ORACLE

Exadata Health and Resource Utilization Monitoring

Exadata Database Machine KPIs

November 2022 Version 3.2

Public

Purpose statement

This document provides an overview of monitoring Exadata key performance indicators. It is intended solely to help you add additional performance monitoring for your Exadata systems.

Table of contents

List of images

List of tables

Purpose

Oracle Exadata allows customers to lower costs by reducing the resources needed for OLTP, analytics, and consolidated Oracle Database workloads. Exadata runs Oracle Database with the highest performance, scale, and availability, allowing customers to improve operating efficiency while simplifying management workloads. Understanding the components and setting up appropriate monitoring are key tasks for Exadata administrators. In this technical brief the components of the Exadata I/O stack will be discussed, key performance indicators identified, and Oracle's world class monitoring solution, Enterprise Manager, will be used to provide a holistic approach to making sure the environment is functioning properly.

Exadata Overview

Exadata

Oracle Exadata is a modern architecture featuring scale-out industry-standard database servers, scale-out intelligent storage servers, and an extremely high speed internal RDMA Network Fabric that connects the database and storage servers. Over the years, the Exadata Database Machine has evolved and depending on the generation, you may see different configurations and components including elastic configurations which allow customer-defined combinations of compute and storage servers. For more details on the hardware configurations available for each Exadata generation see the Exadata Database Machine System Overview documentation. To learn about the architecture in detail visit the Oracle Exadata Machine Technical Architecture.

Enterprise Manager monitors all aspects of the Exadata Infrastructure, taking into account any enhancements that come with newer generations. The Oracle Exadata Database Machine target includes:

- Host
- Oracle Virtual Platform
- Oracle VM Instance
- Exadata Storage Server
- Exadata Storage Server Grid
- Systems Infrastructure Server
- Systems Infrastructure Switch (RoCE or InfiniBand)
- Systems Infrastructure PDU

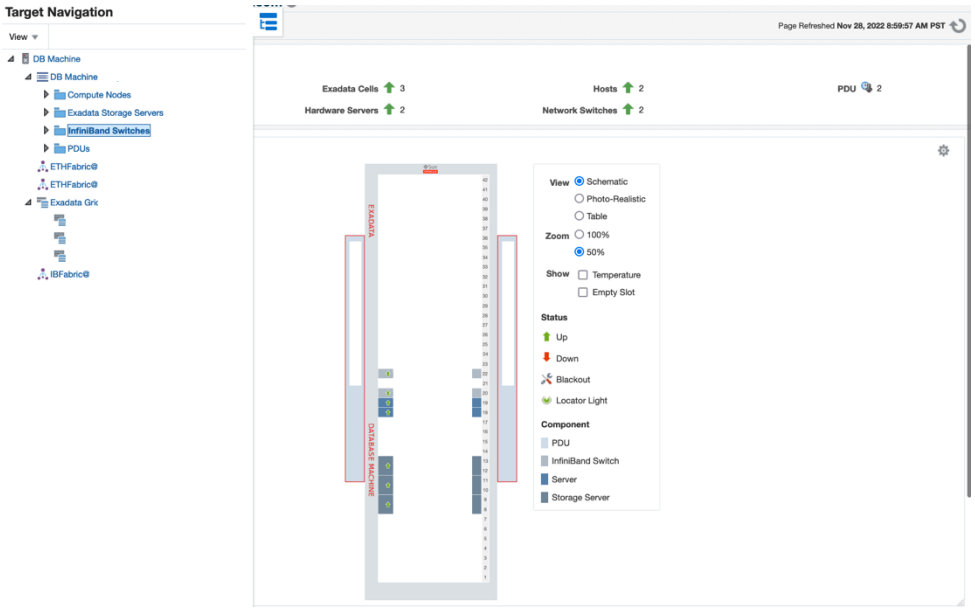

Image 1. Exadata Database Machine target in Enterprise Manager.

For more information on monitoring Exadata Database Machine with Enterprise Manager, see Oracle Exadata Database Machine Getting Started Guide.

Exadata Cloud

Exadata Cloud brings in additional considerations because Oracle Cloud Operations are responsible for monitoring and managing the storage, network and other Exadata infrastructure components, while the customer is primarily responsible for monitoring and managing the VM, cluster, ASM and database components.

Enterprise Manager (EM) organizes these customer-managed targets into the Oracle Exadata Cloud Service target. Of the list of infrastructure targets depicted in the previous section for the on-premises Exadata Database Machine target, the Exadata Cloud Service target includes only the following targets:

- Host
- Exadata Storage Server Grid
- Exadata Storage Server

Of these targets, only the Host targets, which represent the VMs, are customer monitoring responsibilities. The Exadata Storage Server and Exadata Storage Server Grid targets are included to provide visibility to the same performance and configuration information available in EM for on-premises Exadata storage servers.

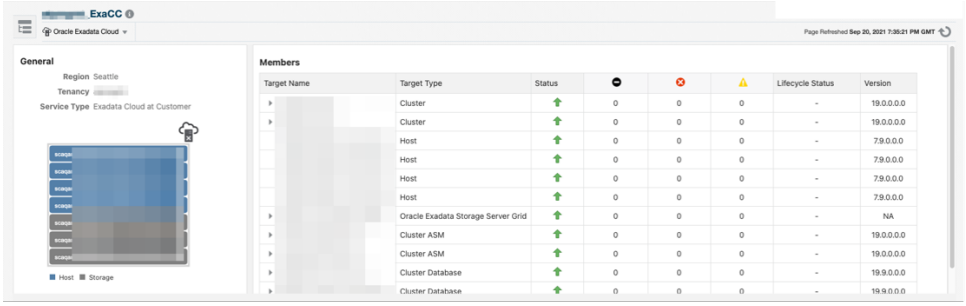

Image 2. Exadata Cloud target in Enterprise Manager.

