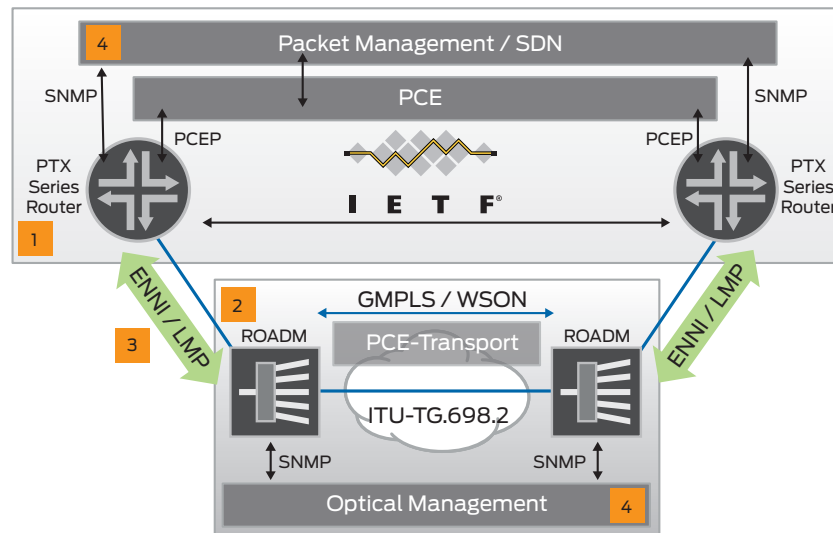


Figure 4: Standards-based management and control plane

<sup>2</sup> Since the communication between the router's control plane and the DWDM system's control plane is enabled by standardized GMPLS-External Network Node Interface (ENNI), custom-tailored choices of router and DWDM system vendors are now possible.

## Colored Optics in Production DWDM Networks

Juniper believes that freedom of DWDM planning allows network operators to pick the optimum solutions for their networks. For this reason, we routinely test our optical solutions against other vendor's equipment, for example ADVA Optical Networking, Nokia Siemens Networks (now Coriant), Ciena and Huawei. The field trial described in this document is part of Juniper's open standards strategy, ensuring easy transmission of 100G over existing DWDM transport infrastructure of ESNet.

### Trial Overview

The PTX5000 was deployed at ESnet's LBNL site and connected to an ESnet DWDM lab network for initial preparatory tests. After validating basic interoperability in the lab, the PTX5000 was connected to ESnet's production BayExpress optical network. Multiple performance tests were then conducted at progressively increasing distances across the BayExpress.

### ESnet BayExpress Network

ESnet's regional optical network in the Bay Area (BayExpress) was chosen for several reasons. It is geographically close to the Juniper R&D headquarters in Sunnyvale, and is exclusively managed by ESnet, which guarantees access to the entire fiber spectrum. Finally, the ESNet fiber topology and number of Reconfigurable Optical Add-Drop Multiplexers (ROADMs) in the path allowed for a test strategy with progressively increasing distances by using loopbacks at different ROADM nodes.

An approximate representation of the BayExpress fiber topology is shown in Fig. 5

