

# Enabling Solutions in Cloud Infrastructure and for Network Functions Virtualization

---

Gateway Use Cases for Contrail Virtual Networks with MX Series Routers

## Table of Contents

Executive Summary .....	3
Introduction.....	4
Introduction to Contrail Virtual Networks .....	4
Key Gateway Features for MX Series Routers.....	6
Use Cases .....	6
Internet Gateway .....	6
Interdomain Gateway .....	7
Data Center Interconnect .....	9
Internetwork Gateway .....	10
Hardware Appliance Insertion.....	11
Virtual Private Cloud for Enterprise VPNs.....	12
Enterprise VPN Service Chains (Virtual CPE).....	13
Service Chaining Gateway for Subscriber Networks.....	14
Conclusion .....	15
About Juniper Networks.....	16

## Table of Figures

Figure 1: MX Series router supporting several gateway use cases for Contrail virtual networks.....	3
Figure 2: Contrail leverages protocols and architectural principles from L3 VPNs as used in WANs.....	5
Figure 3: An MX Series router acting as a gateway to the Internet for a set of Contrail virtual networks.....	7
Figure 4: MX Series platform using VRFs to connect between virtual networks in different Contrail domains .....	8
Figure 5: MX Series platform using VRFs to connect virtual networks in different data centers.....	9
Figure 6: MX Series routers acting as a gateway to bare metal servers connected to VLANs or VMs connected to VLANs .....	10
Figure 7: MX Series platform enabling insertion of a physical appliance between two virtual networks.....	11
Figure 8: MX Series platform acting as a gateway between enterprise VPNs and Contrail virtual networks in a data center .....	12
Figure 9: Virtual appliances can be inserted into enterprise VPNs to support on-demand services.....	13
Figure 10: Virtual appliances connected as a service chain into a VPN using virtual networks.....	14
Figure 11: Service chaining through service delivery gateway with service- and application-aware steering .....	15

## Executive Summary

Cloud automation and Network Functions Virtualization (NFV) are potentially disruptive and transformative networking technologies, and they have captured the attention of network operators and hosting service providers around the world. There are many active programs investigating how best to take advantage of these technologies to enable delivery of new services and improve the quality and operational efficiency of existing ones. The justification for implementing any new technology requires a business case that demonstrates the potential for new revenue, CapEx reduction, or OpEx reduction.

The combination of Juniper Networks® Contrail virtual networking solution and MX Series 3D Universal Edge Routers can deliver value on all three of these axes based on the impact of Moore's law increasing the compute power of the x86 platform (including network processing), the powerful automation features that cloud architectures deliver, and the proven industry-standard protocols supported in both Juniper Networks Junos® operating system and Contrail. Juniper customers using Contrail virtual networking in combination with MX Series routers improve the flexibility of how their networks are architected, bringing opportunities for creating new service combinations that meet the needs of an ever more demanding world.

This document describes how MX Series routers can be deployed to perform gateway functions in a variety of use cases when Contrail is managing virtual networks in service provider and enterprise environments. The key capabilities that the MX Series provides are:

- The ability to interoperate between different virtual network domains (OpenStack, VMware, etc.);
- The ability to bridge geographically distributed virtual network domains;
- The ability to connect virtual networks and physical networks (enterprise VPNs, public Internet, etc.);
- The ability to anchor virtual service chains when network services are applied to specific traffic flows in mobile access networks, broadband/fixed access networks, and enterprise VPNs; and
- The ability to anchor both physical and virtual network elements in a service chaining environment

Contrail virtual networks are implemented as network overlays that can be deployed on existing data center switched networks. Contrail uses standards-based control protocols and encapsulations, including BGP, MPLS, and Virtual Extensible LAN (VXLAN), so interoperating with MX Series routers is straightforward and seamless, and leverages the same Junos OS features that are already in use in today's most demanding networks.

MX Series routers can scale to meet any business need, and they support a rich set of features for resilience and high availability that are not available in software-based gateways. The combination of optimized silicon, systems built using that silicon, and Junos OS software delivers high performance with low power per Gbps of throughput.

This document will also show how MX Series routers can be used as gateways between Contrail domains, between Contrail domains and other physical and virtual networks, and how Contrail domains can be leveraged to extend gateway functionality in physical networks.

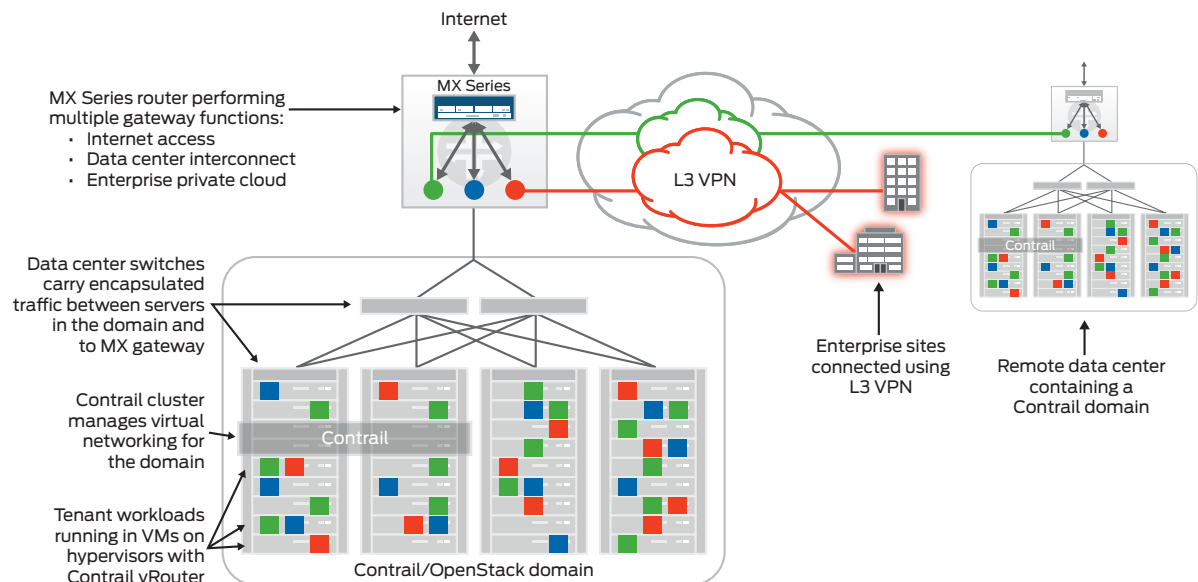


Figure 1: MX Series router supporting several gateway use cases for Contrail virtual networks

## Introduction

The use cases explored in this white paper fall into three general categories:

- **Data Center:** Gateway use cases where virtual networks need to be connected to other networks within the same data center or other data centers, or to the Internet
- **Enterprise VPN Services:** Where cloud infrastructure is used to augment enterprise VPNs with virtual private cloud and network services such as firewall or intrusion detection
- **Service Chaining for Subscriber Networks:** Where traffic flows are directed into chains of virtualized services, such as optimizers and parental control, using subscriber- and application-aware filtering

The use cases described in this document are as follows:

Category	Use Case	Description
Data Center	Internet gateway	Provide Internet connectivity to a set of virtual networks in a data center by linking a gateway router into Contrail virtual networks
	Interdomain gateway	Enable connectivity between virtual networks in different Contrail administrative domains in a data center
	Data center interconnect	Allow geographically separated Contrail domains to easily be linked since the virtual networks inside each data center use the same protocols as the WAN between them
	Internetwork gateway	Provide gateway between Contrail virtual networks and VLAN-based networks, including VMware virtualized servers and physical servers
	Appliance insertion	Insert hardware appliances such as firewalls and load balancers between Contrail virtual networks
Enterprise VPN Services	Virtual private cloud	Enable access to virtual private cloud infrastructure from within a VPN while preserving the security and quality-of-service (QoS) guarantees of the VPN
	Enterprise VPN service chains (vCPE)	Use virtualized service chains to support on-demand network services for enterprise customers without deploying hardware onsite
Subscriber Networks	Service chains	Use subscriber- and application-aware steering into virtualized service chains in mobile and wireline networks to support new services and improve operational efficiency

Each use case illustrates a different facet of what is possible using the Contrail solution in combination with MX Series routers. The flexibility and power of Junos OS in the MX Series routers allows multiple gateway use cases to be combined as needed to satisfy business and technical requirements.

## Introduction to Contrail Virtual Networks

The founding principle of Contrail was to take the protocols and architectural principles that the world relies on in the world's largest physical networks, and apply them inside data centers. This enabled the Contrail team to leverage the fact that many problems of scaling, resilience, and performance in networking had known architectural solutions that had been thought through by leading vendors and operators in standards bodies and were in everyday use in core networks.

Supporting industry-standard protocols such as BGP and MPLS allows Contrail virtual networks to easily interconnect with today's physical networks so that cloud environments using Contrail can seamlessly access other virtual and physical networks. Additionally, the benefits of cloud automation and virtualization can be applied in new ways to networks themselves, which has the potential to radically improve the flexibility of networks and allow a raft of new services to be offered. The use cases described in this document cover both scenarios.

Figure 2 shows how the concepts of route exchange, encapsulation tunnels, and scoped route lookup in virtual routing and forwarding (VRF) have been carried over from the implementation of Layer 3 VPNs in physical networks into the virtual world. A VRF is a virtual instance of a router within which local route lookup is performed, and it is the fundamental building block of VPNs. A physical router may contain thousands of VRFs, each corresponding to a different VPN that is connected to that router.

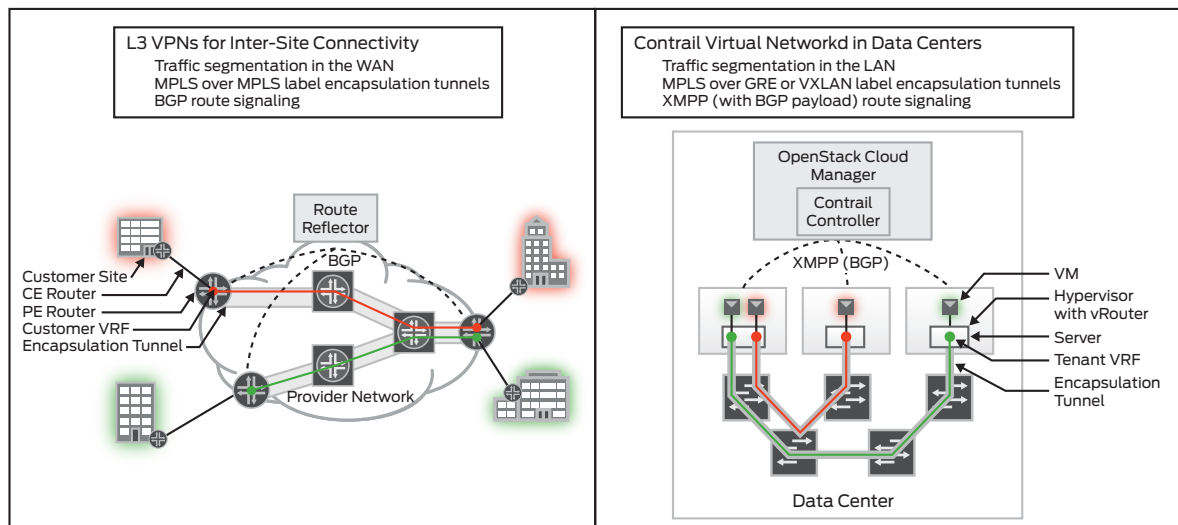


Figure 2: Contrail leverages protocols and architectural principles from L3 VPNs as used in WANs

In the physical network, the route reflector performs exchange of routes between VRFs in different routers with matching VRF identifiers (route targets). The routes may be statically configured, or they may have been learned through peering with other routers or route reflectors. The result is that all sites connected to VRFs with the same route target have routes to each other. Traffic is passed between sites by means of encapsulation tunnels between provider edge (PE) routers that take customer traffic through the provider network without that network having visibility into the customer source and destination.

The Contrail controller is integrated with a cloud management system such as OpenStack, and its function is to ensure that when a virtual machine (VM) is created, it has network connectivity based on the policies applied to it. Virtual machines are associated with virtual networks (groups of IP addresses), and network policies identify which virtual networks should be able to communicate with each other.

The implementation of network policy happens in the Contrail vRouter, which is installed into each hypervisor and provides Layer 3 VPN routing in place of the hypervisor's native Layer 2 switching capability. When a virtual machine is created by the cloud manager, the Contrail controller configures routes in the vRouter connected to the VM just created; these allow the new VM to reach other VMs in the same virtual network, as well as VMs in other virtual networks that network policies allow them to communicate with. Additionally, the controller adds a route to the new VM in all vRouters connected to VMs that can communicate with it, allowing reverse traffic flows.

The Contrail controller effectively performs the function of route reflection in software and publishes the results to vRouters, whereas an actual route reflector may initiate a series of interactions with routers before reaching steady state.

The Contrail controller communicates with vRouters using Extensible Messaging and Presence Protocol (XMPP), which is a robust, publish/subscribe messaging protocol. The messages are essentially a superset of BGP route updates, with the additional ability to create VRFs and install routes to local interfaces. Additionally, Contrail controllers support native BGP, and are able to peer with physical routers and route reflectors.

When a packet is sent from one VM to another, the source vRouter looks up the address of the destination VM and will get back a next hop, which is the address of the server that the destination VM is running on, as well as a label which identifies the interface in the vRouter to which the destination VM is connected. The original packet is encapsulated with a new header that contains the label and destination. The encapsulation protocol can be MPLS over generic routing encapsulation (GRE) or VXLAN. When the encapsulated packet is sent into the data center switching network, the switches only see the outer header and direct the packet to the destination server. When the packet reaches the server, the outer header is stripped off by the vRouter and the label is used to identify the interface of the destination VM. The original packet is then sent to the VM interface.

The concepts of route distribution, tunnel encapsulation, and interface labeling used in Contrail virtual networks are exactly the same as used in today's wide area networks.

Contrail controller software is typically deployed in a distributed fashion for scale and high availability, with multiple instances of each software component ensuring that there is no single point of failure. This is referred to as a Contrail cluster. The group of servers whose virtual networking is managed by a Contrail cluster is referred to in this document as a Contrail domain.

The servers are connected by a set of Layer 2 switches that carry encapsulated traffic between Contrail vRouters running on each server in its hypervisor. Details of the Contrail architecture and its operation can be found in *OpenContrail Architecture Documentation*<sup>1</sup>.

## Key Gateway Features in MX Series Routers

MX Series routers are the latest products to benefit from the rich heritage of Juniper routers and Junos OS. Juniper has been a leader in creating the protocols that run the major networks that form the Internet, as well as service provider and business enterprise networks. MX Series routers support a comprehensive range of protocols and features used in the gateway use cases described later in this document.

The key features of the MX Series routers that are used in the gateway use cases are listed in the table below.

Table 1: MX Series Gateway Features

Use Case	Description
Multiprotocol Border Gateway Protocol (MP-BGP)	Control protocol that allows routes to be exchanged between network domains that may employ different networking protocols. This is used to exchange routes between Contrail domains, and between Contrail domains and physical networks.
Dynamic tunnels using MPLS-over-GRE or VXLAN	Allows network overlay encapsulation to be applied so that user network packets can be carried from point to point in a transport network without the transport network "seeing" the user packet details (source/destination). This technique is used to carry packets between Contrail vRouters within a Contrail domain in a data center, and between vRouters in a Contrail domain and an MX Series router acting as a gateway.
Routing instances (VRF)	A physical router can contain multiple routing instances that allow local routing tables to be applied to network packets. Tunnels between routing instances using an overlay encapsulation allow VPNs to be implemented. The term virtual routing and forwarding (VRF) is used for routing instances that support Layer 3 VPNs using overlay encapsulation.
Routing and policy-based forwarding	This provides the ability to send traffic between routing instances within a router and to routing instances on other physical or virtual routers using signaled routing (via BGP or other control protocol) or forwarding based on rules applied to traffic entering a routing instance.
API for programming	This enables the ability to programmatically change the configuration of a router through a published interface and API. For instance, MX Series routers support Netconf with XML command payloads.

Most or all of these features are used in each of the use cases described in the following sections.

## Use Cases

The following section describes a variety of use cases, including Internet Gateway, Interdomain Gateway, Data Center Interconnect, Internetwork Gateway, Hardware Appliance Insertion, Virtual Private Cloud for Enterprise VPNs, Enterprise VPN Service Chains (Virtual CPE), and Service Chaining Gateway for Subscriber Networks.

### Internet Gateway

In the Internet Gateway use case, connecting to the Internet from a Contrail virtual network is enabled through standards-based route exchange with the MX Series gateway using BGP, and by the MX Series platform natively supporting the same encapsulation that Contrail vRouters use for communicating amongst themselves. From the perspective of Contrail, the MX Series router becomes part of a virtual network; from the point of view of the router, the Contrail vRouters appear as other routers to which packets destined for the VMs attached to them should be sent. The MX Series routers thus integrate seamlessly with the Contrail domain.

Alternative approaches require use of a software gateway that is itself running as a virtual machine. This can require traffic to take additional inefficient paths within the data center, and can also result in the software gateway becoming a bottleneck due to its own packet forwarding capability or CPU contention on the hypervisor on which it is running.

<sup>1</sup> The Contrail software is an open-source distribution available at OpenContrail.Org. The document can be found at: <http://opencontrail.org/opencontrail-architecture-documentation/>

## Benefit

The use of MX Series routers as gateways to Contrail virtual networks brings simplicity and resiliency to cloud infrastructures, which in turn allows service providers to offer strong service-level agreements (SLAs) to their customers while incurring low risk and internal costs. Proven architectures and protocols deliver high performance and resilience.

## Detail

This use case applies to enterprises and service providers (hosting and Infrastructure as a Service) who use Contrail and OpenStack to manage cloud infrastructure and where the virtualized resources need access to the Internet (e.g., Virtual Desktop Infrastructure), or resources need to be accessed from the Internet (web servers). The use case relies on the ability of Contrail to associate publically addressable “floating IP addresses” to interfaces of VMs in the virtual environment.

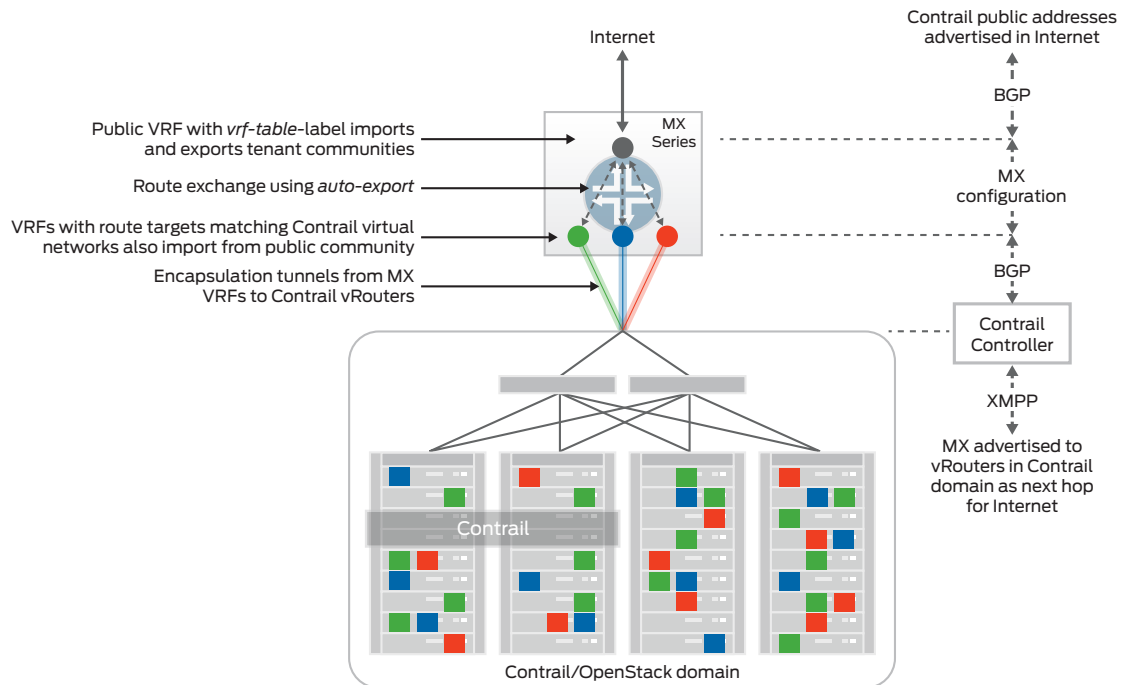


Figure 3: An MX Series router acting as a gateway to the Internet for a set of Contrail virtual networks

The router is configured to peer via BGP to an Internet gateway router and also with the Contrail controller. The Internet gateway router is specified as a next hop for the default route (Internet). A routing instance is created that will be used to aggregate traffic from a set of VRFs created for each Contrail virtual network. The routing instance is configured to advertise the pool of network addresses that can be used by Contrail for floating IP addresses into the Internet. Each VRF has a route target that matches that of a Contrail virtual network, which enables communication between each vRouter and its network's VRF on the MX Series via an encapsulation tunnel. The MX Series router is configured with logical tunnels between each VRF and the Internet routing instance. This provides a route to the Internet for each of the VRFs, and these routes to the Internet, in turn, are advertised into the Contrail virtual networks. Thus, VMs in the data center configured with floating IP addresses have a route to addresses on the Internet.

In most cases, Internet access will be provided by a pair of routers configured either to load-balance traffic, or to act as an active/standby pair. Contrail virtual networking supports either of these models, but for clarity the second router is not shown in Figure 3.

## Interdomain Gateway

When multiple Contrail domains are present in a data center, MX Series routers provide a convenient means to connect virtual networks in different domains, and to allow a virtual network to span multiple Contrail domains. Seamless integration between Contrail domains gives users in the data center the freedom to place workloads where needed, while preserving connectivity between different workloads and their users. The same MX Series router can also be used for connecting to the Internet, reducing costs since multiple separate devices do not need to be purchased.

## Benefit

Interdomain gateway allows seamless expansion of cloud infrastructure without fragmenting into multiple zones. Operational costs are lowered and workload downtime caused by unproductive migrations is eliminated.

## Detail

For administrative and operational reasons, data centers will contain more than one Contrail cluster, each of which forms a separate BGP autonomous system. MX Series routers can be used to allow traffic to flow between virtual networks in different Contrail domains, and can enable a virtual network to span multiple Contrail domains.

The MX Series router is configured to exchange routes with each of the Contrail controllers (using BGP), and VRFs are created that correspond to each of the virtual networks, as shown in Figure 4.

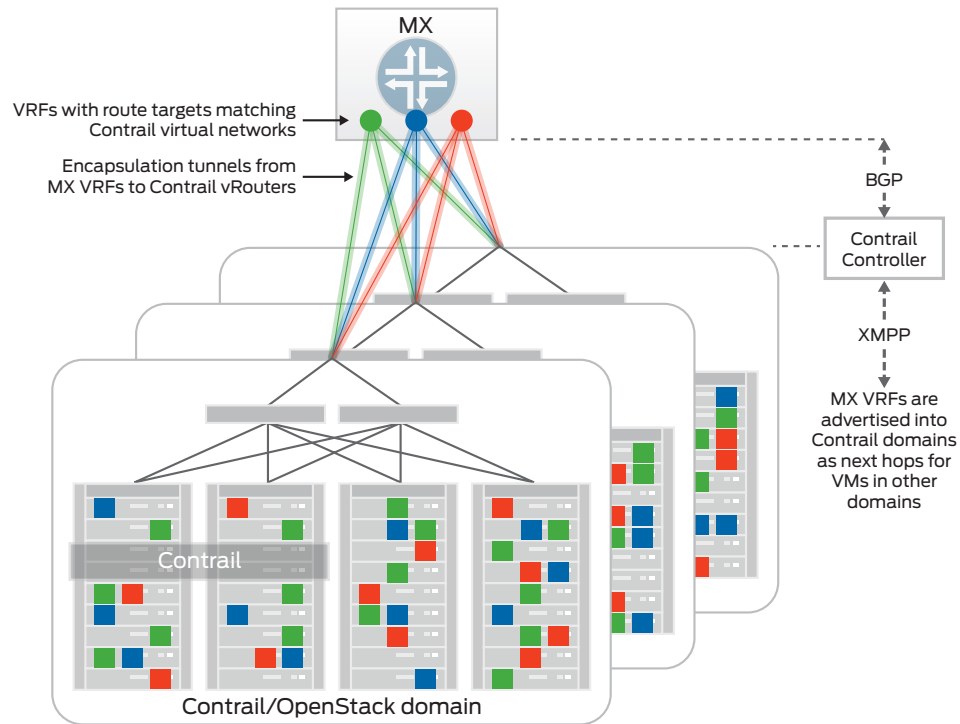


Figure 4: MX Series platform using VRFs to connect between virtual networks in different Contrail domains

The route targets in each VRF are set to correspond to those in the Contrail virtual networks in the Contrail domains. The VRFs can be configured to extend the same virtual network between domains, or to allow VMs in different virtual networks (with different route targets) to communicate. When VMs are created in a virtual network in one Contrail domain, the route to that new VM will be sent as an update to the MX Series platform, and a new route installed in the corresponding VRF. The route will be advertised to the other Contrail controllers and routes to the new VM will be installed in each vRouter with VMs in virtual networks that have policies that allow communication with the new VM. The routes to the new VM will have the MX Series router as next hop with a label that identifies the correct VRF.

In the diagram, a single MX Series router is shown, but normally a pair of routers will be configured to operate in active/active or active/standby configuration in order to provide high availability.

This use case shows how use of standards-based signaling and common encapsulation protocols allows MX Series routers to easily act as a gateway between Contrail domains. This use case can also be combined with the previous use case where a single MX Series router can be an interdomain gateway and provide Internet access.



## Data Center Interconnect

Interconnecting data centers is easily achieved by extending Contrail virtual networks onto VRFs configured on MX Series gateway routers, and by using standard BGP route exchange between gateway routers. Using BGP to exchange IP routes avoids complications that occur when networks are connected at Layer 2, specifically broadcast floods across the interconnect for address resolution (ARP).

### Benefit

Interconnecting Contrail-managed data centers using MX Series routers is simple and efficient, since the connection is formed using the same Layer 3 VPN technology in the WAN that Contrail was initially based on.

### Detail

This use case illustrates how MX Series routers can be used to interconnect multiple data centers, each of which contain multiple Contrail clusters. Using MX Series routers as gateways dramatically reduces control plane traffic between data centers, leading to better stability and performance while simplifying the network architecture and reducing the probability of configuration errors.

Figure 5 shows two data centers with Contrail clusters connected to a pair of MX Series routers in each data center. There are several ways to configure pairs of MX Series routers to provide resilience, including Virtual Chassis technology and multichassis link aggregation, but what is important here is the configuration of VRFs in each one.

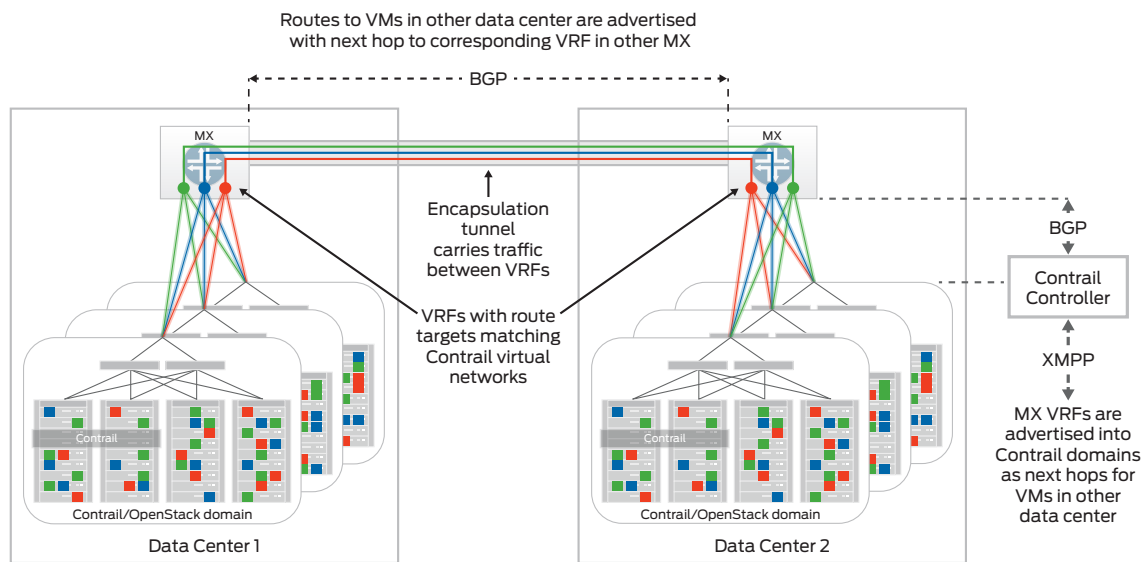


Figure 5: MX Series platform using VRFs to connect virtual networks in different data centers

Gateway routers may be connected via IP, in which case GRE tunnels will be automatically created using the “dynamic tunnels” option; or the connection may be via MPLS, in which case a tunnel can be created dynamically using LDP, or statically using RSVP (which allows control of the bandwidth of the interconnection). The tunnels encapsulate packets traveling between VRFs on different gateways in the same way that packets are encapsulated between Contrail vRouters.

Each MX Series gateway in a data center is configured with a VRF for each virtual network that needs to extend to another data center. Standard BGP route exchange between the MX Series routers allows each gateway to be informed of routes to VMs in the other data center and form the interconnection. A separate VRF is used for each virtual network to ensure that virtual networks with overlapping IP addresses don’t cause conflicts and to ensure isolation of each virtual network’s traffic.

A refinement to the method is to have separate address blocks for virtual networks in each data center, which then allows the blocks to be advertised by each gateway router, rather than all the individual host routes. In large environments, this can be used to reduce control plane traffic between data centers when VMs are created and destroyed. Using address blocks in the gateways corresponds exactly to how VPNs are implemented in carrier networks. The difference here is that the VPNs are extended inside the data centers as Contrail virtual networks, whereas carrier VPNs terminate at the gateway (PE router).

## Internetwork Gateway

Assets that are connected to Layer 2 switch networks using VLANs, including VMware virtualized servers and physical servers, can be seamlessly integrated with Contrail virtual networks by configuring VRFs on MX Series routers to act as gateways between the switched data center network and virtual networks. This allows existing assets to be leveraged in cloud infrastructures, ensuring a smooth evolution from conventional data center networks to overlay virtual networks while preserving the value of legacy assets.

### Benefit

Existing assets are not stranded since applications created in new cloud infrastructures can communicate with systems in legacy environments. This is a critical capability, since most environments are not greenfield.

### Detail

In most data centers, some applications run on physical servers and virtual servers that are connected using VLANs configured on physical switches in the data center. These applications need to be reachable and integrated into the virtual network environment. Examples include legacy database applications and certain video streaming applications that can't be virtualized. Additionally, existing virtualization implementations based on VMware may also need to be accessible from Contrail virtual networks. Such systems are connected to the network via Ethernet and are usually placed in VLANs configured in data center switches.

The MX Series router can act as a gateway between Contrail virtual networks and the VLANs that physical devices and servers are connected to. This is achieved by creating a VRF with the same route target as a Contrail virtual network and by configuring the VLAN interfaces that physical devices are connected to as part of the VRF. Finally, the VRF is configured with one or more static routes that identify which server addresses can be reached on each VLAN. These static routes are advertised on the virtual network so that VMs on the virtual network can now reach the physical devices.

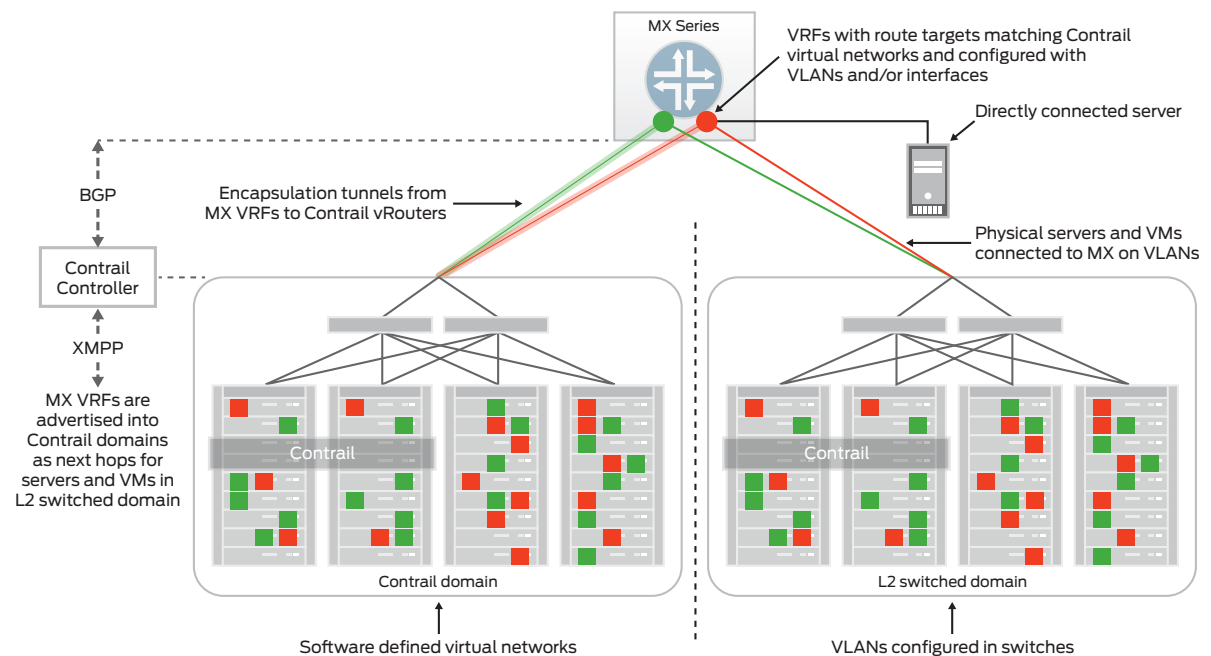


Figure 6: MX Series routers acting as a gateway to bare metal servers connected to VLANs or VMs connected to VLANs

## Hardware Appliance Insertion

Use of static routes in VRFs allows physical devices to be inserted into the data path between virtual networks. This enables existing data center assets, such as firewalls and load balancers, to be leveraged in cloud infrastructure. The solution using MX Series routers and Contrail is simple to deploy and manage because the routes into the physical devices are automatically advertised using BGP. No configuration is required in the Contrail domain, only the VRFs on the MX Series router.

### Benefit

The hardware appliance insertion use case allows continued use of appliances and network devices that provide value in the network, and around which operational procedures and corporate policies have been developed.

### Detail

This use case illustrates how MX Series routers can be used to insert physical devices into the path of traffic flowing between two virtual networks. An example would be a stateful firewall used to control types of traffic flowing between networks, or intrusion detection and prevention (IDP) for traffic coming into the data center from external sources. Most customers have such systems in place in their data centers, and have developed enterprise security policies around the specific capabilities of the devices they own. It is therefore highly desirable that these systems can be integrated into the virtual networks in a cloud infrastructure in order to preserve existing capabilities.

MX Series routers can be configured to direct traffic flowing between two virtual networks to flow through a physical device. This is achieved by creating two VRFs on the router, each with a route target of a virtual network, and with an interface that connects to the device configured in the VRF, as shown in Figure 7. Routes are configured in each VRF such that the next hop to get to the other network is the interface connected to the physical device. These routes are advertised to the rest of each network so that any traffic from a source in one network to a destination in the other will be directed first to the VRF on the MX Series router that is in the source network, then to the physical device, and will then emerge in the VRF in the destination network.

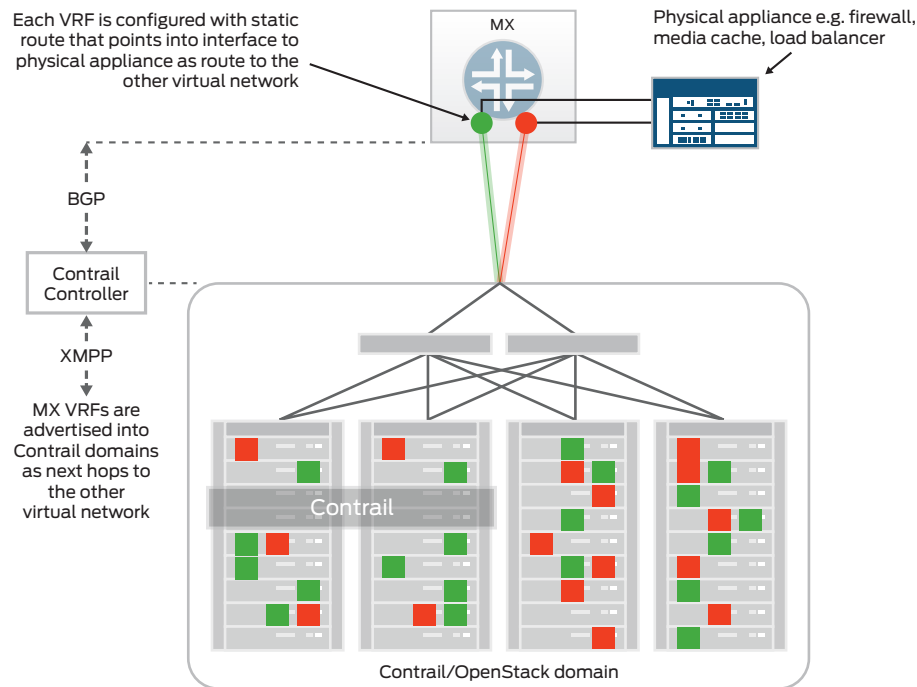


Figure 7: MX Series platform enabling insertion of a physical appliance between two virtual networks

The technique of advertising static routes to an interface in a VRF can be applied when physical devices have other than two interfaces to be connected to the virtual environment. This can be useful, for example, when a physical firewall is connected to multiple virtual networks.

Once again, use of standards-based control and data plane enables MX Series routers to be easily configured in order to bridge the physical and virtual worlds.

## Virtual Private Cloud for Enterprise VPNs

The ability to extend VPNs up to and inside data centers allows network service providers to offer cloud services to their VPN customers while preserving the performance and security SLAs of the VPN from each customer site right up to VMs that may be running critical customer applications. This will provide new business opportunities for network service providers and give customers access to the operational and financial benefits brought by hosting applications on cloud infrastructure.

### Benefit

The Virtual Private Cloud for Enterprise VPNs use case creates new revenue opportunities for network service providers who can capitalize on their existing customer base, installed infrastructure, and a reputation for reliability and security.

### Detail

This use case illustrates how a service provider can extend an enterprise MPLS VPN to include virtual networks created within cloud data centers using Contrail technology.

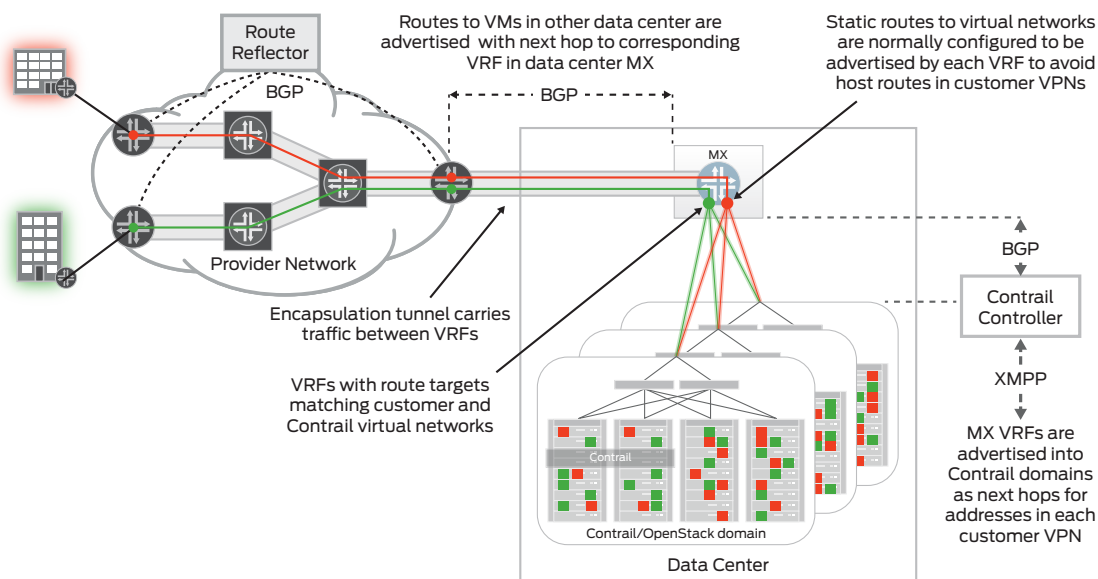


Figure 8: MX Series platform acting as a gateway between enterprise VPNs and Contrail virtual networks in a data center

In this scenario, a provider of MPLS VPN services can offer additional cloud services by configuring customer VRFs on a gateway router that is in a data center where Contrail is managing virtual networks. Similar to the data center interconnect use case, VRFs on the gateway router have route targets which correspond to customer VPN route targets and to those of virtual networks inside the data center. Although the VM addresses could be advertised individually in each VPN, the VRFs are more typically configured to advertise the entire address space of a virtual network as a single route, which is consistent with how PE routers in the VPN network advertise a route to the sites they are connected to. VMs created on a customer virtual network can communicate with destinations in the corresponding VPN, since routes to each site will become present on the VRFs on the gateway router as a result of peering between the data center gateway and the provider network gateway. The VRFs on the data center gateway are implementing the same network isolation technology that is relied on in all major carriers so, in combination with Contrail, a consistent security stance can be maintained end-to-end, all the way up to individual virtual machines.

From the perspective of the provider network, the Contrail domains appear as another provider supporting common VPNs. There are standard ways to make connections between provider networks. The connection method shown in Figure 8 corresponds to Inter AS Option A, where matching VRFs are created on each side of the connection. Alternatively, the MX Series router inside the data center could peer directly with the provider route reflector, and no VRFs would be required in the connecting router. This corresponds to Inter AS Option AB. The connection option depends on the policies of the network provider and whether the data center is provider- or partner-owned.

Typically, customers will be given access to a portal that allows them to create VMs inside the data center on demand. The VMs placed on a virtual network connected to the VRFs on the gateway will be accessible from other sites in the same VPN, but VMs can also be placed in virtual networks in the data center with policies that allow access only to other virtual networks. This allows customers to create applications where, for instance, a web server is accessible from anywhere on the customer network, but the underlying application and database servers are only accessible by administrators through the cloud management portal.

A key advantage that network operators have over public cloud service providers in offering this kind of service is that the network operator is able to offer end-to-end service-level guarantees for network performance and security, since the traffic remains within the operator network and does not have to traverse the Internet at any point. For most large enterprises and for local and national governments, lack of security guarantees, in particular, prevents them from taking advantage of the lower costs and increased flexibility offered by cloud service providers. Extension of VPNs up to and inside data centers solves this problem, and it is a great example of how the use of industry standards in Contrail and on MX Series routers can enable new business solutions.

### Enterprise VPN Service Chains (Virtual CPE)

Virtualizing network services such as firewall, intrusion detection, distributed denial of service (DDoS) detection, SSL gateway, and media cache—which until recently have been deployed as specialized appliances—gives service providers the flexibility to offer new services without having to deploy new equipment at the customer's premise (CPE). Virtualized services are deployed in pools of servers that can be located locally to PE routers, regionally in metro areas, or in centralized data centers, depending on the performance requirements of the service (e.g., latency) and business constraints of the service provider (availability of space, power, and cooling).

By leveraging cloud automation, Contrail virtual networking, and MX Series routers, customers can order services from providers that can be provisioned automatically and inserted into the customer VPNs for immediate use.

#### Benefit

The combination of Contrail and MX Series routers provides a platform that allows network service providers to offer new services to their existing customer base. The compute infrastructure required to support the services can grow with service uptake, and the mix of services deployed can match customers' individual needs.

The use of standards-based BGP signaling and MPLS label encapsulation enables network service providers to reinvent how they deliver services and opens up the possibility of new types of service offerings that align with the needs of their customers, while at the same time being able to turn up these services in minutes rather than weeks.

#### Detail

In this use case, virtualized appliances are inserted into the data path between two physical networks, allowing network service providers to offer additional services to existing Layer 3 VPN customers. Rather than extending customer VPNs to a centralized data center (although this is possible), data centers are more likely to be located close to the PE routers (in the same metro area or colocated with PE routers). The intent is to attach virtualized appliances to customer VPNs to support services such as Internet access, security services (firewall, IDP, UTP, SSL), and media caching (content delivery network or CDN).

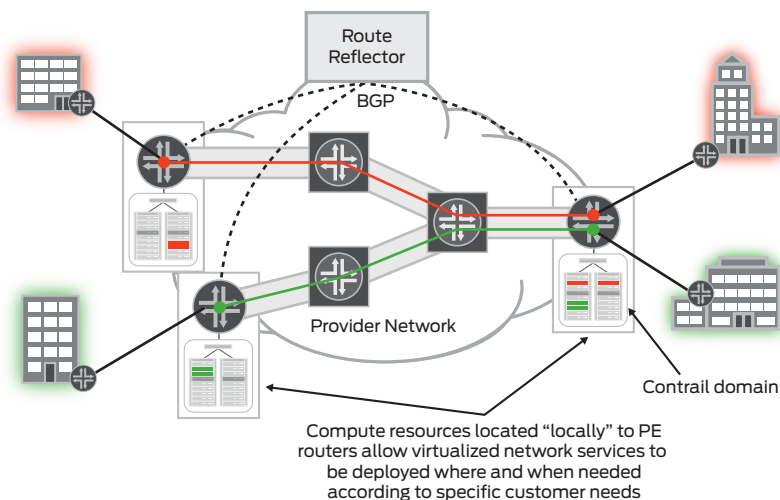


Figure 9: Virtual appliances can be inserted into enterprise VPNs to support on-demand services

To form customer VPNs, each PE router is configured with a VRF for each customer with a connection to that PE, and each VRF contains the interfaces that connect to that customer's sites. BGP is used to exchange routes between customer VRFs on all PE routers in the provider network, so a host on any customer site can reach destinations on any other site belonging to that customer.

When a customer orders a new service, operational systems determine how many virtual machines will be required to support the required performance, and where they should be located, based on which PE routers the customer has connections to. The VMs are inserted into the customer's VPN by advertising new routes into the VPN that divert traffic into the newly created services.

For example, in Figure 10, a customer has requested a combination of Internet access, firewall, and intrusion detection that are implemented in two VMs configured as a service chain between the VPN and the Internet. Contrail is used to deploy service instances in each provider point of presence (POP) where the customer has a connection to a PE router.

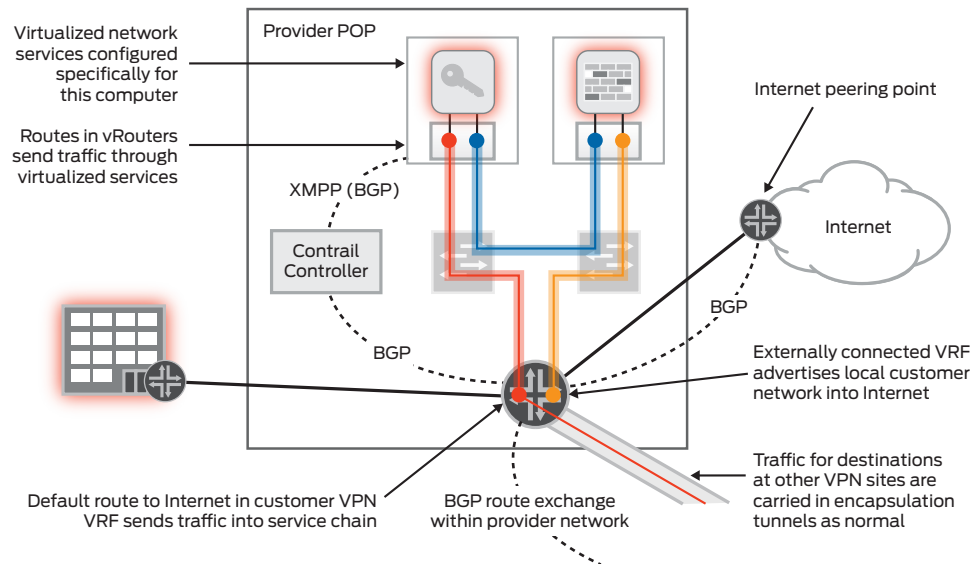


Figure 10: Virtual appliances connected as a service chain into a VPN using virtual networks

The service instances are created from previously defined service templates, which can contain fixed configurations to support specific services, or service instances can be dynamically configured with customer-specific detail after initialization. Contrail configures the VRFs in the vRouters along the service chain with routes that send Internet traffic from the customer side through each service in turn, and traffic destined for addresses in the customer network in the reverse direction. The route target of the VRF at the start of the service chain matches that of the customer VPN, so BGP route exchange causes a route to the Internet via the service chain to be installed in the customer VRF in the PE router. Similarly, the route target of the VRF at the end of the service chain matches that of a VRF on the PE router that is configured with a route to the Internet via some Internet peering point. BGP route exchange causes the public IP addresses for the customer site to be advertised on the Internet. Each customer site gains access to the Internet via a local peering point, and incoming traffic is protected by firewall and intrusion detection services specific to that customer located in the local provider's POP for each site.

### Service Chaining Gateway for Subscriber Networks

Subscriber- and application-aware steering into Contrail service chains allows mobile, cable, and wireline operators to introduce virtualized services into their networks. This is achieved by configuring MX Series routers as service control gateways with standards-based interfaces to policy and charging rules function (PCRF) or authentication, authorization, and accounting (AAA) for policy enforcement, and optionally configured with deep packet inspection (DPI) that can determine websites visited, media being downloaded, and other traffic types. This granular detail allows flows to be steered to a specific service chain based on offerings purchased by individual subscribers and on the details of each traffic flow. Examples include parental control, content optimization, and content cache (CDN).

### Benefit

The combination of MX Series routers configured as a service control gateway, and service chains configured by Contrail from virtualized service instances connected using virtual networking, allows operators to dramatically reduce time to market for new service offerings. It also enables many different offerings to coexist on shared cloud infrastructure. Services can be scaled down for trials and be targeted at small, geographically distributed markets that could not have supported a business case previously. Services can be scaled up for high volume services by simply deploying more virtual service instances. Cloud automation with Contrail virtual networking allows services to be deployed according to business need at the most appropriate location.

## Detail

This use case applies the service chaining method described in the previous use case to both wireline and mobile subscriber networks. In the previous use case, service chains were created to meet the needs of each customer, and traffic was directed into a service chain using routing in a specific VPN. In subscriber networks, individual users have a specific set of services that they have purchased, and steering into a service chain can depend on knowing the subscriber identity, the specific destination being accessed, and the content being accessed. Routing, which is based on destination IP address only, does not have sufficient granularity in this case. Integration with subscriber policy management systems such as PCRF or AAA, in combination with DPI, can allow such fine-grained steering to be performed.

In contrast to the previous use case, where each service instance has to deal with the traffic for a small number of customer sites, in the subscriber gateway use case, the services may see a large fraction of the traffic flow in a significant fraction of a network. Scale and performance become important, and service chains must be designed to handle the expected traffic flows. Important issues such as dynamic scaling of services come into play, ensuring that forward and reverse flows pass through the same instances of stateful services, and supporting reverse steering into the correct service chain must be dealt with. These issues are all solved in the subscriber service chaining architecture using Contrail and MX Series routers, but details are beyond the scope of this document.

An MX Series router configured for subscriber- and application-based steering is called a service control gateway. Such a device is installed with a number of service cards that manage subscriber sessions and also maintain a connection to a policy control system such as PCRF. When a new subscriber flow arrives at the service control gateway, the packet header and the results of DPI inspection of the first part of the payload are used as input to a filter function that matches a set of predefined criteria that defines the destination, traffic type, and content of the flow. The matching criteria, together with source IP address (used to identify the subscriber), are sent to the PCRF whose response identifies which service chain the traffic should be sent to.

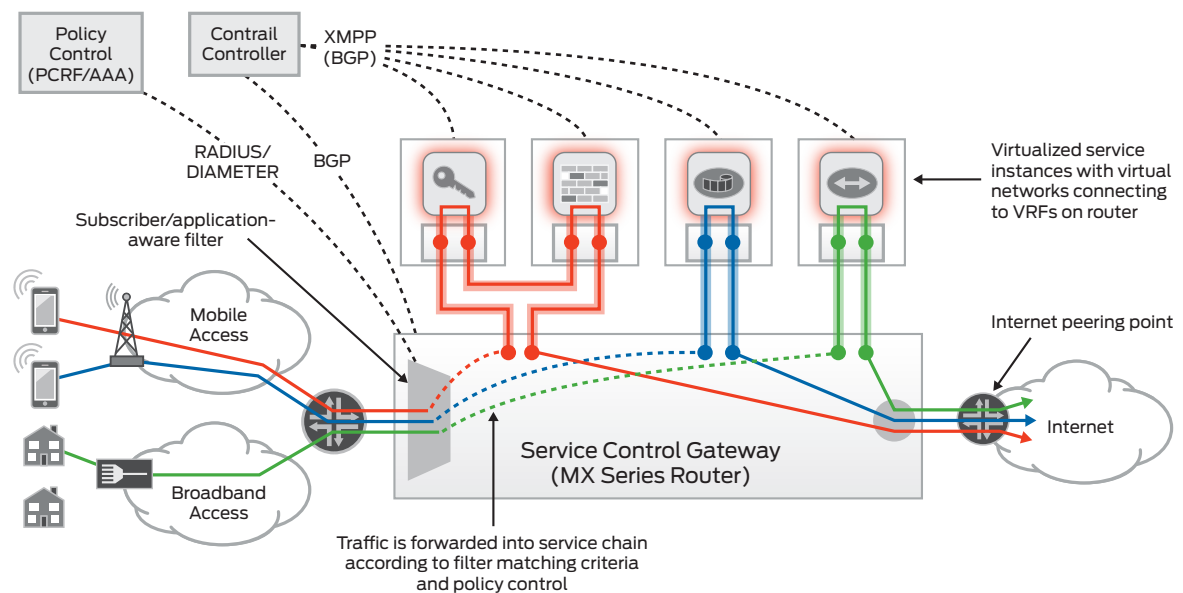


Figure 11: Service chaining through service delivery gateway with service- and application-aware steering

Service chains are deployed in similar fashion to the previous use cases and can be composed of physical and virtual service instances. Instead of routes being exported from the entry and exit VRFs, the filter function sends traffic directly into the VRFs using filter-based forwarding. This allows finer grained steering than is possible using routing alone (based just on the destination IP address).

## Conclusion

Juniper Networks MX Series 3D Universal Edge Routers can be configured as gateways for Contrail virtual networks, meeting a wide variety of use cases. The solutions build on the fact that integration of the MX Series gateway with Contrail virtual networks is easily achieved, since both support industry standards such as BGP and MPLS. The use cases cover a wide variety of situations, including intra-data center, inter-data center, physical and virtual services, business VPNs, and subscriber networks. The use of Layer 3 virtual networks allows the solutions to be simply and easily configured, and the resulting network architectures deliver high scale, resilience, and flexibility of deployment.

## 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.