For more information on monitoring Exadata Cloud with Enterprise Manager, see Oracle Enterprise Manager for Oracle Exadata Cloud.

Monitoring Exadata Performance

Enterprise Manager (EM) is tested and certified for each Exadata hardware generation and software feature release to ensure that incidents are raised for hardware and software faults and that EM provides the necessary monitoring and management capabilities. Each of the Exadata components can be monitored and managed using Enterprise Manager, and the Exadata Database Machine can be monitored as a system. Multiple health and availability metrics are enabled out of the box with recommended thresholds based on product expertise. In addition to basic health and availability, this paper will introduce the Key Performance Indicators (KPIs) recommended for advanced performance monitoring. In earlier versions of Enterprise Manager, some of these KPIs were added with Metric Extensions, however all the KPIs referenced in this paper are provided out of the box in 13.5 (and later releases of 13.4). This paper will guide you through enabling and setting thresholds for these KPIs.

Key Performance Indicators

A Key Performance Indicator (KPI) is a measurement used to define and evaluate successful operation. In the context of this technical brief KPIs are defined and used to evaluate if compute node, storage server and switch performance is within acceptable ranges.

A relatively simple generic example of a KPI would be CPU utilization. If a given system's performance degrades when CPU usage exceeds 95%, then 95% would be the critical threshold for the KPI. Many times, it is useful to have a threshold not only for when levels are critical but before they are critical as well. A warning threshold should be set so that administrators can be notified early enough to correct the issue before it becomes critical. For example, the warning threshold could be 90%. These are only example values to differentiate critical and warning thresholds.

Defining thresholds for metrics can be a challenging task. Threshold values for some metrics will be easily determined as opposed to others which may vary depending on workload, business rules, etc. Although some sample thresholds are provided in the following sections, it will be necessary to carefully evaluate every environment to setup the best possible monitoring.

In the CPU example above, it is relatively easy to set the thresholds because CPU usage is well understood and in most cases the threshold would be close to the same

between servers and environments. Unfortunately, many other KPIs are more difficult to define, especially KPIs that relate to I/O. Not only can the thresholds be difficult to define, but in some instances the KPIs themselves are hard to identify. Looking at an Exadata Storage Server, there can be over 3,000 Storage Server metrics. Sorting through the data and identifying which ones are important can be very challenging!

When looking at a system or subsystem holistically, often one KPI isn't enough to identify issues. Take for example a standard Oracle Linux server. What would be necessary to identify if the server is performing within specification? In the above example, CPU was identified as a KPI. However, there are other areas of server performance that would need to be considered such as memory, paging, disk, etc. The same holds true for the Exadata I/O subsystem. Although there are many metrics that are valid and important, no single metric can identify when the I/O system is at capacity.

For example, one indicator that is frequently considered is I/O per second (IOPS). IOPS shows the number of read and write operations to a disk. It might seem that this would be a finite number upon which it would be easy to base a threshold. However, the nature of the workload can affect the maximum number of IOPS a disk can perform. For example, a disk can perform far more small I/Os than large I/Os in a given period. Combining IOPS with other Storage Server metrics gives us a more comprehensive look at the environment.

When setting thresholds for I/O you can use two approaches. The first would be to leverage the published data sheets and evaluate your historical and current workload to determine appropriate warning and critical thresholds for relevant KPIs to alert when performance deviates from what is acceptable for your workload. This is a starting point and should be continually evaluated to see where the workload falls in these guidelines and how the application performs. Keep in mind that it is both possible to exceed the numbers on the data sheets and to encounter performance concerns for a given workload well below data sheet maximums. Metric history and Adaptive Thresholds functionality can be used to help inform these decisions. The second approach can build on the first by leveraging Adaptive Thresholds functionality to actively manage the thresholds, including for different periods of time, based upon deviation from standard workload during those periods. Refer to the MAA white paper Exadata Health and Resource Utilization Monitoring - Adaptive Thresholds for details on how to define thresholds accurately based on system workload and baselines. This may be a secondary step as the metrics will need to collect for a time before baselines can be established for Adaptive Thresholds.

Enterprise Manager Terms

Enterprise Manager provides extensive monitoring capabilities for Oracle Exadata. The following section will describe the basic concepts of monitoring an Exadata environment in Enterprise Manager. Prior to digging in deeply there are a few EM terms that should be defined. The best practice is to use the latest release and RU available. This document assumes that the Enterprise Manager environment used is version 13.5 RU9 or greater. While the KPIs are available in 13.4 RU18 and higher, the recommendation is to use the most current release and RU update.

