

# Business Drivers and Design Guidelines for Network Convergence and Virtualization of IP/MPLS Core Networks



Network Strategy Partners, LLC

MANAGEMENT CONSULTANTS TO THE NETWORKING INDUSTRY

[www.nspllc.com](http://www.nspllc.com)

June, 2009

**Network Strategy Partners, LLC (NSP)** — Management Consultants to the networking industry — helps service providers, enterprises, and equipment vendors around the globe make strategic decisions, mitigate risk and affect change through custom consulting engagements. NSP's consulting includes business case and ROI analysis, go-to-market strategies, development of new service offers, pricing and bundling as well as infrastructure consulting. NSP's consultants are respected thought-leaders in the networking industry and influence its direction through confidential engagements for industry leaders and through public appearances, whitepapers, and trade magazine articles. Contact NSP at [www.nspllc.com](http://www.nspllc.com).

---

## TABLE OF CONTENTS

<b>INTRODUCTION.....</b>	<b>1</b>
<b>REQUIREMENTS FOR A CONVERGED IP CORE NETWORK .....</b>	<b>2</b>
High Availability .....	2
Resource Allocations must be Separate for each Business Unit .....	3
Security .....	4
Network Management and Service Visibility .....	4
<b>EXAMPLE OF A VIRTUAL CORE IP NETWORK.....</b>	<b>5</b>
Residential Wireline Network .....	5
Business Data Services .....	7
Mobile Data Services.....	9
Wholesale IP Services .....	10
Comparison of Separate Core Networks with a Converged Core Network .....	10
<b>NETWORK DESIGN.....</b>	<b>12</b>
TX Matrix Plus Virtualized Core Routers .....	12
Optical Transport Integration .....	15
Network Management .....	16
<b>CONCLUSION .....</b>	<b>17</b>

## Introduction

Service provider networks are in the process of undergoing a massive network transformation from legacy circuit and packet networks to converged IP/MPLS networks. The vision of IP convergence offers the promise of a single network infrastructure that is scalable, reliable, secure, and cost effective to support new and emerging network services. While the rollout of IP/MPLS networks is moving full steam ahead, however, many service providers have not yet consolidated all network services and business units on a single IP core. Some of the reasons that service providers have maintained separate networks are:

- Service provider business units in residential wireline, wireless, and business services are set up as separate P&L centers and as a result have maintained control of their own network infrastructure
- Separate business units have different network requirements and are reluctant to share one of their most important resources, the core IP backbone with other business units
- The complexity of combining networks has created roadblocks to convergence
- Until recently, routers have not had the necessary scalability and reliability to support the requirements for service and application variability and growth in a converged network
- IP networks raise a real set of security concerns for service providers – the bigger the IP network the greater the security concerns (see the *Network Security Handbook for Service Providers*<sup>1</sup>)

While these are valid reasons for maintaining separate core IP networks, it is clear that over the long term service providers must leverage the economies of scale of a shared IP network to reduce network capital and operating expenses. This is increasingly important as IP traffic continues on its rapid growth path. As trunk bandwidth increases in IP core networks from 10GbE/OC192 to OC768 and eventually to 100GbE it will be necessary to share these high speed trunks among all services and business units to effectively scale the networks in a cost effective manner. Higher speed trunks will result in increased network scalability and performance but will also result in increased capital and operations expenses. In order to fully take advantage of OC768 and future 100GbE trunks it is necessary to maximize statistical multiplexing gain by combining all services and businesses on shared high speed trunks.

The key to effective, secure consolidation of core IP networks is IP router virtualization. A virtual core IP router network has most of the same characteristics of a physically separate IP network. Control plane and routing resources are dedicated to the virtual network, traffic engineering guarantees trunk bandwidth to the virtual network, software upgrades and maintenance windows on each of the virtual networks are independent, virtual trunks are separated at Layer 2 providing security, and failures in one virtual

---

<sup>1</sup> *Network Security Handbook for Service Providers*, Network Strategy Partners & Juniper Networks, <http://www.juniper.net/solutions/literature/misc/710095.pdf>

network will not impact other virtual networks. System virtualization in a core router allows service providers to effectively consolidate IP networks while addressing the concerns of the business units sharing the network.

Clearly, core IP networks need to be consolidated to maximize economies of scale, however, business units must continue to have the flexibility, reliability, and security associated with managing their own IP network. This paper presents the key requirements for a converged IP/MPLS core and explains how a virtual IP network can satisfy these requirements. An example of a hypothetical converged network design for a Tier 1 service provider is presented to aid the reader in understanding how a converged IP virtual network is implemented.

## **Requirements for a Converged IP Core Network**

A service provider's core IP network provides the backbone for many current and future services. In order for separate business units (such as residential, wireless, business data service) inside a large telecommunications company to accept a shared IP backbone the following requirements must be met.

### **High Availability**

The consolidated IP core network must be extremely reliable providing 99.999% availability for each of the business units and associated services sharing the core network. Business units, furthermore, must have complete flexibility in scheduling maintenance windows for software upgrades and should not be impacted by maintenance windows scheduled by other business units. Routers in the core IP network must be designed for high availability in order to meet these requirements. Some of the availability requirements for routers are:

- Redundant routing engines allowing a secondary routing engine to take over for the primary engine in the case of a failure
- Non-stop routing providing for continuous routing updates in the case of a routing engine failure so that neighboring routers do not notice a change in network topology
- In-service software upgrades allowing routing software to be upgraded without disrupting service
- Modular routing software with memory protection which allows software to gracefully recover in the presence of errors
- Redundant switch fabrics
- Redundant system power and cooling

Furthermore, converged IP networks must be architected for resiliency such that alternate paths are available to accommodate link or node failures. Core routers must support standard protocols for resiliency such as:

- MPLS Fast Reroute: Fast protection switching over alternate Label Switched Paths (LSP) similar to SONET or SDH 50 msec ring protection switching

- G.8032 Ethernet Ring Protection: Provides sub-50ms protection for Ethernet traffic in a ring topology and at the same time ensuring that there are no loops formed at the Ethernet layer.
- Virtual Router Redundancy Protocol (VRRP): Uses an IETF standard protocol to advertize a virtual router as a default gateway; if a router fails a second router is assigned to be the gateway.
- Bidirectional Forwarding Detection (BFD): An IETF standard specifying how links or LSP tunnels detect faults between two forwarding engines.