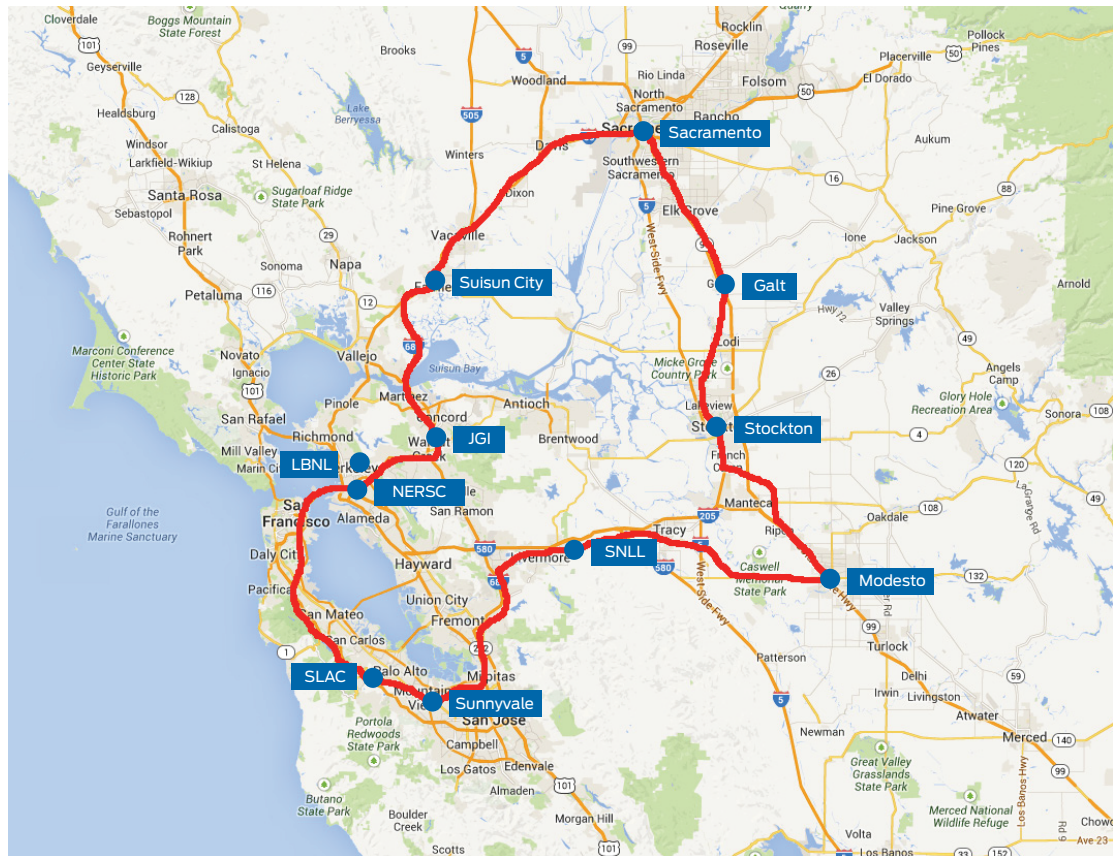


Figure 5: The ESnet BayExpress fiber topology

The BayExpress plant consists of four fiber strands providing physically diverse connectivity between ESNet sites in Sunnyvale and Sacramento. The fiber is mostly Corning LEAF® with small amounts of Corning SMF-28®. The network's DWDM infrastructure includes seven ROADM nodes and five Inline Amplifiers (ILAs). The ROADMs in the Sunnyvale and Sacramento metro hubs connect the regional Bay Area fiber ring to long-haul dark fiber spans.

This fiber layout corresponds to a “logical ring” topology shown below.

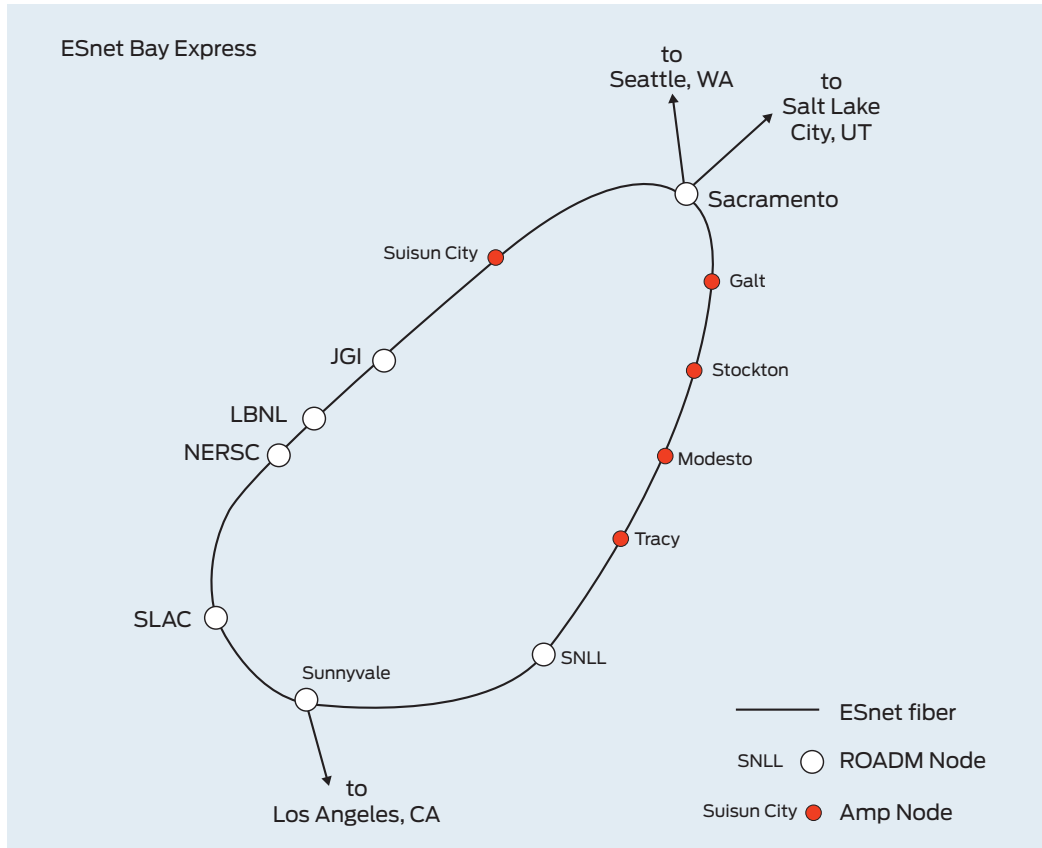


Figure 6: ESnet BayExpress network schematic

The BayExpress network does not use any traditional dispersion compensating fiber (DCF) modules. Instead, both 10G and 100G transponders use electronic dispersion compensation. Currently, there are no 40G wavelengths deployed in the BayExpress network. A detailed map of existing traffic/wavelengths at the time is shown Fig 7. The 100G wavelengths are coherent using a DPQPSK modulation scheme, and the 10G wavelengths are using NRZ encoding.

