
MarkLogic Server

Ops Director Guide

MarkLogic 10
May, 2019

Ops Director Version: 2.0
Last Revised: 10.0-2/2.0.1-1, September 2019

Please see <http://developer.marklogic.com/products/opsdirector> for the most up-to-date documentation and binaries

Table of Contents

Ops Director Guide

1.0	Monitoring MarkLogic with Ops Director	7
1.1	Overview	7
1.2	Terms	8
1.3	New Features in Ops Director	9
1.4	Required Software	10
1.5	Ops Director High-Level Architecture	10
1.6	Ops Director Deployment Configurations	11
1.7	Ops Director Security	12
1.7.1	Communication with Managed Clusters	13
1.7.1.1	Security and Database Dependencies of Managed Clusters	13
1.7.2	Resource Groups	14
1.7.3	Access Inheritance in Resource Groups	14
1.7.4	Execute Privileges	15
1.8	Data Collected by Ops Director	15
1.8.1	Log Data	15
1.8.2	Configuration Data	16
1.8.3	Metering Data	16
1.8.3.1	Explicit Limits on Type and Rates	16
1.8.3.2	Security of Collected Data	17
1.8.3.3	Unattended Operation	17
1.8.3.4	“Best Attempt” SLA	17
1.8.3.5	Data Type Separation	17
1.8.3.6	Data Transmission	18
2.0	Installing and Configuring Ops Director	19
2.1	System Requirements	19
2.2	Installing Ops Director	20
2.2.1	Installing Ops Director version 1.x on an Application Cluster	20
2.2.2	Installing Ops Director 2.0 and Greater	28
2.2.2.1	Before You Install Ops Director	28
2.2.2.2	Ops Director 2.0 Installation	29
2.2.2.3	Ops Director Resources	30
2.3	Connecting a Managed Cluster to Ops Director	31
2.4	Launching Ops Director and Logging In	36
2.5	Switching from Digest to Application-Level Authentication	36
2.6	Configuring the Ops Director Data Retention Policy	38
2.7	Disconnecting a Managed Cluster from Ops Director	40
2.8	Removing Ops Director	42

2.8.1	Removing Ops Director 1.1 and Earlier	42
2.8.2	Removing Ops Director 2.0 and Later	43
2.9	Running Ops Director on Amazon Web Services (AWS)	43
2.10	Securing Ops Director with Externally Signed Certificates	44
2.10.1	Using Externally Signed Certificates with MarkLogic Server 10.0-2 and Later and Ops Director 2.0.1-1	45
2.10.1.1	Configure the Ops Director Cluster	45
2.10.1.2	Configure the Managed Cluster	48
2.10.2	Using Externally Signed Certificates with MarkLogic Server 9.0-9 and Later and Ops Director 2.0.1 and Later	54
2.10.2.1	Configure the Ops Director Cluster	54
2.10.2.2	Configure the Managed Cluster	58
2.10.3	Using Externally Signed Certificates with MarkLogic 9.0-8 and earlier or Ops Director 2.0 and Earlier	64
2.10.3.1	Configure the Ops Director 2.0 and Earlier Cluster	64
2.10.3.2	Manage a Cluster using Ops Director 2.0 and Earlier	68
2.11	Upgrading Ops Director	75
2.11.1	Upgrade Process Overview	75
2.11.2	MarkLogic Server and Ops Director Upgrade Version Compatibility	75
2.11.3	Upgrade Scenarios and Workflows	76
2.11.4	Upgrade Error Prevention and Troubleshooting	79
3.0	Navigating and Filtering Ops Director Views	81
3.1	Main Navigation Bar	81
3.2	Navigating Resource Views	82
3.2.1	View All Resources	82
3.2.2	View Resource Groups	84
3.3	Date and Time Filters	86
3.4	Navigation Icons of Ops Director Views	88
3.5	Preserving View States	89
4.0	MONITOR View	91
4.1	Monitoring Dashboards	91
4.1.1	Key Performance Indicators	91
4.1.2	Cluster Problem Distribution	92
4.1.3	Top Problematic Hosts	93
4.1.4	Busiest Servers	93
4.1.5	Slowest Servers	95
4.1.6	Alerts Panel	96
4.1.7	Filtering by Resource	97
4.2	Overview Pages	98
4.2.1	Clusters Overview	98
4.2.2	Hosts Overview	100
4.2.3	Databases Overview	102
4.2.4	App Servers Overview	105

5.0	MANAGE View	109
5.1	Manage Clusters Tab	110
5.1.1	Cluster Metrics	113
5.1.2	Cluster Information	115
5.1.3	Cluster Tasks	116
5.1.4	Cluster Logs	118
5.1.5	Cluster Hosts	119
5.1.6	Cluster Databases	120
5.1.7	Cluster App Servers	122
5.1.8	Cluster Forests	123
5.2	Manage Hosts Tab	125
5.2.1	Host Metrics	129
5.2.2	Host Information	131
5.2.3	Host Tasks	132
5.2.4	Host Logs	134
5.2.5	Host Databases	135
5.2.6	Host App Servers	137
5.2.7	Host Forests	139
5.3	Manage Databases Tab	141
5.3.1	Database Metrics	144
5.3.2	Database Information	146
5.3.3	Database Tasks	165
5.3.4	Database Hosts	168
5.3.5	Database App Servers	169
5.3.6	Database Forests	171
5.4	Manage App Servers Tab	173
5.4.1	App Server Metrics	175
5.4.2	App Server Information	176
5.4.3	App Server Hosts	185
5.4.4	App Server Databases	186
6.0	ANALYZE View	189
6.1	Configuring and Navigating the ANALYZE View	189
6.2	Performance Charts by Resource	194
6.2.1	Disk Performance Data	195
6.2.2	CPU Performance Data	197
6.2.3	Memory Performance Data	200
6.2.4	Server Performance Data	202
6.2.5	Network Performance Data	204
6.2.6	Database Performance Data	207
7.0	CONSOLE SETTINGS View	215
7.1	Role Based Access Control (RBAC) Settings	215
7.1.1	Roles Tab	216
7.1.2	Resource Access Tab	217

7.1.3	Creating a Resource Group and Assigning it to a Role	217
7.1.4	Editing or Deleting a Role	222
7.2	Resource Groups	223
7.2.1	Creating a Resource Group	224
7.2.2	Editing or Deleting a Resource Group	226
7.2.3	Resource Group Views	227
7.2.3.1	Host Groups	227
7.2.3.2	Database Groups	228
7.2.3.3	App Server Groups	230
7.2.3.4	Cluster Groups	231
7.3	License Information	232
7.3.1	License Information By Host	233
7.3.2	License Information By License	235
7.4	Managed Clusters	236
7.4.1	Viewing and Filtering the List of Managed Clusters	236
7.4.2	Removing Unknown Managed Clusters from the List	238
7.4.3	Reconnecting a Managed Cluster to Ops Director	239
7.5	Configuring Email Notifications	242
7.5.1	Setting Up Email Notifications	242
7.5.2	Setting up and Managing Alerts	244
7.5.2.1	Creating an Email Alert	245
7.5.2.2	Editing an Email Alert	248
7.5.2.3	Viewing and Filtering the List of Alerts	249
7.5.2.4	Deleting an Email Alert	250
8.0	SUPPORT view	251
8.1	System Alerts	251
8.1.1	Alert Levels	255
8.1.2	Point-in-Time Alerts	257
8.2	Event Logs	257
8.3	Task Console	260
9.0	Troubleshooting with Ops Director	265
9.1	Assess Whether MarkLogic Has Adequate Resources	265
9.2	Assess the Overall State of the System	266
9.3	Assess MarkLogic Cluster Performance	267
9.4	Assess Severity of Problems in the System	269
10.0	Technical Support	271
11.0	Copyright	273
11.0	COPYRIGHT	273

1.0 Monitoring MarkLogic with Ops Director

Ops Director enables you to monitor MarkLogic clusters ranging from a single node to large multi-node deployments. A single Ops Director server can monitor multiple clusters. Ops Director provides a unified browser-based interface for easy access and navigation.

Ops Director presents a consolidated view of your MarkLogic infrastructure, to streamline monitoring and troubleshooting of clusters with alerting, performance, and log data. Ops Director provides enterprise-grade security of your cluster configuration and performance data with robust role-based access control and information security powered by MarkLogic Server.

This chapter covers the following topics:

- [Overview](#)
- [Terms](#)
- [New Features in Ops Director](#)
- [Required Software](#)
- [Ops Director High-Level Architecture](#)
- [Ops Director Deployment Configurations](#)
- [Ops Director Security](#)
- [Data Collected by Ops Director](#)

1.1 Overview

Ops Director is designed to accommodate your evolving IT strategy, whether for one or many clusters, on premises or the cloud. Ops Director complements the current MarkLogic admin tools, bringing enterprise IT administrators a simple, flexible, and proactive management experience.

Ops Director decreases your learning curve by using guided visual representations of databases, clusters, and applications. It enables you to identify potential problems and bring them to attention before they occur. It also provides learning and analysis opportunities with centralized data collection, delivery, and storage.

You can use Ops Director for the following tasks:

- To keep track of the day-to-day operations of your MarkLogic Server environment.
- To plan initial capacity and fine-tune 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 and 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 Terms

- *Ops Director Application Cluster* — a MarkLogic cluster that contains the host that runs the Ops Director application.
- *Managed Cluster* — a MarkLogic cluster that is specifically configured to be managed by the Ops Director application.
- *Ops Director Application* — Ops Director application server that is responsible for communication with the browser.
- *Ops Director System* — Ops Director application server that is responsible for inter-cluster communication.
- *Resource Group* — a configuration that represents one or more resources of a certain type, such as hosts, App Servers, databases, and so on.
- *RBAC* — Role-Based Access Control — a method of regulating access to computer or network resources based on the roles of individual users within an enterprise.
- *XDQP* — XML Data Query Protocol — a MarkLogic internal protocol used for communication between nodes in a cluster.
- *CSV* — Comma-Separated Values file stores tabular data in a plain text format, where each line of the file is a data record, each record consists of one or more fields, and the fields are separated by commas.
- *View* — a top-level page of Ops Director UI, which you can access from the menu bar; provides a high-level view of MarkLogic resources.
- *Tab* — a next-level page of Ops Director UI, which you can access from clicking on one of the tabs in a view; enables you to drill down for specific resources or resource types.
- *Report* — a document-centric display of information about one or more assets at a specific date time period.
- *Rates* — The amount of data (MB per second) currently being read from or written to a resource.
- *Loads* — The execution time (in seconds) of current read and write requests on a resource, which includes the time requests spend in the wait queue when maximum throughput is achieved.

- *Color-coded severity* — colors used in graphic representations of alert severity, as indicated in the following table:

Color	Alert Severity
Red	Critical
Yellow	At Risk / Warning
Green	Healthy
Dark Gray	Maintenance
Light Gray	Offline
Blue	Security
Dark Green	Information
White / Hollow	Unknown

1.3 New Features in Ops Director

The following features have been added to Ops Director 2.0.1:

- **Email notification improvements**
 - You are now sent a notification email when you create, are subscribed to, or unsubscribed from an alert.
 - You are now sent a notification email when an alert is disabled or enabled.
 - Email alerts now contain the resource and severity of the trigger.

For more information about email alerts, see [Configuring Email Notifications](#).

- **Streamlined the process of using Ops Director with externally signed certificates**

For more information, see [Securing Ops Director with Externally Signed Certificates](#).

1.4 Required Software

The following table shows which versions of MarkLogic Server are compatible with which versions of Ops Director, as well as which versions of Ops Director can be upgraded under which versions of MarkLogic Server:

If Running MarkLogic Server Version	Can Install Ops Director Version
9.0-4	1.0-0
9.0-5 and 9.0-6	1.0-0 and 1.1-1
9.0-7 or 9.0-8	2.0.0 and 2.0.1
9.0-9 and later	2.0.1 only.
10.0-2	2.0.1-1

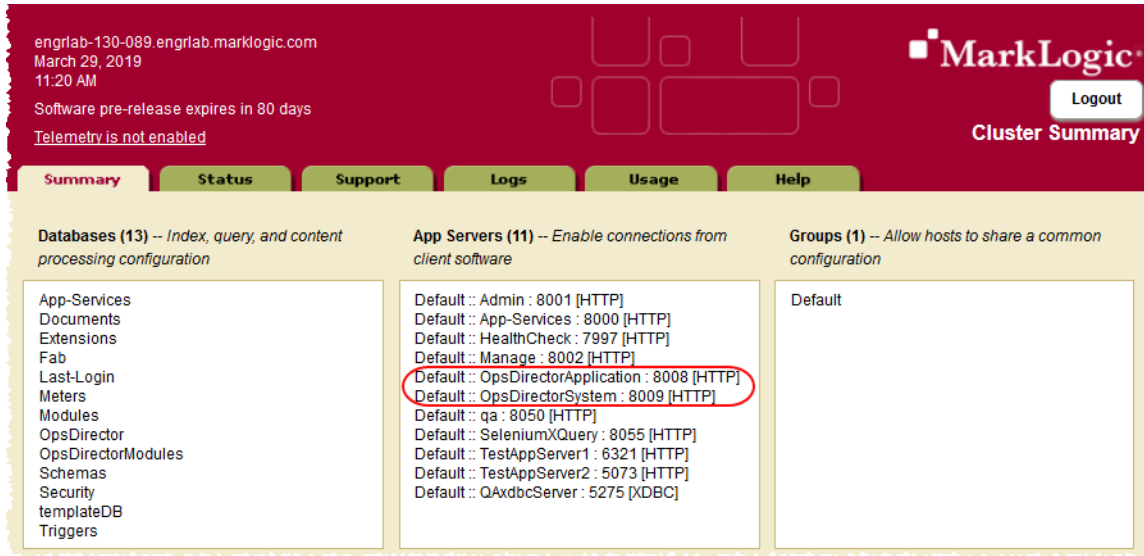
Note: Ops Director 2.0 and greater will not work with versions of MarkLogic Server prior to 9.0-7, and Ops Director 2.0.1 will not work with MarkLogic Server version 10.0 and later.

1.5 Ops Director High-Level Architecture

An Ops Director instance includes two application servers:

- *Ops Director Application* server, which by default runs on port 8008, provides services to the Ops Director browser application.
- *Ops Director System* server, which by default runs on port 8009, receives data transmitted from Managed Clusters and stores it in the Ops Director database.

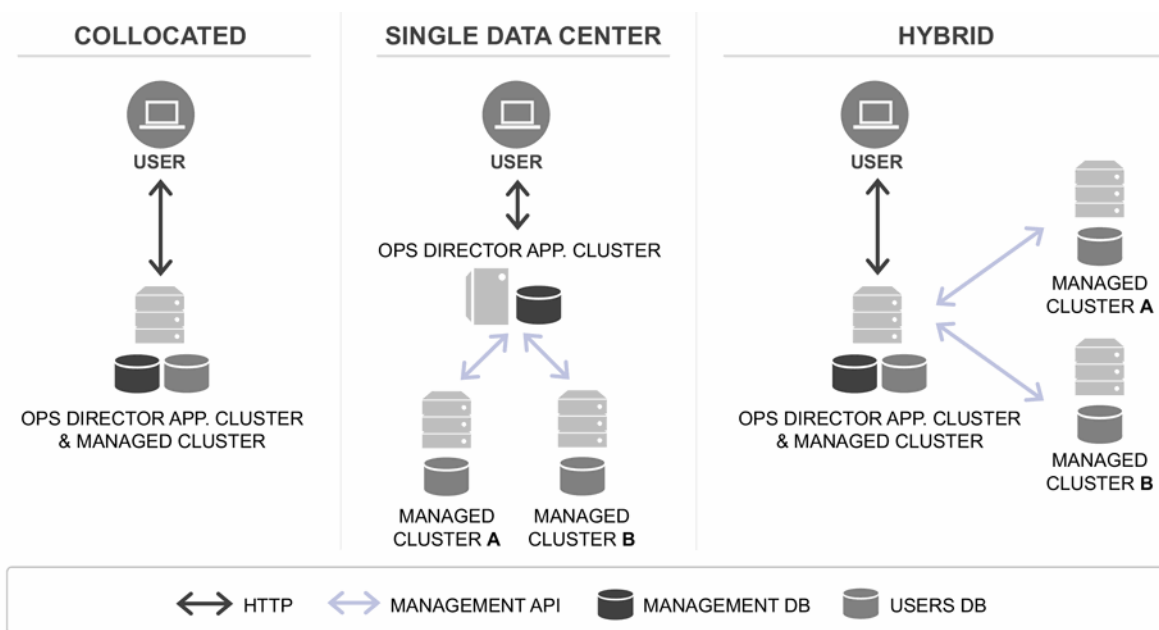
When Ops Director application is installed on a cluster, you can view these two application servers in the Summary page of the Admin Interface.



1.6 Ops Director Deployment Configurations

A single Ops Director instance can manage multiple clusters. Managed clusters collect information about the state of their cluster (error logs, configuration data, and meters) and securely send that data to the Ops Director System. The Ops Director System receives this information and stores it in the Ops Director database. The Ops Director Application executes queries against the Ops Director database to retrieve information for display in the browser.

Ops Director is designed to adapt to a range of MarkLogic deployments, from a single cluster in a single data center to hundreds of clusters. Ops Director has three possible deployment configurations:



- **Collocated:** A cluster serves as both an Ops Director Application Cluster and a Managed Cluster.
- **Single Data Center:** A cluster serves as an Ops Director Application Cluster that communicates with Managed Clusters.
- **Hybrid:** A cluster serves as both an Ops Director Application Cluster and a Managed Cluster that communicates with Managed Clusters.

A multi-cluster configuration can manage local and remote clusters across a network as long as all of the clusters are available and not concealed behind a firewall. All communication between an Ops Director Application Cluster and Managed Clusters is performed securely.

Note: Ops Director can manage only clusters that run either the same version of MarkLogic Server as the version it was shipped with or one of the older versions. This limitation has an impact on colocating Ops Director on clusters with other applications (hybrid deployment configuration). You must not put Ops Director on a cluster with other applications that you would be reluctant to upgrade to the latest version of MarkLogic Server. If you choose to run Ops Director in a hybrid environment, make sure to put Ops Director on a cluster that you are willing and able to upgrade first, because Ops Director will not be able to manage any upgraded clusters until Ops Director Application Cluster has been upgraded.

1.7 Ops Director Security

In a multi-cluster environment, where Ops Director provides alerting, management, and reporting capabilities, different administrators have different goals. To ensure that each type of administrator can achieve their goals without causing harm or accessing information to which they are not supposed to have access, Ops Director employs comprehensive role-based access controls. The assets and views displayed to an administrator are based upon group participation and the group's roles. Each role has specific privileges. The combination of an administrator's role, resource group membership, and the asset determines which capabilities are available to that administrator.

When a Managed Cluster connects to Ops Director, it grants Ops Director the ability to perform actions by making requests back to the Managed Cluster. These requests are protected by certificate-based authentication. User IDs and passwords are not sent across the network. Ops Director enables you to manage clusters on which you have no direct login capability.

This section provides a conceptual overview of Ops Director security. The procedures for configuring Ops Director security are described in “Installing and Configuring Ops Director” on page 19 and “CONSOLE SETTINGS View” on page 215.

The main topics in this section are:

- [Communication with Managed Clusters](#)
- [Resource Groups](#)
- [Access Inheritance in Resource Groups](#)
- [Execute Privileges](#)

1.7.1 Communication with Managed Clusters

When a cluster agrees to be managed by Ops Director, it establishes a long-lived secure communication channel between itself and Ops Director using certificates. Ops Director requests signed with the appropriate certificate can communicate with the Managed Cluster and perform admin tasks on the cluster. The certificate authority for these certificates can be either self-signed, as described in “Installing Ops Director” on page 20, or externally signed, as described in “Securing Ops Director with Externally Signed Certificates” on page 44.

For a particular request, Ops Director determines which roles and privileges are in effect for that request. It then makes the request to the Managed Cluster, using a short-lived certificate, to perform some action. In the request, Ops Director includes an identifier for the request.

The Managed Cluster passes that identifier back to Ops Director, over the secure channel, to obtain the set of roles and privileges that apply for the request. Once that context is established, it is used to perform the request, which will succeed or fail on its own merits.

The procedure for establishing secure certificate-based communication between Ops Director and the Managed Clusters is described in “Installing and Configuring Ops Director” on page 19.

1.7.1.1 Security and Database Dependencies of Managed Clusters

When you connect a managed cluster to Ops Director, a SecureManage application server is created for all groups that contain hosts from this cluster.

Each SecureManage application server is a copy of the ManageServer, with support for SSL and certificate-based authentication. This allows Ops Director to communicate with managed clusters using SSL without knowing any user credentials on the managed cluster.

The SecureManage application server uses the App-Services database and the database's forests. If you disable or delete the SecureManage application server or any of its resources (such as the App-Services database and the forests for that database) on a managed cluster, Ops Director will not be able to retrieve information from this managed cluster, and the cluster will be assigned the Unknown state. Likewise, if you remove an External Certificate from a managed cluster, Ops Director will not be able to retrieve information from this managed cluster, and the cluster will be assigned the Unknown state.

1.7.2 Resource Groups

In a large environment, it is useful to group resources together, for example, “the staging hosts,” or “the production databases.” In Ops Director, these are called *resource groups*. Resource groups are homogenous, which removes complexity from the meaning of “apply this action to this group.” For example, “show me statistics for the production databases group.”

Each resource group consists of the following:

- `name` — The user-visible name for the resource group.
- `type` — The type of resources contained in the resource group.
- `resources` — The list of resources in the resource group.
- `role` — An optional role or roles that control access to this resource group.

The resource group role enables you to establish access to a resource group at a finer and more ad hoc granularity than is provided by the established roles. It is likely that roles defined within the enterprise are fairly coarse-grained and that changing roles (in an external LDAP server, for example), may be considered too heavy weight for ad hoc groupings. For more details on roles and resource groups and how to create them, see “CONSOLE SETTINGS View” on page 215.

You can configure a resource group so it grants additional privileges within the context of that group.

1.7.3 Access Inheritance in Resource Groups

Resource groups consist of a single resource type and an explicitly enumerated list of resources.

Inheritance in resource groups extends the reach of a single resource group across multiple resource types, as follows:

- A cluster in a cluster resource group inherits all of the resources in the cluster (hosts, application servers, databases, and forests).
- A host in a host resource group inherits all of the databases and forests on that host.
- A database in a database resource group inherits all the forests in the database.
- An application server in a server resource group inherits the databases that the server uses.

Inheritance is transitive. If a server resource group gives you access to a database, you have access to its forests. The following inheritance access rules apply in Ops Director:

- If you can access a cluster, you can access all resources of this cluster: hosts, application servers, databases, and forests.
- If you can access a host, you can access all databases and forests of that host.
- If you can access a database, you can access all forests of that database.

- If you can access an application server, you can access all databases that this application server uses and all forests of these databases.

1.7.4 Execute Privileges

The following privileges are specific to Ops Director:

- <http://marklogic.com/xdmp/privileges/opsdir-admin> — Grants Ops Director Administrator privileges.
- <http://marklogic.com/xdmp/privileges/opsdir-license-admin> — Grants privileges to manage licenses in Ops Director.
- <http://marklogic.com/xdmp/privileges/opsdir-user> — Grants privileges to use Ops Director.

These execute privileges are pre-defined and included with every installation of MarkLogic Server. You can view them in the Execute Privileges Summary page of the Admin Interface.

1.8 Data Collected by Ops Director

Once configured, Ops Director begins collecting data from all the hosts in the Managed Clusters and storing it the Ops Director database.

The following types of data are collected:

- [Log Data](#)
- [Configuration Data](#)
- [Metering Data](#)

1.8.1 Log Data

Log messages generated at or above the level specified for delivery to Ops Director will be queued and sent as quickly as possible. Errors generated at a level of critical or higher will block while they are being sent to make sure that a delivery attempt is made before the host restarts.

Server log messages are filtered by the log level. You can configure the minimum log level for log messages sent to Ops Director via the Admin Interface, Ops Director Setup page. For example, if you configure the minimum log level to “error”, then you will see all messages starting from “error” level.

The minimum log level can be configured as:

- fine
- debug
- config

- info
- notice
- warning
- error
- critical
- alert
- emergency

Note: It is not recommended to configure the minimum log level to a finer granularity than info.

1.8.2 Configuration Data

Whenever the configuration of a Managed Cluster changes, Ops Director is notified.

When Ops Director is notified of a configuration change, if the change is more recent than the local data that Ops Director has for a particular configuration, Ops Director retrieves the necessary payloads from the Managed Clusters in the form of *resource documents*. This configuration data includes any changes made to server configuration files. Ops Director leverages this interaction to obtain information about every applicable resource on each Managed Cluster: groups, hosts, databases, application servers, and so forth.

Timestamps allow Ops Director to maintain a history of the configuration of Managed Clusters over time.

Ops Director calls a host on the Managed Cluster to get the modified configurations. Any new or changed configuration is saved with a new timestamped URI. Efficient access to the most recent configuration is managed with properties.

1.8.3 Metering Data

Each Managed Cluster sends metering data (such as documents from the Meters database) to Ops Director. Filtering of metered data is by time (raw, hourly, or daily). The managed hosts introduce a small, random adjustment factor in the actual intervals to avoid the situation where every managed host transmits to Ops Director at exactly the same moment every time.

1.8.3.1 Explicit Limits on Type and Rates

You can configure Managed Cluster to define the type of data sent to Ops Director and its frequency, for example, hourly metering data. By default, Ops Director accepts the configuration as it is defined for the Managed Cluster.

Note: If necessary, you can configure Ops Director to override settings provided by the Managed Cluster, for example, to receive only daily metering data. You can do

this with the internal functions of the XQuery API. If you have an active maintenance contract, you can contact MarkLogic Technical Support for details.

1.8.3.2 Security of Collected Data

Ops Director is expected to operate within an enterprise rather than on the public Internet. Nevertheless, the service is designed to be secure and to protect customer confidential data.

- Local and transient storage of data is secured at the same level or better than the server protects the same data in other use cases.
- Data transmission uses only HTTP/SSL encrypted secure channels in a “point to point” architecture.
- Received data is stored in a MarkLogic database and is thus as secure as any other MarkLogic data.

1.8.3.3 Unattended Operation

The collection service runs unattended and requires no active management by the source administrator. Performance of the cluster is not significantly affected by this service.

1.8.3.4 “Best Attempt” SLA

Ops Director collects and transmits data in a “Best Attempt” SLA (Service Level Agreement). Transmissions are not “transactional” in the sense of database operations and have less priority than all critical and most other MarkLogic services. Since the service requires network connectivity, bandwidth, and local resources that compete with existing services, it is possible that even in periods of continuous functional operation the service may not be able to provide uninterrupted and complete streams. A “Best Attempt” approach is used to provide for periods of resource interruptions, heavy load, and improperly configured systems, while providing the configured data within the SLA.

If the connection between Ops Director and a Managed Cluster goes down, the Managed Cluster keeps a buffer of high-priority data that it will transmit to Ops Director when the connection is restored. However, some data might be lost, depending on the volume of data and the duration of the outage. Priority is given to retaining errors and warnings.

1.8.3.5 Data Type Separation

To decouple the collection, transmission, storage, and access requirements of the different kinds of data, Ops Director maintains a separate logical “stream” for each of the types of data. Each type of data (log, configuration, meter) is collected, configured, prioritized, identified, and delivered independently. This allows for simpler logic and specialization for send, receive, and access use cases.

1.8.3.6 Data Transmission

All data is delivered over HTTP(S) to Ops Director. Data is stored in MarkLogic on the Ops Director Application Cluster. Ops Director operates inside the enterprise; there is no expectation that Ops Director will refuse data from some clusters or be required to guard against denial-of-service or other malicious behavior.

2.0 Installing and Configuring Ops Director

Ops Director provides you with system monitoring, management, analysis, and support capabilities, as well as license auditing and role-based access control (RBAC) management. The Ops Director web-based interface offers graphical and tabular data representation intended to highlight irregular performance or alerts on a given cluster, host, database, or application server. This chapter contains instructions for installing and configuring Ops Director.

This chapter covers the following topics:

- [System Requirements](#)
- [Installing Ops Director](#)
- [Connecting a Managed Cluster to Ops Director](#)
- [Launching Ops Director and Logging In](#)
- [Switching from Digest to Application-Level Authentication](#)
- [Configuring the Ops Director Data Retention Policy](#)
- [Disconnecting a Managed Cluster from Ops Director](#)
- [Removing Ops Director](#)
- [Running Ops Director on Amazon Web Services \(AWS\)](#)
- [Securing Ops Director with Externally Signed Certificates](#)
- [Upgrading Ops Director](#)

2.1 System Requirements

The following table shows installation compatibility requirements for MarkLogic Server and Ops Director:

Installed Version of MarkLogic Server	Can Install Ops Director Version(s)
9.0-4	1.0-0
9.0-6 and earlier	1.0-0 and 1.1-1
9.0-7 or 9.0-8	2.0.0 and 2.0.1
9.0-9 and later	2.0.1 only.
10.0-2	2.0.1-1

Note: Ops Director 2.0 and greater will not work with versions of MarkLogic Server prior to 9.0-7, and Ops Director 2.0.1 will not work with MarkLogic Server version 10.0 and later.

Ops Director 1.1 requires that you have MarkLogic Server Version 9.0-3 or later running on the Application Cluster and MarkLogic Server Version 9.0-2 or later on the Managed Clusters. For definitions of the Application Cluster and Managed Cluster terms, see “Terms” on page 8.

Ops Director 2.0.1 has been tested with the following software versions:

- MarkLogic Server 9.0-7 or greater
- Internet Explorer 11 on Windows 10
- Firefox 62 on Windows 10 and Mac OS
- Chrome 63 on Windows 10 and Mac OS

When a cluster is configured to be managed by Ops Director, a SecureManage App Server is created with SSL enabled.

In general, the hardware requirements for Ops Director are the same as those described in [Memory, Disk Space, and Swap Space Requirements](#) in the *Installation Guide*. However, it is recommended that you have at least 32 GB of memory on a host where Ops Director application runs.

2.2 Installing Ops Director

The installation procedure has changed from Ops Director 1.x to Ops Director 2.x. Please follow the instructions for the version of Ops Director you are installing.

The following installations are available:

- [Installing Ops Director version 1.x on an Application Cluster](#)
- [Installing Ops Director 2.0 and Greater](#)

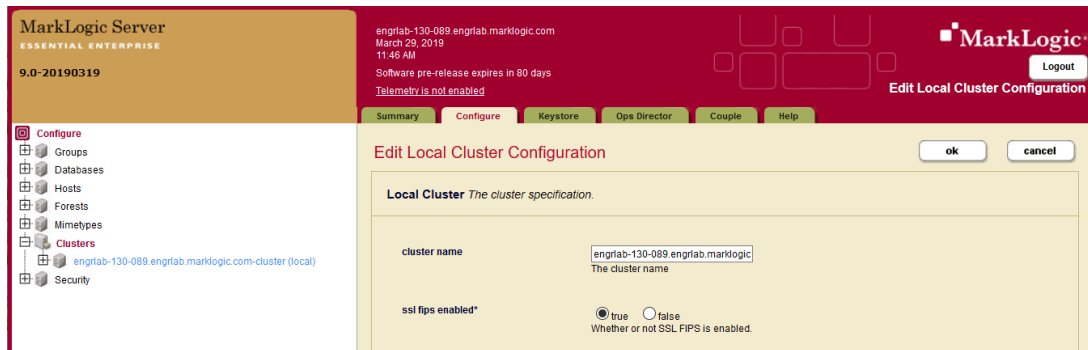
2.2.1 Installing Ops Director version 1.x on an Application Cluster

Ops Director is an application built on top of MarkLogic Server and runs on a designated MarkLogic cluster called the Ops Director Application Cluster. Managed Clusters are connected to the Ops Director Application Cluster. An Ops Director Application Cluster can serve as both an Application Cluster and a Managed Cluster.

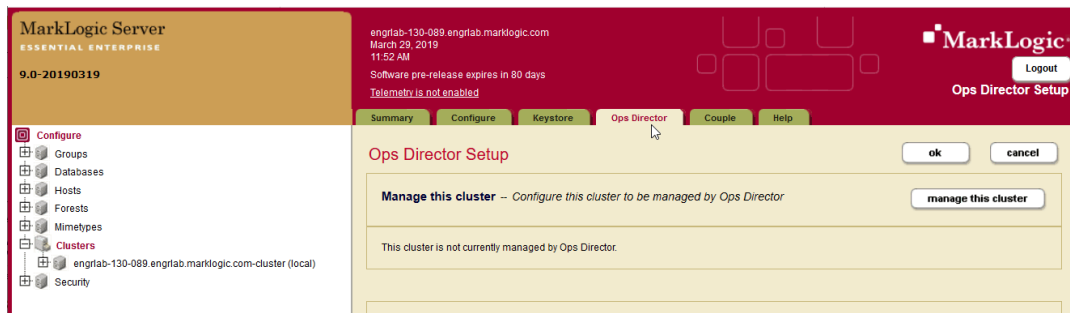
The following procedure describes how to install Ops Director 1.x on an application cluster.

1. Log into the Admin Interface.
2. Navigate to **Configure > Clusters** in the left tree menu.

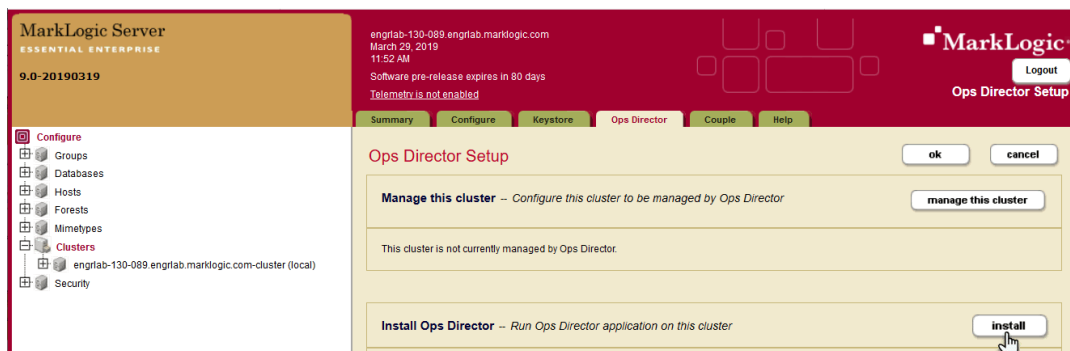
3. Select the local cluster. The Edit Local Cluster Configuration page displays.



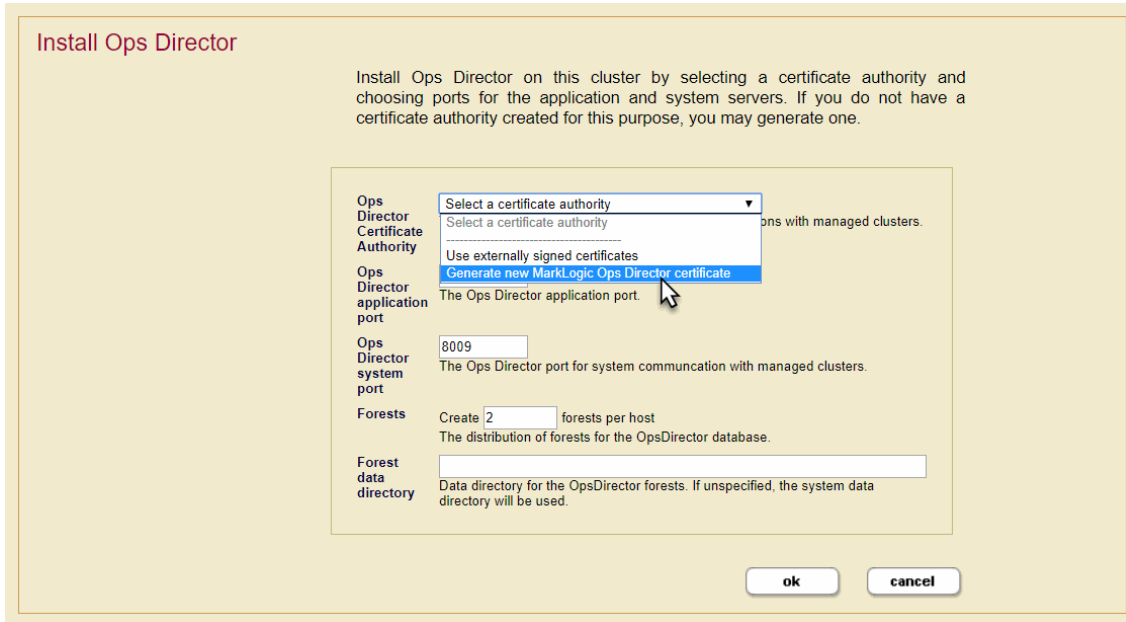
4. Select the **Ops Director** tab at the top of the page.



5. On the Ops Director Setup page, click **install**.



- On the Install Ops Director page, select **Generate new MarkLogic Ops Director certificate** from the Ops Director Certificate Authority menu. Click **ok**.



Note: This procedure secures inter-cluster communication by means of an internally-generated, self-signed certificate. If you plan to use externally signed certificates, see “Securing Ops Director with Externally Signed Certificates” on page 44.

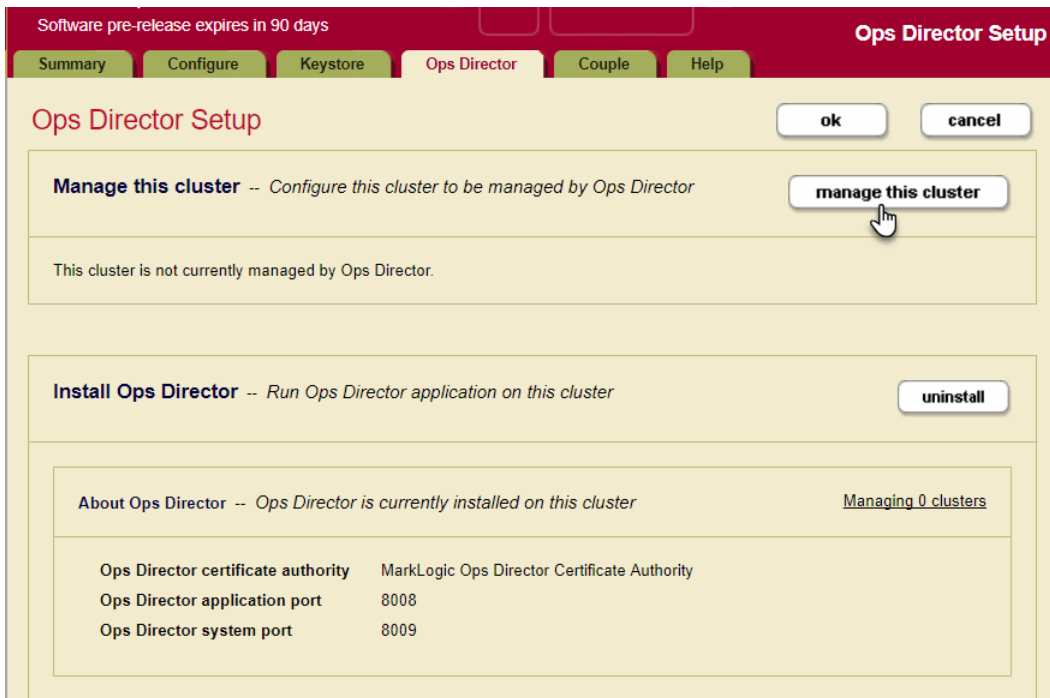
7. A Certificate Authority configuration page displays. Fill in the blank fields with all of your relevant information. Click **ok**.

The screenshot shows a configuration dialog box for the Ops Director Certificate Authority. At the top, there is a dropdown menu set to "Generate new MarkLogic Ops Director certificate" and a subtitle "The certificate authority to use for secure communications with managed clusters." Below this is a section titled "Complete the following fields for the new certificate authority." containing several input fields: "Country" (US), "State or province" (CA), "Locality" (San Carlos), "Organization" (MarkLogic), "Organization unit" (Engineering), and "Email address" (support@marklogic.com). Each field has a descriptive subtitle and a red error message for the Country and Organization fields. Below this section are four more fields: "Ops Director application port" (8008), "Ops Director system port" (8009), "Forests" (Create 2 forests per host), and "Forest data directory" (empty). At the bottom right, there are "ok" and "cancel" buttons, with a mouse cursor pointing at the "ok" button.

Ops Director Certificate Authority	Generate new MarkLogic Ops Director certificate ▼
The certificate authority to use for secure communications with managed clusters.	
Complete the following fields for the new certificate authority.	
Country	US A two character country code (e.g. "US"). Required. You must supply a value for country.
State or province	CA The state or province your server is in.
Locality	San Carlos The city your server is in.
Organization	MarkLogic The organization or company your server belongs to (e.g. MarkLogic). Required. You must supply a value for organization.
Organization unit	Engineering The organizational unit your server belongs to (e.g. Engineering).
Email address	support@marklogic.com The email address to contact regarding your server (e.g. webmaster@yourcompany.com).
Ops Director application port	8008 The Ops Director application port.
Ops Director system port	8009 The Ops Director port for system communication with managed clusters.
Forests	Create 2 forests per host The distribution of forests for the OpsDirector database.
Forest data directory	 Data directory for the OpsDirector forests. If unspecified, the system data directory will be used.

Note: In most cases, you may leave the default values for all other settings. However, in some cases you have to specify values different from the default ones. In particular, for installation on AWS, it is important to change the forest data directory.

8. The Ops Director Setup page displays. If you want this cluster to serve as both the Ops Director Application Cluster and as a Managed Cluster, select manage this cluster. Otherwise click **ok**.



9. If you opted to have the cluster serve as both the Ops Director Application Cluster and as a Managed Cluster, the Configure as a Managed Cluster page displays. Enter the name of the local host in this cluster that contains the OpsDirectorApplication and OpsDirectorSystem application servers. Set the port number for the OpsDirectorSystem server (8009, be default) and select MarkLogic Ops Director Certificate Authority from the Ops Director Certificate Authority pulldown menu.

Set the level for log messages sent to Ops Director, as well as the frequency at which the metering data is collected. For details on the log levels, see [Understanding the Log Levels](#) in the *Administrator's Guide*.

