
MarkLogic Server

Monitoring MarkLogic Guide

MarkLogic 9
May, 2017

Last Revised: 9.0-5, May 2018

Table of Contents

Monitoring MarkLogic Guide

1.0	Monitoring MarkLogic Server	5
1.1	Overview	5
1.2	Selecting a Monitoring Tool	6
1.3	Monitoring Architecture, a High-level View	7
1.4	Monitoring Tools and Security	7
1.5	Guidelines for Configuring your Monitoring Tools	8
1.5.1	Establish a Performance Baseline	8
1.5.2	Balance Completeness Against Performance	8
1.6	Monitoring Metrics of Interest to MarkLogic Server	9
1.6.1	Does MarkLogic Have Adequate Resources?	9
1.6.2	What is the State of the System Overall?	10
1.6.3	What is Happening on the MarkLogic Server Cluster Now?	11
1.6.4	Are There Signs of a Serious Problem?	13
2.0	Using the MarkLogic Server Monitoring Dashboard	14
2.1	Terms Used in this Chapter	14
2.2	Displaying the Monitoring Dashboard	15
2.3	Monitoring Specific Resources	15
2.4	Monitoring Dashboard Sessions	16
2.5	Setting the Sample Interval	16
2.6	Viewing Monitoring Sample Details	17
2.7	Monitoring Query Execution	17
2.8	Monitoring Rates and Loads	19
2.8.1	Overview	20
2.8.2	XDQP Communication	21
2.8.3	Backup/Restore	24
2.9	Monitoring Disk Space	25
2.10	Exporting Monitoring Data	28
3.0	MarkLogic Server Monitoring History	29
3.1	Overview	29
3.2	Enabling Monitoring History on a Group	30
3.3	Setting the Monitoring History Data Retention Policy	32
3.4	Viewing Monitoring History	33
3.5	Viewing Monitoring History by Time Span and Frequency	36
3.6	Labeling Monitoring History Time Spans	38
3.7	Filtering Monitoring History by Resources	41
3.8	Historical Performance Charts by Resource	44

3.8.1	Disk Performance Data	45
3.8.2	CPU Performance Data	49
3.8.3	Memory Performance Data	51
3.8.4	XDQP Server Requests Performance Data	54
3.8.5	Server Performance Data	55
3.8.6	Network Performance Data	58
3.8.7	Database Performance Data	61
3.9	Exporting and Printing Monitoring History	66
4.0	Telemetry	67
4.1	Understanding Telemetry	67
4.2	Configure Telemetry in the Admin UI	68
4.2.1	Enable Telemetry on the Group Configuration Page	68
4.3	Example—Telemetry	70
4.3.1	View Staged Telemetry Files	70
4.3.2	Encryption of Staged Files	71
4.4	Telemetry on the Support Page	72
4.5	Configure Telemetry With XQuery	73
4.6	Baseline System Information	74
4.6.1	Metering Data	74
4.7	Upload a Support Request to Support	75
4.8	APIs for Telemetry	76
4.8.1	Admin APIs	76
4.8.2	REST Management APIs for Telemetry	76
4.9	Interactions With Other MarkLogic Features	77
4.9.1	Encryption at Rest	77
4.9.2	Rolling Upgrades	77
4.9.3	Support Uploads	77
5.0	Using the Management API	78
5.1	Terms used in this Chapter	78
5.2	Overview of the Management API	80
5.3	Security	80
5.4	Management API Requires Writing to the App-Services Database	80
5.5	Resource Addresses	81
5.6	Obtaining the Options Node for a Resource Address	82
5.7	Specifying the Management API Version	83
5.8	Specifying Parameters in a Resource Address	83
5.8.1	Formatting the Monitor Results	84
5.9	Interpreting the Output	85
6.0	Technical Support	87
7.0	Copyright	88

7.0 COPYRIGHT88

1.0 Monitoring MarkLogic Server

MarkLogic Server provides a rich set of monitoring features that include a pre-configured monitoring dashboard and a Management API that allows you to integrate MarkLogic Server with existing monitoring applications or create your own custom monitoring applications.

This chapter includes the following sections:

- [Overview](#)
- [Selecting a Monitoring Tool](#)
- [Monitoring Architecture, a High-level View](#)
- [Monitoring Tools and Security](#)
- [Guidelines for Configuring your Monitoring Tools](#)
- [Monitoring Metrics of Interest to MarkLogic Server](#)

1.1 Overview

In general, you will use a monitoring tool for the following:

- To keep track of the day-to-day operations of your MarkLogic Server environment.
- For initial capacity planning and fine-tuning your MarkLogic Server environment. For details on how to configure your MarkLogic Server cluster, see the *Scalability, Availability, and Failover Guide*.
- To troubleshoot application performance problems. For details on how to troubleshoot and resolve performance issues, see the *Query Performance and Tuning Guide*.
- To troubleshoot application errors and failures.

The monitoring metrics and thresholds of interest will vary depending on your specific hardware/software environment and configuration of your MarkLogic Server cluster. This chapter lists some of the metrics of interest when configuring and troubleshooting MarkLogic Server. However, MarkLogic Server is just one part of your overall environment. The health of your cluster depends on the health of the underlying infrastructure, such as network bandwidth, disk I/O, memory, and CPU.

1.2 Selecting a Monitoring Tool

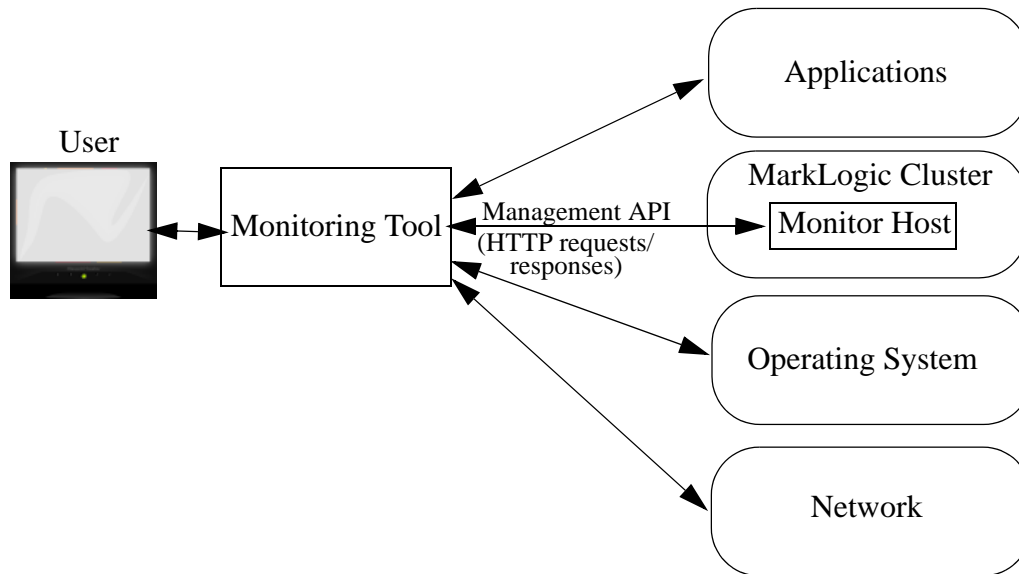
Though this guide focuses on the tools available from MarkLogic that enable you to monitor MarkLogic Server, it is strongly recommended that you select an enterprise-class monitoring tool that monitors your entire computing environment to gather application, operating system, and network metrics alongside MarkLogic Server metrics.

There are many monitoring tools on the market that have key features such as alerting, trending, and log analysis to help you monitor your entire environment. MarkLogic Server includes the following monitoring tools:

- A Monitoring dashboard that monitors MarkLogic Server. This dashboard is pre-configured to monitor specific MarkLogic Server metrics. For details, see “Using the MarkLogic Server Monitoring Dashboard” on page 14.
- A Monitoring History dashboard to capture and make use of historical performance data for a MarkLogic cluster. For details, see “MarkLogic Server Monitoring History” on page 29.
- A RESTful Management API that you can use to integrate MarkLogic Server with existing monitoring application or create your own custom monitoring applications. For details, see “Using the Management API” on page 78.

1.3 Monitoring Architecture, a High-level View

All monitoring tools use a RESTful Management API to communicate with MarkLogic Server. The monitoring tool sends HTTP requests to a monitor host in a MarkLogic cluster. The MarkLogic monitor host gathers the requested information from the cluster and returns it in the form of an HTTP response to the monitoring tool. The Management API is described in “Using the Management API” on page 78.



1.4 Monitoring Tools and Security

To gain access to the monitoring features described in this guide, a user must be assigned the `manage-user` role. Monitoring tools should authenticate as a user with that role. The `manage-user` role is assigned the `http://marklogic.com/xdmp/privileges/manage` execute privilege and provides access to the Management API, Manage App Server, and the UI for the Configuration Manager and Monitoring Dashboard. The `manage-user` role also provides read-only access to all of a cluster's configuration and status information, with the exception of the security settings. For details on assigning roles to users, see [Users](#) in the *Administrator's Guide*.

If you have enabled SSL on the `Manage` App Server, your URLs must start with HTTPS, rather than HTTP. Additionally, you must have a MarkLogic certificate on your browser, as described in [Accessing an SSL-Enabled Server from a Browser or WebDAV Client](#) in the *Security Guide*.

1.5 Guidelines for Configuring your Monitoring Tools

Monitoring tools enable you to set thresholds on specific metrics to alert you when a metric exceeds a pre-specified value.

The topics in this section are:

- [Establish a Performance Baseline](#)
- [Balance Completeness Against Performance](#)

1.5.1 Establish a Performance Baseline

Many metrics that can help in alerting and troubleshooting are meaningful only if you understand normal patterns of performance. For example, monitoring an App Server for slow queries will require a different threshold on an application that spawns many long-running queries to the task server than on an HTTP App Server where queries are normally in the 100 ms range. Most enterprise-class monitoring tools support data storage to support this type of trend analysis. Developing a starting baseline and tuning it if your application profile changes will yield better results for developing your monitoring strategy.

1.5.2 Balance Completeness Against Performance

Collecting and storing monitoring metrics has a performance cost, so you need to balance completeness of desired performance metrics against their cost. The cost of collecting monitoring metrics can differ. In general, the more resources you monitor, the greater the cost. For example, if you have a lot of hosts, server status is going to be more expensive. If you have a lot of forests, database status is going to be more expensive. In most cases, you will use a subset of the available monitoring metrics. And there may be circumstances in which you temporarily monitor certain metrics and, once the issue have been targeted and resolved, you no longer monitor those metrics.

One balancing technique is to measure system performance on a staging environment under heavy load, then enable your monitoring tool and calculate the overhead. You can reduce overhead by reducing collection frequency, reducing the number of metrics collected, or writing a Management API plugin to produce a custom view that pinpoints the specific metrics of interest. Each response from the underlying Management API includes an elapsed time value to help you calculate the relative cost of each response. For details, see “Using the Management API” on page 78.

1.6 Monitoring Metrics of Interest to MarkLogic Server

Environments and workloads vary. Each environment will have a unique set of requirements based on variables including cluster configuration, hardware, operating system, patterns of queries and updates, feature sets, and other system components. For example, if replication is not configured in your environment, you can remove templates or policies that monitor that feature.

This section provides a set of guiding questions to help you understand and identify the relevant metrics. The topics in this section are:

- [Does MarkLogic Have Adequate Resources?](#)
- [What is the State of the System Overall?](#)
- [What is Happening on the MarkLogic Server Cluster Now?](#)
- [Are There Signs of a Serious Problem?](#)

1.6.1 Does MarkLogic Have Adequate Resources?

MarkLogic Server is designed to fully utilize system resources. Many settings, such as cache sizes, are auto-sized by MarkLogic Server at installation.

Some questions to ask are:

- Does MarkLogic Server have enough resources on the host machine? What processes other than MarkLogic Server are running on the host and what host resources do those processes require? When competing with other processes, MarkLogic Server cannot optimize resource utilization and consequently cannot optimize performance.
- Is there enough disk space for forest data and merges? Merges require at least one and one half times as much free disk space as used by the forest data (for details, see [Memory, Disk Space, and Swap Space Requirements](#) in the *Installation Guide*). If a merge runs out of disk space, it will fail.
- Is there enough disk space to log system activity? If there is no space left on the log file device, MarkLogic Server will abort. Also, if there is no disk space available to add messages to the log files, MarkLogic Server will fail to start.
- Is there enough memory for the range indexes? Range indexes improve performance at the cost of memory and increased load/reindex time. Running out of memory for range indexes may result in undesirable memory swapping that severely impacts performance.
- Is swap space configured correctly? At query time, MarkLogic Server makes use of both memory and swap space. If there is not enough of either, the query can fail with SVC-MEMALLOC messages. For details on configuring swap memory, see [Tuning Query Performance in MarkLogic Server](#) in the *Query Performance and Tuning Guide*.
- How many hosts are in the cluster? How are the hosts configured as evaluator and data nodes? How are the hosts organized into groups? For details on configuring MarkLogic

Server clusters, see [Clustering in MarkLogic Server](#) in the *Scalability, Availability, and Failover Guide*.

- What applications use resource-intensive features, such as CPF, replication, and point-in-time recovery? Are the hardware, software, and network resources available and configured to most efficiently support such applications?

1.6.2 What is the State of the System Overall?

Many problems that impact MarkLogic Server originate outside of MarkLogic Server. Consider the health of your overall environment.

Some questions to ask are:

- How efficiently is CPU being used? How much CPU capacity exists at different time slices? What is the execution speed of the current read and write tasks? Can I optimize queries or choose a better time to batch load?
- How efficiently is I/O being used? What amount of data is currently being read from or written to disk? Are there any I/O bottlenecks?
- Is there enough free disk space for each file system?
- Are there any errors or warnings appearing in the logs for the operating system, MarkLogic Server, and applications?
- What is the current state of the network?
- Are there any serious errors in the system log files? Your monitor tool, or an auxiliary tool such as Splunk, should monitor your system logs and report on any detected errors.

1.6.3 What is Happening on the MarkLogic Server Cluster Now?

When you suspect an error or performance problem originates from MarkLogic Server, some questions to ask are:

- Are all of the hosts in the cluster online? Are all of the App Servers enabled? In what states are the forests?
- What are the patterns of queries and updates? Do they appear to be evenly distributed across the hosts in the cluster?
- Are there any long-running queries? Longer than usual query execution times may indicate a bottleneck, such as a slow host or problems with XDQP communication between hosts. Other possible problems include increased loads following a failover or more than the usual number of total requests.
- Is there an increase in the number of outstanding requests? A consistent increase in the total number of outstanding requests may indicate the need to add more capacity and/or load balance. Decreases in total requests may indicate some “upstream” problem that needs to be addressed.
- What is the I/O rates and loads pattern? In this context, *rates* refers to amount of data applications are currently reading from or writing to MarkLogic Server databases (throughput) and *loads* refers to the execution time of the current read and write requests, which includes the time requests spend in the wait queue when maximum throughput is achieved.