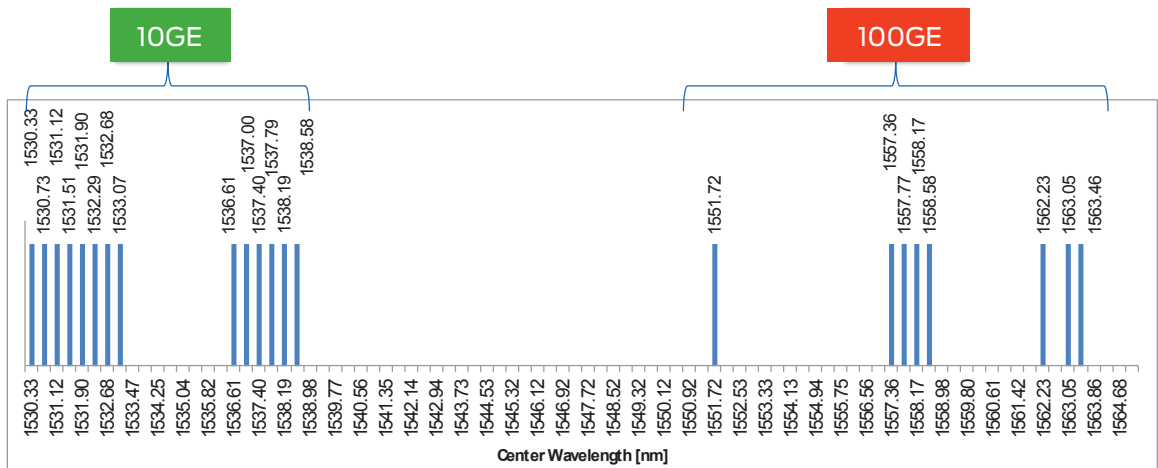


Figure 7: Map of used wavelengths

## Test Scenarios

The main components involved were the PTX5000, an Ixia Optixia XM2 IP Performance Tester, and the ESnet BayExpress production and lab networks. The PTX5000 and Ixia were collocated in the same room at Lawrence Berkeley National Laboratory (LBNL). Access to the BayExpress production and lab networks was via ROADMs located in other rooms at the same LBNL facility.

The PTX5000 was configured to run two logical systems in order to simulate two separate routers using only one piece of equipment. The Ixia IP performance tester was connected to the PTX5000 via two 100GE grey interfaces. The Ixia was used to generate line-rate test traffic and measure the quality of that test traffic after traversing the PTX5000 and DWDM network.

The Ixia IP performance tester was connected to the PTX5000 via two 100GE grey interfaces. The Ixia was used to generate line-rate test traffic and measure the quality of that test traffic after traversing the PTX5000 and DWDM network.

The PTX5000 was first connected to a local lab ROADM with local loopback within the LBNL laboratory for preparatory tests. Next, the PTX5000 was connected to the production BayExpress network for field tests at increasing transmission distances. Three cases (A, B, and C) with loopbacks in different nodes and different resulting transmission distances were tested:

- Case A: LBNL optical loopback in Oakland (NERSC) (~24km)
- Case B: LBNL optical loopback in Stanford (SLAC) (~210km)
- Case C: ESnet BayExpress East-West nearly full optical loop in the reverse direction with loopback at NERSC (~1400km)

Figure 8 shows the general configuration with the PTX5000 configured as two logical routers LS1 and LS2.




Figure 8: Basic test setup(s)

The PTX5000 itself was configured with two FPCs, one for each logical system. The FPC occupied slots 0 and 1. Each FPC housed one two-port 100GbE Physical Interface Card (PIC) and one two-port 100G OTN PIC with integrated DWDM optics. On each of the 100GbE grey optics PICs, one port connected to the IP Performance Tester, and the other was not used. On each of the 100G OTN PICs, one of the two ports was used to connect to each other through the optical network:

1. The lab ROADM in the lab for initial preparatory testing
2. The BayExpress network (cases A, B, and C)






Figure 9: PTX5000 configuration

### Preparatory Lab Tests

The approach was to start with simple tests and increase the complexity and transmission distances after each successfully completed step. This gradual approach was chosen to allow the ESnet and Juniper teams to get acquainted with the PTX5000 and its colored 100G optics, the BayExpress network, the test equipment, and overall setup. This approach was also the most conservative in terms of protecting the existing production traffic on the BayExpress network.

A first basic step was to connect the Ixia XM2 IP Performance Tester to the PTX5000 and perform packet-level verification with loopback on the PTX5000's grey 100GbE ports.