Note: To set the log interval filter in Ops Director to `raw`, as described in “Date and Time Filters” on page 86, you must also set the Ops Director Metering to `raw` here in the Admin Interface.

When finished, Click **ok**.

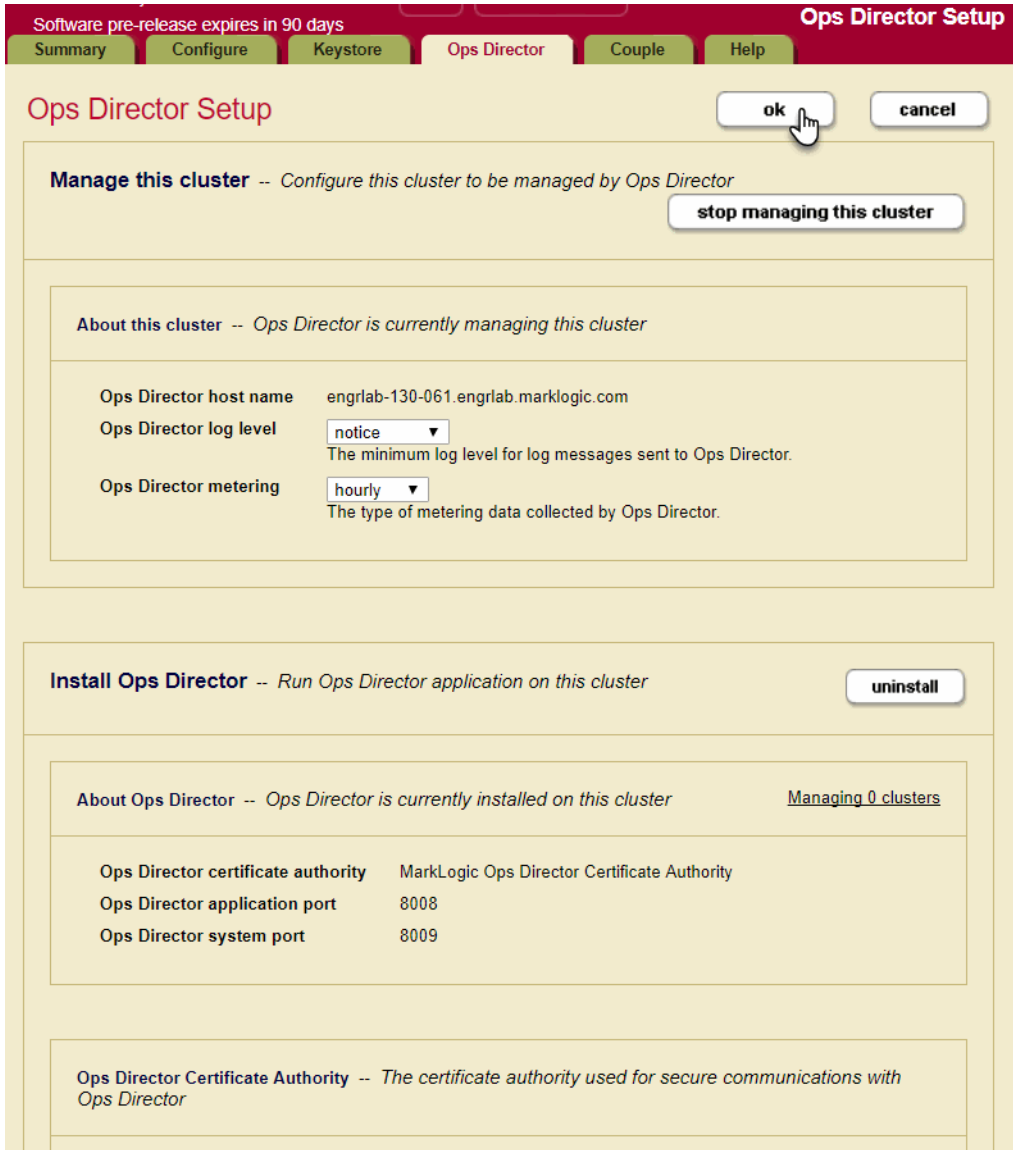
Configure as a Managed Cluster

Use the fields below to identify the Ops Director server already installed within your enterprise. This will allow Ops Director to manage this cluster.

Ops Director hostname	<input type="text" value="engr1ab-130-061.engr1ab.marklogic.com"/> The (resolvable from this host) Ops Director hostname.
Ops Director system port	<input type="text" value="8009"/> The Ops Director port for system communication with managed clusters.
Ops Director log level	<input type="text" value="notice"/> The minimum log level for log messages sent to Ops Director.
Ops Director metering	<input type="text" value="hourly"/> The type of metering data collected by Ops Director.
Ops Director Certificate Authority	<input type="text" value="MarkLogic Ops Director Certificate Authority"/> The certificate authority shared with the Ops Director system.
This hostname	<input type="text" value="engr1ab-130-061.engr1ab.marklogic.com"/> The (resolvable from the Ops Director host) name of this host.

- 10. The Ops Director Setup page displays, with the settings configured during the installation of Ops Director.

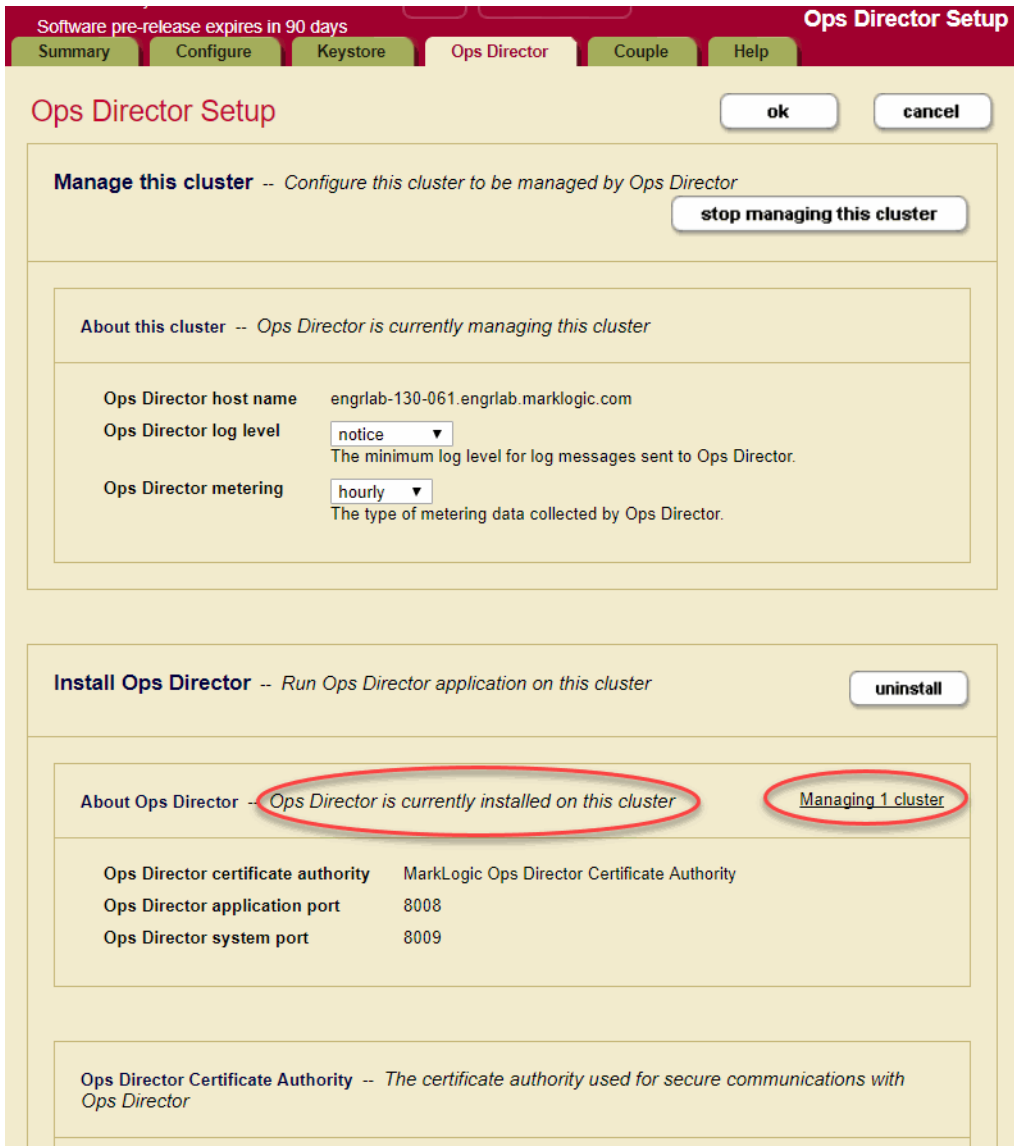
Click **ok**.



11. The Summary page for the local cluster displays.

Select the Ops Director tab again. Notice that the About Ops Director description specifies “*Ops Director is currently installed on this cluster*”.

Also, if you opted to have the cluster serve as both the Ops Director Application Cluster and as a Managed Cluster, the description specifies “*Managing 1 cluster*”.



12. Click Configure on the left tree menu. In the right page, select the Summary tab. The Configuration Summary page displays all of the resources created by installing Ops Director.

The screenshot shows the MarkLogic Server System Summary page. The top navigation bar includes 'Summary', 'Status', 'Support', 'Logs', 'Usage', and 'Help'. The left sidebar shows a tree view with 'Configure' selected. The main content area is divided into several sections:

- Databases (11)** -- Index, query, and content processing configuration
 - App-Services
 - Documents
 - Extensions
 - Fab
 - Last-Login
 - Meters
 - Modules
 - OpsDirector
 - Schemas
 - Security
 - Triggers
- App Servers (7)** -- Enable connections from client software
 - Default :: Admin : 8001 [HTTP]
 - Default :: App-Services : 8000 [HTTP]
 - Default :: HealthCheck : 7997 [HTTP]
 - Default :: Manage : 8002 [HTTP]
 - Default :: OpsDirectorApplication : 8008 [HTTP]
 - Default :: OpsDirectorSystem : 8009 [HTTP]
 - Default :: SecureManage : 8003 [HTTP]
- Groups (1)** -- Allow hosts to share a common configuration
 - Default
- Forests (12)** -- Manage physical content storage for databases
 - App-Services
 - Documents
 - Extensions
 - Fab
 - Last-Login
 - Meters
 - Modules
 - OpsDirector-1
 - OpsDirector-2
 - Schemas
 - Security
 - Triggers
- Security** -- Resources describing the role-based security model
 - Users (7)
 - Roles (93)
 - Execute Privileges (566)
 - URI Privileges (16)
 - Amps (848)
 - Collections (8)
 - Protected Paths (0)
 - Query Rolesets (0)
 - Certificate Authorities (64)
 - Certificate Templates (2)
 - External Security (2)
 - Credentials
 - Secure Credentials (3)
- Hosts (1)** -- Computers belonging to this cluster
 - Default :: engrlab-130-061.engrلاب.marklogic.com
- Clusters (1)** -- Cluster configuration
 - engrلاب-130-061.engrلاب.marklogic.com-cluster (Local Cluster)

2.2.2 Installing Ops Director 2.0 and Greater

Follow these instructions to install Ops Director and greater.

2.2.2.1 Before You Install Ops Director

Make sure the following conditions are in place before you begin the installation process:

- Internet access is required to install Ops Director 2.0. Make sure you can connect to the Internet.
- Make sure you are running MarkLogic Server 9.0-7 or greater.
- Make sure you have access to ports 8000, 8001, 8002, 8008, and 8009 on the Ops Director cluster.
- Start MarkLogic Server and log into the Admin UI (port 8001).

Note: The Ops Director installer is written in ML Gradle. A version of Gradle is included in the Ops Director 2.0 installation directory. For more information about Gradle, visit <https://docs.gradle.org/current/userguide/userguide.html>.

2.2.2.2 Ops Director 2.0 Installation

To install or upgrade to Ops Director 2.0 or greater, follow these instructions:

1. Download the Ops Director zip file from [marklogic.com](https://developer.marklogic.com/products/opsdirector) at:
<https://developer.marklogic.com/products/opsdirector>
2. Unzip the file, and `cd` to the directory containing the files.
3. Either edit the Ops Director settings in the `gradle.properties` file or enter them as options on the command line. The following table contains the host settings.

Setting	Description
<code>mlHost</code>	Host where Ops Director will be installed. Default value: <code>localhost</code> .
<code>mlUsername</code>	Username for admin access to that host. This must be the actual value on the cluster.
<code>mlPassword</code>	Password for admin access to that host. This must be the actual value on the cluster.
<code>opsdirDataExpires</code>	An <code>xs:dayTimeDuration</code> specifying how long to keep historical data in the Ops Director database. Default value: <code>P60D</code>

The following settings determine how the Ops Director cluster communicates with the managed clusters. Choose one of the following Certificate Authority options:

- `opsdirCa=generate`

If you set this value, you also need to add the following values:

Setting	Definition
<code>opsdirCaCountry</code>	The two-letter code for the country in which the server is located (for example, <code>US</code>)
<code>opsdirCaProvince</code>	The two-letter code for the state or province in which the server is located (for example, <code>CA</code>)

Setting	Definition
<code>opsdirCaLocality</code>	The city in which your server is located (for example, San Carlos)
<code>opsdirCaOrgName</code>	The organization or company to which the server belongs (for example, MarkLogic)
<code>opsdirCaOrgUnit</code>	The organizational unit to which the server belongs (for example, Engineering)
<code>opsdirCaEmail</code>	The email address to contact regarding this server (for example, <code>webmaster@yourcompany.com</code>)

- `opsdirCa=external`

Use this setting to use an external certificate for Ops Director.

- `opsdirCa=ML_ID`

Use this setting to use a certificate that is already installed where *ML_ID* is the MarkLogic ID number for the certificate.

4. From the same Ops Director 2.0 installation directory, run the installer:

```
./gradlew mlDeploy
```

Note: If this is the first time you are installing Ops Director 2.0, a lot of periods will display, followed by files being downloaded.

2.2.2.3 Ops Director Resources

The following resources are created on the host where your Ops Director application runs.

Resource Type	Resource Name and Function
Certificate Authority	The Certificate Authority (MarkLogic Ops Director in this example) is used to generate the secure credentials required for the Ops Director application and Managed Clusters to securely communicate.
Certificate Templates	<code>OpsDirector-SSL-Template</code> for SSL on the Application Cluster. <code>OpsDirector Template</code> for SSL, if also configured as a Managed Cluster.

Resource Type	Resource Name and Function
Roles	<p><code>opsdir-guest</code> role is used for a default user of the Ops Director application.</p> <p><code>opsdir-user</code> role is required for any user to be able to access the Ops Director application.</p> <p><code>opsdir-license-admin</code> role grants a user rights to manage licenses for the Ops Director application.</p> <p><code>opsdir-admin</code> role grants a user administrative rights to the Ops Director application.</p>
Execute Privileges	<p><code>opsdir-admin</code> protects administrative functions.</p> <p><code>opsdir-license-admin</code> is required to access license information.</p> <p><code>opsdir-user</code> is required to access the browser application.</p>
Database and Forests	The <code>OpsDirector</code> database and forests hold the Ops Director configuration data.
App Servers	<p><code>OpsDirectorApplication</code> (port 8008) provides the Ops Director browser application and consumes the data stored in the <code>OpsDirector</code> database.</p> <p><code>OpsDirectorSystem</code> (port 8009) receives data transmitted from Managed Clusters and stores it in the <code>OpsDirector</code> database.</p>
External Security Configurations	<p><code>OpsDirector-External-Security</code> for authentication of clients.</p> <p><code>OpsDirectorSystem</code> for secure communication with <code>OpsDirectorSystem</code> server.</p>
Secure Credentials	<p><code>OpsDirector-Credential to sign OpsDirector-SSL-Template</code>.</p> <p><code>MarkLogic-OpsDirector</code> for accessing <code>OpsDirectorSystem</code> server.</p>

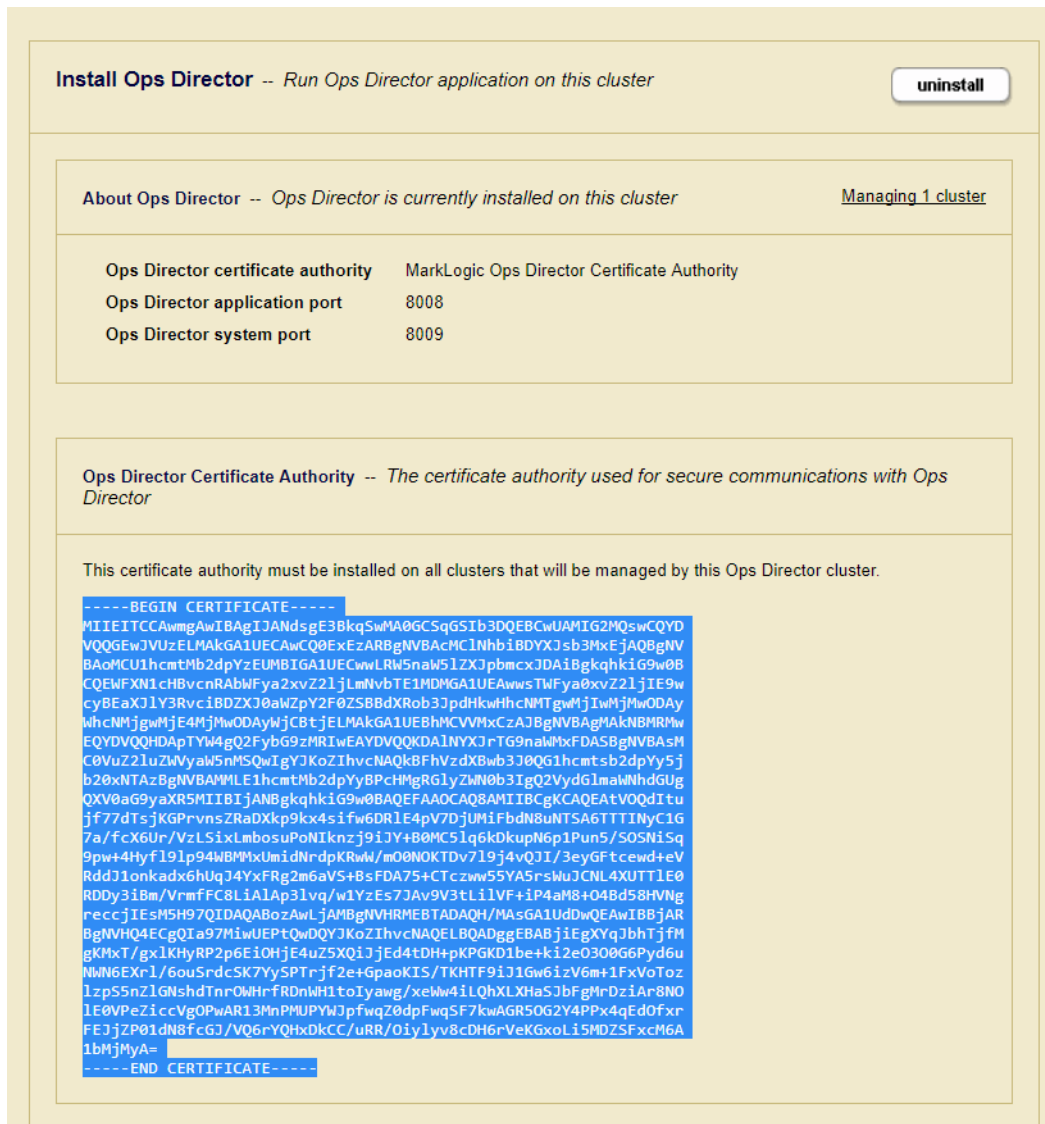
2.3 Connecting a Managed Cluster to Ops Director

The following procedure describes how to configure a cluster to be managed by Ops Director.

Note: Do not connect new managed clusters to Ops Director if you started the process of Ops Director upgrade and have not completed it. Adding managed clusters in the middle of Ops Director upgrade can lead to error conditions. For more details, see “Upgrade Error Prevention and Troubleshooting” on page 79.

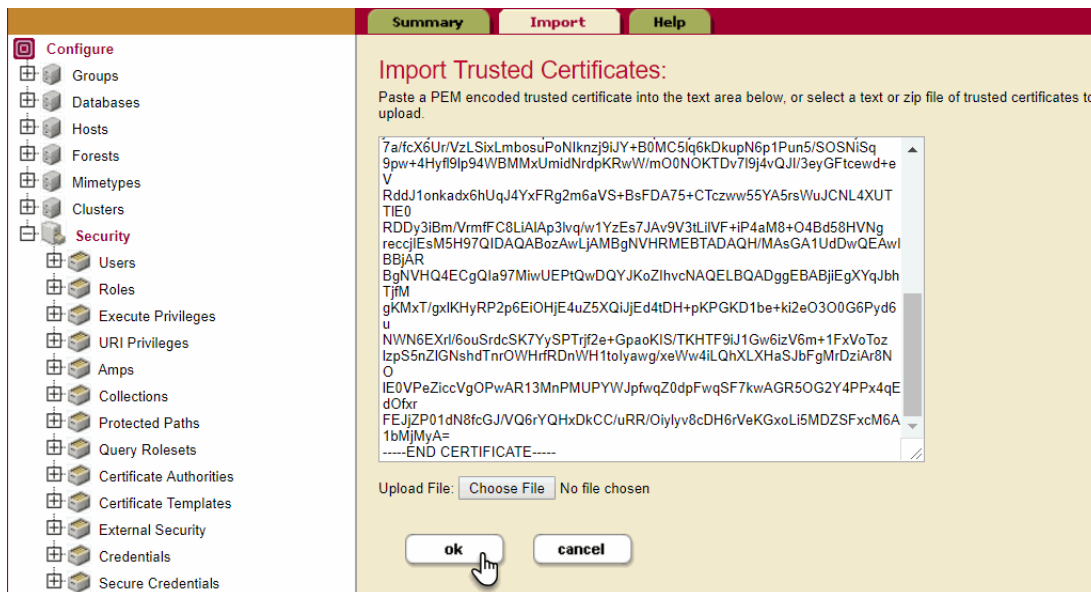
If you are using an internally signed certificate, before you log into the cluster to be managed, you must copy the value of the Ops Director Certificate Authority field in the Admin Interface from the host where Ops Director is installed to a host of your Managed Cluster. You can either use the REST Management API to script the operation, or you can follow these steps:

1. On the host where Ops Director is installed, make sure that MarkLogic Server is running, open the Admin Interface, and click Configure > Clusters in the left tree menu.
2. Select the local cluster. The Edit Local Cluster Configuration page displays. Select the Ops Director tab to display the Ops Director Setup page.
3. Scroll down to the Ops Director Certificate Authority section and copy the certificate by selecting everything between and including -----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----.



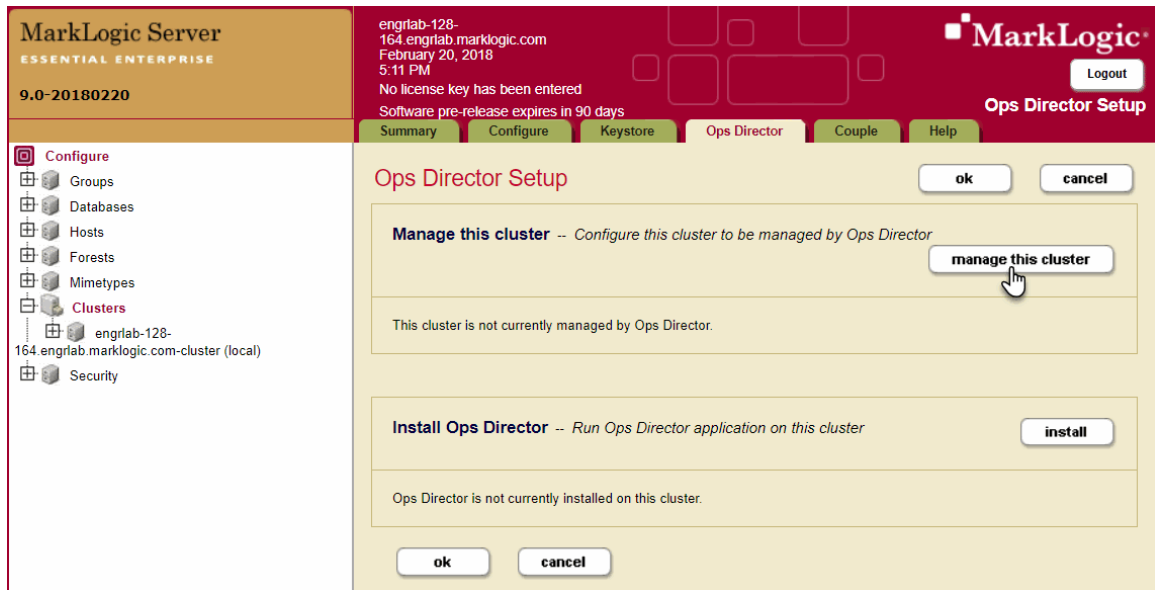
Log on to the Admin Interface of a host on the cluster to be managed by Ops Director and do the following:

1. Select **Configure > Security > Certificate Authorities** on the left menu.
2. On the right page, select the **Import** tab at the top of the Summary page.
3. Paste the certificate copied from the host where Ops Director is installed.
4. Click **ok**.



5. Select **Configure > Clusters** on the left tree menu.
6. Select the local cluster. The Edit Local Cluster Configuration page displays.
7. Select the **Ops Director** tab at the top of the page.

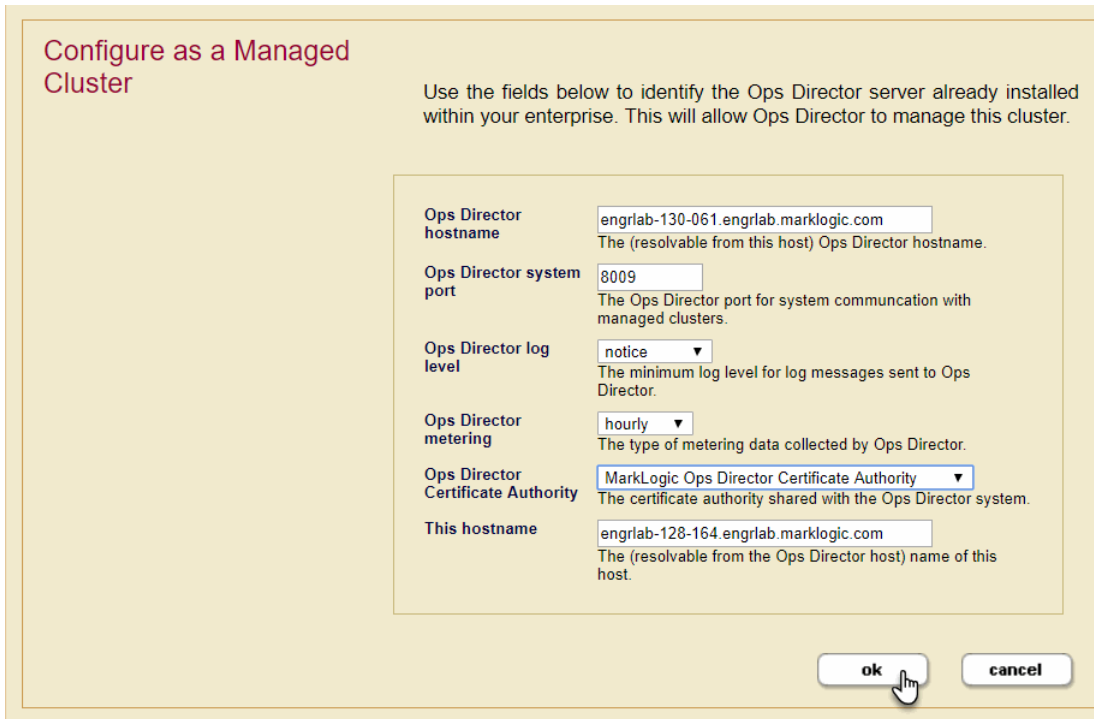
- The Ops Director Setup page displays. Select manage this cluster.



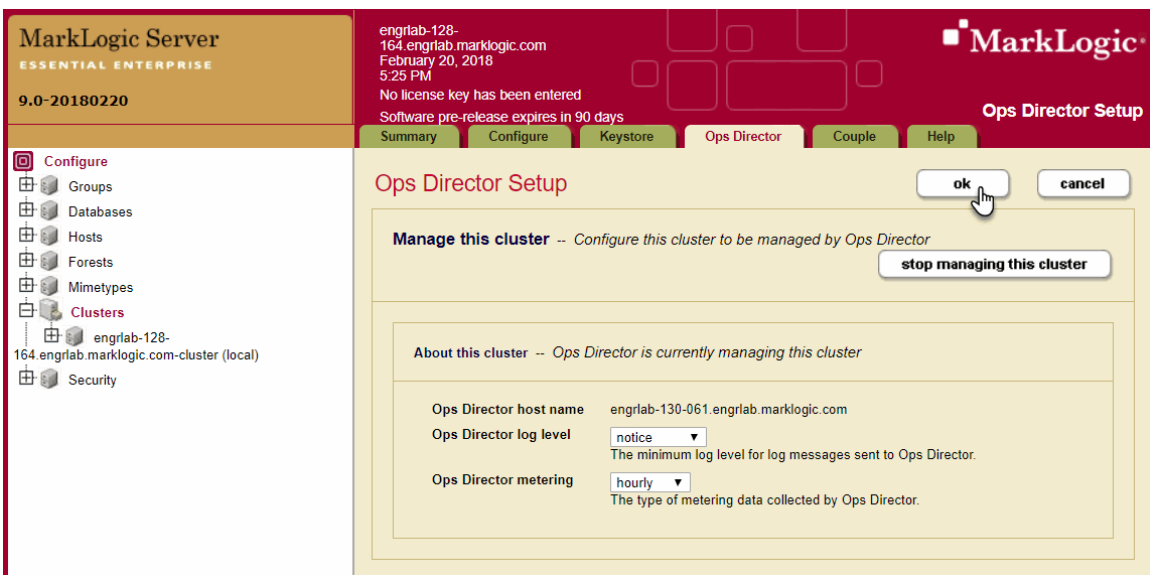
- The Configure as a Managed Cluster page displays. Enter the name of the host where your Ops Director application runs. This is the host you set up in "Installing Ops Director" on page 20.
- Select **MarkLogic Ops Director Certificate Authority** from the Ops Director Certificate Authority menu.

- Set the level for log messages sent to Ops Director, as well as the frequency at which the metering data is collected. For details on the log levels, see [Understanding the Log Levels](#) in the *Administrator's Guide*.

When finished, click **ok**.



- The Ops Director Setup page displays. Click **ok**.



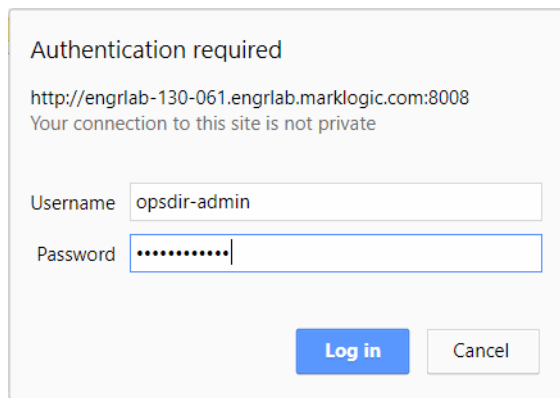
Note: When you designate a Managed Cluster, a SecureManage App Server is created in each group on the managed cluster to provide secure connections with the Ops Director Application Cluster. For more detail, see “Security and Database Dependencies of Managed Clusters” on page 13.

2.4 Launching Ops Director and Logging In

In a browser window, enter the hostname of the host where Ops Director runs and the port number of the OpsDirectorApplication server (by default, 8008). For example:

```
https://englab.marklogic.com:8008
```

You are asked for credentials with a standard authentication dialog box. To log into Ops Director, enter a valid user name and password an existing MarkLogic user with the `opsdir-admin` role and click Log in.



User credentials can come from LDAP or internal security. If internal security is used, the user must exist on the Ops Director Application Cluster.

Note: By default, Ops Director is configured with digest-based authentication. To enable application-level authentication, perform the steps described in “Switching from Digest to Application-Level Authentication” on page 36.


When you use digest-based authentication, make sure to refresh the Ops Director application in the browser after you log into the application with different credentials, which may have different access to resources than the previous user that was logged-in. If you do not refresh the Ops Director application in the browser, the application may show stale information.

2.5 Switching from Digest to Application-Level Authentication

To enable application-level authentication, perform the following steps:

1. Install SSL certificate on the Ops Director application server, as described in [General Procedure for Setting up SSL for an App Server](#) in the *Security Guide*.

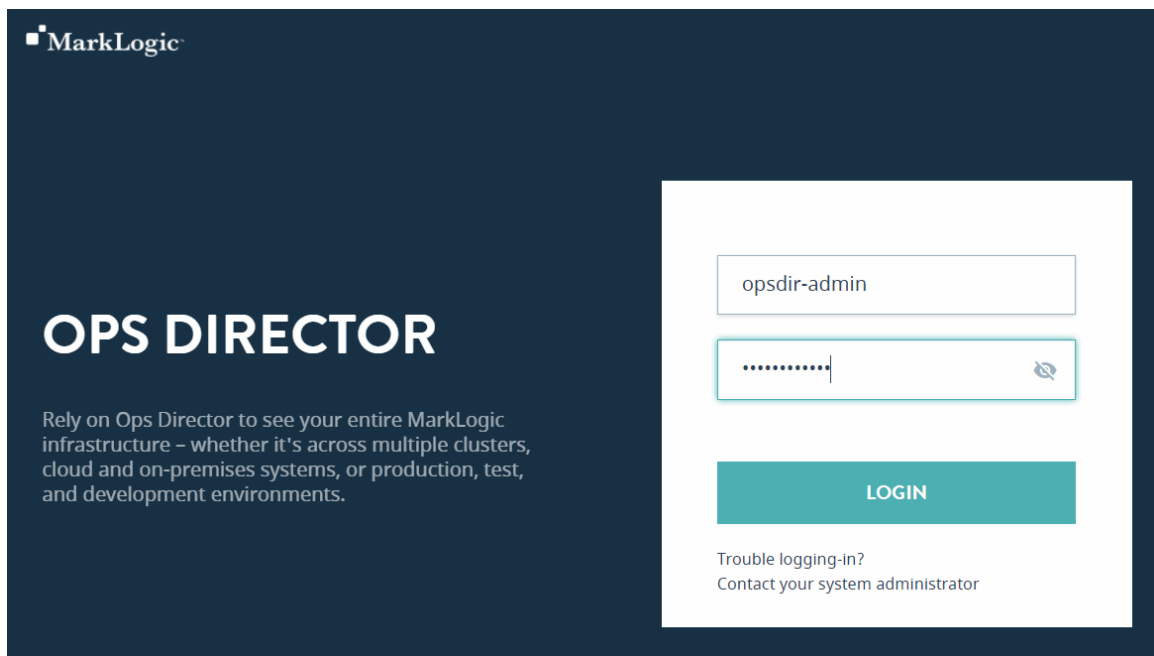
2. Log into the Admin Interface on the host where the Ops Director application is running.
3. In the left menu, select Configure > Groups > Default > App Servers > OpsDirectorApplication. On the right page, select the Configure tab.
4. Scroll down to the authentication menu and change the selection from digest (default) to application-level.



5. Click **ok** and wait for the operation to complete.
6. In a browser window, enter the hostname of the host where Ops Director runs and the port number of the OpsDirectorApplication server (by default, 8008). For example:

```
https://englab.marklogic.com:8008
```

7. The Ops Director login screen displays. Enter a valid user name and password and click LOGIN.



Note: Switching from `digest` authentication to `application-level` authentication without first installing an SSL certificate to enable HTTPS on port 8008 presents a security risk because credentials are sent over HTTP as clear text.

2.6 Configuring the Ops Director Data Retention Policy

The following data accumulates in the `OpsDirector` database:

- Meters data
- Log files
- Configuration data of managed clusters
- System alerts

Ops Director has a built-in auto-cleanup mechanism to delete accumulated data periodically.

Once a day, at 11:00 PM local time, a task runs to delete accumulated Ops Director data based on the defined data retention policy. By default, all types of data—meters, logs, cluster configuration data, and system alerts—is retained for 60 days.

The default data retention period for each type of data is defined in the configuration file `/config/opsdirector.xml` stored in the `OpsDirector` database. Follow these instructions to retrieve the contents of this file.

1. Open Query Console on the host where Ops Director runs:

```
http://<hostname>:8000/qconsole/
```

2. On the Query Console, set Query Type to XQuery and Database to OpsDirector.
3. Execute the following code:

```
xquery version "1.0-ml";
doc("/config/opsdirector.xml")
```

In the response, look for the following elements:

- `max-meters-age` — specifies the maximum age for meters data;
- `max-logs-age` — specifies the maximum age for log files;
- `max-resource-age` — specifies the maximum age for cluster configuration data and system alerts.

By default, the maximum age for each class of data is set to 60 days, as in the following example:

```
<config xmlns="http://marklogic.com/v1/opsdirector/config">
  <max-meters-age>P60D</max-meters-age>
  <max-logs-age>P60D</max-logs-age>
  <max-resource-age>P60D</max-resource-age>
</config>
```

You can modify the default data retention period as follows: edit the configuration file `/config/opsdirector.xml` and replace default `P60D` values with other values of type `xs:dayTimeDuration` to define your preferred data retention periods for corresponding data classes.

Note: For details on `xs:dayTimeDuration` data type, see <https://www.w3.org/TR/xpath-functions/#duration-subtypes>.

Ops Director includes a utility library to safely change values in the configuration file. For example, to retain meters data and log files for 30 days only, run the following query on the OpsDirector database:

```
xquery version "1.0-ml";

declare namespace
  cfg="http://marklogic.com/v1/opsdirector/config";
declare default function namespace
  "http://www.w3.org/2005/xpath-functions";
declare option xdmp:mapping "false";

let $config := doc("/config/opsdirector.xml")/cfg:config
let $_ := if (empty($config))
  then error((), "No config document?")
```

```
else ()

let $expire := xs:dayTimeDuration("P30D")

for $property in ("max-meters-age", "max-logs-age")
let $qname :=
  fn:QName("http://marklogic.com/v1/opsdirector/config", $property)
let $node := $config/*[node-name(.) = $qname]

return
  if (empty($node))
  then xdm:node-insert-child($config, element { $qname } { $expire })
  else xdm:node-replace($node/text(), text { $expire })
```

If you want a specific class of data to be stored forever, set the data retention period for that class to *infinity* in the configuration file `/config/opsdirector.xml`. Note that this setting will cause the OpsDirector database to grow as data accumulates. Make sure that you have sufficient capacity to handle this growth.

2.7 Disconnecting a Managed Cluster from Ops Director

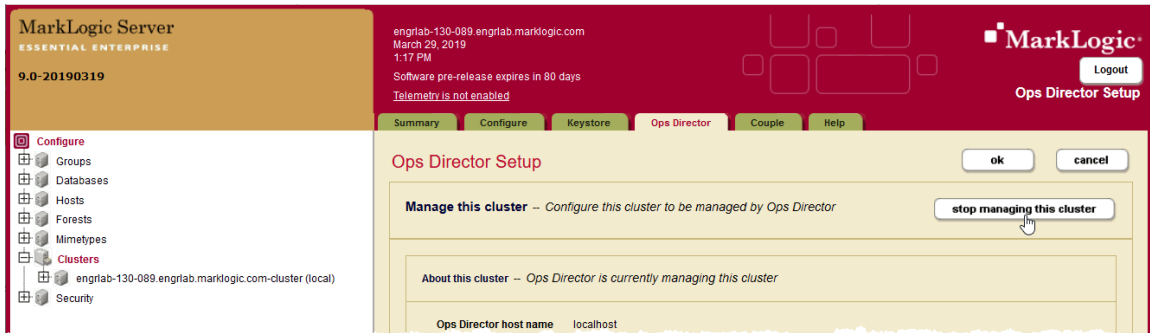
A Managed Cluster can be disconnected from Ops Director. Once disconnected, the Managed Cluster stops sending configuration, metrics, and log data to the Ops Director cluster.

Ops Director will continue to make Management API calls to assess the cluster health as long as there is a valid session for that cluster. After the session expires, Ops Director will stop making calls and the disconnect will be complete.

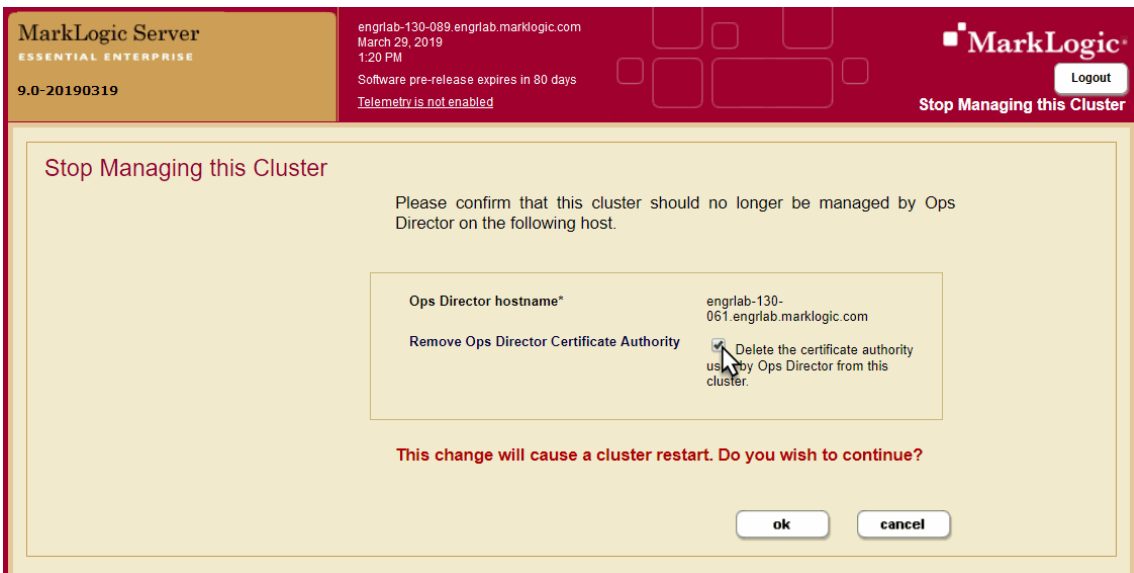
The procedure for disconnecting a Managed Cluster from Ops Director is as follows:

1. Log into the Admin Interface on the Managed Cluster.
2. Select Configure > Clusters on the left tree menu.
3. Select the local cluster. The Edit Local Cluster Configuration page displays.
4. Select the Ops Director tab at the top of the page.

