Juniper® Validated Design

JVD Test Report Brief: Scale-Out Stateful Firewall and CGNAT for SP Edge

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

Introduction

Today 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 highly 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.

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 scale-out 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 its 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 such as EBGP, BFD, Equal Cost Multipath with consistent hashing (ECMP CHASH), SRD, and TLB on MX as a forwarding plane and various security features such as CGNAT, stateful firewall, and MNHA on SRX Series Firewall.

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
	Dual MX (SRD)	CGNAT/NGFW	Yes	No
Traffic Load Balancer [TLB]	Single MX	CGNAT/NGFW	Yes	Yes
	Dual MX	CGNAT/NGFW	Yes	Yes

Test Topology

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

This topology is simple and least redundant. The resiliency is provided at MX router, with a redundant RE, PSU, etc however, there is no protection against MX-node failure. Deployment provides protection against service node failure by redistributing traffic flows between two remaining security nodes. Though there is no session synchronization between the SRX Series Firewalls, which leads to longer restoration time for the affected flows.

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

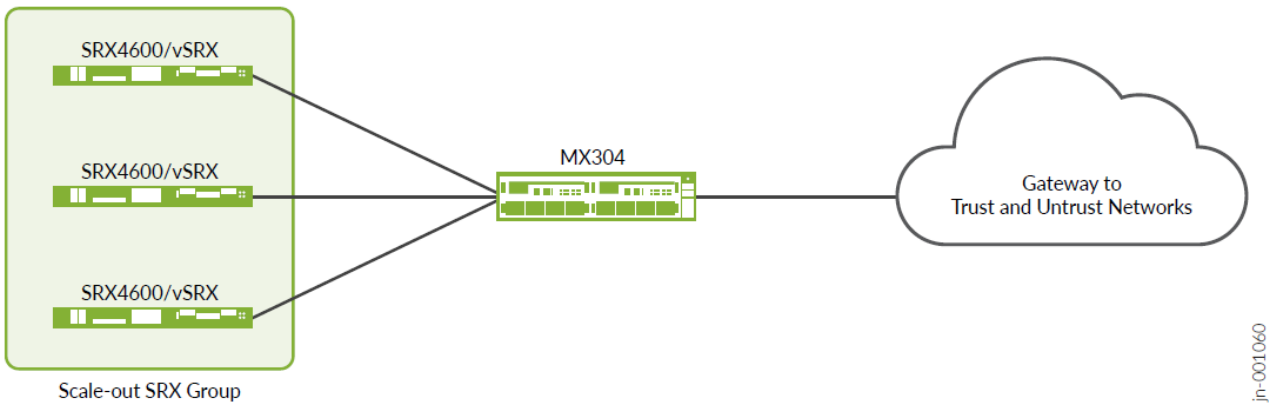

Network operators that are not concerned about stateful failover may want to simply augment security service capacities by adding more SRX Series Firewalls. The application sessions may be short lived anyway (for example, a redundancy mechanism may be handled at the application level so session sync between two different firewalls is not required).

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

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

This topology does offer redundancy at both the MX routers and for each redundant SRX Series Firewalls pair. The redundant pair of MX routers uses an SRD mechanism providing monitoring of physical elements of the network and/or the MX routers itself, as well as any other routing and system events that may need to trigger a failover to the other MX router.