High availability requirements also dictate that core routers must support strong operations, monitoring, and repair capabilities. For example:

- IEEE 802.3ah: Specifies standards for Carrier Ethernet in the First mile which includes OAM and repair capabilities
- IEEE 802.1ag: Defines Ethernet service layer OAM

These are necessary requirements for reliability in the core IP network, however, they are not sufficient conditions to address all potential concerns regarding consolidation. Router system virtualization also is required to separate the control plane, forwarding plane, operations, and administration of each virtual routing network. By separating these functions, network designers can ensure that events in one network do not impact another network. For example, if one virtual network was in the process of upgrading the router software and experienced a problem due to an administrator error, this problem would not impact the other virtual networks because the control planes, forwarding planes, and management planes of these networks are completely separate.

### **Resource Allocations must be Separate for each Business Unit**

One of the primary objections that individual business units have in a large service provider organization is sharing network infrastructure resources with other business units. Shared resources can lead to potential problems with bandwidth allocation, fault management, and network security. Therefore, a key requirement in a converged IP core is that network resources are dedicated to business units operating a virtual network. Network resources include:

- Line cards responsible for packet forwarding
- Control cards responsible for routing updates, system control, and network management
- Network trunk bandwidth – or virtual dedicated trunks

There are two ways to separate system resources – physical separation and logical separation. At the edge of the service provider network, IP service routers support large numbers of enterprise VPNs and, therefore logical separation of resources is typically used with MPLS VPN software<sup>2</sup>. However, the requirements for separation of network resources in the IP core network are different from the service edge requirements. In the

---

<sup>2</sup> MPLS VPNs are standardized by the IETF in RFC-4364 (formerly RFC-2547.)

core, resources must be separated for a relatively small number of business units. Also each business unit operates a significant network in its own right with demanding performance, scalability, service, and security requirements. Network core virtualization, therefore, should use *physical separation* of resources rather than logical separation of resources. More specifically packet forwarding line cards and control plane cards should always be dedicated to a specific virtual network. This will ensure that administrative activities, performance monitoring, or network faults are local to a virtual network. Network trunks should optionally be physically or logically dedicated to a virtual network. For some 10GbE or OC192 trunks, it might make sense to physically dedicate the entire trunk to a virtual network if that network has a high level of trunk utilization. For large trunks such as OC768 or future 100 GbE, however, it makes more sense to logically dedicate a portion of the trunk bandwidth to a virtual network to maximize statistical multiplexing gain on these expensive network transmission resources.

The system components in the router that are shared by all virtual networks include:

- Chassis and common equipment
- System Power and cooling
- Redundant switch fabrics

A combination of physical and logical separation of network resources will ensure an optimal approach to each business unit's network service, scalability, service integrity, service manageability, and security objectives.

## Security

Network security is a top requirement in all networks – especially in service provider IP core networks. Malicious network attacks caused by worms, bots, and other malware could cause network outages resulting in catastrophic service outages. Network security is one of the central concerns that business units have about IP network consolidation. For example, a business unit offering MPLS VPN services to major banks would not want to share an IP network with a business unit offering residential Internet services. It is essential, therefore, that the IP virtual networks have strong separation of network resources so that security is managed independently in each virtual network. It is possible to completely separate the virtual networks by using physical separation of the forwarding and control planes. Trunks are logically shared at Layer 2 which also provides a strong measure of separation and network security. This provides an environment with security equal to that of separate physical networks.

## Network Management and Service Visibility

Each business unit in a large service provider organization is responsible for P&L and therefore must have complete visibility and control over network services. A key requirement in a converged network, therefore, is that each business unit has comprehensive network management capabilities and service visibility over its own virtual IP network. These functions should include:

- Network monitoring and reporting

- Fault management
- Performance management
- Security management
- Administration and configuration management
- IP service layer management

Furthermore, business units should not have visibility into other virtual networks belonging to separate business units. Virtual network management and service visibility allows the business unit owning the virtual network the same benefits of control that it would have operating its own physically separate network.

## Example of a Virtual Core IP Network

This section of the paper presents an example of a Tier 1 service provider using a virtual core IP routing network to consolidate network services for multiple business units. A Juniper T Matrix Plus virtualized multiservice core router is used to consolidate a core IP network for four business units:

1. Residential wireline business unit offering triple play services
2. Business data services business unit offering Carrier Ethernet, MPLS VPN, and FR/ATM services
3. Mobile business unit offering 3G mobile broadband services
4. Wholesale IP business unit offering MPLS VPN and Carrier Ethernet services