Under normal circumstances you will see loads go up as rates go up. As the workload (number of queries and updates) increases, a steadily high rates value indicates the maximum database throughput has been achieved. When this occurs, you can expect to see increasing loads, which reflect the additional time requests are spending in the wait queue. As the workload decreases, you can expect to see decreasing loads, which reflect fewer requests in the wait queue.

If, while the workload is steady, rates decrease and loads increase, something is probably taking away I/O bandwidth from the database. This may indicate that MarkLogic Server has started a background task, such as a merge operation or some process outside of MarkLogic Server is taking away I/O bandwidth.

- What is the journal and save write rates and loads pattern? During a merge, you should see the rates for journal and save writes decrease and the loads increase. Once the merge is done, journal and save writes rates should increase and the loads should decrease. If no merge is taking place, then a process outside of MarkLogic Server may be taking away I/O bandwidth.
- What is the XDQP rates and loads pattern? In this context, *rates* refers to amount of data hosts are currently reading from or writing to other hosts and *loads* refers to the execution time of the current read and write requests, including those in the wait queue. A decrease in rates and an increase in loads may indicate that there is network problem.

- What are the cache hit/miss rates? Lots of cache hits means not having to read fragments off disk, so there is less I/O load. An increasing cache miss rate may indicate a need to increase the cache size, write queries that take advantage of indexes to reduce the frequency of disk reads, or adjust the fragment size to better match that of the queried data.
- How many concurrent updates and reads are in progress? An increase of both updates and reads may indicate that there are queries that are doing too many updates and reads concurrently. The potential problem is lock contention between the updates and reads on the same fragments, which degrades performance.
- How many database merges are in progress? Merges require both I/O and disk resources. If too many database merges are taking place at the same time, it may be necessary to coordinate merges by creating a merge policy or establishing merge blackout periods, as described in [Understanding and Controlling Database Merges](#) in the *Administrator's Guide*.
- How many reindexes are in progress? Database reindexing is periodically done automatically in the background by MarkLogic Server and requires both CPU and disk resources. If there are too many reindexing processes going on at the same time, you may need to adjust when reindexing is done for particular databases, as described in [Text Indexing](#) in the *Administrator's Guide*.
- How many backups and/or restores are in progress? Backup and restore processes can impact the performance of applications and other background tasks in MarkLogic Server, such as merges and indexing. Backups with point-in-time recovery enabled have an even greater impact on performance. If backup and/or restore processes are impacting system performance, it may be necessary to reschedule them, as described in [Backing Up and Restoring a Database](#) in the *Administrator's Guide*.

1.6.4 Are There Signs of a Serious Problem?

If you are encountering a serious problem in which MarkLogic Server is unable to effectively service your applications, some questions to ask are:

- Did MarkLogic Server abort or fail to start? This may indicate that there not enough disk space for the log files on the log file device. If this is the cause, you will need to either add more disk space or free up enough disk space for the log files.
- Is an application unable to update data in MarkLogic Server? This may indicate that you have exceeded the 64-stand limit for a forest. This could be the result of running out of merge space or that merges are suppressed.
- Are queries failing with SVC-MEMALLOC messages? This indicates that there is not enough memory or swap space. You may need to add memory or reconfigure your swap memory, as described in [Tuning Query Performance in MarkLogic Server](#) in the *Query Performance and Tuning Guide*
- Are there any forests in the async replicating state? This state indicates that a primary forest is asynchronously catching up to its replica forest after a failover or that a new replica forest was added to a primary forest that already contains content. If a forest has failed over, see [Scenarios that Cause a Forest to Fail Over](#) in the *Scalability, Availability, and Failover Guide* for possible causes.
- Are there any serious messages in the error logs? The various log levels are described in [Understanding the Log Levels](#) in the *Administrator's Guide*. All log messages at the error level and higher should be investigated, whereas lower-level messages, such as warnings and debug messages are mostly informational. Log messages that indicate a particularly serious problem include:
 - Repeated server restart messages. Possible causes include a corrupted forest, segmentation faults, or some problem with the host's operating system.
 - XDQP disconnect. Possible causes include an XDQP timeout or a network failure.
 - Forest unmounted. Possible causes include the forest is disabled, it has run out of merge space, or the forest data is corrupted.
 - SVC-* errors. These are system-level errors that result from timeouts, socket connect issues, lack of memory, and so on.
 - XDMP-BAD errors. These indicate serious internal error conditions that shouldn't happen. Look at the error text for details and the logs for context and contact MarkLogic Support.

2.0 Using the MarkLogic Server Monitoring Dashboard

This chapter describes how to use the Monitoring Dashboard. The Monitoring Dashboard provides task-based views of MarkLogic Server performance metrics in real time. The Monitoring Dashboard is intended to be used alongside the status pages in the Admin Interface and other monitoring tools that monitor application and operating system performance metrics.

The topics in this chapter are:

- [Terms Used in this Chapter](#)
- [Displaying the Monitoring Dashboard](#)
- [Monitoring Specific Resources](#)
- [Monitoring Dashboard Sessions](#)
- [Setting the Sample Interval](#)
- [Viewing Monitoring Sample Details](#)
- [Monitoring Query Execution](#)
- [Monitoring Rates and Loads](#)
- [Monitoring Disk Space](#)
- [Exporting Monitoring Data](#)

2.1 Terms Used in this Chapter

The following terms are used in this chapter:

- A *Monitoring Session* is the timeframe since the dashboard page was last refreshed. For example, if you navigate from the Query Execution page to the Rates and Loads page, you have ended the Query Execution session and started the Rates and Loads session.
- A *Monitoring Sample* is a bit of information captured during a refresh interval on a graph. For example, one of the candlesticks captured in the Query Execution graph is a single sample.

2.2 Displaying the Monitoring Dashboard


You can display the Monitoring Dashboard by doing the following:

1. Open a browser and enter the URL:


`http://monitor-host:8002/`

where *monitor-host* is a host in the cluster you want to monitor

2. At the top of the page, click on Monitoring and click on Dashboard in the pull-down menu:




3. The Monitoring Dashboard page appears. From the default Monitoring Dashboard page, you can navigate to any of the pages described in this chapter.



2.3 Monitoring Specific Resources

By default the Monitoring Dashboard monitors the entire cluster. You can use the Search box to select a specific resource to monitor. Clicking on the search field produces a pull-down menu in which you can locate the resource. Alternatively, you can directly locate a resource by entering the name of the resource in the search field.



2.4 Monitoring Dashboard Sessions

Each time you navigate to a new Dashboard page, you end the current monitoring session and begin a new one. The monitoring data from the previous session is lost from that point on. If you want to maintain multiple Dashboard sessions, you can open each page in a separate browser tab or window.

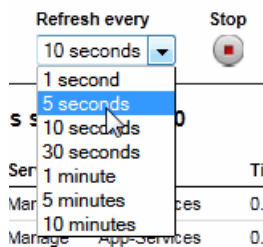
You can freeze the monitoring data for a Dashboard page by clicking on the Stop button in the upper right-hand portion of the page and restart the data by pressing Start. When you stop a page, you will lose any monitoring data between the time the page is stopped and the time it is restarted. If you have multiple Dashboard pages open, the sessions continue on the other pages; so stopping the monitoring data on one page will not stop the data on the other pages. When you start the stopped page, its session will resume at the current timestamp.



2.5 Setting the Sample Interval

The sample interval specifies the frequency in which the selected resource is monitored. By default, the sample interval is every 10 seconds. Use the Refresh pull-down menu to set the sample interval from anything between once every 1 second to every 10 minutes.


If you have multiple Dashboard pages open in separate tabs or windows, changing the sample interval on one page will not change the interval on the other pages. However, if you switch between pages in the same browser tab or window, the interval will be the same for all pages.



2.6 Viewing Monitoring Sample Details

You can hover your mouse on any monitoring sample to view the details of the sample. For example, to view the details of a query execution sample, hover on the bar graphic as shown below.

Query Execution




2.7 Monitoring Query Execution

Query execution data gives you insight into the number of queries currently taking place and the execution time of these queries. Two important query execution metrics to monitor are:


- **Query Execution Time** — Longer than usual query execution times may indicate a bottleneck, such as a slow host or problems with XDQP communication between hosts. Other possible problems include increased loads following a failover or more than the usual number of total requests.
- **Total Requests** — A consistent increase in the total number of outstanding requests may indicate the need to add more capacity and/or load balance. Decreases in total requests may indicate some “upstream” problem that needs to be addressed.

To display monitoring data related to query execution, select the Query Execution tab in the top left-hand portion of the Monitoring Dashboard.



The left side of the Query Execution page displays the maximum execution time (in seconds) of the current queries and the number of requests captured at each sample interval. You can hover a query execution sample to view the mean, maximum, and minimum execution times and the standard deviation from the mean.

Query Execution



The right side of the Query Execution page displays the five longest running queries since the beginning of the session and the longest running queries at the current time.

5 longest running queries since 13:29:24

Host	Server	Module	Time
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	533.47s
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	523.46s
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	518.46s
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	513.45s
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	503.46s

Longest running queries at 13:31:07

Host	Server	Module	Time
gordon-3.marklogic....	App-Servi...	...le/endpoints/eval.xqy	533.47s
gordon-3.marklogic....	Documents...	/apply.xqy	124.65s
gordon-3.marklogic....	TaskServer	...ansaction-manager.xqy	72.71s
gordon-1.marklogic....	TaskServer	...push-local-forest.xqy	65.87s
gordon-1.marklogic....	TaskServer	...push-local-forest.xqy	64.61s


2.8 Monitoring Rates and Loads

In general, rates and loads measure how efficiently data is exchanged between applications and MarkLogic Server. Rates and loads are defined as follows:

- Rates — The amount of data (MB per second) currently being read from or written to MarkLogic Server.
- Loads — The execution time (in seconds) of current read and write requests, which includes the time requests spend in the wait queue when maximum throughput is achieved.

For details on how to interpret rates and loads, see “What is Happening on the MarkLogic Server Cluster Now?” on page 11.

To display monitoring data related to rates and loads, select the Rates and Loads tab in the top left-hand portion of the Monitoring Dashboard.



There are three types of rates and loads monitoring data. Select the type of rates and loads data by clicking on one of the three buttons displayed under Rates and Loads:

Rates and Loads



The monitoring data displayed by each of these buttons is described in the following sections:

- [Overview](#)
- [XDQP Communication](#)
- [Backup/Restore](#)

2.8.1 Overview

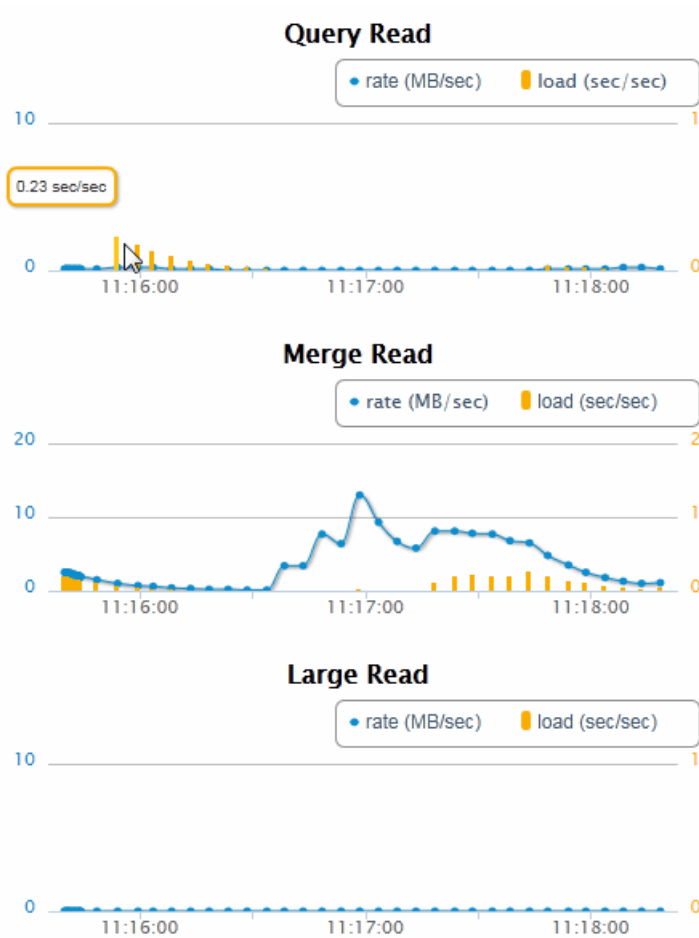
To obtain rates and loads data for queries, merges, and large data, click on the Overview button:

Rates and Loads

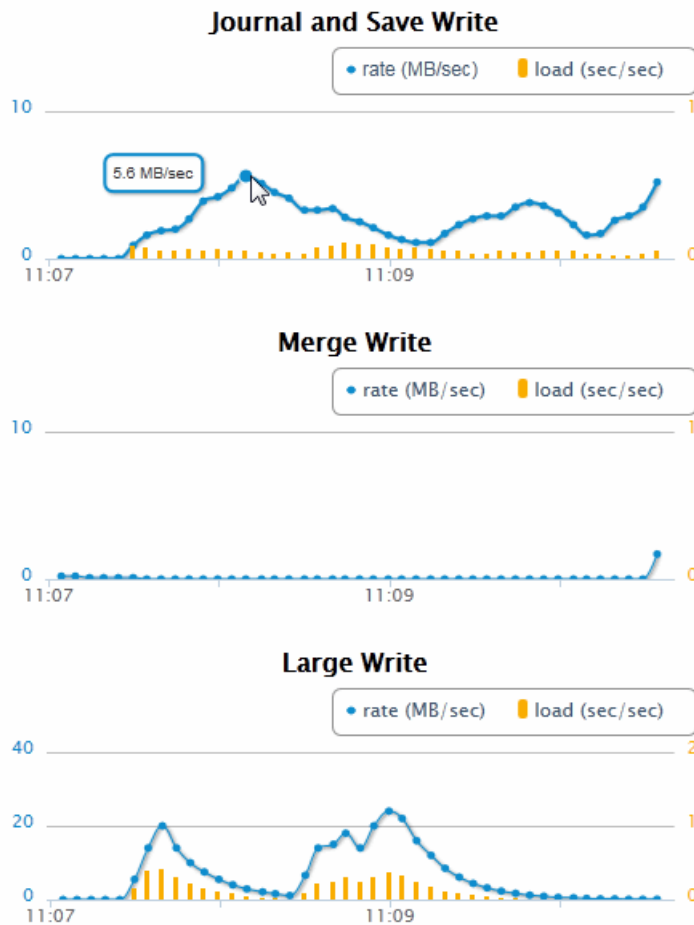


The left-hand side of the Rates and Loads Overview page displays the monitoring data related to query, merge, and large data reads.

Note: For details on Large Data, see [Working With Binary Documents](#) in the *Application Developer's Guide*.



