



Setting Up Reports for Specialized Technologies

The following topics tell you how to set up Prime Performance Manager reports for specialized devices and technologies including Data Center, NetFlow, Star OS Bulk Statistics, Generic CSV Bulk Statistics, OpenStack Ceilometer, Ganglia Reports, Managing Dashboards, Ceph Reports, Custom Report Views, Report Groups and Web Reports:

- [Displaying Data Center Reports, page 8-1](#)
- [Setting Up NetFlow Reports, page 8-17](#)
- [Setting Up StarOS Bulk Statistics Reports, page 8-20](#)
- [Setting Up Generic CSV Bulk Statistics Reports, page 8-29](#)
- [Setting Up Small Cell Reports, page 8-31](#)
- [Setting Up Ganglia Reports, page 8-41](#)
- [Setting Up Cisco Broadband Access Center Reports, page 8-42](#)
- [Ceph and KVM VM Report Notes, page 8-44](#)
- [ONS and CPT Device Report Notes, page 8-45](#)

Displaying Data Center Reports

Prime Performance Manager supports many data center networking, computing, storage, virtualization, and management devices and technologies. Data center report display is covered in the following topics:

- [Supported Data Center Devices and Technologies, page 8-1](#)
- [Displaying ESXi and vCenter Reports, page 8-4](#)
- [Displaying Data Center Tenant Reports, page 8-4](#)
- [Displaying Data Center Resource Allocation and Trend Analysis, page 8-6](#)
- [Setting Up collectd Performance Monitoring For a Single Device, page 8-7](#)
- [Setting Up collectd Performance Monitoring For Multiple Devices, page 8-13](#)

Supported Data Center Devices and Technologies

To display data center reports, from the Performance menu, choose **Views**. Supported data center devices and technologies are displayed in a Data Center view. [Table 8-1](#) lists the data center devices and technologies that Prime Performance Manager supports.

Table 8-1 Supported Data Center Devices and Technologies

Area	Device/Technology	Notes
Network	Citrix NetScaler VPX and SDX	Specifies global and context level LB data; Resources, High Availability (state transitions, failed trans, conf sync failures, heartbeats Tx/Rx).
	Nexus 7000	Specifies new features and performance of Freetown, SUP2E
	GRE	Specifies GRE Tunneling Protocol.
	VM-FEX	
	Cisco ASA 5585	Specifies firewall feature, and four-node clusters
	Cisco ASA 1000v	Specifies FW for tenant edge control.
	Cisco CSR 1000v	Specifies resource, interfaces, MPLS L3VPNs, BGP, IPSEC VPN statistics, LISP, ISIS, FW connections and rate, NAT stats, QoS, NBAR.
	Fabricpath	Specifies L2 multi-path F-tag tree
	OTV	
	LISP	Specifies MR, MS, DB, cache statistics and other details.
	ISIS	Specifies statistics for original ISIS and different variances (for example in OTV).
	vWAAS	
	vNAM	
	IPv6 and Security	
	eBGP	Specifies Cisco CSR XE < eBGP > Cisco ASR 9000
	Cisco ASA 5500	Specifies physical FW and RA VPN termination
	VSM on Nexus 1kv	
	Bare metal UCS servers + Virtual appliances	
	Cisco UCS C without UCSM	
	Cisco ASR 9000-nV	Specifies Cisco ASR 9000 cluster
	Nexus 6000	Specifies Cisco Nexus 6000 series
	Nexus 1100	
Nexus 9000	Specifies Nexus 9300 and 9500/ACI	
SourceFire	Supports NGIPS first and NGFW next	
Compute Logical/Virtual Devices	VMware	Specifies: <ul style="list-style-type: none"> vMotion statistics: vMotion counts, Storage vMotion counts, utilization, top five VMs, allocated space and performances, both host and VM level: IOPS and latency Capacity planning, utilization trend and history, ESXi and VM KPIs: ESXi start time, VM system uptime, VM power usage, VM details, hosts and VM resource report, vCenter clusters and multiple vCenters, statistics at virtual data center and cluster level
	Hypervisor	Specifies host credentials

Table 8-1 Supported Data Center Devices and Technologies (continued)

Area	Device/Technology	Notes
Storage	Storage Area Network (SAN)	Specifies FC and SAN statistics from Cisco MDS 9000 and Nexus 5000
	EMC Storage	Specifies VNX and VMAX devices
	Ceph	Specifies Ceph device details
System	Tenant	Specifies interfaces with Prime Network Service Controller, IAC, and OpenStack
	Support for MSDC and Vinci	
	Dashboard for PC reports	
Application	Monitoring	Specifies CPU, memory, and processes
Integration	Adjust TCA and XL for Prime Network integration	
	Integration with OpenStack	Specifies the Ceilometer services like Nova, Glance, and Swift.
	Integration with IAC/PSC	
	Integration with UCS Director	Specifies integration with Cisco UCS-D, and support storage (NetApp ONTAP 8.2; VMAX and VNX)
Others	Nexus Switches	Collects and pushes data to Prime servers
	DCM	
	OnePK	Uses Java API to collect data and configuration
	REST API scale and performance	Specifies many REST calls from Prime Central.
	VM free space	

In the Prime Performance Manager GUI, data center features are listed under Views and Reports. Under Views, a default Data Center view is provided. It includes:

- Network
 - VDC Devices
 - VSS
 - Routers/Switches
- Compute
 - vCenter
 - ESXi
 - HyperV
 - KVM
 - Xen
 - UCS Clusters (includes the hosts/VMs to blade server mapping)
- Tenants

Under Reports, the following data center report categories are provided:

- Compute

- ESXi
- Ganglia
- Hyper-V
- KVM
- OpenStack
- UCS Clusters
- UCS Standalone
- Xen
- vCenter
- Storage
 - Ceph
 - EMC
 - Fabric Configuration Server
 - Fibre Channel
 - NetApp
 - VSAN Zoning

Some Data Center technologies require special setup procedures, described in the following topics:

- [Displaying ESXi and vCenter Reports, page 8-4](#)
- [Displaying Data Center Tenant Reports, page 8-4](#)
- [Setting Up collectd Performance Monitoring For a Single Device, page 8-7](#)

Displaying ESXi and vCenter Reports

ESXi and vCenter reports have two qualifications:

- Prime Performance Manager supports domain and username to access vCenter and ESXi that are members of an Active Directory domain. For information on adding Telnet and SSH credentials for vCenter and ESXi, see [Adding Device Credentials for Other Protocols, page 5-8](#).
- To display ESXi and vCenter Datastore IOPS, Total Latency and Normalized Latency reports, you must manually enable the Storage I/O Controller option in VMware [datastore] Properties dialog box. (This option is only available in VMware Enterprise Plus.) Enabling Storage I/O Controller allows Prime Performance Manager to poll data from ESXi and vCenter. If the Storage I/O Controller is not enabled, Prime Performance Manager displays 0 in these reports.

Displaying Data Center Tenant Reports

Prime Performance Manager supports multitenant software architectures and technologies including Cisco Network Segmentation Manager (NSM) and OpenStack. NSM integrates VMware vCloud Director 1.6 with the Cisco Nexus 1000V for networking management. OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center.

To display tenant reports, perform the following tasks:

- Import data center devices, such as hypervisors, routers, and switches containing the tenant data. You can import the devices from Cisco Prime Network or Cisco Prime Central, or you can discover the devices from Prime Performance Manager. See [Chapter 5, “Discovering Network Devices.”](#)
- Enable the data center personality in Prime Performance Manager. See [ppm maxhtmlrows, page B-50.](#)
- Import the tenants into Prime Performance Manager. See [Importing Tenants into Prime Performance Manager, page 8-5.](#)
- Display the tenant reports. See [Displaying Tenant Details and Reports, page 8-6.](#)

Importing Tenants into Prime Performance Manager

To import tenants into Prime Performance Manager and display the tenant reports:

Step 1 Verify the data center personality is enabled:

```
ppm manage datacenter status
```

If the status is inactive, complete the following steps:

a. Enable the data center personality:

```
ppm manage datacenter enable
```

b. Restart Prime Performance Manager. See [Restarting Gateways and Units, page 2-5](#)

Step 2 Log into the Prime Performance Manager GUI. See [Launching the Web Interface, page 3-1.](#)

Step 3 From the Administration menu, choose **Tenant Integration**.

Step 4 On the Administration Tenant Integration window, click the tenant type:

- NSM—Cisco Network Segmentation Manager tenant.
- OpenStack—OpenStack tenant.

Step 5 Complete the tenant fields:

- Host Name or IP Address—Enter the tenant host name or IP address.
- Protocol (OpenStack only)—Enter the tenant communications protocol, either HTTP or HTTPS.
- Port—Enter the tenant port number.
- User Name—Enter the tenant administrator username.
- Password—Enter the user password.
- Version—Enter the tenant version. For NSM, this will be 1.0. For OpenStack, this is 2.0.
- Sync Interval—Choose the interval at which you want tenant information in Prime Performance Manager synchronized with the tenant.

Step 6 On the toolbar, click **Import Tenant**.

The Tenant Sync in Progress message appears. The import takes a few minutes to complete. To check on its progress, from the System menu, choose **Messages**. The status is displayed as OpenStack Sync Task and NSM Sync Task messages. Completion is indicated with: OpenStack (or NSM) Inventory Import: Task successfully completed. If the import fails, a dialog appears with the details.

Displaying Tenant Details and Reports

To display tenant details and reports:

-
- Step 1** Verify that the tenants were imported into Prime Performance Manager. See [Importing Tenants into Prime Performance Manager, page 8-5](#).
- Step 2** From the Network menu, choose **Tenants**.
The System Tenants window displays the tenants that are integrated with Prime Performance Manager.
- Step 3** Click a tenant in the navigation tree.
The tenant name, description, type, and last successful import data are displayed on the Tenant Details tab.
- Step 4** To display performance data related to the current tenant, click **Tenant Reports**.
The Tenants Reports tab displays reports specific to the selected tenant.
- Step 5** You can also display tenant reports by choosing **Reports** from the Performance menu, then choosing:
- **Views > Tenants**
Provides a unified structure for each tenant even if tenants were imported from different cloud systems. Tenant performance data is divided into Network and Compute categories. Essentially, Views > Tenants provides a visual map for tenants management.
 - **Reports > Tenant Reports**
Lists all Prime Performance Manager tenant reports. In this release, ESXi, VLAN, VRFs reports are available. Reports list the Top 10 (by default) tenants data.



Note For NSM, only performance data associated with networks, such as VLANs and VRFs are collected. For OpenStack, only performance data associated to VMs is collected.

Displaying Data Center Resource Allocation and Trend Analysis

The data center group report allows you to see a profile of your network data center VMware vCenter servers and perform trend analysis to predict future capacities, availabilities, and other resource allocations. vCenter allocations that you can track and analyze future trends for include:

- Top N hosts by CPU availability (CPU total available, provisioned, consumed).
- Top N hosts by memory availability.
- Top N hosts by storage disk availability (datastores by space available).
- Top VMs with the most alarms (per downtime) availability, including:
 - VM name
 - Severity
 - Alarm count
- Top N hosts with the least resource availability, including:
 - VM name

- CPU used
- Memory used
- Disk usage

To display data center resource allocation:

-
- Step 1** Verify that VMware vCenter devices were added to Prime Performance Manager:
- a. From the Network menu, choose **Credentials Editor**.
 - b. In the device list, vCenter devices are listed with vCenter_HTTPS connection protocol.
 - c. If no vCenter servers are listed, you must add them. See [Prime Performance Manager Device Discovery, page 5-5](#). For the device credential, use vCenter_HTTPS.
- Step 2** From the Administration menu, choose **Group Editor**,
- Step 3** Verify that the following group reports are enabled:
- DatacenterCPU
 - DatacenterDatastore
 - DatacenterMemory
 - VMwareClusterCPU
 - VMwareClusterDatastore
 - VMwareClusterMemory
- These reports are enabled by default. If they are not enabled, check the Enabled checkbox then choose **Enable Selected Groups** from the Actions menu.
- Step 4** To view vCenter resource allocation reports, in the navigation tree, choose: **Grouped Reports > Compute > vCenter > Hosts and Clusters > VMware Clusters**.
- Step 5** To perform Data Center resource trend analysis:
- a. In the navigation tree, choose **Trend Analysis > Data Center** located at the bottom of the Performance Reports navigation tree.
Wait until enough data is aggregated in the group reports. A minimum of 12 data points are needed in these reports.
 - b. Experiment with the Trend Analysis to see how the analysis works. Configure different parameters for example. set Interval as Day or Hour, and change the Sample Period (time period on which the analysis is based) and the Forecast Period.
 - c. After specifying parameters, click **Calculate** to recalculate the trend of resource utilization.
-

Setting Up collectd Performance Monitoring For a Single Device

collectd is a computer background process that collects system performance statistics. These statistics can be used for performance analysis and capacity planning.

To monitor collectd statistics in Prime Performance Manager for a single device:

-
- Step 1** Verify that the following are installed on the computer where you want to gather collectd statistics:
- Net-SNMP
 - RRDTool
 - Collectd
- Step 2** Add the rrdtool binary directory into the environment variable PATH. For example, if the default device shell is BASH, add the following line into ~/.bash_profile:
- ```
PATH=$PATH: /opt/rrdtool-1.6.7/bin
```
- Step 3** Enable the following section in the collectd configuration file.
- ```
LoadPlugin rrdtool
<Plugin rrdtool>
    DataDir "/var/lib/collectd"
    CacheTimeout 120
    CacheFlush 900
    RRARows 1
    RRA Timespan 3600
</Plugin>
```
- The DataDir parameter must be configured with /var/lib/collectd by default. If you want to change the default rrdtool data output directory, change DataDir here and then modify the XMP_PAL.properties under \$PPM_INSTALLATION/properties/. You must modify this file both on gateway and unit, then restart them.
- ```
collectd base directory to store the rrd files
COLLECTD_BASE_DIR = /var/lib/collectd/
```
- Step 4** If you are monitoring an application, configure the appropriate collectd plugin. See the following topics for configuration examples:
- [Apache Plugin Example, page 8-8](#)
  - [MySQL Plugin Example, page 8-9](#)
  - [GenericJMX Plugin Example, page 8-9](#)
  - [Oracle Plugin Example, page 8-11](#)
  - [PostgreSQL Plugin Example, page 8-12](#)
  - [Ceph Plugin Example, page 8-13](#)
- Step 5** Add the SNMP community configured on Net-SNMP in Prime Performance Manager. For information see [Adding SNMP Device Credentials, page 5-6](#).
- Step 6** Add collectd\_SSH in the Credentials Editor, The credential is the operating system SSH username and password. For information see [Adding Device Credentials for Other Protocols, page 5-8](#).
- Step 7** Run device discovery. For information, see [Prime Performance Manager Device Discovery, page 5-5](#)
- 

## Apache Plugin Example

The following shows a sample Apache plugin configuration:

```
LoadPlugin apache
<Plugin apache>
```



```

<Instance "local">
 URL "http://localhost/server-status?auto"
</Instance>
</Plugin>

```

## MySQL Plugin Example

The following shows a sample MySQL plugin configuration:

```

LoadPlugin mysql
<Plugin mysql>
<Database example>
 Host "localhost"
 Socket "/var/lib/mysql/mysql.sock"
 User "root"
 Password ""
 Database "example"
</Database>
</Plugin>

```

## GenericJMX Plugin Example

The GenericJMX collectd plugin reads Managed Beans (MBeans) from an MBeanServer using JMX. The plugin is written in Java and requires the Java plugin.

```

<Plugin "java">
 JVMARG
 "-Djava.class.path=/usr/share/collectd/java/collectd-api.jar:/usr/share/collectd/java/generic-jmx.jar"
 LoadPlugin "org.collectd.java.GenericJMX"

 <Plugin "GenericJMX">
 <MBean>
 ...
 </MBean>

 <Connection>
 Host "crdc-c210-3-6"
 InstancePrefix "ppm_unit-"
 ServiceURL "service:jmx:rmi:///jndi/rmi://10.74.125.84:9011/jmxrmi"
 Collect "memory_pool"
 Collect "classes"
 Collect "compilation"
 Collect "memory"
 Collect "garbage_collector"
 </Connection>
 </Connection>
 ...
</Connection>

</Plugin> </Plugin>
</Plugin>

```

Run your applications with something similar to the following arguments on the computer monitored by GenericJMX. In this example, 10.74.125.84 is used:

```

-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.port=17264
-Dcom.sun.management.jmxremote.ssl=false

```

The MBean configuration includes the following:

- Memory pools

Memory usage by memory pool:

```
<MBean "memory_pool">
 ObjectName "java.lang:type=MemoryPool,*"
 InstancePrefix "memory_pool-"
 InstanceFrom "name"
 <Value>
 Type "memory"
 #InstancePrefix ""
 #InstanceFrom ""
 Table true
 Attribute "Usage"
 </Value>
</MBean>
```

- Memory

Generic heap and nonheap memory usage:

```
<MBean "memory">
 ObjectName "java.lang:type=Memory"
 #InstanceFrom ""
 InstancePrefix "memory"
 # Creates four values: committed, init, max, used
 <Value>
 Type "memory"
 #InstancePrefix ""
 #InstanceFrom ""
 Table true
 Attribute "HeapMemoryUsage"
 InstancePrefix "heap-"
 </Value>
 # Creates four values: committed, init, max, used
 <Value>
 Type "memory"
 #InstancePrefix ""
 #InstanceFrom ""
 Table true
 Attribute "NonHeapMemoryUsage"
 InstancePrefix "nonheap-"
 </Value>
</MBean>
```

- Classes

Number of classes loaded by the JVM:

```
<MBean "classes">
 ObjectName "java.lang:type=ClassLoading"
 InstancePrefix "classes"
 #InstanceFrom ""
 <Value>
 Type "gauge"
 InstancePrefix "loaded_classes"
 #InstanceFrom ""
 Table false
 Attribute "LoadedClassCount"
 </Value>
</MBean>
```

- **Compilation**

Time spent by the JVM compiling or optimizing:

```
<MBean "compilation">
 ObjectName "java.lang:type=Compilation"
 InstancePrefix "compilation"
 #InstanceFrom ""
 <Value>
 Type "total_time_in_ms"
 InstancePrefix "compilation_time"
 #InstanceFrom ""
 Table false
 Attribute "TotalCompilationTime"
 </Value>
</MBean>
```

- **Garbage Collector**

The garbage collector account and time is shown below:

```
<MBean "garbage_collector">
 ObjectName "java.lang:type=GarbageCollector,*"
 InstancePrefix "gc-"
 InstanceFrom "name"

 <Value>
 Type "invocations"
 #InstancePrefix ""
 #InstanceFrom ""
 Table false
 Attribute "CollectionCount"
 </Value>

 <Value>
 Type "total_time_in_ms"
 InstancePrefix "collection_time"
 #InstanceFrom ""
 Table false
 Attribute "CollectionTime"
 </Value>

</MBean>
```

## Oracle Plugin Example

A generic Oracle plugin configuration is shown below:

```
loadPlugin oracle
<Plugin oracle>
 <Query "db_efficiency">
 Statement "SELECT round(sum(decode(METRIC_NAME, 'Database Wait Time Ratio',
value)),2) AS DATABASE_WAIT_TIME_RATIO,
round(sum(decode(METRIC_NAME, 'Database CPU Time Ratio', value)),2) AS
DATABASE_CPU_TIME_RATIO,
'DB_EFFICIENCY' AS DB_EFFICIENCY
FROM SYS.V_$SYSMETRIC
WHERE METRIC_NAME IN ('Database CPU Time Ratio', 'Database Wait Time Ratio')
AND INTSIZE_CSEC = (SELECT max(INTSIZE_CSEC) FROM SYS.V_$SYSMETRIC)"
 <Result>
 Type "efficiency"
 InstancesFrom "DB_EFFICIENCY"
 ValuesFrom "DATABASE_WAIT_TIME_RATIO" "DATABASE_CPU_TIME_RATIO"
 </Result>
```

```

</Query>

<Query "io_per_tablespace">
Statement "SELECT sum(vf.PHYBLKRD)*8192 AS PHY_BLK_R,
 sum(vf.PHYBLKWRT)*8192 AS PHY_BLK_W,
 'tablespace' AS i_prefix,
 dt.tablespace_name
FROM ((dba_data_files dd JOIN v$filestat vf ON dd.file_id = vf.file#
)
 JOIN dba_tablespaces dt ON dd.tablespace_name =
dt.tablespace_name)
 GROUP BY dt.tablespace_name"
<Result>
Type "io_octets"
InstancesFrom "i_prefix" "TABLESPACE_NAME"
ValuesFrom "PHY_BLK_R" "PHY_BLK_W"
</Result>
</Query>

<Database "cisco">
ConnectID "cisco"
Host "10.74.125.76"
Username "cisco"
Password "cisco"
Query "db_efficiency"
Query "io_per_tablespace"
</Database>
</Plugin>

```

If remote mode is enabled add the following dataset specification to `$COLLECTD_HOME/share/collectd/types.db` at both the remote collectd server and the client:

```
efficiency wait:GAUGE:0:100.1, cpu:GAUGE:0:100.1
```

Database blocks define database connections and the queries that should be sent to the database:

- **ConnectID**—Defines the database alias or service name to connect to. Usually, these names are defined in the file named `$ORACLE_HOME/network/admin/tnsnames.ora`.
- **Host**—Hostname to use when dispatching values for this database. Defaults to using the global hostname of the collectd instance.
- **Username**—Username used for authentication.
- **Password**—Password used for authentication.
- **Query**—Associates the query named `QueryName` with this database connection. The query needs to be defined before this statement, that is, all query blocks you want to refer to must be placed above the database block you want to refer to them from.

## PostgreSQL Plugin Example

Query blocks are not used in the PostgreSQL plugin, The plugin uses the default statistics collected from the PostgreSQL statistics collector, which you must enable. The collector is usually enabled by default.

```

loadPlugin postgresql
<Plugin postgresql>
 Database test
 Host "localhost"
 Port "5432"
 User "postgres"
 Password "postgres"
</Database>

```

```

 <Database foo>
 ...
 <Database>
</Plugin>

```

## Ceph Plugin Example

You can base your Ceph plugin configuration on the following example:

```

<LoadPlugin ceph>
 <Daemon "osd.0">
 SocketPath "/var/run/ceph/ceph-osd.0.asok"
 </Daemon>
 <Daemon "mon.a">
 SocketPath "/var/run/ceph/ceph-mon.ceph1.asok"
 </Daemon>
 <Daemon "mds.a">
 SocketPath "/var/run/ceph/ceph-mds.ceph1.asok"
 </Daemon>
</Plugin>

```

Notes:

- The name must start with osd if the SocketPath points to an osd socket file. mon and mds can be done in the same manner.
- The Ceph plugin does not support remote mode, so you cannot use network plugin to do centralized deployment.

## Setting Up collectd Performance Monitoring For Multiple Devices

If you're collecting performance data for multiple devices, data should be in one central location and not across multiple servers. The recommended approach is to designate one server and multiple clients that send their data to the server. For information about collectd, visit the collectd website:

[https://collectd.org/wiki/index.php/Networking\\_introduction](https://collectd.org/wiki/index.php/Networking_introduction)

## Managing Hypervisors on Windows Servers

You can use Prime Performance Manager to display reports on Xen, KVM, and Hyper-V hypervisors for the following Windows operating systems:

- Windows Server 2012 R2
- Windows Server 2012 ST
- Windows Server 2008 R2
- Windows Server 2008 ST

The following topics provide instructions for each hypervisor:

- [Managing Windows Server VMs in KVM, page 8-14](#)
- [Managing Windows Server VMs in Xen, page 8-14](#)
- [Managing Windows Server VMs in Hyper-V, page 8-15](#)
- [Manage Windows Server VMs Using SNMP Credentials, page 8-16](#)

## Managing Windows Server VMs in KVM

You can manage Windows Server VMs in KVM through KVM\_TLS or SNMP credentials.

To manage Windows Server VMs in KVM using KVM\_TLS credentials:

- 
- Step 1** Log into Prime Performance Manager as the system administrator user. For login procedures, see [Launching the Web Interface, page 3-1](#).
- Step 2** From the Network menu, choose **Credentials Editor**.
- Step 3** On the Credentials Editor toolbar, click **Add New Credentials Entry**.
- Step 4** In the Add Credentials Entry dialog box, enter the following:
- Device—Enter the KVM device IP address.
  - Connection Protocol—Enter **KVM\_TLS**.
  - Sub System—Enter **TBD**.
  - User Name—If necessary, enter the KVM device username.
  - Password—If necessary, enter the username password.




---

**Note** Secondary Login Type, Secondary Username, and Secondary Password are not applicable

---

- Step 5** Click **OK**.
- Step 6** From the Network menu, choose **Discovery**,
- Step 7** In the IP Address, Address Range, Subnet, CIDR, or DNS Hostname field, enter the KVM device IP address.
- Step 8** Click **Add**.
- Step 9** From the Network Discovery toolbar, click **Discover Network**.
- Step 10** From the Network menu, choose **Device**,
- Step 11** Verify the KVM device is added to the device list with Device type QEMU/KVM 1.5.0.
- Step 12** Click the KVM device.
- Step 13** Under the device reports, navigate to **Compute > KVM** to view KVM Host and KVM VM reports.




---

**Note** You can also view KVM Host and Windows VM reports by choosing **View > Data Center > Compute > KVM**.

---

## Managing Windows Server VMs in Xen

You can manage Windows Server VMs in Xen through XEN\_TLS or SNMP credentials.