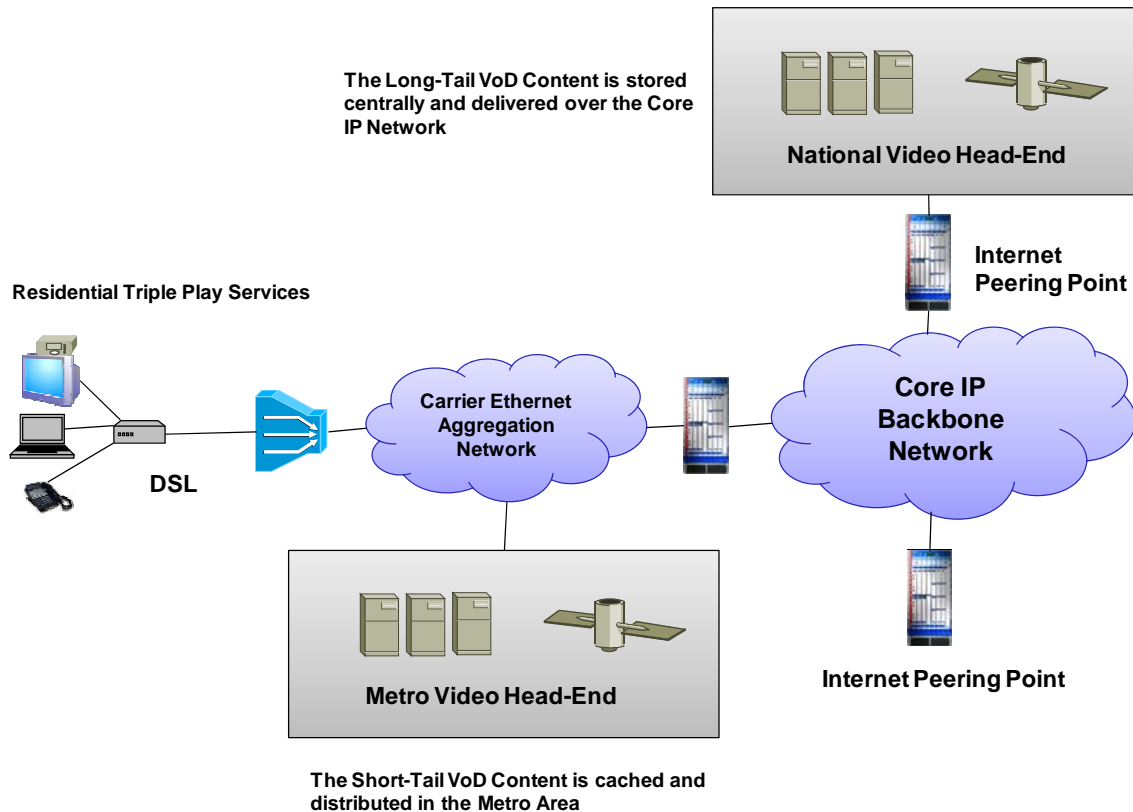
The network architecture is defined for each business unit and a high-level design of the virtualized core IP node is presented.

### Residential Wireline Network

The following set of residential data services are delivered by the Residential Services Business unit over a broadband DSL access network:

- High Speed Internet (HSI)
- Broadcast Television delivered over DSL (IPTV)
- Video-on-Demand delivered over DSL (VoD)
- Voice-over-IP (VoIP)

The architecture of the residential triple play network is presented in Figure 1. All services are delivered over a DSL packet service to the household. The DSL access network is aggregated in a metro area using a Carrier Ethernet aggregation network. Broadcast IPTV is delivered from a metro video head-end and, therefore does not traverse the core IP backbone. Short-Tail VoD is accessed frequently and therefore is cached in the metro area. Short-Tail VoD content does not traverse the core network. Long-Tail VoD content is accessed less frequently and therefore it is stored in a central database and transmitted across the IP core network to all metro areas. All Internet service and VoIP service is carried over the core IP backbone network and distributed to all metro areas.

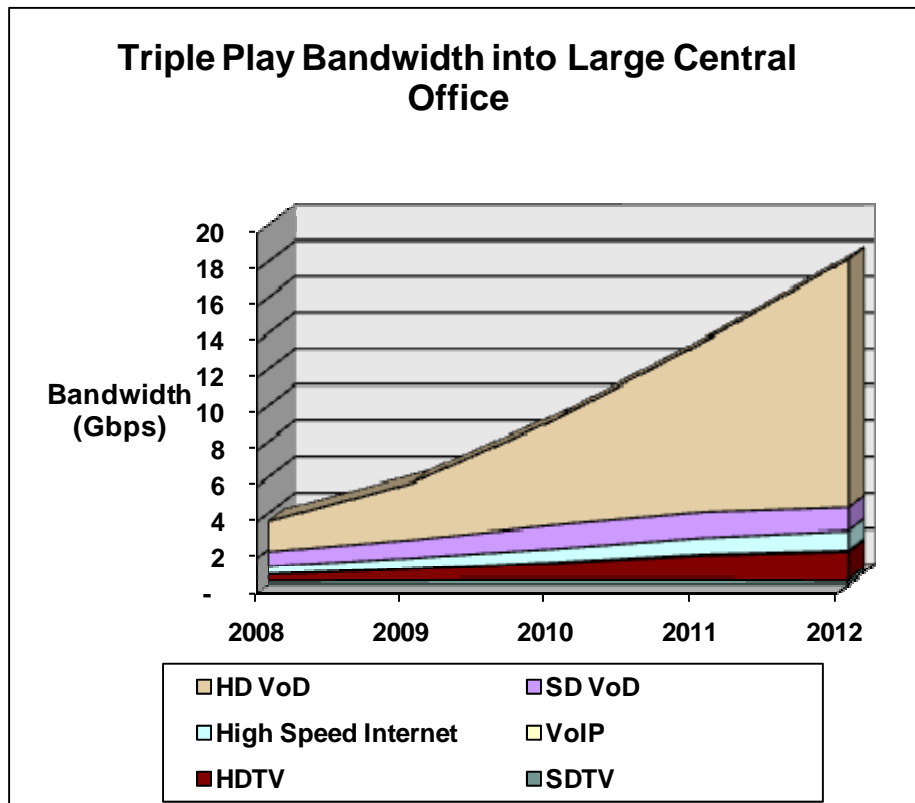


**Figure 1**

Internet traffic is the primary driver of bandwidth for the residential core IP network and is mainly composed of the following applications:

- Web Browsing
- Email
- P2P
- Gaming
- Internet TV
- Internet VoD
- YouTube
- Web Cams and Video Conferencing
- Video upload (SlingBox)

Residential Internet traffic has been increasing at a rapid rate primarily due to emerging multimedia applications such as video download/upload and P2P file sharing. This trend is expected to continue and potentially accelerate. A recent study by Network Strategy Partners LLC has predicted that over the next five years VoD will drive network traffic (see Figure 2). High Definition (HD) Video-on-Demand is the primary engine for growth due to the high data rates associated with HD data streams. VoD can be delivered directly over the service provider's IP network or alternatively, over the Internet.



