

# Junos<sup>®</sup> Space Edge Services Director Release 1.0 Performance and Benchmarking

Release 1.0  
28 September 2015

Junos Space Edge Services Director application enables unified lifecycle management of services and policies on Juniper Networks MX Series routers that function as services delivery gateways in your network.

## Contents

Scalability and Performance of Edge Services Director . . . . .	2
Benchmarking Scenarios for Collection of Performance Management	
Counters . . . . .	2
Variation in the Clustered Setup . . . . .	3
Variation in the Polling Interval . . . . .	3
Variation in the Number of Attributes Being Polled . . . . .	4
Variation in the Monitoring Period of Polled Data . . . . .	4
SDG Discovery and Monitoring Scenarios . . . . .	5
Service Template and Deployment Plan Scenarios . . . . .	6
North-Bound API Test Scenarios . . . . .	7
Serial REST API Calls . . . . .	7
Parallel REST API Calls . . . . .	8
South-Bound API Test Scenarios . . . . .	8
Serial REST API Calls . . . . .	9
Parallel REST Calls . . . . .	9
Backup and Restore of Data . . . . .	9
Temporary Shutdown of the Junos Space Server . . . . .	9
PM Statistics With SDG Switchover . . . . .	9
PM Statistics Collection With Routing Engine Switchover . . . . .	9
Junos Space Documentation and Release Notes . . . . .	9
Documentation Feedback . . . . .	10
Requesting Technical Support . . . . .	10
Self-Help Online Tools and Resources . . . . .	11
Opening a Case with JTAC . . . . .	11
Revision History . . . . .	11

## Scalability and Performance of Edge Services Director

---

This document describes the scalability limits and the performance metrics that have been qualified with Edge Services Director 1.0. These scaling tests and validations enable you to configure your network topology in an optimal and effective manner using Edge Services Director for administration, provisioning, and monitoring of service delivery gateway (SDG) devices. An important aspect of any network management system is to monitor, control, and plan the network infrastructure that comprises a large number of devices and extensive configuration parameters in a streamlined, easy, and cohesive way. The bulk, single-step propagation of settings on large sets of devices without impacting the working efficiency and traffic-handling capacity of the network is a salient objective. With networks constantly increasing in size, heterogeneity, and complexity, effective management and planning for such network becomes more important. The benchmarking tests that have been performed certify the resilience and efficiency of networks with the tested settings.

The important objectives that are considered for such a planning are:

- Identifying the data to be collected.
- Measurement strategy or interpreting the collected data.
- Publish the Threshold events in network.
- Presentation of the data, which helps in analyzing the networks performance.

The following routers, test equipment, and tools have been used to perform the scalability and performance validation tests:

- MX240, MX480, and MX960 routers with a minimum of one MS-DPC card and a maximum of up to six MS-DPC cards.
- Five Junos Space Appliance nodes in a cluster.
- 40 pairs of service delivery gateways (SDGs).
- Real Servers for traffic load balancing (TLB) and application delivery controller (ADC) with valid ADC licenses.
- Third-party test equipment, namely IXIA ports, for traffic generation Junos Space software running on a JA1500 Appliance and a JA2500 Appliance

REST API calls are made for discovery and management of devices, and for collection of performance management statistics. DMI simulators and SNMP simulators are used to simulate the scaling environment.

## Benchmarking Scenarios for Collection of Performance Management Counters

The following scenarios have been evaluated for collection of performance management (PM) statistics and data polled from managed SDG devices:

### Variation in the Clustered Setup

---

A scenario in which Junos Space software manages approximately 40 high-availability pairs of SDG devices and 40 SDG devices is used for this scaling test. The SDG devices are running either Junos OS Release 12.1 or Junos OS Release 14.1. During the test, provisioning on SDG devices is in progress concurrently. Initially, PM collection is performed in a Junos Space cluster containing 5 nodes and also on standalone nodes. When significant performance difference between a standalone and clustered setup is seen, the number of nodes in the clustered topology are modified to achieve an optimal condition. Performance improvements are obtained by varying the number of nodes in the clustered setup. A total of 40 pairs of SDG devices, the number of attributes in each polling cycle, and a polling interval of 15 minutes for both DMI and SNMP channels are not changed and are statically maintained. The following parameters are monitored:

- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

The SDG impact on the JBoss server is monitored and also the hard disk space is assessed. These variables are monitored for a standalone deployment and also a clustered deployment with 5 nodes. No significant difference or degradation in PM collection was observed. The polling interval used was 15 minutes for PM collection with both DMI and SNMP channels. The following table represents the variation of the following variables for different clustered setups and a standalone setup:

- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

### Variation in the Polling Interval

---

A scenario in which Junos Space software manages approximately 40 high-availability pairs of SDG devices and 40 SDG devices is used for this scaling test. The SDG devices are running either Junos OS Release 12.1 or Junos OS Release 14.1. During the test, provisioning on SDG devices is in progress concurrently. The initial polling interval is 15 minutes and the time taken for each polling is monitored. A total of 40 pairs of SDG devices, the number of attributes in each polling cycle, and a clustered setup of 5 nodes are not changed and are statically maintained. The following parameters are monitored:

- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

If time taken for each polling is less 15 minutes, the polling frequency is decreased to 10 minutes and the aforementioned parameters are monitored. The following table represents the variation of the following variables based on the polling interval:

- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

### **Variation in the Number of Attributes Being Polled**

---

A scenario in which Junos Space software manages approximately 40 high-availability pairs of SDG devices and 40 SDG devices is used for this scaling test. During the test, provisioning on SDG devices is in progress concurrently. The attributes that are polled for PM collection depend on the number of real servers and virtual servers for TLB and ADC, and the service Instances of carrier-grade NAT (CGNAT) and stateful firewall (SFW). The test is performed with the following configurations:

- Approximately 250 Real Servers for TLB. Real Server Status KPI for each TLB Real Server and 4 other TLB KPIs.
- Approximately 150 Real Servers for ADC. Real Server Status KPI for each ADC Real Server and 9 other ADC KPIs.
- 7 CGNAT-related KPI attributes.
- 4 SFW-related KPI attributes.
- 15 HA-related KPI attributes.
- 20 MX routers-level KPI attributes.
- Total number of attributes for each SDG pair =  $250 + 150 + 9 + 4 + 7 + 4 + 15 + 20 = 459$ .
- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

### **Variation in the Monitoring Period of Polled Data**

---

A scenario in which Junos Space software manages approximately 40 high-availability pairs of SDG devices and 40 SDG devices is used for this scaling test. During the test, provisioning on SDG devices is in progress concurrently. During ambient polling intervals, few pairs of SDGs are unmanaged and managed. PM collection is observed for various time periods, such as 6 hours, 12 hours, 24 hours, 48 hours, 1 week, and 2 weeks. Perl or shell scripts are used for testing. PM collection on 40 devices is triggered concurrently during each polling cycle. The following table represents the variation of the below server variables based on the number of attributes being polled:

- CPU Percentage
- Memory Utilization of the processes
- Memory Leak Monitor SDG footprint in JBoss
- Hard disk space

## SDG Discovery and Monitoring Scenarios

The following scenarios are validated for discovery, management, and monitoring of SDGs:

- Resynchronization of all the discovered devices using the Junos Space Platform GUI and Edge Services director GUI.
- Configuration updates to device settings are discovered and displayed properly in the Edge Services Director application.
- Concurrently changing all of the 40 SDG devices that are in unmanaged state and synchronized to be managed devices.
- Manage 40 SDGs simultaneously with different SDG groups and KPI templates other the default groups.
- Representation of all of the 40 pairs of SDGs on the Dashboard page by verifying the correct display of statuses and values in the Health Status, Ticker, and Validate Alarm widgets.
- On the Dashboard page, change of status to service or tile with 40 pairs in managed state with 3 parallel users occurs properly.
- In the Monitoring page, all 40 SDG pairs are loaded in the Monitoring tree pane. All the monitoring widget graphs are plotted properly for all the 40 SDG pairs.
- In Fault mode, Alarms and Critical Messages Widgets are updated instantly when new traps or alarms are generated for all 40 SDG pairs
- When a new service instance is created or provisioned, it is properly processed by the respective services in the monitoring workspace for all 40 SDG pairs.
- Verification of all Statistics and Chassis view widgets with 40 pairs of managed SDGs in Monitor mode.
- Validation of the Correlation Engine status computation for multiple devices from 40 SDG pairs.
- For the 40 SDG pairs, PM collection for a period of 2 weeks is validated. PM collection occurs properly for all the intervals for all the SDG pairs during this 2 weeks.
- For the 40 SDG pairs, round-robin database (RRD) and PM collection are validated for all the intervals for all the SDG pairs for 2 weeks.