The right-hand side of the Rates and Loads Overview page displays the monitoring data related to journal and save, merge, and large data writes.



2.8.2 XDQP Communication


Communication between MarkLogic Server hosts within a cluster and between hosts in different clusters is done using the XDQP protocol. Both the rate and load are displayed for each sample interval. Unusually high XDQP loads may indicate a network connection problem.

To monitor the rates and loads related to XDQP communication, click on the XDQP Communication button:


Rates and Loads



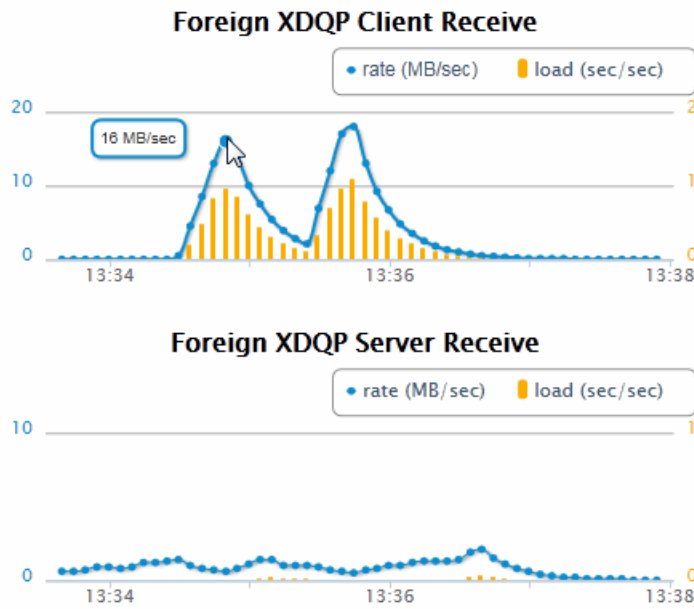
The upper left-hand side of the XDQP Communication page displays the monitoring data related to XDQP data received by the client and server.




The upper right-hand side of the XDQP Communication page displays the monitoring data related to XDQP data sent by the client and server.



The lower left-hand side of the XDQP Communication page displays the monitoring data related to XDQP data received by the client and server from a foreign cluster.



The lower right-hand side of the XDQP Communication page displays the monitoring data related to XDQP data sent by the client and server to a foreign cluster.



2.8.3 Backup/Restore


Backup and restore processes can impact the performance of applications and other background tasks in MarkLogic Server, such as merges and indexing.

To monitor the rates and loads related to backup and restore operations, click on the Backup/Restore button:


Rates and Loads



The left-hand side of the Backup/Restore page displays the monitoring data related to Backup reads and writes.




The right-hand side of the Backup/Restore page displays the monitoring data related to Restore reads and writes.



2.9 Monitoring Disk Space


Disk space usage is a key monitoring metric. In general, forest merges require twice as much disk space than that of the data stored in the forests. If a merge runs out of disk space, it will fail. In addition to the need for merge space on the disk, there must be sufficient disk space on the file system in which the log files reside to log any activity on the system. If there is no space left on the log file device, MarkLogic Server will abort. Also, if there is no disk space available to add messages to the log files, MarkLogic Server will fail to start.

To display monitoring data related to disk space, select the Disk Space tab in the top left-hand portion of the Monitoring Dashboard.



The data displayed on the Disk Space is for a specific host. You can select the host in the upper-left-hand section of the Disk Space page. The hosts in this list are sorted by those with the least available disk space at the top.

MarkLogic Disk Space Available

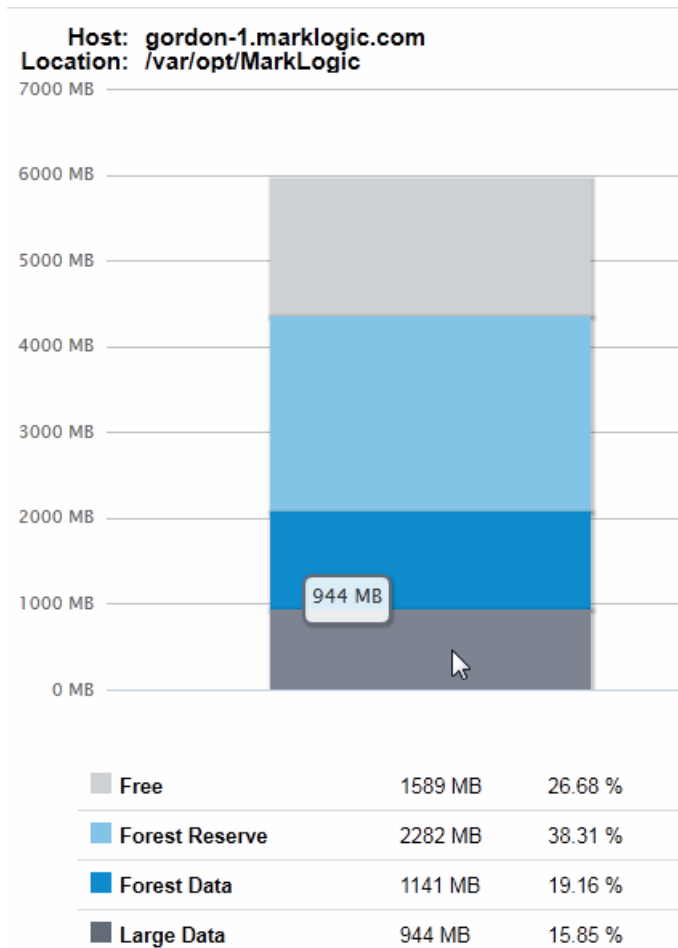


The disk space monitoring metrics are:

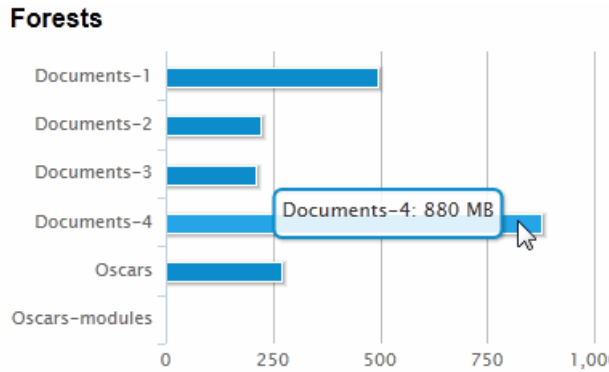
- **Fast Data** — The amount of disk space used by the forests' Fast Data Directory. The Fast Data Directory is typically mounted on a specialized storage device, such as a solid state disk. Fast data consists of transaction journals and as many stands that will fit on the fast storage device. For more information on Fast Data, see [Fast Data Directory on Forests](#) in the *Query Performance and Tuning Guide*.
- **Large Data** — The amount of disk space used by the forests' Large Data Directory. The Large Data Directory contains binary files that exceed the 'large size threshold' property set for the database. Large Data is not subjected to merges so, unlike Forest Data, Large Data does not require any additional Forest Reserve disk space. For more information on Large Data, see [Working With Binary Documents](#) in the *Application Developer's Guide*.
- **Forest Data** — The amount of disk space used by the data in the forest stands. This data is subject to periodic merges.

- Forest Reserve — The amount of free disk space that should be held in reserve to enable MarkLogic Server to merge the Forest Data.
- Free — The amount of free space on the disk that remains after accounting for the Forest Reserved space.


The upper right-hand section of the Disk Space page displays the amount of free space on the disk, along with how much reserve space is reserved for forest merges and the actual amount of space currently used by the forests and large data.



The lower right-hand section of the Disk Space page displays the amount of space on the disk used by the individual forests.




If your disk has less than 15% capacity a warning message is generated, as shown below. If the capacity falls to less than 10%, a critical message is generated.



2.10 Exporting Monitoring Data

Each of the three tabbed Monitoring Dashboard pages (Disk Space, Query Execution, Rates and Loads) has an Export button in its upper right corner, on the same line as the current tab's name. When clicked, it exports the page's data to a local XML file, formatted to be openable in Excel.



The exported files have tab-specific names incorporating a timestamp of when the file was exported. For example:

```
disk-space-20120210-160945.xml
```

indicates that it contains a page of data from the Disk Space tab, exported on February 10th, 2012 (2012 02 10) at 4:09:45 p.m. (16 09 45) (spaces added in this paragraph for clarity).

The exported data is from a JavaScript cache that automatically accumulates data as the page is drawn and refreshed. Two of the tabbed pages, Query Execution and Rates and Loads, accumulate data over time. A maximum 1000 latest data points are cached for each of these pages, no matter how long the monitor page runs.

By default, data is cached every 10 seconds. This rate depends on the polling interval, which is set on the Dashboard page within the Refresh drop-down menu. See “Setting the Sample Interval” on page 16.

When using the Export button, remember these caveats:

- The cache is not in a persistent file, so manually refreshing the browser clears it of all accumulated data. Immediately after a manual browser refresh, there is no data to export.
- Clicking Export returns only the data from the current tab's page. For example, if you are on the Query Execution tab, clicking Export only writes out data from Query Execution and does not write out data from the Rates and Loads or Disk Space tabs. To get the values from all three tabs, you have to go to each tab and click its Export button, resulting in three separate files.
- However, when clicking Rates and Loads' Export button, the file does contain the data from all three of Rates and Loads' sub-tabs (Overview, XDQP Communication, and Backup/Restore).

Previously, you had to turn on caching this data with a `debug=true` parameter in the browser URL. Now, data is cached by default.

3.0 MarkLogic Server Monitoring History

This chapter describes how to use the Admin Interface and Monitoring History dashboard to capture and make use of historical performance data for a MarkLogic cluster. These same Monitoring History operations can also be done using the XQuery and REST APIs, as described in *XQuery and XSLT Reference Guide* and the *MarkLogic REST API Reference*.

Note: All MB and GB metrics described in this chapter are base-2.

The main topics in the chapter are:

- [Overview](#)
- [Enabling Monitoring History on a Group](#)
- [Setting the Monitoring History Data Retention Policy](#)
- [Viewing Monitoring History](#)
- [Viewing Monitoring History by Time Span and Frequency](#)
- [Labeling Monitoring History Time Spans](#)
- [Filtering Monitoring History by Resources](#)
- [Historical Performance Charts by Resource](#)
- [Exporting and Printing Monitoring History](#)

3.1 Overview

The Monitoring History feature allows you to capture and view critical performance data from your cluster. Once the performance data has been collected, you can view the data in the Monitoring History pages. The top-level Monitoring History page provides an overview of the performance metrics for all of the key resources in your cluster. For each resource, you can drill down for more detail. You can also adjust the time span of the viewed data and apply filters to view the data for select resources to compare and spot exceptions.

By default, the performance data is stored in the Meters database. Monitoring history capture is enabled at the group level. Typically you have one group per cluster. You can also configure a consolidated Meters database that captures performance metrics from multiple groups. The group configuration defines which database is used to store performance metrics for that group (defaulting to a shared Meters database per cluster), as well as all configuration parameters for performance metrics, such as the frequency of data capture and how long to retain the performance data. The Meters database can participate in all normal database replication, security, and failover operations.

3.2 Enabling Monitoring History on a Group

In order to collect monitoring history data for your cluster, you must enable performance metering for your group.

1. Log into the Admin Interface.
2. Click the Groups icon on the left tree menu.
3. Locate the Performance Metering Enabled field toward the bottom of the Group Configuration page and click on `true`.

The screenshot displays the configuration page for a group in the MarkLogic Admin Interface. The page has a light beige background and contains several configuration fields:

- metering enabled**: A radio button selection with `true` selected and `false` unselected. Below it is the text "Enable Usage metering."
- performance metering enabled**: A radio button selection with `true` selected and `false` unselected. Below it is the text "Enable performance monitoring history." A mouse cursor is pointing at the `true` radio button.
- meters database**: A dropdown menu currently showing "Meters". Below it is the text "The metering and performance history database."
- performance metering period**: A text input field containing the number "1". Below it is the text "Historic performance metering period (minutes)."
- performance metering retain raw**: A text input field containing the number "7". Below it is the text "Retain raw performance metering data in days."
- performance metering retain hourly**: A text input field containing the number "30". Below it is the text "Retain hourly performance metering data in days."
- performance metering retain daily**: A text input field containing the number "90". Below it is the text "Retain daily performance metering data in days."

You can configure the parameters for collecting monitoring history data, as described in the table below.

Parameter	Description
meters database	The database in which performance monitoring history data and usage metrics documents are stored. By default, historical performance and usage metrics are stored in the Meters database.
performance metering period	<p>The performance metering period, in minutes. Performance data is collected at each period. The period can be any value of 1 minute or more.</p> <p>Note: If you are collecting monitoring history for multiple groups, you should either set the same period for each group or configure your filter to view the history data for one group at a time.</p>
performance metering retain raw	The number of days raw performance monitoring history data is retained. See “Setting the Monitoring History Data Retention Policy” on page 32 for details.
performance metering retain hourly	The number of days hourly performance monitoring history data is retained. See “Setting the Monitoring History Data Retention Policy” on page 32 for details.
performance metering retain daily	The number of days daily performance monitoring history data is retained. See “Setting the Monitoring History Data Retention Policy” on page 32 for details.

3.3 Setting the Monitoring History Data Retention Policy

The retention policy (for raw, hourly, daily) is a value set in days. If performance metering is enabled, then all data that is older than that many days for the specified period (raw, hour, day) is deleted. The retention policy is set at a group level, so different groups can have different retention policies. For example, GroupA may have raw set to 1 day and GroupB may have raw set to 10 days. The cleanup code follows this retention value on a per-group basis.

There are cases where metering data may become orphaned, so it may no longer belong to an existing group. Some examples of when this could occur are:

- Deleting a group
- Importing metering data from another cluster

Any metering data that no longer belongs to any active group in the current cluster is deleted. To avoid this, turn off metering or avoid deleting groups and instead move hosts out of the group but keep the group in the cluster configuration.

Note: Loading older monitoring history data (for example, by restoring a backup of the Meters database) will be immediately affected by data retention policy. So, you should turn off performance metering prior to restoring any data that is older than the time specified by your retention policy.

Deletion of data older than the retention policy occurs no sooner than the retention policy, but may, for various reasons, still be maintained for an unspecified amount of time.

Note: Changing the retention policy from smaller to larger values does not restore data that has already been deleted.

The default data retention policy settings are as shown in the table below. To maximize efficiency, it is a best practice to retain raw data for the least number of days and the daily data for the most number of days.

Period	Retention Period
Raw	7 Days
Hourly	30 Days
Daily	90 Days

3.4 Viewing Monitoring History


You can display the Monitoring History dashboard by doing the following:

1. Open a browser and enter the URL:

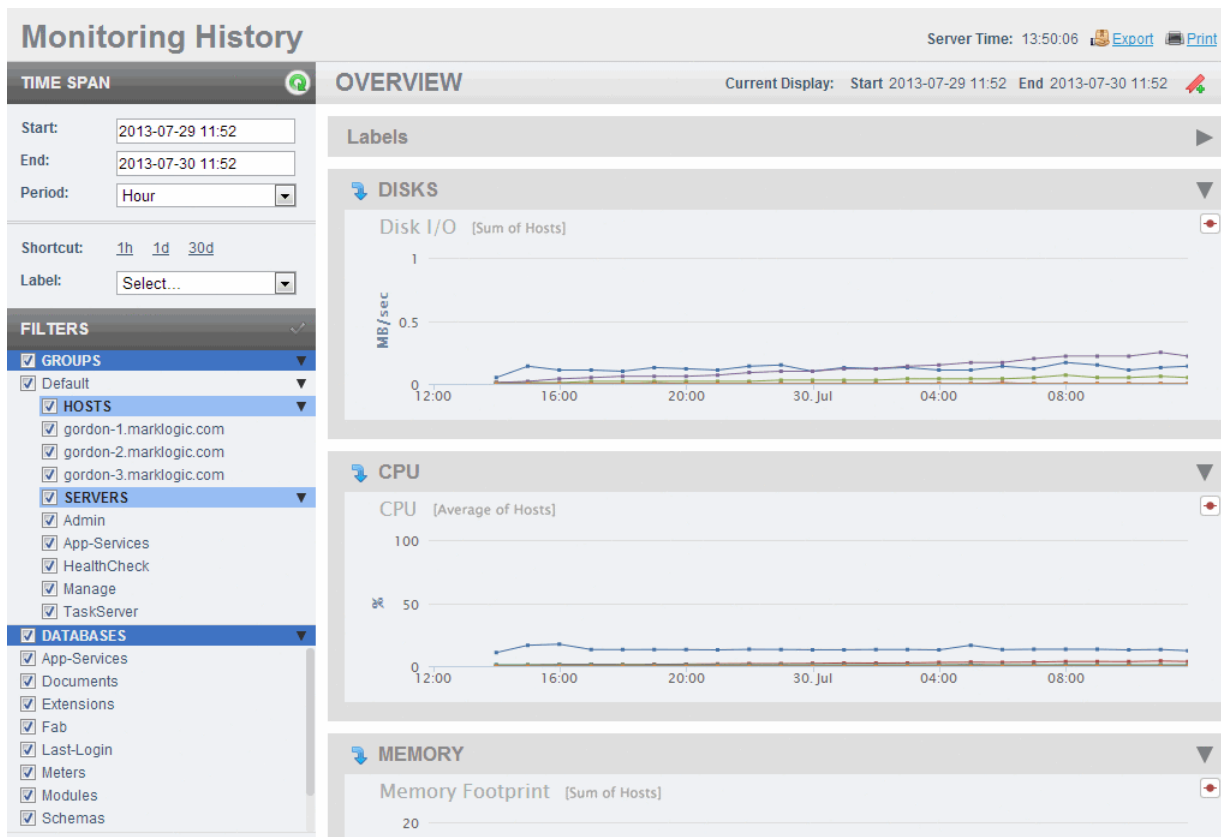
`http://monitor-host:8002/`

where *monitor-host* is a host in the cluster you want to monitor

2. At the top of the page, click on Monitoring and click on History in the pull-down menu:



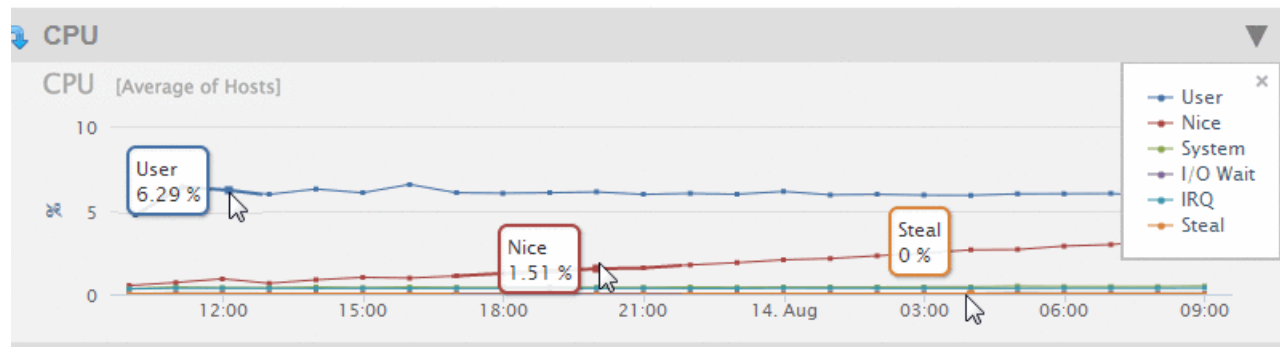
3. The Monitoring History page appears. From the Monitoring History Overview page, you can navigate to any of the pages described in this chapter.




Each line in a chart represents a metric for the resource. In the Overview page, the lines represent an aggregate of the metrics for all of the cluster resources. In each Details page, the lines represent the metric for each specific resource.

Chart titles on the Overview page include bracketed information specifying how chart data gathered across multiple resources is aggregated. For example, [Sum of Hosts] means that the data retrieved from one or more hosts is summed for display as points on the chart.


Each point on a line represents a period in which the performance data was captured. Hovering over a chart point displays the name of the resource metric, along with the performance value for the metric at that point in time.




The displayed metrics (in MegaBytes per second) are color coded. You can display a legend that indicates which colors represent which metrics by clicking on the red dot in the upper right-hand section of the graph. To close the legend, click on the 'x' in the upper right-hand portion of the legend window.



To simplify the view of charts on a page, you can collapse a chart or a group of charts for a resource by clicking on the triangle in the upper right-hand portion of the chart or chart group.



To expand a collapsed chart view, click on the triangle in the upper right-hand portion of the collapsed chart.



3.5 Viewing Monitoring History by Time Span and Frequency

As described in “Enabling Monitoring History on a Group” on page 30, the frequency in which performance metrics are captured is configurable, in minute intervals. The snapshots of performance metrics for each host are rolled up into a summary document that contains aggregate calculations on the values for that host.

You can configure your view of the captured performance data by time span and frequency.


The Time Span settings are located in the upper left-hand corner of the Monitoring History page.

There are three basic settings you can adjust to control how the data is displayed:

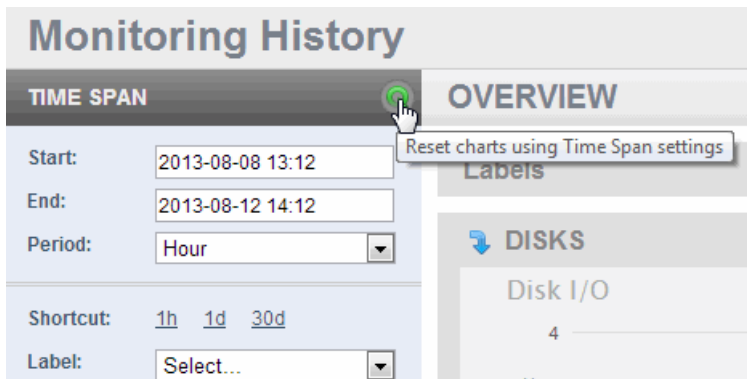
- A date/time range, down to the granularity of a minute, that determines the time span of the displayed data. (By default, this is the last 24 hours.)
- A period interval that determines the frequency of the displayed data. The possible intervals are shown in the following table.

Period	Description
Raw	Display the performance data just as it was captured with the set frequency.
Hour	Display the performance data, in aggregate form, per hour. (This is the default.)
Day	Display the performance data, in aggregate form, per day.

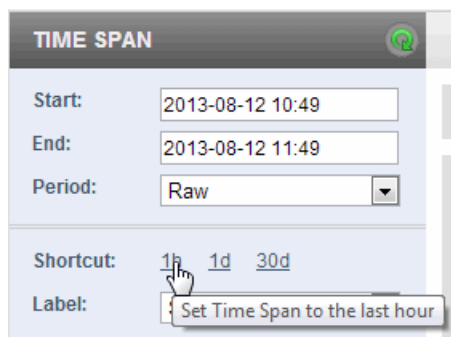
You can “zoom in” to display part of the timespan by selecting the begin time of your “zoom” on any chart and click and hold your left mouse button and drag it to the end “zoom” time. The selected timeframe is highlighted and the zoomed-in time is displayed for all of the charts in the page. Navigating to another Monitoring History page resets all of the charts to the timespan selected in the TIME SPAN panel.



After changing either the time span and/or the period, click on refresh to display the updated charts. Clicking refresh will also update any changes you've made to the Filters settings. For details about filters, see “Filtering Monitoring History by Resources” on page 41. If you have zoomed into a portion of a timespan, refresh will redisplay the charts using the timespan selected in the TIME SPAN panel.



You can use the Shortcut links to display either the last hour, day or 30 days of performance data. Selecting a Shortcut link will automatically refresh the displayed charts.



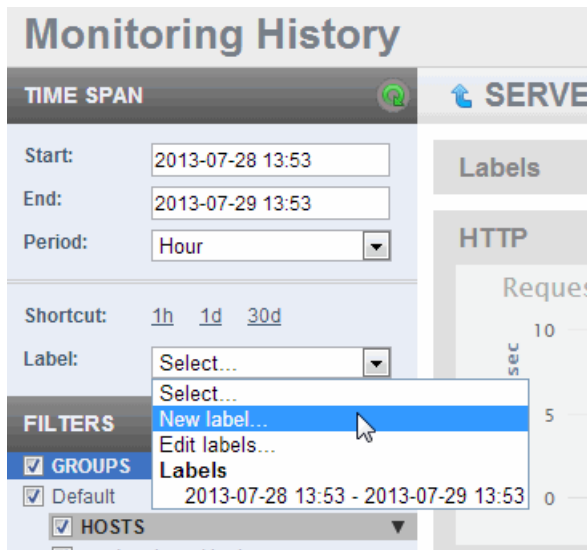
Each Shortcut also sets the Period value, as shown in the table below.

Shortcut	Period
1h	Raw
1d	Hour
30d	Day


3.6 Labeling Monitoring History Time Spans

You can use the Label feature to capture and tag metrics for the set time span. You can store any number of labels. These labels can be used to identify events, instances, and periods of time. Labels can be added, updated or deleted at any time. Labels themselves are not stored with the raw metric data. They are only used for reporting purposes.


1. To create a label for your current view of the Monitoring History, select New Label from the Label pull-down menu.




2. In the Create a New Label popup window, the name of the label is the time span of the currently displayed charts, by default.



3. You can keep the default name for the label, or change it to be more descriptive. Click Save.



4. You can edit your label names or delete labels by selecting Edit Labels from the Labels pull-down menu.




5. In the Edit Labels popup window, you can either edit the label name or delete the label. To delete a label, hover over the label and a click on the garbage can icon to the right. When finished editing, click Close.




Note: If you edit a label and, before closing the Edit Labels window, decide not to save your edits, press the Esc key to terminate the edits and keep the original labels.

- You can view all of the labels that have data within the currently selected timespan by clicking on the triangle to the right of the Labels section at the top of the Monitoring History page to expand the Labels chart.




- Each label appears as a timeline. Hover over a timeline to display the label name. Click on a timeline to update the view to the time span associated with the label. Selecting a timeline is functionally equivalent to selecting a label from the Label menu in that it updates the view with the start and end times in the TIME SPAN panel.




Note: If your labeled data has been purged from the Meters database, as the result of the retention policy or some other reason, the label will remain but there will be no data associated with that label.

- You can click on the label icon at the top right-hand portion of the page to create a label for the currently displayed time span. Follow the same procedure as described in steps [2](#) and [3](#) to finish creating the label.



If the data for a label does not fall within the currently displayed timespan, the label will not be displayed in the Labels chart. To display the charts for such labels, select the label from the Label pull-down menu.

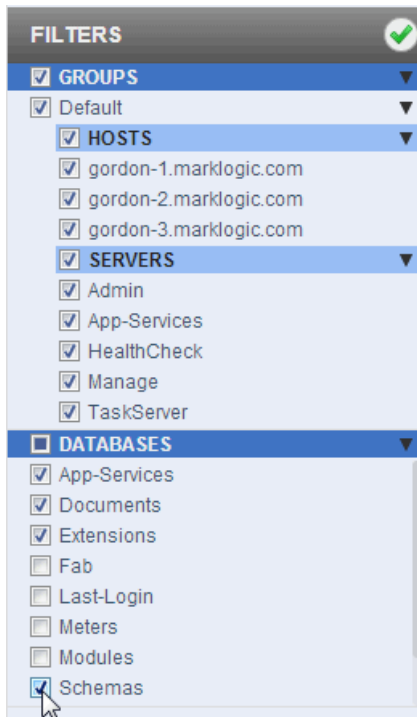


3.7 Filtering Monitoring History by Resources


You can set filters for select resources to display only the stored performance metrics for those resources. You can filter by groups and databases. And in each group, by hosts and servers. By default, the metrics for all of the resources in the cluster are displayed.

Filter types that are active for the current view have headings highlighted in blue. For example, on the Overview page, all filters are active while on the Databases Detail view, only database resources are active.


In the filters panel, you can check or uncheck a resource to display or not display the performance metrics for that resource.



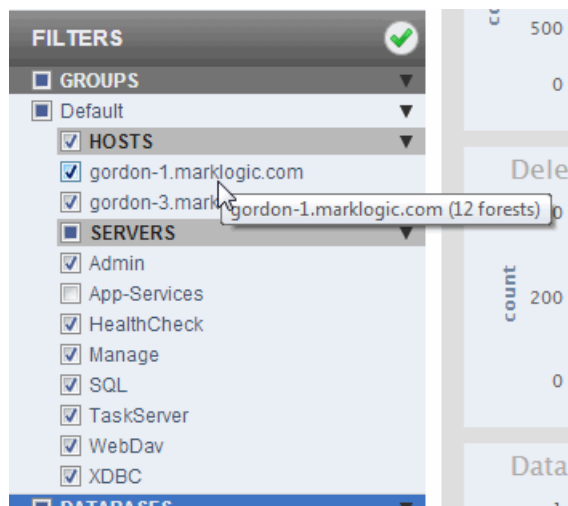
In order to focus on the resources of interest, you can collapse a category by clicking on the triangle in the right-hand section of the panel. The number of resources for the collapsed category are displayed.



Clicking the checkmark updates the charts with the current filter settings. It does not apply any changes that may have been made to the above TIME SPAN settings.




You can mouse over the resource names in the filter list to get extra information about the resources. For example, mousing over a host name shows the number of forests associated with the host and mousing over a server name shows the server type.




3.8 Historical Performance Charts by Resource

From the Monitoring History dashboard, you can view Overview and Detailed performance metrics in graph form for each resource in the cluster. In the Overview page, the lines on a graph represent an aggregate of the metrics for all of the cluster resources of that type. In each Details page, the lines represent the metric for each specific resource in the cluster.

To view the Detail page for a resource, click on the down arrow at the upper left-hand section of the resource graph on the Overview page.



To return to the Overview page from a Detail page, click on the up arrow at the upper left-hand section of the resource graph on the Detail page.




This section describes the Overview and Detail pages for the following resources:

- [Disk Performance Data](#)
- [CPU Performance Data](#)
- [Memory Performance Data](#)
- [XDQP Server Requests Performance Data](#)
- [Server Performance Data](#)
- [Network Performance Data](#)
- [Database Performance Data](#)

3.8.1 Disk Performance Data


The Overview page displays a graph of the aggregate I/O performance data for the disks used by the hosts selected in the filter.



As described in “Viewing Monitoring History” on page 33, you can hover on a period point to view what disk operation was taking place at that point in time. Each performance metric is described in the table below.

Metric	Description
Writes	The disk I/O performance during journal and save write operations. This is the sum of journal-write-rate, save-write-rate, and large-write-rate.
Query Traffic	The disk I/O performance during a query or queries. This is the sum of query-read-rate and large-read-rate.
Merge Reads	The disk I/O performance during a merge read operation.
Merge Writes	The disk I/O performance during a merge write operation.
Backup Reads	Throughput of reading backup data from disk, in megabytes per second.
Backup Writes	Throughput of writing data for backups, in megabytes per second.
Restore Reads	Disk read throughput for restore, in megabytes per second.
Restore Writes	Disk writing throughput for restore, in megabytes per second.

Click on the  arrow in the upper left-hand section of the DISKS graph in the Overview page to view charts that present more detailed disk performance metrics.



The metrics displayed by the charts on the DISKS DETAIL page are described in the table below.

Chart	Definition of Displayed Metric
Journal Write Rate	The moving average of data writes to the journal.
Save Write Rate	The moving average of data writes to in-memory stands.
Query Read Rate	The moving average of reading query data from disk
Merge Read Rate	The moving average of reading merge data from disk
Merge Write Rate	The moving average of writing data for merges
Backup Read Rate	Throughput of reading backup data from disk, in megabytes per second.
Backup Write Rate	Throughput of writing data for backups, in megabytes per second.
Restore Read Rate	Disk read throughput for restore, in megabytes per second.
Restore Write Rate	Disk writing throughput for restore, in megabytes per second.
Large Binary Read Rate	The moving average of reading large documents from disk.
Large Binary Write Rate	The moving average of writing data for large documents to disk.

By default, Host data is viewed in aggregated form and must be viewed that way if multiple hosts are selected. When in the DISKS DETAIL page, you can rollover any Host filter to reveal the Select and Expand button. This will deselect all of the other Hosts across all Groups, and apply all pending filter changes. The expanded charts display the data for each forest in that host as separate line in each chart.


The screenshot shows the 'Monitoring History' interface for 'DISKS DETAIL'. On the left, the 'TIME SPAN' section shows a start time of 2013-08-25 09:34 and an end time of 2013-08-26 09:34, with a period set to 'Hour'. Below this, there are shortcuts for '1h', '1d', and '30d', and a 'Label' dropdown set to 'Select...'. The 'FILTERS' section is expanded to show a tree view: 'GROUPS' (unchecked), 'HOSTS' (checked), and 'SERVERS' (checked). Under 'HOSTS', two hosts are listed: 'gordon-1.marklogic.com' (checked) and 'gordon-3.marklogic.com' (unchecked). A tooltip is displayed over the 'gordon-1.marklogic.com' host, containing the text: 'Select only this Host and show its Forest data in charts.' The right side of the interface features three line charts: 'Journal Write Rate', 'Save Write Rate', and 'Query Read Rate'. Each chart has a y-axis labeled 'MB/sec' ranging from 0 to 1 and an x-axis with markers at 12:00 and 15:00. The data points in all three charts are very low, near the 0 line.

To return to the aggregate view, click on Aggregate button on an expanded Host. Doing so will also apply all pending filter changes to the displayed charts.

The screenshot displays the 'Monitoring History' interface. On the left, the 'TIME SPAN' section includes input fields for 'Start' (2013-08-25 09:34), 'End' (2013-08-26 09:34), and 'Period' (Hour). Below this is a 'Shortcut' section with buttons for '1h', '1d', and '30d', and a 'Label' dropdown menu. The 'FILTERS' section is expanded, showing a tree view with 'GROUPS' (Default), 'HOSTS' (gordon-1.marklogic.com, gordon-3.marklogic.com), and 'SERVERS' (Admin, App-Services, HealthCheck, Manage, SQL, TaskServer, WebDav, XDBC). A tooltip points to the 'gordon-3.marklogic.com' host, stating 'Show aggregated Forest data for this Host in charts.' The main area shows 'DISKS DETAIL (FORESTS)' with three charts: 'Journal Write Rate', 'Save Write Rate', and 'Query Rate'. Each chart has a y-axis labeled 'MB/sec' and an x-axis with markers at 12:00 and 15:00.

3.8.2 CPU Performance Data

The Overview page displays a graph of the aggregate performance data for the CPUs used by the hosts selected in the filter.



As described in “Viewing Monitoring History” on page 33, you can hover on a period point to view what CPU operation was taking place at that point in time. Each performance metric in the CPU Overview chart is described in the table below.

Metric	Description
User	Total percentage of CPU used running user processes that are not niced.
Nice	Total percentage of CPU used running user processes that are niced.
System	Total percentage of CPU used running the operating system kernel and its processes.
I/O Wait	Total percentage of CPU time spent waiting for I/O operations to complete.
IRQ	Total percentage of CPU utilization for servicing soft interrupts.
Steal	Total percentage of CPU ‘stolen’ from this virtual machine by the hypervisor for other tasks (such as running another virtual machine).


Click on the  arrow in the upper left-hand section of the CPU graph in the Overview page to view graphs that present more detailed CPU performance metrics. The charts on the CPU DETAIL page are described in the table below.

Chart	Description
I/O Wait	The percentage of CPU used waiting for I/O operations to complete for each host.
User	The percentage of CPU used running user processes that are not niced for each host.
System	The percentage of CPU used running the operating system kernel and its processes for each host.
Nice	The percentage of CPU used running user processes that are niced for each host.
Steal	The percentage of CPU 'stolen' from this virtual machine by the hypervisor for other tasks (such as running another virtual machine) for each host.
Idle	The percentage of CPU that is not doing any work for each host.
IRQ	The percentage of CPU servicing soft interrupts for each host.

3.8.3 Memory Performance Data

The Overview page displays a graph of the aggregate performance data for the Memory used by the hosts selected in the filter.



As described in “Viewing Monitoring History” on page 33, you can hover on a period point to view what CPU operation was taking place at that point in time. Each chart and associated performance metrics are described in the table below.

Chart	Description
Memory Footprint	The total amount (in GB) of memory consumed by all of the hosts in the cluster. The displayed metrics are: <ul style="list-style-type: none"> • RSS: The total amount of GB of Process Resident Size (RSS) consumed by the cluster. • Anon: The total amount of GB of Process Anonymous Memory consumed by the cluster.
Memory Size	The amount of space forest data files take up in memory.

Chart	Description
Memory I/O	<p>The number of pages per second moved between memory and disk.</p> <p>The displayed metrics are:</p> <ul style="list-style-type: none"> • Page-In Rate: The page-in rate (from Linux /proc/vmstat) for the cluster in pages/sec. • Page-Out Rate: The page-out rate (from Linux /proc/vmstat) for the cluster in pages/sec. • Swap-In Rate: The swap-in rate (from Linux /proc/vmstat) for the cluster in pages/sec. • Swap-Out Rate: The swap-out rate (from Linux /proc/vmstat) for the cluster in pages/sec.
Virtual Memory	<p>Size of virtual memory used by different objects and processes; includes:</p> <ul style="list-style-type: none"> • Data Files: Size of virtual memory mapped to data files. • Forests: Size of virtual memory used by forests. • Unclosed Stands: Size of virtual memory used by unclosed stands. • Caches: Size of virtual memory used by caches. • Registered Queries: Size of virtual memory used to store registered queries. • Joins: Size of virtual memory used for join processing.


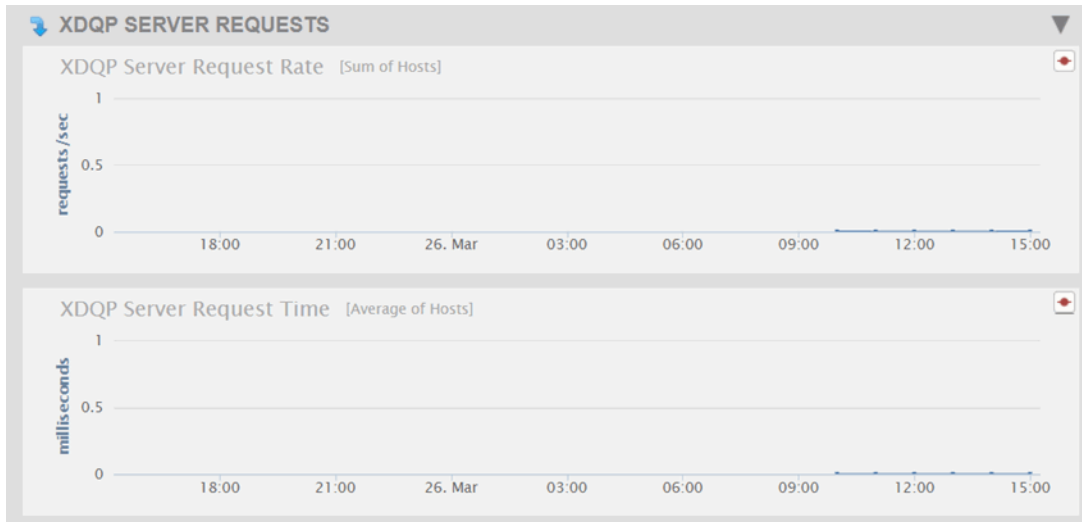
Click on the  arrow in the upper left-hand section of the MEMORY graph in the Overview page to view graphs that present more detailed MEMORY performance metrics. The charts on the MEMORY DETAIL page are described in the table below. The displayed metrics are drawn from `/proc/vmstat`.

Chart	Description
RSS	The amount of GB of Process Resident Size (RSS) for each host in the cluster.
Anon	The amount of GB of Process Anonymous Memory for each host in the cluster.
Process Size	The number of MB of total process virtual memory for the MarkLogic process.
Huge Pages	The size of huge pages for the MarkLogic process in MB. Available on Linux platform. Sum of Sizes after <code>/anon_hugepage</code> in <code>/proc/[MLpid]/smaps</code> .
System Free Memory	The free system memory in MB. MemFree from <code>/proc/meminfo</code> on Linux, <code>m.ullAvailPhys</code> from <code>GlobalMemoryStatusEx(m)</code> on Windows.
Page-In Rate	The page-in rate (in pages/sec) for each host in the cluster.
Page-Out Rate	The page-out rate (in pages/sec) for each host in the cluster.
Swap-In Rate	The swap-in rate (in pages/sec) for each host in the cluster.
Swap-Out Rate	The swap-out rate (in pages/sec) for each host in the cluster.
Data File Memory	Size of virtual memory mapped to data files.
Forest Memory	Size of virtual memory used by forests.
Unclosed Stand Memory	Size of virtual memory used by unclosed stands.
Cache Memory	Size of virtual memory used by caches.
Registered Query Memory	Size of virtual memory used to store registered queries.
Join Memory	Size of virtual memory used for join processing.

3.8.4 XDQP Server Requests Performance Data

The Overview page displays a graph of the aggregate performance data for the XDQP Server Requests processed by the hosts selected in the filter.



Each chart and associated performance metrics are described in the table below.

Chart	Description
XDQP Server Request Rate	Number of XDQP requests processed per second.
XDQP Server Request Time	Average response time to XDQP requests from other nodes.


Click on the  arrow in the upper left-hand section of the XDQP SERVER REQUESTS graph in the Overview page to view graphs that present more detailed performance metrics. The charts on the XDQP SERVER REQUESTS DETAIL page are described in the table below.

Chart	Description
XDQP Server Request Rate	Number of XDQP requests processed per second.
XDQP Server Request Time	Average response time to XDQP requests from other nodes.

3.8.5 Server Performance Data

The Overview page displays graphs of the aggregate performance data for the App Servers selected in the filter.




The Overview page displays the charts described in the table below.


Chart	Description
App Server Request Rate	The total number of queries being processed per second, across all of the App Servers.
App Server Latency	The average time (in seconds) it takes to process queries, across all of the App Servers.
Task Server Queue Size	The number of tasks in the Task Server queue.
Expanded Tree Cache Hits/Misses	The number of times per second that queries could use (Hits) and could not use (Misses) the expanded tree cache.

With the exception of the Task Server Queue Size chart, which only displays the queue size for the one task server, the color-coded metrics for the server charts are as shown in the table below.


Metric	Description
HTTP	The metrics for the HTTP servers.
ODBC	The metrics for the ODBC servers.
WebDAV	The metrics for the WebDAV servers.
XDBC	The metrics for the XDBC servers.
Task	The metrics for the Task server.

Click on the  arrow in the upper left-hand section of the SERVERS graph in the Overview page to view graphs that present more detailed performance metrics for each App Server. The charts displayed on the SERVERS DETAIL page are described in the table below.

Note: If there are multiple groups defined, server names have the group that they are associated with in square brackets in the legend and rollovers.



The number of servers displayed out of the number of servers of each type in the cluster (for example, HTTP) is shown in the upper right-hand section of each server type group.



The following detailed charts are displayed for each type of App Server:

Chart	Description
Request Rate	The number of queries being processed per second by each App Server.
Latency	The average time it takes each App Server to process queries.
Expanded Tree Cache Rate Hits	The number of times queries could use the expanded tree cache on each App Server.
Expanded Tree Cache Rate Misses	The number of times queries could not use the expanded tree cache on each App Server.
Queue Size (Task Server only)	The number of tasks in the Task Server queue on each host.


Chart	Description
Send Rate (for any type of App Server except Task Server)	Throughput of application servers of that type sending data, in megabytes per second.
Receive Rate (for any type of App Server except Task Server)	Throughput of application servers of that type receiving data, in megabytes per second.

3.8.6 Network Performance Data

The network performance data graphs display performance in terms of XDQP reads and writes. XDQP is the protocol MarkLogic uses for internal host-to-host communication on port 7999.

The Overview page displays various XDQP performance as the sum of XDQP activity across the cluster. High XDQP rates are usually not an issue unless they are so high as to saturate your internal network. Higher usage occurs during data load and query execution. Merges do not involve XDQP.

Note: If XDQP is excessively high during loads, running the MarkLogic Content Pump (m1.cp) with fast forest placement will minimize XDQP communication needs. For details on the MarkLogic Content Pump, see [Loading Content Using MarkLogic Content Pump](#) in the *Loading Content Into MarkLogic Server Guide*.



The Overview page displays a chart with the metrics described in the table below.

Metric	Description
Network	<p>The network traffic between nodes in the cluster. Heavy queries or ingestion will create a spike.</p> <p>The displayed metrics are:</p> <ul style="list-style-type: none"> • XDQP Read: The total volume of all XDQP reads between hosts in the cluster. This is the sum of <code>xdqp-client-receive-rate</code> and <code>xdqp-server-receive-rate</code>. • XDQP Write: The total volume of all XDQP writes between hosts in the cluster. This is the sum of <code>xdqp-client-send-rate</code> and <code>xdqp-server-send-rate</code>. • Foreign XDQP Read: The total volume of all XDQP reads by the hosts in the cluster from a foreign cluster. This is the sum of <code>foreign-xdqp-client-receive-rate</code> and <code>foreign-xdqp-server-receive-rate</code>. • Foreign XDQP Write: The total volume of all XDQP writes by the hosts in the cluster to a foreign cluster. This is the sum of <code>foreign-xdqp-client-send-rate</code> and <code>foreign-xdqp-server-send-rate</code>.
External KMS Request Rate	Number of requests per second to the external key management server.
External KMS Request Time	Average round-trip time for a request to an external key management server.



Click on the  arrow in the upper left-hand section of the NETWORK graph in the Overview page to view graphs that present more detailed performance metrics for each host in the cluster. The charts displayed on the NETWORK DETAIL page are described in the table below.

Chart	Description
XDQP Read Rate	The amount of data (in MB/sec) read over XDQP by each host in the cluster. This is the sum of foreign-xdqp-client-receive-rate and foreign-xdqp-server-receive-rate.
XDQP Write Rate	The amount of data (in MB/sec) written over XDQP by each host in the cluster. This is the sum of foreign-xdqp-client-send-rate and foreign-xdqp-server-send-rate.
XDQP Read Load	The execution time (in seconds) of read requests by each host in the cluster. This is the sum of xdqp-client-receive-load and xdqp-server-receive-load.
XDQP Write Load	The execution time (in seconds) of write requests by each host in the cluster. This is the sum of xdqp-client-send-load and xdqp-server-send-load.
Foreign XDQP Read Rate	The amount of data (in MB/sec) read over XDQP by each host in the cluster from a foreign cluster. This is the sum of foreign-xdqp-client-receive-rate and foreign-xdqp-server-receive-rate.
Foreign XDQP Write Rate	The amount of data (in MB/sec) written over XDQP by each host in the cluster to a foreign cluster. This is the sum of foreign-xdqp-client-send-rate and foreign-xdqp-server-send-rate.
Foreign XDQP Read Load	The execution time (in seconds) of read requests by each host in the cluster from a foreign cluster. This is the sum of foreign-xdqp-client-receive-load and foreign-xdqp-server-receive-load.
Foreign XDQP Write Load	The execution time (in seconds) of write requests by each host in the cluster to a foreign cluster. This is the sum of foreign-xdqp-client-send-load and foreign-xdqp-server-send-load.
External KMS Request Rate	Number of requests per second to the external key management server.
External KMS Request Time	Average round-trip time for a request to an external key management server.

3.8.7 Database Performance Data

The Overview page displays graphs of the aggregate performance data for all of the databases in the cluster.



The table below describes the charts displayed in the Databases section of the Overview page.

Chart	Description
Fragments	<p>Displays the aggregate number of fragments in all of the databases in the cluster.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Active Fragments: The fragments available to queries. • Deleted Fragments: The fragments to be deleted during the next merge operation.
Storage Footprint	<p>The total disk capacity (in GBs) used by all of the databases in the cluster.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Data Size: The amount of data in the forest data directories. • Fast Data Size: The amount of data in the forest fast data directories. • Large Data Size: The amount of data in the forest large data directories.

Chart	Description
Lock Rate	<p>The number of locks set per second across all of the databases in the cluster.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Read: The number of read locks set per second. • Write: The number of write locks set per second. • Deadlock: The number of deadlocks per second.
Lock Wait Load	<p>The aggregate time (in seconds) transactions wait for locks;</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Read: The time transactions wait for read locks. • Write: The time transactions wait for write locks.
Lock Hold Load	<p>The aggregate time (in seconds) locks are held.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Read: The time read locks are held. • Write: The time write locks are held.
Deadlock Wait Load	<p>The aggregate time (in seconds) deadlocks remain unresolved.</p>
Database Replication	<p>The amount of data (in MB per second) sent by and received from this cluster and foreign clusters.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Database Replication Send: The amount of data sent to foreign clusters. • Database Replication Receive: The amount of data received from foreign clusters.

Chart	Description
<p>List Cache Hits/Misses</p>	<p>The number of times per second that queries could use (Hits) and could not use (Misses) the list cache.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • List Cache Hit Rate: The average number of hits on the list cache. • List Cache Miss Rate: The average number of misses on the list cache.
<p>Compressed Tree Cache Hits/Misses</p>	<p>The number of times per second that queries could use (Hits) and could not use (Misses) the compressed tree cache.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Compressed Tree Cache Hit Rate: The average number of hits on the compressed tree cache. • Compressed Tree Cache Miss Rate: The average number of misses on the compressed tree cache.
<p>Triple Cache Hits/Misses</p>	<p>The number of times per second that queries could use (Hits) and could not use (Misses) the triple cache.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Triple Cache Hit Rate: The average number of hits on the triple cache. • Triple Cache Miss Rate: The average number of misses on the triple cache.
<p>Triple Value Cache Hits/Misses</p>	<p>The number of times per second that queries could use (Hits) and could not use (Misses) the triple value cache.</p> <p>The displayed lines are:</p> <ul style="list-style-type: none"> • Triple Value Cache Hit Rate: The average number of hits on the triple value cache. • Triple Value Cache Miss Rate: The average number of misses on the triple value cache.


Click on the  arrow in the upper left-hand section of the DATABASES graph in the Overview page to view graphs that present more detailed performance metrics for each database. The charts displayed on the DATABASES DETAIL page are described in the table below. The metrics for each database in the cluster are displayed as a separate line.


Chart	Description
Active Fragments	The number of active fragments (the fragments available to queries) in each database.
Deleted Fragments	The number of deleted fragments (the fragments to be removed by the next merge operation) in each database.
Data Size	The amount of data in the data directories of the forests attached to each database.
Fast Data Size	The amount of data in the fast data directories of the forests attached to each database.
Large Data Size	The amount of data in the large data directories of the forests attached to each database.
Read Lock Rate	The number of read locks set per second on each database.
Write Lock Rate	The number of write locks set per second on each database.
Deadlock Rate	The number of deadlocks per second on each database.
Read Lock Wait Load	The time (in seconds) transactions wait for read locks on each database.
Write Lock Wait Load	The time (in seconds) transactions wait for write locks on each database.
Deadlock Wait Load	The aggregate time (in seconds) deadlocks remain unresolved on each database.
Read Lock Hold Load	The time (in seconds) read locks are held on each database.
Write Lock Hold Load	The time (in seconds) write locks are held on each database.
Database Replication Send Rate	The amount of replication data (in MB per second) sent by each database to foreign clusters.
Database Replication Receive Rate	The amount of replication data (in MB per second) received by each database from foreign clusters.

Chart	Description
Database Replication Send Load	The time (in seconds) it takes each database to send replication data to foreign clusters.
Database Replication Receive Load	The time (in seconds) it takes each database to receive replication data from foreign clusters.
Database Replication Lag	The amount of time, in seconds, that the replica database is lagged behind the master database.
List Cache Hit Rate	The number of times per second that queries could use (Hits) the list cache. The average number of hits on the list cache.
List Cache Miss Rate	The number of times per second that queries could not use (Misses) the list cache. The average number of misses on the list cache.
Compressed Tree Cache Hit Rate	The number of times per second that queries could use (Hits) the compressed tree cache. The average number of hits on the compressed tree cache.
Compressed Tree Cache Miss Rate	The number of times per second that queries could not use (Misses) the compressed tree cache. The average number of misses on the compressed tree cache.
Triple Cache Hit Rate	The number of times per second that queries could use (Hits) the triple cache. The average number of hits on the triple cache.
Triple Cache Miss Rate	The number of times per second that queries could not use (Misses) the triple cache. The average number of misses on the triple cache.
Triple Value Cache Hit Rate	The number of times per second that queries could use (Hits) the triple value cache. The average number of hits on the triple value cache.
Triple Value Cache Miss Rate	The number of times per second that queries could not use (Misses) the triple value cache. The average number of misses on the triple value cache.
Reindex Refragment Rate	The rate of reindexing and refragmenting.
Rebalance Rate	The rate of rebalancing.

3.9 Exporting and Printing Monitoring History


You can export and print your monitoring history data.

To export the monitoring history data to an Excel Spreadsheet file, click the Export at the upper-right portion of the Monitoring History page.



The metrics are displayed in separate tabs at the bottom of the spreadsheet.

To print out the charts displayed on the current page, click Print. This will open the printer dialog page from which you can print the charts.



4.0 Telemetry

The MarkLogic telemetry feature provides faster, more complete communication with MarkLogic Support to facilitate the resolution of issues. If enabled, the telemetry feature collects, encrypts, packages, and sends diagnostic and system-level usage information about MarkLogic clusters with minimal impact to performance. Telemetry sends information about your MarkLogic Servers to a protected and secure location where it can be accessed by MarkLogic Support to facilitate troubleshooting and monitor performance. No application data is collected or sent.

Telemetry data is collected from:

- System Error Logs
- Metering Data
- Configuration Data

This chapter describes telemetry and includes the following sections:

- [Understanding Telemetry](#)
- [Configure Telemetry in the Admin UI](#)
- [Example—Telemetry](#)
- [Configure Telemetry With XQuery](#)
- [Baseline System Information](#)
- [Upload a Support Request to Support](#)
- [APIs for Telemetry](#)
- [Interactions With Other MarkLogic Features](#)

4.1 Understanding Telemetry

If telemetry is enabled, MarkLogic Server registers with a well-known endpoint. Data is collected from each host in a cluster: log records from the Error log (`ErrorLog.txt` which contains system logs), monitoring history and usage data (XML documents from the Meters database), and application and host config files from `Data/*.xml`, including license key feature and entitlement info. Personally Identifiable Information (PII) and Application data have been physically split from System data (into application-specific log files), and are not collected. Only System data is collected for Telemetry. Telemetry data is sent to secured cloud storage in a protected location. See “Viewing Monitoring History” on page 33 for more about monitoring history and the Meters database.

The type and granularity of data sent via telemetry is configurable. You can:

- Enable/disable log data: to collect log data; you can select values for the level of logging from fine through debug to emergency levels
- Enable/disable metering history data: to collect metering history data at the level you choose
- Enable/disable host/cluster configuration data: to collect any changes to configuration files

The MarkLogic telemetry feature enables you to send relevant historical data to MarkLogic Support to facilitate troubleshooting and resolution of issues.

4.2 Configure Telemetry in the Admin UI

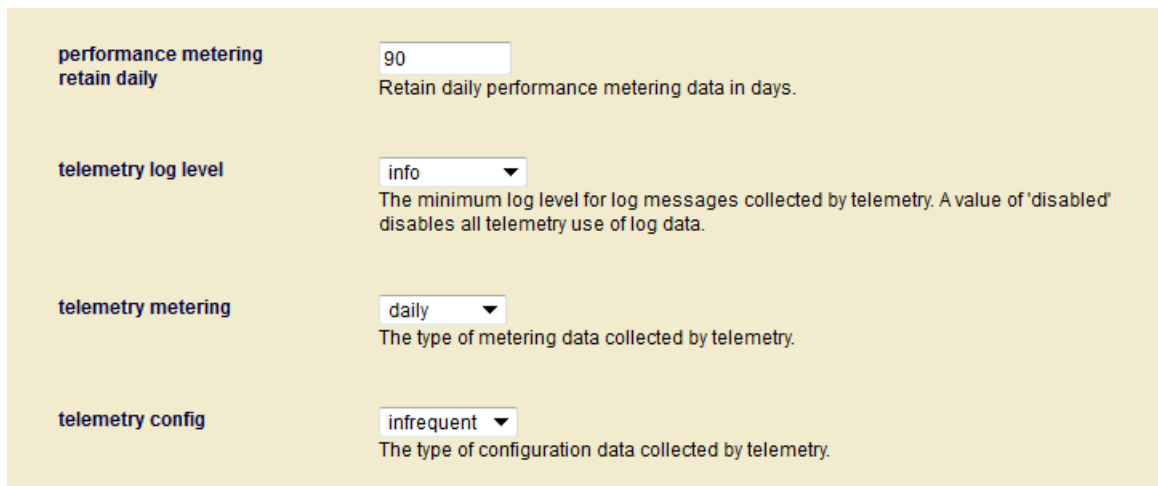
To configure telemetry for your MarkLogic cluster, see the following sections:

- [Enable Telemetry on the Group Configuration Page](#)
- [Telemetry on the Support Page](#)

4.2.1 Enable Telemetry on the Group Configuration Page

To enable telemetry from the Admin UI, navigate to the Group Configuration page and set the telemetry levels for logs, metering, and configuration changes.

1. In the Admin UI, click Groups in the left tree menu.
2. On the Group Configuration page, scroll down to the Telemetry options.



The screenshot shows four configuration options for telemetry on a light yellow background:

- performance metering retain daily**: A text input field containing the value "90". Below it is the text "Retain daily performance metering data in days."
- telemetry log level**: A dropdown menu with "info" selected. Below it is the text "The minimum log level for log messages collected by telemetry. A value of 'disabled' disables all telemetry use of log data."
- telemetry metering**: A dropdown menu with "daily" selected. Below it is the text "The type of metering data collected by telemetry."
- telemetry config**: A dropdown menu with "infrequent" selected. Below it is the text "The type of configuration data collected by telemetry."

Each of these types of telemetry can be configured individually. For example, you can configure telemetry for metering and not configure telemetry for log level or config.

Select the type and level of MarkLogic telemetry you want to enable (logging, metering, or configuration information) from the drop-down menu next to the option.

Telemetry log level: Enable and configure the log messages collected by telemetry. Possible values are finest, finer, fine, debug, config, info, notice, warning, error, critical, alert, emergency, or disabled. The default is disabled. A value of “disabled” disables all telemetry use of log data. We recommend setting log level to debug or config, as these values will allow MarkLogic Support to better assist and troubleshoot tickets.

Telemetry metering: Enable telemetry for metering information. Possible values are raw, daily, hourly, or disabled. The default is disabled. A value of “disabled” disables all telemetry use of metering data. We recommend a setting of hourly or daily, as these values will allow MarkLogic Support to better assist and troubleshoot tickets.

Telemetry config: Enable telemetry for any configuration changes in the cluster. Possible values are frequent, infrequent, or disabled. The default is disabled. A value of “disabled” disables all telemetry use of configuration data. We recommend a setting of infrequent, as these values will allow MarkLogic Support to better assist and troubleshoot tickets.

Note: The level set for each telemetry setting affects the amount of data collected and has a small impact on storage and network use. Telemetry may use up to a maximum of 20GB of storage in the Stage directory, though typically the volume will be less than 100MB.

The system baseline data, which is sent regardless of what type of telemetry is enabled, includes the following information: system data (hostid, clusterid, license) and Standard HTTP Headers added by the system and internet routers (IP address, content type, and so on). Data containing Personally Identifiable Information (PII) is not collected.

So, if telemetry metering is enabled, the Usage and Feature Metrics data will be sent. The actual data in the Usage and Feature Metrics files is very small and consists of two types of data - Usage data (collected once an hour), which includes a few core values like OS, Memory size, License Key, and so on, and Feature Metrics (also collected once an hour), which is a set of name-counter pairs for feature usage.

A feature usage example would be measuring to see if the Optic API is being used. This helps MarkLogic understand the adoption of new features and guide us on how to improve the product.

See [Log Files](#) in the *Administrator's Guide* for more about log files and log level descriptions. See “Monitoring MarkLogic Server” on page 5 for more about metering information.

4.3 Example—Telemetry

This simple example shows how telemetry works by having you configure telemetry and then viewing the information that would be sent, either locally or in a browser.

For this example, perform these steps:

1. Configure telemetry - see “Enable Telemetry on the Group Configuration Page” on page 68 for details.
2. Let your system run for 10-15 minutes to generate data and then change the log level or make another configuration change to your system.
3. To view the output, you can see the staged messages being collected to be uploaded in `/var/opt/MarkLogic/Stage` (Config, Logs, Meters). See “View Staged Telemetry Files” on page 70.

Note: Checking the locally staged files will not work if you have encryption at rest enabled. Encryption settings apply to all staged files, if enabled See [Encryption at Rest](#) in the *Security Guide* for more information.

4.3.1 View Staged Telemetry Files

To see the files being staged locally for telemetry, navigate to `/var/opt/MarkLogic/Stage/Logs`. Type `tail -f ErrorLog.txt |grep -i telemetry` and press return. When telemetry is enabled this will show a running log of the data being that is being collected and uploaded every 5 minutes.


This example shows when data that has been sent to telemetry:

```
[localhost Logs]$ tail -f ErrorLog.txt |grep -i telemetry
2017-02-14 10:40:47.001 Info: Uploaded 47 records, 1 MB of Config Files
to Telemetry
2017-02-14 10:40:47.972 Info: Uploaded 4 records, 1 MB of Meters Data
to Telemetry
2017-02-14 10:40:48.794 Info: Uploaded 64 records, 1 MB of Error Logs
to Telemetry
```

You can also view the data that is being collected in the log files using Query Console. Type the following into the Query Console:

```
xdmp:logfile-scan("/var/opt/MarkLogic/Logs/ErrorLog.txt")
```

The output will look similar to the following:



Note: Checking the locally staged log files will not work if you have encryption at rest enabled for log files. See [Configuration File and Log File Encryption Options](#) in the *Security Guide* for more information.

To see the log data being collected for telemetry (which may differ from the file log `ErrorLog.txt` data due to different log file settings), use this query:

```
xdmp:logfile-scan("/var/opt/MarkLogic/Stage/Logs/ErrorLog.txt")
```

Or more generally:

```
xdmp:logfile-scan(xdmp:data-directory() || "/Stage/Logs/ErrorLog.txt")
```

4.3.2 Encryption of Staged Files

If you have enabled encryption at rest for your log files and/or configuration files, these files will be encrypted while they are staged, prior to being uploaded. The staged files can be found in `/var/opt/MarkLogic/Stage`. The encryption process is transparent to the user. See [Encryption at Rest](#) in the *Security Guide* for more information about file encryption.

4.4 Telemetry on the Support Page

To verify the MarkLogic telemetry data being collected or see a sample of what is being sent, check the Telemetry section on the Support page. To navigate to the Support page in the Admin UI, click the Support tab.

The screenshot shows the MarkLogic Admin UI interface. At the top, there are navigation tabs: Summary, Status, Support (selected), Logs, Usage, and Help. Below the tabs are 'ok' and 'cancel' buttons. The main content area is divided into several sections:

- support request** – Collect information about your system for MarkLogic Support.
 - scope**: host cluster
 - detail**: status only status and system logs status and all logs
 - version**: latest all
 - destination**: browser file upload to MarkLogic Secure Storage
- Telemetry** — For faster response to your call to **MarkLogic Support**, configure **telemetry-log-level** in the **Group Settings page**.
 - Error Logs**: Telemetry is **enabled** at log-level: **info**. Your log is now available to MarkLogic Support when you open a case. Consider setting **telemetry-log-level** in the **Group Settings page** to **debug** so MarkLogic Support has more detail available.
 - Metering**: Telemetry is **enabled** at level: **daily**. Your metering data is now available to MarkLogic Support when you open a case. Consider setting **telemetry-metering** in the **Group Settings page** to **raw** with at least 7 days retention to provide more detail.
 - Configuration**: Telemetry is **enabled** at level: **infrequent**. Your configuration data is available to MarkLogic Support when you open a case.
- Upgrade** — The upgrade logic runs automatically when MarkLogic Server is upgraded, and there is normally no need to run the upgrade logic manually.
 - Upgrade** button

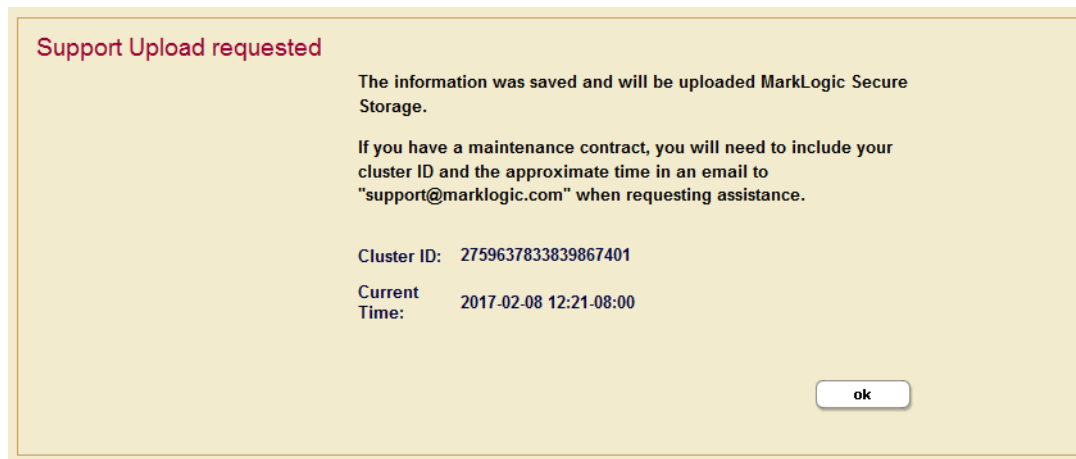
At the bottom of the form are 'ok' and 'cancel' buttons.

The top part of this page is used to create Support Request information, and is independent of Telemetry data files, which are sent continuously once enabled. The Support Request data is a snapshot of the system information that can be sent to MarkLogic secure storage on demand.

In the Telemetry section, you can verify the information that telemetry has collected to send on a scheduled basis, and check the current telemetry status for the different types of data. To change the configuration, click the hyperlink next to each telemetry option. You can also:

- Enable or disable collection of different types of telemetry on the Group Settings page by clicking the link next to each type.
- Check the levels for the types of data to collect for telemetry: log messages, metering data, and configuration information. You can set these on the Group Settings page. See “Enable Telemetry on the Group Configuration Page” on page 68 for details.

After you have configured your telemetry settings on the Group Setting page, the information will be displayed on this page, and collected and sent automatically in the background. You can manually upload support request data to MarkLogic Support by selecting the radio button next to “upload to MarkLogic Secure Storage” and clicking the ok button. The Support Request confirmation screen will be displayed when the upload is complete.



The confirmation includes the cluster ID for the Support Request information and a timestamp. Include this information in your email to MarkLogic Support. Support will also need to know the time of the incident or the when the issue was first noticed. See “Upload a Support Request to Support” on page 75 for more information.

Note: Telemetry data is collected only after you enable the telemetry settings. We recommend that you enable telemetry as a normal practice, so that you have the system information available as part of a Support request should you have an issue.

4.5 Configure Telemetry With XQuery

You can also configure telemetry using the Telemetry APIs. See “APIs for Telemetry” on page 76 for the complete list. Here are a few examples of configuring telemetry using XQuery. You can run these examples in the Query Console. To set the telemetry log level with XQuery:

```
xquery version "1.0-ml";
import module namespace admin = "http://marklogic.com/xdmp/admin"
  at "/MarkLogic/admin.xqy";

let $config := admin:get-configuration()
let $groupid := admin:group-get-id($config, "Default")
let $value := "finest"
```

```

let $tlogconfig := admin:group-set-telemetry-log-level($config,
$groupid, $value)
return
admin:save-configuration($tlogconfig)

```

To check the telemetry log configuration, use this query:

```

xquery version "1.0-ml";
import module namespace admin = "http://marklogic.com/xdmp/admin"
  at "/MarkLogic/admin.xqy";

let $config := admin:get-configuration()
let $groupid := admin:group-get-id($config, "Default")
return
admin:group-get-telemetry-log-level($config, $groupid)
=>
finest

```

To see the type of metering data that is being collected by telemetry, you can use the `admin:group-get-telemetry-metering` function.

```

xquery version "1.0-ml";
import module namespace admin = "http://marklogic.com/xdmp/admin"
  at "/MarkLogic/admin.xqy";

let $config := admin:get-configuration()
return
admin:group-get-telemetry-metering($config,
  admin:group-get-id($config, "Default"))
=>
daily

```

4.6 Baseline System Information

Baseline system information is sent whenever telemetry is on. The system baseline data, which is sent regardless of what type of telemetry is enabled, includes the following information: system data (hostid, clusterid, license) and Standard HTTP headers added by the system and internet routers (IP address, content type, and so on). Data containing Personally Identifiable Information (PII) is not collected.

If telemetry metering is enabled, metering data will be collected as part of the baseline information that is sent.

4.6.1 Metering Data

To view the metering data that is being collected as part of the baseline system information, you can use the Query Console. To see the data do the following:

1. In the Query Console, select the Meters database.
2. Click the Explore button.

3. You will see files ending in “-raw.xml”, “-hourly.xml”, and “-day.xml”. Filenames ending in “-raw.xml” are collected every minute, if telemetry is enabled. Filenames ending in “-hour.xml” are collected hourly, and those ending in “-day.xml” are collected daily. There may also be files ending in “features.xml”.

Note the size of the data that is returned.

- Hour = 60 * minutes
- Day = 24 * hours * 3600 * minutes.

Hourly and Daily metering is very small, while raw data is much larger.

4.7 Upload a Support Request to Support

When you contact Support (<https://help.marklogic.com/>) you may be asked to create a Support Request (Support Dump) for your MarkLogic installation to help diagnose and solve the issue. The Support Request can be configured to send information about the host, or the whole cluster. You can select the level of detail to provide and which version (current or all) of the status and/or system logs to send. The Support Request is a snapshot of a point in time in your MarkLogic environment, and includes status information, log files, configuration files, and other information. Depending on what options you select, the Support Request will collect the latest, or all versions of the selected options (status, logs, config files). See “Telemetry on the Support Page” on page 72 for option details. These files are read from disk and can be reviewed by opening the data directory located at `/var/opt/MarkLogic OR C:\Program Files\MarkLogic\Data`.

Telemetry data files, in addition to the Support Request data mentioned in the previous paragraph, will help MarkLogic Support determine root causes for issues. Telemetry, if enabled, sends a packet of current information on a scheduled basis about a MarkLogic cluster (node) to the secured cloud storage. MarkLogic Support can then retrieve historic information about the MarkLogic environment from the day telemetry was enabled. This helps Support to identify at what point in time the reported issue started and determine the possible root cause. The Support Request snapshot and telemetry stream are independent of one another. However, since telemetry data captures information about your system as snapshots of “current” data on a regular basis, enabling telemetry early before issues arise, is the key to leveraging this feature.

When providing information for a Support request, you should only provide to MarkLogic information that is required to provide Support and which is cleared of confidential or other sensitive information. MarkLogic does not require Protected Health Information (PHI), Payment Card Industry (PCI) information, or Personally Identifiable Information (PII) to provide Support Services and you should not forward any of such types of information to MarkLogic in connection with a Support request. At all times, information provided to MarkLogic in the course of Support will be handled in accordance with the Privacy Policy available here: <http://www.marklogic.com/privacy-policy/>. Similarly, any information collected via telemetric functionality enabled by users of MarkLogic Server will be handled in accordance with the above-referenced Privacy Policy.

To create and upload a Support Request, navigate to the Support tab. When selecting the Support information, click the “upload to MarkLogic Secure Storage” radio button. The contents will be uploaded to a secure server where MarkLogic Support can access the information.

In addition, you will be asked to send an email to Support with your cluster ID and the approximate time and date of the Support Request. This information is available to you as part of the Support Request confirmation screen. Support will also need to know the time the incident first occurred or first was noticed. See “Telemetry on the Support Page” on page 72 for an example of the Support Request Saved confirmation screen, which includes the cluster ID, along with time and date information. You can also find the cluster ID through the Admin UI by clicking on Clusters on the left tree menu. The cluster ID is listed on the Cluster Summary View page.

4.8 APIs for Telemetry

These APIs are available for managing the telemetry functionality, both Admin APIs and REST Management APIs.

4.8.1 Admin APIs

- `admin:group-set-telemetry-log-level`
- `admin:group-get-telemetry-log-level`
- `admin:group-set-telemetry-config`
- `admin:group-get-telemetry-config`
- `admin:group-set-telemetry-metering`
- `admin:group-get-telemetry-metering`
- `admin:group-set-telemetry-session-endpoint`
- `admin:group-get-telemetry-session-endpoint`

4.8.2 REST Management APIs for Telemetry

You can use the REST API for Telemetry management. The REST Management APIs provide the same functionality as the XQuery APIs covered in “Admin APIs” on page 76. The security endpoint is used to by the REST Management APIs to manage telemetry.

```
GET:/manage/v2/security/properties
```

```
PUT:/manage/v2/security/properties
```

```
GET:/manage/v2/logs
```

```
POST:/manage/v2/logs
```

4.9 Interactions With Other MarkLogic Features

The following section describes possible feature interactions between telemetry and other MarkLogic features.

4.9.1 Encryption at Rest

Telemetry follows the encryption settings for the group for all temporary data that is stored on disk as configured for:

“User Data” --> Metering

“Config Data” --> Config Data

“Log Data” --> Error Logs

See “Encryption of Staged Files” on page 71 for more details about the encryption of telemetry data. For more information about encryption, see [Encryption at Rest](#) in the *Security Guide*.

