SolarWinds

Scalability Engine Guidelines for SolarWinds Products

Technical Reference

Scalability Engine Guidelines for SolarWinds Products, Version 2015.1.3, 3/31/2016

Table of Contents

Using Orion Scalability Engines

Orion scalability engines, including Additional Polling Engines and Additional Web Servers, can extend the monitoring capacity of your SolarWinds installation.

Requirements and recommendations will vary from product to product. Refer to the Administrator Guide for your specific product for more information.

Scalability Engine Requirements

Scalability engine requirements are generally the same as the requirements for a primary polling engine.

> (i) SNMP access must be allowed to all SolarWinds polling engines. For more information, see the installation instructions in the Administrator Guide for your SolarWinds product.

Scalability Engine Guidelines by Product

The following sections provide guidance for using scalability engines to expand the capacity of your SolarWinds installation.

> (i) Requirements and recommendations will vary from product to product. Refer to the Administrator Guide for your specific product for more information.

- DameWare in Centralized Mode
- Database Performance Analyzer (DPA)
- Engineer's Toolset on the Web
- Enterprise Operations Console (EOC)
- IP Address Manager (IPAM)
- Log and Event Manager (LEM)
- NetFlow Traffic Analyzer (NTA)
- Network Configuration Manager (NCM)
- Network Performance Monitor (NPM)
- Patch Manager (SPM)
- Quality of Experience (QoE)
- Server & Application Monitor (SAM)
- Serv-U FTP Server and MFT Server
- Storage Resource Monitor (SRM)
- User Device Tracker (UDT)
- Virtualization Manager (vMan)
- VoIP & Network Quality Manager (VNQM)
- Web Performance Monitor (WPM)

DameWare in Centralized Mode

DAMEWARE SCALABILITY ENGINE GUIDELINES	
Scalability Options	150 concurrent Internet Sessions per Internet Proxy 5,000 Centralized users per Centralized Server 10,000 Hosts in Centralized Global Host list 5 MRC sessions per Console

Database Performance Analyzer (DPA)

DATABASE PERFORMANCE ANALYZER (DPA) SCALABILITY ENGINE GUIDELINES	
Scalability Options	Less than 20 database instances monitored on a system with 1 CPU and 1 GB RAM 20 - 50 database instances monitored on a system with 2 CPU and 2 GB RAM 51 - 100 database instances monitored on a system with 4 CPU and 4 GB RAM 101 - 250 database instances monitored on a system with 4 CPU and 8 GB RAM More than 250 database instances monitored through Central Server mode See "Link together separate DPA servers" in the SolarWinds DPA Administrator Guide

Engineer's Toolset on the Web

ENGINEER'S TOOLSET ON THE WEB SCALABILITY ENGINE GUIDELINES	
Scalability Options	45 active tools per Engineer's Toolset on the Web instance 3 tools per user session 1 active tool per mobile session 10 nodes monitored at the same time per tool 48 interfaces monitored at the same time per tool 12 metrics rendered at same time per tool

Enterprise Operations Console (EOC)

ENTERPRISE OPERATIONS CONSOLE (EOC) SCALABILITY ENGINE GUIDELINES	
Scalability Options	Maximum 1 million elements
WAN and/or Bandwidth Considerations	Minimal monitoring traffic is sent between the EOC server and any remote Orion servers or APEs.

IP Address Manager (IPAM)

IP ADDRESS MANAGER (IPAM) SCALABILITY ENGINE GUIDELINES	
Scalability Options	3 million IPs per SolarWinds IPAM instance

Log and Event Manager (LEM)

LOG AND EVENT MANAGER (LEM) SCALABILITY ENGINE GUIDELINES	
Scalability Options	Maximum 120 million events per day
	10,000 rule hits per day

NetFlow Traffic Analyzer (NTA)

NETFLOW TRAFFIC ANALYZER (NTA) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	No
Poller Remotability Available?	No
Primary Poller Limits	50k FPS per poller
	For more information, see Network Performance Monitor (NPM).
Scalability Options	5 APEs up to 300k FPS
	For more information, see Network Performance Monitor (NPM).
WAN and/or Bandwidth Considerations	1.5 - 3% of total traffic seen by exporter