- Agent – A process that runs on a host to monitor the status, health, and performance of all managed components (also referred to as targets) on that host. In Exadata, the agents run only on the compute nodes. The Storage Servers are monitored via the compute node agents.
- Plug-In – A Plug-in is a target specific UI supported by plug-in specific files (such as target definition files, collection scripts to collect metrics from targets). In the following discussion, the Exadata Plug-In will be used extensively. Plug-Ins are deployed to the OMS(s) and to the agent(s) requiring them.
- Release Update (RU) – A release update is a collection of OMS, Agent and Plug-in enhancements and fixes intended to be a complete set of updates. Release on a monthly or bi-monthly frequency.
- Target – A component monitored by Enterprise Manager through an agent.
- Metric – Measurement used to monitor target conditions or state.
- Threshold – A value defined for a metric. This value sets the level at which notifications or actions occur for the metric. There are two levels of thresholds: warning and critical.
- Occurrences – The number of consecutive collections that must exceed a threshold before a violation is triggered.

To learn more about Enterprise Monitoring review the Enterprise Manager Cloud Control Administrator's Guide.

Exadata Storage KPIs

Enterprise Manager will display availability, incidents, and performance information for each Exadata Storage Server target, and as an aggregated target at the Exadata Storage Server Grid level, which shows the overall health. The target home pages provide health status and summarize key performance information.

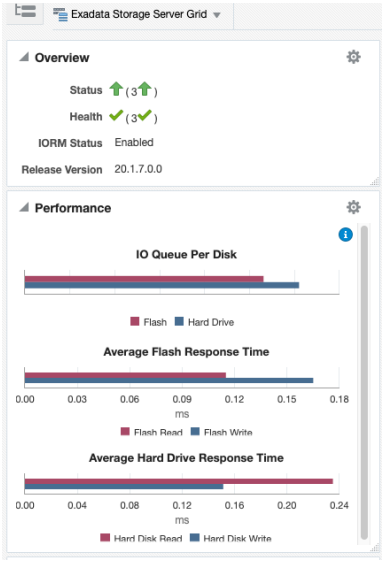

Image 3. Example of Exadata Storage Server Grid home page.

In addition to the EM availability and performance metrics for Exadata Storage Servers, there are a few key metrics that aggregate data to give an idea of whether the Exadata Storage Server infrastructure is at capacity.

Monitoring Exadata Storage Server capacity necessitates that all the Storage Servers are evaluated as a complete Exadata Storage Server Grid in addition to monitoring individual Storage Servers. The following 8 KPIs (4 flash disk and 4 hard disk) are available at both the Storage Server and Storage Server Grid levels and allow the storage servers to be evaluated individually and holistically together to determine the environment's health.

- Total Flash Disk IOPS – Aggregated total read and write IOPS of all flash disks
- Total Hard Disk IOPS – Aggregated total read and write IOPS of all hard disks
- Total Flash Disk Throughput – Aggregated total read and write throughput (MB/s) of all flash disks
- Total Hard Disk Throughput – Aggregated total read and write throughput (MB/s) of all hard disks
- Average Flash Disk Response Time – Average read and write latency (ms/request) across all flash disks
- Average Hard Disk Response Time – Average read and write latency (ms/request) across all hard disks
- Average Flash Disk IO Load – Average IO load (disk queue length) across all flash disks
- Average Hard Disk IO Load – Average IO load (disk queue length) across all hard disks

Collections are disabled by default and can be enabled by following the steps in the Configuring Metrics in Enterprise Manager section. Once metric collection is enabled, metric data should be monitored during peak usage to determine settings specific to the environment and updated with appropriate values. In environments with multiple Database Machines, thresholds cannot be assumed to be consistent between environments. Refer to Exadata Health and Resource Utilization Monitoring - Adaptive Thresholds for more advanced methods on setting thresholds. The warning and critical thresholds will vary depending on many factors including rack size, Exadata generation and version, application workload, etc. Adjusting thresholds may be necessary as different Exadata environments will have different thresholds. For example, IOPS thresholds for an X8M would be greater than for an X5 Database Machine. Using Enterprise Manager templates to specify settings for each DB Machine version may help standardize thresholds.

Health Exceptions

In addition to monitoring multiple targets, to complete the holistic approach to monitoring the Storage Servers and Storage Server Grid, many times one metric exceeding its threshold is not significant enough to conclude that one or more Storage Servers are having issues or are at capacity. By looking at multiple metrics and how they interact with each other, a more accurate picture starts to develop that allows for an accurate diagnosis and reduces the number of false alerts raised. The

Health Exceptions KPIs track the number of Flash or Hard Drive KPIs on the Storage Server that exceed their critical thresholds, and at the Storage Server Grid Level, they track the number of Storage Servers that have Flash or Hard Drive Health Exceptions:

- Flash Disk IO Health Exceptions - Number of members whose Exadata Key Performance Indicator Flash Disk IO Health Exceptions exceed their critical thresholds
- Hard Disk IO Health Exceptions - Number of members whose Exadata Key Performance Indicator Hard Disk IO Health Exceptions exceed their critical thresholds.

For these metrics to trigger, there must be thresholds set on the corresponding other 8 KPI metrics. As an example, setting Warning threshold to 2, and Critical threshold to 3 on the Flash Disk IO Health Exception at the Storage Server level would mean that if the Average Flash Disk Response Time and Total Flash Disk Throughput metrics on the Storage Server crossed their critical thresholds, a warning alert would trigger for the Flash Disk IO Health Exception metric for that Storage Server. If two Storage Servers triggered Health Exception alerts, and if the Warning threshold on the correlating Storage Server Grid Health Exception was set to 2, this would then alert at the Storage Server Grid level as well. Since these composite metrics are meant to summarize overall health, it is suggested that alerting for incidents be enabled for these metrics only. This will require customizing the out of box Incident Rule Set to exclude the Storage Server KPIs other than the Health Exceptions and include Warning and Critical events. If external communication (email, SNMP, ticketing, etc.) is required, ensure that the appropriate notifications are setup. Information on setting up notifications can be found in the Enterprise Manager Cloud Control Administrator's Guide.