Once the two PICs with 2x100G DWDM OTN interfaces each were installed in the PTX5000, the optical parameters were confirmed via an optical spectrum analyzer. The results are displayed in Figure 10 and Figure 11. The wavelength was centered at 1547.697 nm with -7.82 dBm signal power, an optical signal-to-noise ratio (OSNR) of 19.22 dB with a noise of -27.4 dBm, and a reference bandwidth (RWB) of 0.1 nm.

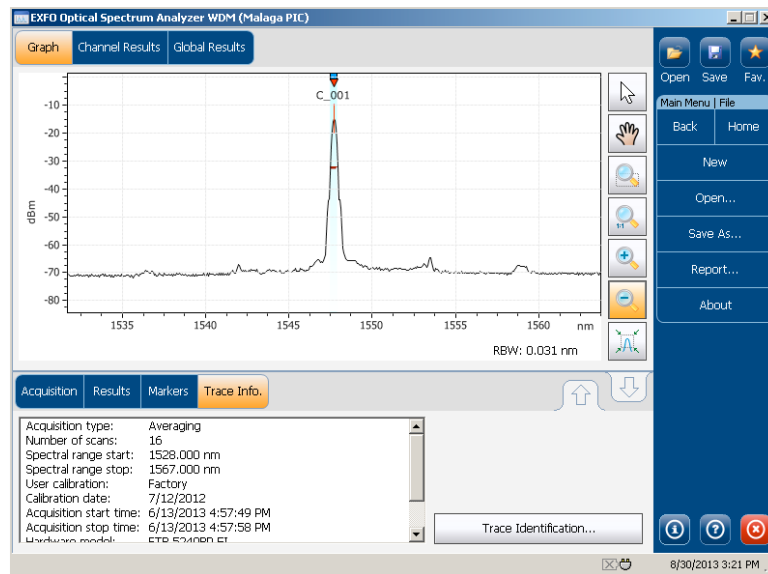


Figure 10: Wavelength verification on optical spectrum analyzer

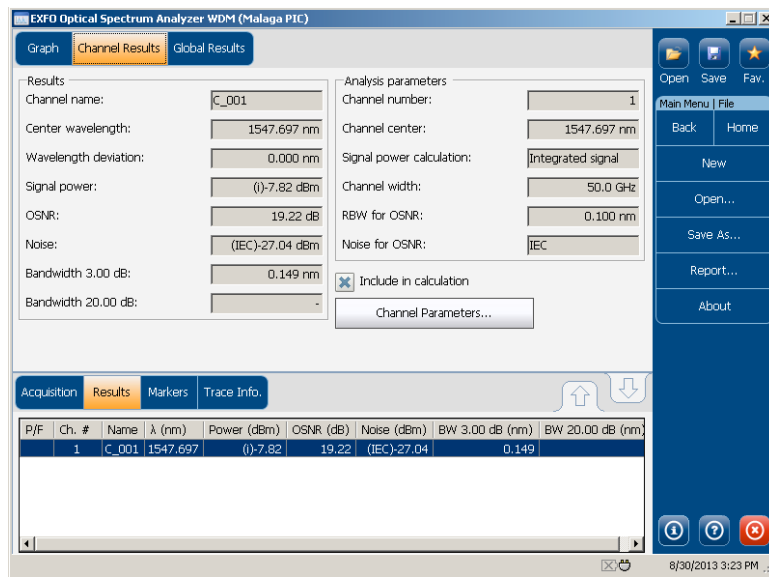


Figure 11: Optical channel data via optical spectrum analyzer

Then a local loopback was performed between the colored optics ports via patch cables and optical attenuators. This step confirmed the functioning of the local setup.

In order to better understand the interaction between the type of ROADM gear in the BayExpress and the PTX Series colored optics, we connected the PTX5000 to a standalone ROADM in the LBNL lab and did a loopback at the DWDM multiplex module on the second shelf of the ROADM (Figure 12 and Figure 13).




Figure 12: Architecture of ROADM




Figure 13: Setup for preparatory lab test

### Test Case A

After understanding the interaction of the PTX5000 optical interfaces and the ROADM and its amplifiers, the team felt comfortable adding the colored optics to the production network. Due to the short fiber distances involved and the continuous presence of onsite operational personnel at both LBNL and NERSC, the LBNL-NERSC span was chosen to be tested on the production network.

Before performing the loopback in NERSC, the span's per-channel optical powers and bit error rates (BERs) were measured. After that, the wave was added to the production network with a loopback at NERSC. The optical performance was excellent as expected with no bit errors recorded on the PTX5000 wavelength and no impact to existing production traffic. The resulting distance was 2x12 km or 24 km.