To manage Windows Server VMs in Xen using XEN\_TLS credentials and check host and VM reports:

- 
- Step 1** Log into Prime Performance Manager as the system administrator user. For login , see [Launching the Web Interface, page 3-1](#).

- Step 2** From the Network menu, choose **Credentials Editor**.
- Step 3** On the Credentials Editor toolbar, click **Add New Credentials Entry**.
- Step 4** In the Add Credentials Entry dialog box, enter the following:
- Device—Enter the Xen device IP address.
  - Connection Protocol—Enter **XEN\_TLS**.
  - Sub System—Enter **TBD**
  - User Name—If necessary, enter the Xen device username.
  - Password—If necessary, enter the username password.




---

**Note** Secondary Login Type, Secondary Username, and Secondary Password are not applicable

---

- Step 5** Click **OK**.
- Step 6** From the Network menu, choose **Discovery**,
- Step 7** In the IP Address, Address Range, Subnet, CIDR, or DNS Hostname field, enter the Xen device IP address.
- Step 8** Click **Add**.
- Step 9** From the Network Discovery toolbar, click **Discover Network**.
- Step 10** From the Network menu, choose **Devices**,
- Step 11** Verify the Xen device is added to the device list with Device type Xen 3.1.
- Step 12** Click the Xen device.
- Step 13** Under the device reports, navigate to **Compute > XEN** to view Xen Host and Xen VM reports.




---

**Note** You can also view Xen host and Windows VM reports by choosing **View > Data Center > Compute > XEN**.

---

## Managing Windows Server VMs in Hyper-V

You can manage Windows Server VMs in Hyper-V through XEN\_TLS or SNMP credentials.

To manage Windows Server VMs in Hyper-V using XEN\_TLS credentials and check host and VM reports:

- Step 1** Log into Prime Performance Manager as the system administrator user. For login, see [Launching the Web Interface, page 3-1](#).
- Step 2** From the Network menu, choose **Credentials Editor**.
- Step 3** On the Credentials Editor toolbar, click **Add New Credentials Entry**.
- Step 4** In the Add Credentials Entry dialog box, enter the following:
- Device—Enter the Hyper-V device IP address.
  - Connection Protocol—Enter **XEN\_TLS**.

- Sub System—Enter TBD
- User Name—If necessary, enter the Hyper-V device username.
- Password—If necessary, enter the username password.




---

**Note** Secondary Login Type, Secondary Username, and Secondary Password are not applicable

---

- Step 5** Click **OK**.
- Step 6** From the Network menu, choose **Discovery**,
- Step 7** In the IP Address, Address Range, Subnet, CIDR, or DNS Hostname field, enter the Hyper-V device IP address.
- Step 8** Click **Add**.
- Step 9** From the Network Discovery toolbar, click **Discover Network**.
- Step 10** From the Network menu, choose **Device**,
- Step 11** Verify the Hyper-V device is added to the device list with Device type HyperV6.2.0.
- Step 12** Click the Hyper-V device.
- Step 13** Under the device reports, navigate to **Compute > Hyper-V** to view Hyper-V Host and Hyper-V VM reports.




---

**Note** You can also view Hyper-V Host and Windows VM reports by choosing **View > Data Center > Compute > HyperV**.

---

## Manage Windows Server VMs Using SNMP Credentials

To manage Windows Server VMs using SNMP credentials:

- 
- Step 1** Enable SNMP Service on the Windows VM Servers.
- Step 2** Log into Prime Performance Manager as the system administrator user. For information, see [Launching the Web Interface, page 3-1](#).
- Step 3** From the Network menu, choose **SNMP Editor**.
- Step 4** On the Network SNMP Editor toolbar, click **Add New SNMP Entry**.
- Step 5** In the Add SNMP Entry dialog box, enter the following parameters:
- IP Address Range or Hostname—Enter the Windows VM server IP address.
  - SNMP Version
  - Read Community
  - Max Table Varbind
  - Port
  - User Name
  - Authentication Protocol



- Authentication Password
  - Privacy Protocol
  - Privacy Password
- Step 6** Click **OK**.
- Step 7** From the Network menu, choose **Discovery**,
- Step 8** In the IP Address, Address Range, Subnet, CIDR, or DNS Hostname field, enter the Windows server IP address.
- Step 9** Click **Add**.
- Step 10** From the Network Discovery toolbar, click **Discover Network**.
- Step 11** From the Network menu, choose **Device**,
- Step 12** Verify the Windows VM is added to the list with the device type, WindowsNT Server.
- Step 13** Click the device and view the Windows VM reports.
- 

## Setting Up NetFlow Reports

Prime Performance Manager can generate NetFlow data reports. NetFlow is a push-model Cisco network protocol that collects IP traffic information. Devices are configured to export IP packets to the NetFlow collector on the Prime Performance Manager unit. The unit processes the packets and generates report data. Prime Performance Manager supports:

- NetFlow v5 and v9
- Flexible NetFlow IP Flow Information Export (IPFIX)
- NetFlow records exported using User Datagram Protocol (UDP).

When configuring NetFlow for reports, keep the following in mind:

- To support UDP, configure the NetFlow collector IP address and the UDP destination port on the sending router. The NetFlow collector IP address is the unit server IP Address. The destination UDP port needs an available unit port. All configured devices must send NetFlow packets to the same unit UDP port.
- To change the default UDP port on the unit, use the `ppm netflowport` command. See [ppm netflowport](#), page B-57.
- To accurately calculate statistics, set the device active cache timeout to one minute: **ip flow-cache timeout active 1**.

### NBAR2 Support

Network-Based Application Recognition 2 (NBAR2), or Next Generation NBAR, is a revised Cisco NBAR architecture. NBAR is a method through which Cisco routers and switches identify data flows to determine the flow traffic category.

Prime Performance Manager provides the following NBAR2 application reports:

NetFlow > NetFlow Applications:

- Conversations
- Destinations

- Interfaces
- Source Destinations
- Sources
- ToS

An Application column is added to the table view of existing NetFlow All Flows reports. When defining the flow record on the managed device, verify the application name matches as one flow record key, for example:

```
cisco_281(config-flow-record)#match application name
```

This statement allows NBAR data to be included in the NetFlow records.

### Flow Start and End Times

Prime Performance Manager uses the time values sent in the packet header and in the flow to calculate the exact flow start and end times. The flow start and end times help Prime Performance Manager place the flow in the right reporting interval.

For NetFlow Version 1-9, the flow sysUpTime and UTC seconds headers and flowStartSysUpTime and flowEndSysUpTime are used for calculation.

For IPFIX Version 10, several field combinations can be sent in the flow. Prime Performance Manager uses these fields and the UTC seconds header to calculate the flow start and end times.

Valid IPFIX combinations include:

- systemInitTimeMilliseconds, flowStartSysUpTime, flowEndSysUpTime



#### Note

If start and end uptime fields are sent in the flow, the systemInitTimeMilliseconds field is also required to calculate the reporting interval.

- flowStartSeconds, flowEndSeconds
- flowStartMilliseconds, flowEndMilliseconds
- flowStartMicroseconds, flowEndMicroseconds
- flowStartNanoseconds, flowEndNanoseconds
- flowStartDeltaMicroseconds, flowEndDeltaMicroseconds

### Missed Flow Sequence Numbers Report (MFSNs)

The NetFlow Metrics Missing Flow Sequence Numbers (MFSNs) report captures the NetFlow header Sequence Number fields in the NetFlow Header. The Sequence Number value has a different meaning, depending on the NetFlow version:

- Version 5—The total flows sent by the device. For example, if the number of flows is 20, the sequence increments by 20 with every packet sent from the device. Reports for this version help identify flows sent by the device but not received by the collector.
- Version 9—A running sequence of all export packet sent by the device, starting with 1 and incrementing by 1 when each packet is sent from the device. Reports for this version help identify missing packets sent by the device but not received by the collector.
- Version 10—The total flows sent by the device. For example, if the number of flows is 20, the sequence increments by 20 with every packet sent from the device. Reports for this version help identify total flows sent by the device.

To set up Prime Performance Manager for NetFlow reports:

- 
- Step 1** Log into the Prime Performance Manager GUI. See [Launching the Web Interface, page 3-1](#).
- Step 2** From the Network menu, choose **Devices**.
- Step 3** Click the link of the first device configured for NetFlow.
- Step 4** On the device window, click **Data Collection**.
- Step 5** In the Collector Status area, the NetFlow item displays one of the following statuses:
- Active—The device is configured to export NetFlow and the collector is receiving the flows regularly.
  - Not Active—The device is configured for NetFlow but it might not be receiving flows recently.
  - Not Configured—The device is not configured for NetFlow export.
- Step 6** Repeat Steps 2 through 5 to identify all the NetFlow devices.
- Step 7** Enable the NetFlow reports:
- a. From the Performance menu, choose **Reports**.
  - b. Click **Report/Group Status**.
  - c. Navigate to the NetFlow reports and enable the ones you want to see.



---

**Note** Some NetFlow reports display detailed NetFlow stream information. These reports are available only for the lowest enabled interval and are not aggregated to higher intervals. These reports can be viewed only from the device level, for example, the NetFlow All Flows reports.

---

## Setting Up NetFlow Reports For IP Addresses

You can create NetFlow reports for specified IP addresses or address ranges assigned to logical entities that you want to monitor, for example, a customer, building, department, or other logical entity for which you want to monitor data.

To set up NetFlow reports for IP addresses:

- 
- Step 1** Log into Prime Performance Manager using the CLI.
- Step 2** Navigate to the IP group definition file:
- ```
/opt/CSCOppm-gw/etc/IPGroupSchema
```
- Step 3** Open **IPGroupSchema** with a text editor.
- Step 4** Add the IP address(es) and/or IP address range(s):
- You can list addresses separately or as IP ranges.
 - Each line represents one unique IP group definition. If two lines are entered for the same group, the second definition overrides the first definition.
 - Use commas to separate IP addresses (or IP ranges).

- List IPv6 addresses individually and not in a range format. IPv4 addresses can be listed in a range format.
- IPv4 group definition examples:
 - groupA = 10.10.10.10,192.168.0.3,10.10.11-13.5
 - groupB = 20.20.20.10-20, 20.20.21.*
- IPv6 group definition examples:
 - groupV6A = 2001::cafe:1,2001::cafe:2,2001::cafe:3
 - groupV6B = 2012:20bf:30cf:40df:50ef:60ff:beef:1,2013:0:0:0:32e4:dbff:fe32:f4c0

Step 5 Configure NetFlow on the device and verify that NetFlow records are exported successfully.

Step 6 To view IP address reports:

- Log into the Prime Performance Manager GUI.
- From the Performance menu, choose **Reports**.
- In the report navigation tree, click **Reports > NetFlow > NetFlow IP Group**.

The IP group reports are divided by source and destination.

Setting Up NetFlow Reports For Top XX Entries

By default, Prime Performance Manager displays the top 10 report items for any report. This number can be changed at the system and individual user level. For example, the system setting could be 25, which allows individual users to change their Top XX setting to any number below 25.

To implement this feature for NetFlow reports, add a seriesLimit tag with the value MAX_CHART_SERIES_NETFLOW into the GraphSummary section of the NetFlow webreport.

```
<GraphView>
  <GraphSummary title="gstNetFlowSrcASAggStats"
seriesLimit="MAX_CHART_SERIES_NETFLOW" />
  <Graph title="gtNetFlowSrcASAggBytes" showTotal="true">
    <Column name="bytes" showPercentageColumn="true">inBytes</Column>
  </Graph>
```

For information about editing Prime Performance Manager reports, see [Creating Web Reports, page 7-54](#) or the [Cisco Prime Performance Manager 1.6 Integration Developer Guide](#).

Setting Up StarOS Bulk Statistics Reports

Prime Performance Manager retrieves report data for most devices using SNMP to poll the device MIBs containing the performance data. Some devices, such as the Cisco ASR 5000 and Cisco ASR 5500, provide less SNMP support and few MIBs, so few statistics can be gathered using SNMP. However, you can generate reports for these devices using bulk statistics. Bulk statistics are collected in a groups, called schemas, at regular intervals. The device sends the schemas to a specified location as comma separated value (CSV) files. Prime Performance Manager retrieves the files and generates the reports. Each schema contains many performance data variables called counters.

The Cisco ASR 5000 and ASR 5500 devices can be configured to collect bulk statistics and FTP them to a collection server (remote folder). To generate reports from the bulk statistics, you must configure the device to generate the bulk statistics in the specific format expected by Prime Performance Manager and set up Prime Performance Manager so it can read the CSV files generated by the device.

To set up Prime Performance Manager for bulk statistics:

-
- Step 1** Following instructions in the device documentation, configure devices to FTP bulk statistics files to either the Prime Performance Manager unit server or a SAN directory.
- Step 2** Complete the [“Creating the StarOS Device Bulk Statistics Configuration” procedure on page 8-23](#) to generate the device configuration. The procedure generates the bulk statistics in the format Prime Performance Manager expects.
- Step 3** Copy the generated configuration to all the required devices. Prime Performance Manager does not configure the devices automatically, so you must copy the configuration file manually.