Compute Node KPIs

Exadata on-premises can be configured as physical or virtual. In a virtual environment Enterprise Manager discovers the Virtual Platform, VM Instance, Host and ILOM targets and has different metrics depending on the target. Whether the Exadata is configured using KVM or Xen virtualization, it will have the same targets in Enterprise Manager. For simplification, throughout this document Hypervisor will refer to either the Xen dom0 or KVM Host and VM guest will refer to Xen domU or KVM Guest.

In addition, Exadata Cloud leverages a dual operational model in which access to Compute Nodes is segregated, with Oracle operations teams responsible for monitoring the physical infrastructure of the Exadata including the Compute Node ILOMs and Virtual Platform, and customers having access only to VM guest displayed as a Host target type.

The home page for each target will also display utilization specific to the physical or virtual target.

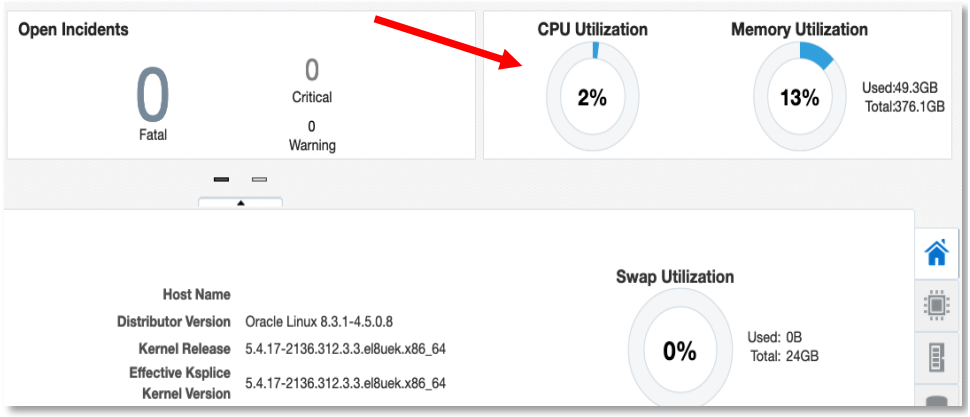

Image 4. Physical host home page

For the Virtual Platform target, the home page displays the Physical and VM Guest CPU performance information.

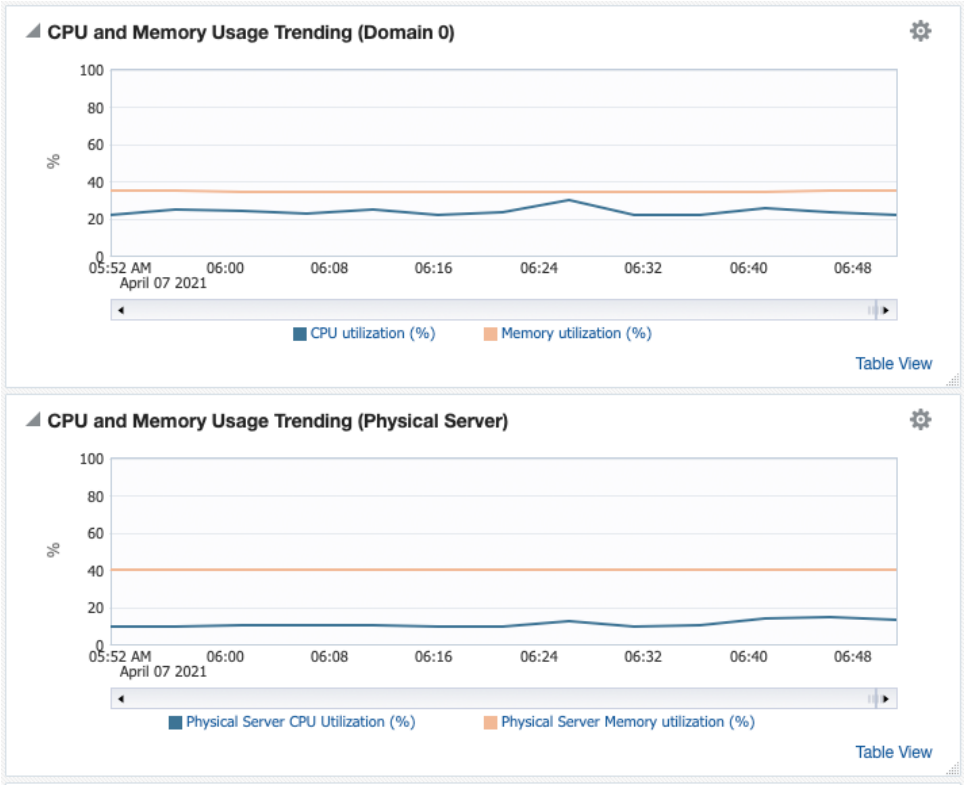

Image 5. Virtual Platform home page

From a customer perspective, the core Compute Node KPIs are the same whether the Compute Node is part of an on-premises Exadata configured as physical or virtual or whether the Compute Node is part of an Exadata Cloud Service or Exadata Cloud at Customer service.

Compute Nodes are monitored in much the same way as any Oracle Linux server. Below are the Compute Node KPIs that are recommended by the MAA team:

- CPU Utilization
- Memory Utilization
- Load Average

- Swap Utilization

These KPIs will provide a base set of monitoring that will indicate when server utilization is at or near capacity.