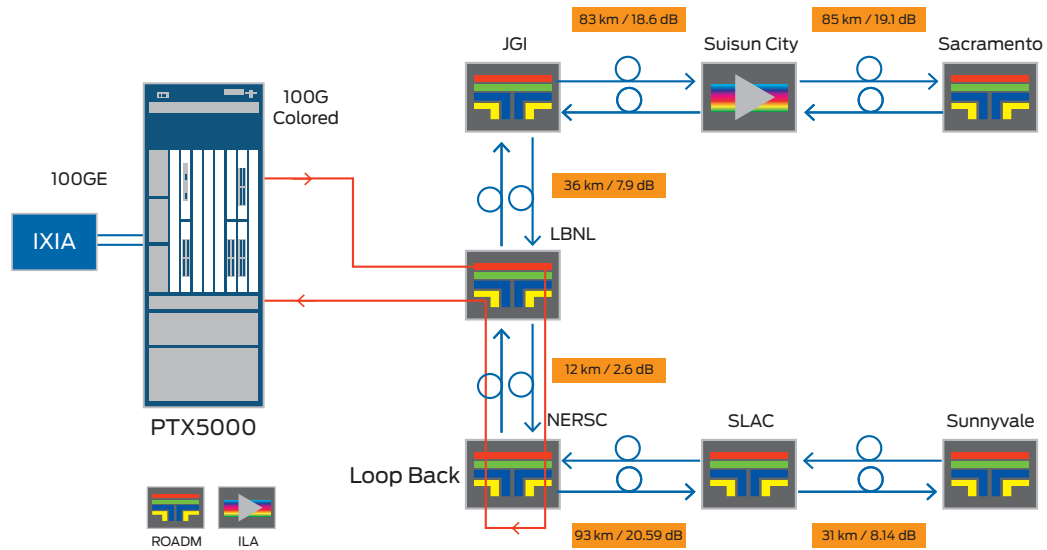


Figure 14: Test Case A with loopback at NERSC

### Test Case B

As the test in Case A was successful without any impact on the operation of the existing traffic, the team decided to step up the challenge. NERSC was switched to a “pass-through” configuration for the PTX5000 wavelength, and it was then looped back at the next ROADM in SLAC.

The resulting distance was 2\*(12+93) = 210 km going through each of three ROADMs twice for a total of 6 concatenated ROADMs. Again, the results were as expected without any impact on operations.

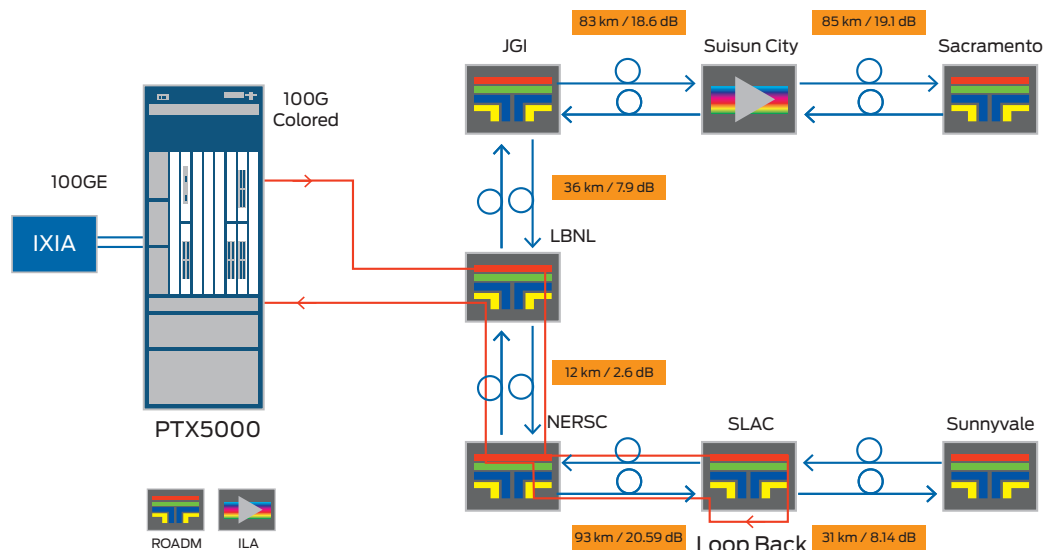


Figure 15: Test Case B with loopback at SLAC

### Test Case C

With full confidence in the setup, the level of challenge was increased. This time, the PTX5000 wavelength was routed in the reverse direction toward JGI around the whole BayExpress ring up to NERSC, where it was looped back and then all the way back to LBNL. The total transmission distance was 1400 km, and the number of ROADMs passed was 14. An additional tax on the optical budget was the fact that the signal had to traverse four patch panels and the associated patch cables between the PTX5000 in the lab room and the ROADM in the production network’s transport room. In a permanent install of the PTX5000 with integrated colored optics, the PTX5000 would be located near the ROADM, and the additional penalty of the connection between two different rooms would be avoided.