**Figure 2**

Clearly the IP core network must be sufficiently scalable and flexible to support rapidly growing bandwidth requirements and distributed traffic patterns.

### Business Data Services

A Business Data Services business unit is the second business unit in the design example. This organization is responsible for selling and operating wireline data services to enterprises and SMBs. A network architecture for the business IP network is depicted in Figure 3. The services offered include:

- Legacy Frame Relay and ATM (FR/ATM) services
- Carrier Ethernet transport services
- MPLS VPN services
- Internet Service
- New age services such as Cloud Computing
- Software As a Services
- HD video conferencing

Frame Relay and ATM services are connection oriented services – enterprises lease Frame Relay or ATM PVCs and use them for Layer 2 data transport. While many enterprises are migrating to newer MPLS VPN and Carrier Ethernet services, there is still a very large installed base of FR/ATM. Transport of FR/ATM services in the IP core

network uses MPLS Pseudo wires – frames and cells are encapsulated in MPLS packet and sent across an MPLS Label Switched Path (LSP), thus emulating a FR/ATM PVC.

Carrier Ethernet services are next generation connection oriented transport services. Carrier Ethernet is similar in many respects to FR/ATM, however, while Frame Relay typically operates over T1 or T3 circuits, Carrier Ethernet runs over Ethernet physical connections at data rates of 10 Mbps, 100 Mbps, 1 Gbps, and 10 Gbps. Carrier Ethernet services are defined by the Metro Ethernet Forum's standards.<sup>3</sup>

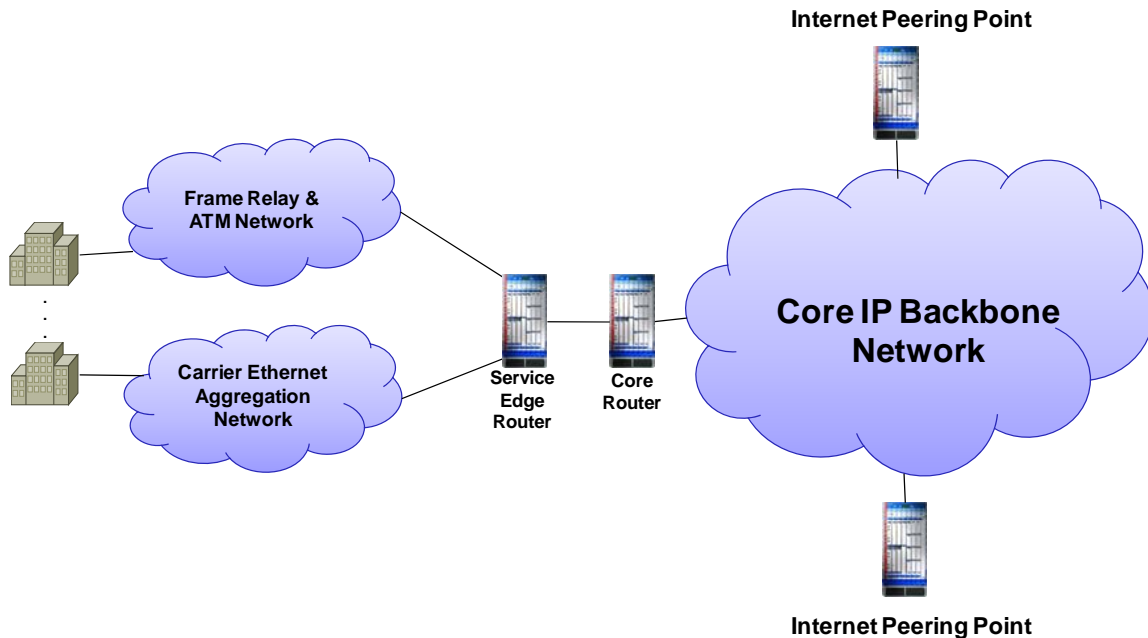
- E-Line – An Ethernet service type that is based on a Point-to-Point Ethernet Virtual Connection. Two E-Line services are defined:
  - EPL (Ethernet Private Line) – This is a very simple point-to-point service characterized by low Frame Delay, Frame Delay Variation and Frame Loss Ratio.
  - EVPL (Ethernet Virtual Private Line) – This is a point-to-point service wherein service multiplexing (more than one Ethernet virtual circuit) is allowed.
- E-LAN – An Ethernet service type that is based on a Multipoint-to-Multipoint Ethernet Virtual Connection.

Carrier Ethernet can be used as an Internet access technology or as a private Layer 2 enterprise transport technology. If it is used as an Internet access technology then the Layer 2 Ethernet connections will terminate on the services edge network which provides full access to the Internet. If it is used as an end-to-end transport technology Pseudo Wires are used to encapsulate the Ethernet frames and transport them across the IP Core network. Pseudo Wires provide traffic engineering, QoS, and Layer 2 connection oriented capabilities in a similar manner to native Carrier Ethernet.