CPU Utilization

Common practice is to keep actual utilization below 85% for response time sensitive workloads. If utilization is over 85%, use IORM or instance caging to prioritize workloads appropriately and prevent CPU starvation of essential processes. It is common and appropriate to have CPU or I/O utilization reach 100% for large parallel workloads that seek maximum batch or reporting performance. Data Warehouses are an important example of this. In mixed-use systems, batch and reporting performance needs to be traded off against interactive response times. If interactive response time is important, then the degree of parallelism used for batch and reporting will need to be restricted so that batch and reporting don't consume excessive resources on the system. Setting the IORM objective to 'auto' provides protection for latency sensitive small IOs. IORM automatically switches into high throughput mode if the majority of IOs are from data warehouse workloads. This provides better throughput when only batch workloads are running. For more information on IORM refer to the Exadata System Software Users Guide.

CPU utilization is available in Enterprise Manager for Compute Node related targets for both physical and virtual Exadata on premise and virtual Exadata Cloud environments. In virtual environments, CPU Utilization for the Compute Node is measured both physically and virtually.

- In physical environments, CPU Utilization is available as part of the host target representing the Compute Node.
- In virtual environments, physical CPU Utilization and Hypervisor CPU Utilization is available as part of the Virtual Platform target. VM guest CPU Utilization is available as part of the host and Oracle VM Instance targets for each VM guest.
- In Exadata Cloud environments, CPU Utilization is available on the host.

CPU Metrics

ENVIRONMENT TYPE	COMPONENT	EM TARGET	METRIC SCOPE	METRIC
Physical	Compute Node	Host	Physical	Load > CPU Utilization %
Virtual	Compute Node	Oracle Virtual Platform	Physical	Server Load > Physical Server CPU Utilization %
Virtual	Hypervisor	Oracle Virtual Platform	Virtual	Dom0 Load > CPU Utilization %
Virtual	VM Guest	Host	Virtual	Load > CPU Utilization%
Virtual	VM Guest	VM Instance	Virtual	Load > CPU Utilization of the virtual machine (%)
Exadata Cloud	VM Guest	Host	Virtual	Load > CPU Utilization %

Table 1. Compute Node CPU KPI Metrics.

Memory

Memory usage is dependent on individual system characteristics, including how many databases, memory allocated, hugepages and more. Identifying the memory utilization baseline before setting thresholds will help reduce false alerts. Linux will free memory from cache when necessary, so cache can be regarded as part of free memory. The Exadata databases do not use the Linux page cache for database I/Os, so a relatively small Linux page cache is needed.

Memory utilization is available in Enterprise Manager for Compute Node related targets for both physical and virtual Exadata on premise and virtual Exadata Cloud environments. In virtual environments, Memory Utilization for the Compute Node is measured both physically and virtually.

- In physical environments, Memory Utilization is available as part of the host target representing the Compute Node.
- In virtual environments, physical and Hypervisor Memory Utilization is available as part of the Virtual Platform target. Virtual Memory Utilization is available on the host targets for each guest.
- For Exadata Cloud environments, Memory Utilization is available on the host.

It is important to note that in virtual environments, the Physical Server Memory Utilization % metric on the Oracle Virtual Platform considers any memory allocated to the VM Guest as used. The actual memory utilization should be evaluated within the Host target.

Memory Metrics

ENVIRONMENT TYPE	COMPONENT	EM TARGET	METRIC SCOPE	METRIC
Physical	Compute Node	Host	Physical	Load > Memory Utilization %
Virtual	Compute Node	Oracle Virtual Platform	Physical	Server Load > Physical Server Memory Utilization %
Virtual	Hypervisor	Oracle Virtual Platform	Virtual	Dom0 Load > Memory Utilization %
Virtual	VM Guest	Host	Virtual	Load > Memory Utilization%
Exadata Cloud	VM Guest	Host	Virtual	Load > Memory Utilization%

Table 2. Compute Node Memory KPI Metrics

Swap Utilization

Swap is the process in which inactive pages of memory are moved to disk in the event the amount of physical memory is full. While Exadata systems contain a large amount of memory, it is still important to avoid swapping. If you observe consistent, excessive swapping, memory should be evaluated.

Swap Utilization is available in Enterprise Manager for Compute Node related targets for both physical and virtual Exadata on premise and virtual Exadata Cloud environments.

- In physical environments, Swap Utilization is available as part of the host target representing the Compute Node.
- In virtual environments, the Virtual Platform target for the Hypervisor and on the Host targets for each guest.
- In Exadata Cloud environments, Swap Utilization is available on the Host target.

Swap Metrics

ENVIRONMENT TYPE	COMPONENT	EM TARGET	METRIC SCOPE	METRIC
Physical	Compute Node	Host	Physical	Load > Swap Utilization (%)
Virtual	Hypervisor	Oracle Virtual Platform	Virtual	Dom0 Load > Swap space utilization (%)
Virtual	VM Guest	Host	Virtual	Load > Swap Utilization (%)
Exadata Cloud	VM Guest	Host	Virtual	Load > Swap Utilization (%)

Table 3. Compute Node Swap KPI Metrics

Load Average

Load Average is a measure of processes that are waiting (either in a running, runnable or uninterruptible state) for CPU time. The point at which Load Average becomes concerning is dependent on the available CPUs on the systems. For example, a Load Average of 80 on a 64-core server would result in a process count of 1.25 per core which would be acceptable. However, a Load Average of 80 on a 12-core server would result in a process count of 6.67 per core which would be indicative of a problem. Load Average isn't necessarily caused by a lack of CPU resources but could also be an indicator of other issues such as a saturated or stalled I/O subsystem.

