Juniper® Validated Design

JVD Test Report Brief: Scale-Out Stateful Firewall and Source NAT for Enterprise

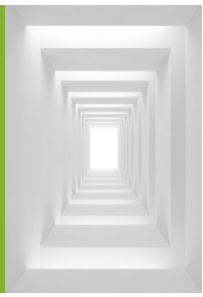

test-report-brief-mse-cgnat-offbox-01-01

Introduction

Today the existing firewall solutions market trends and moves towards an offering with a smaller footprint, significantly increases its power efficiency to align with their companies' green initiatives making it high scalable with high throughput. Whereas the traditional firewall solutions offer fixed slots for specific purposes, power inefficient and require pre allocated rack space until the device is used at maximum scale. Scaling further in the same approach is costly and less efficient deployment.

The security and network require a new form of deployment which is highly scalable, offers a higher grade of redundancy and flexibility, higher efficiency and is based on the existing Juniper devices and software. This is where a distributed firewall architecture increases scalability and improves performance without adding any complexity and management overhead. In terms of security, a distributed firewall architecture provides a more robust and resilient security solution which reduces downtime and ensures continuous protection and operational stability. It is maintained with a single pane of glass avoiding unnecessary configuration challenges resulting in significant cost savings over time.

The CSDS architecture is a solution which combines the available Juniper forwarding architecture devices with service plane capabilities of SRX/vSRX Series Firewalls or instances. The service card capabilities are available outside the forwarding chassis itself and is connected to the forwarding layer directly or indirectly through the distribution layer if needed with this solution. It offers a new form of redundancy of the forwarding path and remote service layer by segregating both the planes into multiple groups depending on the use case and chosen deployment method effectively eliminating a single point of failure. In conclusion, CSDS is the future of firewall architecture. With improved scalability, performance, power efficiency, flexibility, security, and cost-effectiveness, a distributed firewall system is a solution that businesses of all sizes should consider.

CSDS leverages existing features like EBGP, BFD, ECMP CHASH, SRD, and TLB on MX as a forwarding plane and various security features such as CGNAT, IPSEC, stateful firewall, and MNHA on SRX Series Firewall.

This JVD test plan is modified to have only use cases related to JTMS Test plan TPI. The details are:

- TPI-128508 - CGNAT/SFW with vSRX and MX304 platforms - MX doing ECMP CHASH Load balancing or RE TLB based load balancing

Other JTMS JVD test plans created for other features and platforms are:

- TPI-128244 - CGNAT/SFW with SRX4600 and MX304 platforms - MX doing ECMP CHASH load balancing or RE TLB based load balancing
- TPI-128507 - IPsec with SRX4600 and MX304 platforms - MX doing RE TLB based load balancing TPI-129168 - IPsec with vSRX and MX304 platforms - MX doing RE TLB based load balancing

Table 1: CSDS Solution Matrix

MX Load-Balancer Component	MX Load-Balancer Redundancy (Single MX or Dual MX (SRD))	Security Features	SRX/vSRX in MNHA mode	SRX/vSRX in Standalone mode
ECMP CHASH	Single MX	CGNAT/NGFW	No	Yes
		IPSEC	No	Yes
	Dual MX (SRD)	CGNAT/NGFW	Yes	No
		IPSEC	No	No
Traffic Load Balancer [TLB]	Single MX	CGNAT/NGFW	Yes	Yes
		IPSEC	Yes	Yes
	Dual MX	CGNAT/NGFW	Yes	Yes
		IPSEC	Yes	Yes

Test Topology

Topology 1 – ECMP CHASH – Single MX Series Router with Scaled Out Standalone SRXs (Multiple Individual SRX Series Firewalls)

This topology is simple and least redundant. There is no backup of the MX Series Router and there are no sessions or Internet Key Exchange (IKE) synchronization between the SRX Series Firewalls.