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

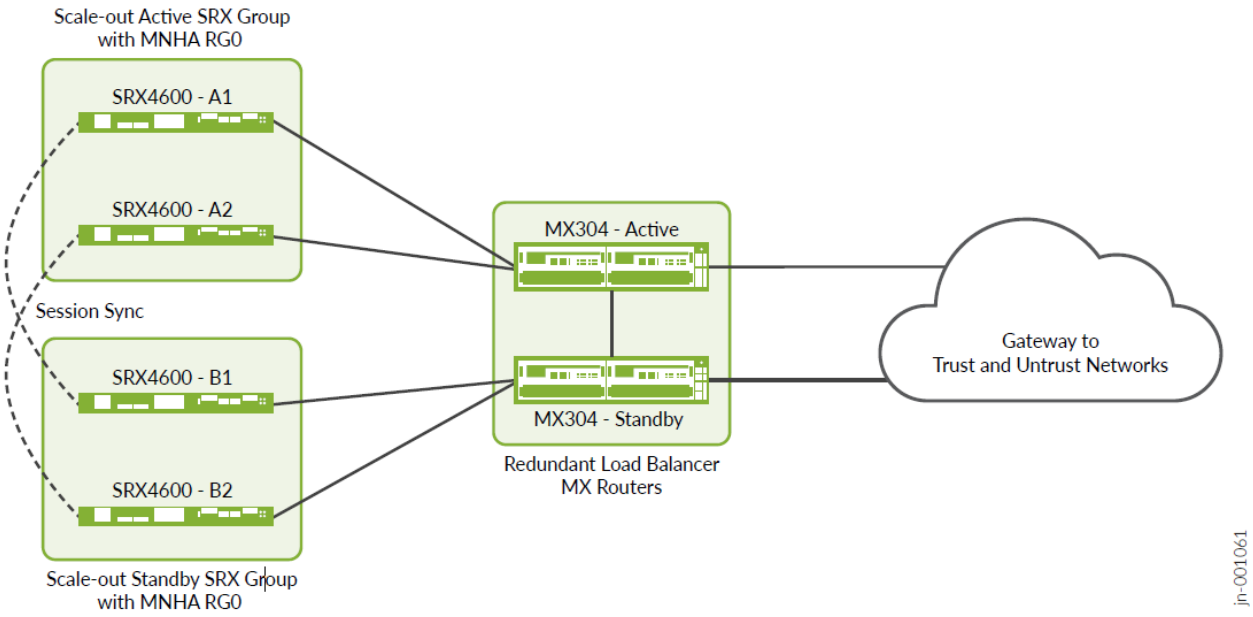

In case of a network failure detected by the active MX router, the second MX router takes over the active role and all traffic is then redirected to this active MX router. It means that traffic is sent to the previous backup SRX Series Firewalls, becoming master of the MNHA pair. This architecture only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

On the SRX Series Firewalls side, MNHA allows both SRX Series Firewalls to handle and synchronize sessions and support any requested security services on both firewalls. Since this topology uses SRG0 as cluster mode, there is no need of failing over a SRX Series Firewalls to the other firewall in case of any failure detected by the MX router (only when detected by SRX Series Firewalls itself). The session synchronization allows any traffic coming from the MX router (at SRD level) to process traffic for existing sessions, and any new sessions coming to it.

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

Deployment Scenario 3 – TLB – Single MX Scaled Out MNHA SRX Pairs (Multiple Pairs of SRX Series Firewall)

This topology does offer redundancy for the SRX Series Firewall however, not for the MX router, though this one may have a second Routing Engine (RE) installed in the appropriate slot and is not using two MX chassis in that case.

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

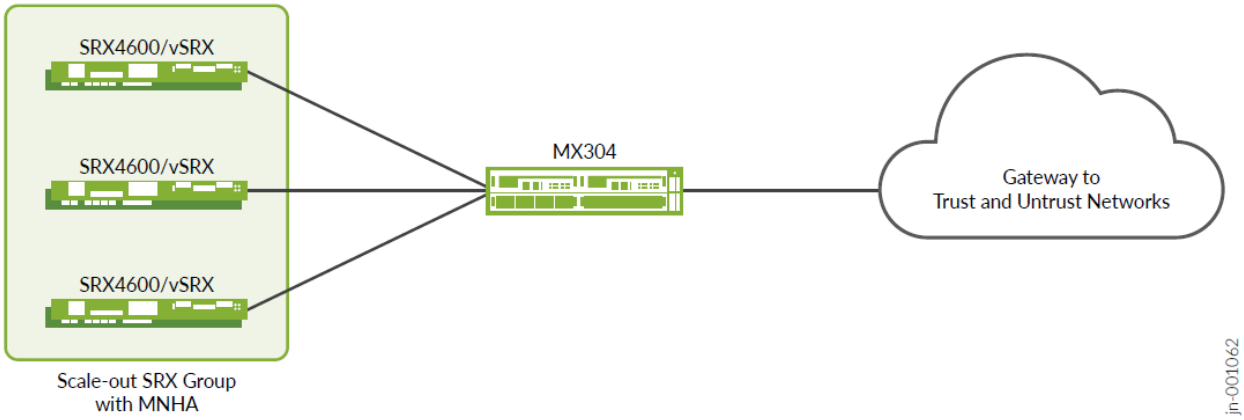

MNHA offers sessions synchronization within a cluster and help with any failure scenario.