Enterprise Manager makes understanding Load Average easier by taking the overall Load Average and dividing by the number of CPUs. This gives the average number of runnable processes per CPU.

Load Average is available in Enterprise Manager for Compute Node related targets for both physical and virtual Exadata on-premises and virtual Exadata Cloud environments.

- In physical environments, Load Average is available as part of the host target representing the Compute Node.
- In virtual environments, Load Average is available as part of the Virtual Platform target for the Hypervisor and on the Host targets for each guest.
- In Exadata Cloud environments, Load Average is available on the Host target.

Load Utilization

ENVIRONMENT TYPE	COMPONENT	EM TARGET	METRIC SCOPE	METRIC
Physical	Compute Node	Host	Physical	Load > Run Queue Length (15 minute average, per CPU)
Virtual	Hypervisor	Oracle Virtual Platform	Virtual	Dom0 Load > Run Queue Length (15 minute average, per core)
Virtual	VM Guest	Host	Virtual	Load > Run Queue Length (15 minute average, per CPU)
Exadata Cloud	VM Guest	Host	Virtual	Load > Run Queue Length (15 minute average, per CPU)

Table 4. Compute Node Load KPI Metrics

Systems Infrastructure Switch KPIs

Depending on the Exadata version, Enterprise Manager provides detailed information about performance and throughput metrics on the RoCE or InfiniBand switches under the Systems Infrastructure Switch target. X8M and higher will have RoCE switches and metrics, while X8 and older systems will have InfiniBand switches and metrics. In additional to the throughput and performance metrics, there are additional key performance metrics that help get the best possible insight into potential problems that could occur on the switch. These metrics include CPU, memory, storage and SSH sessions.

- **CPU Utilization** - Monitors the average CPU utilization (user and kernel). Ideally, CPU usage should be below 80%.
- **Memory Utilization** - Measures the free memory on the switch.
- **Root Filesystem Usage** - Provides the percent of free filesystem space on the switch.
- **SSH Session Count** - Number of SSH sessions logged into the switch. These sessions are usually administrative connections for monitoring, configuration, etc. On InfiniBand switches where this is of particular concern, more than 10-12 sessions can create issues due to either PID limits or CPU/memory constraints.

These metrics are disabled out of the box and will need to be enabled for collections. Note that these metrics are not relevant on Exadata Cloud targets as the switches are not monitored by customers.

Switch Metrics

ENVIRONMENT TYPE	COMPONENT	EM TARGET	METRIC
Physical/Virtual	Switch	Systems Infrastructure Switch	CPU Utilization > Kernel CPU Utilization (%)
Physical/Virtual	Switch	Systems Infrastructure Switch	CPU Utilization > User CPU Utilization (%)
Physical/Virtual	Switch	Systems Infrastructure Switch	Memory Utilization > Free Memory (%)
Physical/Virtual	Switch	Systems Infrastructure Switch	Storage Summary > Used Space (%)
Physical/Virtual	Switch	Systems Infrastructure Switch	SSH Session Count > Number of SSH Sessions

Table 5. Switch KPI Metrics

Configuring Metrics in Enterprise Manager

Enterprise Manager allows you to customize each target with specific metric collections, frequencies, thresholds, and actions. Typically, this is done in a template and then mass applied to a group of targets. Details on how to customize metrics at the target level are shown in the section below. For more detailed instructions on monitoring targets from Enterprise Manager, please see the Cloud Control Administrator's Guide.

Enabling Metrics

While the metrics are available out of the box, some of the Exadata Storage Server, Exadata Storage Server Grid and Systems Infrastructure Switch KPIs discussed in this paper are disabled by default as threshold and alerting considerations need to be evaluated before starting to send notifications.

To begin collecting these metrics, enable the metric collection from the All Metrics page. For example, from the Exadata Storage Server target menu, select Monitoring > All Metrics. Select the metric group and click on the Modify (pencil icon) for the Collection Schedule. In the pop-up window, click on Enabled and select OK to save your changes.