Figure 1: Topology 1 ECMP/CHASH -Single MX with Scaled out Standalone SRX

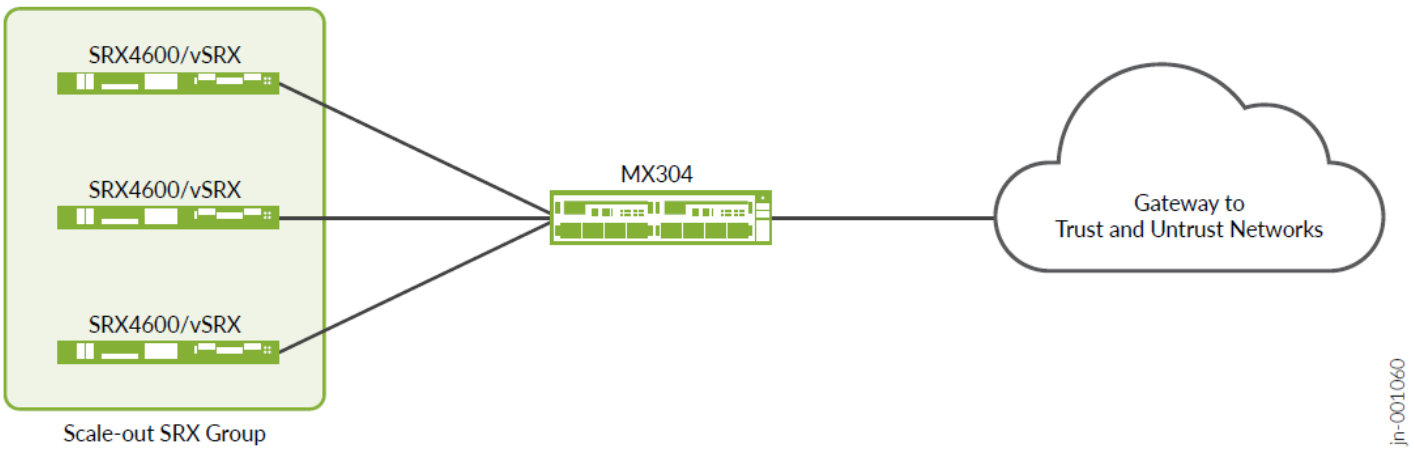

However, it helps to understand how this architecture works. Typically, you can opt for more redundancies. If you are not concerned about stateful failover and may want to augment security service capacities by adding more SRX Series Firewalls, then the application sessions may be short lived (a redundancy mechanism may be handled at an application level not requiring any session sync between two different firewalls).

- Pros: Simplicity and scaling with each individual SRX Series Firewall
- Cons: No redundancy

Topology 2 – ECMP CHASH – Dual MX Series Routers with Scaled Out MNHA SRX Pairs (Multiple Pairs of SRX Series Firewalls)

This topology does offer redundancy for the MX Series Routers and for each SRX Series Firewall. The dual MX Series Routers uses an SRD mechanism to monitor the physical elements of the network and/or the MX Series Routers itself, as well as any other routing and system event that may need to trigger a failover to the other MX Series Routers.