NETFLOW TRAFFIC ANALYZER (NTA) SCALABILITY ENGINE GUIDELINES	
Other Considerations	See "Section 4 – Deployment Strategies" of NetFlow Basics and Deployment Strategies

Network Configuration Manager (NCM)

NETWORK CONFIGURATION MANAGER (NCM) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	No
Poller Remotability Available?	No
Primary Poller Limits	~10k devices
Scalability Options	One APE for every 10k devices, for NCM 7.1 and later Maximum of 30k devices per primary SolarWinds NCM instance (i.e. NCM server + 2 NCM APEs) Integrated standalone mode
WAN and/or Bandwidth Considerations	None
Other Considerations	None

Network Performance Monitor (NPM)

NETWORK PERFORMANCE MONITOR (NPM) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	Yes. Up to three total polling engines may be installed on a single server (i.e. one primary NPM polling engine with one or two additional polling engines or three additional polling engines on the same server). ⓘ A stack requires only 1 IP address, regardless of the number of APEs
Poller Remotability Available?	Yes, for NPM 10.4 and later ⓘ Poller remotability is a feature that enables the local storage, using MSMQ, of up to ~1 GB of polled data per poller in the event that the connection between the polling engine and the database is temporarily lost.

NETWORK PERFORMANCE MONITOR (NPM) SCALABILITY ENGINE GUIDELINES	
Primary Poller Limits	12k elements per polling engine at standard polling frequencies: • Node and interface up/down: 2 minutes/poll • Node statistics: 10 minutes/poll • Interface statistics: 9 minutes/poll 25-50 concurrent Orion Web Console users SNMP Traps: ~500 messages per second (~1.8 million messages/hr) Syslog: 700-1k messages/second (2.5 - 3.6 million messages/hr) ⓘ If you are monitoring more than ~100,000 elements, consider using SolarWinds Enterprise Operations Console.
Scalability Options	One polling engine for every ~12k elements Maximum of 100k elements per primary SolarWinds NPM server (i.e. 1 NPM server + 9 APEs). For more information about licensing, see "How is SolarWinds NPM licensed?"
WAN and/or Bandwidth Considerations	Minimal monitoring traffic is sent between the primary NPM server and any APEs that are connected over a WAN. Most traffic related to monitoring is between an APE and the SolarWinds database.
Other Considerations	How much bandwidth does SolarWinds require for monitoring? For hardware requirements, see "Orion Server Hardware Requirements" in the SolarWinds Orion NPM Administrator Guide.

Patch Manager (SPM)

PATCH MANAGER (SPM) SCALABILITY ENGINE GUIDELINES	
Scalability Options	1,000 nodes per automation server 1,000 nodes per SQL Server Express instance (SQL Server does not have this limitation)

Quality of Experience (QoE)

QUALITY OF EXPERIENCE (QOE) SCALABILITY ENGINE GUIDELINES	
Scalability Options	1,000 QoE sensors 50 application per sensor

Server & Application Monitor (SAM)

SERVER & APPLICATION MONITOR (SAM) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	Yes, for SAM 6.2 and later Two polling engines can be installed on a single server
Poller Remotability Available?	Yes, for SAM 5.5 and later ⓘ Poller remotability is a feature that enables the local storage, using MSMQ, of up to ~1 GB of polled data per poller in the event that the connection between the polling engine and the database is temporarily lost.
Primary Poller Limits	~8-10k component monitors per polling engine 25-50 concurrent Orion Web Console users
Scalability Options	One APE for every 8-10k component monitors Maximum of 150k component monitors per primary SolarWinds SAM installation (i.e. 1 SAM server + 4 APEs). You can use up to 14 APEs. For more information about licensing, see Why are you licensing by monitors instead of by servers?
WAN and/or Bandwidth Considerations	Minimal monitoring traffic is sent between the primary SAM server and any APEs that are connected over a WAN. Most traffic related to monitoring is between an APE and the SolarWinds database. Bandwidth requirements depend on the size of the relevant component monitor. Based on 67.5 kB / WMI poll and a 5 minute polling frequency, the estimate is 1.2 Mbps for 700 component monitors. For more information, see How do SNMP and WMI polling compare? ⓘ WMI is best suited for environments where latency is < 100ms.
Other Considerations	WMI Security Blog