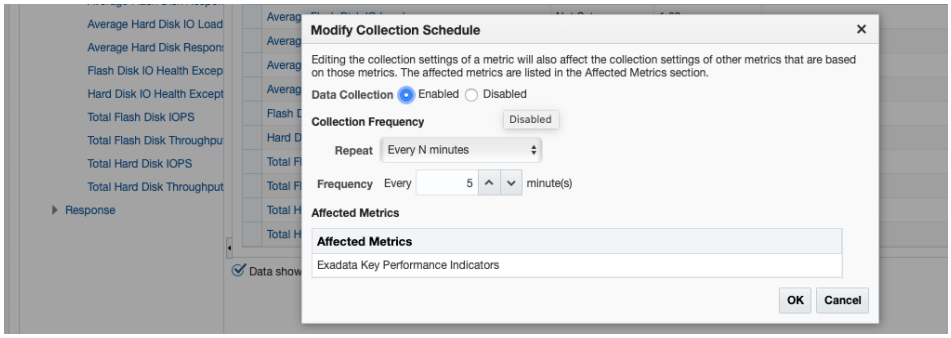

Image 6. Enabling metric collection

Change Collection Frequency

If more granular data is required, the frequency can be adjusted to a lower value in the target's metric collection setting page. Select Monitoring > All Metrics, click the Collection Schedule of the metric group, and adjust the minutes as seen in the figure below. Be careful that for correlation, related metrics need to be collected at the same frequency, for example KPI metrics for the Exadata Storage Server and Exadata Storage Server Grid should be set to the same frequency as the Exadata Cell Disk metrics that they aggregate. In addition, the **Storage Server - Aggregated Exadata FlashDisk and HardDisk Metric** and **Storage Server Grid – Key Performance Indicators** collections need to be set to the same collection schedule to ensure proper rollup. Changing the collection schedule for the **Storage Server - Aggregated Exadata FlashDisk and HardDisk Metric** will also change additional metrics to collect at the same interval as they are all part of the same collection script.

Be aware that increasing the collection frequency faster than the default values can place additional load on the agent deployed on the Exadata environment as well as the Enterprise Manager Repository. Setting this value to anything under 5 minutes is not recommended. Always test these changes in a non-production environment.

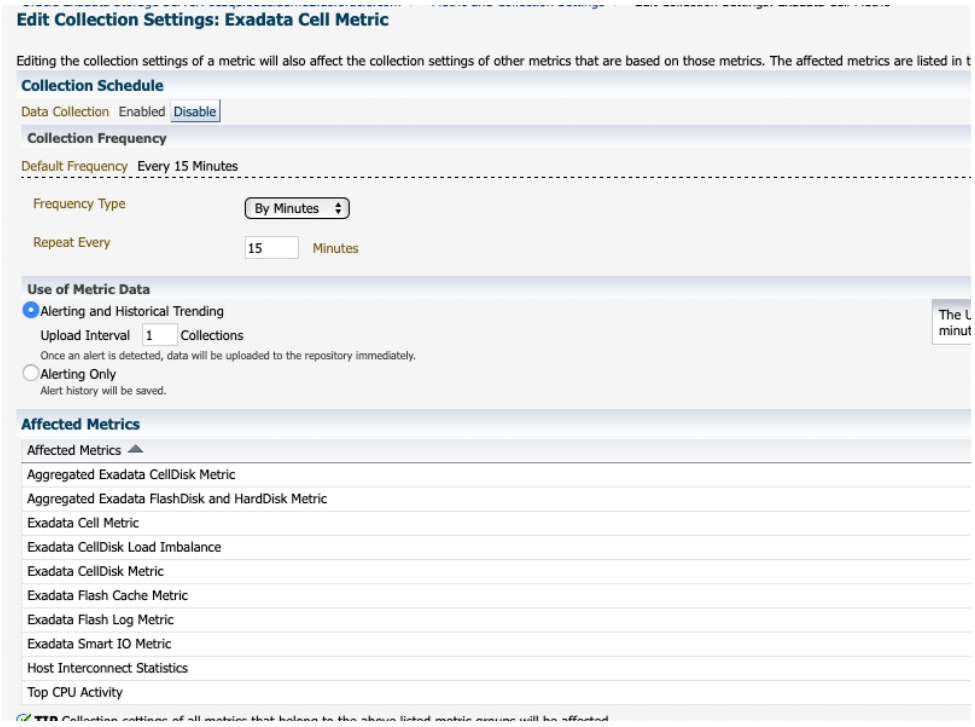

Image 7. Changing collection frequency

Setting Thresholds

To set the metric threshold, navigate from the target home page for the metric referenced in this paper. For example, for the host target, to set the metric threshold for CPU Utilization select Host > Monitoring > All Metrics

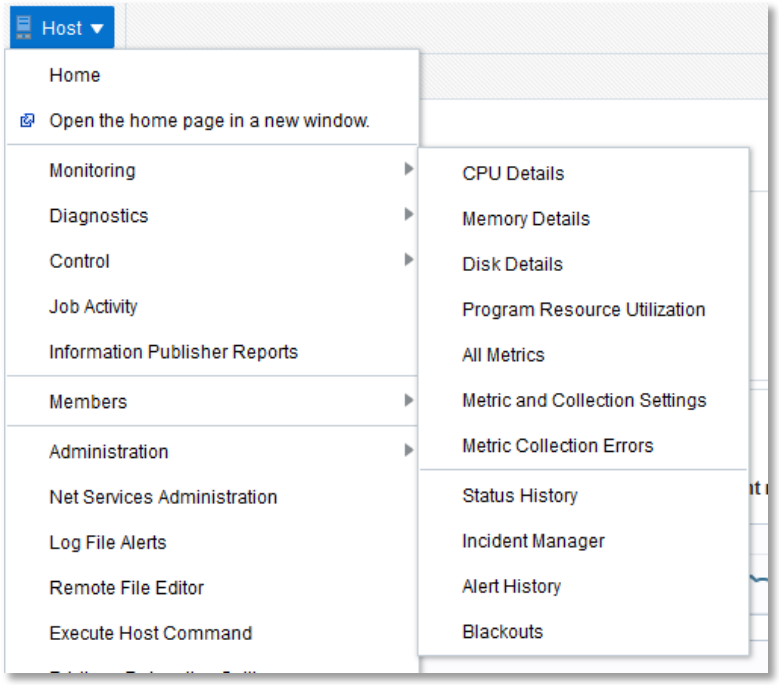

Image 8. Metric and Collection Settings menu.