Figure 2: Topology 2 ECMP CHASH -Dual MX [SRD] with Scaled out MNHA SRX Pairs

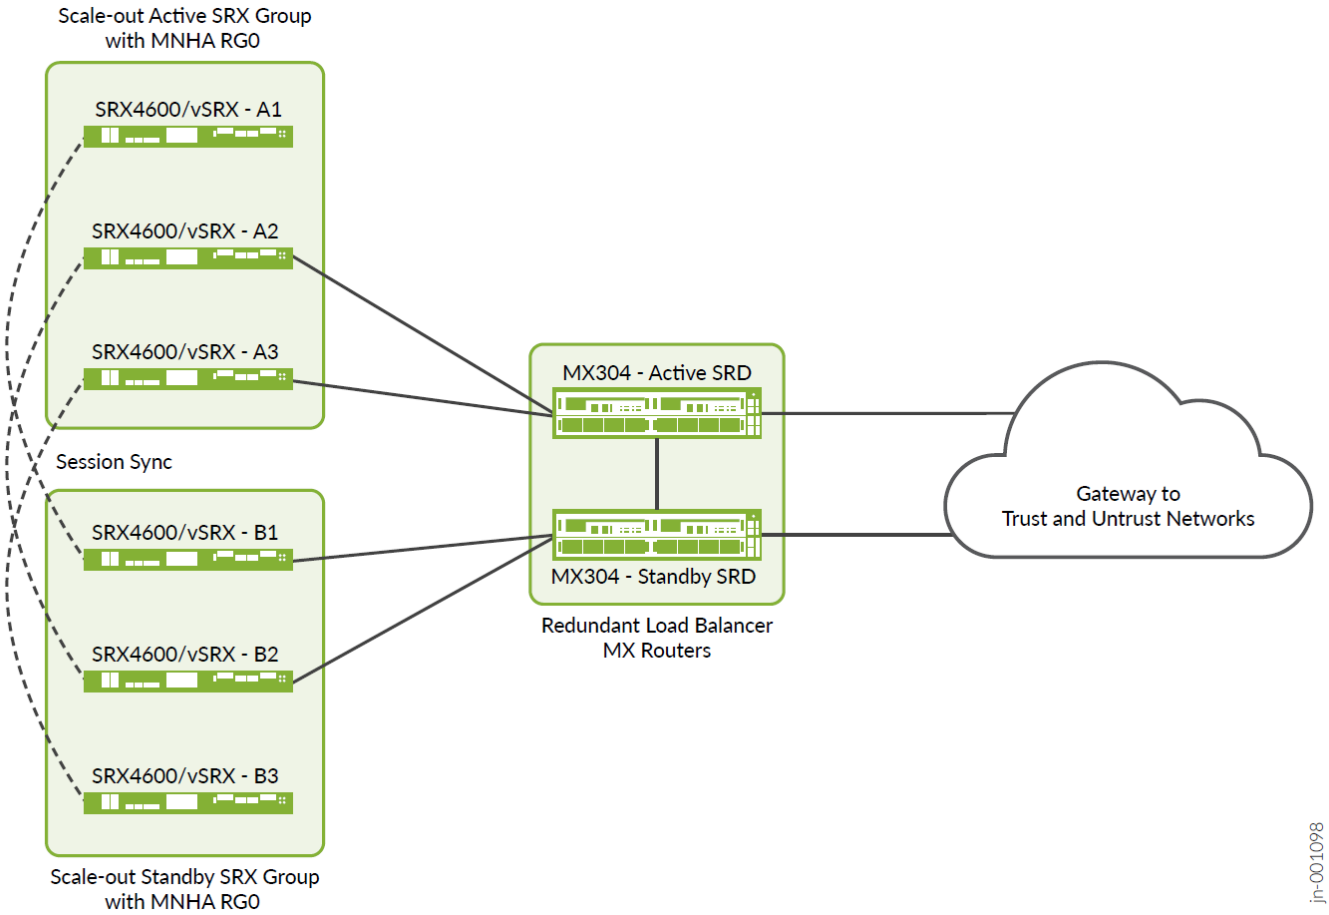

In case of a network failure detected by an active MX Series Routers, the second MX Series Routers takes over the active role and all traffic is redirected to this active MX Series Routers. It means the traffic sent to the previously backup SRX Series Firewall is becoming master of the MNHA pair. This architecture allows the use of only one SRX Series Firewall of a pair at a time, basically the SRX Series Firewall connected to the same MX Series Routers. However, in case of any failover, the traffic continues the second node of each MNHA pair.

On the SRX Series Firewall side, Multi-Node High Availability (MNHA) allows both SRX Series Firewall to handle and synchronize the sessions and offer any requested security services on both the firewalls. Since this topology uses SRG0 (active/active) as cluster mode, there is no need to failover the MNHA SRX Series Firewall pair to the redundant SRX Series Firewall when the MX Series Router detects a failure. The session synchronization in the MNHA pair ensures that the redundant SRX Series Firewall assumes responsibility for the sessions previously processed by the other SRX Series Firewall while maintaining session state.

Note that, when an SRX Series Firewall detects a failure, a failover occurs in the MNHA pair.

- Pros: Simple redundancy and scaling with each SRX Series Firewall pair
- Cons: half of the architecture is active at a time

Topology 3 – TLB – Single MX Series Router Scaled Out MNHA SRX Pairs (Multiple Pairs of SRX Series Firewalls)

This topology does offer redundancy for the SRX Series Firewalls and not for the MX Series Routers, though this one may have a second Routing Engine (RE) installed in the appropriate slot. In that case, this solution does not use two MX Series Routers.

Figure 3: Topology 3- TLB - Single MX with Scaled out SRX MNHA Pairs

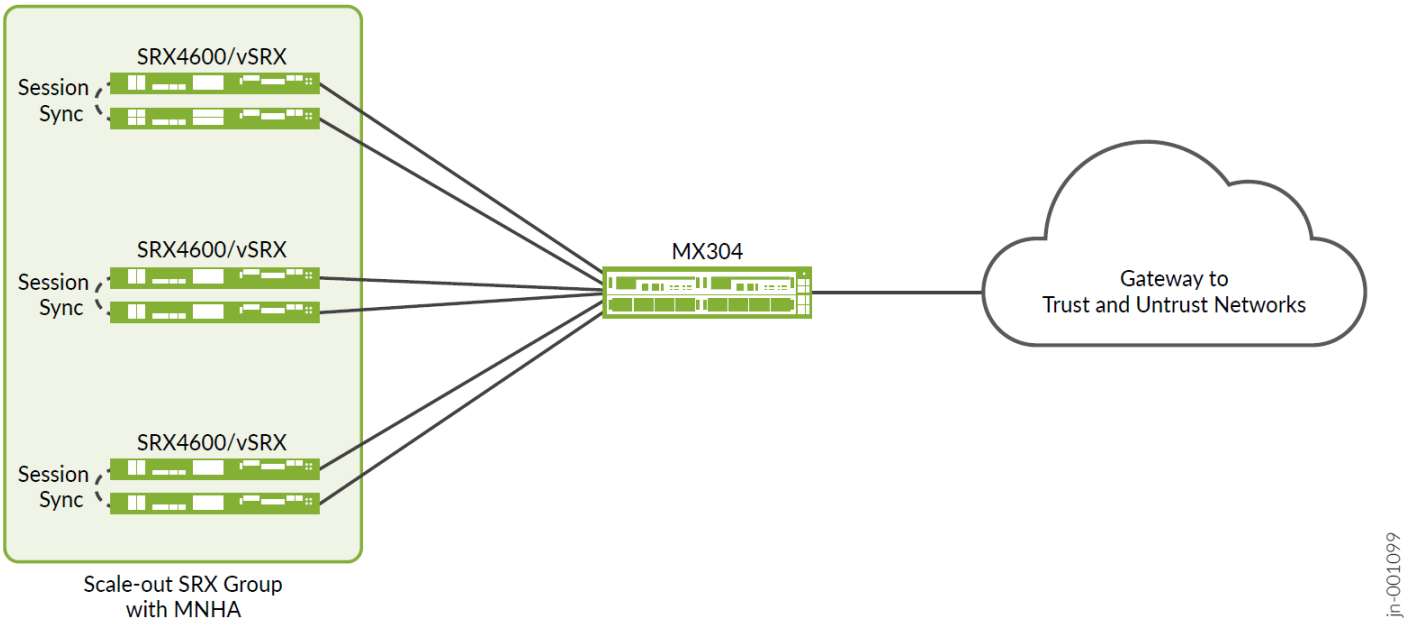

MNHA offers session synchronization within a cluster and helps with any failure scenario.

- Pros: Redundancy and scaling with each SRX Series Firewall pair
- Cons: No redundancy on the router (except using dual RE)

Topology 4 – TLB – Dual MX Series Routers Scaled Out MNHA SRX Pairs (Multiple Pairs of SRX Series Firewalls)

This topology offers redundancy for both the MX Series Routers and SRX Series Firewalls and takes advantage of having all the components used at the same time. Any failover scenario can be covered.

Figure 4: Topology 4-TLB - Dual MX with Scaled out SRX MNHA Pairs

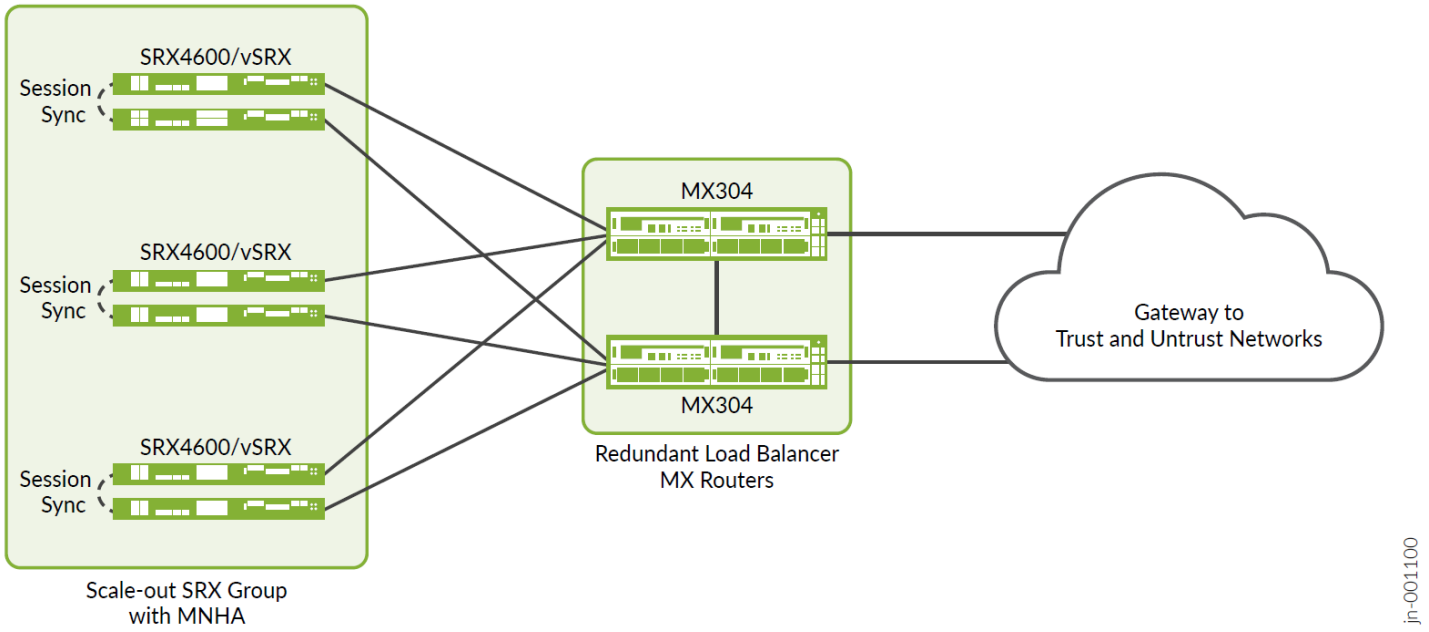

The MX Series Routers can handle traffic on any of the two routers, while SRX Series Firewall can be used either in the Active or Backup role and even in the Active-Active role, making use of both nodes at the same time. This augments the capacity of the network during normal operation. However, this leaves one node active at a time when a failure occurs (in case of single MNHA cluster).

All other SRX Series Firewall pairs may have an independent failover from the other SRX Series Firewall pairs and MX Series Routers if any one node fails within a cluster.

- Pros: Full redundancy and scaling for MX Series Routers and SRX Series Firewall pairs
- Cons: none

Platforms Tested

This JVD is tested on the following platforms.

Table 2: Platforms Tested

Name Convention	Supported Platforms	OS
Forwarding Node	MX304	Junos OS Release 23.4R2
Service Node	SRX4600	Junos OS Release 23.4R2
Service Node	vSRX	Junos OS Release 23.4R2 running on Linux-KVM

Version Qualification History

This JVD has been qualified in Junos OS Release 23.4R2.

Scale and Performance Data

This section shares information about the parameters that are scaled for the devices under test.

Table 3: Scale Numbers for the Devices Under Test (DUTs)

Traffic Profile	Throughput/MNHA-Pair	Session Count/MNHA-Pair	Traffic Type	File Size	Session Type
1	100Gbps	1000000	UDP	IMIX	Long lived (PPS)
2	100Gbps	1000000	TCP	4k	Long lived (PPS)
3	N/A	100000	TCP	1byte	Short lived (CPS)

According to scale numbers for the devices under test, a total of 2.1 million sessions with all the three traffic profiles run together.

The performance details are as follows:

Table 4: Performance Details

Platform	CPS/MNHA-Pair	Throughput/MNHA-Pair	CPU/vSRX
SRX4600	100K CPS	200Gbps	90%

According to the performance details,

- TCP - 200G throughput is generated using two million long lived sessions, unique source IP address, and source ports going to two destination HTTP servers.
- TCP - 100K connections per second (100K TCP session create and delete happening at the same time with 1 byte http transaction for each TCP session).

Hence, it is decided to use /8 source prefix for each of these traffic profiles and advertise routes between MX Series Routers and SRX Series Firewall. Route scaling is not tested as part of this JVD.

Event Testing

The following SRX Series Firewall failure events have been tested:

- MX Series Routers to SRX Series Firewall link failures
- SRX Series Firewall reboot
- SRX Series Firewall power off
- Complete MNHA pair power off

The following MX Series Router failure events have been tested:

- Reboot MX Series Router
- Restart routing process
- Restart traffic-dird daemon
- Restart Network-monitor daemon
- Restart sdk-process
- GRES

- TLB next-hop addition/deletion [adding/deleting new scale out SRX Series Firewall MNHA pair]
- SRD based CLI switchover between MX Series Routers

Traffic recovery is validated after all failure scenarios. With TCP traffic profiles ixia retry is configured with [1sec * 3]. With this no resets should be seen during basic MNHA failovers. <1% resets can be seen during failure events testing.

UDP traffic generated using IxNetwork for all the failure related test cases is used to measure the failover convergence time.

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: +31.207.125.700
Fax: +31.207.125.701