Serv-U FTP Server and MFT Server

SERV-U FTP SERVER AND MFT SERVER SCALABILITY ENGINE GUIDELINES	
Scalability Options	500 simultaneous FTP and HTTP transfers per Serv-U instance

SERV-U FTP SERVER AND MFT SERVER SCALABILITY ENGINE GUIDELINES	
	50 simultaneous SFTP and HTTPS transfers per Serv-U instance
	For more information, see the Serv-U Distributed Architecture Guide.

Storage Resource Monitor (SRM)

STORAGE RESOURCE MONITOR (SRM) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	No. One APE instance can be deployed on a single host.
Poller Remotability Available?	Yes Note: Poller remotability is a feature enabling the local storage, using MSMQ, of up to ~1 GB of polled data per poller in case the connection between the polling engine and the database is temporarily lost.
Primary Poller Limits	20k disk per poller 25-50 concurrent Orion Web Console users
Scalability Options	1 APE per 10k disk polled
WAN and/or Bandwidth Considerations	Minimal monitoring traffic is sent between the primary SRM server and any APEs that are connected over a WAN. Most traffic related to monitoring is between an APE and the SolarWinds database.

User Device Tracker (UDT)

USER DEVICE TRACKER (UDT) SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	No
Poller Remotability Available?	No
Primary Poller Limits	100k ports
Scalability Options	One APE per 100k additional ports Maximum of 500k port per instance (1 main poller and 4 additional)

USER DEVICE TRACKER (UDT) SCALABILITY ENGINE GUIDELINES	
WAN and/or Bandwidth Considerations	None
Other Considerations	UDT version 3.1 supports the ability to schedule port discovery. In UDT version 3.1 the Max Discovery Size is 2,500 nodes/150,000 ports

Virtualization Manager (vMan)

VIRTUALIZATION MANAGER (VMAN) SCALABILITY ENGINE GUIDELINES	
Scalability Options	3000 VMs* 700 Hosts 75 Clusters 1800 Datastores *By using federated collectors, you can monitor 10,000 or more VMs. For information about federated collectors, see the Virtualization Manager documentation.

VoIP & Network Quality Manager (VNQM)

VOIP & NETWORK QUALITY MANAGER SCALABILITY ENGINE GUIDELINES	
Stackable Pollers Available?	No
Poller Remotability Available?	No
Primary Poller Limits	~5,000 IP SLA operations ~200k calls/day with 20k calls/hour spike capacity
Scalability Options	One APE per 5,000 IP SLA operations and 200,000 calls per day Maximum of 15,000 IP SLA operations and 200,000 calls per day per SolarWinds VNQM instance (i.e. SolarWinds VNQM + 2 VNQM APEs)
WAN and/or Bandwidth Considerations	Between Call Manager and VNQM: 34 Kbps per call, based on estimates of ~256 bytes per CDR and CMR and based on 20k calls per hour

VoIP & Network Quality Manager Scalability Engine Guidelines	
Other Considerations	None

Web Performance Monitor (WPM)

Web Performance Monitor (WPM) Scalability Engine Guidelines	
Stackable Pollers Available?	No
Poller Remotability Available?	No, but recordings may be made from multiple locations
Primary Poller Limits	Dozens of recordings per player
Scalability Options	One APE per dozens additional recordings, with the complexity of transactions determining the limits per player
WAN and/or Bandwidth Considerations	None
Other Considerations	None

Scalability Engine Deployment Options

The following sections discuss common scalability engine deployment options:

- Centralized Deployment
- Distributed Deployment
- Centralized Deployment with Remote Polling Engines

Centralized Deployment

This is the simplest deployment option, as there is only one SolarWinds Orion server, and software is only installed in the Primary region. This option is well suited to environments where most of the monitored nodes are located in a single, primary region and where other regional offices are much smaller. This deployment is optimal when the following conditions apply:

1. The remote office is not large enough to require a local SolarWinds Orion server instance or polling engine.
2. There are not enough monitored nodes to require a local SolarWinds Orion server instance or polling engine.
3. You prefer to have a central point of administration for the SolarWinds Orion server.
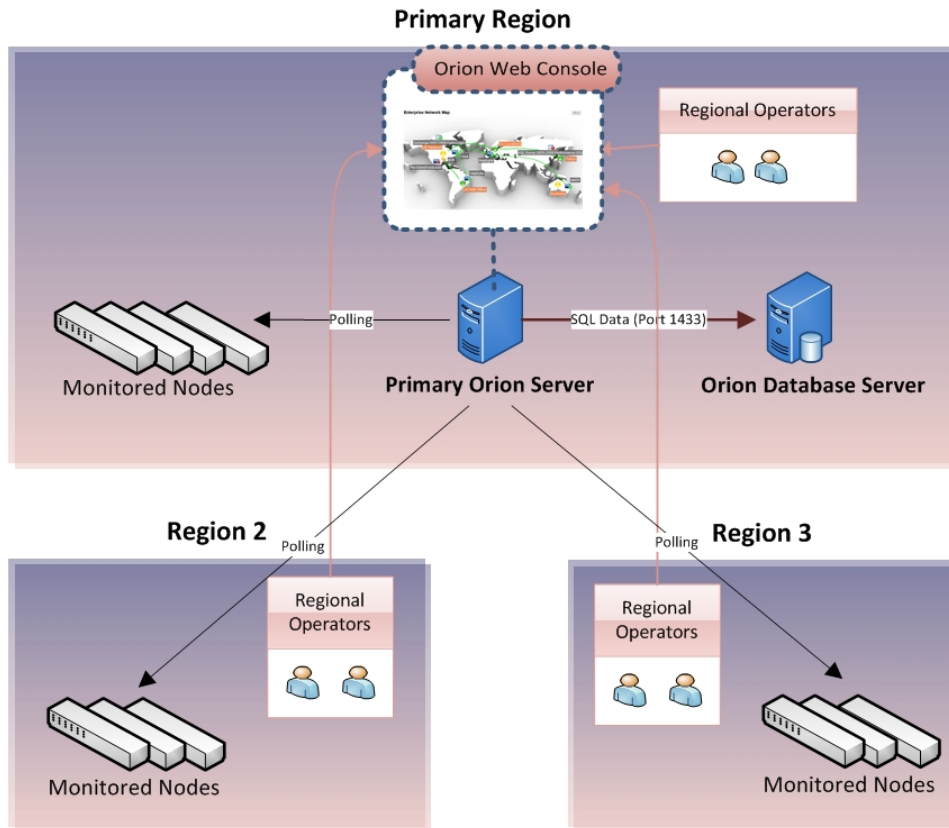
In a typical centralized deployment, the primary SolarWinds Orion server polls all data that is then stored centrally in the database server. Both the primary SolarWinds Orion server and the database server are in the Primary Region. To view data Regional Operators in each region must log into the Orion Web Console in the primary region, where your Orion Platform products are installed. Additional Web Servers are available and may be installed in secondary regions. If an Additional Web Server is deployed, a Regional Operator can log into a local web console to view all network data.

A reliable static connection is required between the primary region and all monitoring regions. This connection continually transmits monitoring data. The quantity of bandwidth consumed will depend on many factors, including the type and number of SolarWinds Orion Platform products that are installed and the types and quantity of monitored elements. It is difficult to precisely estimate bandwidth requirements, as each SolarWinds monitoring environment is unique.

ⓘ   ■   All nodes are polled from a single SolarWinds Orion server instance in the Primary Region, and all data is stored centrally on the database server in the primary region.

ⓘ   ■ Each installed module will need to have enough available licenses to cover all regions.

       ■ All KPI's, such as Node Response Times, will be calculated from the perspective of the Primary Orion Server. For example, the response time for a monitored node in Region 2 will be equal to the round trip time from the Primary Orion Server to that node.

Distributed Deployment

This is the traditional SolarWinds Orion distributed deployment option, comprising separate instances of SolarWinds Orion Platform products installed locally in each region with the Enterprise Operations Console (EOC) available as a top level dashboard to access data across all related instances.

This option is well suited to organizations with multiple regions or sites where the quantity of nodes to be monitored in each region would warrant both localized data collection and storage. It works well when there are regional teams responsible for their own environments, and when regional teams need autonomy over their monitoring platform. This option gives regional operators this autonomy and the ability to have different modules and license sizes installed in each region to match individual requirements. While the systems are segregated between regions, all data can still be accessed from the centrally located Enterprise Operations Console (EOC).
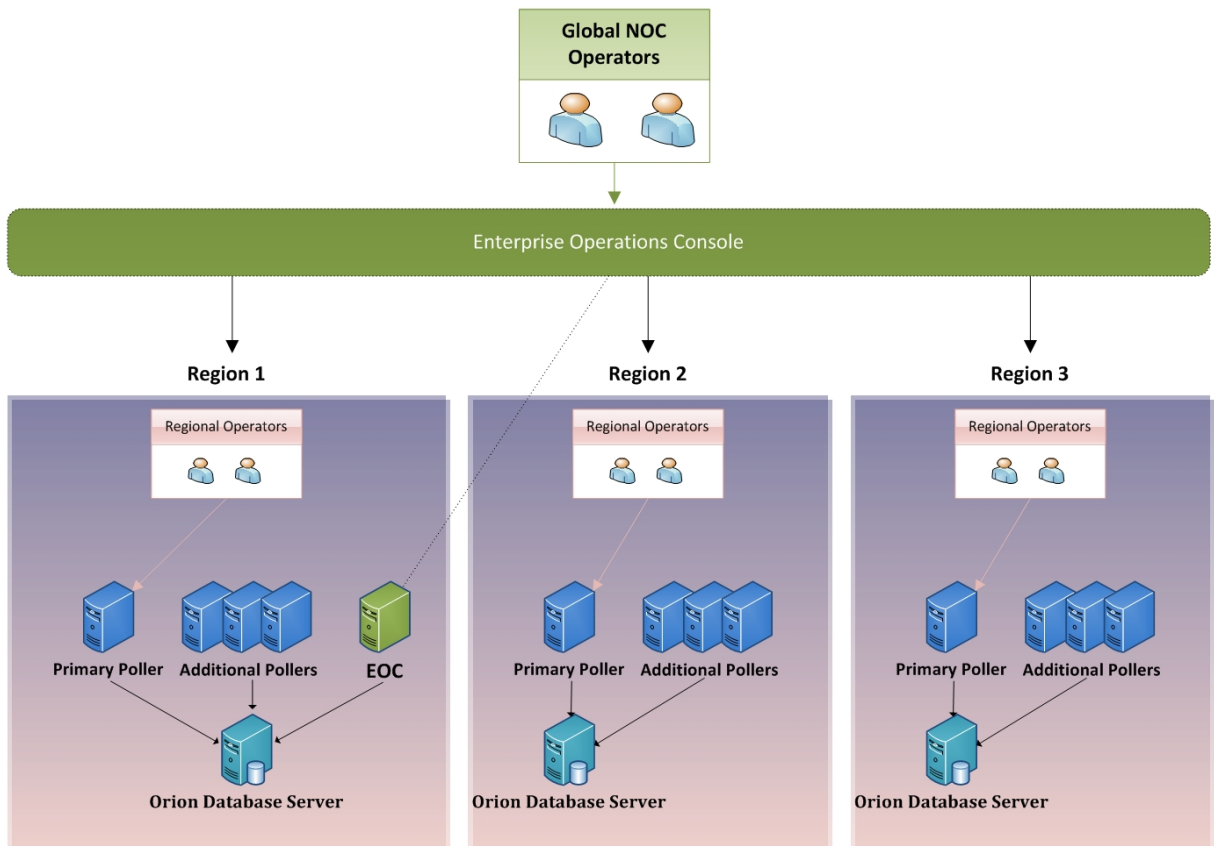
Each region is licensed independently, and data are polled and stored locally in each region. Modules and license sizes may be mixed and matched accordingly. In the example provided,

- Region 1 has deployed NPM SLX, SAM AL1500, UDT 50,000, and three additional polling engines
- Region 2 has deployed NPM SL500, NTA for NPM SL500, UDT 2500, and three additional polling engines
- Region 3 has deployed NPM SL100 only and three additional polling engines

As in this example, if EOC is used as a centralized dashboard to access data stored regionally, the following considerations apply:

- A reliable static connection is required between EOC and all monitoring regions.
- Each SolarWinds Orion server is incrementally polled for current status and statistics only. EOC does not store historical data. Because it only performs incremental polling for current status and statistics, the bandwidth used by EOC is not considered to be significant.

> ⓘ ■ Each region is managed, administered, and upgraded independently. For

(i)   example, node, user, alert, and report creation, deletion and modification are performed separately in each region. Certain objects, such as alert definitions, Universal Device Pollers, and Server and Application Monitor templates can be exported and imported between instances.

- Each region can scale independently by adding additional polling engines as required.

Centralized Deployment with Remote Polling Engines

This option combines the benefits of a centralized Orion instance with the flexibility of localized data collection. Management and administration is done centrally on the primary server. This is well suited to organizations that require centralized IT management and localized collection of monitoring data.
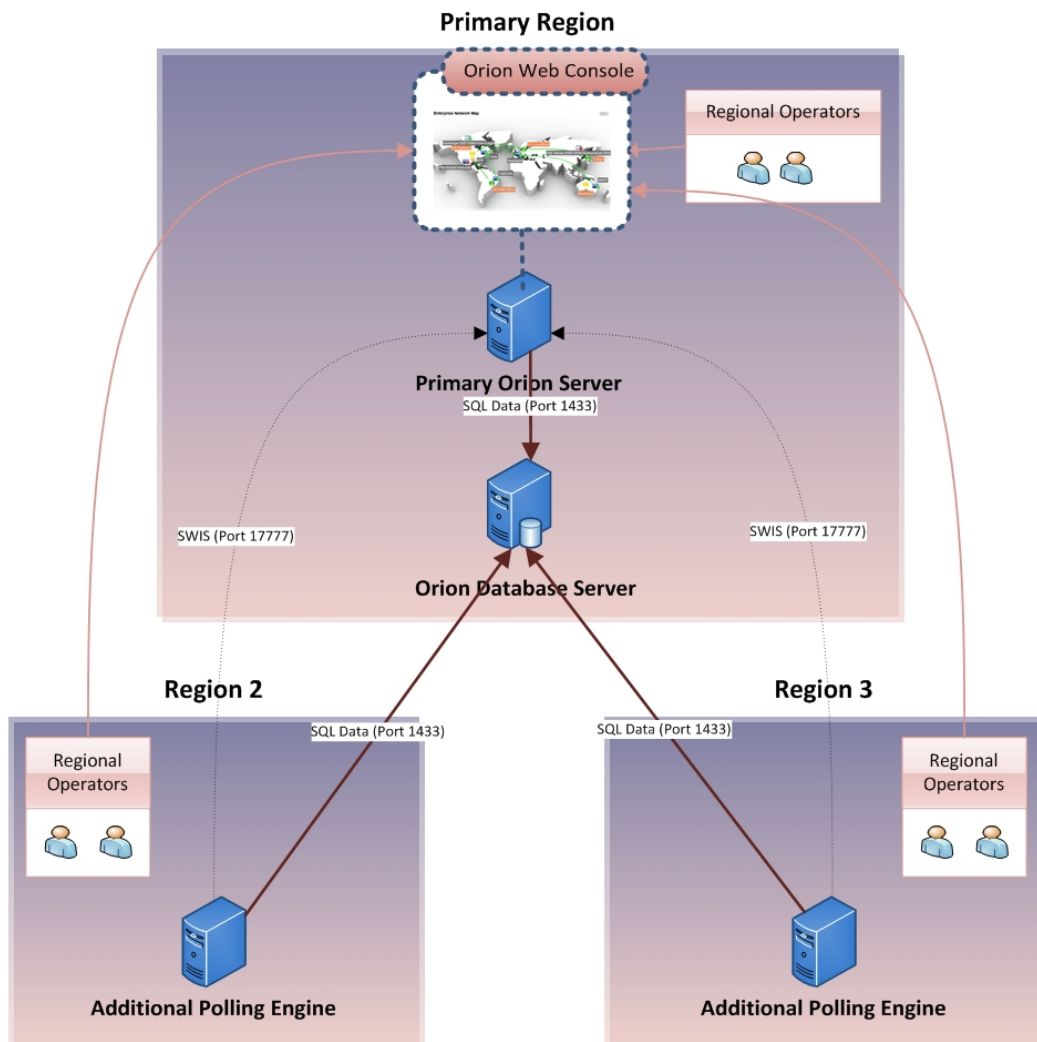
In a centralized deployment with remote polling engines, additional polling engines poll data locally in each region, and the polled data is then stored centrally on the database server in the primary region. Regional Operators in each region log into the Orion Web Console in the Primary Region where the primary SolarWinds Orion server is installed to view data.

Additional Web Servers are available and may be installed in secondary regions. Using an Additional Web Server, a Regional Operator can then log into a local web console to view all network data.

Notes:

- The combination of the Primary Orion Server, database server and all remotely deployed polling engines is considered to be a single SolarWinds Orion instance.
- This single instance is being managed and administered centrally. For example, node, user, alert, and report creation, deletion and modification is performed centrally on the Primary Orion Server only.
- When nodes are added, the user selects the polling engine to which the node is assigned. All data collection for that node is then performed by that polling engine, and nodes can be re-assigned between polling engines, as required.
- A reliable static connection must be available between each region.
    - This connection will be continually transmitting MS SQL Data to the Orion Database Server; it will also communicate with the Primary Orion Server.
    - The latency (RTT) between each additional polling engine and the database server should be below 300ms. Degradation may begin around 200ms, depending on your utilization. In general, the remote polling engine is designed to handle connection outages, rather than high latency. The ability to tolerate connection latency is also a function of load. Additional polling engines polling a large number of elements may be potentially less tolerant of latency conditions.
    - To calculate the bandwidth requirement for a remote polling engine, consider the following example. If the additional polling engine polls 800 SNMP nodes, each node containing 12 interfaces and two volumes, then the data flow between the polling engine and the database server is approximately 300 KB/s. This calculation only considers the polling activity with disabled topology, and does not take into account the bandwidth requirement associated with syslogs, traps and alerts.
- Each polling engine uses Microsoft Message Queuing (MSMQ).
    - This allows data to be cached locally on the additional polling engine servers in the event of a connection outage to the Orion Database Server.
    - The amount of data that can be cached depends on the amount of disk space available on the polling engine server. The default storage space is 1 GB. A general guideline is that up to one hour of data can be cached. When the connection to the database is restored, the Orion Database Server is updated with the locally cached data. The synchronization occurs in a FIFO order, meaning that the oldest data is processed first. This means that after the connection is restored, a period of time elapses before the most up-to-date polling data appears instantly in the database.
    - If the database connection is broken for a longer time and the collector queue becomes full, the newest data is discarded until a connection to the database is

>> re-established.
- Data queuing is supported for modules that use the collector.
- Regional Operators in each region will log into the Orion Web Console in the Primary Region where you SolarWinds Orion Platform products are installed to view data.
- An optional Additional Web Server is available, and it can be installed in secondary regions. Regional operators can then log into their local web consoles.
- All KPIs, such as Node Response Times, will be calculated from the perspective of each regional Additional Polling Engine. For example, the response time for a monitored node in Region 2 will be equal to the round trip time from the Additional Polling Engine in Region 2 to that node.

Install Additional Polling Engines

Installing and configuring an Additional Polling Engine is identical to installing a primary SolarWinds polling engine, with the following considerations:

- The most recent installer is available in your SolarWinds Customer Portal under My Downloads > View downloads for: Orion Additional Polling Engine.
- To monitor or manage devices polled by an additional polling engine in any Orion platform product, install the Additional Polling Engine for the appropriate product. For more information, see the documentation for your product.
- Up to three total polling engines can be installed on a single server. In this case, individual licenses must be activated for each polling engine.
- If you configured an alert with a Send Email action to trigger on a node monitored by an additional polling engine, confirm that the additional polling engine can access your SMTP server. Otherwise, the emails will not be sent.

To install an Orion Additional Polling Engine:

1. Extract the .zip file you downloaded, and run the executable file.

   > (i)
   > - The extracted folder contains Additional Polling Engine installers for all Orion products that support Additional Polling Engines.
   > - Launch the installer that corresponds to the SolarWinds product installed on your primary SolarWinds Orion server.
   > - If you have multiple Orion products installed on your primary SolarWinds Orion server, install the additional polling engine for each product.

2. On the Welcome window of the Compatibility Check, provide the following information:
   - The host name or IP address of your primary SolarWinds Orion server.
   - The user name and password of a user with administrative privileges to the Orion Web Console on your Main Polling Engine.
3. Click Next to complete the installation the same way as on a primary SolarWinds Orion server.

Activating Stackable Poller Licenses

When installing additional polling engines in a stacked poller installation, licenses must be activated using the Smart Bundle installer.

Choose from the scenarios below:

1. Main Orion Server + Stackable Pollers
   a. Use Full Installer to install main poller.
   b. Use Smart Bundle to install stackable poller(s). The license will be initiated on the first screen.
2. Additional Polling Engine + Stackable Pollers
   a. Use Smart Bundle to install Additional Polling Engine. This will install bits for all needed modules/products (unless something is not included by Smart Bundle).
   b. Use Smart Bundle to install stackable poller(s). The license will be asked on the first screen
3. Add stackable Pollers to existing APE (could be installed by previous Smart Bundle or regular APE installer).
   a. Use Smart Bundle to install any missing or out-of-date modules. If some modules are not included in the Smart Bundle, it will provide instruction for you to download a regular Additional Poller installer and install the missing parts.

Frequently Asked Questions

The following questions address some common issues encountered when using scalability engines with a SolarWinds installation.

Does each module have its own polling engine?

No, any additional polling engine may have all relevant modules installed on it, and it will perform polling for all installed modules. An additional polling engine essentially works in the same way as your primary polling engine on your main server.

If I am monitoring with both NPM and SAM, do I need to install a NPM polling engine and a separate SAM polling engine?

No, any additional polling engine may have all relevant modules installed on it, and it will perform polling for all installed modules. An additional polling engine essentially works in the same way as your primary polling engine on your main server.

Are polling limits cumulative or independent? For example, can a single polling engine poll 12k NPM elements AND 10k SAM monitors together?

Yes, a single polling engine can poll up to the limits of each module installed, providing sufficient hardware resources are available.

Are there different size license available for the Additional Polling Engine?

No, the Additional Polling Engine is only available with an unlimited license.

Can you add an Additional Polling Engine to any size module license?

Yes, you can add an Additional Polling Engine to any size license.

> ⓘ Adding an Additional Polling Engine does not increase your license size. For example, if you are licensed for an NPM SL100, adding an additional polling engine does not increase the licensed limit of 100 nodes/interfaces/volumes, but the polling load is spread across two polling engines instead of one

Will an Additional Polling Engine allow me to monitor overlapping IP's?

Yes, you will be able to add nodes with the same IP Address to separate polling engines allowing you to monitor overlapping IP Addresses.