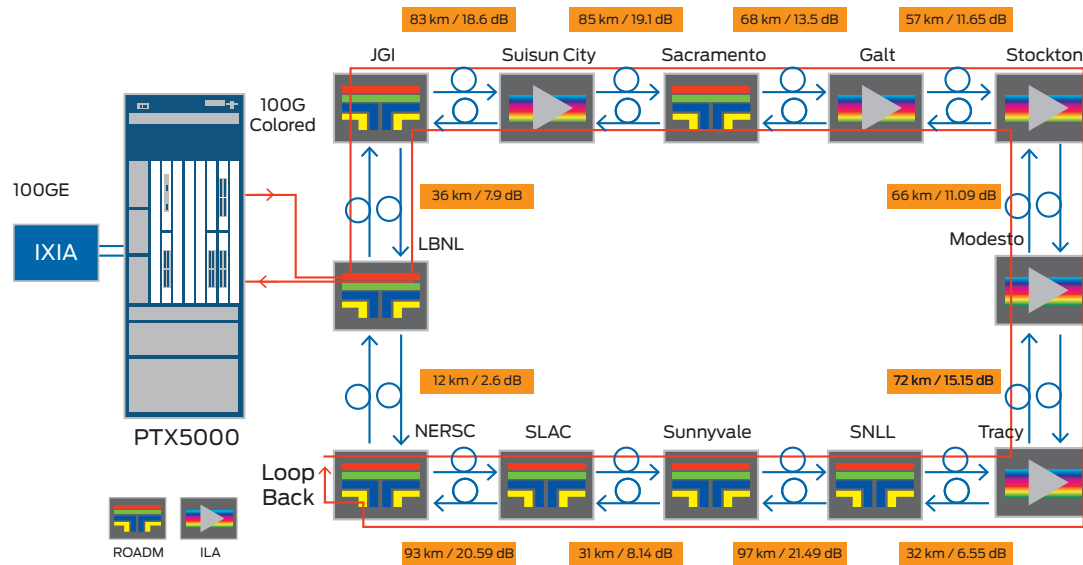


Figure 16: Test Case C with full loop and loopback at NERSC

### Results

#### Measured Optical Parameters – PTX5000 Wavelength

The measured parameters showed a solid OSNR on the PTX5000 wavelength. The parameters are displayed in Figure 17. The chromatic dispersion of 6040 ps/nm was due to the large amount of LEAF fiber in the network. The reported Pre-FEC BER shows a margin >1dBQ over FEC cliff. Line-system fiber launch power can be further optimized to increase margin. Also, placing the PTX5000 closer to the ROADM and avoiding the 4 patch panels and associated cables and connectors would result in even better margins.

#### Malaga Optical Monitoring Statistics:

TX output power	:	-3.01 dBm
RX input power	:	-9.51 dBm
Line chromatic dispersion	:	<b>6040 ps/nm</b>
RX Q factor value	:	<b>8.00 dB</b>
Carrier frequency offset	:	-682 MHz

#### Malaga OTN FEC statistics:

Corrected Errors	4883917085509
Uncorrected Words	0
<b>Corrected Error Ratio (6039 sec average)</b>	<b>8e-3</b>

Figure 17: Some optical parameters measured for Case C

## Measured Optical Parameters – Native Wavelengths

An important part of the test for ESnet was to understand the impact of the colored optics/alien wavelength on the other wavelengths on the BayExpress and on the system itself. In order to verify proper operation of the existing native wavelengths, Pre-FEC bit error rates (BERs) were collected for each wavelength on the system before and after the addition of the alien wavelength. The Pre-FEC BERs represent the BER on the DWDM line-side transceivers before forward error correction is applied. In the case of signal impairment, Pre-FEC BERs increase even though FEC might continue to correct the errors before client handoff.

Pre-FEC BER values for a representative subset of wavelengths in test Case C are included in Figure 18. After adding the alien wavelength, the values vary only slightly relative to their initial readings. Thus, no negative impacts on the DWDM line side were observed as a result of the alien wave addition.

ROADM Node	Wavelength (nm)	Pre-FEC BER Before	Pre-FEC BER After
SACR	1531.12	1.00E-06	1.00E-06
SACR	1531.90	1.00E-06	1.00E-06
SACR	1530.33	1.10E-06	1.00E-06
SACR	1537.40	1.00E-06	1.00E-06
SACR	1538.58	1.10E-06	1.10E-06
SACR	1536.61	1.00E-06	1.00E-06
SACR	1532.68	1.00E-06	1.00E-06
NERSC	1530.33	1.00E-06	1.10E-06
NERSC	1530.72	1.10E-06	1.00E-06
NERSC	1538.19	1.00E-06	1.00E-06
NERSC	1551.72	3.00E-04	3.10E-04
NERSC	1558.17	2.80E-07	3.00E-07

Figure 18: Some Pre-FEC BERs measured for Case C

## Measured Tx Power Levels – Case C

While the quality (Pre-FEC BER) of the existing wavelengths was not affected by the PTX5000 colored optics wavelength, the DWDM system's algorithms to automatically adjust overall power levels without manual intervention was still to be confirmed.