MPLS VPN services are rapidly replacing FR/ATM services. This service provides an enterprise with a private virtual IP network. The network uses private IP addresses and runs separate instances of routing protocols to separate logical routers. MPLS LSPs are used across all physical connections to provide virtual trunks with QoS for the VPNs. The Services edge network must run the separate instances of routing and forwarding specified by RFC-4364 (formerly RFC-2547.) The core network provides connectivity between all enterprise MPLS VPNs and also provides Internet connectivity to VPNs.

The final service offered by the data services business unit is Internet service. The Internet is directly connected to the core IP network and is accessed by Frame Relay/ATM, Carrier Ethernet, or IP MPLS connections. Some Internet sites are local to the service provider's IP backbone and other sites are accessed by traversing the Internet Peering Points.

<sup>3</sup> See MEF 6 – Ethernet Service Definitions – Phase I, Metro Ethernet Forum, June 2004 for the technical specification of E-Line and E-LAN service types.

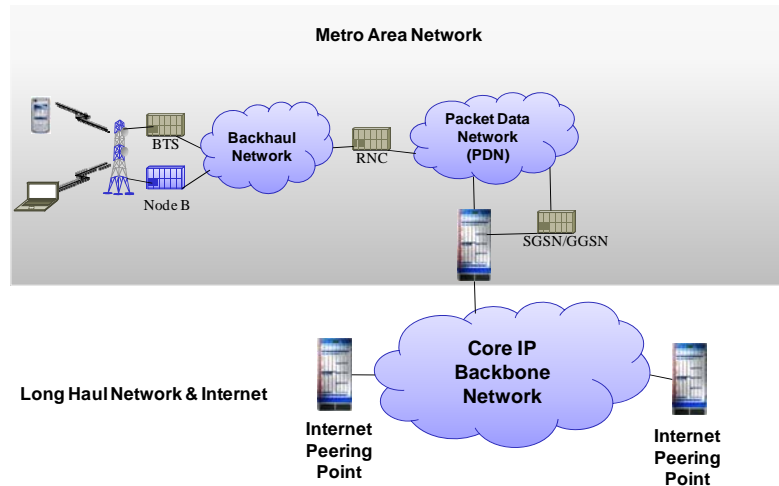


**Figure 3**

### Mobile Data Services

A business unit offering mobile data services to customers also is considered in the design example. These services are delivered over 3G/4G mobile networks. An example of such a network is depicted in Figure 4. Wireless end devices (phones, PDAs, laptops) are connected to the Radio Access Network (RAN) using 3G or 4G radio technology. Data and voice traffic is sent across a backhaul network (typically T1/E1, ATM, or Ethernet). The Radio Network Controller (RNC) then separates voice and data traffic sending the data traffic across the service providers Packet Data Network (PDN). The PDN is typically a metro network providing transport from RNCs to SGSNs and GGSNs. The SGSN/GGSNs then provide data services and a gateway to the Internet. The core IP network is used to interconnect SGSN/GGSNs in multiple metro areas and provide Internet connectivity.

In 4G networks such as LTE and WiMax, all voice traffic is sent as VoIP and therefore the backbone is a pure IP network. For both 3G and 4G networks, the key requirements of the core IP network are to provide reliable, scalable IP transport with QoS guarantees for voice and video.



**Figure 4**

## Wholesale IP Services

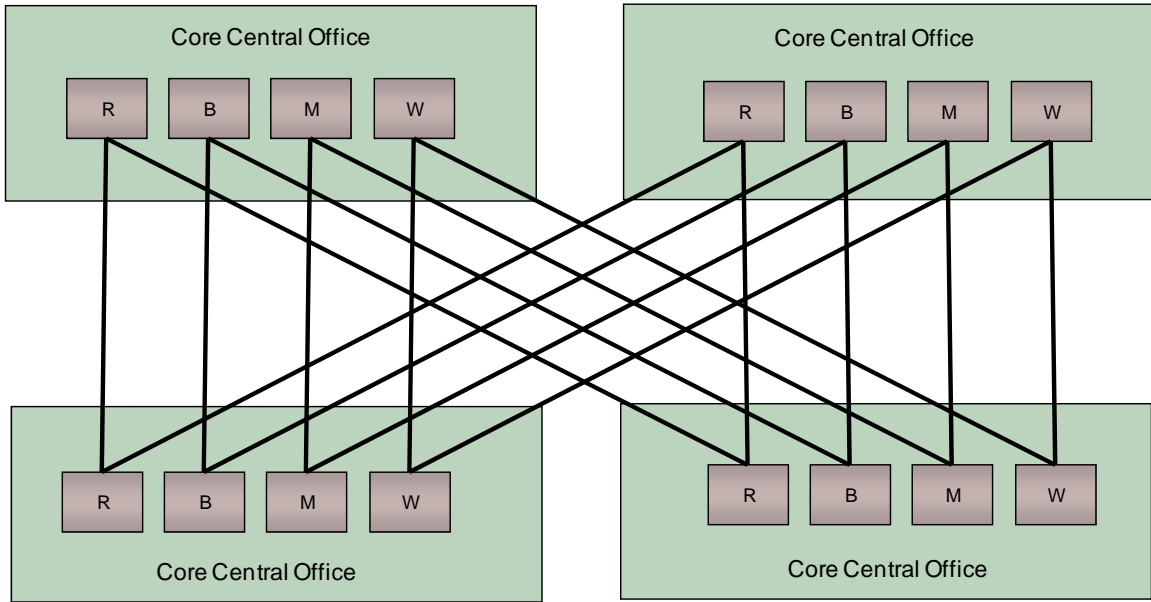
The final business unit considered in the example is a wholesale services group. This business unit sells wholesale packet services to other service providers. The wholesale services that use the core IP network are:

- Carrier Ethernet Services
- MPLS VPN Services
- Wholesale Internet services
- Wholesale Voice transit service

These services are similar in nature to business services offered to enterprise and SMB customers with the primary difference being that wholesale services are sold to other service providers who in turn offer residential or business retail services to their own customers.