## Service Template and Deployment Plan Scenarios

The following configuration scenarios are examined for the proper functioning of service templates and deployment plans in a scaled environment that contains a Junos Space cluster of 5 nodes and 40 pairs of SDG devices:

- Simultaneous creation or modification of ADC Service Templates by importing configuration from an existing device or object builder by three users.
- Simultaneous creation or modification of TLB Service Templates by importing configuration from an existing device or object builder by three users.
- Simultaneous creation or modification of CGNAT Service Templates by importing configuration from an existing device or object builder by three users.
- Simultaneous creation or modification of SFW Service Templates by importing configuration from an existing device or object builder by three users.
- Simultaneous creation or modification of ADC deployment plans by three users.
- Simultaneous creation or modification of TLB deployment plans by three users.
- Simultaneous creation or modification of CGNAT deployment plans by three users.
- Simultaneous creation or modification of SFW deployment plans by three users.
- Simultaneous preview, validation, and scheduling of ADC deployment plans by three users.
- Simultaneous preview, validation, and scheduling of TLB deployment plans by three users.
- Simultaneous preview, validation, and scheduling of CGNAT deployment plans by three users.
- Simultaneous preview, validation, and scheduling of SFW deployment plans by three users.
- Commissioning of 5 ADC deployment plans (including Service Edit Deployment plans) as part of a single transaction.
- Commissioning of 5 TLB deployment plans (including Service Edit Deployment plans) as part of a single transaction.
- Commissioning of 5 CGNAT deployment plans (including Service Edit Deployment plans) as part of a single transaction.
- Commissioning of 5 SFW deployment plans (including Service Edit Deployment plans) as part of a single transaction.
- Simultaneous commissioning of 15 ADC deployment plans (including Service Edit Deployment plans), each as part of a single transaction by 5 users.
- Simultaneous commissioning of 5 TLB deployment plans (including Service Edit Deployment plans), each as part of a single transaction (deployment plan with 15 SDG pairs) by 2 users.

- Simultaneous commissioning of 15 CGNAT deployment plans(including Service Edit Deployment plans), each as part of a single transaction by 3 users.
- Simultaneous commissioning of 15 SFW deployment plans(including Service Edit Deployment plans), each as part of a single transaction by 3 users.
- Simultaneous editing of CGNAT, SFW, ADC, and TLB services and sending them for deployment by 3 users.
- Simultaneous editing of Packet Filter and sending it for deployment and provisioning by 3 users.
- Verification of statistics and import of objects to Object Builder with 40 pairs of managed SDGs.

## North-Bound API Test Scenarios

The following north-bound API test scenarios are validated:

### Serial REST API Calls

Serial REST API calls for same or different KPIs for the same or different pair of SDG devices and for the same or different time ranges are made. A scenario in which Edge Services Director manages approximately 40 pairs of SDG devices with polling every 15 minutes is deployed. These REST API calls are sequentially made to the performance manager module for KPIs for different time ranges such as 1 hour, 12 hours, 24 hours, 1 week, and 2 weeks. REST calls to the PM module for KPIs return the following output:

Time	Duration	Data	Points
1	Hour		6
12	Hours		12
24	Hours		24
1	Week		7
2	Weeks		5

The following scenarios are validated with serial REST API calls:

- Serial REST calls for the same KPI attributes for the same SDG pair and same time range from the same client three times.
- Serial REST calls for the same KPI attributes for the same SDG pair and same time range from different clients 3 times.
- Serial REST calls for the same KPI attributes for the same SDG pair and different time range from the same client three times.
- Serial REST calls for the same KPI attributes for the same SDG pair and different time range from different clients three times.
- Serial REST calls for the same KPI attributes for different SDG pair and different time range from the same client three times.
- Serial REST calls for the same KPI attributes for different SDG pair and different time range from different clients three times.

- Serial REST calls for different KPIs for different SDG pairs and different time ranges from the same client three times.
- Serial REST calls for different KPIs for different SDG pairs and different time ranges from different clients three times.

### Parallel REST API Calls

---

Parallel REST API calls for same or different KPIs for the same or different pair of SDG devices and for the same or different time ranges are made. A scenario in which Edge Services Director manages approximately 150 pairs of SDG devices with polling every 15 minutes is deployed. Also, simulation of data is performed for a period of 6 months to one year and REST API queries are triggered. These REST API calls are made in parallel to the performance manager module for KPIs for different time periods, such as 1 hour, 12 hours, 24 hours, 1 week, and 2 weeks. REST calls to the PM module for KPIs return the following output:

Time Duration	Data Points
1 Hour	6
12 Hours	12
24 Hours	24
1 Week	7
2 Weeks	5

The following scenarios are validated with parallel REST API calls:

- Parallel REST calls for same KPIs for the same SDG pair and same time range from the same client three times.
- Parallel REST calls for the same KPI attributes for the same SDG pair and the same time range from different clients three times.
- Parallel REST calls for the same KPI attributes for the same SDG pair and different time ranges from the same client three times.
- Parallel REST calls for the same KPI attributes for the same SDG pair and different time ranges from different clients three times.
- Parallel REST calls for the same KPI attributes for different SDG pairs and different time ranges from the same client three times.
- Parallel REST calls for the same KPI attributes for different SDG pairs and different time ranges from different clients three times.
- Parallel REST calls for different KPIs for different SDG pairs and different time ranges from the same client three times.
- Parallel REST calls for different KPIs for different SDG pairs and different time ranges from different clients three times.

### South-Bound API Test Scenarios

The following south-bound test scenarios are validated:



### Serial REST API Calls

---

A scenario in which Edge Services Director manages approximately 40 pairs of SDG devices with polling every 15 minutes is deployed. Serial REST API calls for fetching values from the MySQL database tables for different object IDs (OIDs) are followed by SNMP Get operation for the same OID from the Junos Space server to the device. The outputs from both these operations are compared and saved to a CSV file that contains columns for KPI name, date, and time. A script is used to read all the entries from the CSV file and a report is saved in .txt or html format. It is verified that all the entries match properly.

### Parallel REST Calls

---

A scenario in which Edge Services Director manages approximately 40 pairs of SDG devices with polling every 15 minutes is deployed. Parallel REST API calls for fetching values from the MySQL database tables for different object IDs (OIDs) are followed by SNMP Get operation for the same OID from the Junos Space server to the device. The outputs from both these operations are compared and saved to a CSV file that contains columns for KPI name, date, and time. A script is used to read all the entries from the CSV file and a report is saved in .txt or html format.

### Backup and Restore of Data

---

A scenario in which a Junos Space software installation is performed followed by the installation of Edge Services Director is used. After the install, the backed up data is restored on the Junos Space server. It is verified that all the managed SDG pairs of devices display correctly and it is also observed that collection of PM statistics starts again correctly.

### Temporary Shutdown of the Junos Space Server

---

A scenario in which the Junos Space server is shut down when it contains the collection data of already managed SDG pairs (10 pairs). When the Junos Space server is shut down and restarted after an hour, it is verified that collection of PM statistics starts again correctly.

### PM Statistics With SDG Switchover

---

A scenario that contains a managed real router pair of SDG devices along with a simulator pair of SDG devices is used. A switchover of the SDG device is performed and it is verified that collection of PM statistics restarts or continues.

### PM Statistics Collection With Routing Engine Switchover

---

A scenario that contains a managed real router pair of SDG devices along with a simulator pair of SDG devices is used. A Routing Engine switchover is performed and it is verified that collection of PM statistics restarts or continues.

---

## Junos Space Documentation and Release Notes

---

For a list of related Junos Space documentation, see  
<https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the *Junos Space Release Notes*.

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <https://www.juniper.net/books>.

---

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <https://www.juniper.net/documentation/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/documentation/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

---

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://www.juniper.net/support/requesting-support.html>.

## Revision History

---

28 September 2015—Revision 1, Junos Space Edge Services Director, Release 1.0

Copyright © 2018 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be 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.