To this end, power data was extracted from selected ROADMs along the wavelength path before and after addition of the alien wavelength. The first measurement was taken right after the wavelength passed through the Sacramento ROADM toward SNLL. The second measurement was taken after passing through Sunnyvale toward SLAC and the third right after the signal looped back in NERSC and headed back toward SLAC. The fourth measurement was taken right after the wavelength passed through Sunnyvale on its way back toward SNLL, and the fifth and final measurement was in Sacramento on the way back to LBNL (Figure 19).

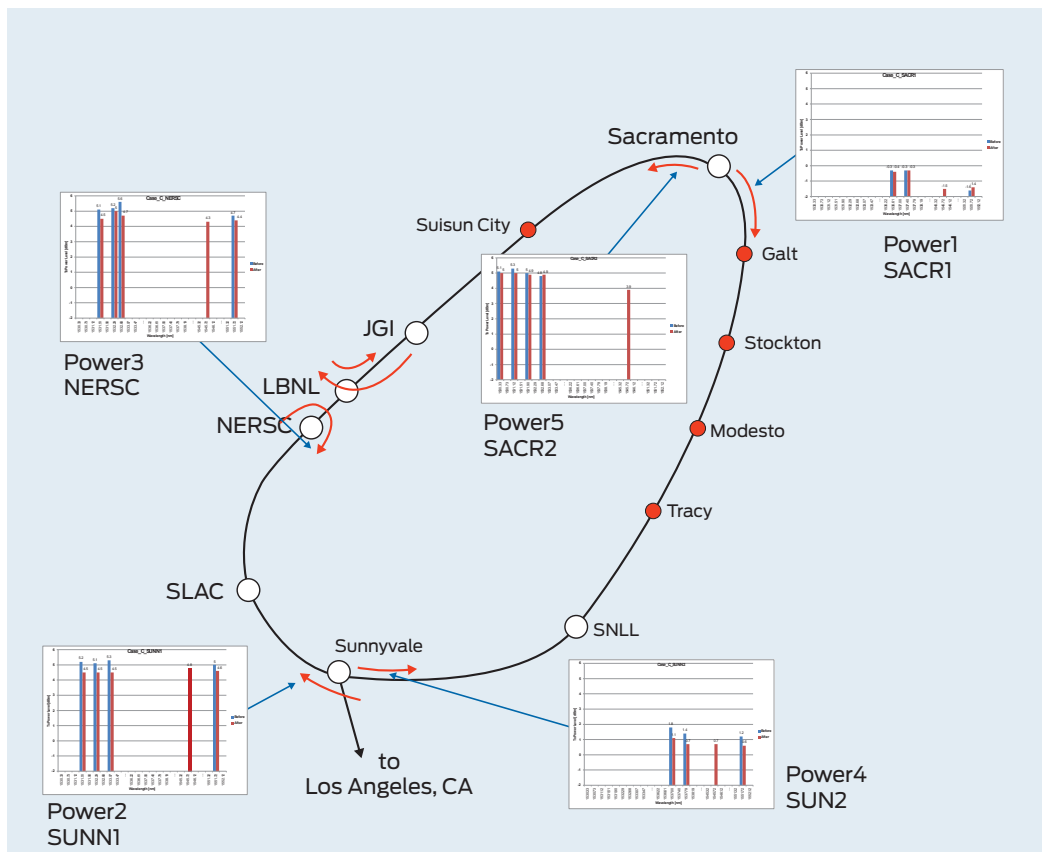


Figure 19: Test Case C Tx power-level measurement locations

Figure 20 shows the power measurements along the measuring points in sequence before (blue bars) and after (red bars) addition of the alien wavelength. We can see minor changes in the power levels of the existing traffic. We also see the add-drop of existing wavelengths along the path.

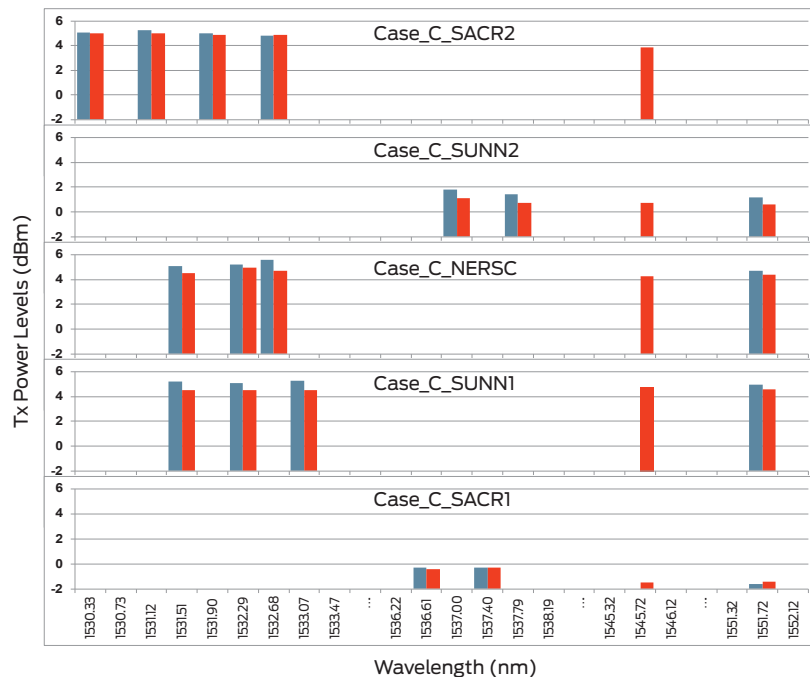


Figure 20: Comparison of Tx power levels through network

Detailed measurements are shown in Figure 21 for the fifth and final measurement in Sacramento.

The per-channel powers vary from one span to another. This is expected behavior of the optical system, and the values are influenced by a number of factors such as the span distance, number of wavelengths in use, and the type of wavelengths. The important aspect of these measurements is the relative optical powers of different wavelengths on any individual span. The optical system attempts to balance power across all wavelengths within approximately 1dB of each other regardless of their input power. As these results show, the optical system was able to successfully balance the alien and native wavelengths with no negative impacts.

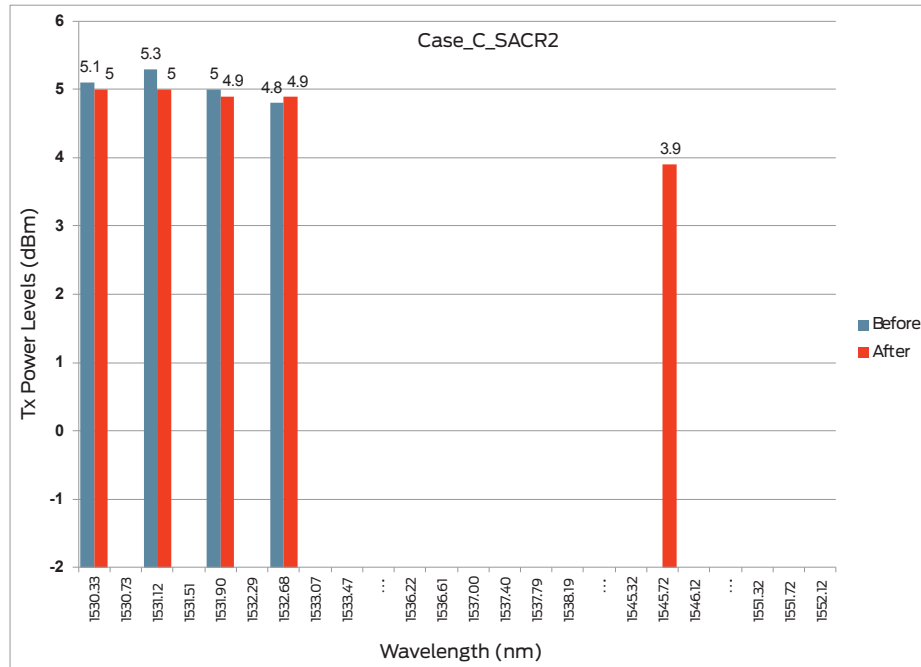


Figure 21: Last Tx power levels after passing through Sacramento

## Conclusions

As a result of the trial, ESnet and Juniper were able to demonstrate the following in a production network:

1. The integrated optics provided error-free transmission over an existing DWDM system over 1400km of mixed LEAF and SMF-28 fiber.
2. The alien wavelength from the router port did not affect the existing wavelengths provided by the transponders on the DWDM system, and no errors were reported on the existing traffic.
3. The alien wavelength was easy to provision and deprovision, and no manual intervention to provide power balancing and adjustments of amplifiers and transponders was needed, as the third-party DWDM system provided these adjustments completely automatically.

With optical integration of DWDM interfaces into routers at a density comparable to grey optics, and the simplification of DWDM network design because of the use of DP-QPSK and digital signal processing for dispersion compensation, operators can now get the CapEx, OpEx, and service benefits of the new technology.

## About ESnet

The Energy Sciences Network (ESnet) is a high-performance, unclassified national network built to support scientific research. Funded by the U.S. Department of Energy's Office of Science (SC) and managed by Lawrence Berkeley National Laboratory, ESnet provides services to more than 40 DOE research sites, including the entire National Laboratory System. ESnet's fifth-generation network, ESnet5, was built in collaboration with Internet2 and launched in November 2012. The two organizations equally share an 8.8 terabit-per-second optical transport platform to serve their respective communities. Additional information can be found at [www.es.net](http://www.es.net).

## 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](http://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](http://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

Copyright 2015 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos and QFabric 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.