- In the Ops Director Setup page, click stop managing this cluster.



- On the Stop Managing this Cluster page, select the Remove Ops Director Certificate Authority box and click **ok**.



Note: If you plan to have this cluster managed again by the same instance of Ops Director at some point in the future, you may leave the Remove Ops Director Certificate Authority box unchecked.

2.8 Removing Ops Director

This section describes how to remove the Ops Director application.

Note: Before you remove Ops Director, disconnect all Managed Clusters, as described in “Disconnecting a Managed Cluster from Ops Director” on page 40.

2.8.1 Removing Ops Director 1.1 and Earlier

The following procedure describes how to remove Ops Director 1.1 and earlier from the MarkLogic host.

1. Log into the Admin Interface.
2. Select **Configure > Clusters** on the left tree menu.
3. Select the local cluster. The Edit Local Cluster Configuration page displays.
4. Select the **Ops Director** tab at the top of the page.
5. In the Install Ops Director section, click **uninstall**.

The screenshot displays the MarkLogic Admin Interface. The top header shows the MarkLogic logo and system information: "engrlab-130-061.engrlab.marklogic.com", "February 20, 2018 4:58 PM", and "No license key has been entered. Software pre-release expires in 90 days". The left navigation menu is expanded to "Clusters", showing the local cluster "engrlab-130-061.engrlab.marklogic.com-cluster (local)". The main content area is titled "Ops Director Setup" and has tabs for "Summary", "Configure", "Keystore", "Ops Director", "Couple", and "Help". The "Ops Director" tab is selected. The page contains several sections: "Manage this cluster" with a "stop managing this cluster" button; "About this cluster" with a note that Ops Director is currently managing the cluster; "Install Ops Director" with an "uninstall" button (highlighted by a mouse cursor) and a note that Ops Director is currently installed; and "About Ops Director" with a link to "Managing 1 cluster" and a table of configuration details.

Ops Director certificate authority	MarkLogic Ops Director Certificate Authority
Ops Director application port	8008
Ops Director system port	8009

6. The Uninstall Ops Director page displays and asks for confirmation: “This change will cause a cluster restart. Do you wish to continue?” Click **ok**.

Note: When you use this method of uninstalling Ops Director, the Ops Director database is not uninstalled.

2.8.2 Removing Ops Director 2.0 and Later

The following procedure describes how to remove Ops Director 2.0 and later from the MarkLogic host.

1. Change to the Ops Director installation directory (for example, `opsdirector-2-0-0`).
2. If you used the default value for `mlGroupName` in the `gradle.properties` file, enter the command:

```
gradlew mlUndeploy -Pconfirm=true -PconfirmManaged=true
```

3. If you used a different value for `mlGroupName` in the `gradle.properties` file, enter the command (on one line):

```
gradlew mlUndeploy -PmlGroupName=mlGroupName -Pconfirm=true  
-PconfirmManaged=true
```

where *mlGroupName* is the value from the `gradle.properties` file.

2.9 Running Ops Director on Amazon Web Services (AWS)

When AWS hosts are stopped and restarted, they are assigned new host names and new IP addresses. They do not get new MarkLogic host IDs, as those are persistent. Ops Director will recover from AWS host renaming under the following circumstances:

- If a Managed Cluster host is renamed, when Ops Director receives information from it, the fact that an existing host id has appeared with a new host name will be used to update the Ops Director configuration. At that point, things return to normal.
- If an Ops Director host is renamed, the Managed Cluster will be unable to communicate with it. But the next time Ops Director polls that Managed Cluster, it will detect that out-of-date session endpoint and update it. At that point, things return to normal.

However, if you stop both the Ops Director host and the Managed Cluster hosts at the same time, then all hosts will come back with new host names. There is no way to fix the problem automatically, but you can fix the problem with a script. In this example, Ops Director used to be AMAZONOLD and now is AMAZONNEW. On each Managed Cluster that is still trying to talk to AMAZONOLD, use the following script on the Query Console:

```
xquery version "1.0-ml";

import module namespace admin = "http://marklogic.com/xdmp/admin"
  at "/MarkLogic/admin.xqy";

let $config := admin:get-configuration()
let $config := admin:cluster-set-opsdirector-session-endpoint($config,
  "https://AMAZONNEW:8009/v1/opsdirector/session")

return
  admin:save-configuration($config)
```

2.10 Securing Ops Director with Externally Signed Certificates

The procedures described in “Installing Ops Director” on page 20 make use of an internally-generated, self-signed certificate.

To use externally signed certificates with any version of Ops Director, you must get the Certificate Signing Requests (CSRs) from MarkLogic, pass the CSRs to the Certificate Authority and receive signed certificates, and then install the signed certificates on the cluster. You must ensure that you have secured a certificate authority (CA) and can get signed certificates issued by that CA for each of the nodes involved.

In this example, we will be working with a two-node Ops Director cluster with hosts named “node1” and “node2,” and a two-node managed cluster with hosts named “node3” and “node4.” In each step, a code example is presented showing how the REST Management API (RMA) can be used to accomplish the task. Where the RMA cannot be used, XQuery scripts are shown.

Note: Your MarkLogic server hosts must have fully-qualified names that can be resolved by DNS.

If the signing authority is Verisign, Thawte, or another well established authority, MarkLogic will already have the appropriate Certificate Authority installed. If you are using a less well-known authority, you must install the appropriate Certificate Authority on each MarkLogic cluster.

Note: Before you begin this procedure, make sure you have a way of obtaining certificates.

The following sections contain instructions for configuring Ops Director to use externally signed certificates.

- [Using Externally Signed Certificates with MarkLogic Server 10.0-2 and Later and Ops Director 2.0.1-1](#)
- [Using Externally Signed Certificates with MarkLogic Server 9.0-9 and Later and Ops Director 2.0.1 and Later](#)
- [Using Externally Signed Certificates with MarkLogic 9.0-8 and earlier or Ops Director 2.0 and Earlier](#)

2.10.1 Using Externally Signed Certificates with MarkLogic Server 10.0-2 and Later and Ops Director 2.0.1-1

Note: Ops Director 2.0.1-1 is not compatible with MarkLogic Server 10.0-1.

The steps involved for configuring Ops Director to use externally signed certificates are:

- [Configure the Ops Director Cluster](#)
- [Configure the Managed Cluster](#)

In the following examples, we use different CAs for client and server.

2.10.1.1 Configure the Ops Director Cluster

Once you have obtained the CA file, follow these steps to configure the Ops Director cluster:

1. If necessary, use the following `curl` commands to install the Certificate Authority on the cluster, replacing `admin:admin` with the credentials for your Ops Director installation, `node1` with the name of your bootstrap host, `/tmp/server.pem` with the location and name of your server CA file, and `/tmp/client.pem` with the name of your client CA file:

```
curl --digest -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/server.pem \
  "http://node1:8002/manage/v2/certificate-authorities?tag=opsdir:ssl-ca"
```

```
curl --digest -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/client.pem \
  "http://node1:8002/manage/v2/certificate-authorities?tag=opsdir:client-ca"
```

Note: You are not required to use different CAs for the client and server.

2. Install Ops Director:

```
./gradlew -PopsdirCa=external mlDeploy
```

You may need to use additional options, such as `mlHost`, depending on your configuration.

3. Create a certificate template for the hosts on the cluster:

```
{
  "template-name": "OpsDirector-SSL-Template",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",
  "key-options": {
    "key-length": "2048"
  },
  "req": {
    "version": "0",
    "subject": {
      "countryName": "US",
```

```

    "stateOrProvinceName": "California",
    "localityName": "San Carlos",
    "organizationName": "MarkLogic Corporation"
  }
}
}

```

Note: Our examples use MarkLogic-specific values. You may use other values, but you must use them *consistently*.

4. Use `curl` to post this certificate template to the REST Management API, replacing `/tmp/od-template.json` with the location and name of your template filename:

```

curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-template.json \
  http://node1:8002/manage/v2/certificate-templates

```

5. Create a JSON file (in this example, we call it `/tmp/od-extsec.json`) that creates an external authentication configuration object:

```

{
  "external-security-name": "OpsDirectorSystem",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
}

```

6. Use `curl` to POST this file on the REST Management API (port 8002), replacing `node1` with the name of your Ops Director bootstrap host:

```

curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-extsec.json \
  http://node1:8002/manage/v2/external-security

```

7. Create an XML file (in this example, we call it `/tmp/od-server.xml`) to update the server configuration to modify some of the properties of the `OpsDirectorSystem` application server. In this example, replace *Insert_the_certificate_here* with your certificate:

```

<http-server-properties xmlns="http://marklogic.com/manage">
  <enabled>true</enabled>
  <authentication>application-level</authentication>
  <external-securities>
    <external-security>OpsDirectorSystem</external-security>
  </external-securities>
  <ssl-certificate-template>OpsDirector-SSL-Template</ssl-certificate-
template>
  <ssl-client-certificate-pems>
    <ssl-client-certificate-pem>Insert_the_certificate_here</ssl-client-

```

```
certificate-pem>
  </ssl-client-certificate-pems>
</http-server-properties>
```

8. Use curl to PUT this file on the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X PUT -H \
  "content-type:application/xml" \
  --data-binary @/tmp/od-server.xml \
  http://node1:8002/manage/v2/servers/OpsDirectorSystem/properties?gro
  up-id=Default
```

Note: If you are not using the Default group for the Ops Director servers, update `group-id` accordingly.

9. Use the following shell script to generate certificate requests for each of the hosts, replacing `node1` with the name of a host on your cluster and `host_list` with the names of the hosts in your cluster separated by spaces:

```
for host in host_list; do
  echo "Generating CSR for $host to /tmp/$host.csr ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:application/json" -d "{\"operation\": \
    \"generate-certificate-request\", \"common-name\": \
    \"$host\", \"dns-name\": \"$host\" }" \
    //node1:8002/manage/v2/certificate-templates/OpsDirector-SSL-
  Template \
    > /tmp/$host.csr
done
```

10. Get a signed server certificate for each CSR from the server CA in step [1](#).
11. To install the certificates, run the following code, replacing `node1` with the name of a host on your cluster and `host_list` with the names of the hosts in your cluster separated by spaces, and store the certificate in `/tmp/hostname.crt` for each host:

```
for host in host_list; do
  "Installing server certificate for $host from /tmp/$host.crt ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:text/plain" --data-binary @/tmp/$host.crt \
    "http://node1:8002/manage/v2/certificates?tag=opsdir:host-
  certificate"
done
```

Note: If you prefer, you can put all of the certificates in a single ZIP file and post that to the `/manage/v2/certificates` endpoint.

12. Verify that the certificates have been installed. In the Admin UI, navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.

13. Create a JSON file (in this example, we call it `/tmp/od-user.json`) that creates a user called `opsdirector-system-user`:

```
{
  "user-name": "opsdirector-system-user",
  "description": "Ops Director system user",
  "role": [ "opsdir-admin-internal" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}
```

Note: Make sure the external name *exactly* matches the client certificates generated in later steps.

14. Use `curl` to POST this to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-user.json \
  http://node1:8002/manage/v2/users
```

2.10.1.2 Configure the Managed Cluster

When using external certificates, configuring a new managed cluster requires changes on both the managed cluster *and* the Ops Director Cluster.

- [Setting up the Managed Cluster](#)
- [Setting Up the Ops Director Cluster](#)

Note: These instructions use "node3" as the name of the host on the managed cluster you are using to do the installation; replace that with the name of a host on the cluster you want to manage.

Setting up the Managed Cluster

Follow these instructions on the managed cluster to set up the managed cluster to work with Ops Director.

1. If necessary, use the following `curl` commands to install the CA on the cluster, replacing `admin:admin` with the server login credentials, `node3` with the name of your bootstrap host, `/tmp/server.pem` with the location and name of your server CA, and `/tmp/client.pem` with the location and name of your client CA:

```
curl --digest -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/server.pem \
  "http://node3:8002/manage/v2/certificate-authorities?tag=opsdir:ssl-ca"

curl --digest -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/client.pem \
  "http://node3:8002/manage/v2/certificate-authorities?tag=opsdir:client-ca"
```


Note: You are not required to use different CAs for the client and server.

2. Use one of the following commands to obtain the ID of the managed cluster:

- XQuery:

```
xquery version "1.0-ml";
xdmp:cluster()
```

- REST (via `curl`):

```
curl -s --anyauth -u admin:admin http://builder:8002/manage/v2 \
  | grep "<id>" | sed "s/<id> //" | sed "s/<.*> //"
```

3. Create the following JSON file to create a certificate template for the hosts on the cluster (in this example, we call it `/tmp/md-template.json`), replacing `cluster_ID` with the managed cluster ID from the previous step:

```
{
  "template-name": "OpsDirector Template cluster_ID",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",
  "key-options": {
    "key-length": "2048"
  },
  "req": {
    "version": "0",
    "subject": {
      "countryName": "US",
      "stateOrProvinceName": "California",
      "localityName": "San Carlos",
      "organizationName": "MarkLogic Corporation"
    }
  }
}
```

4. Use `curl` to POST this JSON file to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-template.json \
  http://node3:8002/manage/v2/certificate-templates
```

5. Run the following code, substituting the names of each of the hosts in your cluster, separated by spaces, for `host_list`:

```
for host in host_list; do
  echo "Generating CSR for $host to /tmp/$host.csr ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:application/json" \
    -d "{\"operation\": \"generate-certificate-request\", \"
```

```

    \"common-name\": \"$host\", \"dns-name\": \"$host\" }" \
    "http://node3:8002/manage/v2/certificate-templates/OpsDirector
    Template cluster_ID" > /tmp/$host.csr
done

```

6. Get a signed server certificate for each CSR from the server CA in step [1](#).
7. To install these certificates, run the following code, substituting the names of each of the hosts in your cluster, separated by spaces, for *host_list*, and replacing `/tmp/host.crt` with the location and name of your saved certificate:

```

for host in host_list; do
  echo "Installing server certificate for $host from /tmp/$host.crt
  ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:text/plain" --data-binary @/tmp/$host.crt \
    "http://node3:8002/manage/v2/certificates?tag=opsdir:host-
    certificate"
done

```

Note: If you prefer, you can put all of the certificates in a single ZIP file and post that to the `/manage/v2/certificates` endpoint.

8. Verify that the certificates have been installed. Navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.
9. Use one of the following commands to obtain the ID of the CA from step [1](#), replacing `OpsDirector CA` with the name of your CA:

- XQuery:

```

xquery version "1.0-m1";
import module namespace pki = "http://marklogic.com/xdmp/pki" at
"/MarkLogic/pki.xqy";
xdmp:invoke-function(function () {
  /pki:certificate[pki:authority=true()][pki:host-name='OpsDirector
  CA']/pki:certificate-id/data()
}, map:entry("database", xdmp:security-database()))

```

- REST (via curl):

```

curl -s --anyauth -u admin:admin
"http://node3:8002/manage/v2/certificate-authorities?format=json" | jq
"\"certificate-authorities-default-list\".\"list-items\".\"list-
item\"[] | select(.nameref==\"OpsDirector CA\") | .idref" | tr -d ' '

```

Note: The `https://stedolan.github.io/jq/` tool must be on your system for the REST command to work.

In this example, the Certificate Authority is `5955670036362913372`.

10. Create the following JSON file (in this example, we use `/tmp/md-extsec.json`) to create an external security module, replacing `cluster_ID` with the name of your cluster:

```
{
  "external-security-name": "OpsDirectorSystem-cluster_ID",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
}
```

11. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-extsec.json \
  http://node3:8002/manage/v2/external-security
```

12. Create the following JSON file (in this example, we use `/tmp/md-client.json`) to generate a CSR for the Ops Director client certificate, matching the parameters to those from step [14](#) in the previous section:

```
{
  "operation": "client-certificate-request",
  "countryName": "US",
  "organizationName": "MarkLogic Corporation",
  "commonName": "OpsDirector"
}
```

13. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/md-client.json \
  http://node3:8002/manage/v2/credentials/secure \
  > /tmp/md-client.csr
```

The server responds by generating a private/public key pair, storing the relevant private key in the Security database and returning a client certificate request.

14. Get a signed *client* certificate for the CSR from the CA in step [1](#).
15. Create an XML file (in this example we use `/tmp/md-secure-credential.xml`) that constructs a secure credential for accessing the Ops Director cluster, replacing `cluster_ID`

with the cluster ID, *certificate* with the certificate you obtained in the previous step, and 8009 with the Ops Director system port:

```
<secure-credential-properties
xmlns="http://marklogic.com/manage/secure-credential/properties">
  <name>MarkLogic-OpsDirector-cluster_ID</name>
  <description>Secure credential to access the Ops Director
cluster</description>
  <certificate>certificate</certificate>
  <targets>
  <target>
    <uri-pattern>https://.*:8009/.*</uri-pattern>
    <authentication>basic</authentication>
  </target>
  </targets>
  <signing>>false</signing>
  <permissions>
  <permission>
    <role-name>admin</role-name>
    <capability>read</capability>
  </permission>
  </permissions>
</secure-credential-properties>
```

16. Use `curl` to POST this credential to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/xml" \
  --data-binary @/tmp/md-secure-credential.xml \
  http://node3:8002/manage/v2/credentials/secure
```

17. Create the following JSON file (in this example, we use `/tmp/md-user.json`) to construct a user called `opsdirector-system-user-cluster_ID`, replacing *cluster_ID* with the ID of the cluster, and replacing the parameters in the `external-name` field with the external name from step [12](#):

```
{
  "user-name": "opsdirector-system-user-cluster_ID",
  "description": "Ops Director system user",
  "role": [ "manage-admin", "security" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}
```

Note: The external name must *exactly* match what you put in the client certificate!

18. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
```

```
-d @/tmp/md-user.json \
http://node3:8002/manage/v2/users
```

Setting Up the Ops Director Cluster

Follow these instructions on the Ops Director cluster to set up the managed cluster to work with Ops Director.

1. Create a JSON file (in this example, we use `/tmp/od-client.json`) to generate a CSR for the Ops Director managed cluster client certificate, making sure to match the parameters in step [12](#):

```
{
  "operation": "client-certificate-request",
  "countryName": "US",
  "organizationName": "MarkLogic Corporation",
  "commonName": "OpsDirector"
}
```

2. Use `curl` to POST the request to the REST Management API on the Ops Director cluster:

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-client.json \
  http://node1:8002/manage/v2/credentials/secure \
  > /tmp/od-client.csr
```

The server responds by generating a private/public key pair, storing the relevant private key in the Security database, and returning a client certificate request.

3. Get a signed *client* certificate for that CSR from the CA.
4. Create an XML file (in this example, we use `/tmp/od-secure-credential.xml`) that constructs a secure credential for accessing the managed cluster, replacing *cluster_ID* with the name of your cluster ID and *certificate* with the certificate obtained in the previous step:

```
<secure-credential-properties
xmlns="http://marklogic.com/manage/secure-credential/properties">
  <name>opmdir-cluster_ID</name>
  <description>Secure credential to access the cluster_ID
cluster</description>
  <certificate>certificate</certificate>
  <targets>
    <target>
      <uri-pattern>https://.*:8003/.*</uri-pattern>
      <authentication>basic</authentication>
    </target>
  </targets>
  <signing>>false</signing>
```

```

    <permissions>
    <permission>
      <role-name>admin</role-name>
      <capability>read</capability>
    </permission>
  </permissions>
</secure-credential-properties>

```

5. Use `curl` to POST this to the REST Management API (port 8002):

```

curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/xml" \
  --data-binary @/tmp/md-secure-credential.xml \
  http://node1:8002/manage/v2/credentials/secure

```

Now you can manage the cluster using external certificates from the Admin UI.

2.10.2 Using Externally Signed Certificates with MarkLogic Server 9.0-9 and Later and Ops Director 2.0.1 and Later

The steps involved for configuring Ops Director to use externally signed certificates are:

- [Configure the Ops Director Cluster](#)
- [Configure the Managed Cluster](#)

2.10.2.1 Configure the Ops Director Cluster

Once you have obtained the CA file (in this example, it is `/tmp/opsdir.ca`), follow these steps to configure the Ops Director cluster:

1. If necessary, use `curl` to install the Certificate Authority on the cluster, replacing `/tmp/opsdir.ca` with the location and name of your CA file:

```

curl --anyauth -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/opsdir.ca \
  http://node1:8002/manage/v2/certificate-authorities

```

2. Install Ops Director:

```
./gradlew -PopsdirCa=external mlDeploy
```

You may need to specify additional options, such as `mlHost`, depending on your configuration.

3. Create a certificate template for the hosts on the cluster:

```

{
  "template-name": "OpsDirector-SSL-Template",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",

```

```

    "key-options": {
      "key-length": "2048"
    },
    "req": {
      "version": "0",
      "subject": {
        "countryName": "US",
        "stateOrProvinceName": "California",
        "localityName": "San Carlos",
        "organizationName": "MarkLogic Corporation"
      }
    }
  }
}

```

Note: Our examples use MarkLogic-specific values. You may use other values, but you must use them *consistently*.

4. Use `curl` to post this certificate template to the REST Management API, replacing `/tmp/od-template.json` with the location and name of your template filename:

```

curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-template.json \
  http://node1:8002/manage/v2/certificate-templates

```

5. Obtain the ID of the Certificate Authority from Step 1, replacing `OpsDirector CA` with the common name of your certificate of the Certificate Authority. The following XQuery code accomplishes this:

```

xquery version "1.0-m1";
import module namespace pki = "http://marklogic.com/xdmp/pki"
  at "/MarkLogic/pki.xqy";
xdmp:invoke-function(function () {
  /pki:certificate[pki:authority=true()] [pki:host-
name='OpsDirector CA']/pki:certificate-id/data()
}, map:entry("database", xdmp:security-database()))

```

If you prefer, the following `curl` command will do it with the REST Management API. Note that this command relies on the <https://stedolan.github.io/jq/> tool, which must be on your system.

```

curl -s --anyauth -u admin:admin \
  "http://node1:8002/manage/v2/certificate-authorities?format=json" \
  | jq ".\"certificate-authorities-default-list\".\"list-
items\".\"list-item\"[ ] \
  | select(.nameref==\"OpsDirector CA\") | .idref" \
  | tr -d '"'

```

Either of these techniques will give you the ID of the certificate authority.

6. Create a JSON file (in this example, we call it `/tmp/od-extsec.json`) that creates an external authentication configuration object, replacing `CA_ID` with the CA ID you obtained in the previous step:

```
{
  "external-security-name": "OpsDirectorSystem",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
  "ssl-client-certificate-authority": [ "CA_ID" ]
}
```

7. Use curl to POST this file on the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-extsec.json \
  http://node1:8002/manage/v2/external-security
```

8. Create an XML file (in this example, we call it `/tmp/od-server.xml`) to update the server configuration to modify some of the properties of the `OpsDirectorSystem` application server. In this example, replace *Insert_the_certificate_here* with your certificate:

```
<http-server-properties xmlns="http://marklogic.com/manage">
  <enabled>true</enabled>
  <authentication>application-level</authentication>
  <external-securities>
    <external-security>OpsDirectorSystem</external-security>
  </external-securities>
  <ssl-certificate-template>OpsDirector-SSL-Template</ssl-certificate-
template>
  <ssl-client-certificate-pems>
    <ssl-client-certificate-pem>Insert_the_certificate_here</ssl-client-
certificate-pem>
  </ssl-client-certificate-pems>
</http-server-properties>
```

9. Use curl to PUT this file on the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X PUT -H \
  "content-type:application/xml" \
  --data-binary @/tmp/od-server.xml \
  http://node1:8002/manage/v2/servers/OpsDirectorSystem/properties?gro
up-id=Default
```

Note: If you are not using the Default group for the Ops Director servers, update `group-id` accordingly.

10. Use the following shell script to generate certificate requests for each of the hosts, replacing *host_list* with the names of the hosts in your cluster separated by spaces:

```
for host in host_list; do
  echo "Generating CSR for $host to /tmp/$host.csr ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:application/json" -d "{\"operation\": \
    \"generate-certificate-request\", \"common-name\": \
    \"$host\", \"dns-name\": \"$host\" }" \
    //node1:8002/manage/v2/certificate-templates/OpsDirector-SSL-
  Template \
    > /tmp/$host.csr
done
```

11. Get a signed server certificate for each CSR from the CA in step 1.
12. To install the certificates, run the following code, replacing *host_list* with the names of the hosts in your cluster separated by spaces, and store the certificate in */tmp/hostname.crt* for each host:

```
for host in host_list; do
  "Installing server certificate for $host from /tmp/$host.crt ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:text/plain" --data-binary @/tmp/$host.crt \
    http://node1:8002/manage/v2/certificates
done
```

Note: If you prefer, you can put all of the certificates in a single ZIP file and post that to the */manage/v2/certificates* endpoint.

Warning Uploading a ZIP file of certificates fails if the ZIP file contains an empty directory entry.

13. Verify that the certificates have been installed. In the Admin UI, navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.
14. Create a JSON file (in this example, we call it */tmp/od-user.json*) that creates a user called *opsdirector-system-user*:

```
{
  "user-name": "opsdirector-system-user",
  "description": "Ops Director system user",
  "role": [ "opsdir-admin-internal" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}
```

Note: Make sure the external name *exactly* matches the client certificates generated in later steps.

15. Use `curl` to POST this to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-user.json \
  http://node1:8002/manage/v2/users
```

2.10.2.2 Configure the Managed Cluster

When using external certificates, configuring a new managed cluster requires changes on both the managed cluster *and* the Ops Director Cluster.

- [Setting up the Managed Cluster](#)
- [Setting Up the Ops Director Cluster](#)

Note: These instructions use "node3" as the name of the host on the managed cluster you are using to do the installation; replace that with the name of a host on the cluster you want to manage.

Setting up the Managed Cluster

Follow these instructions on the managed cluster to set up the managed cluster to work with Ops Director.

1. If necessary, use the following `curl` command to install the CA on the cluster, replacing `/tmp/opsdir.ca` with the location and name of your CA:

```
curl --anyauth -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/opsdir.ca \
  http://node3:8002/manage/v2/certificate-authorities
```

2. Use one of the following commands to obtain the ID of the managed cluster:

- XQuery:

```
xquery version "1.0-m1";
xdmp:cluster()
```

- REST (via `curl`):

```
curl -s --anyauth -u admin:admin http://builder:8002/manage/v2 \
  | grep "<id>" | sed "s/<id> //" | sed "s/<.*> //"
```

3. Create the following JSON file to create a certificate template for the hosts on the cluster (in this example, we call it `/tmp/md-template.json`), replacing `cluster_ID` with the managed cluster ID from the previous step:

```
{
  "template-name": "OpsDirector Template cluster_ID",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",
  "key-options": {
    "key-length": "2048"
  },
  "req": {
    "version": "0",
    "subject": {
      "countryName": "US",
      "stateOrProvinceName": "California",
      "localityName": "San Carlos",
      "organizationName": "MarkLogic Corporation"
    }
  }
}
```

4. Use `curl` to POST this JSON file to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-template.json \
  http://node3:8002/manage/v2/certificate-templates
```

5. Run the following code, substituting the names of each of the hosts in your cluster, separated by spaces, for `host_list`:

```
for host in host_list; do
  echo "Generating CSR for $host to /tmp/$host.csr ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:application/json" \
    -d "{\"operation\": \"generate-certificate-request\", \
    \"common-name\": \"$host\", \"dns-name\": \"$host\" }" \
    "http://node3:8002/manage/v2/certificate-templates/OpsDirector
  Template cluster_ID" > /tmp/$host.csr
done
```

6. Get a signed server certificate for each CSR from the CA in step 1.
7. To install these certificates, run the following code, substituting the names of each of the hosts in your cluster, separated by spaces, for `host_list`, and replacing `/tmp/host.crt` with the location and name of your saved certificate:

```
for host in host_list; do
  echo "Installing server certificate for $host from /tmp/$host.crt
  ..."
  curl -s --anyauth -X POST -u admin:admin -H \
    "content-type:text/plain" --data-binary @/tmp/$host.crt \
    http://node3:8002/manage/v2/certificates
done
```

Note: If you prefer, you can put all of the certificates in a single ZIP file and post that to the `/manage/v2/certificates` endpoint.

Warning Uploading a ZIP file of certificates fails if the ZIP file contains an empty directory entry.

8. Verify that the certificates have been installed. Navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.

9. Use one of the following commands to obtain the ID of the CA from step 1, replacing `OpsDirector CA` with the name of your CA:

- XQuery:

```
xquery version "1.0-m1";
import module namespace pki = "http://marklogic.com/xdmp/pki" at
"/MarkLogic/pki.xqy";
xdmp:invoke-function(function () {
/pki:certificate[pki:authority=true()] [pki:host-name='OpsDirector
CA']/pki:certificate-id/data()
}, map:entry("database", xdmp:security-database()))
```

- REST (via `curl`):

```
curl -s --anyauth -u admin:admin
"http://node3:8002/manage/v2/certificate-authorities?format=json" | jq
".\"certificate-authorities-default-list\".\"list-items\".\"list-
item\"[] | select(.nameref==\"OpsDirector CA\") | .idref" | tr -d ' "'
```

Note: The <https://stedolan.github.io/jq/> tool must be on your system for the REST command to work.

In this example, the Certificate Authority is 5955670036362913372.

10. Create the following JSON file (in this example, we use `/tmp/md-extsec.json`) to create an external security module, replacing `cluster_ID` with the name of your cluster and 5955670036362913372 with your CA:

```
{
  "external-security-name": "OpsDirectorSystem-cluster_ID",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
  "ssl-client-certificate-authority": [ "5955670036362913372" ]
}
```

11. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-extsec.json \
  http://node3:8002/manage/v2/external-security
```

12. Create the following JSON file (in this example, we use `/tmp/md-client.json`) to generate a CSR for the Ops Director client certificate, matching the parameters to those from step [14](#) in the previous section:

```
{
  "operation": "client-certificate-request",
  "countryName": "US",
  "organizationName": "MarkLogic Corporation",
  "commonName": "OpsDirector"
}
```

13. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/md-client.json \
  http://node3:8002/manage/v2/credentials/secure \
  > /tmp/md-client.csr
```

The server responds by generating a private/public key pair, storing the relevant private key in the Security database and returning a client certificate request.

14. Get a signed *client* certificate for the CSR from the CA in step [1](#).
15. Create an XML file (in this example we use `/tmp/md-secure-credential.xml`) that constructs a secure credential for accessing the Ops Director cluster, replacing *cluster_ID* with the cluster ID, *certificate* with the certificate you obtained in the previous step, and 8009 with the Ops Director system port:

```
<secure-credential-properties
xmlns="http://marklogic.com/manage/secure-credential/properties">
  <name>MarkLogic-OpsDirector-cluster_ID</name>
  <description>Secure credential to access the Ops Director

cluster</description>
  <certificate>certificate</certificate>
  <targets>
    <target>
      <uri-pattern>https://.*:8009/.*</uri-pattern>
      <authentication>basic</authentication>
    </target>
  </targets>
  <signing>>false</signing>
  <permissions>
```

```

    <permission>
      <role-name>admin</role-name>
      <capability>read</capability>
    </permission>
  </permissions>
</secure-credential-properties>

```

16. Use `curl` to POST this credential to the REST Management API (port 8002):

```

curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/xml" \
  --data-binary @/tmp/md-secure-credential.xml \
  http://node3:8002/manage/v2/credentials/secure

```

17. Create the following JSON file (in this example, we use `/tmp/md-user.json`) to construct a user called `opsdirector-system-user-cluster_ID`, replacing `cluster_ID` with the ID of the cluster, and replacing the parameters in the `external-name` field with the external name from step [12](#):

```

{
  "user-name": "opsdirector-system-user-cluster_ID",
  "description": "Ops Director system user",
  "role": [ "manage-admin", "security" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}

```

Note: The external name must *exactly* match what you put in the client certificate!

18. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```

curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/md-user.json \
  http://node3:8002/manage/v2/users

```

Setting Up the Ops Director Cluster

Follow these instructions on the Ops Director cluster to set up the managed cluster to work with Ops Director.

1. Create a JSON file (in this example, we use `/tmp/od-client.json`) to generate a CSR for the Ops Director managed cluster client certificate, making sure to match the parameters in step [12](#):

```

{
  "operation": "client-certificate-request",
  "countryName": "US",
  "organizationName": "MarkLogic Corporation",
  "commonName": "OpsDirector"
}

```

2. Use `curl` to POST the request to the REST Management API on the Ops Director cluster:

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-client.json \
  http://node1:8002/manage/v2/credentials/secure \
  > /tmp/od-client.csr
```

The server responds by generating a private/public key pair, storing the relevant private key in the Security database, and returning a client certificate request.

3. Get a signed *client* certificate for that CSR from the CA.
4. Create an XML file (in this example, we use `/tmp/od-secure-credential.xml`) that constructs a secure credential for accessing the managed cluster, replacing *cluster_ID* with the name of your cluster ID and *certificate* with the certificate obtained in the previous step:

```
<secure-credential-properties
xmlns="http://marklogic.com/manage/secure-credential/properties">
  <name>opmdir-cluster_ID</name>
  <description>Secure credential to access the cluster_ID
cluster</description>
  <certificate>certificate</certificate>
  <targets>
  <target>
    <uri-pattern>https://.*:8003/.*</uri-pattern>
    <authentication>basic</authentication>
  </target>
  </targets>
  <signing>>false</signing>
  <permissions>
  <permission>
    <role-name>admin</role-name>
    <capability>read</capability>
  </permission>
  </permissions>
</secure-credential-properties>
```

5. Use `curl` to POST this to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/xml" \
  --data-binary @/tmp/md-secure-credential.xml \
  http://node1:8002/manage/v2/credentials/secure
```

Now you can manage the cluster using external certificates.

2.10.3 Using Externally Signed Certificates with MarkLogic 9.0-8 and earlier or Ops Director 2.0 and Earlier

The following sections describe how to configure MarkLogic Server 9.0-8 and earlier and Ops Director 2.0 and earlier to use externally signed certificates.

The steps involved for configuring Ops Director to use externally signed certificates are:

- [Configure the Ops Director 2.0 and Earlier Cluster](#)
- [Manage a Cluster using Ops Director 2.0 and Earlier](#)

2.10.3.1 Configure the Ops Director 2.0 and Earlier Cluster

To configure the Ops Director 2.0 and earlier cluster:

1. Use `curl` to install the Certificate Authority on the cluster:

```
curl --anyauth -u admin:admin -X POST -H "content-type:text/plain" \  
  --data-binary @/tmp/opsdir.ca \  
  http://node1:8002/manage/v2/certificate-authorities
```

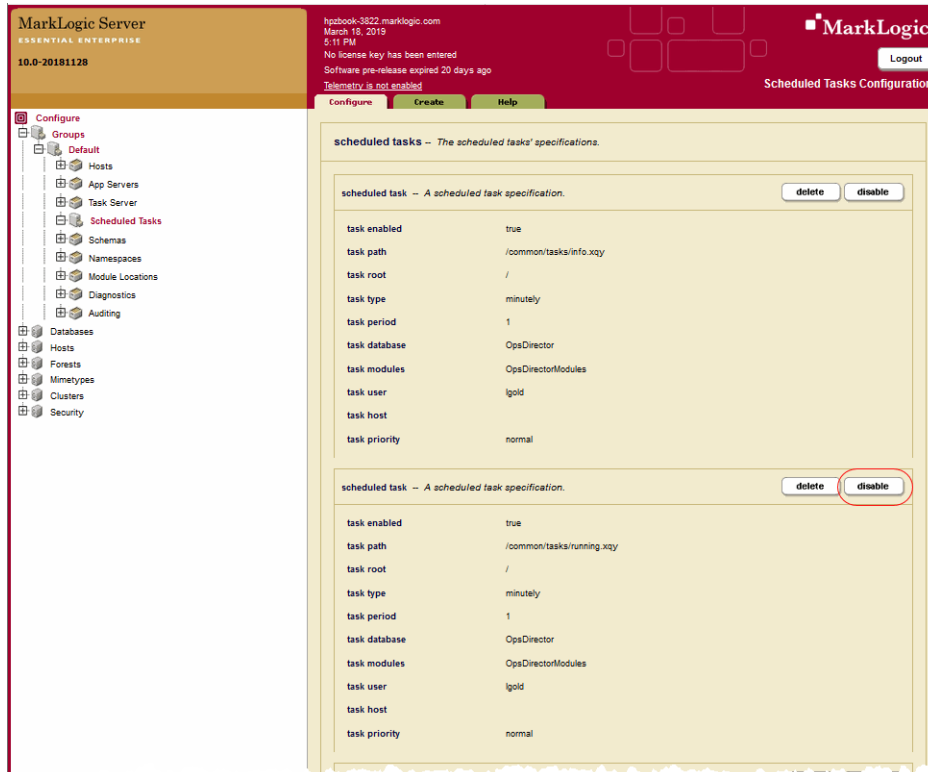
2. Install Ops Director.

- If you are using MarkLogic 9.0-7 or earlier, use the Admin UI. See “Installing Ops Director” on page 20 for details.
- If you are using MarkLogic 9.0-8 or later, use the Gradle installer for Ops Director 2.0-0:

```
./gradlew -PopsdirCa=external mlDeploy
```

Note: You may need to specify additional options, such as `mlHost`, depending on your configuration.

3. In the Admin UI, navigate to Configure > Groups > Default. On the Scheduled Tasks tab, disable the task “/common/tasks/running.xqy”:



4. Create the following JSON file (in this example, we use /tmp/od-template.json) to create a certificate template for the hosts on the cluster:

```
{
  "template-name": "OpsDirector-SSL-Template",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",
  "key-options": {
    "key-length": "2048"
  },
  "req": {
    "version": "0",
    "subject": {
      "countryName": "US",
      "stateOrProvinceName": "California",
      "localityName": "San Carlos",
      "organizationName": "MarkLogic Corporation"
    }
  }
}
```

5. Use `curl` to POST this to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-template.json \
  http://node1:8002/manage/v2/certificate-templates
```

6. Obtain the ID of the Certificate Authority from Step 1. If you are using a different certificate, replace `OpsDirector CA` with the common name of your certificate. The following XQuery code accomplishes this:

```
xquery version "1.0-m1";
import module namespace pki = "http://marklogic.com/xdmp/pki" at
"/MarkLogic/pki.xqy";
xdmp:invoke-function(function () {
  /pki:certificate[pki:authority=true()] [pki:host-name='OpsDirector
CA']/pki:certificate-id/data()
}, map:entry("database", xdmp:security-database()))
```

If you prefer, the following `curl` command will do it with the REST Management API. Note that this command relies on the <https://stedolan.github.io/jq/> tool, which must be on your system.

```
curl -s --anyauth -u admin:admin \
  "http://node1:8002/manage/v2/certificate-authorities?format=json" \
  | jq ".\"certificate-authorities-default-list\".\"list-
items\".\"list-item\"[ ] \
  | select(.nameref==\"OpsDirector CA\") | .idref" \
  | tr -d '''
```

Either of these techniques will give you the ID of the `OpsDirector CA` certificate authority (in this example, it is 5955670036362913372).

7. Create a JSON file (in this example, we call it `/tmp/od-extsec.json`) where the external security name is called `OpsDirectorSystem`, replacing the 5955670036362913372 with the number you got when running the previous step:

```
{
  "external-security-name": "OpsDirectorSystem",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
  "ssl-client-certificate-authority": [ "5955670036362913372" ]
}
```

8. Use `curl` to post the JSON code to the external security endpoint using the REST Management API (in this example, `http://node1:8002/manage/v2/external-security`):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/od-extsec.json \
  http://node1:8002/manage/v2/external-security
```

9. To update the server configuration, create an XML file (in this example, we call it `/tmp/od-server.xml`) to update the server configuration to modify some of the properties of the `OpsDirectorSystem` application server. In this example, the CA is `ssl-client-certificate-pem`, and replace *certificate* with your certificate:

```
<http-server-properties xmlns="http://marklogic.com/manage">
  <enabled>true</enabled>
  <authentication>application-level</authentication>
  <external-securities>
    <external-security>OpsDirectorSystem</external-security>
  </external-securities>
  <ssl-certificate-template>OpsDirector-SSL-Template</ssl-certificate-
template>
  <ssl-client-certificate-pems>
    <ssl-client-certificate-pem>certificate</ssl-client-certificate-pem>
  </ssl-client-certificate-pems>
</http-server-properties>
```

10. Use `curl` to PUT this file to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X PUT -H \
  "content-type:application/xml" \
  --data-binary @/tmp/od-server.xml \
  http://node1:8002/manage/v2/servers/OpsDirectorSystem/properties?gro
up-id=Default
```

Note: If you are not using the Default group for the Ops Director servers, update the `group-id` accordingly.

11. Generate server certificate requests for each of the hosts:
- Navigate to `Configure > Security > Certificate Templates > OpsDirector-SSL-Template`. In the right pane, select the Request tab.
 - Generate CSRs for all hosts.
 - On the Status tab, download the Pending Certificate Requests.
 - Get them signed as server certificates, and put all of the signed certificates in a ZIP file at `/tmp/od-certs.zip`.

12. Use `curl` to upload the certificates:

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/xml" \
  --data-binary @/tmp/od-certs.zip \
  http://node1:8002/manage/v2/certificates
```

Warning Uploading a ZIP file of certificates fails if the ZIP file contains an empty directory entry.

13. Verify that the certificates have been installed. Navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.

If any nodes still show a pending certificate request or a temporary certificate, go back to step <HyperlinkSerif>11. Make sure you select to generate requests for “All” nodes and repeat the steps to download, sign, and upload the certificates.

14. In the Admin UI, navigate to Configure > Groups > Default. On the Scheduled Tasks tab, enable the task “/common/tasks/running.xqy”:
15. Create a JSON file (in this example, we call it /tmp/od-user.json) that creates a user called `opsdirector-system-user`, replacing *password* with a password of your choice:

```
{
  "user-name": "opsdirector-system-user",
  "password": "password",
  "description": "Ops Director system user",
  "role": [ "opsdir-admin-internal" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}
```

16. Use curl to POST this to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-user.json \
  http://node1:8002/manage/v2/users
```

Note: The external name must *exactly* match the subject of the client certificates that you generated when setting up a managed cluster.

2.10.3.2 Manage a Cluster using Ops Director 2.0 and Earlier

When using external certificates, configuring a new managed cluster requires changes on both the managed cluster *and* the Ops Director Cluster.

- [Setting Up the Managed Cluster](#)
- [Setting Up the Ops Director Cluster](#)

Setting Up the Managed Cluster

Follow these instructions on the managed cluster to set up the managed cluster to work with Ops Director 2.0 and earlier.

1. Use the following `curl` command to install the CA on the cluster, replacing `/tmp/opsdir.ca` with your CA location and name:

```
curl --anyauth -u admin:admin -X POST -H "content-type:text/plain" \
  --data-binary @/tmp/opsdir.ca \
  http://node3:8002/manage/v2/certificate-authorities
```

2. Use one of the following commands to obtain the ID of the managed cluster:

- XQuery:

```
xquery version "1.0-m1";
xdmp:cluster()
```

- REST (via `curl`):

```
curl -s --anyauth -u admin:admin http://builder:8002/manage/v2 \
  | grep "<id>" | sed "s/<id> //" | sed "s/<.*> //"
```

3. Create the following JSON file to create a certificate template for the hosts on the cluster (in this example, we call it `md-template.json`), replacing `cluster_ID` with the managed cluster ID from the previous step:

```
{
  "template-name": "OpsDirector Template cluster_ID",
  "template-description": "Ops Director SSL Template",
  "key-type": "rsa",
  "key-options": {
    "key-length": "2048"
  },
  "req": {
    "version": "0",
    "subject": {
      "countryName": "US",
      "stateOrProvinceName": "California",
      "localityName": "San Carlos",
      "organizationName": "MarkLogic Corporation"
    }
  }
}
```

4. Use `curl` to POST this JSON file to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-template.json \
  http://node3:8002/manage/v2/certificate-templates
```

5. Use one of the following commands to obtain the ID of the CA from step 1:

- XQuery:

```
xquery version "1.0-m1";
import module namespace pki = "http://marklogic.com/xdmp/pki" at
"/MarkLogic/pki.xqy";
xdmp:invoke-function(function () {
  /pki:certificate[pki:authority=true()] [pki:host-name='OpsDirector
CA']/pki:certificate-id/data()
}, map:entry("database", xdmp:security-database()))
```

- REST (via `curl`):

```
curl -s --anyauth -u admin:admin
"http://node3:8002/manage/v2/certificate-authorities?format=json" | jq
"\"certificate-authorities-default-list\".\"list-items\".\"list-
item\"[] | select(.nameref==\"OpsDirector CA\") | .idref" | tr -d ' '
```

Note: The <https://stedolan.github.io/jq/> tool must be on your system for the REST command to work.

Either of these techniques will give you the ID of the `OpsDirector` CA certificate authority.

6. Create the following JSON file (in this example, we use `/tmp/md-extsec.json`) to create an external security module, replacing `cluster_ID` with the ID of the Ops Director cluster and `CA_ID` with your CA ID:

```
{
  "external-security-name": "OpsDirectorSystem-cluster_ID",
  "description": "Ops Director System",
  "authentication": "certificate",
  "cache-timeout": 300,
  "authorization": "internal",
  "ssl-client-certificate-authority": [ "CA_ID" ]
}
```

7. Use `curl` to POST this certificate template to the REST Management API (port 8002):

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/json" \
  -d @/tmp/md-extsec.json \
  http://node3:8002/manage/v2/external-security
```

8. Run the following XQuery code to generate *server* certificate requests for each of the hosts:

```
xquery version "1.0-m1";

import module namespace pki="http://marklogic.com/xdmp/pki"
  at "/MarkLogic/pki.xqy";

declare namespace x509="http://marklogic.com/xdmp/x509";

let $tid := pki:get-template-by-name("OpsDirector
Template")/pki:template-id
for $hid in xdmp:hosts()
  let $hostname := xdmp:host-name($hid)
  let $csr := pki:generate-certificate-request($tid, $hostname, (), ())
return
  $csr
```

9. Get a signed *server* certificate for each CSR from the CA in step 1.
10. Put all of the signed certificates in a ZIP file (in this example, we use `/tmp/mc-certs.zip`).
11. Use `curl` to upload the certificates:

```
curl --anyauth -u admin:admin -X POST -H \
  "content-type:application/zip" \
  --data-binary @/tmp/mc-certs.zip \
  http://node3:8002/manage/v2/certificates
```

Warning Uploading a ZIP file of certificates fails if the ZIP file contains an empty directory entry.

12. Verify that the certificates have been installed. Navigate to Configure > Security > Certificate Templates > OpsDirector-SSL-Template. In the right pane, select the Status tab. Make sure each node has a signed certificate and that it is not Temporary.
13. Run the following XQuery code to generate a *client* CSR:

Note: The subject must *exactly* match the external name of the `opsdirector-system-user` on the Ops Director cluster.

```
xquery version "1.0-m1";

declare namespace x509 = "http://marklogic.com/xdmp/x509";

let $country := "US"
let $state := ()
let $city := ()
let $org := "MarkLogic Corporation"
let $unit := ()
```

```

let $email := ()
let $common := "OpsDirector"

let $keys := xdmp:rsa-generate()
let $privkey := $keys[1]
let $pubkey := $keys[2]
let $req := <x509:req>
  <x509:version>0</x509:version>
  <x509:subject>
    { $country ! <x509:countryName>{.}</x509:countryName> }
    { $state !
<x509:stateOrProvinceName>{.}</x509:stateOrProvinceName> }
    { $city ! <x509:localityName>{.}</x509:localityName> }
    { $org ! <x509:organizationName>{.}</x509:organizationName>
}
    { $unit !
<x509:organizationalUnitName>{.}</x509:organizationalUnitName> }
    { $email ! <x509:emailAddress>{.}</x509:emailAddress> }
    { $common ! <x509:commonName>{.}</x509:commonName> }
  </x509:subject>
  <x509:publicKey>{$pubkey}</x509:publicKey>
</x509:req>
let $csr := xdmp:x509-request-generate($req, $privkey)
return
  ($csr, $privkey)

```

14. Get the CSR signed as a *client* certificate.
15. Run the following XQuery code to create a credential using the client certificate and its private key, replacing *certificate* with the client certificate and *private_key* with the private key for the certificate:

```

xquery version "1.0-m1";

import module namespace sec="http://marklogic.com/xdmp/security"
  at "/MarkLogic/security.xqy";

let $credname := "MarkLogic-OpsDirector"
let $opts := <options xmlns="xdmp:document-
get"><format>text</format></options>
let $cert := "certificate"
let $privkey := "private_key"
let $uri-pattern := "https://.*:8009/.*"
return
  xdmp:invoke-function(function() {
    sec:create-credential($credname,
      "A credential for accessing .*:8009",
      (), (), $cert, $privkey, fn:false(),
      sec:uri-credential-target($uri-pattern, "basic"),
      xdmp:permission("security", "read"))
  }, map:entry("database", xdmp:security-database()))

```


16. Create a JSON file (in this example, we call it `/tmp/od-user.json`) that creates a user called `opsdirector-system-user`, replacing *password* with a password of your choice:

```
{
  "user-name": "opsdirector-system-user",
  "password": "password",
  "description": "Ops Director system user",
  "role": [ "opsdir-admin-internal" ],
  "external-name": [ "C=US,O=MarkLogic Corporation,CN=OpsDirector" ]
}
```

Note: The external name must *exactly* match the subject of the client certificates that you generated when setting up a managed cluster.

17. Use curl to POST this to the REST Management API (port 8002):

```
curl -s --anyauth -X POST -u admin:admin -H \
  "content-type:application/json" \
  -d @/tmp/od-user.json \
  http://node3:8002/manage/v2/users
```

Setting Up the Ops Director Cluster

Follow these instructions on the Ops Director cluster to set up the managed cluster to work with Ops Director 2.0-0 and earlier.

1. Run the following XQuery code to generate a client CSR:

```
xquery version "1.0-ml";

declare namespace x509 = "http://marklogic.com/xdmp/x509";

let $country := "US"
let $state := ()
let $city := ()
let $org := "MarkLogic Corporation"
let $unit := ()
let $email := ()
let $common := "OpsDirector"

let $keys := xdmp:rsa-generate()
let $privkey := $keys[1]
let $pubkey := $keys[2]
let $req := <x509:req>
  <x509:version>0</x509:version>
  <x509:subject>
    { $country ! <x509:countryName>{.}</x509:countryName> }
    { $state !
      <x509:stateOrProvinceName>{.}</x509:stateOrProvinceName> }
    { $city ! <x509:localityName>{.}</x509:localityName> }
    { $org ! <x509:organizationName>{.}</x509:organizationName> }
```

```

    { $unit !
  <x509:organizationalUnitName>{.}</x509:organizationalUnitName> }
    { $email ! <x509:emailAddress>{.}</x509:emailAddress> }
    { $common ! <x509:commonName>{.}</x509:commonName> }
  </x509:subject>
  <x509:publicKey>{$pubkey}</x509:publicKey>
</x509:req>
let $csr := xdmp:x509-request-generate($req, $privkey)
return
  ($csr, $privkey)

```

Note: The subject must *exactly* match external name of the `opsdirector-system-user` on the managed cluster.

2. Get the CSR signed as a *client* certificate.
3. Run the following XQuery code to create a credential using the client certificate and its private key, replacing *cluster_ID* with the cluster ID, *certificate* with the client certificate, and *private_key* with the private key:

```

xquery version "1.0-ml";

import module namespace sec="http://marklogic.com/xdmp/security"
  at "/MarkLogic/security.xqy";

let $credname := "opsdir-cluster_ID"
let $opts := <options xmlns="xdmp:document-get"><format>text</format></options>
let $cert := "certificate"
let $privkey := "private_key"
let $uri-pattern := "https://.*:8003/.*"
return
  xdmp:invoke-function(function() {
    sec:create-credential($credname,
      "A credential for accessing SecureManage on the cluster",
      (), (), $cert, $privkey, fn:false(),
      sec:uri-credential-target($uri-pattern, "basic"),
      xdmp:permission("security", "read"))

    }, map:entry("database", xdmp:security-database()))

```

You can now manage the cluster from the Admin UI using Ops Director 2.0-0 and earlier.

2.11 Upgrading Ops Director

This section describes the upgrade process for the Ops Director application. It covers the following topics:

- [Upgrade Process Overview](#)
- [MarkLogic Server and Ops Director Upgrade Version Compatibility](#)
- [Upgrade Scenarios and Workflows](#)
- [Upgrade Error Prevention and Troubleshooting](#)

2.11.1 Upgrade Process Overview

Ops Director 1.0-0 and 1.1-1 were shipped with MarkLogic Server. This meant that a server upgrade automatically upgraded the version of Ops Director, from the perspective of *executable code*.

Starting with Ops Director 2.0, upgrading MarkLogic Server does not automatically upgrade Ops Director; you must upgrade Ops Director separately.

The MarkLogic Server upgrade process updates neither the *databases* that Ops Director uses nor the *configurations* of Ops Director application servers.

2.11.2 MarkLogic Server and Ops Director Upgrade Version Compatibility

The Ops Director application was part of the MarkLogic Server release up to Ops Director 1.1-1 and MarkLogic Server 9.0-6; the version of Ops Director depended on the version of the server installed:

- Ops Director version 1.0-0 was shipped with MarkLogic Server releases 9.0-4.
- Ops Director version 1.1-1 was shipped with MarkLogic Server release 9.0-5 and 9.0-6.
- Starting with MarkLogic Server 9.0-7 and Ops Director 2.0, Ops Director was separated from MarkLogic Server.

Ops Director Version	Can Be Installed on Clusters Running MarkLogic Releases
1.0-0, 1.1-1	9.0-4 to 9.0-6
2.0	9.0-7 to 9.0-8
2.0.1	9.0-7 and later versions of 9.0
2.0.1-1	9.0-9 and later, including 10.0-2 and later

Note: Ops Director cannot be installed on MarkLogic Server 10.0-1.

When you are upgrading Ops Director from one version to another, the MarkLogic Server version needs to be compatible with the versions listed in the preceding version compatibility matrix. If you are upgrading to a version of Ops Director that needs a different MarkLogic Server version, first upgrade the MarkLogic Server, and then upgrade Ops Director.

The following version compatibility matrix explains which versions of Ops Director can be upgraded on which versions of MarkLogic Server.

If Running Ops Director Version	Can Upgrade To Ops Director Version
1.0-0	1.1-1
1.1-1	2.0, 2.0.1
2.0	2.0.1
2.0.1	2.0.1-1

Note: When you are upgrading from Ops Director 1.1-1, you must first upgrade the MarkLogic Server version before you install Ops Director 2.x.

2.11.3 Upgrade Scenarios and Workflows

Possible scenarios and workflows of upgrading Ops Director are described in this section.

Note: The Ops Director upgrade procedure requires that you have the admin role. It is not sufficient to login as opsdir-admin. You must login with administrator credentials of the MarkLogic Server.

Scenario 1: Single host, MarkLogic Server upgraded, Ops Director databases need update.

In this scenario, the Ops Director Application Cluster is a single-node cluster. The single host has already been upgraded to the latest MarkLogic Server release. The Ops Director databases and configuration need to be updated.

Perform the following steps:

1. Launch the Ops Director application on the Ops Director host: in a browser window, enter the hostname and the port number of the OpsDirectorApplication server (by default, 8008).

2. Login with admin credentials. The following message displays:

```
Ops Director update required. Ops Director is being updated on this cluster.  
Please wait....
```

The waiting indicator shows that the update is in progress.

3. Once the update has completed and the Ops Director application is ready to run, the following message displays:

```
Ops Director update completed. Ops Director has been updated to version version.
```

4. Click **ok** to launch the Ops Director application.

Note: If you encounter errors in the Ops Director application after upgrade, clear the browser cache and reload the Ops Director application.

Scenario 2: Multiple hosts, one or more need upgrade of MarkLogic Server release.

In this scenario, the Ops Director Application Cluster is a multi-node cluster. Some of the hosts have already been upgraded to the latest MarkLogic Server release, but one or more hosts still have to be upgraded.

Perform the following steps:

1. Launch the Ops Director application on a host of the Ops Director cluster: in a browser window, enter the hostname and the port number of the OpsDirectorApplication server (by default, 8008).

2. Login with admin credentials. The following message displays:

```
Cluster upgrade must be completed. The following host(s) must be upgraded before  
Ops Director can run on this host: [hostname1], [hostname2], [hostname3]. Retry  
when upgrades have been completed.
```

3. For each host from the list of hosts provided in the message on the previous step, upgrade to the latest MarkLogic Server release. Follow the steps described in [Upgrading from Previous Releases](#) in the *Installation Guide*.

Note: After you installed the latest MarkLogic Server release on the host and started MarkLogic Server, make sure to open the Admin Interface in a browser (<http://hostname:8001/>), and, when the Admin Interface prompts you to upgrade the databases and the configuration files, click the button to confirm the upgrade. Also, make sure to reload Ops Director application in the browser (<http://hostname:8008/>) after the server is upgraded.

4. While hosts are being upgraded, you may check the status of the upgrade by clicking the **Retry** button on the message box from step 2.

Once all hosts have been upgraded, the following message displays: "Ops Director update required. Ops Director is being updated on this cluster. Please wait...".

The waiting indicator shows that the update is in progress.

Once the update is completed and the Ops Director application is ready to run, the following message displays:

```
Ops Director update completed. Ops Director has been updated to version version.
```

5. Click **ok** to launch the Ops Director application.

Note: If you encounter errors in the Ops Director application after the upgrade, clear the browser cache and reload the Ops Director application.

Scenario 3: Multiple hosts, all upgraded, Ops Director databases' update is in progress.

In this scenario, the Ops Director Application Cluster is a multi-node cluster. All hosts in the cluster have already been upgraded to the latest MarkLogic Server release. The Ops Director databases update is currently in progress (due to being initiated from a different host).

Perform the following steps:

1. Launch the Ops Director application on a host of the Ops Director cluster: in a browser window, enter the hostname and the port number of the OpsDirectorApplication server (by default, 8008).
2. Login with admin credentials. The following message displays:

```
Ops Director update in progress. Ops Director is being updated on this cluster. Please wait....
```

The waiting indicator shows that the update is in progress. Once the update is completed and the Ops Director application is ready to run, the following message displays:

```
Ops Director update completed. Ops Director has been updated to version [version].
```

3. Click **ok** to launch the Ops Director application.

Note: If you encounter errors in the Ops Director application after upgrade, clear the browser cache and reload the Ops Director application.

2.11.4 Upgrade Error Prevention and Troubleshooting

Before you begin to upgrade the Ops Director Application Cluster hosts to the latest MarkLogic Server release, make sure to close your browser containing the Ops Director application. If you do not close your browser, you may see errors in the browser window.

If you have started an Ops Director upgrade, do not connect new managed clusters until the upgrade process has completed. Adding managed clusters in the middle of an upgrade can cause error conditions.

If upgrading any host in the cluster to the latest MarkLogic Server release fails, address the problem before trying to launch Ops Director.

Undelivered data from managed clusters might be lost during the upgrade process, just like during network outage.

3.0 Navigating and Filtering Ops Director Views


This chapter describes the general navigation features of Ops Director. The navigation features that are specific to a view are described in its respective chapter.

This chapter covers the following topics:

- [Main Navigation Bar](#)
- [Navigating Resource Views](#)
- [Date and Time Filters](#)
- [Navigation Icons of Ops Director Views](#)
- [Preserving View States](#)

3.1 Main Navigation Bar

In the context of this document, *view* is a top-level page of Ops Director UI, which you can access from the menu bar. The menu bar appears at the top of every Ops Director view.



[MONITOR View](#) is an aggregate view that shows the summary of the state of resources in your enterprise.

[MANAGE View](#) shows details for a selected MarkLogic resource or resource category, in the form of charts, tables, and other displays.

[ANALYZE View](#) presents a comprehensive set of detailed charts that allow to analyze utilization and performance of system resources, such as disks, CPU, memory, network, databases, and servers.

[SUPPORT view](#) shows alerts for application servers in your enterprise, allows you to view and filter server logs, and displays tasks that run on your managed clusters.

[CONSOLE SETTINGS View](#) allows you to configure role-based access control to resources, manage user accounts, manage licensing, and define telemetry settings.



Click to view a display of latest system alerts. The Alerts icon is persistent across the application. It displays a count of alerts that require attention. To glance through the new alerts quickly in the internal screens of the application, select the alert icon to display the alerts panel as an overlay. Alerts are filtered by severity (All, Critical, Warning, and Information). Each alert has a menu with a list of relevant actions (Acknowledge, View Logs). Click **View Alert Details** to view the Support / System Alerts screen.



Click to provide help.



Click the User icon for a display with the name of the user currently logged in and an option to log out of Ops Director. Administrators are advised to log out of Ops Director when not in use to guard against unauthorized agents that could compromise the security of the system.

In general, clicking on a region of an Ops Director page that represents a resource will enable you to drill down to a more detailed representation of that resource. The context of the detailed representation is determined by the current view. For example, when displaying the Monitor view, clicking **Clusters** displays a color-coded view of the health of each cluster from which you can select to drill down into more detail on the health metrics for that cluster, such as CPU, memory, and network loads. When displaying the Manage view, clicking **Clusters** displays a list of clusters from which you can select to drill down into more detail on the metrics for that cluster.

For each Ops Director view, there is a dedicated chapter in this guide that provides detailed description of the view.

3.2 Navigating Resource Views

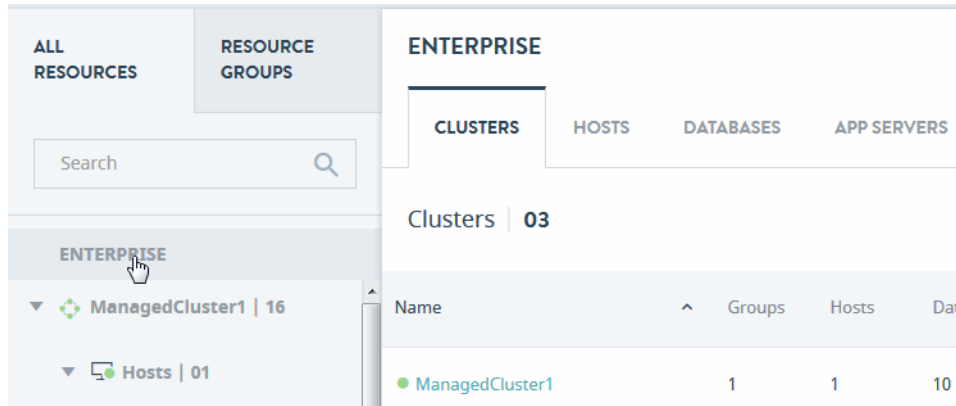
The Manage and ANALYZE views provide you with an ability to modify your view of the resources. The left navigation panel of these views enables navigating across individual resources or groups of resources. The main resource view options are:

- [View All Resources](#)
- [View Resource Groups](#)

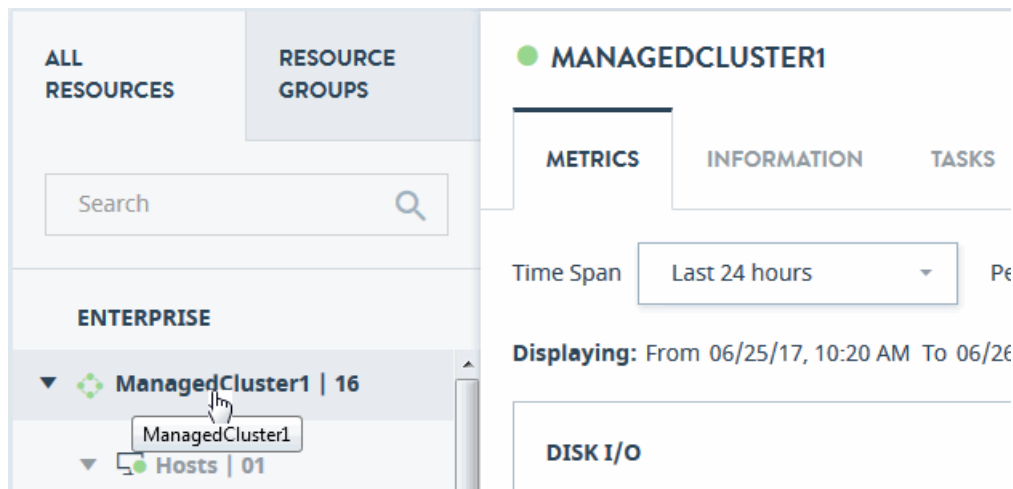
3.2.1 View All Resources

Select **All Resources** to display a consolidated view of all of the resources in your enterprise. The lists can be expanded and collapsed by clicking on the name. Each Cluster will show Hosts, Databases, and App Servers that belong to that particular cluster.

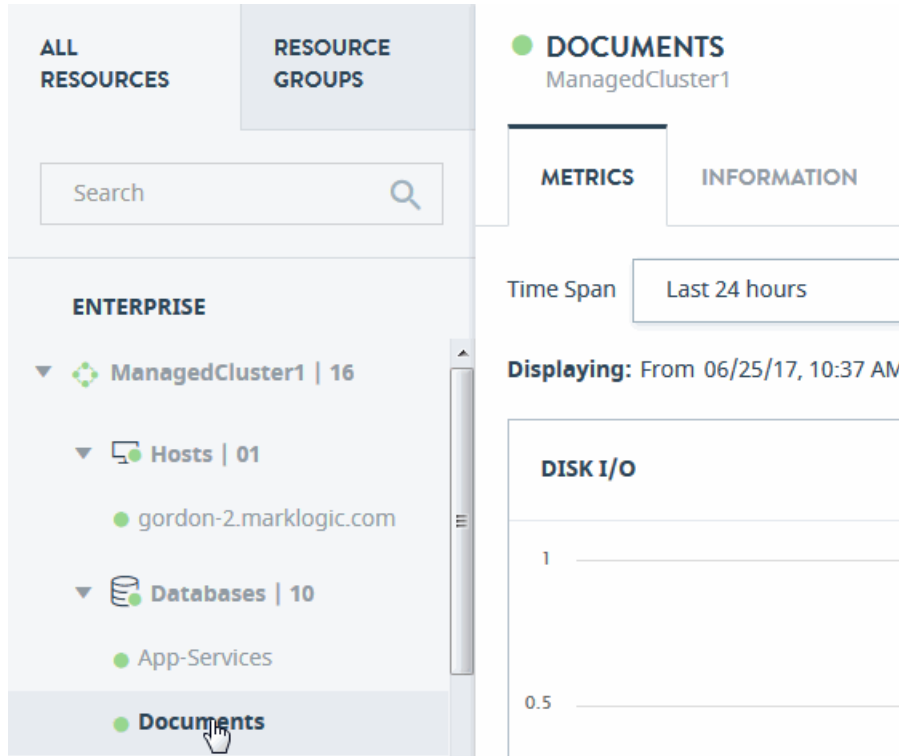
The health status of resources are shown at both parent and child level. The colors used to represent resource status are described in “Terms” on page 8.



You can drill down to more specific views of a particular cluster or a group of similar resource types in your cluster.

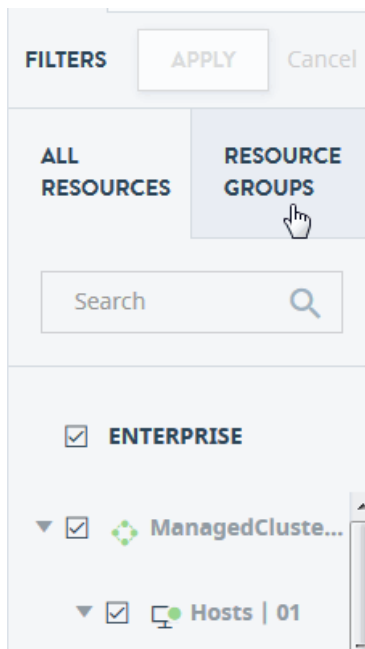


You can also drill down to a specific resource within a resource type, for example a specific database or a specific application server.



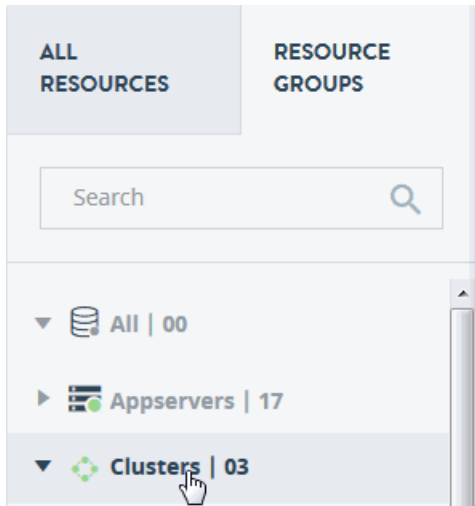
3.2.2 View Resource Groups

Select the **Resource Groups** tab to display a consolidated view of all the defined resource groups. If a group contains Clusters, collapsing or expanding each cluster hides or reveals members of the resource group.

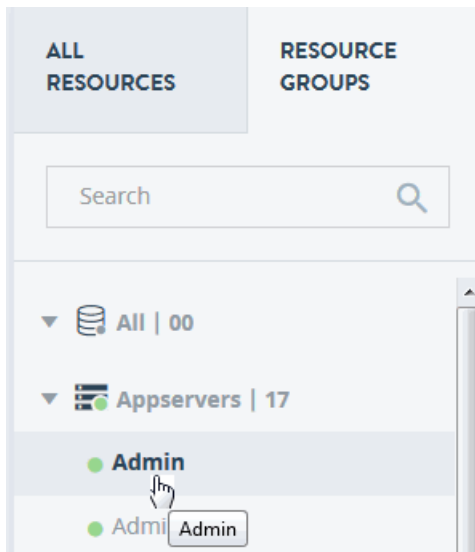


For more information on resource groups, see “Resource Groups” on page 223.

You can drill down to more specific views of a resources in a particular group or a specific resource within a resource group.



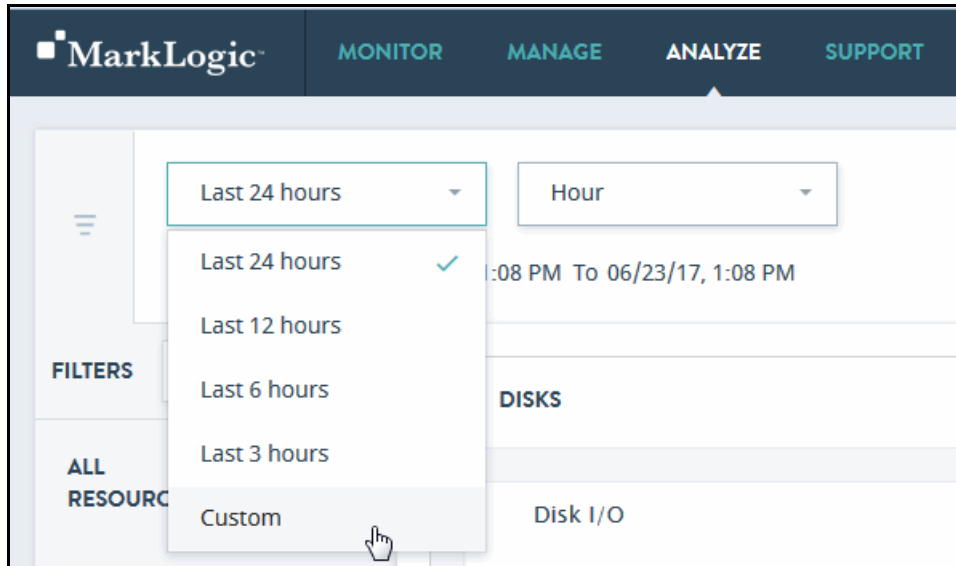
Click on an object or resource in the navigation tree to display relevant information in the content area. The selection is highlighted in the navigation tree.



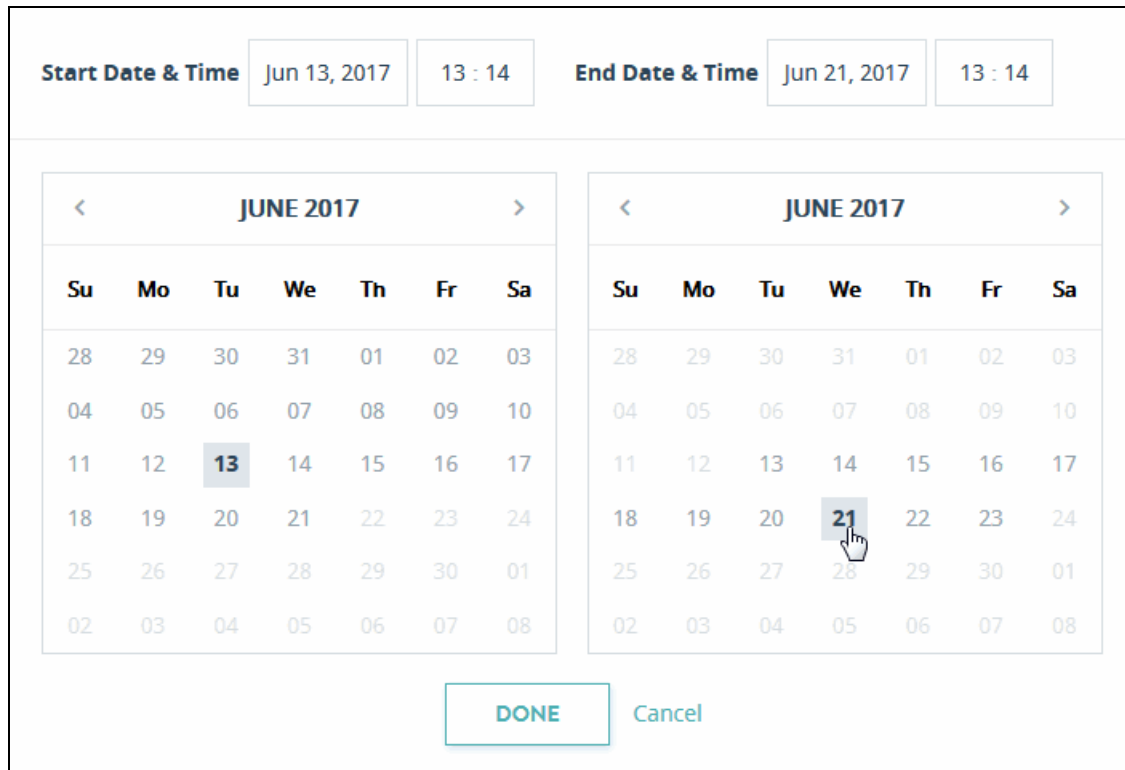
3.3 Date and Time Filters

The Analyze and Support views of Ops Director, as well as some tabs of the Manage view, provide a date and time filter that enable you to modify your data views.


Select from a list of commonly used time periods, or select **Custom**.



The **Custom** time filter allows you to choose a specific start and end date and time.



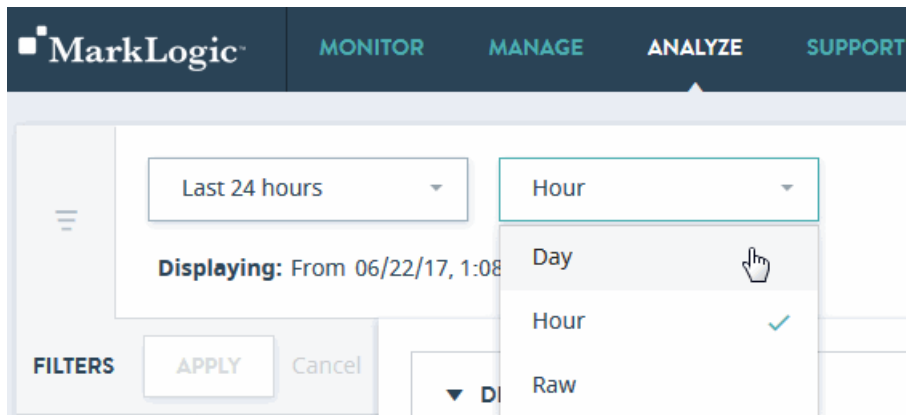
The date and time the graphs were initially and last updated is shown under the time filters.



Select the Refresh icon to refresh the widget with up-to-date information.

You can use the interval filter to display the data in daily or hourly intervals, or select **Raw** to display all of the collected data.

Note: The interval filter controls how the data is displayed in Ops Director and will not display less frequently collected data. For example, to set the log interval filter in Ops Director to **Raw**, you must also set Ops Director Metering to **raw** in the Admin Interface, as described in Step 9 in “Installing Ops Director” on page 20.



3.4 Navigation Icons of Ops Director Views

Each view allows you to navigate and filter the displayed data. Common navigation and filter icons are shown in the following table. Not all views have all of the icons.



Enter the name of the resource or object to locate in this view. Wildcards are supported.



Click to refresh the view with up-to-date information, based on the current set of filter settings.



Click to export data as a CSV (Comma Separated Values) file ([Resource Groups](#), [MANAGE View](#), [SUPPORT view](#)).



All records that satisfy the set filter parameters are exported with all the columns, regardless of whether they are visible or hidden in the table.

Click to expand and isolate graph ([ANALYZE View](#)).



Click to provide details on this metric ([ANALYZE View](#)).



Click to sort the listed results. Default sorting is always enabled on the first column of the table. The up/down arrow is used in the column header to indicate the sorting state. You can click on the same column again to switch between sorting a column in ascending or descending manner.




Click to view the graph legend. Below is an example of the legend for the CPU graph.




Click to select which resource rows to display in the table.


Hover over an item in the legend to highlight the corresponding line in the graph.



Hover on a period point to view what operation was taking place at that point in time.



The following is an example of the resource rows for the Cluster group. Select to display rows in the table.



3.5 Preserving View States

Main views of Ops Director preserve their state when you navigate between them. For example, if you select a resource in the Manage view, drill down to a detailed view for that resource, and after that navigate to the ANALYZE view, when you navigate back to the Manage view, you will see the detailed view of that previously selected resource. This allows you to bookmark the state of any view and return to it later.

In particular, the state of the resource navigation panel in the Manage view, including search, filter, and selection, is preserved when leaving and then returning to the view. Also, the state of the ANALYZE view is preserved when leaving and returning to it, which includes filter selections, collapsed or expanded state of the charts, and selected time span.

4.0 MONITOR View

The Monitor view is an aggregate view that shows the summary of the state of resources in your enterprise.

This chapter covers the following topics:

- [Monitoring Dashboards](#)
- [Overview Pages](#)

4.1 Monitoring Dashboards

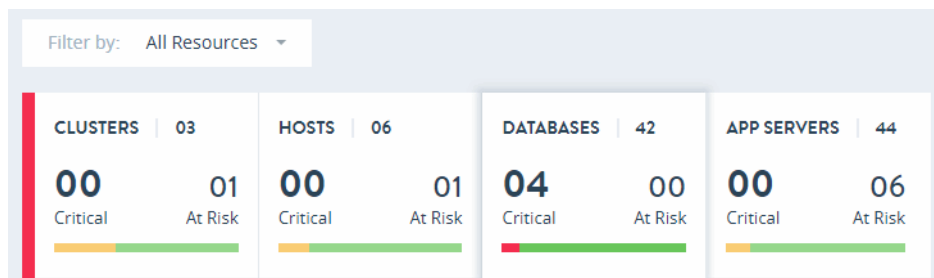
The default Monitor view displays the following dashboards, graphs, and panels:

- [Key Performance Indicators](#)
- [Cluster Problem Distribution](#)
- [Top Problematic Hosts](#)
- [Busiest Servers](#)
- [Slowest Servers](#)
- [Alerts Panel](#)
- [Filtering by Resource](#)


4.1.1 Key Performance Indicators

The key performance indicators panel is a set of active tabs for each resource type: clusters, hosts, databases, and application servers.

Selecting a resource type tab leads to a heat map view for that resource type. These tabs provide at-a-glance statistics, including overall system health and resource health for clusters, hosts, databases, and application servers.



Overall system health is indicated with a color-coded vertical bar at the top left side of the Monitor view, where red indicates a Critical condition, yellow indicates At Risk, and green indicates Healthy.




The overall health of the resources is calculated from a score generated for each alert type (only critical, at-risk, and unknown alerts are considered). The alert type that scores highest represents overall health of the system. The color of the vertical bar is the same as the alert type with highest score. Since there are no alerts for healthy, its score is calculated by getting a count of all healthy resources.

When only a subset of alerts are displayed, the overall health is calculated from the visible alerts. If there are no alerts, then the overall health is represented by green.

The overall health for the individual resources (clusters, hosts, databases, and application servers) is indicated by a combination of graphical and textual elements. For each resource type, the number of resources is presented; a horizontal bar depicts the distribution of health status as Critical (red), At Risk (yellow), Offline (light gray), Maintenance (dark gray), and Healthy (green); and the number of resources that are Critical and At Risk for a selected resource group.

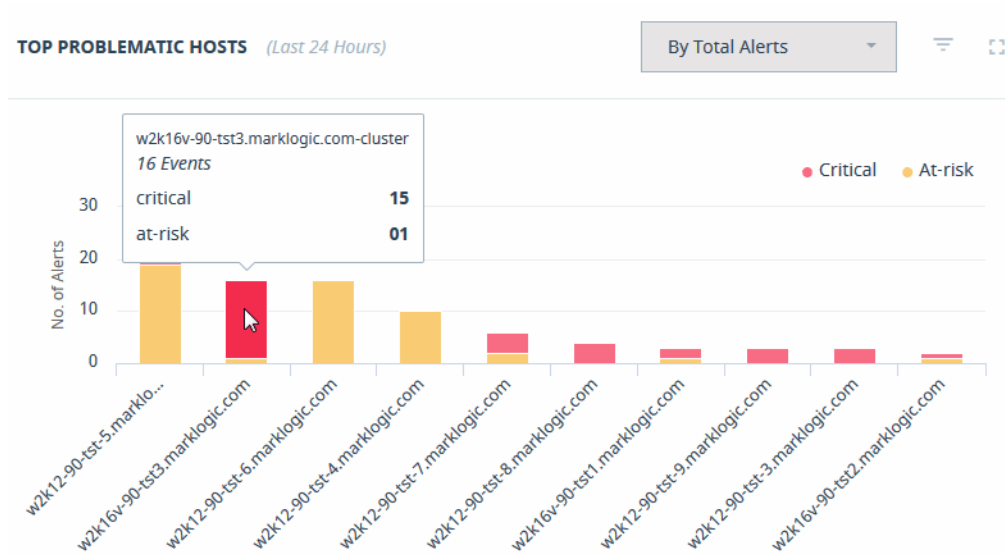
4.1.2 Cluster Problem Distribution

The Cluster Problem Distribution graph provides a visualization of the time and date of cluster events over the last seven days. Hover over the event icon to view details.




4.1.3 Top Problematic Hosts

The Top Problematic Hosts graph provides a visualization of the hosts in your enterprise that have experienced the greatest number of events over the last 24 hours. Hover over the host bar to view details.



The Top Problematic Hosts graph can be sorted by a count of Total Alerts, or weighted based on the severity of the messages.




4.1.4 Busiest Servers

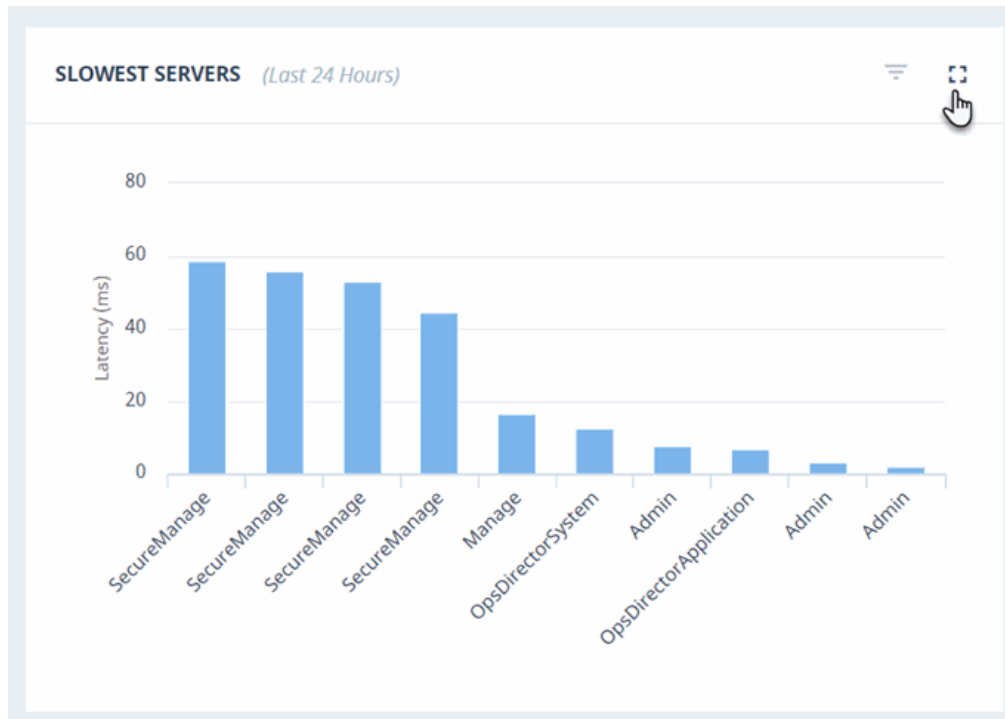
The Busiest Servers graph provides a visualization of the App Servers in your enterprise that have experienced the greatest request rate activity over the last 24 hours. Hover over the App Server bar to view details.

4.1.5 Slowest Servers

The Slowest Servers graph displays the App Servers in your enterprise that have the greatest latency the most activity over the last 24 hours. Hover over the App Server bar to view details.

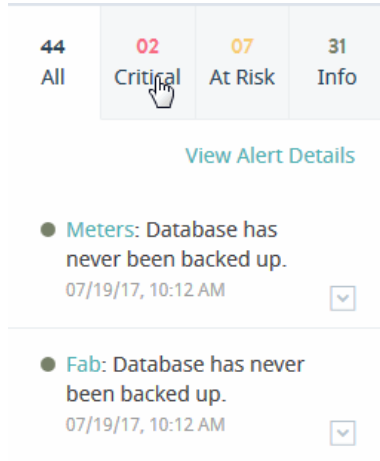


All graphs can be switched to a full-size window by clicking the maximize icon. You may click on the minimize icon or click anywhere on the screen to get back to the previous state.



4.1.6 Alerts Panel


The Alerts panel on the right-hand side lists the most recent system alerts for your enterprise. The Alerts panel lists up to 50 latest unacknowledged, open, and closed alerts. Every 30 seconds, the list is refreshed with latest 50 alerts.



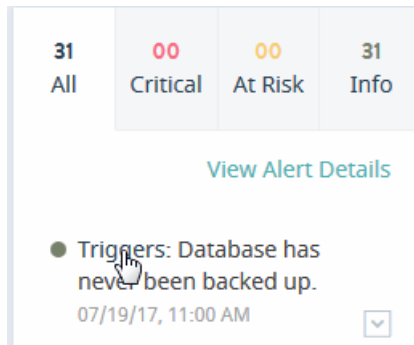
The Alerts panel contains tabs for Critical, At Risk, and Info alert types, as well as the All tab that shows all the above alert types. There are no tabs for other alert types, such as Offline and Maintenance. You can navigate the tabs to list events by their level of severity.

Tab	Description
All	Display all alerts.
Critical	Display only Critical alerts.
At Risk	Display only At Risk alerts.
Info	Display only Info alerts.

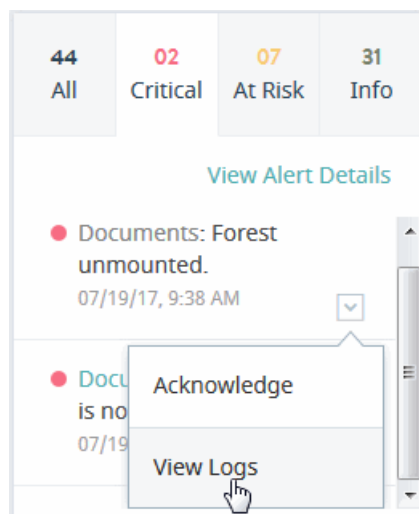
Click **View Alert Details** to navigate to the System Alerts page in the Support view. This page is described in “System Alerts” on page 251.



Click on the alert to navigate to the Metrics page for the resource in the Manage view. These pages are described in the following sections: “Cluster Metrics” on page 113, “Host Metrics” on page 129, “Database Metrics” on page 144, and “App Server Metrics” on page 175.



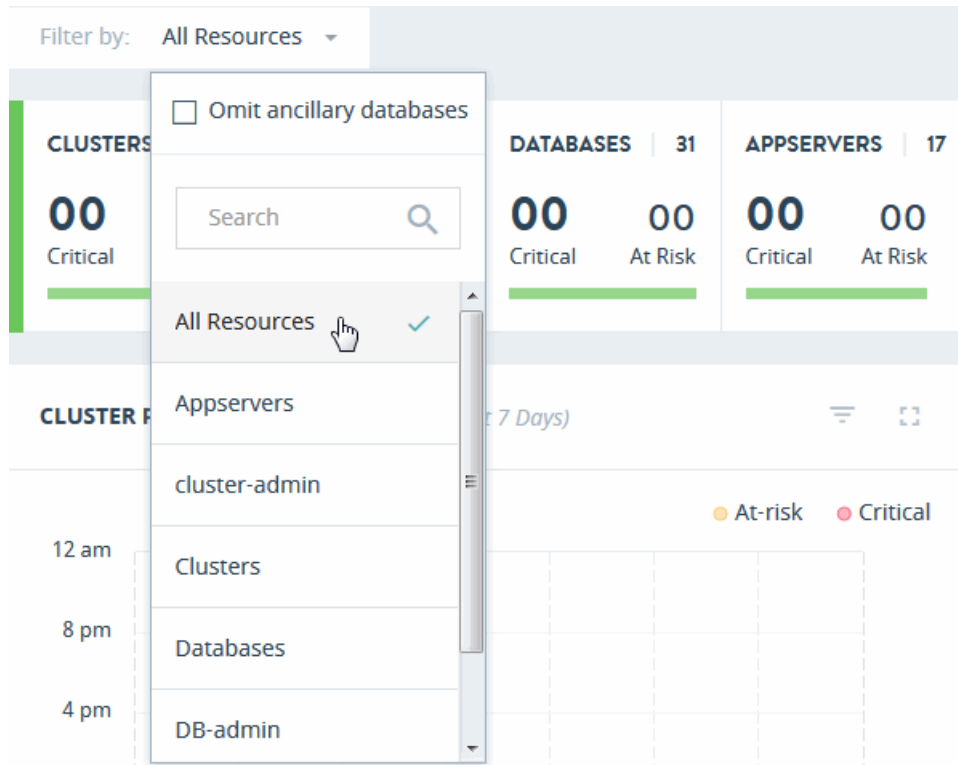
Next to each alert in the list is a menu where you can mark the alert as acknowledged, which suppresses the alert from further notifications, as described in “System Alerts” on page 251. Selecting **View Logs** displays the list of log files for the resource associated with the alert, as described in “Event Logs” on page 257.



4.1.7 Filtering by Resource

Use the **Filter by** menu to specify the Resource Groups to monitor relevant information in this view.

Either select **All Resources**, or choose from the available Resource Groups. Select **Omit ancillary databases** to omit system databases, such as App-Services, Extensions, Fab, Last-Login, Meters, Modules, Schemas, Security, and Triggers, from this view.



4.2 Overview Pages

The following sections describe each overview page in detail:

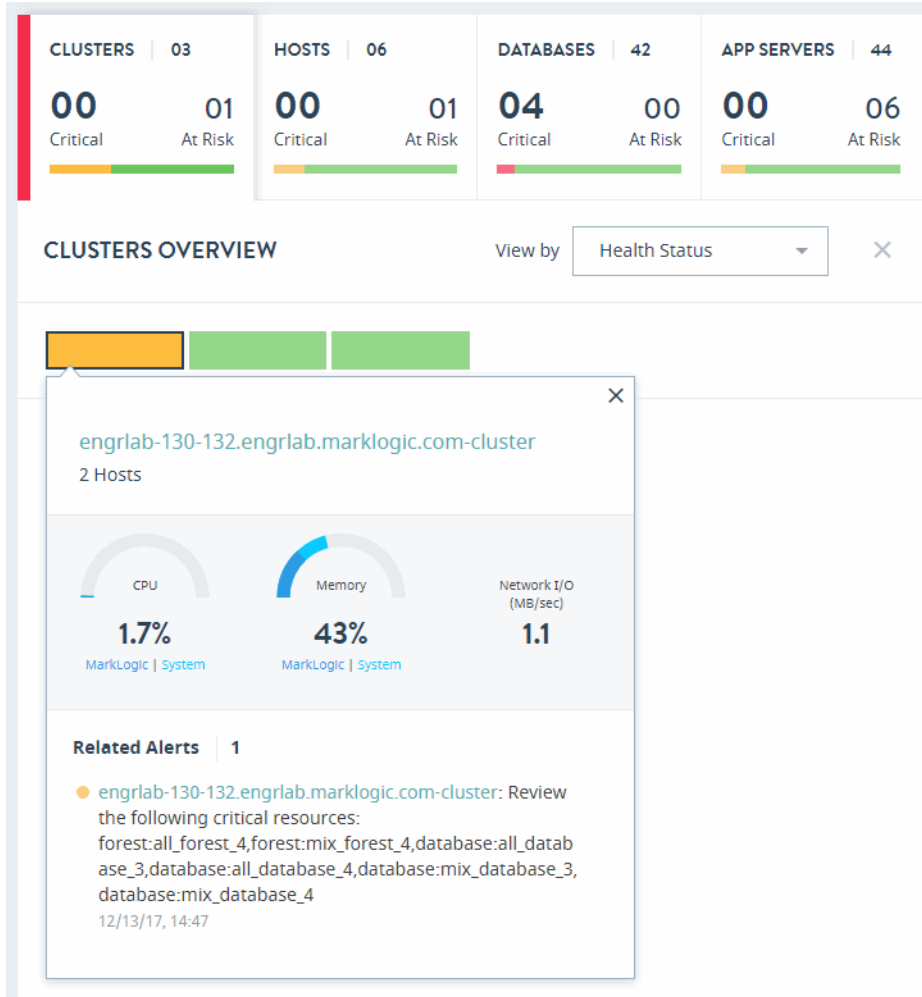
- [Clusters Overview](#)
- [Hosts Overview](#)
- [Databases Overview](#)
- [App Servers Overview](#)

4.2.1 Clusters Overview


Click on the **Clusters** tab in the key performance indicators panel to open the heat map view for clusters.

The CLUSTERS OVERVIEW page displays colored rectangles, each representing a cluster. You can quickly grasp the health status of large number of clusters and comprehend the overall impact at the enterprise or group level.


Clicking on a cluster displays the key performance metrics and any alerts for that cluster. The gauges represent resource consumption with percentage values.



In the CLUSTERS OVERVIEW page, you can sort the resources either by **Health Status** or **Name**.



To close the CLUSTERS OVERVIEW page and return to the top-level Monitor view, click the **X** to the right of the “View by” menu.

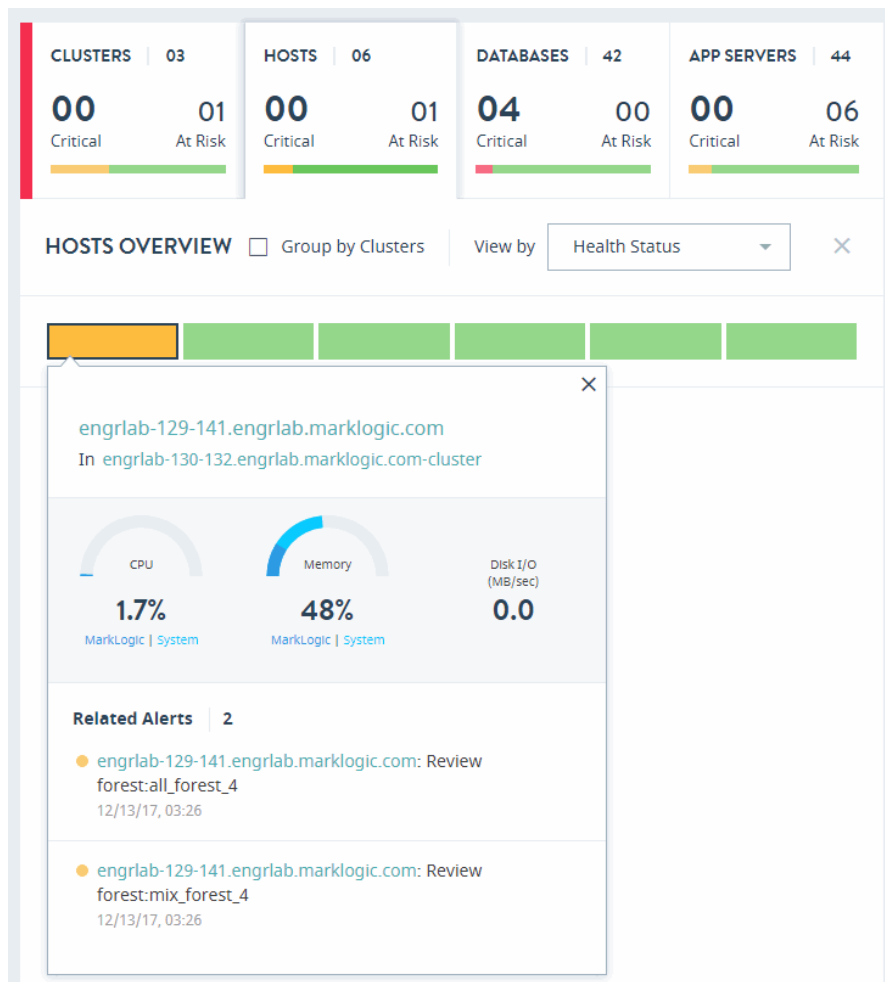


4.2.2 Hosts Overview


Click on the **Hosts** tab in the key performance indicators panel to open the heat map view for hosts.

The HOSTS OVERVIEW page displays colored rectangles, each representing a host. You can quickly grasp the health status of large number of hosts and comprehend the overall impact at the enterprise/group level.

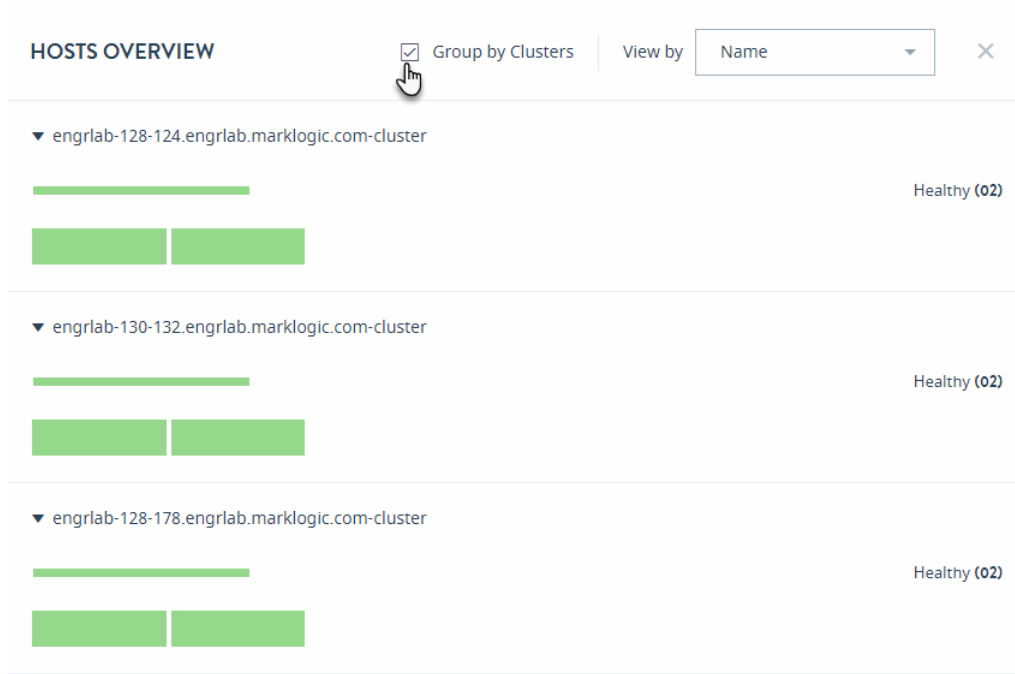
Click on a host to display the key performance metrics and any alerts for that host. The gauges represent resource consumption with percentage values.



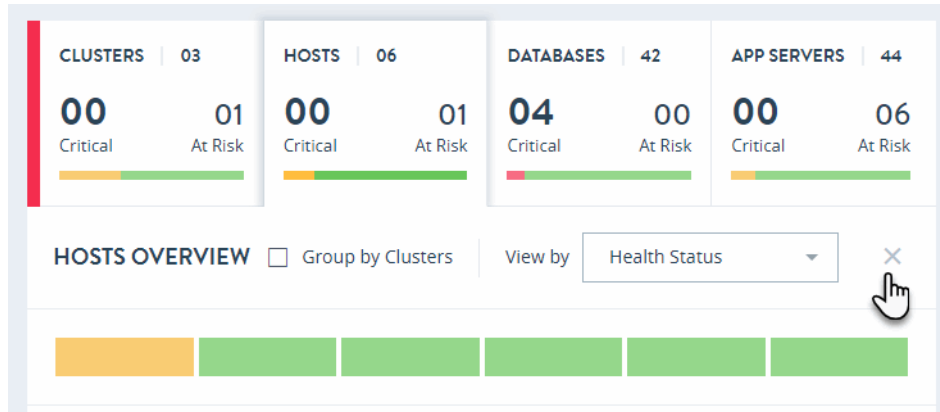
In the HOSTS OVERVIEW page, you can sort the resources by either **Health Status** or **Name**.



Select **Group by Clusters** to sort your view of the hosts by their clusters.



To close the HOSTS OVERVIEW page and return to the top-level Monitor view, click the **X** to the right of the “View by” menu.



4.2.3 Databases Overview

Click on the DATABASES tab in the key performance indicators panel to open the heat map view for databases.

The DATABASES OVERVIEW page displays colored rectangles, each representing a database. You can quickly grasp the health status of large number of databases and comprehend the overall impact at the enterprise or group level.

Clicking on a database displays the key performance metrics and any alerts for that database. The gauges represent resource consumption with percentage values.

CLUSTERS | 03 HOSTS | 06 DATABASES | 42 APP SERVERS | 44

00 Critical 01 At Risk 00 Critical 01 At Risk 04 Critical 00 At Risk 00 Critical 06 At Risk

DATABASES OVERVIEW Group by Clusters View by: Health Status

all_database_4
For engr1ab-130-132.engr1ab.marklogic.com-cluster
Encryption OFF Total Forest Size: 0 MB in 1 Forest

Avg. Forest Size (MB): 0.0 Largest: 0.0 MB
Lowest Forest Free Space (MB): 0
Cache Ratio: 0.0%

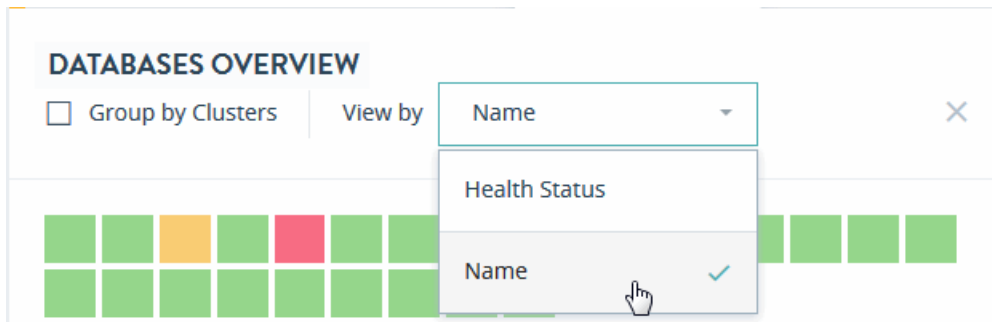
Related Alerts | 3

- all_database_4: Database has never been backed up. 12/13/17, 03:26
- all_database_4: Database is offline. 12/13/17, 03:26
- all_database_4: Database is not available. 12/13/17, 03:26

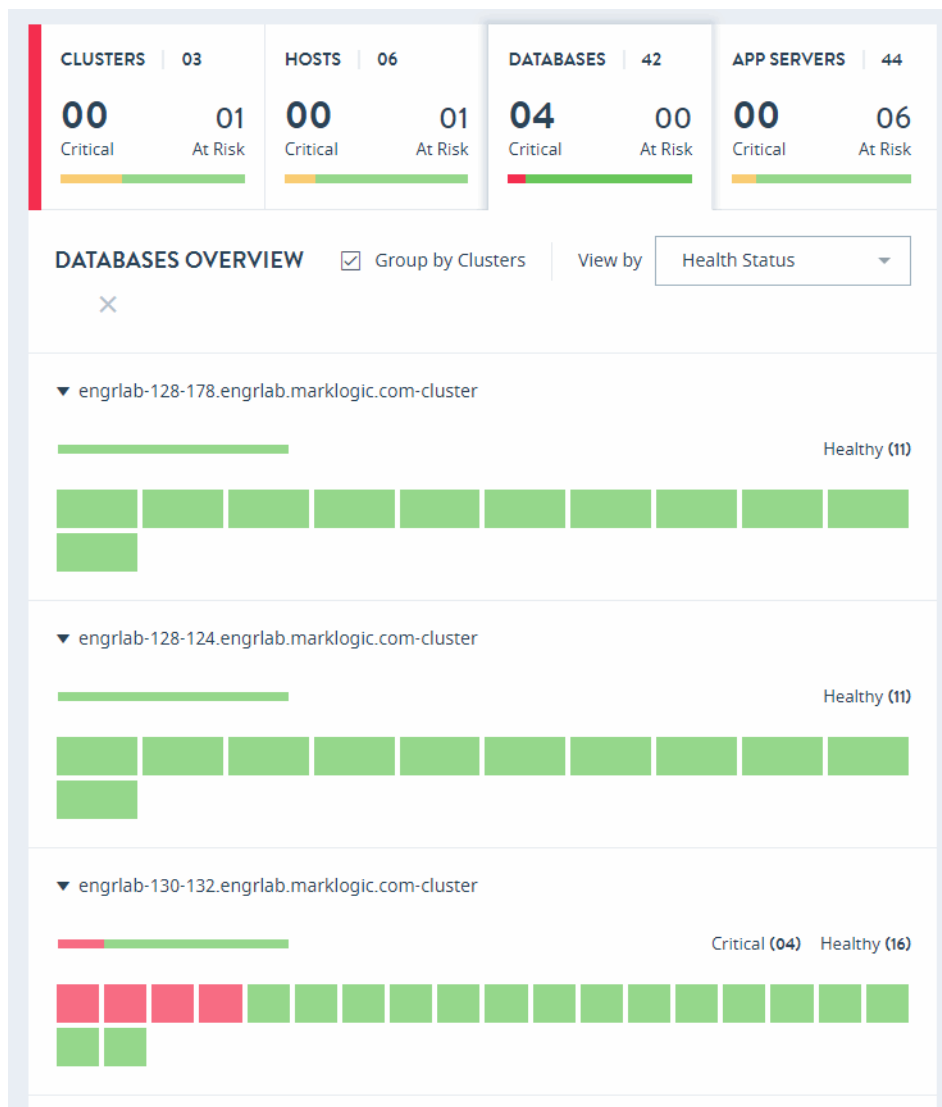
In the DATABASES OVERVIEW page, you can sort the resources either by **Health Status** or **Name**.

DATABASES OVERVIEW

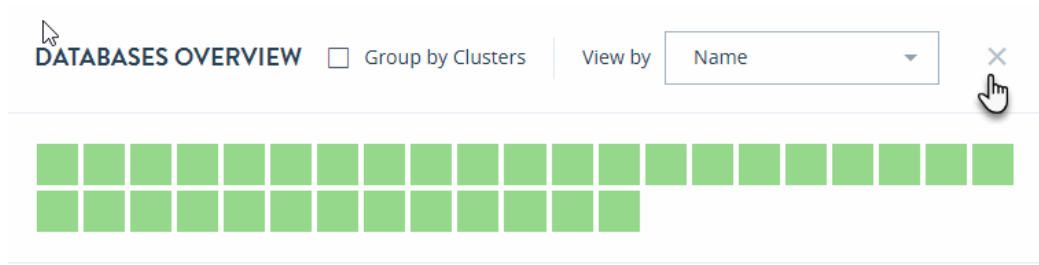
Group by Clusters View by: Health Status



Select **Group by Clusters** to sort your view of the databases by their clusters.



To close the DATABASES OVERVIEW page and return to the top-level Monitor view, click the **X** to the right of the “View by” menu.

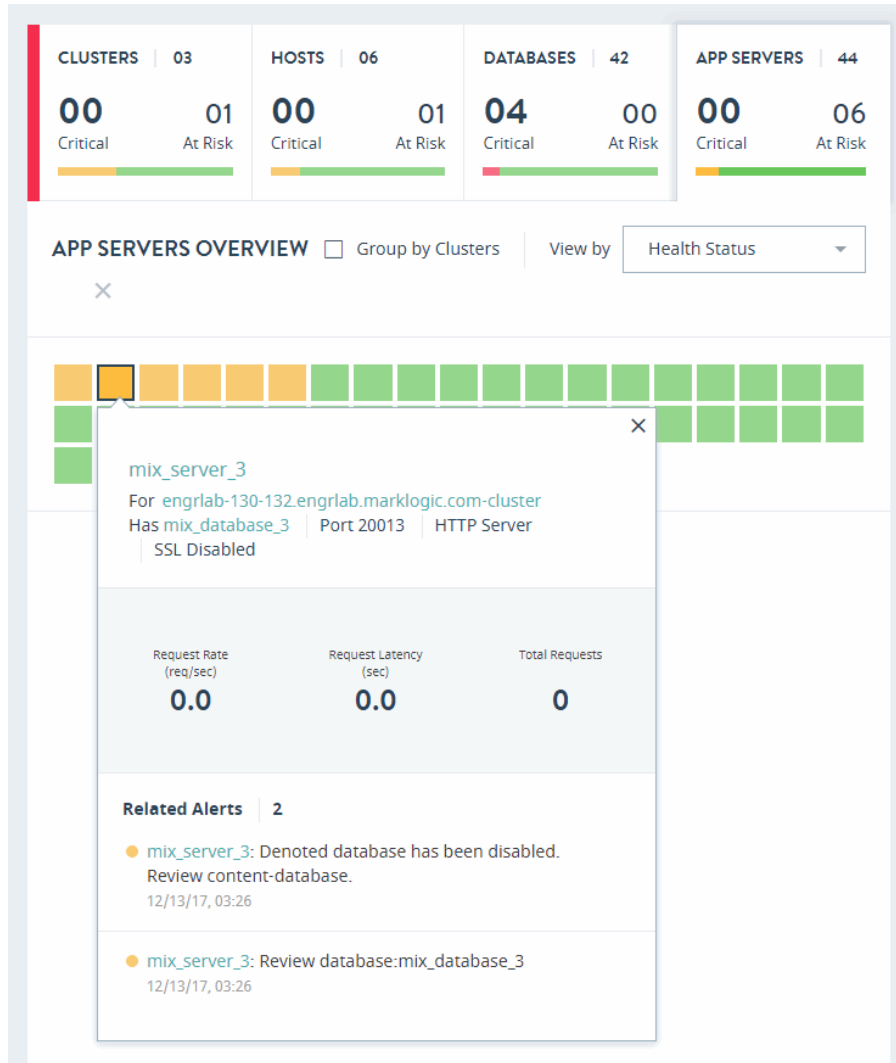


4.2.4 App Servers Overview


Click on the APP SERVERS tab in the key performance indicators panel to open the heat map view for application servers.

The APP SERVERS OVERVIEW page displays colored rectangles, each representing an application server. Administrators can quickly grasp the health status of large number of application servers and comprehend the overall impact at the enterprise or group level.

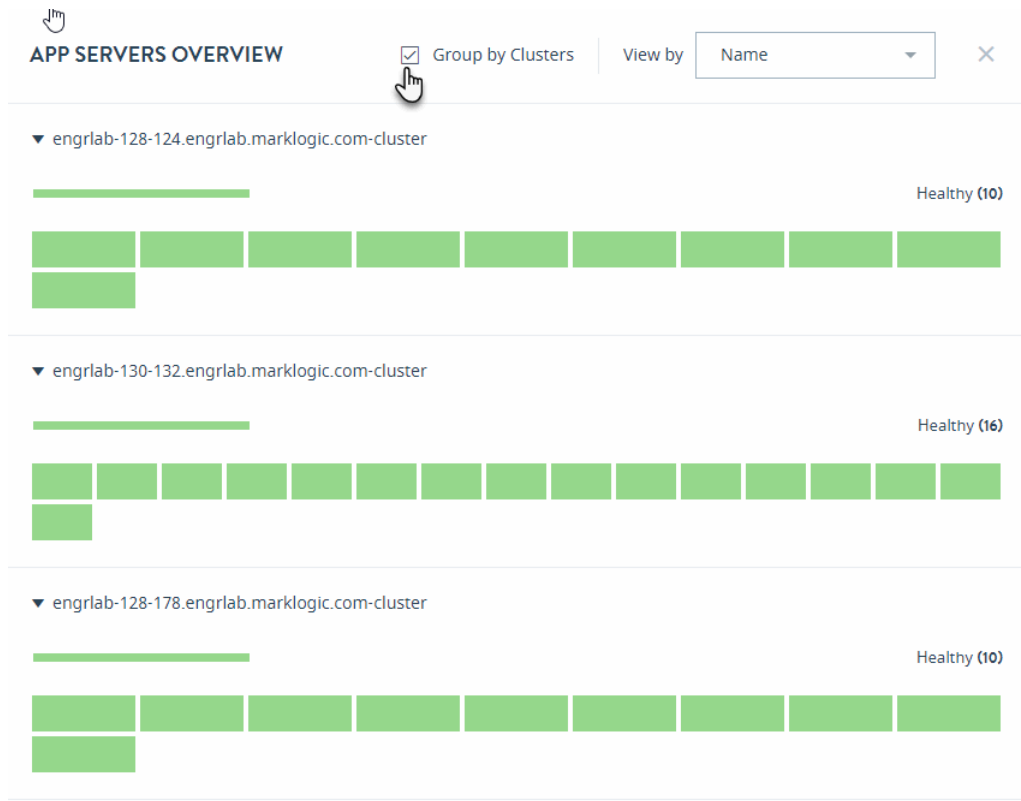
Click on an application server to display the key performance metrics and any alerts for that server. The gauges represent resource consumption with percentage values.




In the App Servers overview page, you can sort the resources either by **Health Status** or **Name**.



Select **Group by Clusters** to sort your view of the application servers by their clusters.



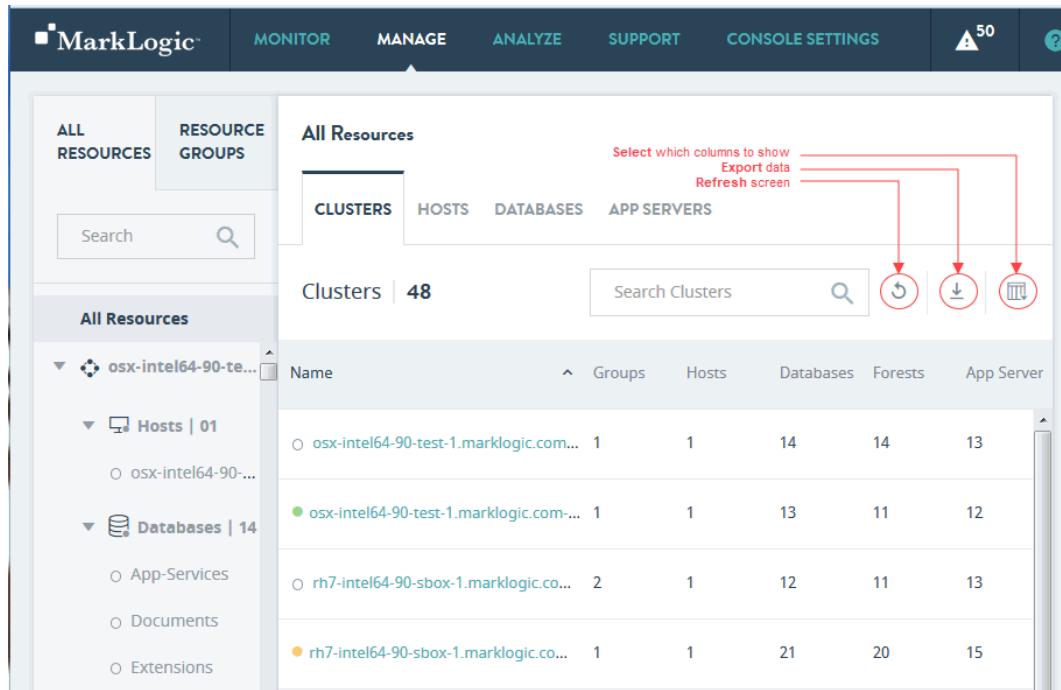
To close the APP SERVERS OVERVIEW page and return to the top-level Monitor view, click the **X** to the right of the “View by” menu.



5.0 MANAGE View

The MANAGE view shows details for a selected MarkLogic resource or resource category in the form of charts, tables, and other displays.

The default MANAGE view displays an inventory list of resources within an enterprise. It displays the list of clusters across an enterprise, along with their key properties, in a data grid.



The panel on the left side is a navigation tree with scrolling lists of all resources displayed according to the object hierarchy. The title of the content area displays the object name selected in the navigation tree. For individual resources, this area displays the resource name and key information related to the resource.

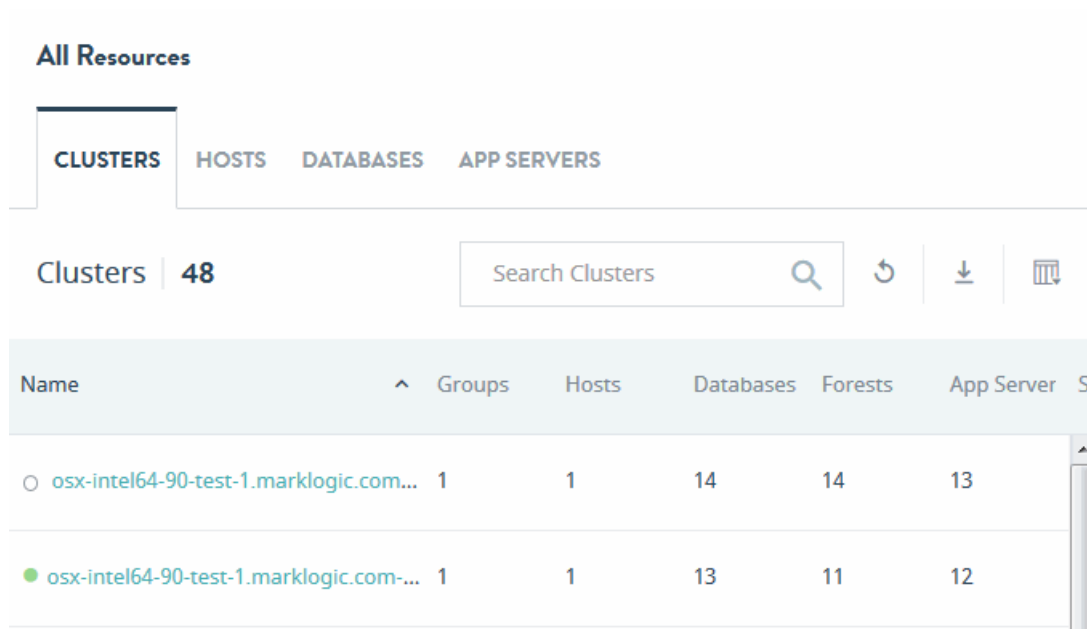
The content area of this view presents content tabs in different categories for the selected object or resource. In the Enterprise case, there are four tabs, one for each resource type: **CLUSTERS**, **HOSTS**, **DATABASES**, and **APP SERVERS**. Each tab presents a table that lists the corresponding resources within the enterprise that has the resource name and status indicator in first column, followed by other properties.

This chapter covers the following topics:

- [Manage Clusters Tab](#)
- [Manage Hosts Tab](#)
- [Manage Databases Tab](#)
- [Manage App Servers Tab](#)

5.1 Manage Clusters Tab

The **CLUSTERS** tab displays the list of clusters in your enterprise.



The columns displayed in the manage **CLUSTERS** tab are described in the following table.

Column	Description
Name	The name of the cluster.
Groups	The number of groups in the cluster.
Hosts	The number of hosts in the cluster.
Databases	The number of databases in the cluster.
Forests	The number of forests in the cluster.
App Servers	The number of App Servers in the cluster.
Server Version	The version of MarkLogic Server running on the cluster's hosts.
OS	The name and version of the operating system on the cluster's hosts.
Uptime	The duration (Days Hrs:Min) the cluster has been available.
Encryption	Specifies whether data encryption is turned on or off in the cluster.

You can export data from the Manage Clusters tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Manage Clusters table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

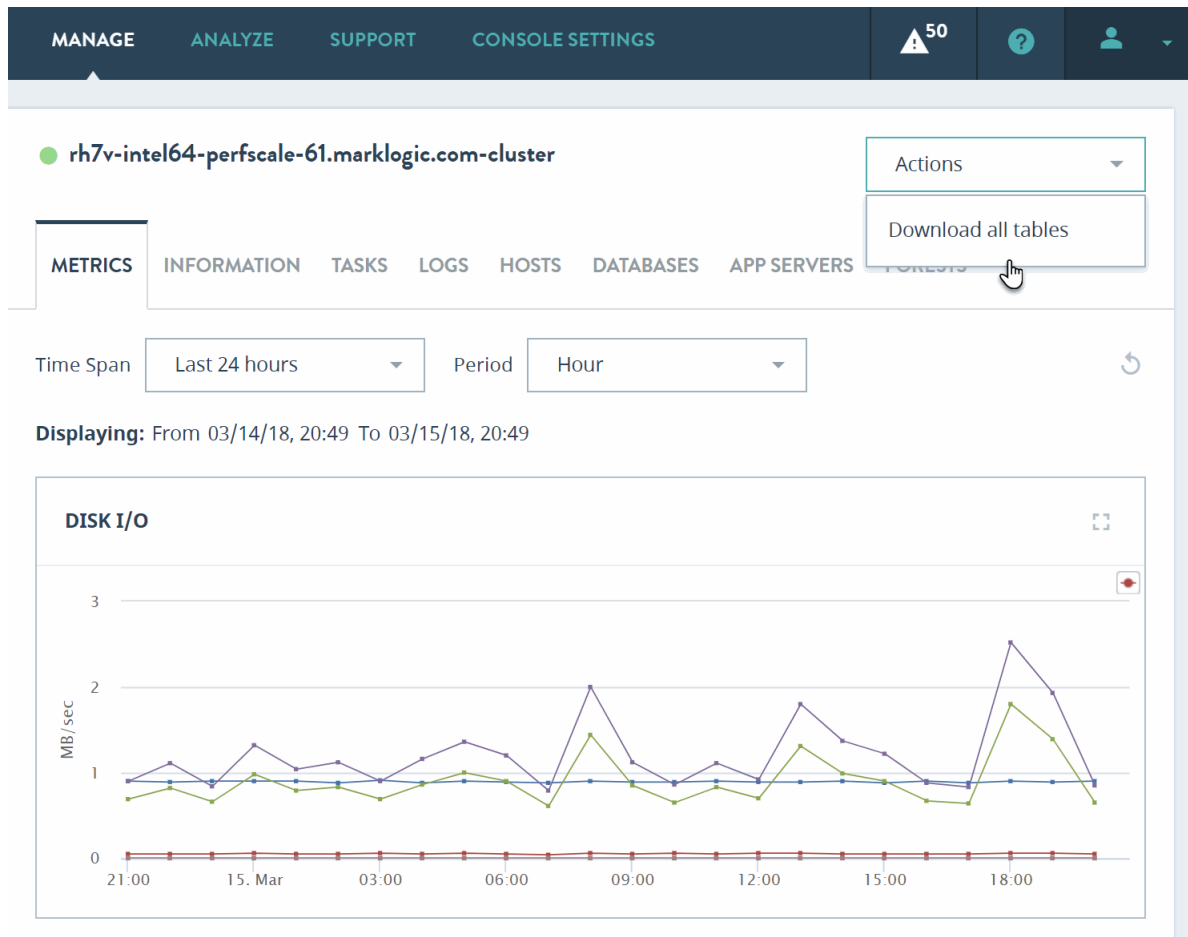
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

To drill down for a particular cluster, click on the cluster name in the table. You will see the following content tabs, each one representing an information category for the selected cluster:

- [Cluster Metrics](#)
- [Cluster Information](#)
- [Cluster Tasks](#)
- [Cluster Logs](#)
- [Cluster Hosts](#)
- [Cluster Databases](#)
- [Cluster App Servers](#)
- [Cluster Forests](#)

You can export all tables from the content tabs for a particular cluster. When you select a specific cluster (either in the left side resource navigation panel or in the content area of the view), **Actions** menu becomes available in the upper right corner.

From the **Actions** menu, select **Download all tables** to export all tables for this particular cluster.



A zip file with all exported tables for this cluster will be downloaded to your computer. Each table is represented by the corresponding CSV file. The zip file will contain the following CSV files for the cluster:

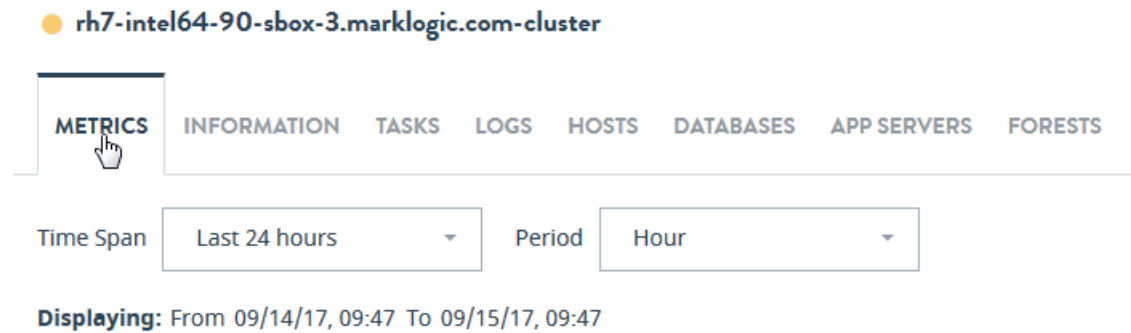
- Tasks CSV file (see the [Cluster Tasks](#) section for details on the contents)
- Hosts CVS file (see the [Cluster Hosts](#) section for details on the contents)
- Databases CVS file (see the [Cluster Databases](#) section for details on the contents)
- App Servers CVS file (see the [Cluster App Servers](#) section for details on the contents)
- Forests CVS file (see the [Cluster Forests](#) section for details on the contents)

Note: For a cluster, exported tables will not include Logs table. This table can be exported from the **Logs** tab of that cluster, but not as part of the **Download all tables** operation.

5.1.1 Cluster Metrics

METRICS is the first tab displayed for any resource type. This tab displays key indicators enabling you to determine the health of the selected resource.

To filter the data used for rendering the graphs, select a pre-defined time period or specify a custom time period, as described in “Date and Time Filters” on page 86.



The metrics displayed by charts on the cluster METRICS tab are described in the following table.

Chart	Definition of Displayed Metric
Disk I/O	Disk I/O in MB per second. For descriptions of the lines displayed on this graph, see “Disk Performance Data” on page 195.
App Server Request Rate	The total number of queries being processed per second, across all of the App Servers in the cluster. For descriptions of the lines displayed on this graph, see “Server Performance Data” on page 202.
App Server Latency	The average time (in seconds) it takes to process queries, across all of the App Servers in the cluster. For descriptions of the lines displayed on this graph, see “Server Performance Data” on page 202.
CPU	The aggregate I/O performance data for the CPUs in the cluster. For descriptions of the lines displayed on this graph, see “CPU Performance Data” on page 197.

Chart	Definition of Displayed Metric
Memory Footprint	<p>The total amount (in MB) of memory consumed by all of the hosts in the cluster.</p> <p>The displayed metrics are:</p> <ul style="list-style-type: none"> • RSS: The total amount of MB of Process Resident Size (RSS) consumed by the cluster. • Anon: The total amount of MB of Process Anonymous Memory consumed by the cluster.
Memory Size	<p>The amount of space (in MB) forest data files in this cluster use in memory.</p>
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 <code>/proc/vmstat</code>) for the cluster in pages per second. • Page-Out Rate: The page-out rate (from Linux <code>/proc/vmstat</code>) for the cluster in pages per second. • Swap-In Rate: The swap-in rate (from Linux <code>/proc/vmstat</code>) for the cluster in pages per second. • Swap-Out Rate: The swap-out rate (from Linux <code>/proc/vmstat</code>) for the cluster in pages per second. <p>For more information, see “Memory Performance Data” on page 200.</p>
Network	<p>Various XDQP (XML Data Query Protocol) performance metrics, such as the sum of XDQP activity across the cluster. XDQP is a MarkLogic internal protocol used for communication between nodes in a cluster. For more detail, see “Network Performance Data” on page 204.</p>

5.1.2 Cluster Information

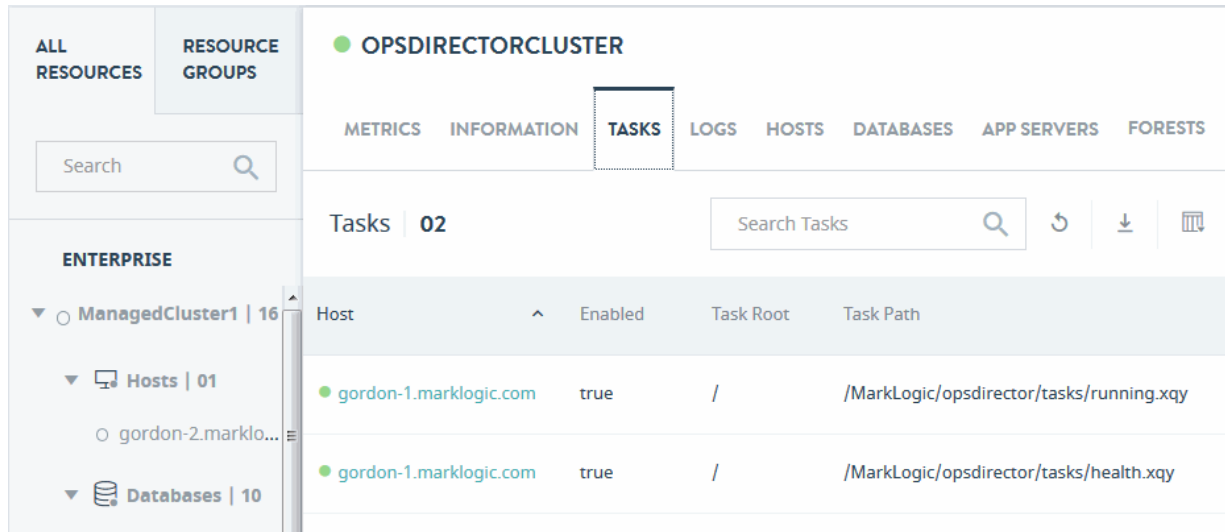
Use the **INFORMATION** tab to access properties of the selected cluster.

The properties displayed in the cluster INFORMATION tab are described in the following table.

Field	Description
Bootstrap Host	<p>The connection attributes for the bootstrap host on this cluster.</p> <ul style="list-style-type: none"> • Bootstrap Connect Port: The bind port for the bootstrap host. • Bootstrap Host Name: The name of the bootstrap host. • Bootstrap Host ID: The ID used to identify the bootstrap host. <p>For details, see Bootstrap Hosts in the <i>Concepts Guide</i>.</p>
Cluster ID	The ID used to identify this cluster.
Cluster Name	The name used to identify this cluster.
Data Directory	The MarkLogic default data directory. For details, see Installing MarkLogic in the <i>Installation Guide</i> .
Effective Version	The effective version of MarkLogic Server that is installed on each host in this cluster. For details, see Effective version and software version in the <i>Administrator's Guide</i> .
Filesystem Directory	The MarkLogic installation directory. For details, see Installing MarkLogic in the <i>Installation Guide</i> .
Role	The role (local or foreign) of the cluster in the current environment. For details, see Coupling Clusters in the <i>Administrator's Guide</i> .
Security Version	The MarkLogic Server version number.
SSL Fips Enabled	Whether FIPS-capable OpenSSL is enabled (true) or disabled (false). For details, see OpenSSL FIPS 140-2 Mode in the <i>Administrator's Guide</i> .
Version	The software version of MarkLogic Server that is installed on each host in this cluster. For details, see Effective version and software version in the <i>Administrator's Guide</i> .
XDQP SSL Certificate	The SSL Certificate used for secure communication between the clusters. For details, see “Ops Director Security” on page 12 and Inter-cluster Communication in the <i>Concepts Guide</i> .

5.1.3 Cluster Tasks

Use the **TASKS** tab for an up-to-date report on tasks across the hosts within the selected cluster. The report shows tasks that are currently running, are scheduled to run, or have been delayed.



The columns displayed in the cluster **TASKS** tab are described in the following table.

Column	Description
Cluster	Cluster on which the task host is located.
Host	The hostname of the host computer on which the scheduled module is to be invoked.
Enabled	Whether the task is enabled (true) or disabled (false).
Task Root	The root directory (filesystem) or URI root (database) that contains the module.
Task Path	The module the task is to invoke.

Column	Description
Task Type	<p>The task type:</p> <ul style="list-style-type: none"> • minutely specifies how many minutes between each invocation of the module. • hourly specifies how many hours and minutes between each invocation of the module. • daily specifies how many days between each invocation of the module and the time of day (in 24:00 notation). • weekly specifies how many weeks between each invocation of the module, check one or more days of the week, and the time of day (in 24:00 notation) for the task to start. • monthly specifies how many months between each invocation of the module, select one day of the month (1-31), and the time of day (in 24:00 notation) for the task to start. • one-time specifies the start date in MM/DD/YYYY notation (for example, 07/29/2009 for July 29, 2009) and time in 24:00 notation.
Task Period	How often the module is to be invoked (every <i>N</i> months, weeks, days, hours, or minutes).
Created on	The date and time when the task was created.
Database	The database to which the scheduled module connects for query execution.
Task Modules	The name of the database in which the scheduled module locates the application code. If set to (filesystem), any files in the specified task root directory are executable (given the proper permissions). If set to a database, any documents in the database whose URI begins with the specified task root directory are executable.
User	The user with permission to invoke the module.
Priority	<p>The priority of the task:</p> <ul style="list-style-type: none"> • normal specifies that the task is queued with normal priority. • higher specifies that the task is queued with higher priority.

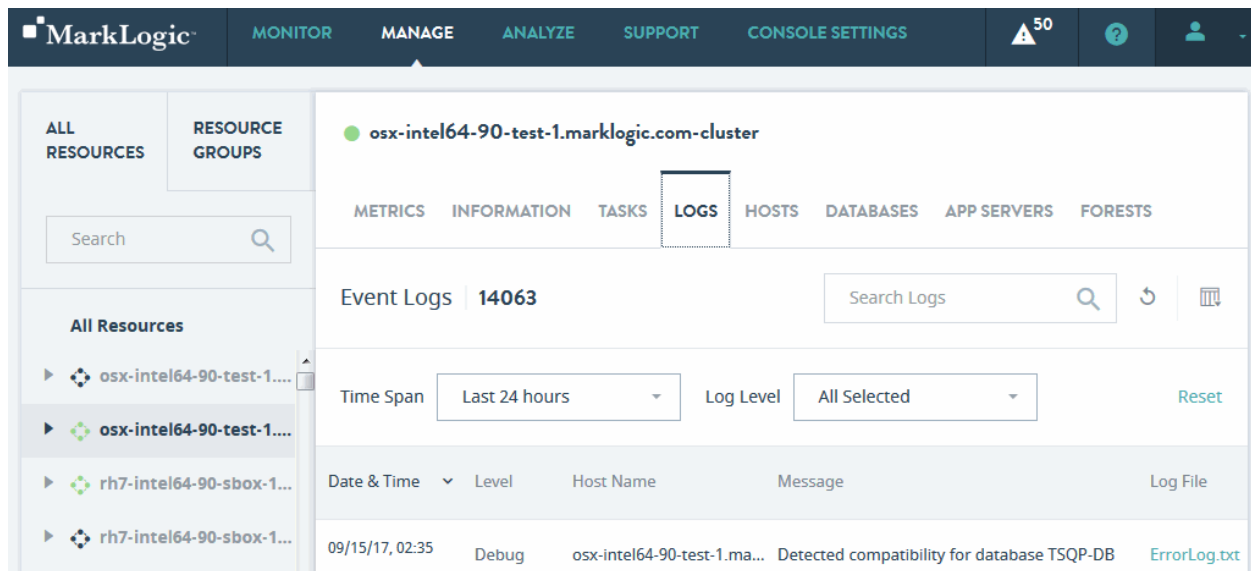
You can export data from the cluster **TASKS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the cluster **TASKS** table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.1.4 Cluster Logs

Use the **LOGS** tab to access a tabular listing of log records that have been logged by various resources within the cluster.



The columns displayed in the cluster **LOGS** tab are described in the following table.

Column	Description
Date & Time	Datetime of the logged event.
Level	The log level of the event. For a description of the log levels, see Understanding the Log Levels in the <i>Administrator's Guide</i> .

Column	Description
Cluster Name	The name of the cluster on which the logged event occurred.
Host Name	The name of the host on which the logged event occurred.
Message	The logged error message.
Log File	The full text of the logged event.

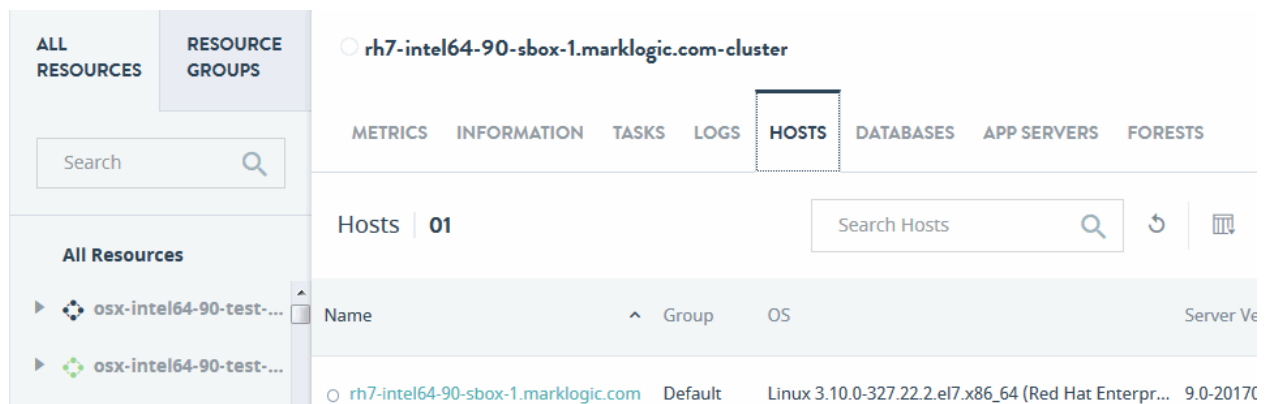
You can export data from the cluster **LOGS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Cluster Logs table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.1.5 Cluster Hosts

Use the **HOSTS** tab to get the list of all hosts within the selected cluster. This tab shows host name as the first column, followed by other properties. Click the host name to navigate to the manage view of the selected host.



The columns displayed in the cluster **HOSTS** tab are described in the following table.

Column	Description
Name	The hostname of each host.
Group	The name of the group that contains each host.
OS	The name and version of the operating system on which each host runs.
Server Version	The version of MarkLogic Server on each host.
Forests	The number of forests on each host.
Databases	The number of databases on each host.
App Servers	The number of App Servers on each host.
Disk Space (MB)	The amount of disk space (in MB) used on each host.
Uptime	The duration (Days Hrs:Min) each host has been available.
Maint. Mode	The host maintenance mode (normal or maintenance) for each host. For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .
Zone	The Amazon Web Services (AWS) zone in which each host resides, if applicable. For details, see <i>MarkLogic Server on Amazon Web Services (AWS) Guide</i> .

You can export data from the Cluster **HOSTS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Cluster Hosts table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.1.6 Cluster Databases

Use the **DATABASES** tab to get a list of all the databases within the cluster, with database name as the first column followed by other properties. Click on the database name to navigate to the manage view of the selected database.

The screenshot shows the MarkLogic MANAGE View interface. The top navigation bar includes 'MONITOR', 'MANAGE', 'ANALYZE', 'SUPPORT', and 'CONSOLE SETTINGS'. The main content area is titled 'rh7-intel64-90-sbox-1.marklogic.com-cluster' and has tabs for 'METRICS', 'INFORMATION', 'TASKS', 'LOGS', 'HOSTS', 'DATABASES', 'APP SERVERS', and 'FORESTS'. The 'DATABASES' tab is active, showing a table with 12 databases. The table columns are: Name, Forests, Disk Size (MB), Documents, Last Backup, Encryption, and HA. Two databases are visible: 'App-Services' (1 forest, 16 MB, 1 document, off encryption, no HA) and 'Documents' (1 forest, 15 MB, 0 documents, off encryption, no HA).

The columns displayed in the cluster **DATABASES** tab are described in the following table.

Column	Description
Name	Name of the database.
Forests	The number of forests used by the database.
Disk Size (MB)	The amount of disk space used by the database forests, in megabytes.
Documents	The number of documents in the database.
Last Backup	The date and time of the last backup of the database. No value, if the database has never been backed up. For details on backing up a database, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Encryption	Specifies whether or not encryption at rest is enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .
HA	Specifies whether or not shared disk failover is enabled. For details, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled (On/Off). For details, see the <i>Database Replication Guide</i> .
Replication Status	Specifies whether or not database replication is configured for the database.
Security DB	The name of the security database used by the database. For details, see Administering Security in the <i>Security Guide</i> .

Column	Description
Schemas DB	The name of the schemas database used by the database. For details, see Understanding and Defining Schemas in the <i>Administrator's Guide</i> .
Triggers DB	The name of the schemas database used by the database. For details, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .

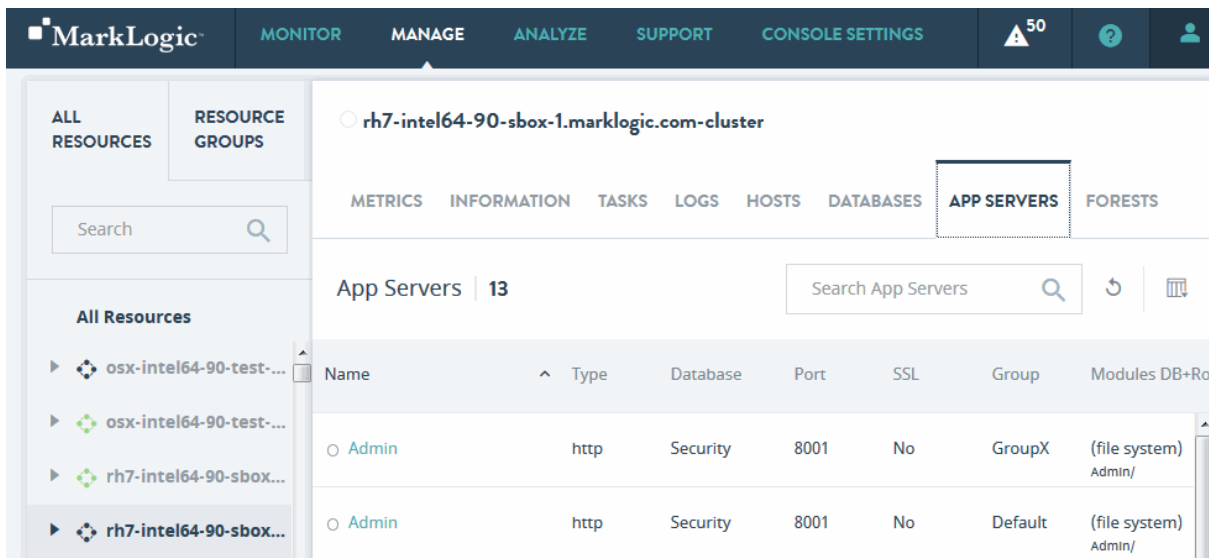
You can export data from the cluster **DATABASES** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Cluster Databases table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. in the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.1.7 Cluster App Servers

Use the **APP SERVERS** tab to get a list of all the application servers within the cluster, with App Server name as the first column followed by other properties. Click on the App Server name to navigate to the manage view of the selected application server.



The columns displayed in the cluster **APP SERVERS** tab are described in the following table.

Column	Description
Name	The name of the App Server.
Type	The App Server type (HTTP, WebDAV, XDBC, ODBC).
Database	The name of the App Server content database.
Port	The port number used to access the App Server.
SSL	Whether the App Server has SSL enabled (yes) or disabled (no). For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Group	The name of the group to which the App Server belongs
Modules DB+Root	The name of the modules database, or if filesystem, the root directory.
Security	The type of security (internal or external).

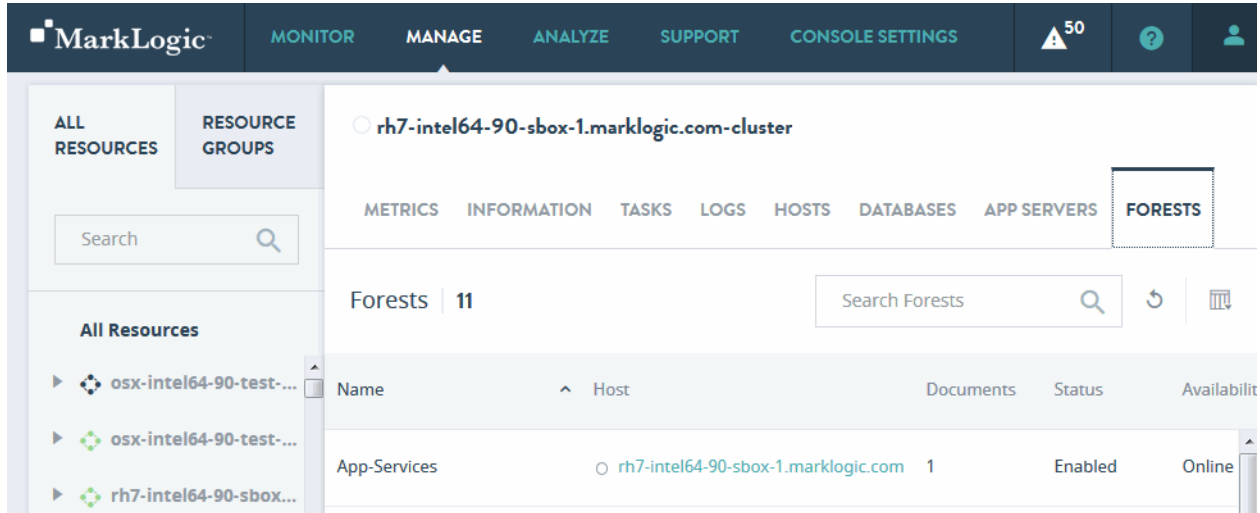
You can export data from the cluster **APP SERVERS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Cluster App Servers table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.1.8 Cluster Forests

Use the **FORESTS** tab to get a list of all the forests within the cluster, with forest name as the first column followed by other properties. Click on the forest name to navigate to the manage view of the selected forest.



The columns displayed in the cluster **FORESTS** tab are described in the following table.

Column	Description
Name	The name of the forest.
Host	The forest host.
Documents	The number of documents in the forest.
Status	The status of the forest. Enabled or Disabled.
Availability	The availability of the forest. Online or Offline.
Fragments	The number of active fragments (the fragments available to queries) in the forest.
Deleted Fragments	The number of deleted fragments (the fragments to be removed by the next merge operation) in the forest.
Stands	The number of stands in the forest. For more information on stands, see Databases, Forests, and Stands in the <i>Concepts Guide</i> .
Size (MB)	The size of the forest, in MB.
Encrypted Size (MB)	The amount of encrypted data in the forest. For details on data encryption, see Encryption at Rest in the <i>Security Guide</i> .
Free Space (MB)	The number of MB of free space on this forest.

Column	Description
Large Data Size (MB)	The amount of data in the large data directories of the forest. For more information on Large Data, see Working With Binary Documents in the <i>Application Developer's Guide</i> .
Fast Data Size (MB)	The amount of data in the fast data directories of the forest. For more information on Fast Data, see Fast Data Directory on Forests in the <i>Query Performance and Tuning Guide</i> .
Failover Enabled	Whether failover is enabled for the forest. For more information on stands, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled for this forest. For details, see the <i>Database Replication Guide</i> .

You can export data from the cluster **FORESTS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

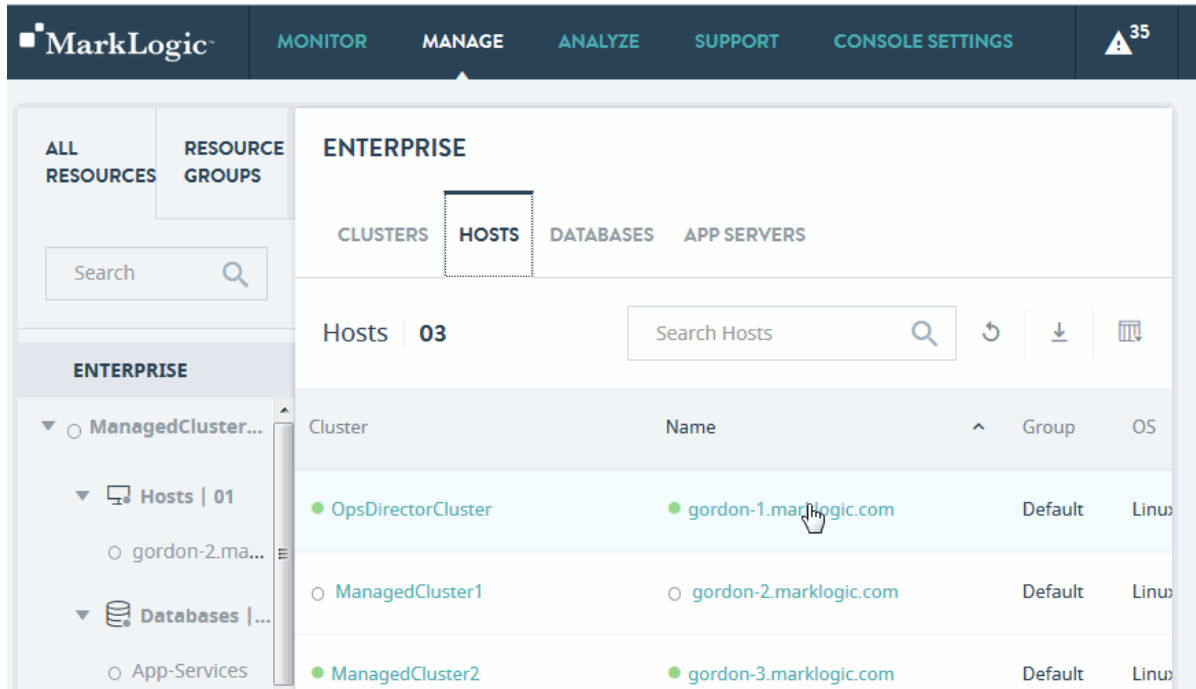
- The resulting CSV file will have the same columns as the cluster FORESTS table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.2 Manage Hosts Tab

The **HOSTS** tab displays the list of hosts in your enterprise.

Note: You can only see the hosts to which you have access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the manage **HOSTS** tab are described in the following table.

Column	Description
Name	The hostname of the host.
Cluster	The name of the cluster on which the host resides.
Group	The name of the group that contains the host.
OS	The name and version of the operating system on which the host runs.
Server Version	The version of MarkLogic Server on the host.
Forests	The number of forests on the host.
Databases	The number of databases on the host.
App Servers	The number of App Servers on the host.
Disk Space (MB)	The amount of disk space (in MB) used on the host.
Uptime	The duration (Days Hrs:Min) the host has been available.
Maint. Mode	The host maintenance mode (normal or maintenance). For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .

Column	Description
Zone	The Amazon Web Services (AWS) zone in which the host resides, if applicable. For details, see <i>MarkLogic Server on Amazon Web Services (AWS) Guide</i> .

You can export data from the manage **HOSTS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Manage Hosts table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

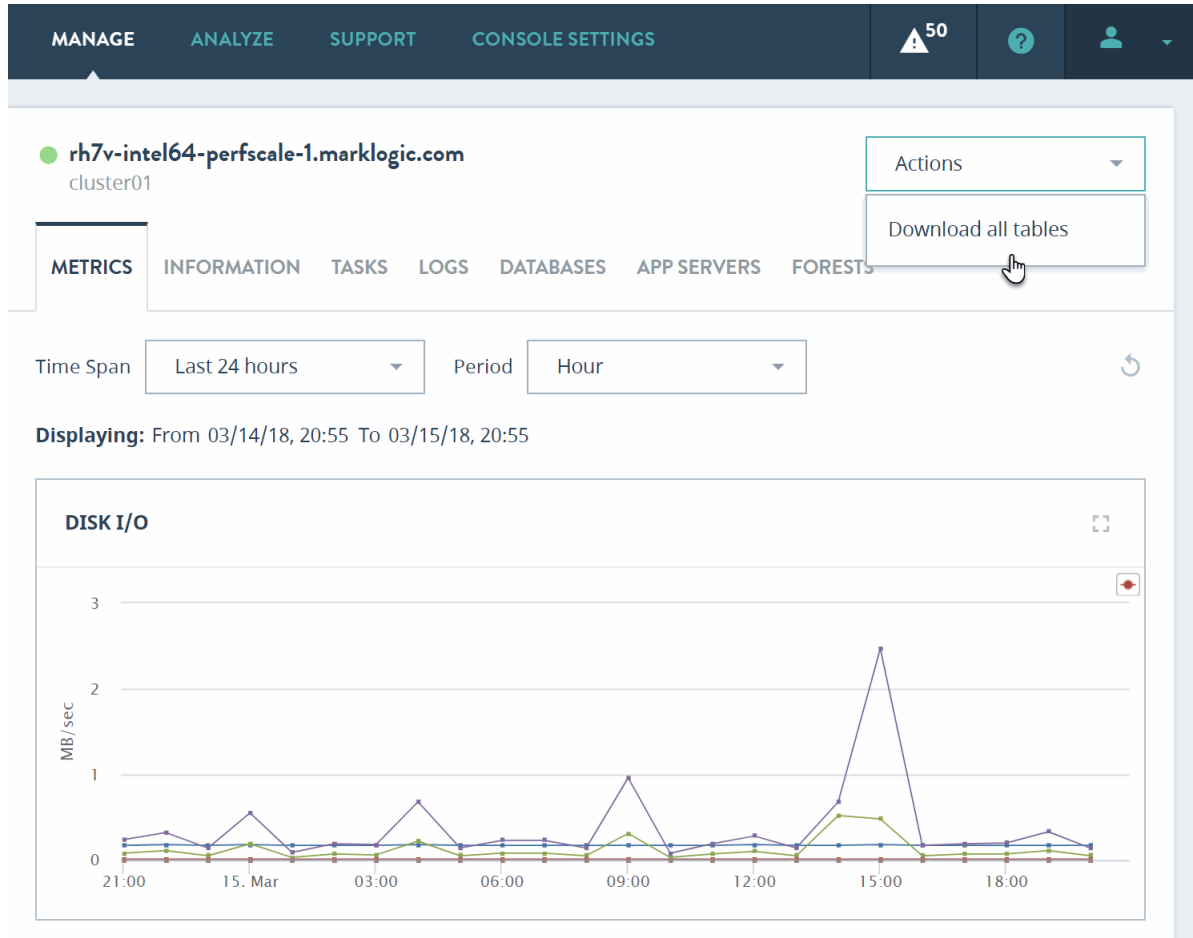
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

To drill down for a particular host, click on the host name in the host table. You will see the following content tabs, each one representing an information category for the selected host:

- [Host Metrics](#)
- [Host Information](#)
- [Host Tasks](#)
- [Host Logs](#)
- [Host Databases](#)
- [Host App Servers](#)
- [Host Forests](#)

You can export all tables from the content tabs for a particular host. When you select a specific host (either in the left side resource navigation panel or in the content area of the view), the **Actions** menu becomes available in the upper right corner.

From the **Actions** menu, select **Download all tables** to export all tables for this particular host.



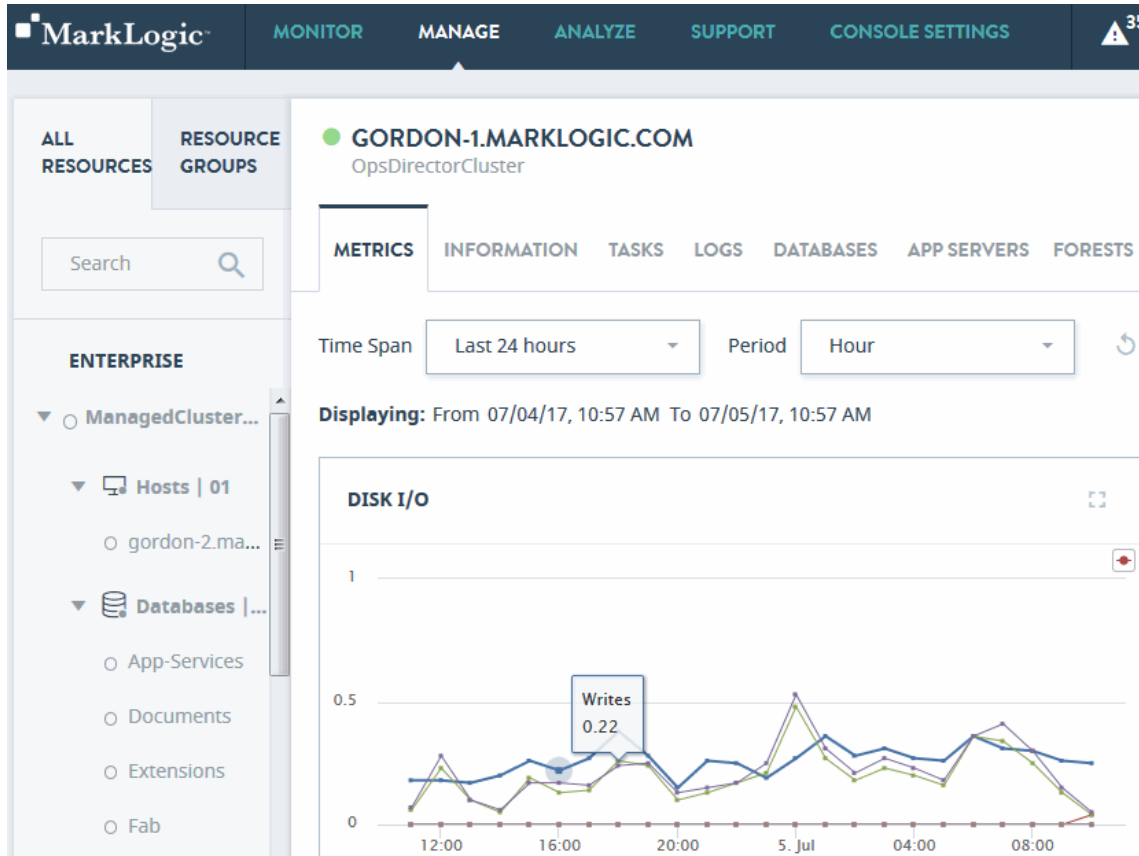
A zip file with all exported tables for this host will be downloaded to your computer. Each table is represented by a corresponding CSV file. The zip file will contain the following CSV files for the host:

- Tasks CSV file (see the [Host Tasks](#) section for details on the contents)
- Databases CVS file (see the [Host Databases](#) section for details on the contents)
- App Servers CVS file (see the [Host App Servers](#) section for details on the contents)
- Forests CVS file (see the [Host Forests](#) section for details on the contents)

Note: For a host, exported tables will not include Logs table. This table can be exported from the **LOGS** tab of that host, but not as part of the **Download all tables** operation.

5.2.1 Host Metrics

Select a single host from the navigation tree. The **METRICS** tab displays key indicators allowing administrators to determine the health of the selected resource.



To filter the data used for rendering the graphs, select a pre-defined time period or specify a custom time period, as described in “Date and Time Filters” on page 86.

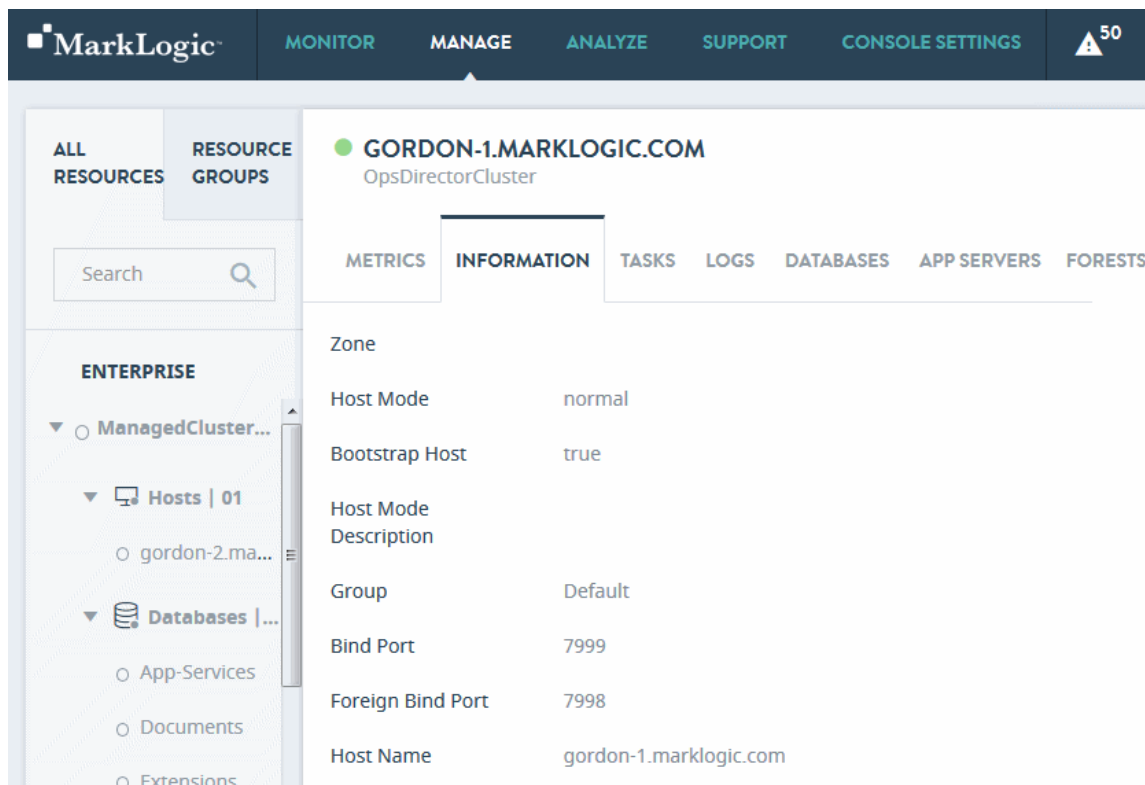
The metrics displayed by charts on the host **METRICS** tab are described in the following table.

Chart	Definition of Displayed Metric
Disk I/O	Disk I/O in MB/sec. For details, see “Disk Performance Data” on page 195.
CPU	The aggregate I/O performance data for the CPUs in the host. For details, see “CPU Performance Data” on page 197.

Chart	Definition of Displayed Metric
Memory Footprint	<p>The total amount (in MB) of memory consumed by this host.</p> <p>The displayed metrics are:</p> <ul style="list-style-type: none"> • RSS: The total amount of MB of Process Resident Size (RSS) consumed by the host. • Anon: The total amount of MB of Process Anonymous Memory consumed by the host.
Memory Size	<p>The amount of space (in MB) forest data files for this host take up in memory.</p>
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 host in pages per second. • Page-Out Rate: The page-out rate (from Linux /proc/vmstat) for the host in pages per second. • Swap-In Rate: The swap-in rate (from Linux /proc/vmstat) for the host in pages per second. • Swap-Out Rate: The swap-out rate (from Linux /proc/vmstat) for the host in pages per second.
Network	<p>Various XDQP performance metrics as the sum of XDQP activity for this host. For details, see “Network Performance Data” on page 204.</p>

5.2.2 Host Information

Use the **INFORMATION** tab to access properties of the selected host.



The properties displayed in the host **INFORMATION** tab are described in the following table.

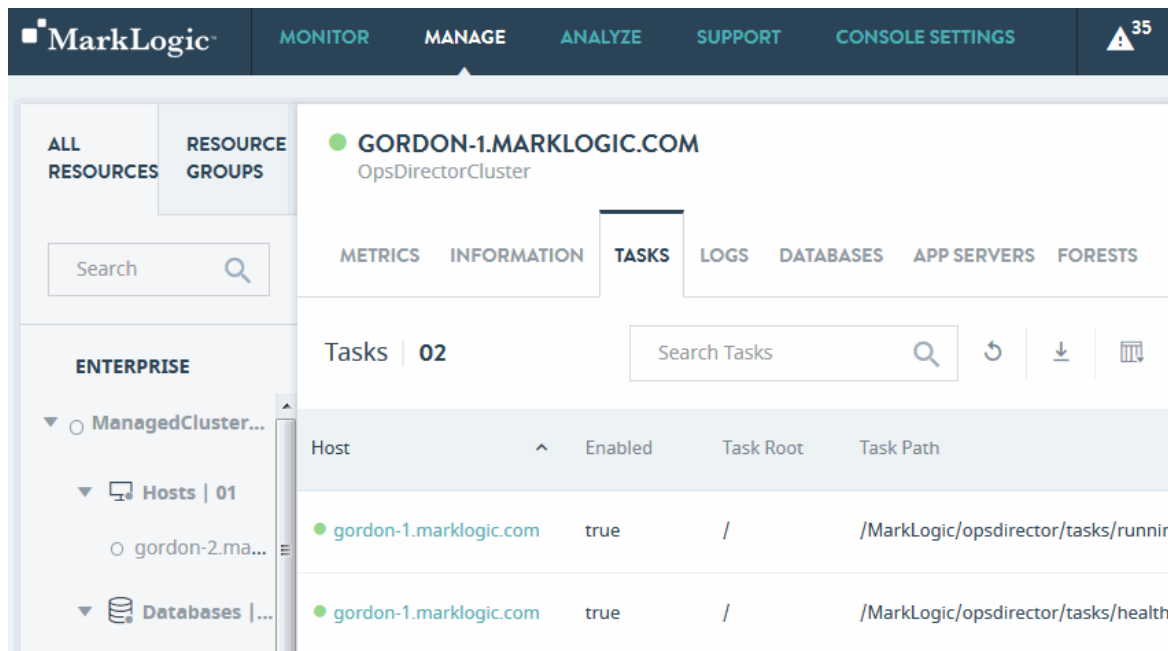
Field	Description
Bind Port	The port on which the host to handle inter-host communication within this cluster. For details, see Inter-cluster Communication in the <i>Concepts Guide</i> .
Bootstrap Host	Whether this host is a bootstrap host (true). If not, the value is false. For details, see Bootstrap Hosts in the <i>Concepts Guide</i> .
Foreign Bind Port	The port used by the host to handle XDQP communication with foreign clusters. For details, see Bootstrap Hosts in the <i>Concepts Guide</i> .
Group	The group to which this host belongs. For details on groups, see Groups in the <i>Administrator's Guide</i> .
Host Mode	The host maintenance mode (normal or maintenance). For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .

Field	Description
Host Mode Description	Description of the host mode, if set.
Host Name	The name of the host.
Zone	The Amazon Web Services (AWS) zone in which the host resides, if applicable. For details, see <i>MarkLogic Server on Amazon Web Services (AWS) Guide</i> .

5.2.3 Host Tasks

Use the **TASKS** tab for an up-to-date report on tasks that are currently running or scheduled to run on this host.

Note: You can only see the tasks for the hosts to which you have access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the Host **TASKS** tab are described in the following table.

Column	Description
Cluster	Cluster on which the task host is located.

Column	Description
Host	The hostname of the host computer on which the scheduled module is to be invoked.
Enabled	Whether the task is enabled (true) or disabled (false).
Task Root	The root directory (filesystem) or URI root (database) that contains the module.
Task Path	The module the task is to invoke.
Task Type	<p>The task type:</p> <ul style="list-style-type: none"> • minutely specifies how many minutes between each invocation of the module. • hourly specifies how many hours and minutes between each invocation of the module. • daily specifies how many days between each invocation of the module and the time of day (in 24:00 notation). • weekly specifies how many weeks between each invocation of the module, check one or more days of the week, and the time of day (in 24:00 notation) for the task to start. • monthly specifies how many months between each invocation of the module, select one day of the month (1-31), and the time of day (in 24:00 notation) for the task to start. • one-time specifies the start date in MM/DD/YYYY notation (for example, 07/29/2009 for July 29, 2009) and time in 24:00 notation.
Task Period	How often the module is to be invoked (every <i>N</i> months, weeks, days, hours, or minutes).
Created on	The datetime the task was created.
Database	The database to which the scheduled module connects for query execution.
Task Modules	The name of the database in which the scheduled module locates the application code. If set to (filesystem), any files in the specified task root directory are executable (given the proper permissions). If set to a database, any documents in the database whose URI begins with the specified task root directory are executable.
User	The user with permission to invoke the module.

Column	Description
Priority	<p>The priority of the task:</p> <ul style="list-style-type: none"> • normal specifies the task is queued with normal priority. • higher specifies the task is queued with higher priority.

You can export data from the Host **TASKS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

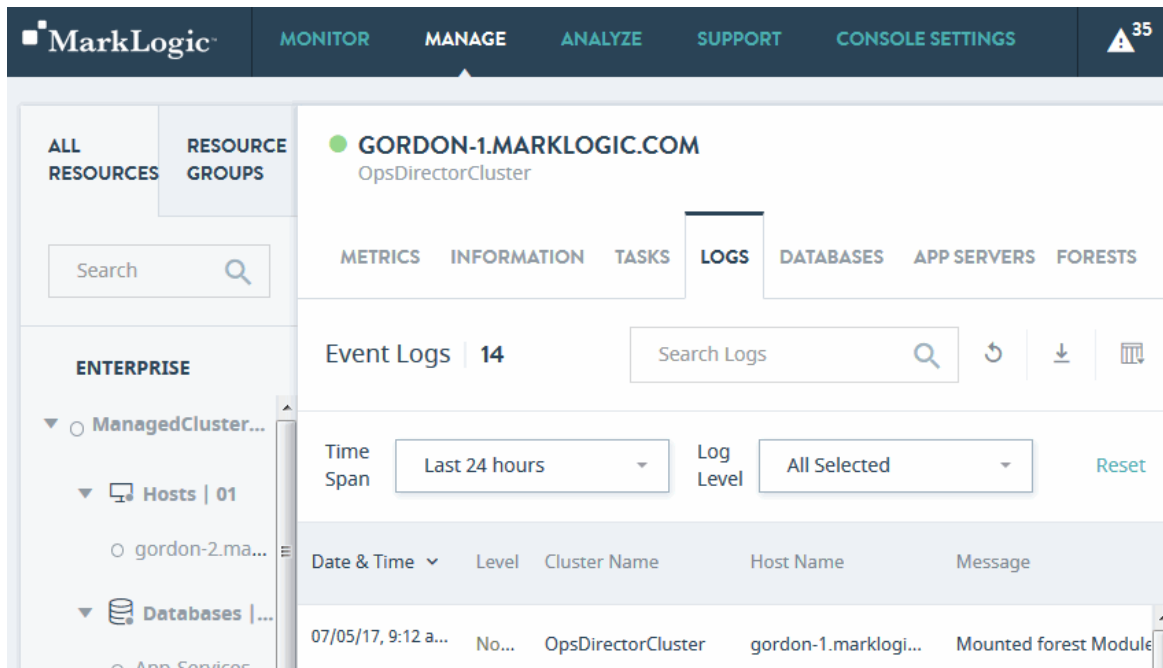
- The resulting CSV file will have the same columns as the Host Tasks table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.2.4 Host Logs

Use the **LOGS** tab to access a tabular listing of log records of the host.

Note: You can only see the logs from the hosts to which you have access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the host **LOGS** tab are described in the following table.

Column	Description
Date & Time	Datetime of the logged event.
Level	The log level of the event. For a description of the log levels, see Understanding the Log Levels in the <i>Administrator's Guide</i> .
Cluster Name	The name of the cluster on which the logged event occurred.
Host Name	The name of the host on which the logged event occurred.
Message	The logged error message.
Log File	The full text of the logged event.

You can export data from the host **LOGS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Host Logs table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.2.5 Host Databases

Use the **DATABASES** tab to get a list of all the databases within the host, with database name as the first column followed by other properties. Click on the database name to navigate to the **METRICS** view of the selected database.

Note: Access to databases of a host is implicit if you have access to that host. For access rules, see “Access Inheritance in Resource Groups” on page 14.

The screenshot shows the MarkLogic MANAGE View for a host named GORDON-1.MARKLOGIC.COM. The 'DATABASES' tab is selected, displaying a table of databases. The table has the following columns: Name, Forests, Disk Size (MB), Documents, Last Backup, and Replication. The data rows are:

Name	Forests	Disk Size (MB)	Documents	Last Backup	Replication
App-Services	1	1	256		N/A
Documents	1	10	10		N/A
Extensions	1		0		N/A

The columns displayed in the host **DATABASES** tab are described in the following table.

Column	Description
Name	Name of the database.
Forests	The number of forests used by the database.
Disk Size (MB)	The amount of disk space used by the database forests, in megabytes.
Documents	The number of documents in the database.
Last Backup	The data-time of the last backup of the database. No value, if the database has never been backed up. For details on backing up a database, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Encryption	Specifies whether or not encryption at rest is enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .
HA	Specifies whether or not shared disk failover is enabled. For details, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled. For details, see the <i>Database Replication Guide</i> .

Column	Description
Replication Status	Specifies whether or not database replication is configured for the database.
Security DB	The name of the security database used by the database. For details, see Administering Security in the <i>Security Guide</i> .
Schemas DB	The name of the schemas database used by the database. For details, see Understanding and Defining Schemas in the <i>Administrator's Guide</i> .
Triggers DB	The name of the schemas database used by the database. For details, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .

You can export data from the host **DATABASES** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Host Databases table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.2.6 Host App Servers

Use the **APP SERVERS** tab to get the list of all the application servers within the host, with App Server name as the first column followed by other properties. Click on the App Server name to navigate to the manage view of the selected application server.

Note: Access to App Servers is not implicit with host access. For access rules, see “Access Inheritance in Resource Groups” on page 14.

The screenshot shows the MarkLogic MANAGE View for the host GORDON-1.MARKLOGIC.COM. The 'APP SERVERS' tab is active, displaying a table with the following data:

Name	Type	Database	Port	SSL	Group
Admin	http	Security	8001	No	Default
App-Services	http	Documents	8000	No	Default

The columns displayed in the host **APP SERVERS** tab are described in the following table.

Column	Description
Name	The name of the App Server.
Type	The App Server type (HTTP, WebDAV, XDBC, ODBC).
Database	The name of the App Server content database.
Port	The port number used to access the App Server.
SSL	Whether the App Server has SSL enabled (yes) or disabled (no). For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Group	The name of the group to which the App Server belongs
Modules DB+Root	The name of the modules database, or if filesystem, the root directory.
Security	The type of security (internal or external).

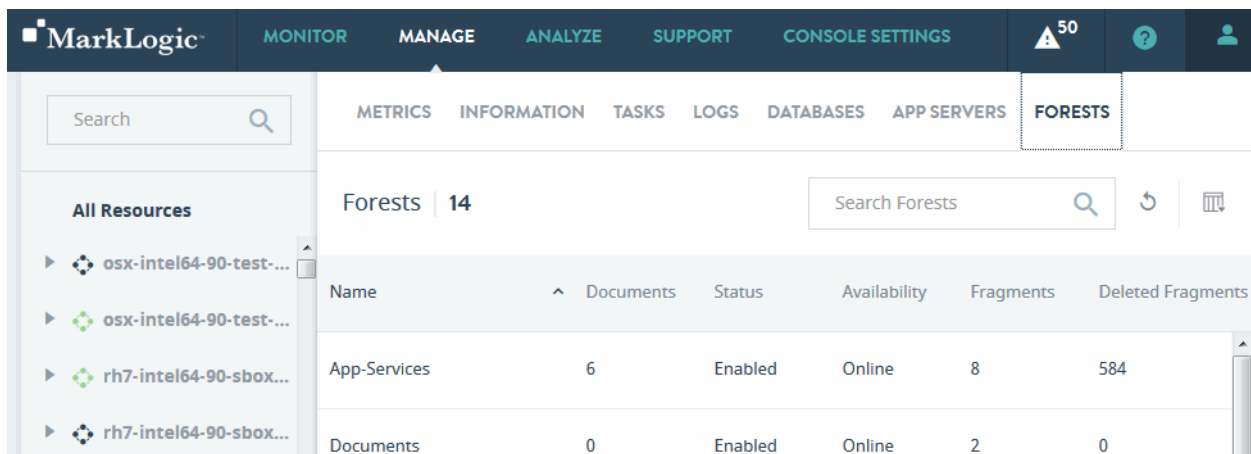
You can export data from the host **APP SERVERS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Host App Servers table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.2.7 Host Forests

Use the **FORESTS** tab to get a list of all the forests within the host, with forest name as the first column, followed by other properties. Click on the forest name to navigate to the manage view of the selected forest.



The columns displayed in the host **FORESTS** tab are described in the following table.

Column	Description
Name	The name of the forest.
Documents	The number of documents in the forest.
Status	The status of the forest. Enabled or Disabled.
Availability	The availability of the forest. Online or Offline.

Column	Description
Fragments	The number of active fragments (the fragments available to queries) in the forest.
Deleted Fragments	The number of deleted fragments (the fragments to be removed by the next merge operation) in the forest.
Stands	The number of stands in the forest. For more information on stands, see Databases, Forests, and Stands in the <i>Concepts Guide</i> .
Size (MB)	The size of the forest, in MB.
Encrypted Size (MB)	The amount of encrypted data in the forest. For details on data encryption, see Encryption at Rest in the <i>Security Guide</i> .
Free Space (MB)	The number of MB of free space on this forest.
Large Data Size (MB)	The amount of data in the large data directories of the forest. For more information on Large Data, see Working With Binary Documents in the <i>Application Developer's Guide</i> .
Fast Data Size (MB)	The amount of data in the fast data directories of the forest. For more information on Fast Data, see Fast Data Directory on Forests in the <i>Query Performance and Tuning Guide</i> .
Failover Enabled	Whether failover is enabled for the forest. For more information on stands, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled for this forest. For details, see the <i>Database Replication Guide</i> .

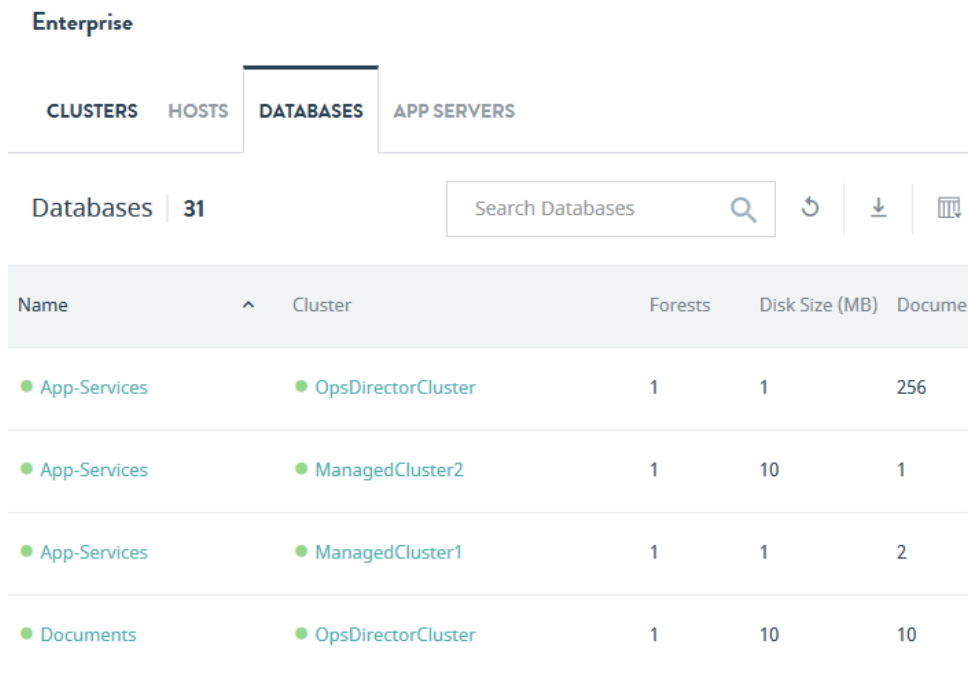
You can export data from the host **FORESTS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the host FORESTS table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.3 Manage Databases Tab

The **DATABASES** tab displays a list of databases in your enterprise.



The columns displayed in the manage **DATABASES** tab are described in the following table.

Column	Description
Name	Name of the database.
Cluster	Cluster on which the database resides.
Forests	The number of forests used by the database.
Disk Size (MB)	The amount of disk space used by the database forests, in megabytes.
Documents	The number of documents in the database.
Last Backup	The data-time of the last backup of the database. No value, if the database has never been backed up. For details on backing up a database, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Encryption	Specifies whether or not encryption at rest is enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .

Column	Description
HA	Specifies whether or not shared disk failover is enabled. For details, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled. For details, see the <i>Database Replication Guide</i> .
Security DB	The name of the security database used by the database. For details, see Administering Security in the <i>Security Guide</i> .
Schemas DB	The name of the schemas database used by the database. For details, see Understanding and Defining Schemas in the <i>Administrator's Guide</i> .
Triggers DB	The name of the schemas database used by the database. For details, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .

You can export data from the Manage Databases tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Manage Databases table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

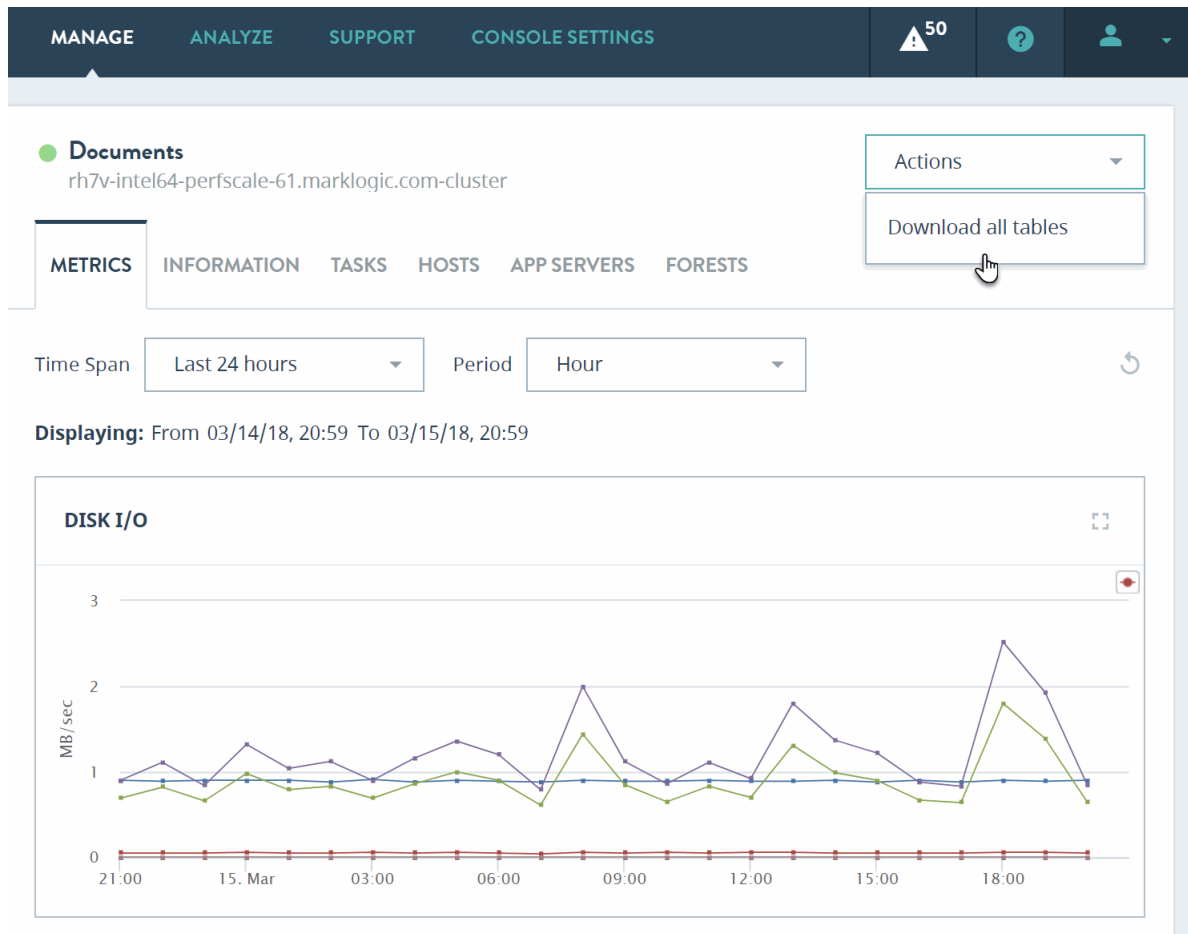
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

To drill down for a particular database, click on the database name in the database table. You will see the following content choices, each one representing an information category for the selected database:

- [Database Metrics](#)
- [Database Information](#)
- [Database Tasks](#)
- [Database Hosts](#)
- [Database App Servers](#)
- [Database Forests](#)

You can export all tables from the database content tabs for a particular database. When you select a specific database (either in the left resource navigation panel or in the content area of the view), the **Actions** menu becomes available in the upper right corner.

From the **Actions** menu, select **Download all tables** to export all tables for this particular database.

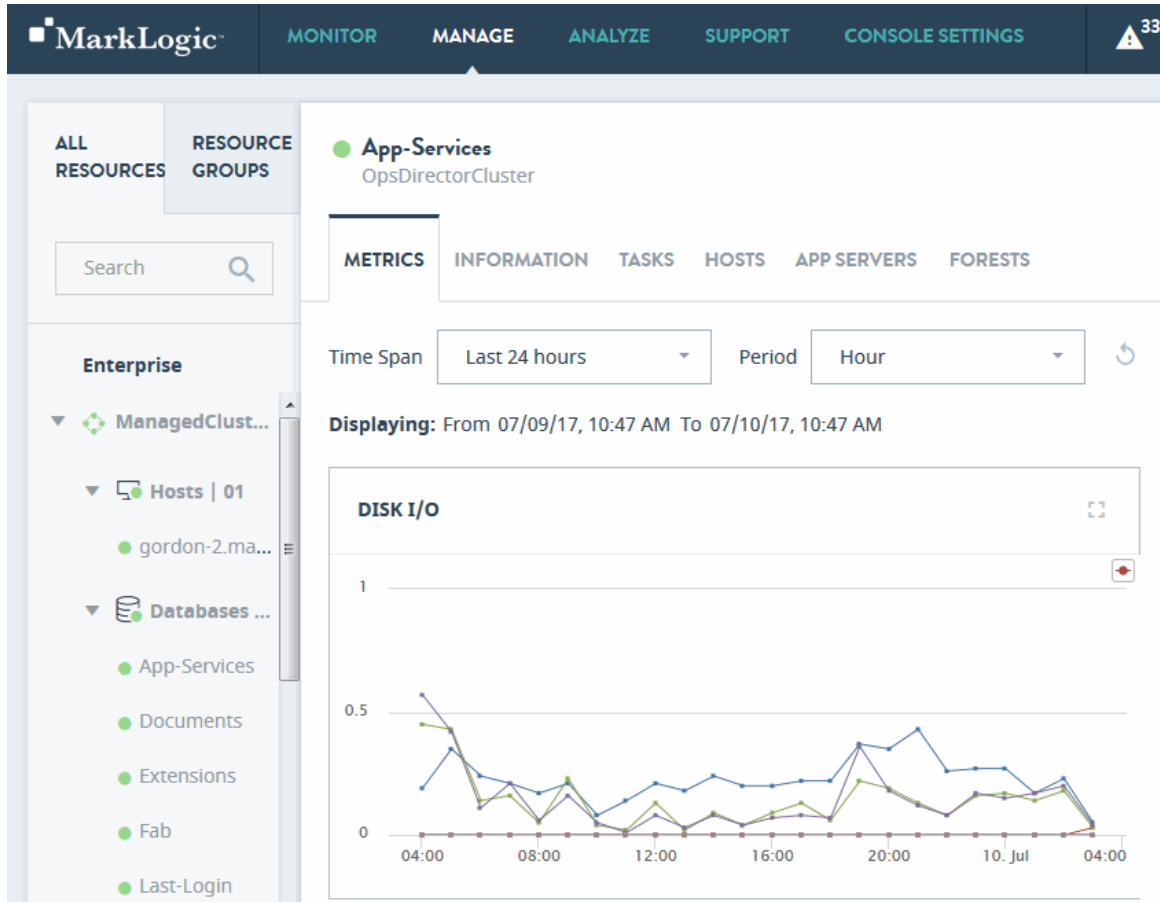


A zip file with all exported tables for this database will be downloaded to your computer. Each table is represented by the corresponding CSV file. The zip file will contain the following CSV files for the database:

- Tasks CSV file (see the [Database Tasks](#) section for details on the contents)
- Hosts CVS file (see the [Database Hosts](#) section for details on the contents)
- App Servers CVS file (see the [Database App Servers](#) section for details on the contents)
- Forests CVS file (see the [Database Forests](#) section for details on the contents)

5.3.1 Database Metrics

Select a single database or all databases from the navigation tree. The **METRICS** tab displays key indicators enabling you to determine the health of the selected resource or resource group.



To filter the data used for rendering the graphs, select a pre-defined time period or specify a custom time period, as described in “Date and Time Filters” on page 86.

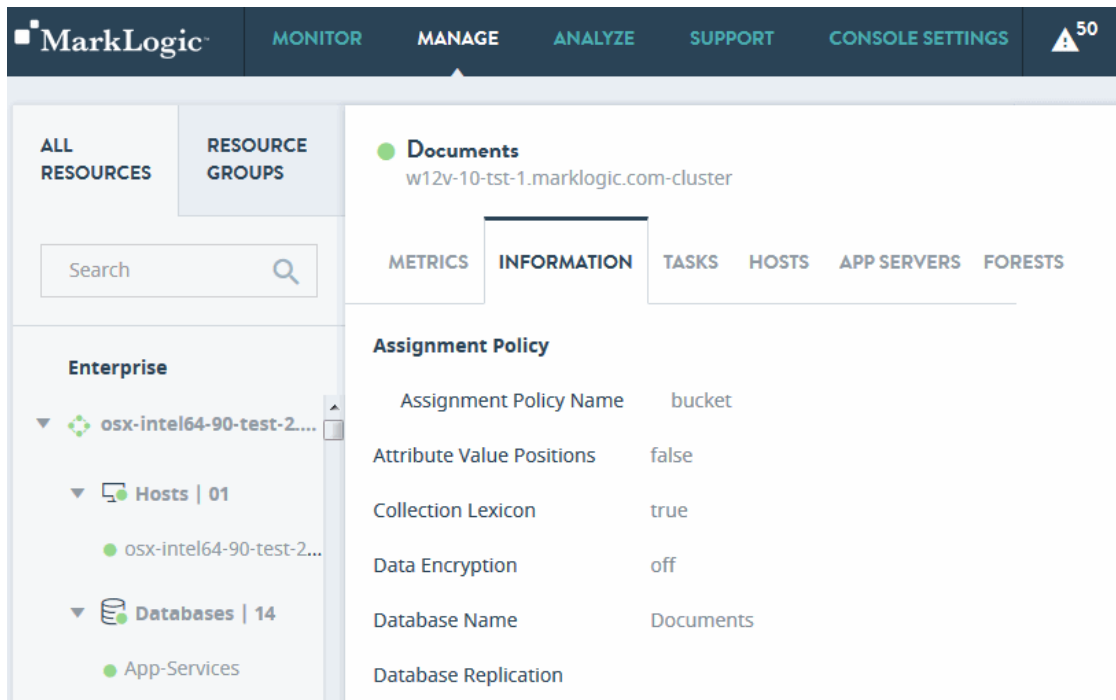
The metrics displayed by the charts on the database **METRICS** tab are described in the following table.

Chart	Definition of Displayed Metric
Disk I/O	Disk I/O in MB/sec. The displayed metrics are described in “Disk Performance Data” on page 195.

Chart	Definition of Displayed Metric
Memory Footprint	<p>The total amount (in MB) of memory consumed by the database.</p> <p>The displayed metrics are:</p> <ul style="list-style-type: none"> • RSS: The total amount of MB of Process Resident Size (RSS) consumed by the database. • Anon: The total amount of MB of Process Anonymous Memory consumed by the database.
Memory Size	<p>The amount of space (in MB) forest data files for the database take up in memory.</p>
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 database in pages per second. • Page-Out Rate: The page-out rate (from Linux /proc/vmstat) for the database in pages per second. • Swap-In Rate: The swap-in rate (from Linux /proc/vmstat) for the database in pages per second. • Swap-Out Rate: The swap-out rate (from Linux /proc/vmstat) for the database in pages per second.
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.

5.3.2 Database Information

Use the **INFORMATION** tab to access properties of the selected database.



The properties displayed in the database **INFORMATION** tab are described in the following table.

Field	Description
Assignment Policy	Database rebalancer assignment policy. For details, see Rebalancer Document Assignment Policies in the <i>Administrator's Guide</i> .
Attribute Value Positions	Specifies whether index data is included which speeds up the performance of proximity queries that use the <code>cts:field-value-query</code> function. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time.
Collection Lexicon	Specifies whether to create a lexicon of all of the collection URIs in the database. The collection lexicon lets you tolist all of the collection URIs in the database and perform lexicon-based queries on the URIs.
Data Encryption	Specifies whether or not encryption at rest is enabled. For details, see Encryption at Rest in the <i>Security Guide</i> .

Field	Description
Database Name	The name of the database.
Database Replication	Specifies whether or not database replication is enabled.
Directory Creation	<p>Specifies whether directories are automatically created in the database when documents are created. The default for a new database is manual. The settings are:</p> <ul style="list-style-type: none"> • automatic: Specifies that a directory hierarchy is automatically created to match the URI of a document or a directory that is created. This is the recommended setting, especially if you are accessing the database with a WebDAV Server or if you are using it as a Modules database. • manual: Specifies that directories must be manually created. No directory hierarchy is enforced. • manual-enforced: The same as manual, except it raises an error if the parent directory does not exist when creating a document or directory. For example, to create a document with the URI <code>http://marklogic/file.xml</code>, the directory <code>http://marklogic/</code> must first exist.
Element Value Positions	Specifies whether index data is included for faster element-based phrase and <code>cts:near-query</code> searches that use <code>cts:element-value-query</code> .
Element Word Positions	Specifies whether index data is included in the database files to enable proximity searches (<code>cts:near-query</code>) within specific XML elements or JSON properties. You must also enable word positions to perform element position searches. When this setting is true, positional searches are possible within an XML element or JSON property, but document loading is slower and the database files are larger.

Field	Description
Element Word Query Through	<p>The element markup to be searched through in searches that use the <code>cts:element-word-query</code> constructor.</p> <p>Words contained in text node children of a phrase through or Element Word Query Through element node match words in a <code>cts:element-word-query</code> for both the element and for the element's parent. If the parent is also a Phrase Through or Element Word Query Through element, the words also match words in a <code>cts:element-word-query</code> for the grandparent.</p> <ul style="list-style-type: none"> • Namespace URI: The namespace for an element or an attribute. • Localname: The local name for an XML element or attribute, or the name of a JSON property. <p>For details, see Element Word Query Throughs in the <i>Administrator's Guide</i>.</p>
Enabled	Whether data encryption is enabled. For details, see Encryption at Rest in the <i>Security Guide</i> .
Encryption Key ID	Data encryption key ID. For details, see Encryption at Rest in the <i>Security Guide</i> .
Expunge Locks	<p>Specifies if MarkLogic Server automatically expunges any lock fragments created using <code>xdmp:lock-acquire</code> with specified timeouts. The possible values are:</p> <ul style="list-style-type: none"> • automatic: Cleans up the locks as they expire. This is the default value and is recommended for most installations. • none: The lock fragments remain in the database after the locks expire (although they will no longer be locking any documents) until they are explicitly removed with <code>xdmp:lock-release</code>. This setting is only recommended to speed cluster startup time for extremely large clusters.
Fast Case Sensitive Searches	Specifies whether index terms are included in the database files to support fast case-sensitive searches. When this setting is true, case-sensitive searches are faster, but document loading is slower and the database files are larger.

Field	Description
Fast Diacritic Sensitive Searches	Specifies whether index terms are included in the database files to support fast diacritic-sensitive searches. When this setting is true, diacritic-sensitive searches are faster, but document loading is slower and the database files are larger.
Fast Element Character Searches	Specifies whether wildcard searches are enabled to speed up element-based wildcard searches. For details about wildcard searches, see Understanding and Using Wildcard Searches in the <i>Search Developer's Guide</i> .
Fast Element Phrase Searches	Specifies whether index terms are included in the database files to enable fast element-phrase searches. When set to true, element-phrase searches are faster, but document loading is slower and the database files are larger.
Fast Element Trailing Wildcard Searches	Specifies whether index terms are included in the database files to enable element trailing wildcard searches and faster character-based XQuery predicates. When set to true, element-trailing-wildcard searches are faster, but document loading is slower and the database files are larger.
Fast Element Word Searches	Specifies whether index terms are included in the database files to support fast element-word searches. When set to true, element-word searches are faster, but document loading is slower and the database files are larger.
Fast Phrase Searches	Specifies whether index terms are included in the database files to support fast phrase searches. When set to true, phrase searches are faster, but document loading is slower and the database files are larger.
Fast Reverse Searches	Specifies whether index terms are included in the database files to support fast reverse searches. When set to <code>true</code> , <code>cts:reverse-query</code> searches are faster, but document loading is slower and the database files are larger.
Field	The configuration of any fields in the database. For details on fields, see Fields Database Settings in the <i>Administrator's Guide</i> .
Field Value Positions	Specifies whether index data is included which speeds up the performance of proximity queries that use the <code>cts:field-value-query</code> function. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time.

Field	Description
Field Value Searches	Specifies whether index data is included which speeds up the performance of field value queries that use the <code>cts:field-value-query</code> function. Turn this index off if you are not interested in field value queries and if you want to conserve disk space and decrease loading time.
Forest	The names of the forests used by the database.
Format Compatibility	The version compatibility that MarkLogic Server applies to the indexes for the database during request evaluation. A value other than <code>automatic</code> specifies that all forest data has the specified on-disk format, and it disables the automatic checking for index compatibility information. The automatic detection occurs during database startup and after any database configuration changes, and can take some time and system resources for very large forests and for very large clusters. The default value of <code>automatic</code> is recommended for most installations.
In-memory Geospatial Region Index Size	The amount of cache and buffer memory to be allocated for managing geospatial region index data for an in-memory stand.
In-memory Limit	The maximum number of fragments in an in-memory stand. An in-memory stand contains the latest version of any new or changed fragments. Periodically, in-memory stands are written to disk as a new stand in the forest. Also, if a stand accumulates a number of fragments beyond this limit, it is automatically saved to disk by a background thread.
In-memory List Size	The amount of cache and buffer memory allocated for managing termlist data for an in-memory stand.
In-memory Range Index Size	The amount of cache and buffer memory allocated for managing range index data for an in-memory stand.
In-memory Reverse Index Size	The amount of cache and buffer memory allocated for managing reverse index data for an in-memory stand.
In-memory Tree Size	The size, in megabytes, of the in-memory tree storage. The In-memory Tree Size must be at least one or two megabytes larger than the largest text or small binary document you plan on loading into the database. The largest small binary file size is always constrained by the Large Size Threshold database configuration setting.

Field	Description
In-memory Triple Index Size	The amount of cache and buffer memory allocated for managing triple index data for an in-memory stand.
Index Detection	<p>Specifies whether to auto-detect index compatibility between the content and the current database settings. This detection occurs during database startup and after any database configuration changes, and can take some time and system resources for very large forests and for very large clusters. The possible values are:</p> <ul style="list-style-type: none"> • automatic: Auto-detect index. This is the default value and is recommended for most installations. • none: Causes queries to use the current database index settings, even if some settings have not completed reindexing.
Inherit Collections	When set to true, documents and directories automatically inherit collection settings from their parent directory (if collections are not set explicitly when creating the document or directory). If there are any default collections on the user who is creating the document or directory, those permissions are combined with any inherited collections.
Inherit Permissions	When set to true, documents and directories automatically inherit permissions from their parent directory (if permissions are not set explicitly when creating the document or directory). If there are any default permissions on the user who is creating the document or directory, those permissions are combined with any inherited permissions.
Inherit Quality	When set to true, documents and directories automatically inherit any quality settings from their parent directory (if quality is not set explicitly when creating the document or directory).
Journal Count	The number of journal files for the database.

Field	Description
Journal Size	<p>The size, in megabytes, of each journal file. The system uses journal files for recovery operations if a transaction fails to complete successfully. The default value is sufficient for most systems; it is calculated at database configuration time based on the size of your system. If you change the other memory settings, however, the journal size must equal the sum of the In-memory List Size and the In-memory Tree Size. Additionally, you must add space to the journal size if you use range indexes (particularly if you use a lot of range indexes or have extremely large range indexes), as range index data can take up journal space. Also, if your transactions span multiple forests, you may also need to add journal size, as each journal must keep the lock information for all of the documents in the transaction, not just for the documents that reside in the forest in which the journal exists.</p> <p>When you change the journal size, the next time the system creates a new journal, it will use the new size limit; existing journals will continue to use the old size limit until they are replaced with new ones (for example, when a journal fills up, when a forest is cleared, or when the system is cleanly shutdown and restarted).</p>
Journaling	<p>Specifies how robust transaction journaling is. The possible values are:</p> <ul style="list-style-type: none"> • strict: The journal protects against MarkLogic Server process failures, host operating system kernel failures, and host hardware failures. • fast: The journal protects against MarkLogic Server process failures but not against host operating system kernel failures or host hardware failures. • off: The journal does not protect against MarkLogic Server process failures, host operating system kernel failures, or host hardware failures.
Language	<p>The default language for content in the database. Any content without an <code>xml:lang</code> attribute is indexed in the language specified here.</p>

Field	Description
Large Size Threshold	The size, in kilobytes, beyond which large binary documents are stored in the Large Data Directory instead of directly in a stand. Binaries smaller than or equal to the threshold are considered small binary files and stored in stands. Binaries larger the threshold are considered large binary files and stored in the Large Data Directory.
Locking	<p>Specifies how robust transaction locking is. The possible values are:</p> <ul style="list-style-type: none"> • strict: Locking enforces mutual exclusion on existing documents and on new documents. • fast: Locking enforces mutual exclusion on existing and new documents. Instead of locking all the forests on new documents, it uses a hash function to select one forest to lock. In general, this is faster than strict. However, for a short period of time after a new forest is added, some of the transactions need to be retried internally. • off: Locking does not enforce mutual exclusion on existing documents or on new documents; only use this setting if you are sure all documents you are loading are new (a new bulk load, for example), otherwise you might create duplicate URIs in the database.
Maintain Directory Last Modified	Specifies whether to include a timestamp on the properties for each directory in the database. If set to true, update operations on documents in a directory also update the directory last-modified timestamp, which can cause some contention when multiple documents in the directory are being updated. If your application is experiencing contention during these type of updates (for example, if you see <code>deadlock-detected</code> messages in the error log), set this property to false . The default is false .
Maintain Last Modified	Specifies whether to include a timestamp on the properties document for each document in the database. The default is true .

Field	Description
Merge Max Size	<p>The maximum size, in megabytes, of a stand that will result from a merge. If a stand grows beyond the specified size, it will not be merged. If two stands would be larger than the specified size if merged, they will not be merged together. If you set this to smaller sizes, large merges (which may require more disk and CPU resources) will be prevented. The default is 48 GB (49152 MB), which is recommended because it provides a good balance between keeping the number of stands low and preventing very large merges from using large amounts of disk space. Set this to 0 to allow any sized stand to merge. Use care when setting this to a non-zero value lower than the default value, as this can prevent merges which are ultimately required for the system to maintain performance levels and to allow optimized updates to the system.</p> <p>For details on controlling merges, see Understanding and Controlling Database Merges in the <i>Administrator's Guide</i>.</p>
Merge Min Ratio	<p>A positive integer indicating the minimum ratio between the number of fragments in a stand and the number of fragments in all of the other smaller stands (the stands with fewer fragments) in the forest. Stands with a fragment count below this ratio relative to all smaller stands are automatically merged with the smaller stands. For an example, see If You Want to Reduce the Number of 'Large' Merges in the <i>Administrator's Guide</i>.</p>
Merge Min Size	<p>The minimum number of fragments that a stand can contain. Two or more stands with fewer than this number of fragments are automatically merged.</p> <p>For details on controlling merges, see Understanding and Controlling Database Merges in the <i>Administrator's Guide</i>.</p>

Field	Description
Merge Priority	<p>Specifies the CPU scheduler priority at which merges run. The settings are:</p> <ul style="list-style-type: none"> • normal: Specifies the same CPU scheduler priority as for requests. • lower: Specifies a lower CPU scheduler priority than for requests. <p>Merges always run with normal priority on forests with more than 16 stands.</p> <p>For details on controlling merges, see Understanding and Controlling Database Merges in the <i>Administrator's Guide</i>.</p>
Merge Timestamp	<p>The timestamp stored on merged stands. This is used for point-in-time queries, and determines when space occupied by deleted fragments and old versions of fragments may be reclaimed by the database. If a fragment is deleted or updated at a time after the merge timestamp, the old version of the fragment is retained for use in point-in-time queries. Set this to 0 (the default) to let the system reclaim the maximum amount of disk space during merge activities. A setting of 0 removes all deleted and updated fragments when a merge occurs. Set this to 1 before loading or updating any content to create a complete archive of the changes to the database over time. Set this to the current timestamp to preserve all versions of content from this point on. Set this to a negative number to specify a window of timestamp values, relative to the last merge, at ten million ticks per second. The timestamp is a number maintained by MarkLogic Server that increments every time a change occurs in any of the databases in a system (including configuration changes from any host in a cluster). For details on point-in-time queries, see Point-In-Time Queries in the <i>Application Developer's Guide</i>.</p>
One Character Searches	<p>Specifies whether wildcard searches are enabled so that the search pattern contains a single non-wildcard characters (for example, "a*"). This index is not needed if you have Three Character Searches and a word lexicon. For details about wildcard searches, see Understanding and Using Wildcard Searches in the <i>Search Developer's Guide</i>.</p>

Field	Description
Phrase Around	<p>Specifies the element markup phrased around in searches. Searches for phrases can skip around element boundaries when a phrase-around is configured for that element. For example, if there is a Phrase Around configured on a tag, the following XML snippet will return a match for a search for the phrase “to be or not to be”:</p> <pre data-bbox="773 537 1284 569" style="text-align: center;"><p>to be or not to be</p></pre> <p>The phrase indexer needs to know phrases may cross these markup boundaries at load time and reindex time so that it includes the right information in the indexes supporting phrase search.</p> <p>Each phrase-around has the following settings:</p> <ul style="list-style-type: none"> • Namespace URI: The namespace for an element or an attribute. • Localname: The local name for an XML element or attribute. <p>For more detail, see Phrasing and Element-Word-Query Boundary Control in the <i>Administrator’s Guide</i>.</p>

Field	Description
Phrase Through	<p>Specifies the element markup to be phrased through in searches. Searches for phrases can cross element boundaries when a phrase-through is configured for that element. For example, if there is a phrase-through configured on a <code></code> tag, the following XML snippet returns a match for a search for the phrase “to be or not to be”:</p> <pre data-bbox="773 537 1284 569" style="text-align: center;"><p>to be or not to be</p></pre> <p>The phrase indexer needs to know phrases may cross these markup boundaries at load time and reindex time so that it includes the right information in the indexes supporting phrase search.</p> <p>Words contained in text node children of a phrase-through or Element Word Query Through element node match words in a <code>cts:element-word-query</code> for both the element and for the element's parent. If the parent is also a phrase-through or Element Word Query Through element, the words also match words in a <code>cts:element-word-query</code> for the grandparent.</p> <p>If you delete a phrase-through, the system will not automatically reindex away the indexed phrase-throughs, even if reindexing is enabled; to reindex deleted phrase-throughs, force a reindex (for example, by clicking the reindex button on the database configuration page).</p> <p>Each phrase-through has the following settings:</p> <ul style="list-style-type: none"> • Namespace URI: Specifies the namespace for an element or an attribute. • Localname: Specifies the local name for an XML element or attribute. <p>For more detail, see Phrasing and Element-Word-Query Boundary Control in the <i>Administrator's Guide</i>.</p>

Field	Description
Positions List Max Size	The maximum size, in megabytes, of the position list portion of the index for a given term. If the position list size for a given term grows larger than the limit specified, the position information for that term is discarded. The default value is 256, the minimum value is 1, and the maximum value is 512. For example, position queries (<code>cts:near-query</code>) for frequently occurring words that have reached this limit (words like a, an, the, and so on) are resolved without using the indexes. Even though those types of words are resolved without using the indexes, this limit helps improve performance by making the indexes smaller and more efficient to the data actually loaded in the database.
Preallocate Journals	Has no effect as of MarkLogic, Version 8.0-4.
Preload Mapped Data	Specifies whether memory mapped data (for example, range indexes and word lexicons) is loaded into memory when a forest is mounted to the database. Preloading the memory mapped data improves query performance, but uses more memory, especially if you have a lot of range indexes and/or lexicons. Also, it will cause a lot of disk I/O at database startup time, slowing the system performance during the time the mapped data is read into memory. If you do not preload the mapped data, it will be paged into memory dynamically when a query requests data that needs it, slowing the query response time.
Preload Replica Mapped Data	Specifies whether memory mapped data (for example, range indexes and word lexicons) are loaded immediately into memory when a stand is opened. The setting of <code>preload-replica-mapped-data</code> is ignored if Preload Mapped Data is set to false .