4.9.2 Rolling Upgrades

During a rolling upgrade, before the cluster has been committed to the new version of MarkLogic (9.0-1 or later), any node that has been upgraded to the newer version, will be in read-only mode until the upgrade is committed. This is so that no configuration file changes can be made in a mixed cluster (for example 8.0-7 ->9.0-1) before the upgrade is completed.

4.9.3 Support Uploads

Telemetry settings do not interfere with manually initiated uploads of Support Requests using the Support tab in the Admin UI. These uploads will work whether or not telemetry is enabled.

5.0 Using the Management API

The Management API is a REST-based API that allows you to access MarkLogic Server instrumentation with no provisioning or set-up. The API provides the ability to easily capture detailed information on MarkLogic Server objects and processes, such as hosts, databases, forests, App Servers, groups, transactions, and requests, from a wide variety of tools. The Monitoring Dashboard described in this guide is implemented on top of the Management API.

This chapter describes how to use the Management API to obtain monitoring data from MarkLogic Server. This chapter includes the following sections:

- [Terms used in this Chapter](#)
- [Overview of the Management API](#)
- [Security](#)
- [Management API Requires Writing to the App-Services Database](#)
- [Resource Addresses](#)
- [Obtaining the Options Node for a Resource Address](#)
- [Specifying the Management API Version](#)
- [Specifying Parameters in a Resource Address](#)
- [Interpreting the Output](#)

5.1 Terms used in this Chapter

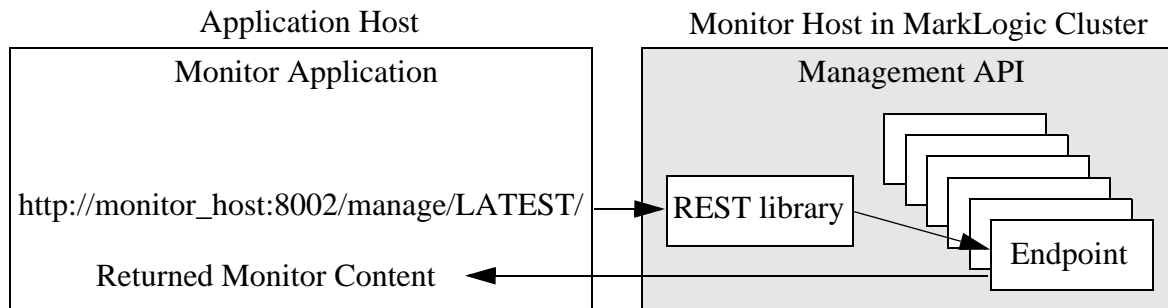
- An *Object* is a component of interest in MarkLogic Server, such as a cluster, host, App Server, or database.
- A *Process* is a request or transaction in MarkLogic Server.
- *Monitor Content* is the XML, HTML, or JSON structure that represents the data returned by the Management API.
- A *Monitor Application* can be any application that requests and makes use of monitoring data, such as a Web browser, a plugin for an existing monitoring tool, or the Monitoring Dashboard described in “Using the MarkLogic Server Monitoring Dashboard” on page 14.
- The *Monitor Host* is the host in the MarkLogic Server cluster that exposes the Management API to respond to requests for monitoring content from a monitor application.
- The *Manage App Server* is the App Server on the Monitor Host that is configured to handle monitor requests. The Manage App Server is bound to port 8002 and is the App Server used by the Monitoring Dashboard.

- *REST* stands for *Representational State Transfer*, which is an architectural style that, in the context of monitoring MarkLogic Server, describes the use of HTTP to make calls between a monitoring application and monitor host.
- A *Resource* is an abstraction of a MarkLogic Server object, as represented by the REST architecture.
- A *Resource Address* is a URL that identifies a MarkLogic Server resource. The resource addresses are described in “Resource Addresses” on page 81.
- A *View* is the returned monitoring information about a resource. You can have different views of the same resource. A view can be for a single resource (known as an *item view*) or a number of resources (known as a *list view*).
- A *Representation* is a view of a resource in a particular format, such as XML, HTML, or JSON.
- A *Parameter* is an addition to the end of a resource address to filter and/or format the view returned from MarkLogic Server. Parameters are expressed as query strings in the URL and are described in “Specifying Parameters in a Resource Address” on page 83.
- An *Endpoint* is an XQuery module on MarkLogic Server that is invoked by and responds to an HTTP request for monitoring information.
- A *Plugin* is an XQuery module that provides extension capabilities using the Plugin framework described in the [System Plugin Framework](#) chapter in the *Application Developer’s Guide*.

5.2 Overview of the Management API

The Management API is implemented on top of the REST Library described in [Creating an Interpretive XQuery Rewriter to Support REST Web Services](#) in the *Application Developer's Guide*. Requests to monitor an object in MarkLogic Server are made by means of a resource address that returns a view containing the monitor data for the object. The view can be returned in various formats, such as XML, HTML, or JSON.

Every resource address in the Management API invokes a monitoring endpoint, which is an XQuery module on the target MarkLogic Server host. The monitoring endpoints are invoked by a resource address in an application, such as a browser. The Management API uses the REST library to validate the request, authorize the user, and rewrite the resource address to one understood by the monitoring framework before invoking the endpoint module. The endpoint module returns the monitoring data for the resource to the application.



5.3 Security

As described in “Monitoring Tools and Security” on page 7, client access to a Management API and endpoints requires that they authenticate as a user with the `manage-user` role. If custom Plugin code requires additional privileges, you can create and assign a custom role to users of the Management API to enable that functionality.

If you have enabled SSL on the `manage` App Server, your resource address must start with HTTPS, rather than HTTP, and you must have a MarkLogic certificate authority on your browser, as described in [Accessing an SSL-Enabled Server from a Browser or WebDAV Client](#) in the *Security Guide*.

5.4 Management API Requires Writing to the App-Services Database

The Management API sometimes writes documents to the App-Services database for internal purposes, and it therefore assumes that the App-Services database is writable. On each cluster in which the Management API runs (which might include a replica cluster), it requires a writable view of the App-Services database. The App-Services database is used to store data used by various MarkLogic applications, such as Query Console. For these reasons, MarkLogic recommends that you do not replicate the App-Services database using database replication. If the App-Services database is not available, it falls back to the schemas-database configured for the schemas-database of the App Server under which the Management API is running.

5.5 Resource Addresses

This section provides an overview of the structure and capabilities of the resource addresses provided by the Management API. For details on each resource address, see the *MarkLogic REST API Reference*.

Resource addresses fall into the categories listed in the table below. The output from each type of resource address is described in “Interpreting the Output” on page 85.

Type of Resource Address	Returns
List	A list of resources. For example, as list of the forests in the cluster: <code>http://localhost:8002/manage/LATEST/forests</code>
Item	A specific resource. For example, a specific forest in the cluster: <code>http://localhost:8002/manage/LATEST/forests/Documents</code>

A resource address takes the form of a URL that includes a host name and a port number. The most basic resource address returns summary information for the entire cluster:

```
http://host:port/manage/LATEST/
```

The Management API version, `LATEST` in this release, is specified in every resource address to maintain compatibility with future revisions of the Management API.

You can optionally include the name of a resource and parameters in a resource address as follows:

```
http://host:port/manage/version/resource?param=value&param=value
```

5.6 Obtaining the Options Node for a Resource Address

As described in [Creating an Interpretive XQuery Rewriter to Support REST Web Services](#) in the *Application Developer's Guide*, the REST Library uses an `options` node to map incoming requests to endpoints. The `options` node contains information about the communication options available on the request/response chain for the resource address, such as which parameters can be specified with the resource address.

You can use the `xdmp:http-options` function to output the `options` node for any resource address. For example, you can enter the following query in Query Console to display the `options` node for the `/manage/LATEST/transactions` resource address:

```
xdmp:http-options (
  "http://localhost:8002/manage/LATEST/transactions",
  <options xmlns="xdmp:http">
    <authentication method="digest">
      <username>admin</username>
      <password>admin</password>
    </authentication>
  </options>)
```

The output will include a `request` element that defines the options associated with the GET and HEAD methods for the resource address. From this, you can determine the supported parameters and values. For example, the above resource address supports the `view`, `seconds-min`, `host-id`, `fullrefs`, and `format` parameters, as shown below.

```
<rest:http method="GET">
  <rest:param name="view" values="default" default="default"/>
  <rest:param name="seconds-min" as="string"/>
  <rest:param name="host-id" as="string"/>
  <rest:param name="fullrefs" as="boolean" required="false"/>
  <rest:param name="format" as="string" values="xml|json|html"/>
  <rest:or>
    <rest:accept>application/xml</rest:accept>
    <rest:accept>application/json</rest:accept>
    <rest:accept>text/html</rest:accept>
    <rest:accept>application/x-javascript</rest:accept>
  </rest:or>
</rest:http>
```

5.7 Specifying the Management API Version

To guarantee stable behavior of the Management API as new versions are released, each resource address in the Management API includes a version number. The examples in this chapter show the version as `LATEST`, which means to use the latest version of the API. However, you can also specify the version number to use a specific version of the API, using the format:

```
v#
```

Where # is the version number. For example, in the initial version of the API:

```
http://localhost:8002/manage/LATEST/databases
```

is the same as:

```
http://localhost:8002/manage/v2/databases
```

If you want to update your clients when you choose, use the explicit version number. If you want to update your clients to the most recent version of the Management API, use `LATEST`.

Note: The version number is only updated when resource addresses and/or parameters have changed. It is not updated when resource addresses and/or parameters are added or removed.

5.8 Specifying Parameters in a Resource Address

Resource addresses can take parameters to do the following:

- Specify the format of the returned view
- Return a filtered view

To specify multiple parameters, use the ‘?’ sign before the first parameter and the ‘&’ sign before any additional parameters:

```
http://host:port/manage/LATEST/resource?param1=value&param2=value...
```

Some resource addresses support optional parameters that are specific to that resource address. For example, to return monitoring information on the forests used by the `Documents` database, you can use the `database-id` parameter with the `/forests` resource as follows:

```
http://monitor_host:8002/manage/LATEST/forests?database-id=Documents
```

The remainder of this section describes the `format` parameter in more detail.

5.8.1 Formatting the Monitor Results

The application that issues a request to the Management API specifies the format for the returned view. For example, most Web browsers specify the default return format as HTML. If no return format is specified by the application, the view is formatted as XML. You can explicitly specify the view format by means of the `format` parameter at the end of the resource address:

```
format=value
```

Where *value* is either HTML, JSON, or XML.

The XML and JSON formats provide a rich set of data for your monitoring application. For example, to return an XML view of the entire cluster, you can enter the following in a browser:

```
http://monitor_host:8002/manage/LATEST?format=xml
```

This will return a `cluster` view in XML format. For example:

```
<cluster-default xsi:schemaLocation=
  "http://marklogic.com/manage/clusters manage-clusters.xsd">
  <meta>
    <uri>/manage/LATEST</uri>
    <current-time>2011-06-30T16:00:06.81-07:00</current-time>
    <elapsed-time units="sec">0.012</elapsed-time>
  </meta>
  <relations>
    <relation-group array="true">
      <uriref>/manage/LATEST/databases</uriref>
      <typeref>databases</typeref>
      <relation-count>13</relation-count>
    </relation-group>
    <relation-group array="true">
      <uriref>/manage/LATEST/forests</uriref>
      <typeref>forests</typeref>
      <relation-count>13</relation-count>
    </relation-group>
    <relation-group array="true">
      <uriref>/manage/LATEST/groups</uriref>
      <typeref>groups</typeref>
      <relation-count>1</relation-count>
    </relation-group>

    <relation-group array="true">
      <uriref>/manage/LATEST/hosts</uriref>
      <typeref>hosts</typeref>
      <relation-count>1</relation-count>
    </relation-group>
    <relation-group array="true">
      <uriref>/manage/LATEST/requests</uriref>
      <typeref>requests</typeref>
      <relation-count>1</relation-count>
    </relation-group>
```

```

<relation-group array="true">
  <uriref>/manage/LATEST/servers</uriref>
  <typeref>servers</typeref>
  <relation-count>8</relation-count>
</relation-group>
<relation-group array="true">
  <uriref>/manage/LATEST/transactions</uriref>
  <typeref>transactions</typeref>
</relation-group>
</relations>
<related-views>
  <related-view array="true">
    <view-type>item</view-type>
    <view-name>query</view-name>
    <view-uri>/manage/LATEST/query</view-uri>
  </related-view>
  <related-view array="true">
    <view-type>item</view-type>
    <view-name>status</view-name>
    <view-uri>/manage/LATEST?view=status</view-uri>
  </related-view>
</related-views>
</cluster-default>

```

5.9 Interpreting the Output

The reference documentation for the Management API in the *MarkLogic REST API Reference* describes each element in the XML output for each Management API resource address.

The main elements in the output from each resource address for an item or item list are shown in the table below.

Type of Resource Address	Element	Description
Item and Item View	id	The item id number.
Item and Item View	name	The item name (not available for requests and transactions).
Server Item	server-kind	The type of App Server (http, WebDAV, or XDBC)
Item, Item List, and View	meta	Metadata that describes: <ul style="list-style-type: none"> The URI of the resource. The current timestamp. The number of seconds it took to execute the resource address.

Type of Resource Address	Element	Description
View	<code>view-properties</code>	The properties of the item or item list view.
Item and Item List	<code>relations</code>	The items that are related to this item or item list.
Item List	<code>list-items</code>	The items in this item list
Item, Item List, and View	<code>related-views</code>	The views related to this item or item list. If an item, the item list view is also included.

6.0 Technical Support

MarkLogic provides technical support according to the terms detailed in your Software License Agreement or End User License Agreement.

We invite you to visit our support website at <http://help.marklogic.com> to access information on known and fixed issues, knowledge base articles, and more. For licensed customers with an active maintenance contract, see the [Support Handbook](#) for instructions on registering support contacts and on working with the MarkLogic Technical Support team.

Complete product documentation, the latest product release downloads, and other useful information is available for all developers at <http://developer.marklogic.com>. For technical questions, we encourage you to ask your question on [Stack Overflow](#).

7.0 Copyright

MarkLogic Server 9.0 and supporting products.
Last updated: April 28, 2018

COPYRIGHT

Copyright © 2018 MarkLogic Corporation. All rights reserved.
This technology is protected by U.S. Patent No. 7,127,469B2, U.S. Patent No. 7,171,404B2, U.S. Patent No. 7,756,858 B2, and U.S. Patent No 7,962,474 B2, US 8,892,599, and US 8,935,267.

The MarkLogic software is protected by United States and international copyright laws, and incorporates certain third party libraries and components which are subject to the attributions, terms, conditions and disclaimers set forth below.

For all copyright notices, including third-party copyright notices, see the Combined Product Notices for your version of MarkLogic.