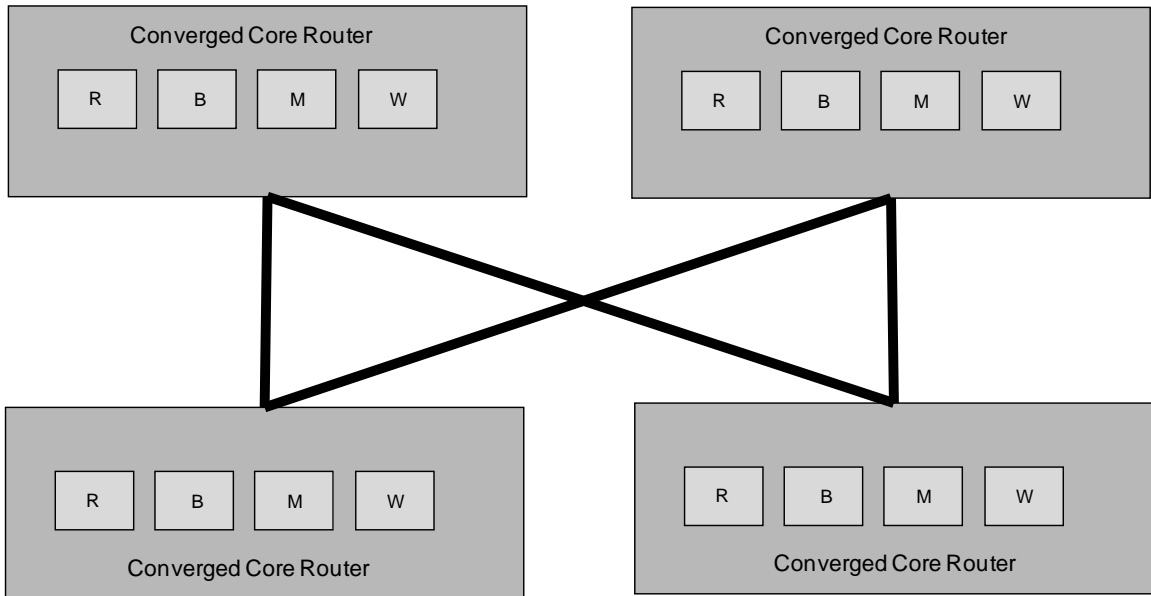
## Comparison of Separate Core Networks with a Converged Core Network

A comparison of separate core networks with a converged core network is made in Figure 5 and Figure 6. In Figure 5 each of the separate business units (residential, business, mobile, and wholesale services) use separate routers and links to create four physically separate networks. This results in a large number of trunks and routing platforms that must be supported by each business unit. In Figure 6 the same configuration is depicted using converged core routers with virtual routers dedicated to each business unit. This architecture results in fewer chassis and trunks. Furthermore, as a result of statistical multiplexing gain, less total bandwidth is required to interconnect each core central office.



R = Residential Services Core Router  
 B = Business Services Core Router  
 M = Mobile Services Core Router  
 W = Wholesale Services Core Router

**Figure 5**



R = Residential Services Virtual Core Router  
 B = Business Services Virtual Core Router  
 M = Mobile Services Virtual Core Router  
 W = Wholesale Services Virtual Core Router

**Figure 6**

## Network Design

This section of the paper provides an overview of a converged, virtual network design providing core IP services for each of the business units and network architectures described in the previous section. The design example uses a Juniper Networks TX Matrix Plus multi-chassis virtual router.

### TX Matrix Plus Virtualized Core Routers

The TX Matrix Plus multi-chassis virtualized Core router consists of three main subsystems:

1. Two or more T1600 Chassis
2. A JCS1200 Chassis with routing engines
3. Switching Card Chassis

The key virtualization features used in this example are:

- FPCs in the T1600s are assigned either to a virtual system or are assigned as shared uplink cards to be shared by multiple virtual routers
- Routing engines in the JCS1200 manage multiple FPCs in the Line Card Chassis.
- Trunk interfaces on the T1600 can be either dedicated or shared by multiple virtual routers

In this example, the core network is designed with five virtual networks:

1. Internet Virtual Network (Shared by each business unit)
2. Residential Wireline Virtual Network
3. Business Services Virtual Network
4. Mobile Services Virtual Network
5. Wholesale Services Virtual Network

An example of a converged core IP node is presented in Figure 7. Each converged node assigns two redundant routing engines in the JCS1200 to each virtual router, and one or more line cards in the T1600 to each virtual router. The JCS1200 and the T1600s are physically connected to each other by the TX Matrix Plus redundant switch chassis. The colored lines in Figure 7 represent logical connections between the routing engines in the JCS1200 and forwarding cards in the T1600, however, the physical connections are represented by the black lines connecting the JCS1200 and T1600 to the TX Matrix Plus.

An arbitrary number of forwarding cards in the T1600 can be assigned to each virtual router based on the network traffic requirements of the business unit and virtual network. Each virtual router has dedicated line cards with access ports, packet forwarding engines, and dedicated routing engines. Trunks between core routers can be either dedicated or shared as virtual trunks. A possible approach to migration of separate networks to a converged virtual core is to start out with dedicated physical trunks for each of the networks. In this scenario each virtual router would have complete dedicated physical resources including trunks. Over time as the network grows, virtual trunks are introduced

to share optical transport bandwidth and therefore reduce transport expenses. Sharing of trunks will become especially important as trunks migrate to OC768 and 100 GbE.

In this example physical links between core routers contain five virtual trunks depicted in Figure 8:

- Residential Services Virtual Trunk
- Business Services Virtual Trunk
- Mobile Data Services Virtual Trunk
- Wholesales Services Virtual Trunk
- Internet Virtual Trunk

Each virtual trunk is separated from the other virtual trunks using VLAN/DLCI identifiers (similar to Carrier Ethernet and Frame Relay). The virtual trunks also are allocated guaranteed bandwidth over the link using the router's QoS capabilities and Weighted Fair Queuing. The combination of Layer 2 separation of bandwidth and QoS provides the same level of security and performance in virtual trunks that organizations have come to expect from physical trunks.

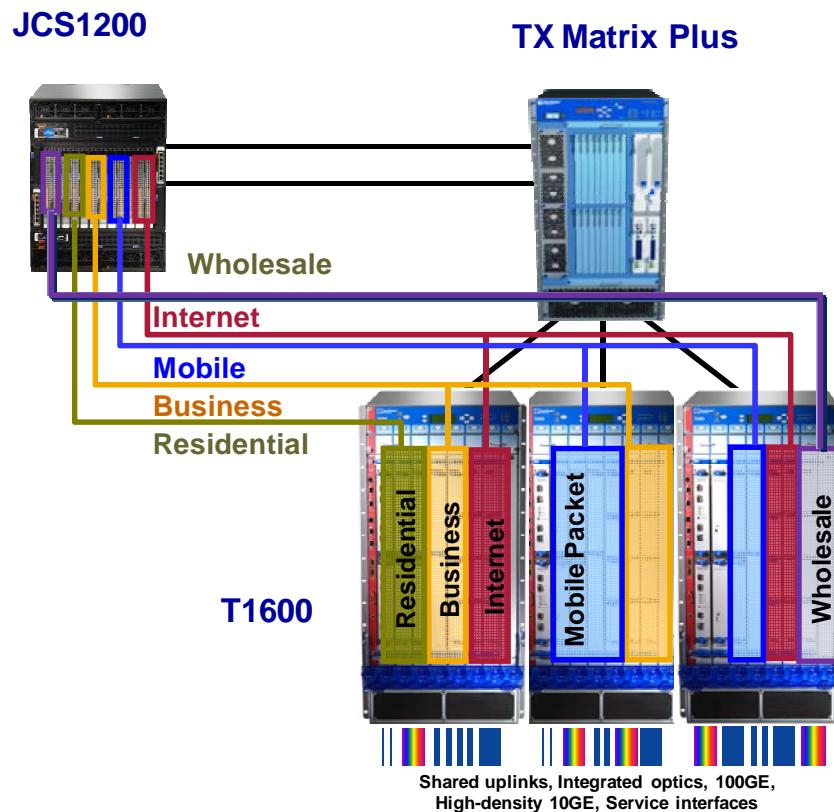



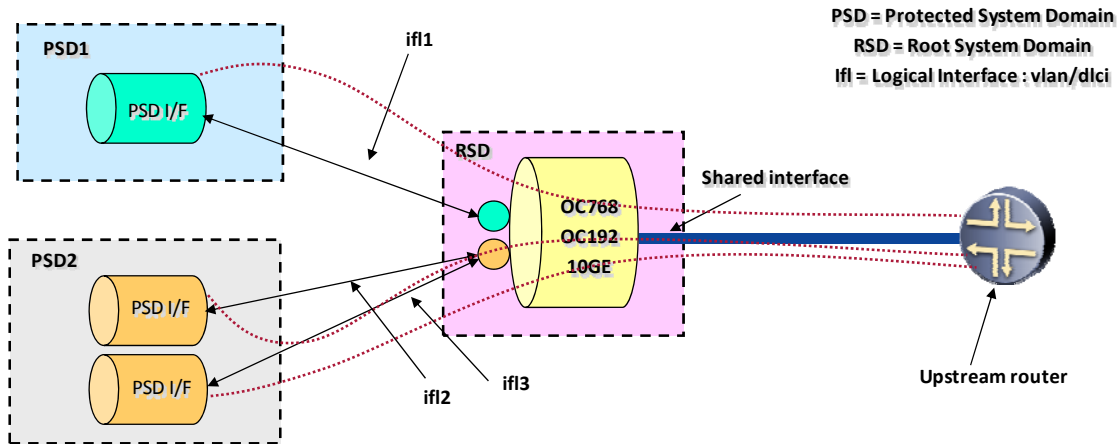
Figure 7



Physical Trunk with embedded Virtual Trunks

**Figure 8**

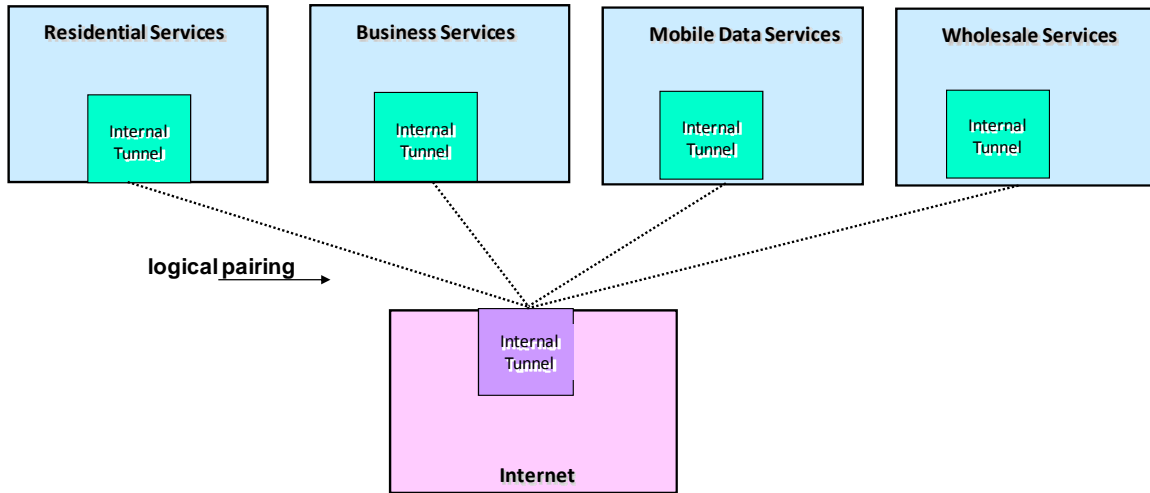
Each of the virtual routers inside the core routing node is part of a Protected System Domain (PSD). Forwarding engines and routing engines that are part of a PSD have full interconnectivity with each other but not with other PSDs. The shared physical trunk interfaces belong to the Root System Domain (RSD). The RSD interfaces are controlled by the local routing engines on the T1600 chassis and provide shared links for all the virtual routers. Figure 9 provides an overview of how the PSDs interconnect to the RSD providing the shared trunk. In order to connect to the shared trunk, the forwarding card that is part of a PSD (for example the residential services PSD) uses an internal tunnel across the switching fabric. This is a hardware based tunnel that is secure and provides guaranteed bandwidth. The tunnel provides a connection to the virtual trunk that is identified at Layer 2 using a combination of a VLAN/DLCI. This mechanism ensures both performance and security in the virtual trunks.