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

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

This last topology offers the most redundancy for both MX and SRX Series Firewalls nodes and takes advantage of having all 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

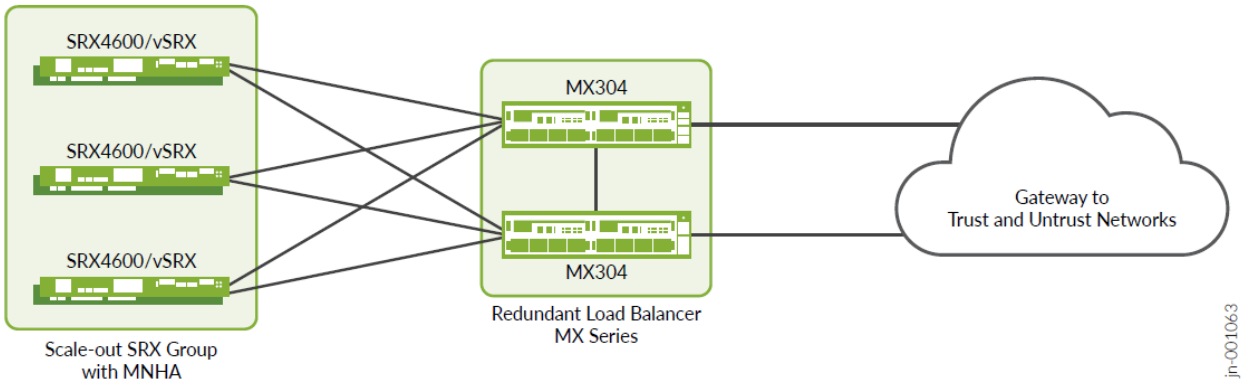

MX Series Routers handle traffic on any of the two routers, while SRX Series Firewalls can be used either in active or backup role or in active or active role, making use of both nodes at the same time. This augments the capacity of the network during normal operation, however this leaves one active role at a time when a failure occurs (consider a single MNHA cluster).

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

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

Platforms Tested

This JVD is tested on the following platforms.

Table 2: Platforms Tested

Role	Platform	OS
EDGE	MX304	23.4R2
EDGE	SRX4600	23.4R2
EDGE	vSRX	23.4R2

Version Qualification History

This JVD has qualified in Junos OS Release 23.4R2.

Scale and Performance Data

This section shares information about the parameters that are scaled along with the scale figures.

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

SRX4600	1	SFW/CGNAT
vSRX	1	SFW/CGNAT
Platform	SRX Count	Services

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, using 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 Router and SRX Series Firewalls. Route scaling is not tested as part of this JVD.

High Level Features Tested

The following features have been tested.

Table 5: Tested Traffic Profiles

Platform	SRX Count	Services	Traffic Type
SRX 4600	1	SFW/CGNAT	TCP
vSRX	1	SFW/CGNAT	TCP

Packet size is an Internet mix, which is an average packet size of ~700 bytes.

Table 6: Packet Size Details

Packet Size	Weight
64	8
127	36
255	11
511	4
1024	2
1518	39

Event Testing

The following SRX Series Firewalls failure events have been tested:

- MX Series Routers to SRX Series Firewalls link failures
- SRX Series Firewall reboot
- SRX Series Firewall power off
- Complete MNHA pair power off
- Restart SUB/PUB broker process

The following MX Series Router failure events have been tested:

- Reboot MX Series Router
- Restart routing process
- Restart traffic-dird deamon
- Restart Network-monitor deamon
- 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.

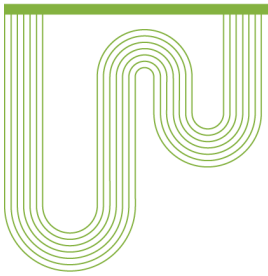

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

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

Send feedback to: design-center-comments@juniper.net V1.0/240912