Find the metric group, expand to display, and select the individual metric. In this example, the metric group is Load and the metric is CPU Utilization (%).

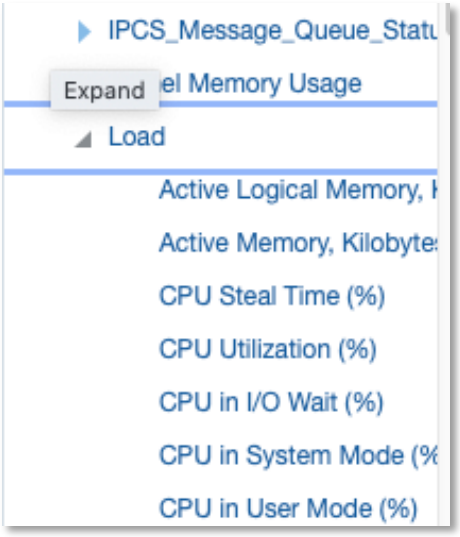

Image 9. Metric details.

Use the Modify Thresholds button to select the warning and critical thresholds. Select how many of the occurrences of the threshold being violated need to occur before an incident is created to reduce the occurrence of alerts during isolated spikes. Press Save Thresholds to save the changes.

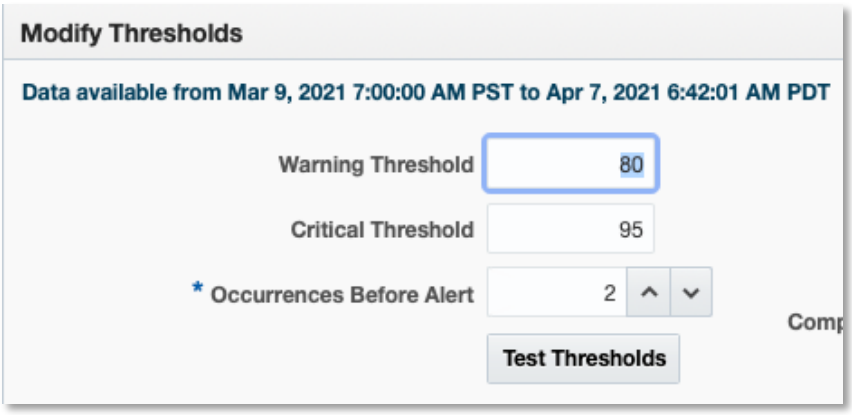

Image 10. Adjust thresholds and number of occurrences.

Threshold Suggestions

Enterprise Manager can be used to evaluate the past performance of a metric and visualize threshold violations that would have occurred in the past.

From the target menu, go to Monitoring > Metric and Collection Settings, then click the Edit icon for that metric. Expand the section Threshold Suggestion to see the metric history.

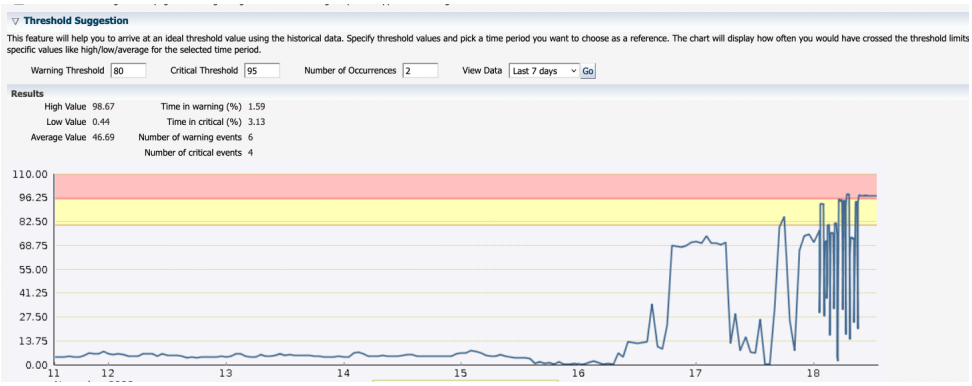

Image 11. Viewing threshold suggestions.

Add To Templates

For mass deployment it's best to use the Enterprise Manager Templates. With the KPIs you will need to edit or create a template and add the new metrics from a sample target that you have previously enabled the metrics on. This will enable and the designated collection frequency into the template and will then enable against any targets you deploy to.

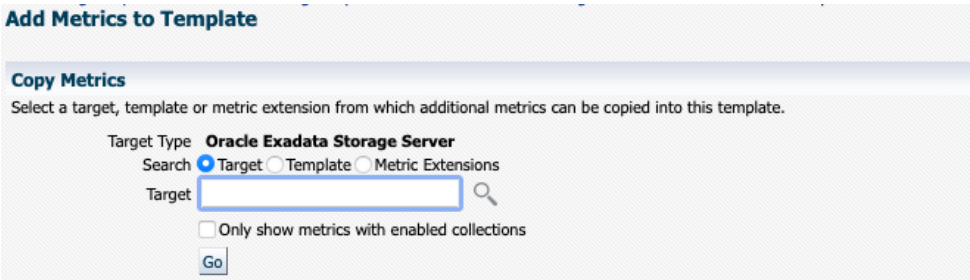

Image 12. Copy Metrics from Target.

Conclusion

With the use of these metrics, Exadata administrators have an overall picture in Enterprise Manager that enables evaluation of the state of the Exadata Database Machine. By setting up appropriate alerts and thresholds, administrators will also be proactively notified of potential issues before they impact business service level agreements.

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle