Building Continuously Available and Simplified Networks with MX Series Virtual Chassis

Virtual Chassis Technology on MX Series 3D Universal Edge Routers Highlighting Key Benefits in Real-World Deployment Scenarios

Table of Contents

Executive Summary 3
Introduction 3
Introducing MX Series Virtual Chassis 3
MX Series Virtual Chassis Overview 4
Deployment Scenarios with MX Series Virtual Chassis 5
Extending Resiliency for Subscriber Aggregation 5
Simplified and Scalable Data Center Deployment 6
Conclusion 7
About Juniper Networks 8

List of Figures

Figure 1: Routing Engine roles in a 2-member Virtual Chassis 4
Figure 2: Converged infrastructure for large-scale subscriber services with MX Series Virtual Chassis technology 6
Figure 3: Data center architecture with traditional L2 design as compared with MX Series Virtual Chassis 7

Executive Summary

Sustained broadband subscriber growth, exponential bandwidth growth, unpredictable traffic patterns, and a vast increase in global mobility have challenged traditional service provider broadband network architectures. Business and technical decision makers are looking for cost-effective ways to scale network and service resources, increase service resiliency, improve operations, and enhance end user experience.

Juniper Networks® MX Series 3D Universal Edge Routers are purpose-built to enable a converged edge infrastructure that cost-effectively integrates multiple network services such as broadband edge, business edge, data center edge, and mobile services core onto a single platform and operating system. As a key element in a converged service edge architecture, the MX Series router is designed to flexibly increase service capacity "on demand," while addressing the zero tolerance policy for network failures or outages that could affect service-level agreements (SLAs). Virtual Chassis technology on MX Series routers is an innovative and powerful addition to the various technologies available on the MX Series 3D Universal Edge Routers that address today's network growth and service availability requirements.

This paper provides an overview of MX Series Virtual Chassis technology and illustrates how this feature can simplify operations, provide unparalleled resiliency, and increase scale and service avaibility.

Introduction

Service providers are struggling to keep up with consumer demand for higher speeds, rich service suites, and continuous connectivity. The concurrent need to increase scale, maximize uptime, minimize cost, and deliver revenue generating services has put more emphasis on the requirement for continuously available and stateful networks. To address this challenge, service providers are investigating creative network designs and alternative technologies to increase network availability, while supporting IP service solutions that scale to match broadband growth.

Traditional redundancy schemes focused on intra-chassis mechanisms are no longer sufficient to meet the levels of availability needed for increasingly stringent SLAs that require both geographic redundancy and chassis-level redundancy to protect against catastrophic outages. These inter-chassis redundancy (ICR) schemes must be able to protect against access link failures, uplink failures, and whole chassis failures—without visible disruption to attached subscribers.

The following are key requirements that must be met with any resiliency scheme offering inter-chassis redundancy:

- The resiliency scheme should address multiple types of failures including link, port, line card, Routing Engine, and complete chassis failure.
- The failover or backup chassis can be colocated or located in a different POP. Resiliency solutions must allow the backup chassis to be remote to ensure that catastrophic or environmental events in one location will not impact the continuity of the network or service.
- The resiliency scheme should interoperate with standards-based adjacent equipment such as access nodes, aggregation nodes, other edge equipment, and backbone nodes.
- The resiliency scheme should be service and subscriber aware. In the event of a Routing Engine switchover, the resiliency scheme should be able to preserve subscriber state to ensure nonstop operation without any need for subscriber reauthentication.
- The resiliency scheme should not introduce any operational complexity on the control plane or the management plane. Ideally it should reduce operational complexity by reducing the number of managed devices.

Key business requirements driving advanced network resiliency:

- Need to rapidly deploy new services and applications with improved scalability and performance
- Improve business continuity by providing continuously available networks for voice, video, and data services
- Reduce operational costs by improving the efficiency of the network and back-office systems
- Reduce latency and enhance the end user experience
- Improve sustainability by efficiently utilizing network resources and reducing facility power and cooling demands

Introducing MX Series Virtual Chassis

The MX Series with Virtual Chassis technology protects end users from outages by providing complete system-level redundancy with sub-second failover to standby resources, ensuring uninterrupted network connectivity. MX Series Virtual Chassis acts as a network firewall by gracefully handling failures in adjacent equipment, thereby reducing or eliminating the need to propagate route table changes or failover to standby links which can impact services.

MX Series Virtual Chassis allows service providers to handle catastrophic events without service interruptions. Virtual Chassis extends the router's control and forwarding planes to remote facilities, and it preserves critical service-related state information on the remote Virtual Chassis members in real time. This allows the remote MX Series Virtual Chassis member to provide continuous, uninterrupted end user services when the site housing the other MX Series Virtual Chassis member is no longer capable of supporting services.

In addition to increasing system-level redundancy, MX Series Virtual Chassis allows operators to logically extend the capacity of a single chassis without extending the router's control plane. From a network view, multiple systems become a single virtual system where element-level capacity can be logically extended beyond the confines of the chassis hardware. This allows service providers to support both physical connectivity growth and service capacity increases without introducing additional routing elements or control overhead into the network.

MX Series Virtual Chassis provides a solid foundation on which service providers can confidently build their products and services, while at the same time providing an unparalleled service experience for their end users. Juniper understands that a service provider's business success and reputation rely on providing uninterrupted connectivity to "the new network" for their customers. MX Series Virtual Chassis meets this challenge, providing a resilient, scalable, and cost-effective solution for the next generation of networks.

MX Series Virtual Chassis provides the following advantages to service providers:

- Reduced administrative overhead as two or more discrete routers become one Virtual Chassis
- More efficient link utilization by eliminating Layer 2 loops, which allows more active links without additional cost
- Simplified manageability by providing a unified control plane and single configuration for the Virtual Chassis
- Improved resource utilization by intelligently utilizing line-card interfaces and service modules on separate chassis, letting customers benefit from a "pay-as-you-grow" model
- Protection of user sessions across chassis, line card, or port failure by using stateful redundancy (failures are transparent to end users)
- Minimized network reconvergence of adjacent devices during failure events (these features are extended across the members of the virtual chassis)

MX Series Virtual Chassis Overview

MX Series Virtual Chassis allows multiple MX Series routers to be grouped together into a single virtual router without any additional specialized hardware. Virtual Chassis technology extends all of the features of a single virtual router to each of the member routers to operate as a single unified system. The dedicated interfaces interconnecting the member routers are called Virtual Chassis Control Protocol (VCCP) ports. The VCCP control protocol, which is based on IS-IS with extensions, is the control protocol for the Virtual Chassis system that creates a loop-free topology, provides for dynamic master election, and establishes the routing table for inter-chassis routing within the Virtual Chassis configuration.

Following are the roles for the different Routing Engines in a Virtual Chassis:

- Virtual Chassis Master (VC-Mm)—Master Routing Engine for the Virtual Chassis that is elected using the VCCP protocol
- Virtual Chassis Backup (VC-Bm)—The backup Routing Engine that maintains a stateful backup of the Master Routing Engine and becomes the Master Routing Engine upon failure of the current VC-Mm

Each physical chassis has a standby that backs up the local master (i.e., VC-Ms and VC-Bs). Figure 1 shows these Routing Engines and their roles in detail.

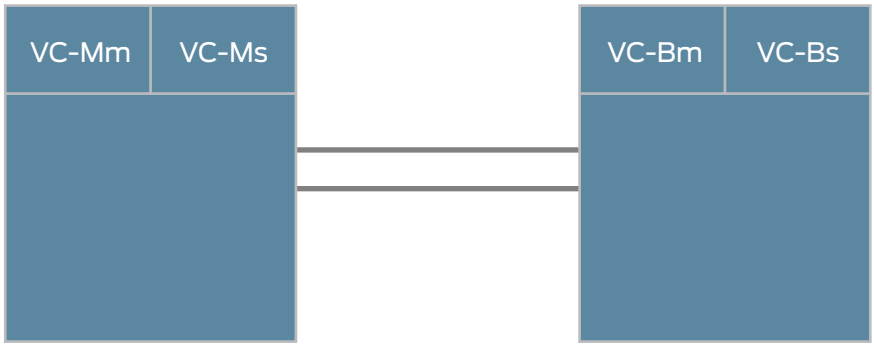

Figure 1: Routing Engine roles in a 2-member Virtual Chassis

The VC-Mm Routing Engine in the master router for the Virtual Chassis runs chassis management processes and control protocols. It manages the master and all member routers. This Routing Engine receives and processes all incoming control plane traffic destined for the Virtual Chassis, and is also responsible for propagation of Virtual Chassis configuration to all Virtual Chassis members.

If the VC-Mm fails, the VC-Bm becomes the master for the Virtual Chassis, and VC-Ms becomes VC-Bms. Support for stateful redundancy ensures that subscriber sessions stay active if a master Routing Engine failover occurs. By separating the control plane and data plane, forwarding continues uninterrupted and there are no changes in network topology with peer devices, thereby providing additional network stability.

Inter-chassis redundancy on the MX Series Virtual Chassis meets all of the requirements for continuously available routing and switching. These translate into tangible and unique benefits for network design and operations in a variety of deployments.

- Provides the ability to manage all control and management plane functions of the member routers as a single system. The MX Series Virtual Chassis appears as a single system for all control plane updates and management plane tasks, including configuration and interaction with backend systems such as RADIUS authentication and accounting for subscriber services.
- MX Series Virtual Chassis leverages MX Series and Junos OS features such as graceful Routing Engine switchover (GRES), unified ISSU, nonstop bridging (NSB), and nonstop active routing (NSR).
- Scaling and performance is enhanced with the ability to extend the forwarding plane beyond the physical ports available in a single chassis.
- There is no custom hardware needed for MX Series Virtual Chassis; the resiliency scheme is software-based using programmable Trio chipset Modular Port Concentrators (MPCs).
- Enterprise Internet edge networks can reduce the number of peering points with their Internet service providers (ISPs) by consolidating these links onto the MX Series Virtual Chassis to provide continuously available geographically dispersed connectivity.
- VoIP application layer call routing employed by session border controllers (SBCs) typically selects between multiple routers to route calls to the called party. With MX Series Virtual Chassis, SBC application layer call routing can be simplified by reducing the number of routers in the VoIP routing tables.
- In subscriber edge BNG deployments, MX Series Virtual Chassis extends the resiliency benefits across the network, simplifying the access and aggregation as well as back-office operations considerably.
- In data center scenarios, the increased scale and enhanced resiliency of MX Series Virtual Chassis simplifies network operations and architecture design, and maximizes network resource utilization.

Deployment Scenarios with MX Series Virtual Chassis

This section illustrates the scale, operational simplicity, and enhanced resiliency benefits that can be achieved with MX Series Virtual Chassis routers in two deployment scenarios—as a BNG for subscriber aggregation, and as a data center edge solution.

Extending Resiliency for Subscriber Aggregation

The customer in this example needs to roll out wholesale residential edge services at large scale. The key business requirements for the next-generation network buildout are:

- Centralized edge platform to deliver multiple services over a converged infrastructure
- Lower capital and operational expenses
- “Always available” service uptime
- Simplified operations and management of subscriber services

To achieve these goals, the customer must:

- Deliver large-scale subscriber services on an Ethernet optimized platform that can be leveraged for business edge and carrier Ethernet
- Improve service resiliency without complicating operations
- Use a subscriber access model that minimizes static configuration
- Optimize bandwidth usage and minimize L2/L3 broadcasts
- Ensure flexibility to add and deploy new services on the centralized edge platform
- Support dual stack IPv4/IPv6 and native IPv6 services

Juniper Networks has been a leading provider of broadband solutions for well over a decade, and it maintains its leadership with extensive BNG feature support on the MX Series routers. To support next-generation subscriber aggregation requirements, not only do the MX Series routers provide a scale and performance benefit, but they also offer the Virtual Chassis technology needed to improve SLA guarantees and provide enhanced resiliency while maintaining operational simplicity.

Figure 2 illustrates a scalable and resilient architecture to support a residential broadband network.

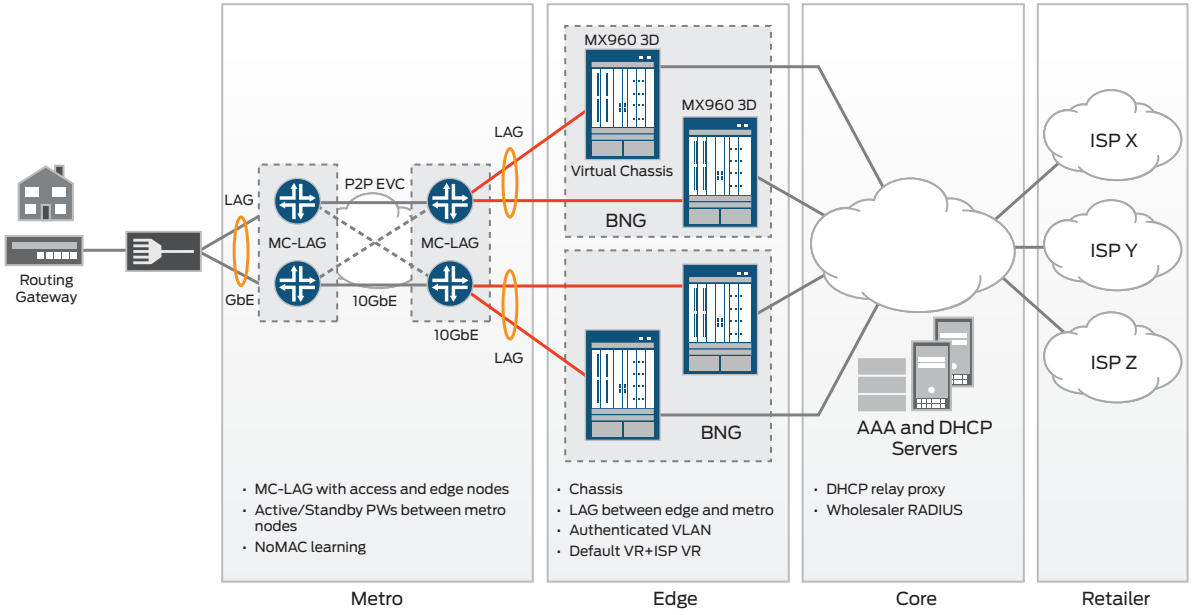

Figure 2: Converged infrastructure for large-scale subscriber services with MX Series Virtual Chassis technology

Key features in this architecture include:

- The MX Series is an Ethernet optimized platform and supports extensive BNG features and scale.
- The MX Series Virtual Chassis technology can simultaneously be deployed to serve business and residential subscribers.
- MX Series Virtual Chassis provides enhanced resiliency for complete chassis failures and maintains operational simplicity without the need for any additional hardware.
- MX Series Virtual Chassis inter-chassis redundancy can be extended to member routers that may be colocated or in geographically separate locations.
- In the event of link/port/Routing Engine/chassis failure on member routers within the Virtual Chassis, there is no impact to subscribers or any need for RADIUS reauthentication.
- With the use of multi-chassis link aggregation (MC-LAG) in the metro and using P2P pseudowires, services can be restored in less than a second. All of this is achieved without adding any complexity to the network design or day-to-day operational tasks.
- Service availability is enhanced and network design is simplified as the number of nodes is reduced.

Simplified and Scalable Data Center Deployment

Modern day global enterprise businesses continue to raise the scale demands and availability targets on today's data center deployments. Consolidation and virtualization are the key design avenues to deliver increased scale and service availability. MX Series Virtual Chassis technology meets both requirements and offers many benefits to data center network design.

By combining physical nodes into one virtual system that can be managed as a single entity, MX Series with Virtual Chassis technology helps reduce the number of management entities on the network. A reduced number of nodes to manage translates into a simpler network design. Day-to-day tasks such as service provisioning are made easier, since the touchpoints that need to be configured are reduced. The combined benefits of simpler network design and less complex operations ultimately lower operational overhead.

Compared to traditional L2/L3 designs, data center network design with MX Series Virtual Chassis can be simplified and optimized for maximum network utilization. At Layer 2, spanning tree can be completely eliminated (see Figure 3) by setting all links in active state. Without MX Series Virtual Chassis in play, 50% of the links would be blocked via Spanning Tree Protocol (STP) loop detection.

The use of Virtual Router Redundancy Protocol (VRRP) or Gateway Load Balancing Protocol (GLBP) for Layer 3 redundancy can also be eliminated, and traffic can be load-balanced on all active links. Hybrid L2/L3 connections from compute/storage resources towards the core are typical in data center deployments. The single control plane for the MX Series Virtual Chassis provides significant simplicity as the Virtual Chassis maintains a single Address Resolution Protocol (ARP) table, Dynamic Host Configuration Protocol (DHCP) server, snooping database, and multicast

membership table. Traditional node redundancy with the use of L2 and L3 protocol techniques cannot provide these benefits. With increased port density and bandwidth availability, Virtual Chassis design also dramatically improves network utilization and efficiency.

Figure 3 depicts aggregation and distribution layers, and it illustrates a comparison between the poor link utilization using STP and MX Series Virtual Chassis, where all links are active.

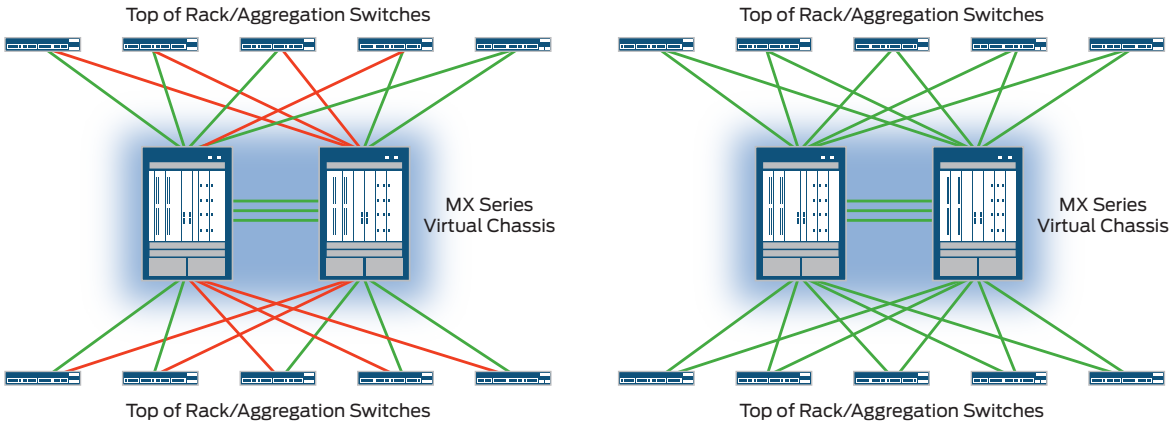

Figure 3: Data center architecture with traditional L2 design as compared with MX Series Virtual Chassis

The overall service availability in the data center is improved for the following reasons:

1. Failures within the Virtual Chassis are transparently handled.
    - Nonstop operation is supported with GRES/NSR/NSB on the MX Series Virtual Chassis just like a single chassis system.
    - In Juniper Networks Junos® operating system, all processes support graceful restart. If they crash, they are restarted and automatically come back up.
    - In case of any hardware failure on the Routing Engine, a switchover is triggered and the backup Routing Engine takes over. This is transparently handled with support for nonstop operation and reduces the churn caused by network convergence.
2. No dependency on slow convergence from STP; fast convergence can be achieved with the use of Layer 2 pseudowire technology and MC-LAG.
3. Reduced L2 broadcast via a single ARP table for the entire MX Series Virtual Chassis.
4. Support for disaster recovery planning: MX Series Virtual Chassis members can be deployed in geographically diverse locations.

Conclusion

Today's broadband service providers must cost effectively increase subscriber scale, maximize uptime, minimize cost, and deliver more revenue generating services. These business imperatives have placed increased emphasis on continuously available and stateful networks to maximize revenue and meet customer availability requirements.

Juniper Networks MX Series 3D Universal Edge Routers with Virtual Chassis technology provides a powerful tool that network architects can use to create continuously available network services while reducing OpEx costs and complexity. Interchassis redundancy is a high availability feature that help prevent network outages and protects routers against access link failures, uplink failures, and wholesale chassis failures.

These systems provide a solid foundation on which providers can build their products and services, while offering an unparalleled service experience for end users. With MX Series Virtual Chassis as the foundation for the next generation of high-performance networks, the new network is here.

About Juniper Networks

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

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

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

2000487-002-EN Feb 2016