Note You do not need to configure all the schema types generated by the command. However the reports will have data only for the schemas that are configured.



Note If the drop directory is a SAN directory, the Prime Performance Manager unit must have read and write permissions to it.

- Step 4** Verify the bulk statistics samples are received at the drop directory.
- The drop location for the files is configured in the device config using the CLI remotefile format. Verify the devices actually FTP the bulk statistics CSV file to this location. If files are not received at the drop directory, verify the directory has the proper permissions and the login user name and password provided in the device configuration is valid.



Tip Perform a save configuration from the device to the drop directory. If save configuration works successfully, the bulk statistics samples should FTP without issues.

- Step 5** Open the bulk statistics samples and verify the counter names are complete.

The example schema definitions below show incomplete counters at the end.

```
PPM,system,systemSch55,1349867400,20121010,071000,1702713,0,,%d
PPM,system,systemSch56,1349867400,20121010,071000,1702713,,0,%disc-reason
PPM,system,systemSch6,1349867400,20121010,071000,1702713,,,,,0,606734,%sess-s
```



Tip Incomplete counters in the bulk statistics samples often occur when you copy too many configuration lines to the device at the same time. Devices have a buffer limit on the number characters that can be pasted to it at one time. If you experience incomplete counters, configure a smaller number of lines at a time.

Bulk statistics samples with complete counter names enclosed within % symbols are acceptable. These are counters that are obsolete and not supported in the current Star OS version. Prime Performance Manager ignores these counters and uses the default data type value for processing. For example, %cpu3-cpuused-user% is an obsolete or unsupported counter in the following:

```
PPM,card,cardSch3,1346768700,20120904,142500,928325,1,0,0,0,%cpu3-cpuused-user%,0.00,0.00,
0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
```

Step 6 Set the drop location on the unit server, by entering:

```
/opt/CSCOppm-unit/bin/ppm csvdrop [dir]
```

When you install the unit, the default drop directory location is `/opt/CSCOppm-unit/csvdrop/`. Use the `ppm csvdrop` command to point to the directory where bulk statistics samples are received from the Cisco ASR 5000 and Cisco ASR 5500 devices. The command updates the `CSV_DROP_DIR` property in `BulkStats.properties` located in `/opt/CSCOppm-unit/properties`. Restarting the unit is not necessary; the property changes takes effect automatically.

Use the same drop directory for all Cisco ASR 5000 and Cisco ASR 5500 devices in the network. The `%host%` variable in the filename helps to uniquely identify each device.

Step 7 Discover devices and check for alarms in the Alarms/Events window. For information about device discovery, see [Chapter 5, “Discovering Network Devices.”](#) For information about the Prime Performance Manager alarms management, see [Chapter 10, “Managing Network Alarms and Events.”](#)

Prime Performance Manager might raise bulk statistics alarms. See [Bulk Statistics Alarms, page 8-25](#) for alarm descriptions and actions to resolve them. The unit console and message log files in `/opt/CSCOppm-unit/logs/` are also be a good place to view possible errors.

Step 8 Verify the Prime Performance Manager CSV Bulk Stats status for each device:

- a. From the Network menu, choose **Devices**.
- b. Click the link for each device configured for bulk statistics.
- c. On the device window, click **Data Collection**.
- d. In the Collector Status area, the CSV Bulk Stats item displays one of the following statuses:
 - Active—The device is configured for bulk statistics and the unit is receiving the files regularly.
 - Not Active—The device is configured for bulk statistics but it might not be discovered or some files are skipped.
 - Not Configured—The device is not configured properly or bulk statistics are not supported for the device type.

Prime Performance Manager processes new files in the drop directory when they arrive. If the file name format is valid, Prime Performance Manager processes the new files for enabled Star OS reports.

If parameters are missing in the file name or the headers and footers in the bulk statistics samples do not match the property file values, marker files are created in the drop directory. The files have zero size and have the same name as the file in error but are appended with an error extension.

If multiple sample files are collected in the drop directory, Prime Performance Manager processes from the oldest to latest sample. Prime Performance Manager processes about 50 files at a time until it catches up with the recent files. Daily statistics can be looked up for older time range if necessary.



Note To process more than 50 files, edit the `BulkStats.properties` file and change `MAX_FILE_COUNT`. Increasing the number might cause performance issues depending on the number of devices discovered on the unit.

- Step 9** Verify the device reports by choosing **Reports** from the Performance menu and scrolling to the Mobile StarOS Statistics reports in the navigation tree. If the Star OS reports lack data, complete the following steps:
- Verify that the report is enabled at the network and device level. To check the network level, from the Administration menu, choose **Report Status**. To check at the device level, select a device and click the **Report Status** tab.
 - Verify the particular schema is configured on the device. Prime Performance Manager server might also be processing a backlog of files in the drop directory, check for data for longer durations.
 - Verify the device StarOStm software version. The StarOS version for Cisco ASR 5000 and ASR 5500 devices is taken from the bulk statistics files. You can view the version by displaying the device, then clicking the **Details** tab. If the version is Unknown, verify if you have the system schema 71 configured on the device. This is the schema line containing the version information. The schema systemSch71 format is:

```
PPM,system,systemSch71,%epochtime%,%localdate%,%localtime%,%uptime%,%disc-reason-490%,
%disc-reason-491%,%disc-reason-492%,%disc-reason-summary%,,%swversion%,%peak-cpuusa
ge%,%peak-memusage%,,%system-capacity-usage%,%session-capacity%,%session-capacity-usag
e%,%npu-capacity%,%npu-capacity-usage%,%sess-max-lastreset-time%,,%
PPM,system,systemSch71,1349964000,20121011,100000,3595,0,0,0,,,,,14.0,15.61,11291412.
00,,0.06,10137600,6095,1207959552,608208,Never,,,,
```

Creating the StarOS Device Bulk Statistics Configuration

The ppm starbuild command builds the StarOS bulk statistics device configuration with the desired schemas and counters. Prime Performance Manager requires bulk statistics files to be in a very specific format. This command generates the device configuration in the format that Prime Performance Manager expects. The command generates the configuration that you must copy to the Cisco ASR 5000 and Cisco ASR 5500 devices. The command format is:

```
/opt/CSCOppm-gw/bin/ppm starbuild <schemafilename>
```

The input, *schemafilename*, is the CSV file with the bulk statistics schema and counters that need to be configured on the devices. Prime Performance Manager includes a full StarOS schema file. It is located in the install directory:

```
/opt/CSCOppm-gw/install/ASR5K_BulkStats_StarOS_Schema_Counters.csv
```

Prime Performance Manager expects the filenames to be in the following format:

```
%host%_bulkstats_%localdate%_%localtime%_%localtz%_5_5.csv
```

%host%, *%localdate%*, *%localtime%*, *%localtz%* are common bulk statistics counter variables.

Prime Performance Manager uses the *%host%* variable to identify the device. This is usually the device sysName. The "_5_5" at the end of the file name is the sample and transfer interval set on the device.

To build the StarOS device configuration:

Step 1 Log into the gateway as an administrator user.

Step 2 Enter the ppm starbuild command using the provided StarOS schema file for the *schemafilename*:

```
/opt/CSCOppm-gw/bin/ppm starbuild
/opt/CSCOppm-gw/install/ASR5K_BulkStats_StarOS_Schema_Counters.csv
```

For information, see [ppm starbuild](#), page B-91.

- Step 3** Enter the IP address of the Prime Performance Manager unit where the Cisco ASR 5000 and Cisco ASR 5500 devices are discovered, or the SAN folder location IP address:

Enter the IP Address of Prime Performance Manager unit To Send Files To:

- Step 4** Enter the path to the file directory where you want the device to drop the CSV files. This is usually a folder mounted on SAN. The bulk statistics files are retained for 14 days by default in this directory. The full path to the folder is provided here.

Enter File Directory On Prime Performance Manager unit To Drop Files To:



Note If you need to change the default 14 days bulk statistics age use the ppm bulkstatsage command. For information, see [ppm bulkstatsage](#), page B-18.

- Step 5** Enter the output file name. This can be any valid filename. The output configuration file is created in /opt/CSCOppm-gw/bin/ folder by default.

Enter Output Filename To Write StarOS onfig To: staros-bulkstats-config.txt



Note You can press **Enter** to each prompt and change the values by editing the generated output file later.

Removing StarOS Bulk Statistics Device Configurations

To remove bulk statistics configurations from a device on a per schema basis:

- Step 1** Log into the gateway as an administrator user.

- Step 2** Enter the following command:

```
/opt/CSCOppm-gw/bin/ppm starbuild
/opt/CSCOppm-gw/install/ASR5K_BulkStats_StarOS_Schema_Counters.csv -no
```

The output configuration file is created in /opt/CSCOppm-gw/bin/ folder by default. The file contains the delete CLIs for all the available schemas.

Adding New StarOS Bulk Statistics Schemas or Counters

If a configuration must be generated for a smaller subset of schemas or if new schemas need to be added, modify the spreadsheet included in the install directory.

```
/opt/CSCOppm-gw/install/ASR5K_BulkStats_StarOS_Schema_Counters.xlsx
```

After the changes are complete, update the CSV file in the install folder by saving the spreadsheet in CSV format and executing the ppm starbuild command to generate the revised configuration.

**Note**

If you must remove counters within a schema, mark the Export column as “no” in the spreadsheet. This generates the configuration with consecutive commas (,,,) to maintain the counter position in the schema.

Updating the Prime Performance Manager StarOS Bulk Statistics Schema File

To add a new schema or counter to an existing Prime Performance Manager report or to a new report, update the gateway schema file. Prime Performance Manager looks up this schema file to identify the format in which the devices are configured.

**Note**

You only need to generate the Prime Performance Manager schema file when new counters must be used in reports.

To generate the schema file for Prime Performance Manager:

Step 1 Log into the gateway as the administrator user.

Step 2 Enter the ppm starbuild command with the -ppm option:

```
/opt/CSCOppm-gw/bin/ppm starbuild
/opt/CSCOppm-gw/install/ASR5K_BulkStats_StarOS_Schema_Counters.csv -ppm
```

Step 3 At the prompt, enter the output file name:

```
Enter Output Filename To Write PPM Schema: /opt/CSCOppm-gw/etc/bulkstatsschema.csv
PPM Schema File Written To: /opt/CSCOppm-gw/etc/bulkstatsschema.csv
```

Step 4 Reload the new schema file:

```
/opt/CSCOppm-gw/bin/ppm reloadbulkstats
```

Bulk Statistics Alarms

Bulk statistics alarms you might see are listed below. You can view BulkStats.properties in /opt/CSCOppm-unit/properties to see values shown in the alarm generation.

- BulkStatsInfo Alarm - Files are available but node not discovered.

This is an Informational alarm. Examples:

```
Unit: ppm-ucs-vm13 - Bulk Statistics available for RTPZ5SVCGW01 but device not
discovered.
```

- BulkStatsInfo Alarm—The file name parameter is missing or invalid.

This is an informational alarm. If parameters are missing in the file name, Prime Performance Manager stops processing the file. New files with zero size are created in the drop directory with the same name as the original file appended with extension, skipped. Examples:

```
Unit: ppm-ucs-vm13 - Missing parameters in filename:
_bulkstats_20120419_152500_EDT_5_5.csv Missing Hostname.
```

```
Unit: ppm-ucs-vm13 - Missing parameters in filename:
Prime5k_bulkstats_20sds120419_155500_EDT_5_5.csv Unparseable date:
"20sds120419155500EDT" format:yyyyMMddHHmmssz value: 20sds120419155500EDT.
```

```
Unit: ppm-ucs-vm13 - Missing parameters in filename:
Prime5k_bulkstats_20120419_160000_EDT.csv Missing Sample Interval.
```

If necessary, you can modify the following properties in BulkStats.properties

```
FILENAME_SUBSTR = _bulkstats_
FILENAME_DELIMITER = _
DATE_VARIABLE_FORMAT = yyyyMMdd
TIME_VARIABLE_FORMAT = HHmmss
```

- BulkStatsInfo Alarm—Indicates files have no header or footer information.

The header is the first line in the bulk statistics samples. A missing header indicates the file is incomplete. Cisco ASR 5000 and 5500 devices maintain a buffer while collecting the samples until they can transfer the CSV files to the drop directory. The header is the first information stored in the buffer. If the buffer allocated on the device is too small for the sample and transfer interval, the old data in the buffer is overwritten, so the header and some collected data might be removed.

The footer is the last line in the bulk statistics samples. A missing footer also indicates the file is incomplete. The file might be in the transfer process or FTP transfer issues might exist. The header and footer are required for processing. If necessary, you can modify the following properties in BulkStats.properties.

```
HEADER_LINE_PREFIX = Version
FOOTER_LINE_PREFIX = EndOfFile
```

If header and footer is missing in the filename, Prime Performance Manager stops processing the files. New files of zero size are created in drop directory with same name as the original file appended with extension, noheader or nofooter. Examples:

```
Unit: ppm-ucs-vm13 - Bulk statistics skipped. File
Prime5k_bulkstats_20120419_161500_EDT_5_5.csv received with no header information.
Unit: ppm-ucs-vm13 - Bulk statistics skipped. File
Prime5k_bulkstats_20120419_161500_EDT_5_5.csv received with no footer information.
```

If a footer is unavailable, Prime Performance Manager waits for a specified interval to see if the transfer gets completed. If no footer exists after the specified interval, Prime Performance Manager creates the zero size file and raises the NoFooter alarm. If necessary, you can edit the following property in BulkStats.properties to control the wait duration:

```
FOOTER_WAIT_TIME = 3
```

- BulkStatsError Alarm—The device is discovered but no files are available.

This is a major alarm. Check if the drop directory is correct in BulkStats.properties. Also check device configurations. Example:

```
Device 172.18.53.231 has no bulk statistics.
```

- BulkStatsError Alarm—The device is missing 1-5 files.

This is a minor alarm. Check if device is being reloaded. Example:

```
Device 172.18.20.166 failed to receive 4 bulk statistics. Last received time is Apr
19, 2012 5:12:00 PM.
```

- BulkStatsError Alarm—The device is missing more than 5 files.

This is a major alarm. Examples:

```
Device 172.18.20.166 failed to receive 17 bulk statistics. Last received time is Apr
19, 2012 5:39:57 PM.
```

- BulkStatsError Alarm—Indicates files are received after a period of failure.

This is a normal alarm. Examples:

```
Device 172.18.53.231 receives bulk statistics as of Apr 19, 2012 5:12:00 PM.
```

- BulkStatsError Alarm—Several devices are missing more than five files.

This is a major alarm. Examples:

```
Unit ppm-ucs-vm13 - 2 devices failed to receive bulk statistics. Devices are:
Prime5k, RTPZ5SVCGW02,
```

If necessary, modify the following properties in BulkStats.properties to control when to raise alarms.

```
MINOR_ALARM_COUNT = 1
```

```
MAJOR_ALARM_COUNT = 5
```

```
NODES_FAIL_COUNT = 5
```

Setting Up StarOS Quantum Virtual Packet Core Reports

Prime Performance Manager supports the following virtual cloud architectures:

- Quantum Virtual Packet Core - Single Instance (QvPC-SI)—QvPC-SI is essentially StarOS running within a virtual machine (VM). The Single Instance architecture is best suited for low capacity scenarios. Each QvPC-SI VM takes on the roles of an entire StarOS system. The only interfaces exposed outside the VM are those for external management and service traffic. Each QvPC-SI is managed independently.
- Quantum Virtual Packet Core - Distributed Instance (QvPC-DI)—QvPC-DI addresses the scaling and redundancy limitations of QvPC-SI by extending the StarOS boundaries beyond a single VM. QvPC-DI allows multiple VMs to act as a single StarOS instance with shared interfaces, shared service addresses, load balancing, redundancy, and a single point of management.

QvPC-DI operates as a fully distributed network of multiple VMs grouped to form a StarOS instance with following major components:

- QvPC-DI Control Function CF VMs —Two CF VMs act as an active:standby (1:1) redundant pair. The active CF is responsible for controller tasks, local context MGMT, system boot image, out of band management.
- QvPC-DI Service Function SF VMs—SF VMs provide service context (user I/O ports) and handle protocol signaling and session processing tasks. A QvPC-DI instance can contain up to 46 SF VMs. A minimum configuration for a QvPC-DI instance requires four SFs - two active and two in standby mode.
- QvPC-DI Network—In order for the VMs within a QvPC-DI instance to communicate with each other, each QvPC-DI instance must have a private L2 network that interconnects the VMs. The QvPC-DI network must be for the exclusive use of a single QvPC-DI instance. No other devices might be connected to this network.
- Quantum Virtual Packet Core - Virtualized Services Module (VSM) (QvPC-VSM) on Cisco® Aggregation Services Router (ASR) 9000 Series—QvPC-VSM consists of a single StarOS instance running in a VM on a Cisco ASR 9000 VSM. The VM is represented as a virtual card with a single CPU subsystem.

Install StarOS in the virtualized environment using the StarOS product documentation. After you install StarOS, complete the following configuration steps specific for Prime Performance Manager:

-
- Step 1** Configure the SNMP read community string on the device.
 - Step 2** Configure the system host name to a unique value across the network. This value must match the `%host%` bulk statistics counter value.
 - Step 3** Verify the required number of cards are active using the "show card table" CLI on the device. For SI and VSM, one virtual card should be up with an operation state as Active. For DI, a minimum of six virtual cards must be up: one active CF, one standby CF, two active SFs, and two standby SFs.
 - Step 4** Verify the clock is set correctly on the device. (You can use the **show clock** command.)
 - Step 5** Configure the bulk statistics receivers and Prime Performance Manager bulk statistics. Verify the bulk statistics files drop to the Prime Performance Manager drop directories.
After you complete these steps, you can discover QvPC devices. For information, see [Chapter 5, "Discovering Network Devices."](#)
 - Step 6** Enable the QvPC report following steps in the ["Customizing Report Display" procedure on page 7-13](#). The report is located in Reports > Availability > QvPC - DI - VM. It shows all the VMs that are grouped together to form a StarOS DI instance along with the specific slot number, UUID for each distributed VM, the VM type SF/CF and the Operation State of each VM.
-

Converting StarOS Bulk Statistics CSV Input Files to 3GPP XML Exports

Prime Performance Manager supports the direct conversion of StarOS bulk statistics CSV input files to 3rd Generation Partnership Project (3GPP) XML exports. You have the option to export 3GPP XML files with delta calculations for counter type bulk statistics variables.

After you enable the export, each input StarOS bulk statistics CSV file has a corresponding converted export file. These files are created after the input CSV is parsed and placed in the user-specified drop directory accessible from the unit server. The directory can be mounted on SAN storage. The files collected in the export drop directory are automatically cleaned using the Bulk Stats Export Age value set in Administration menu > System Settings > Report Settings.

The converted files can also be created in CSV in the same format and filename as the input bulk statistics CSV, but with delta calculations for counter type bulk statistics variables. The converted files can be zipped in the end if you set the `ZIP_EXPORT_FILES` property in `/opt/CSCOppm-unit/properties/BulkStats.properties` to true.

The export is not enabled by default. To enable it, use the following commands:

- `ppm starexp` (See [ppm starexp](#), page B-92.)
- `ppm starexpdropdir {dir}` (See [ppm starexpdropdir](#), page B-92.)
- `ppm starexprules` (See [ppm starexprules](#), page B-93.)
- `ppm starepxmlformat` (See [ppm starepxmlformat](#), page B-93.)
- `ppm statreps bulkstatsexpage` (See [ppm statreps bulkstatsexpage](#), page B-93.)

Setting Up Generic CSV Bulk Statistics Reports

The generic CSV bulk statistics collection framework allows you to customize and define the CSV file format used to parse the CSV bulk statistics files. End users write the report extensible markup language (XMLs) to retrieve the counter metrics from the CSV bulk statistics files using the GenericCsvPoll macro definition.

For sample template definition files and report XML usage of the GenericCsvPoll macro, see examples in `/opt/CSCOppm-gw/samples/csvstats`. The properties files are the template definitions that need to be copied to `/opt/CSCOppm-gw/etc/csvstats/user/` after you customize them. The XML files are report poller definitions that you must copy to `/opt/CSCOppm-gw/etc/pollers/user/` after customizations are completed.

To set up generic CSV bulk statistics reports, complete the following procedures:

- [Defining a Generic CSV Bulk Statistics Template, page 8-29](#)
- [Setting Up Devices to Drop Files in the Drop Directory, page 8-30](#)
- [Writing Report XML Definitions to Retrieve Metrics, page 8-31](#)

Defining a Generic CSV Bulk Statistics Template

A template property file has properties that help Prime Performance Manager read and parse the CSV bulk statistics file. One template file is written for each generic CSV collection type. The template is saved in `/opt/CSCOppm-gw/etc/csvstats/system` or the user folder with properties filename extension. The filename or the template name must be unique in both the `csvstats/system` and `csvstats/user` folders. The templates in `/opt/CSCOppm-gw/etc/csvstats` automatically synchronizes to all the connected units. Key properties include:

- `header`—Comma-separated fields occurring in the CSV in the same order as expected in the CSVs.



Note All statistical data row in the CSV file must match with the fields defined in header property.

- `checkHeader`—If set to true, Prime Performance Manager validates the first CSV line starting with the first token in the header property. Prime Performance Manager raises alarm if `checkHeader` is true and the first CSV file line does not match the header property. The file is skipped. If `checkHeader` is false and CSV has a header line matching the header property, the line is skipped.
- `footer`—Defines the text for the last CSV file line. Enable this property only when CSV files have a specific footers.
- `checkFooter`—Operates the same as `checkHeader` but but checks the footer.
- `dropDir`—The folder path where CSV files are dropped from the devices. The same folder is used for all the devices using the same collection type. Folder can be mounted on a storage area network (SAN) or local disk. Prime Performance Manager must have read and write access to the folder used for collection.



Note Only one drop can exist for one specific CSV file format. If a different CSV format must be monitored and processed, then you must define a new template with its own drop directory location.

- **skipLines**—Indicates other CSV non-data lines that might be skipped during CSV file parsing. The lines can occur anywhere in the CSV, not necessarily at the beginning or end.

The format of the CSV filenames is always `nodeid_constant_datetime.extension`.

The filename starts with Node or Device Identifier. This can be `sysname`, `displayname`, `customname`, `syncname` or `primaryIP`. The middle of the filename is a constant that can be made up of one or more strings. Date and Time follow the constant and indicate whether a file is ready to age, sort files to process oldest file first, and so on. The date and time format must include year, month, day, hour, minutes, seconds and, timezone information. The following properties further define the filename. Extensions can be `csv`, `gz` or `zip`.

- **filenameNodeID**—Identifies the attribute Prime Performance Manager uses to identify the device. Valid values are `sysname`, `displayname`, `customname`, `syncname` or `primaryIP`. For example, if you choose `sysname`, Prime Performance Manager uses the device system name to look up the device and associate the CSV file to that device.



Note Each CSV file has statistics for only one device. The device must be identified by one of the `filenameNodeID` attributes.

- **filenameDelimiter**—Delimits the various filename sections.
- **filenameConstant**—Identifies the text used in the constant part of the filename. This can be one or more strings separated by the `filenameDelimiter`.
- **filenameExtension**—The `filenameExtension` is `csv`, `gz`, or `zip`.

Each CSV statistics row must contain a date and time field that helps calculate the start and end reporting interval for that row. The following fields help calculate the reporting interval and the format used to parse the field. The index is the field position in the header property starting with zero for the first field.

- **startTimeIndex**—The field index number used for reporting the interval start time.
- **startTimeFormat**—The measurement or format used to parse the start time field. The values can be `epochsecs`, `epochmillisecs`, `epochmins`, or custom date and time format.
- **endTimeIndex** and **endTimeFormat**—If the start time cannot be specified, you can use the `endTimeIndex` and `endTimeFormat` properties to indicate the end time of the reporting interval.
- **durationIndex**—Helps calculate the start or end times when only one of them can be indicated in the CSV file. The applicable `durationFormat` are `secs`, `millisecs`, or `minutes`. If `duration` cannot be specified, the `sampleInterval` property is specified.



Note See the property files defined in `/opt/CSCOppm-gw/etc/csvstats/system` for complete list of applicable property definitions and usage.

Setting Up Devices to Drop Files in the Drop Directory

Setting up devices to drop files in the drop directory is similar to other bulk statistics device setups that send CSV files to the Prime Performance Manager unit where the device is discovered. The mechanism that sends files is implementation-specific and left to you to choose. Multiple templates are available, so be sure to send files to the correct drop directories indicated in the template file.

Prime Performance Manager monitors all drop directories based on the templates in the `/opt/CSCOppm-gw/etc/csvstats/system` or user folder. After new files arrive, Prime Performance Manager starts processing the CSV bulk stats files specific to the collector type.

Prime Performance Manager can raise bulk statistics alarms when files are skipped or when filename attributes are missing or incorrect. See [Bulk Statistics Alarms, page 8-25](#) for more information on bulk statistics alarms.

Writing Report XML Definitions to Retrieve Metrics

To collect metrics from the generic bulk statistics framework you must write your report in extensible markup language (XMLs) using the `GenericCsvPoll` macro. The template file name without the properties extension is passed as the first argument of the `GenericCsvPoll` followed by the fields to pull from CSV files.



Note

See the **GenericCsvPoll** macro description in the [Cisco Prime Performance Manager 1.6 Integration Developer Guide](#) for more details.

Setting Up Small Cell Reports

To generate small cell reports in Prime Performance Manager, you must add small cell devices to Prime Performance Manager and perform a number of configurations. and configure the following small cell devices:

- RAN Management System (RMS) Central node
 - Provisioning and Management Gateway (PMG) component
 - Broadband Access Center (BAC) Regional Data Unit (RDU) component
- RMS Serving node
 - BAC Device Provisioning Engine (DPE) component
- RMS Upload Server
- 3G Access Points (APs)

Concepts and procedures for discovering and configuring small cell devices are provided in the following topics:

- [Adding Central RMS Nodes to Prime Performance Manager, page 8-31](#)
- [Adding Upload Servers to Prime Performance Manager, page 8-33](#)
- [Setting Up Broadband Access Center Reports, page 8-36](#)
- [Setting Up AP Reports, page 8-37](#)
- [Collecting Small Cell Access Point Inventory Data, page 8-39](#)

Adding Central RMS Nodes to Prime Performance Manager

The RMS central node PMG component provides its performance statistics in CSV files. Currently Prime Performance Manager only supports the `pmg-perf-periodic.csv` file, which is located in `/rms/log/pmg/` by default.

pmg-perf-periodic.csv is archived regularly as gzip files. After each archive the pmg-perf-periodic.csv contains only the lines added after the last archive. Prime Performance Manager unit regularly pulls the pmg-perf-periodic.csv file from PMG by SCP over SSH, then finds the lines of statistics added after the last pull. Prime Performance Manager processes the new lines and stores the statistics in the database.

PMG might archive the pmg-perf-periodic.csv file between two continuous pulls by Prime Performance Manager. When this occurs, Prime Performance Manager goes to the latest archived .csv.gz file to find the last line it processed.

```
[rms-aio-central] /rms/log/pmg $ ls -l pmg-perf-periodic.csv
pmg-perf-periodic.csv
```

```
[rms-aio-central] /rms/log/pmg $ ls -l pmg-perf-periodic-*.csv.gz
-rw-r--r--. 1 ciscorms ciscorms 125 Nov  9 00:03 pmg-perf-periodic-11-08-2014.0.csv.gz
-rw-r--r--. 1 ciscorms ciscorms 125 Nov 10 00:03 pmg-perf-periodic-11-09-2014.0.csv.gz
-rw-r--r--. 1 ciscorms ciscorms 125 Nov 11 00:03 pmg-perf-periodic-11-10-2014.0.csv.gz
```

```
[rms-aio-central] /rms/log/pmg $ more pmg-perf-periodic.csv
```

Server restarted.

```
period end,summary period sec,msg name,avg response time ms,max response time ms,min
response time ms,num msgs,num errors
2013-05-24T00:00:00.519Z,1913,GetAllFRMPools,69,87,53,16,0
2013-05-24T00:00:00.520Z,1913,SetFRMPools,115,130,100,2,0
2013-05-24T00:00:00.520Z,1913,GetFRMGGroupType,78,84,73,5,0
2013-05-24T00:00:00.520Z,1913,GetAllFRMGGroupTypes,88,124,69,23,0
2013-05-24T00:00:00.521Z,1913,GetAllFRMPoolTypes,87,377,54,34,0
2013-05-24T00:00:00.521Z,1913,GetFRMPools,53,53,53,2,0
2013-05-24T00:00:00.522Z,1913,GetAllFRMGGroups,122,205,86,28,0
2013-05-24T01:00:00.517Z,3599,SetFRMGGroups,88,108,69,13,0
2013-05-24T01:00:00.518Z,3599,GetFRMGGroupType,66,90,39,21,0
2013-05-24T01:00:00.518Z,3599,SetFRMPoolType,90,160,57,6,0
2013-05-24T01:00:00.518Z,3599,Register,92,184,70,11,0
2013-05-24T01:00:00.519Z,3599,SetFRMGGroupType,129,186,85,5,0
2013-05-24T01:00:00.519Z,3599,GetFRMPoolType,56,59,52,6,0
2013-05-24T01:00:00.519Z,3599,GetFRMGGroups,80,126,53,26,0
2013-05-24T01:00:00.520Z,3599,GetAllFRMPoolTypes,69,93,46,37,0
2013-05-24T01:00:00.520Z,3599,GetAllFRMGGroups,82,147,47,47,0
2013-05-24T01:00:00.520Z,3599,GetAllFRMGGroupTypes,74,96,54,23,0
2013-05-24T01:00:00.521Z,3599,GetAllFRMPools,69,86,55,30,0
```

Server restarted.

To add the Central RMS node to Prime Performance Manager:

-
- Step 1** Verify that Network Time Protocol (NTP) is synchronized between Prime Performance Manager unit and the small cell devices.
 - Step 2** Configure the RMS node SNMP credential. See [Adding SNMP Device Credentials, page 5-6](#).
 - Step 3** Configure the RMS node SSH credential. See [Adding Device Credentials for Other Protocols, page 5-8](#).
 - Step 4** Complete the RMS node discovery. See [Prime Performance Manager Device Discovery, page 5-5](#).
 - Step 5** Verify the etc/csvPull/system/pmg-perf.properties file settings on the gateway are consistent with the PMG configuration. The gateway synchronizes the file to all units. If settings are inconsistent, make the appropriate changes:
 - a. Verify the activeFileDir and activeFileName so Prime Performance Manager knows the pmg-perf-periodic.csv absolute file path.
 - b. Verify the rollOverFileDir and rollOverFilePattern so Prime Performance Manager knows where the archived csv files are located in PMG.

- c. Verify the `headerPrefix` matches the `pmg-perf-periodic.csv` header line. This setting tells Prime Performance Manager which line is the CSV file header. Configure the fixed part that the header line starts with, such as `period end`. Do not configure the whole header line because the header line might change across different versions.
- d. Verify the `skipLines`. PMG creates a line `Server Restarted` in the CSV file when it is started or restarted. Prime Performance Manager ignores this line by configuring it in the `skipLines` field.
- e. Verify the `endTimeFormat`. This is the date and time format for the first field `period end` in the CSV file.

Sample `pmg-perf.properties`:

```
[root@crdc-c240-176 CSCOppm-unit]# more etc/csvPull/system/pmg-perf.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
#
# The active file directory: absolute path
activeFileDir = /rms/log/pmg/
# The active file name
activeFileName = pmg-perf-periodic.csv
# Archived file directory: absolute path
rollOverFileDir = /rms/log/pmg/

# Archived file name pattern
rollOverFilePattern = pmg-perf-periodic-*.csv.gz

# File content delimiter
contentDelimiter = ,

# The fixed part that the header line starts with, used by Prime Performance Manager to
identify the header line.
# Don't configure the whole header line here because the header line may be different
across versions.
headerPrefix = period end

# Specify non-data lines to be skipped in CSV file.
# Separate multiple lines with comma, for each line just specify the intial part the line
starts with.
# PMG creates a line "Server restarted." in its CSV file when it's started/restarted.
skipLines = Server restarted.

# Format of the "period end" field in the CSV file.
endTimeFormat = yyyy-MM-dd'T'HH:mm:ss.SSS
```



Note

You only need to configure the `etc/csvPull/system/properties/pmg-perf.properties` file. The `etc/csvstats/system/pmg-perf.properties` file, which used to configure how Prime Performance Manager pulls CSV and generate intermediate CSV files, is removed in this release.

Adding Upload Servers to Prime Performance Manager

Upload Servers provide their performance statistics (non-AP statistics) in CSV files. Currently Prime Performance Manager only supports the `upload-perf-periodic.csv` file located under `/opt/CSCOul/logs/` by default. `upload-perf-periodic.csv` is archived regularly as `gzip` files. After each archive, `upload-perf-periodic.csv` contains only the lines added after the last archive.

Prime Performance Manager units regularly pull the upload-perf-periodic.csv file from the Upload Server using SCP over SSH and finds the statistics lines added after the last pull. The new lines are processed and the statistics are stored in the database.

The Upload Server might archive the upload-perf-periodic.csv file between two continuous pulls by Prime Performance Manager. If this occurs, Prime Performance Manager goes to the last archived .gz file to find the last line it processed.

```
[admin1@rms-aio-upload logs]$ ls /opt/CSCOUls/logs/upload-perf-periodic.csv
/opt/CSCOUls/logs/upload-perf-periodic.csv[admin1@rms-aio-upload ~]$ more
/opt/CSCOUls/logs/upload-perf-periodic.csvperiod end,period duration (sec),uptime ms,
SB uploads completed,SB uploads failed,SB status sent,SB status failure,SB bytes rcvd,
SB bytes sent,SB session timeout,NB downloads completed,NB downloads failed,NB bytes
rcvd,NB bytes sent,NB session timeout,Bytes before compression,Bytes after
compression,Bytes deleted,Bytes archived,Number of files listed,Number of files
deleted,Number of files archived,Total time compressing,Total time archiving,Total time
deleting,Total time listing,Number of archives,Total stat files,Total stat bytes,
Total demand files,Total demand bytes,Total unknown files,Total unknown bytes,max
concurrent uploads,max concurrent downloads
2014-11-12T00:00:00.000,3599,1705636932,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,7,2,0,0,0,0,0,0,0,0,0,0
2014-11-12T01:00:00.000,3599,1709236932,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,8,1,0,0,0,0,0,0,0,0,0,0
2014-11-12T02:00:00.000,3600,1712836932,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,6,1,0,0,0,0,0,0,0,0,0,0
2014-11-12T03:00:00.000,3600,1716436932,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,7,1,0,0,0,0,0,0,0,0,0,0
2014-11-12T04:00:00.001,3600,1720036933,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,8,2,0,0,0,0,0,0,0,0,0,0
2014-11-12T05:00:00.001,3600,1723636933,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,7,2,0,0,0,0,0,0,0,0,0,0
2014-11-12T06:00:00.000,3599,1727236932,0,0,0,0,0,0,160,0,169872,77520,0,0,0,0,0,80,0,0,
0,0,7,1,0,0,0,0,0,0,0,0,0,0

[admin1@rms-aio-upload ~]$ ls -l /opt/CSCOUls/server-perf-archives/daily_archives/
-rw-r--r--. 1 ciscorms ciscorms 572 Nov 9 00:00 upload-perf-periodic-2014-11-08.0.gz
-rw-r--r--. 1 ciscorms ciscorms 549 Nov 10 00:00 upload-perf-periodic-2014-11-09.0.gz
-rw-r--r--. 1 ciscorms ciscorms 562 Nov 11 00:00 upload-perf-periodic-2014-11-10.0.gz
```

To add Upload Servers to Prime Performance Manager:

-
- Step 1** Verify that Network Time Protocol (NTP) is synchronized between Prime Performance Manager unit and the small cell devices.
 - Step 2** Verify the Linux find utility on the Upload Server is Findutils-4.4.2 or later. If not, upgrade it following procedures in the Linux documentation. (This is required for AP reports.)
 - Step 3** Configure the Upload Server SNMP credential. See [Adding SNMP Device Credentials, page 5-6](#).
 - Step 4** Configure the Upload Server SSH and ULS_HTTP credentials. See [Adding Device Credentials for Other Protocols, page 5-8](#).
 - Step 5** Discover the Upload Servers. See [Prime Performance Manager Device Discovery, page 5-5](#).
 - Step 6** Verify the etc/csvPull/system/uls-perf.properties file settings are consistent with the Upload Server configuration. The gateway synchronizes this file to all units. If settings are inconsistent, make the appropriate changes:
 - a. Verify the activeFileDir and activeFileName so Prime Performance Manager knows the upload-perf-periodic.csv absolute file path.
 - b. Verify the rollOverFileDir and rollOverFilePattern so Prime Performance Manager knows the archived CSV file location on the Upload Server.

- c. Verify the headerPrefix to match the header line of upload-perf-periodic.csv. This setting tells Prime Performance Manager the CSV file header line. Configure the fixed part that begins the header line, such as “period end”. Do not configure the whole header line because the header line might differ across different versions.
- d. Verify the skipLines. The Upload Server might add server restarted lines in the CSV file when it is started or restarted. Prime Performance Manager ignores these lines by adding them to the skipLines field.
- e. Verify the endTimeFormat. This is the date and time format for the first field "period end" in the CSV file.
- f. For multiple Upload Servers, if CSV files are located in different directories or the CSV file name or pattern varies, configure the following per-device settings:
 - activeFileDir. <node_primary_IP>
 - activeFileName. <node_primary_IP>
 - rollOverFileDir. <node_primary_IP>
 - rollOverFilePattern. <node_primary_IP>

For example:

```
activeFileDir.192.168.0.100 = /aa/bb/cc
activeFileName.192.168.0.100 = upload-perf-periodic.csv
rollOverFileDir.192.168.0.100 = /dd/ee/ff
rollOverFilePattern. 192.168.0.100 = upload-perf-periodic-*.gz
```

Per-device settings override the default ones if both exists.

Sample upload-perf-periodic.csv:

```
[root@crdc-c240-176 CSCOppm-gw]# more etc/csvPull/system/uls-perf.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
#
# The active file directory: absolute path
activeFileDir = /opt/CSCOuls/logs/

# The active file name
activeFileName = upload-perf-periodic.csv

# Archived file directory: absolute path
rollOverFileDir = /opt/CSCOuls/server-perf-archives/daily_archives/

# Archived file name pattern
rollOverFilePattern = upload-perf-periodic-*.gz

# File content delimiter
contentDelimiter = ,

# The fixed part that the header line starts with, used by Prime Performance Manager to
identify the header line.
# Don't configure the whole header line here because the header line might be different
across versions.
headerPrefix = period end

# Specify non-data lines to be skipped in CSV file.
# Separate multiple lines with comma, for each line just specify the intial part the line
starts with.
# Upload Server might create a line "Server restarted." in its CSV file when the Upload
Server is started or restarted.
skipLines = Server restarted.
```

```
# Format of the "period end" field in the CSV file.
endTimeFormat = yyyy-MM-dd'T'HH:mm:ss.SSS
```



Note You only need to configure the `etc/csvPull/system/properties/uls-perf.properties` file. The `etc/csvstats/system/uls-perf.properties` file, which used to configure how Prime Performance Manager pulls CSV and generates intermediate CSV files, was removed in this release.

Setting Up Broadband Access Center Reports

Cisco Broadband Access Center performance reports are supported. BAC consists of a BAC RDU component running on an RMS Central Node, and a BAC DPE component running on RMS serving nodes.

The BAC RDU and DPE provide performance counters log files in CSV format. Prime Performance Manager retrieves the log files by SCP over SSH and parses them to generate reports that can be viewed in the Prime Performance Manager GUI. Currently Prime Performance Manager only the `perfstat.log` file is supported. By default:

- BAC RDU `perfstat.log` is located in `/rms/data/CSCObac/rdu/logs/statistics`.
- BAC DPE `perfstat.log` is located in `/rms/data/CSCObac/dpe/logs/statistics`.

`perfstat.log` file is archived daily by renaming the current `perfstat.log` file to `perfstat.1.log`, and renaming the old `perfstat.1.log` (if it exists) to `perfstat.2.log`, and so on. The last 30 days of `perfstat.log` files are retained. `perfstat.1.log` is always the last day before the current day. This is a fixed file name pattern handled automatically by Prime Performance Manager; you do not need to configure the file name pattern in the Prime Performance Manager properties file.

The BAC RDU and DPE might archive the active `perfstat.log` file between two continuous pulls by Prime Performance Manager. If this occurs, Prime Performance Manager goes to the archived `perfstat.xx.log` files to find the last line it processed.

```
[rms-aio-central] ~ $ ls -lt /rms/data/CSCObac/rdu/logs/statistics
total 18388
-rw-rw-r--. 1 root root 223180 Nov 12 08:02 perfstat.log
-rw-rw-r--. 1 root root 662640 Nov 11 23:56 perfstat.1.log
-rw-rw-r--. 1 root root 662640 Nov 10 23:56 perfstat.2.log
-rw-rw-r--. 1 root root 660340 Nov 9 23:55 perfstat.3.log
-rw-rw-r--. 1 root root 662640 Nov 8 23:59 perfstat.4.log
-rw-rw-r--. 1 root root 662640 Nov 7 23:58 perfstat.5.log
-rw-rw-r--. 1 root root 662640 Nov 6 23:57 perfstat.6.log
-rw-rw-r--. 1 root root 662640 Nov 5 23:57 perfstat.7.log
-rw-rw-r--. 1 root root 662640 Nov 4 23:56 perfstat.8.log
-rw-rw-r--. 1 root root 660340 Nov 3 23:55 perfstat.9.log
-rw-rw-r--. 1 root root 662640 Nov 2 23:59 perfstat.10.log
-rw-rw-r--. 1 root root 662640 Nov 1 23:58 perfstat.11.log
-rw-rw-r--. 1 root root 662640 Oct 31 23:57 perfstat.12.log
-rw-rw-r--. 1 root root 662640 Oct 30 23:57 perfstat.13.log
-rw-rw-r--. 1 root root 662640 Oct 29 23:56 perfstat.14.log
-rw-rw-r--. 1 root root 660340 Oct 28 23:55 perfstat.15.log
-rw-rw-r--. 1 root root 662640 Oct 27 23:59 perfstat.16.log
-rw-rw-r--. 1 root root 662640 Oct 26 23:58 perfstat.17.log
-rw-rw-r--. 1 root root 662640 Oct 25 23:57 perfstat.18.log
-rw-rw-r--. 1 root root 662641 Oct 24 23:57 perfstat.19.log
-rw-rw-r--. 1 root root 662640 Oct 23 23:55 perfstat.20.log
-rw-rw-r--. 1 root root 662641 Oct 22 23:55 perfstat.21.log
-rw-rw-r--. 1 root root 662640 Oct 21 23:55 perfstat.22.log
```

```
-rw-rw-r--. 1 root root 662643 Oct 20 23:55 perfstat.23.log
-rw-rw-r--. 1 root root 662641 Oct 19 23:55 perfstat.24.log
-rw-rw-r--. 1 root root 296837 Oct 18 23:55 perfstat.25.log
-rw-rw-r--. 1 root root 241786 Oct 17 08:36 perfstat.26.log
-rw-rw-r--. 1 root root 667228 Oct 16 23:55 perfstat.27.log
-rw-rw-r--. 1 root root 669552 Oct 15 23:59 perfstat.28.log
-rw-rw-r--. 1 root root 669554 Oct 14 23:59 perfstat.29.log
```

You must enable the performance log (perfstat.log) for BAC RDU and DPE. See [Setting Up Cisco Broadband Access Center Reports, page 8-42](#) to complete this and remaining report setup steps.

Setting Up AP Reports

Access Points (APs) send performance files regularly to the Upload Server in 3GPP XML format. Upload Servers are added to Prime Performance Manager as devices, and Prime Performance Manager pulls the AP performance files from the Upload Server local disk to generate reports.

Prime Performance Manager units regularly scan the Upload Server AP file directory through SSH to find the files uploaded to it since last scan. The file names are inserted to the Prime Performance Manager AP file list queue. Prime Performance Manager polling retrieves the file names from the file list queue, composes an HTTP URL for each file and pulls the AP file from the Upload Server by HTTP GET.

You must configure the `etc/apStats/system/RMS-ULS.properties` file, and verify the setting is consistent with Upload Server configuration. Prime Performance Manager gateway synchronizes this file to all units.

- `deviceCapability` is an internal Prime Performance Manager property, do not modify it.
- Enter the top AP fileDirectory absolute path so Prime Performance Manager knows the Upload Server AP performance files location. Multiple Upload Servers might have different AP file directories. To configure per Upload Server settings, enter: `fileDirectory.<node_primary_IP>`. For example:

```
fileDirectory.192.168.0.100 = /opt111/CSCOuls/files/uploads/stat
```

The per-device setting overrides the default setting if both exist.

- Configure the `endTimeFormat` to match the `endTime` attribute of the `<granPeriod>` in the AP performance file.

```
# endTime format used by the 3GPP XML file
# <granPeriod duration="PT900S" endTime="1970-01-01T01:15:00+01:00"/>
endTimeFormat = yyyy-MM-dd'T'HH:mm:ssXXX
```

- `maxFileSize` controls the file size that Prime Performance Manager unit pulls from the Upload Server. This is for performance tuning with assistant by Cisco TAC.

```
[root@crdc-c240-176 CSCOppm-gw]# more etc/apStats/system/RMS-ULS.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
#
# What capability a device should have to enable this template
deviceCapability = RMS_ULS_AP

# The 3GPP XML file directory (top directory)
fileDirectory = /opt/CSCOuls/files/uploads/stat

# endTime format used by the 3GPP XML file
# <granPeriod duration="PT900S" endTime="1970-01-01T01:15:00+01:00"/>
endTimeFormat = yyyy-MM-dd'T'HH:mm:ssXXX
```

```
# Pull and parse this amount of file size from device each time, unit: KB
maxFileSize = 81920
```

To add Upload Servers to Prime Performance Manager for AP reports:

-
- Step 1** Verify that Network Time Protocol (NTP) is synchronized between Prime Performance Manager unit and the small cell devices.
 - Step 2** Verify the Linux find utility on the RMS server is Findutils-4.4.2 or later. If not, upgrade it following procedures in the Linux documentation.
 - Step 3** Configure the Upload Server SNMP credential in Prime Performance Manager. See [Adding SNMP Device Credentials](#), page 5-6.
 - Step 4** Configure the Upload Server SSH credential in Prime Performance Manager. See [Adding Device Credentials for Other Protocols](#), page 5-8.
 - Step 5** Configure the Upload Server ULS_HTTP credential for APs in Prime Performance Manager. See [Adding Device Credentials for Other Protocols](#), page 5-8.

The fileDirectory property in the RMS-ULS.properties file is the Upload Server absolute and top AP file directory. The ULS_HTTP credential Subsystem field is the relative part of the fileDirectory in terms of the Upload Server HTTP API URL path. For example, if fileDirectory=/opt/CSCOuls/files/uploads/stat, configure the subsystem as uploads/stat in the ULS_HTTP credential form. If fileDirectory=/opt/CSCOuls/files/uploads/stat/xxx, configure the subsystem as uploads/stat/xxx in the ULS_HTTP credential form.

- Step 6** For Cisco RAN Management System 4.1 or later, verify the archives/stat directory is configured in RMS-ULS.properties:

```
#####
# The 3GPP XML file directory: the absolute and top AP file directory.
#
# RMS Upload Server provides compressed XML AP files in its archives/stat
# directory since RMS 4.1.
# The compressed files are in gzip format with .xml.gz file extension.
# PPM identifies the format automatically by the file extension .xml.gz,
# no need to configure this format.
#####
fileDirectory = /opt/CSCOuls/files/archives/stat
```



Note The RAN Management System 4.1 Upload Server archives/stat directory contains compressed XML AP files in gzip format with an .xml.gz file extension. Prime Performance Manager identifies the format automatically so you do not need to configure this format.

- Step 7** In residential SmallCell deployments where the AP data is uploaded once a day, additional system properties must be set to trigger threshold events on the aggregate reports in Prime Performance Manager. In \$GW/properties/System.properties, add/update the following properties:

```
MetaDataStrategy = MAX_DELAY
AGG_MAX_DELAY = 9000000
```

- Step 8** Restart the gateway. See [Restarting Gateways and Units](#), page 2-5.

The MetaDataStrategy attribute tells what strategy to use when waiting for contributors to an aggregation. MAX_DELAY tells Prime Performance Manager to wait AGG_MAX_DELAY milliseconds after the first instance of data for a given timestamp.

Step 9 Discover the Upload Servers. See [Prime Performance Manager Device Discovery, page 5-5](#).

Collecting Small Cell Access Point Inventory Data

Small cell access point (AP) inventory data is static, and not performance KPIs for which you can create threshold crossing alerts in Prime Performance Manager. Collecting AP inventory data can be used to group AP performance reports. For example, you might want to group AP performance reports by a particular AP site or Home Node B Gateway (HNBGW).

To collect AP inventory data, you must provide information to the `ApInv.properties` file located in the `/opt/CSCOppm-gw/etc/apInventory/template/` directory. `ApInv.properties` collects AP inventory from:

- Get Device Data Tool (GDDT). This script is included in the Provisioning Management Gateway (PMG) component of the RAN Management System (RMS).
- DNPrefix. DNPrefix is contained in normal AP performance file polled by other collectors.

Information you will provide includes:

- The field name that represents the unique AP ID.
- The fields to extract from the exported GDDT CSV file.
- The path of the file exported by the GDDT script.
- The default GDDT output file extension is `.csv`.
- The local directory where the exported GDDT file will be copied.
- Interval in (hours) at which Prime Performance Manager retrieves the GDDT exported file from PMG. The default is 24 hours.
- The length of the time, in days, that Prime Performance Manager should keep copied CSV files. The default is three days. Files are archived and deleted after this time period.
- Indicator telling Prime Performance Manager how to treat the field, `Role`, in inventory data. If this property is set to `false`, Prime Performance Manager assigns `Role` with the value contained in the GDDT output identified by the field, `ENTID`, or DNPrefix, identified by the field, `EnterPriseID`. Otherwise, Prime Performance Manager treats `Role` as a binary value, namely `Enterprise` or `Residential`.
- The name of the DNPrefix output field to reference for the `Role` value. The default is `EnterPriseID`.
- The name of the GDDT output field to reference for the `Role` value. The default is `ENTID`,
- The number of attempts allowed to retrieve the GDDT output file before considering it a failure. The default is 3.
- The interval, in minutes, between each adjacent attempt to retrieve the GDDT output file, The default is 3 minutes.

Prime Performance Manager uses a minimum set of fixed field names, shown in [Table 8-2](#). To generate small cell reports, the constructed AP inventory data must have corresponding fields names. The preferred method is to configure RMS to have field names that match the names used by Prime Performance Manager for each inventory item, whether it is from GDDT output or DNPrefix. For example, when configuring the inventory item, `site`, (for DNPrefix or GDDT output), the name, `SITEID`, should be used to match the site name used by Prime Performance Manager. If you cannot name inventory items on the RMS side, you must do additional field name mapping on the Prime Performance

Manager side so Prime Performance Manager knows what field to look for a certain inventory items. The field mapping can be configured in /opt/CSCOppm-gw/etc/properties. After you modify FieldName.properties, restart the gateway and units. See [Restarting Gateways and Units, page 2-5](#).

Table 8-2 AP Inventory Field Names

Field Name Used by Prime Performance Manager	Description
HNB-GW	Unique ID for the Cisco ASR 5000 chassis at discretion of the SP. This tag is only provided for HNB.
HNB-GW-SRV	HNB-GW service name within a chassis. Same value as in ASR 5000 CLI. This tag is only provided for HNB.
HeNB-GW	Unique ID for Cisco ASR 5000 chassis at discretion of SP. This tag is only provided for HeNB.
HeNB-GW-SRV	HeNB-GW service name within a chassis. Same value as in Cisco ASR 5000 CLI. This tag is only provided for HeNB.
SeGW	Unique ID for SeGW chassis at discretion of SP. In case of Cisco ASR 5000 SeGW, the value is the same as for HNB-GW tag.
SeGW-SRV	Unique ID for SeGW service within the chassis. In case of Cisco ASR 5000 SeGW, the value is the same as for HNB-GW-SRV tag.
Area	Geographical area of Small Cell. Optional in some deployments - empty value or no tag.
ENTID	Numeric enterprise ID for Small Cell. Empty value or no tag for Resi APs.
SITEID	Numeric enterprise Site ID for Small Cell. Empty value or no tag for Resi APs.
Chassis	The ID of the chassis of the multi-stack AP. Empty value or no tag for single-stack AP.
RNC ID	The HNB ID or HeNB ID.
secondary ID	The optional secondary ID for the AP, at SP discretion. Tag can be missing or have empty value.
SW Version	
HW Version	
RAT	3G or LTE
Manufacturer	
Live HW Model	
Role	
RF Profile	

After you modify ApInv.properties, it is automatically sent to all connected units. You must restart the units. A sample ApInv.properties file is shown below:

```
# What capability a device should have to enable this template.
deviceCapability = RMS_PMG_PERF

# Name of the field representing the unique id of Access Point node.
apNodeIdField = EID
```



```

# A list of fields users want to extract from GDDT exported csv file. Empty
# value or non-explicitly specified value means users want to extract all fields
# included in the csv file.
header =

# Full path of the file exported by GDDT script.
gddtExportFile = /rms/ops/GetDeviceData-reports/latest/device-data.csv

# File extension of the GDDT output file, will use '.csv' as default
# if not explicitly specified.
fileExtension = .csv

# Local directory which the GDDT exported file will be copied to.
dropDir = /tmp/apinv/

# Interval(hours) in which PPM regularly pull GDDT exported file from RMS-PMG.
updateInterval = 24

# Length of the time(days) PPM will keep those copied csv files. Those pulled
# files will be archived and deleted after this period of time.
ageOutInterval = 3

# Toggle telling PPM how to treat the field 'Role' in inventory data. If this
# property is explicitly specified with 'false' value, PPM will assign 'Role'
# with the value contained in GDDT output(identified by field 'ENTID') or
# DNPrefixed(identified by field 'EnterpriseID'). Otherwise, PPM will treat 'Role'
# as a binary value, namely Enterprise or Residential.
roleAttrBinary = true

# Name of the field in DNPrefixed which we refer to for the value of 'Role',
# default with 'EnterpriseID'
entidInDnprefix = EnterpriseID

# Name of the field in GDDT output which we refer to for the value of 'Role',
# default with 'ENTID'
entidInGddt = ENTID

# Number of attempts allowed to retrieve GDDT output file before considered failure.
retryCount = 3

# Interval between each adjacent attempt to retrieve GDDT output file, in minutes.
retryInterval = 3

```

GDDT and DNPrefixed must use consistent field names. If field names are not consistent, provide the DNPrefixed to GDDT mapping in the FieldNames.properties file located in /opt/CSCOppm-gw/etc/properties. The FieldNames.properties file is also located in /opt/CSCOppm-gw/etc/apInventory/template/.

Setting Up Ganglia Reports

The Ganglia Monitoring System is a scalable distributed monitoring system for high-performance computing systems such as clusters and grids. The Ganglia hierarchical organization is designed for cluster federations. It uses XML for data representation, External Data Representation (XDR) for portable data transport, and the Round Robin Databast (RRD) tool for data storage and visualization. For additional information, visit the Ganglia website: <http://ganglia.info>.

To generate Ganglia performance reports:

-
- Step 1** Complete the “[Adding SNMP Device Credentials](#)” procedure on page 5-6 to add the SNMP credentials for the server hosting Ganglia.
- Step 2** Complete the “[Adding Device Credentials for Other Protocols](#)” procedure on page 5-8 to add the Ganglia information. In the Add Credentials Entry dialog box,
- Device Name—Enter the device name.
 - Connection Protocol—Choose **GMOND_SOCKET**.
 - Port—Port 8649 is the default.
 - Sub System—Leave blank.
 - User Name—Leave blank.
 - Password—Leave blank.
- Step 3** Click **OK**.
- Step 4** On the dialog asking if you want to save without a username or password, click **OK**.
- Step 5** Complete the “[Running Device Discovery from Prime Performance Manager](#)” procedure on page 5-13 to add the server that hosts Ganglia.
- Step 6** To view performance data for remote Gmond hosts:
- a. From the Performance menu, choose **Reports**,
 - b. In the report navigation tree choose **Compute > Ganglia > Gmond**.
- The report lists all remote hosts collected by the gmond server including the local host.
- Step 7** To view performance data for the local gmond host:
- a. From the Performance menu, choose **Resources**,
 - b. Choose a subcategory, for example Buffers, CPU, or Disk.
- The report displays the gmond node performance data. It does not include remote hosts.
-

Setting Up Cisco Broadband Access Center Reports

Cisco Broadband Access Center (BAC) provides information about the traffic between customer premises equipment (CPE) and the BAC device provisioning engine (DPE). This data provides visibility into traffic flows that might be causing network issues. This traffic profiling provides statistics on the following:

- The number of handled CPE WAN Management Protocol (CWMP) sessions.
- The number of rejected devices.
- The number of handled HTTP file requests.
- The Home Provisioning Group redirection status.
- Identification of traffic caused by chatty clients.

The periodic statistics provides details, including the name of each Remote Procedure Call (RPC) and the specific information message types. The following RPC methods are monitored and reported:

- GetRPCMethods
- SetParameterValues

- GetParameterValues
- SetParameterAttributes
- GetParameterAttributes
- AddObject
- DeleteObject
- Download
- Reboot
- Inform
- TransferComplete
- AutonomousTransferComplete
- GetQueuedTransfers
- ScheduleInform
- SetVouchers
- GetOptions
- Upload
- FactoryReset
- GetAllQueuedTransfers
- Kicked
- RequestDownload

The BAC DPE and RDU provide performance counters log files in CSV format. Prime Performance Manager retrieves and parses the log files for reports that can be viewed in the Prime Performance Manager GUI.

To enable reports based on BAC performance logs:

-
- Step 1** Log into the BAC RDU.
- Step 2** From the user interface, choose **Configuration > Defaults > System Defaults**.
- Step 3** For Performance Statistics Collection, click **Enabled**.
- Step 4** To enable traffic statistics on the DPE, from the DPE CLI in the enabled mode, enter **debug dpe statistics**.
- Step 5** Verify that Network Time Protocol (NTP) is synchronized between the Prime Performance Manager unit and the small cell devices.
- Step 6** Enable the SNMP and SSH on the BAC system by configuring the SNMP and SSH credentials.



Note For additional information about Steps 1 through 6, see the [Cisco Broadband Access Center Administration Guide 3.8](#).

- Step 7** For BAC statistics performance reports, verify the parameters are correctly configured on the Prime Performance Manager gateway:
- For RDU performance statistics, see the etc/bacStats/system/bac-rdu-perf.properties file.
 - For DPE performance statistics, see the etc/bacStats/system/bac-dpe-perf.properties file.

The properties file contains the log file path. It supports configuring different fileDirectory for source log files in the properties file by activeFileDir.<node_primary_IP>. It shows the default directory. If per-device configuration exists, the new directory overrides the default.

For example:

```
activeFileDir.192.168.0.100 = /aa/bb/cc
activeFileName.192.168.0.100 = perfstat.log
```

Per-device setting overrides the default one if both exists.

```
[root@crdc-c240-176 CSCOppm-gw]# more etc/bacStats/system/bac-rdu-perf.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
#
# The active file directory: absolute path
activeFileDir = /rms/data/CSCObac/rdu/logs/statistics

# The active file name
activeFileName = perfstat.log
[root@crdc-c240-176 CSCOppm-gw]# more etc/bacStats/system/bac-dpe-perf.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
# The active file directory: absolute path
activeFileDir = /rms/data/CSCObac/rdu/logs/statistics

# The active file name
activeFileName = perfstat.log

[root@crdc-c240-176 CSCOppm-gw]# more etc/bacStats/system/bac-dpe-perf.properties
# Copyright (c) 2014 by Cisco Systems, Inc.
#
# The active file directory: absolute path
activeFileDir = /rms/data/CSCObac/dpe/logs/statistics

# The active file name
activeFileName = perfstat.log
```

- Step 8** Enable the relevant BAC performance statistics reports (Small Cell Statistics > RMS > RMS System). For information, see [Customizing Individual Report Settings, page 7-25](#).

Ceph and KVM VM Report Notes

Ceph is a distributed object store and file system designed to provide performance, reliability and scalability. If you enabled Ceph reports (see [Ceph Device Support, page 5-22](#)), they appear under the Storage report category.



Note

If no Ceph data is collected, verify that the Ceph device time is synchronized with the Prime Performance Manager server. If necessary, add an NTP server to the Ceph device.

Hypervisor-level disk space monitors the capacity and used space from a layer above the OpenStack instance where it is deployed. For local storage, this means disk space usage as viewed by the KVM hypervisor. For certain block device formats, the KVM hypervisor can allocate the entire disk capacity, even though disk capacity might not actually be used one level lower (guest level).

Prime Performance Manager provides two KVM VM disk space reports, one for the hypervisor-level storage, and one for the lower guest-level storage:

- KVM VM Hypervisor-Level Disk Space (formerly KVM VM Disk Bytes)—Supports only local storage.
- KVM VM Guest-Level Disk Space—Gathers data for both local and network (Ceph) storage back ends.

Guest-level disk space monitors the capacity and used space as it appears inside the OpenStack instance. This is the usable space in the OpenStack instance. Guest-level disk space is calculated using the `libguestfs` library. To monitor guest-level disk space in KVM, you must add a new device credential. Previously, only a `KVM_TLS` user credential was required to get hypervisor-level disk space statistics. For guest-level storage disks, a `SSH_V2` user credential for the KVM device is required. The `libguestfs` (specifically the `virt-df` command) version installed on the KVM device must be 1.26.0 or later.

To get Ceph Storage and KVM VM Guest-Level Disk Space report data, the device user credential must have access to the Ceph configuration file. The default configuration file location is `/etc/ceph/$clusterName.conf`, where `$clusterName` is the name of the cluster. The default is `ceph`. If the user credential added to Prime Performance Manager is a non-root user, the user must have read permission in the Ceph configuration file. (If the device is not getting Ceph Storage and KVM VM Guest-Level Disk Space report data, this might be the problem.) For KVM VM Guest-Level Disk Space reports, if you follow the documentation for integrating Ceph with OpenStack, the Ceph configuration file should have the required permission. If it does not, you can set it using the `chmod +r /etc/ceph/$clusterName.conf` command.

ONS and CPT Device Report Notes

Prime Performance Manager provides reports at 15-minute and 1-day intervals for Cisco Optical Network Service (ONS) and Carrier Packet Transport (CPT) devices. You can access these reports in Reports > Transport Statistics > CPT/ONS. CPT and ONS device report data include:

- Ethernet interfaces—Bytes, Unicast Packets, Multicast Packets, Total Packets, Interface Discards, Interface Errors.
- G.709 Section—Errored Seconds, Severely Errored Seconds, Unavailable Seconds, Background Block Errors, Failure Counts, Errored Seconds Ratio, Severely Errored Seconds Ratio, Background Block Errors Ratio.
- G.709 Path—Errored Seconds, Severely Errored Seconds, Unavailable Seconds, Background Block Errors, Failure Counts, Errored Seconds Ratio, Severely Errored Seconds Ratio, Background Block Errors Ratio

Additionally, Prime Performance Manager supports HTTP DT FEC and DWDM client statistics.