**Figure 9**

It is also possible for different virtual routers to interconnect to each other using internal system tunnels. In this example, all of the virtual networks run by the business units need to interconnect to the Internet. Therefore every virtual router (or PSD) needs to interconnect to the Internet PSD at every backbone node. This is illustrated in Figure 10 – each virtual router tunnels across the switch fabric to directly access the Internet virtual router. This allows Internet connectivity for all virtual routers at each node in the backbone network. The Internet virtual network connects directly to various Internet web hosting sites as well as interconnects to other service providers via Internet Peering

Points. Typically at least two Peering Points in different locations are used for redundancy.



**Figure 10**

## Optical Transport Integration

The converged core IP network in this example also has tight integration with the optical transport network using G.709 and GMPLS. Traditional core IP routers and transport infrastructure are separate networks with completely different network planning, management, and administration. Core optical transport networks are composed of DWDM network elements with transponders, ROADMs, OADMs, and optical amplifiers designed for metro and long haul transport of services such as SONET, SDH, Ethernet, and Fiber Channel. The Juniper TX Matrix Plus provides tight integration between the core router and the optical transport network by integrating DWDM transponders directly into the router line cards. Trunks are carried on separate DWDM wavelengths using OTN, a high speed digital framing hierarchy established in the G.709 ITU-T standard. Furthermore optical channels and optical network paths can be automatically established by the router with the GMPLS; an IETF standard that uses many of the protocols developed for MPLS to automatically set up and tear down paths across an optical transport network. Key benefits of integrating the core router with the optical transport network are:

### Improved OAM&P

Integration of the core router with the optical network allows for circuit provisioning automation using GMPLS. G.709 also provides rich optical channel performance and fault monitoring capabilities that are similar in nature to SONET and SDH. Improved performance and fault monitoring can improve the core IP network availability.

### High Speed fault recovery using FRR

Router based fast reroute (FRR) is used to allow for very fast recovery in the case of network trunk failures. FRR switches traffic to standby MPLS LSPs in the case of a primary LSP failure allowing for SONET/SDH like failure recovery speeds. Failure recovery using MPLS based FRR is more cost effective and flexible than using channel protection at the optical layer. By integrating the router with the transport layer, the router uses G.709 OAM fault and performance monitoring to make fast reroute decisions.

### Network Management

One of the concerns regarding converging separate IP networks on a shared IP backbone is that each business unit will lose control of its network. In this example each business unit operates as a separate P&L center and therefore has concerns about turning over network operations to a central authority. Some of the business unit concerns regarding central network operations are:

- Loss of control over network operations including configuration, fault and performance management, and software upgrades
- Competition for technical resources with other business units
- Lack of control over charge-back rates for operations and technical services carried out by a central management organization

The TX Matrix Plus along with JCS 1200 addresses these concerns by allowing each business unit to manage its own virtual router while a central organization manages the shared network infrastructure such as the chassis, environment (power, cooling, space), shared trunks, and the optical transport network. Virtual routers can be independently managed because each virtual router uses physically separate routing and forwarding engines. Management functions that are business unit responsibilities include:

- Software upgrades
- IP router configuration
- IP router fault management
- IP router performance management
- IP router security management

Management functions that are central operations group responsibilities include:

- Environmental operations (power, cooling, space, and physical security)
- Common equipment management (chassis, switch fabric, shared trunks)

In this example the separate business units running residential, business, wholesale, and mobile networks manage their core IP networks separately. However, it is also possible for a central organization to take over all network management responsibilities if that mode of operations is a better fit with a service provider's strategic objectives and operations strategy.

## Conclusion

Large service providers must evolve network architectures to a converged IP core in order to minimize network capital and operations costs while maximizing service revenue and profitability. While most organizations understand the basic economics of a converged IP core, there is still some resistance to convergence due to differences in requirements and operations among different business units. By using scalable and secure virtual routers in the core network many of these barriers can be overcome. This paper presents a concrete example to illustrate how a typical Tier 1 service provider organization with multiple business units can use a virtual core IP network to converge its IP network while satisfying security and operations concerns of the business units.

In summary, the key benefits of the Juniper TX Matrix Plus virtual network architecture are:

- Complete separation of the forwarding plane and control plane for each virtual router
- Dedication of routing engine and forwarding engine cards to individual virtual routers (Physical separation vs. logical separation)
- Sharing of physical trunks with Layer 2 logical separation and QoS guaranteeing performance and security
- Ability to run different versions of software on each virtual router and independently upgrade software on routers
- Ability to independently manage each virtual router