Field	Description
Range Element Attribute Index	<p>The range indexes on attributes in XML elements.</p> <p>The fields are as follows:</p> <ul style="list-style-type: none"> • Range Value Positions: Specifies whether index data is included that speeds up the performance of proximity queries involving range queries. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time. • Parent Namespace URI: The namespace URI of the parent element. • Parent Localname: The local name of the parent XML element or the name of the parent JSON property. • Invalid Values: Specifies whether the server allows insertion of documents that contain XML elements or JSON properties on which range index is configured and their contents cannot be coerced to the index data type. It can be configured to either ignore or reject. By default server rejects insertion of such documents. However, if you set invalid values to ignore, these documents can be inserted. This setting does not change the behavior of queries on invalid values after documents are inserted in the database. Performing an operation on an invalid value at query time can still result in an error. • Namespace URI: The namespace for an element or an attribute. • Collation: The collation to use in the comparison operations. • Scalar Type: The ordered domain for the index. Each of the types correspond with an XQuery type. • Localname: The local name of an XML element or attribute, or the name of a JSON property. <p>For details, see Defining Attribute Range Indexes in the <i>Administrator's Guide</i>.</p>

Field	Description
Range Element Index	<p>The range indexes on XML elements or JSON properties.</p> <p>The fields are as follows:</p> <ul style="list-style-type: none"> • Range Value Positions: Specifies whether index data is included that speeds up the performance of proximity queries involving range queries. Turn this index off if you are not interested in proximity queries and if you want to conserve disk space and decrease loading time. • Invalid Values: Specifies whether the server allows insertion of documents that contain XML elements or JSON properties on which range index is configured and their contents cannot be coerced to the index data type. It can be configured to either ignore or reject. By default server rejects insertion of such documents. However, if you set invalid values to ignore, these documents can be inserted. This setting does not change the behavior of queries on invalid values after documents are inserted in the database. Performing an operation on an invalid value at query time can still result in an error. • Namespace URI: The namespace for an element or an attribute. • Collation: The collation to use in the comparison operations. • Scalar Type: The ordered domain for the index. Each of the types correspond with an XQuery type. • Localname: The local name for an XML element or attribute, or the name for a JSON property. <p>For details, see Defining Element Range Indexes in the <i>Administrator's Guide</i>.</p>
Range Index Optimize	<p>Specifies how range indexes are to be optimized. The values are defined as follows:</p> <ul style="list-style-type: none"> • facet-time: Range indexes are optimized to minimize the amount of CPU time used. • memory-size: Range indexes are optimized to minimize the amount of memory used.

Field	Description
Rebalancer Enable	<p>Specifies whether rebalancing are automatically performed in the background after configuration settings are changed. When set to true, the database rebalancer automatically redistributes the content across the database forests. When set to false, rebalancing is disabled.</p> <p>For details, see Configuring the Rebalancer on a Database in the <i>Administrator's Guide</i>.</p>
Rebalancer Throttle	<p>The priority of system resources devoted to rebalancing. Higher numbers give rebalancing a higher priority.</p> <p>For details, see How the Rebalancer Moves Documents in the <i>Administrator's Guide</i>.</p>
Reindexer Enable	<p>Specifies whether indexes are automatically rebuilt in the background after index configuration settings are changed. When set to true, index configuration changes automatically initiate a background reindexing operation on the entire database. When set to false, any new index settings take effect for future documents loaded into the database; existing documents retain the old settings until they are reloaded or until you set reindexer enable to true.</p> <p>For details, see Understanding the Reindexer Enable Settings in the <i>Administrator's Guide</i>.</p>
Reindexer Throttle	<p>The priority of system resources devoted to reindexing. Reindexing occurs in batches, where each batch is approximately 200 fragments. When set to 5 (the default), the reindexer works aggressively, starting the next batch of reindexing soon after finishing the previous batch. When set to 4, it waits longer between batches, when set to 3 it waits longer still, and so on until when it is set to 1, when it waits the longest. Therefore, higher numbers give reindexing a higher priority and use the most system resources.</p> <p>For details, see Understanding the Reindexer Enable Settings in the <i>Administrator's Guide</i>.</p>

Field	Description
Reindexer Timestamp	<p>The timestamp of fragments to force a reindex or refragment operation. If Reindex Enable is set to true, a reindex and refragment operation are done on all fragments in the database that have a timestamp equal to or less than the displayed timestamp. Note that when you restore a database that has a timestamp set, if there are fragments in the restored content that are older than the specified content, they will start to reindex as soon as they are restored.</p>
Retain Until Backup	<p>Specifies whether the deleted fragments are retained since the last full or incremental backup. When enabled, Retain Until Backup supersedes Merge Timestamp. Deleted fragments are not merged until backups are finished, regardless of the Merge Timestamp setting. Enabling Retain Until Backup is same to setting the Merge Timestamp to the timestamp of the last backup.</p> <p>For more information, see Incremental Backup with Journal Archiving in the <i>Administrator's Guide</i>.</p>
Retired Forest Count	<p>The number of forests used by the database to be retired.</p> <p>For more information, see Retiring a Forest from the Database in the <i>Administrator's Guide</i>.</p>
Security Database	<p>The security database used by the database.</p> <p>For more information, see Overview of the Security Database in the <i>Security Guide</i>.</p>
Stemmed Searches	<p>The level of stemming applied to word searches. Stemmed searches match not only the exact word in the search, but also words that come from the same stem and mean the same thing (for example, a search for <code>be</code> will also match the term <code>is</code>). For details on stemmed searches, see Understanding and Using Stemmed Searches in the <i>Search Developer's Guide</i>.</p>

Field	Description
TF Normalization	<p>Specifies the term-frequency normalization. The values are described as follows:</p> <ul style="list-style-type: none"> • scaled-log: The default term-frequency normalization, which specifies the maximum scaling of the document based on document size. • unscaled-log: No scaling based on document size. • weakest-scaled-log, weakly-scaled-log, moderately-scaled-log, strongly-scaled-log: Increased degrees of scaling in between the least and the most scaling. <p>For more information, see Term Frequency Normalization in the <i>Search Developer's Guide</i>.</p>
Three Character Searches	<p>Specifies whether to enable wildcard searches where the search pattern contains three or more consecutive non-wildcard characters (for example, <code>abc*x</code>, <code>*abc</code>, <code>a?bcd</code>). When combined with a codepoint word lexicon, speeds the performance of any wildcard search (including searches with fewer than three consecutive non-wildcard characters). MarkLogic recommends combining the three character search index with a codepoint collation word lexicon.</p> <p>For details about wildcard searches, see Understanding and Using Wildcard Searches in the <i>Search Developer's Guide</i>.</p>
Three Character Word Positions	<p>Specifies whether index data is included in the database files to enable proximity searches (<code>cts:near-query</code>) within wildcard queries. You must also enable Three Character Searches to perform wildcard position searches. When set to <code>true</code>, positional searches are possible within a wildcard query, but document loading is slower and the database files are larger.</p> <p>For details about wildcard searches, see Understanding the Wildcard Indexes in the <i>Search Developer's Guide</i>.</p>

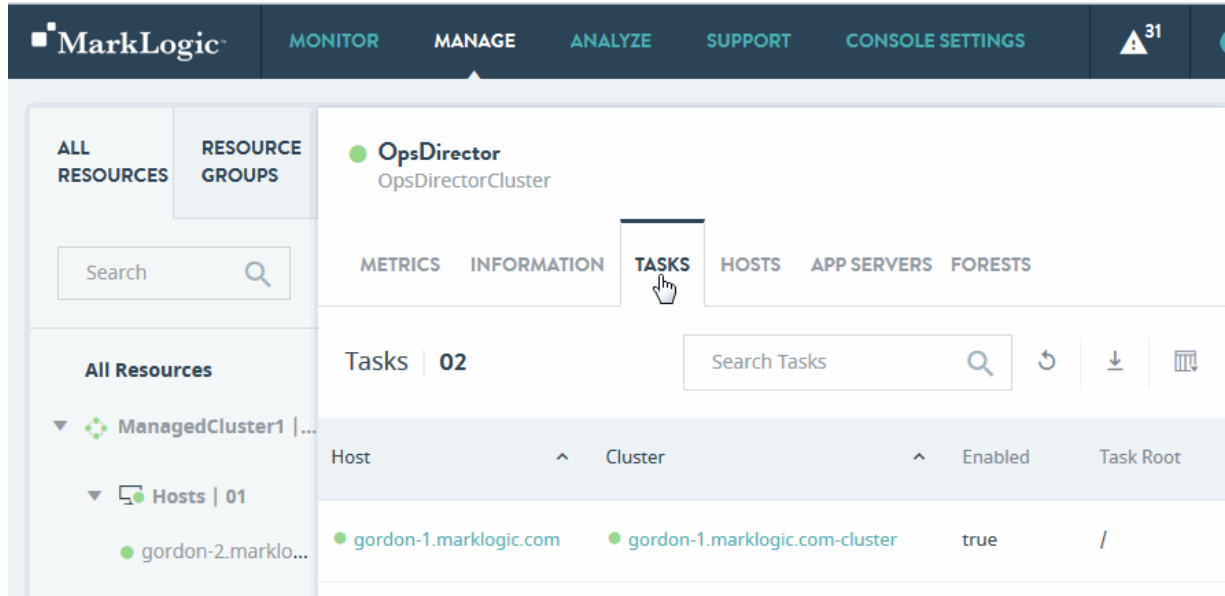
Field	Description
Trailing Wildcard Searches	<p>Specifies whether indexes are created to enable wildcard searches where the search pattern contains one or more consecutive non-wildcard characters at the beginning of the word, with the wildcard at the end of the word (for example, abc*). When this setting is true, character searches are faster, but document loading is slower and the database files are larger.</p> <p>For details about wildcard searches, see Understanding and Using Wildcard Searches in the <i>Search Developer's Guide</i>.</p>
Trailing Wildcard Word Positions	<p>Specifies whether index data is included in the database files to enable proximity searches (<code>cts:near-query</code>) within trailing wildcard queries. You must also enable trailing wildcard searches to perform trailing wildcard position searches. When this setting is true, positional searches are possible within a trailing wildcard query, but document loading is slower and the database files are larger.</p> <p>For details about wildcard searches, see Understanding the Wildcard Indexes in the <i>Search Developer's Guide</i>.</p>
Triple Index	<p>Specifies whether the RDF triple index is enabled to support SPARQL execution over RDF triples. When set to <code>true</code>, <code>sem:sparql</code> can be used, but document loading is slower and the database files are larger. This index must also be enabled to support Optic and SQL queries.</p> <p style="text-align: center;">Note: This feature requires a valid semantics license key.</p> <p>For details, see Triple Index Overview in the <i>Semantics Developer's Guide</i>.</p>
Triple Positions	<p>Specifies whether index data is included which speeds up the performance of proximity queries that use the <code>cts:triple-range-query</code> function. Triple positions also improve the accuracy of the item-frequency option of <code>cts:triples</code>.</p> <p>For details, see Triple Positions in the <i>Semantics Developer's Guide</i>.</p>

Field	Description
Two Character Searches	<p>Specifies whether to enable wildcard searches where the search pattern contains two or more consecutive non-wildcard characters (for example, <code>ab*</code>). This index is not needed if you have Three Character Searches and a word lexicon.</p> <p>For details about wildcard searches, see Understanding and Using Wildcard Searches in the <i>Search Developer's Guide</i>.</p>
URI Lexicon	<p>Specifies whether to create a lexicon of all of the URIs used in a database. The URI lexicon speeds up queries that constrain on URIs. It is like a range index of all of the URIs in the database. To access values from the URI lexicon, use the <code>cts:uris</code> or <code>cts:uri-match</code> functions.</p> <p>For details, see URI and Collection Lexicons in the <i>Search Developer's Guide</i>.</p>
Word Positions	<p>Specifies whether index data is included in the database files to enable proximity searches (<code>cts:near-query</code>). When set to true, positional searches are possible, but document loading is slower and the database files are larger.</p> <p>For details, see Positions Indexes Can Help Speed Phrase Searches in the <i>Query Performance and Tuning Guide</i>.</p>
Word Searches	<p>Specifies whether index terms are included in the database files to support fast word searches. When this setting is <code>true</code>, word searches are faster, but document loading is slower and the database files are larger.</p> <p>For details, see the <i>Search Developer's Guide</i>.</p>

5.3.3 Database Tasks

Use the **TASKS** tab to get the list of the database tasks.

Note: You can only see the tasks for the databases to which you have access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the database **TASKS** tab are described in the following table.

Column	Description
Cluster	Cluster on which the task host is located.
Host	The hostname of the host computer on which the scheduled module is to be invoked.
Enabled	Whether the task is enabled (true) or disabled (false).
Task Root	The root directory (filesystem) or URI root (database) that contains the module.
Task Path	The module the task is to invoke.

Column	Description
Task Type	<p>The task type:</p> <ul style="list-style-type: none"> • minutely specifies how many minutes between each invocation of the module. • hourly specifies how many hours and minutes between each invocation of the module. • daily specifies how many days between each invocation of the module and the time of day (in 24:00 notation). • weekly specifies how many weeks between each invocation of the module, check one or more days of the week, and the time of day (in 24:00 notation) for the task to start. • monthly specifies how many months between each invocation of the module, select one day of the month (1-31), and the time of day (in 24:00 notation) for the task to start. • one-time specifies the start date in MM/DD/YYYY notation (for example, 07/29/2009 for July 29, 2009) and time in 24:00 notation.
Task Period	How often the module is to be invoked (every n months, weeks, days, hours, or minutes).
Created on	The datetime the task was created.
Database	The database to which the scheduled module connects for query execution.
Task Modules	The name of the database in which the scheduled module locates the application code. If set to (filesystem), any files in the specified task root directory are executable (given the proper permissions). If set to a database, any documents in the database whose URI begins with the specified task root directory are executable.
User	The user with permission to invoke the module.
Priority	<p>The priority of the task:</p> <ul style="list-style-type: none"> • normal specifies the task is queued with normal priority. • higher specifies the task is queued with higher priority.

You can export data from the Database Tasks tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

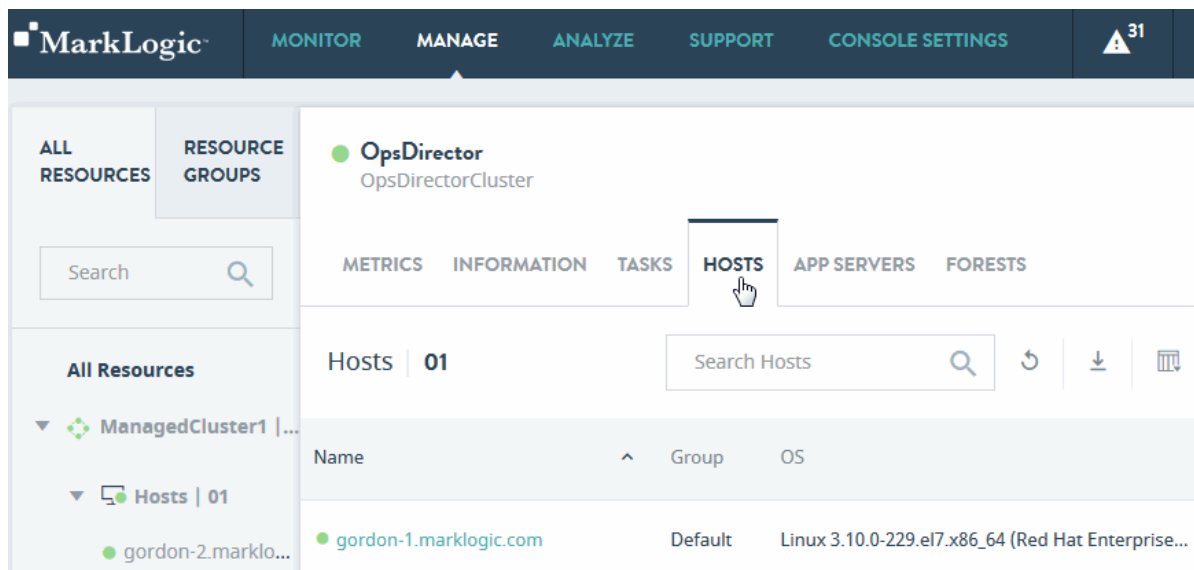
- The resulting CSV file will have the same columns as the Database Tasks table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.3.4 Database Hosts

Use the **HOSTS** tab to get the list of hosts on which the database is deployed. This tab shows host name as the first column, followed by other properties. Click on the host name to navigate to the manage view of the selected host.

Note: Access to hosts is not implicit with database access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the Database Hosts tab are described in the following table.

Column	Description
Name	The hostname of the host.

Column	Description
Group	The name of the group that contains the host.
OS	The name and version of the operating system on which the host runs.
Server Version	The version of MarkLogic Server on each host.
Forests	The number of forests on the host.
Databases	The number of databases on the host.
App Servers	The number of App Servers on the host.
Disk Space (MB)	The amount of disk space (in MB) used on the host.
Uptime	The duration (Days Hrs:Min) the host has been available.
Maint. Mode	The host maintenance mode (normal or maintenance). For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .
Zone	The Amazon Web Services (AWS) zone in which the host resides, if applicable. For details, see <i>MarkLogic Server on Amazon Web Services (AWS) Guide</i> .

You can export data from the Database Hosts tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

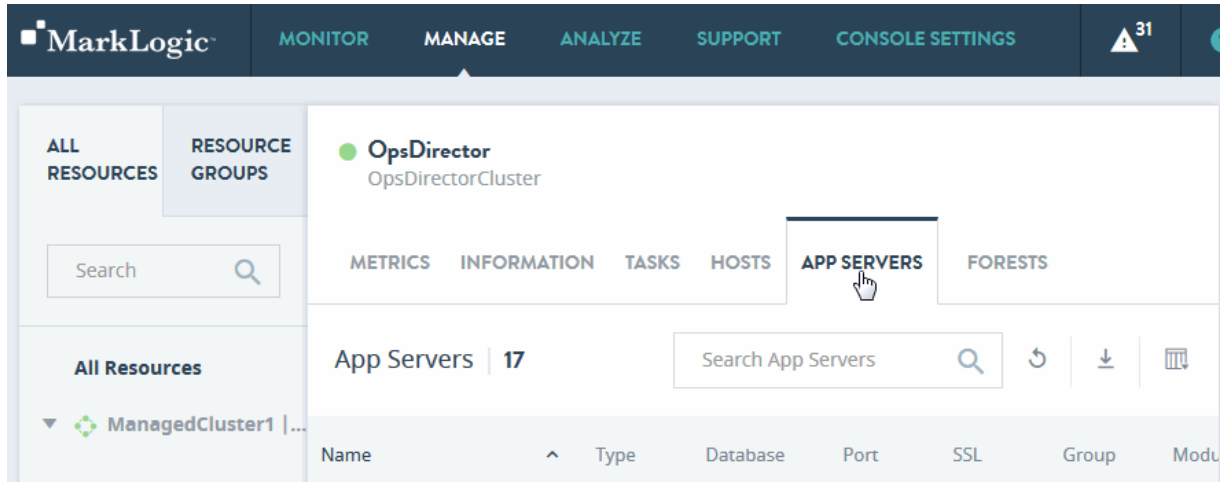
- The resulting CSV file will have the same columns as the Database Hosts table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.3.5 Database App Servers

Use the **APP SERVERS** tab to get the list of all the application servers that use the database, with App Server name as the first column followed by other properties. Click on the App Server name to navigate to the METRICS view of the selected application server.

Note: Access to the APP SERVERS tab is not implicit with database access. For access rules, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed in the database **APP SERVERS** tab are described in the following table.

Column	Description
Name	The name of the App Server.
Type	The App Server type (HTTP, WebDAV, XDBC, ODBC).
Database	The name of the App Server content database.
Port	The port number used to access the App Server.
SSL	Whether the App Server has SSL enabled (yes) or disabled (no). For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Group	The name of the group to which the App Server belongs
Modules DB+Root	The name of the modules database, or if <code>filesystem</code> , the root directory.
Security	The type of security (<code>internal</code> or <code>external</code>).

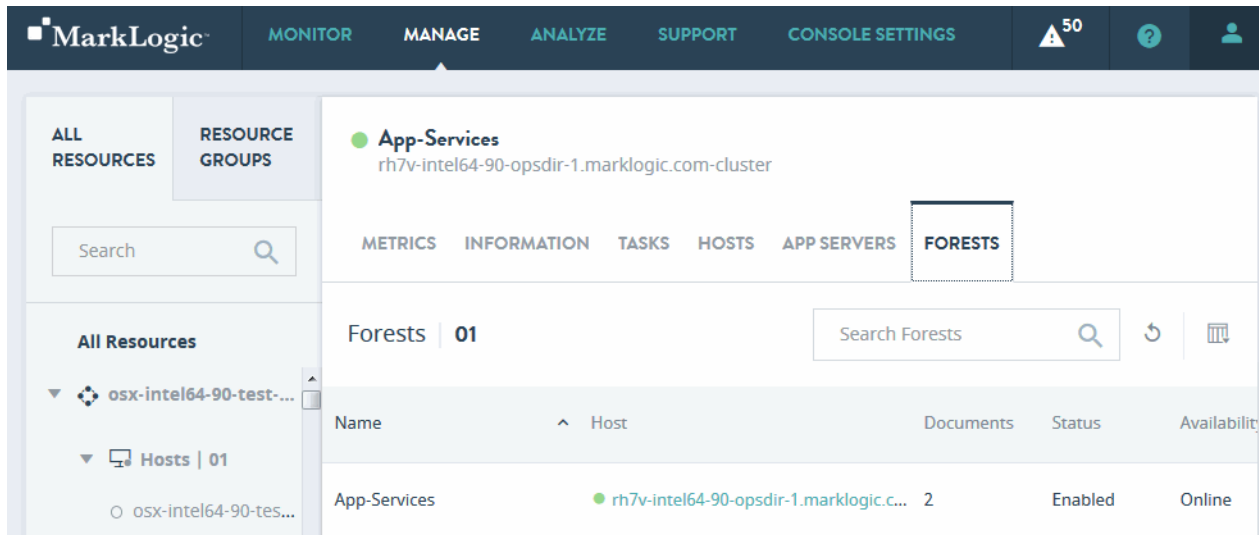
You can export data from the database **APP SERVERS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Database App Servers table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.3.6 Database Forests

Use the **FORESTS** tab to list all the forests within the database, with forest name as the first column, followed by other properties. Click on the forest name to navigate to the manage view of the selected forest.



The columns displayed in the database **FORESTS** tab are described in the following table.

Column	Description
Name	The name of the forest.
Host	The forest host.
Documents	The number of documents in the forest.
Status	The status of the forest. Enabled or Disabled .
Availability	The availability of the forest. Online or Offline .
Fragments	The number of active fragments (the fragments available to queries) in the forest.
Deleted Fragments	The number of deleted fragments (the fragments to be removed by the next merge operation) in the forest.
Stands	The number of stands in the forest. For more information on stands, see Databases, Forests, and Stands in the <i>Concepts Guide</i> .

Column	Description
Size (MB)	The size of the forest, in MB.
Encrypted Size (MB)	The amount of encrypted data in the forest. For details on data encryption, see Encryption at Rest in the <i>Security Guide</i> .
Free Space (MB)	The number of MB of free space on this forest.
Large Data Size (MB)	The amount of data in the large data directories of the forest. For more information on Large Data, see Working With Binary Documents in the <i>Application Developer's Guide</i> .
Fast Data Size (MB)	The amount of data in the fast data directories of the forest. For more information on Fast Data, see Fast Data Directory on Forests in the <i>Query Performance and Tuning Guide</i> .
Failover Enabled	Whether failover is enabled for the forest. For more information on stands, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled for this forest. For details, see the <i>Database Replication Guide</i> .

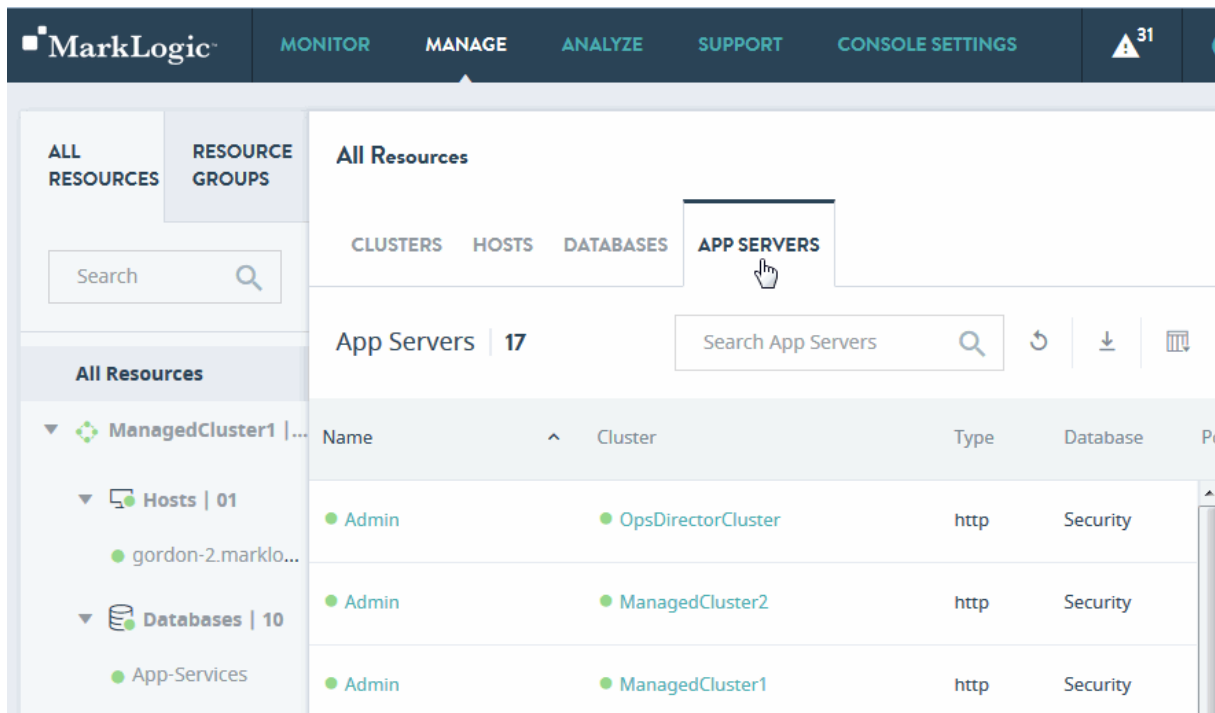
You can export data from the Database Forests tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Database Forests table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.4 Manage App Servers Tab

The **APP SERVERS** tab displays a list of App Servers in your enterprise.



The columns displayed in the manage **APP SERVERS** tab are described in the following table.

Column	Description
Name	The name of the App Server.
Cluster	The name of the cluster that hosts the App Server.
Type	The App Server type (HTTP, WebDAV, XDBC, ODBC).
Database	The name of the App Server content database.
Port	The port number used to access the App Server.
SSL	Whether the App Server has SSL enabled (yes) or disabled (no). For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Group	The name of the group to which the App Server belongs
Modules DB+Root	The name of the modules database, or if <code>filesystem</code> , the root directory.
Security	The type of security (internal or external).

You can export data from the Manage App Servers tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Manage App Servers table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, *Yes/No* in the UI corresponds to `TRUE/FALSE` in the CSV file.

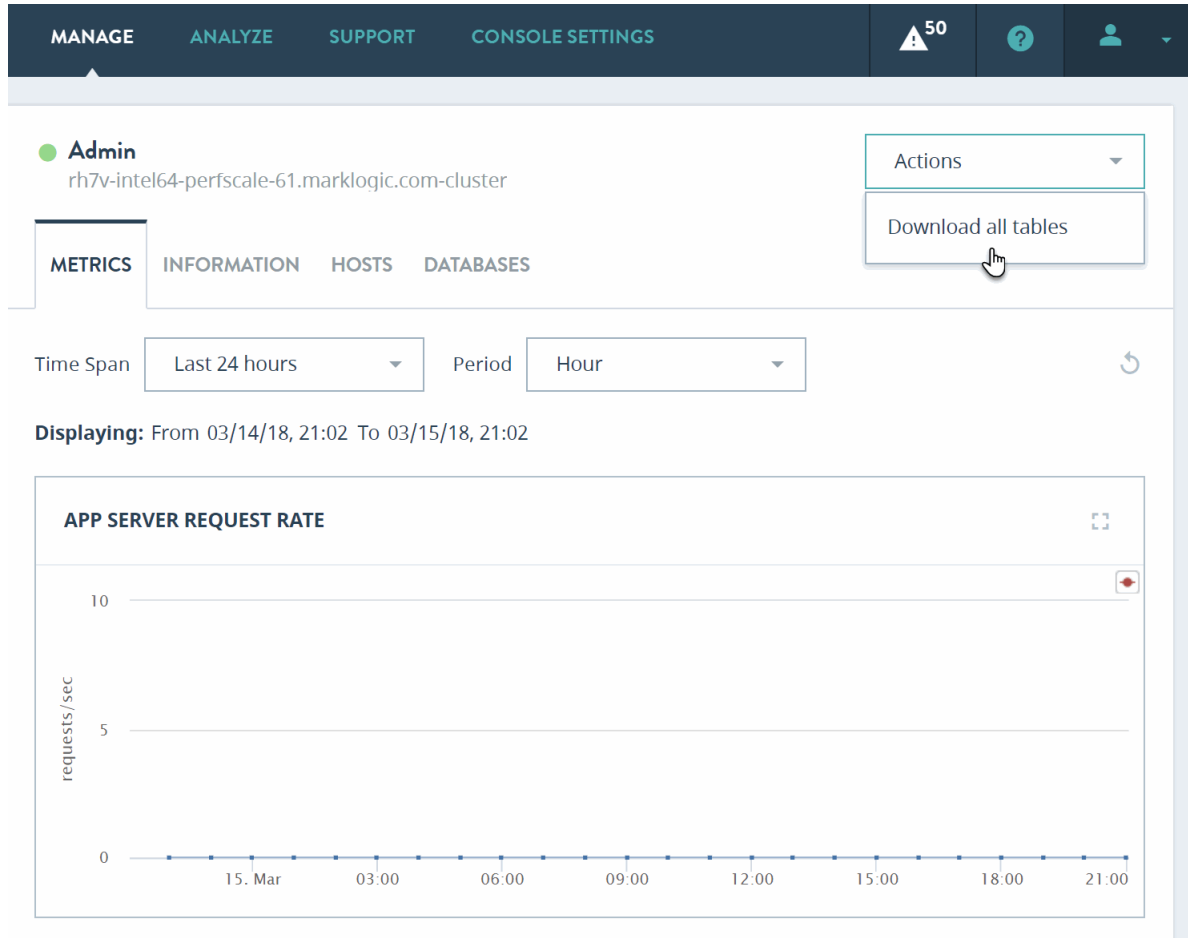
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

To drill down for a particular App Server, click on the App Server name in the App Servers table. The following content tabs each represent an information category for the selected App Server:

- [App Server Metrics](#)
- [App Server Information](#)
- [App Server Hosts](#)
- [App Server Databases](#)

You can export all tables from the App Server content tabs for a particular App Server. When you select a specific App Server (either in the left side resource navigation panel or in the content area of the view), **Actions** menu becomes available in the upper right corner.

From the **Actions** menu, select **Download all tables** to export all tables for the selected App Server.

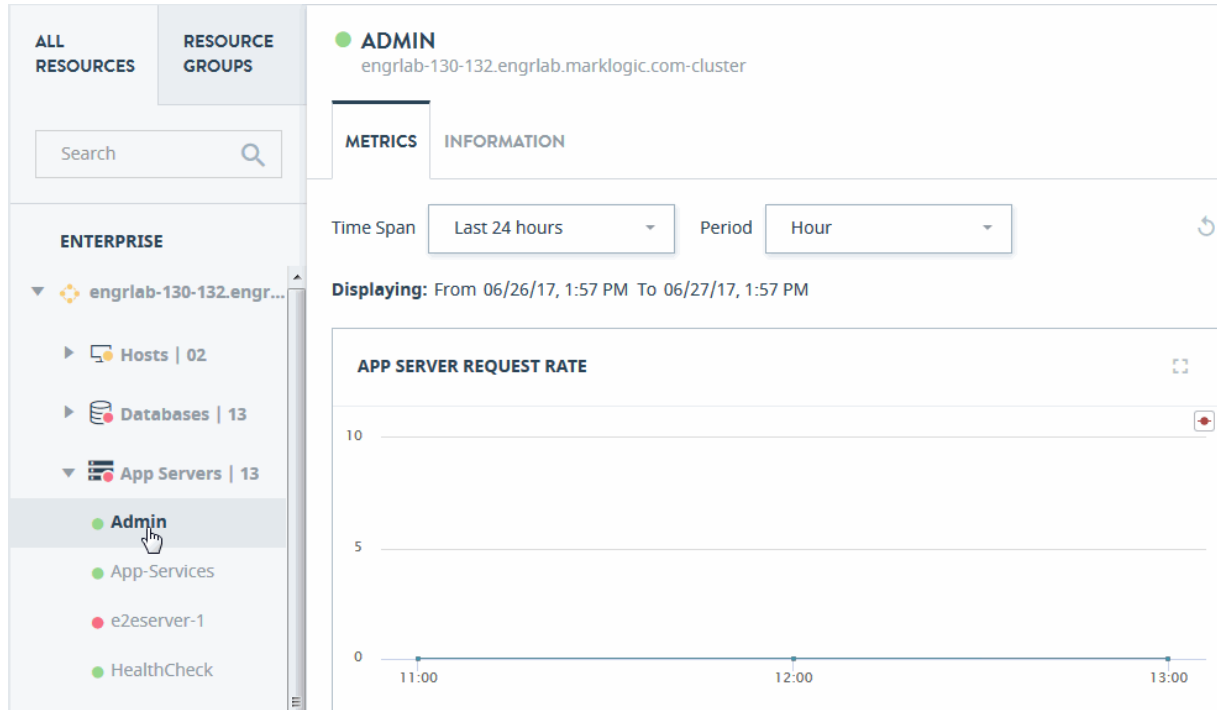


As the result, a zip file with all exported tables for this App Server will be downloaded to your computer. Each table is represented by the corresponding CSV file. The zip file will contain the following CSV files for the App Server:

- Hosts CVS file (see the [App Server Hosts](#) section for details on the contents)
- Databases CVS file (see the [App Server Databases](#) section for details on the contents)

5.4.1 App Server Metrics

Select a single App Server or all App Servers from the navigation tree. The **METRICS** tab displays key indicators allowing administrators to determine the selected resource or resource group's health.



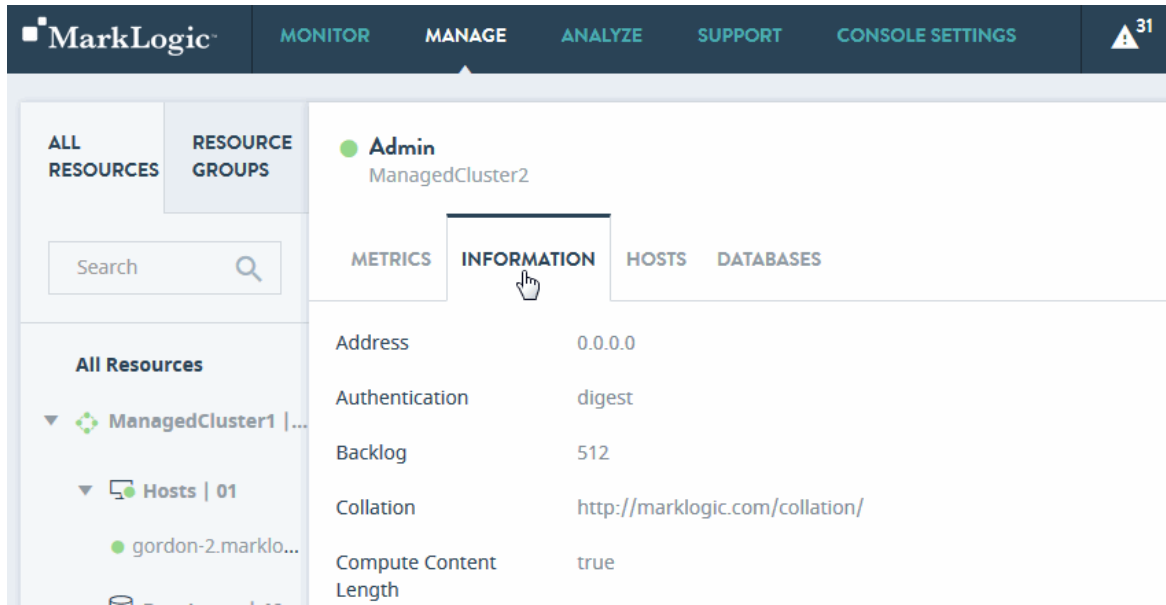
To filter the data used for rendering the graphs, select a pre-defined time period or specify a custom time period, as described in “Date and Time Filters” on page 86.

The metrics displayed by charts on the App Server Metrics tab are described in the following table.

Chart	Definition of Displayed Metric
App Server Request Rate	The total number of queries being processed per second, across all of the App Servers in the cluster.
App Server Latency	The average time (in seconds) it takes to process queries, across all of the App Servers in the cluster.
Expanded Tree Cache Hits/Misses	The aggregate I/O performance data for the CPUs in the cluster.
Network	Various XDQP performance metrics as the sum of XDQP activity across the cluster. (For further detail, see Cluster Metrics.)

5.4.2 App Server Information

Use the Information tab to access properties of the selected application server.



The properties displayed in the App Server Information tab are described in the following table.

Note: Most of these properties are described in detail in the [HTTP Servers](#), [XDBC Servers](#), [WebDAV Servers](#), and [ODBC Servers](#) chapters in the *Administrator's Guide*.

Field	Description
Address	The IP address for the App Server.
Authentication	The authentication scheme used by this App Server. The authentication scheme can be one of: <code>digest</code> , <code>basic</code> , <code>digestbasic</code> , <code>certificate</code> , <code>kerberos-ticket</code> , or <code>application-level</code> . For details, see Types of Authentication in the <i>Security Guide</i> .
Backlog	The maximum number of pending connections allowed on the App Server socket.
Collation	The default collation for queries run in this App Server. This is the collation used for string comparison and sorting if none is specified in the query. For details, see Encodings and Collations in the <i>Search Developer's Guide</i> .
Compute Content Length	Specifies whether to compute content length when using a webDAV server. A value of <code>true</code> indicates to compute content length; otherwise <code>false</code> .

Field	Description
Concurrent Request Limit	The maximum number of requests any user may have been running at a specific time. 0 indicates no maximum. For details, see Managing Concurrent User Sessions in the <i>Administrator's Guide</i> .
Content Database	The content database for this App Server.
Coordinate System	The default coordinate system for queries run in this App Server. This will be the coordinate system used for geospatial operations if none is specified in the query. For details, see Understanding Coordinate Systems in the <i>Search Developer's Guide</i> .
Debug Allow	Specifies whether to allow requests against this App Server to be stopped for debugging, using the MarkLogic Server debugging APIs. A value of <code>true</code> allows the App Server to be stopped; otherwise <code>false</code> .
Default Error Format	<p>The default format for protocol errors for this server. The value can be:</p> <ul style="list-style-type: none"> • <code>html</code>: Errors formatted as HTML. • <code>xml</code>: Errors formatted as XML. • <code>json</code>: Errors formatted as JSON. • <code>compatible</code>: Match as closely as possible the format used in prior releases for the type of request and error. <p>For more information, see set-error-format in the <i>Application Developer's Guide</i>.</p>
Default Inference Size	The amount of memory available to use for semantic inference queries. By default the amount of memory available for inference is 100mb (<code>size=100</code>). For details, see Memory Available for Inference in the <i>Semantics Developer's Guide</i> .
Default Time Limit	The default value for any request's time limit (in seconds), when otherwise unspecified. A request can change its time limit using the <code>xdmp:set-request-time-limit</code> function. The time limit, in turn, is the maximum number of seconds allowed for servicing a query request. The App Server gives up on queries that take longer, and returns an error.

Field	Description
Default User	The user used as the default user in application-level authentication. Setting the <code>admin</code> user as the default user is equivalent to turning security off. For details, see Application Level in the <i>Security Guide</i> .
Default XQuery Version	The default XQuery language for this App Server if an XQuery module does explicitly declare its language version.
Display Last Login	Specifies whether the <code>xdmp:display-last-login</code> API returns true or false in the <code>display-last-login</code> element. For details, see Storing and Monitoring the Last User Login Attempt in the <i>Administrator's Guide</i> .
Distribute Timestamps	<p>Specifies how the latest timestamp is distributed after updates. This affects performance of updates and the timeliness of read-after-write query results from other hosts in the group. The possible values are:</p> <ul style="list-style-type: none"> • fast: Updates return as quickly as possible. No special timestamp notification messages are broadcasted to other hosts. Instead, timestamps are distributed to other hosts when any other message is sent. The maximum amount of time that could pass before other hosts see the timestamp is one second, because a heartbeat message is sent to other hosts every second. • strict: Updates immediately broadcast timestamp notification messages to every other host in the group. Updates do not return until their timestamp has been distributed. This ensures timeliness of read-after-write query results from other hosts in the group, so requests made to this app server on other hosts in the group will see immediately consistent results. • cluster: Updates immediately broadcast timestamp notification messages to every other host in the cluster. Updates do not return until their timestamp has been distributed. This ensures timeliness of read-after-write query results from any host in the cluster, so requests made to any app server on any host in the cluster will see immediately consistent results.

Field	Description
Enabled	Whether this App Server is enabled (true) or disabled (false).
Error Handler	The page to internally redirect to in case of any 400 or 500 errors. For details, see Controlling App Server Access, Output, and Errors in the <i>Application Developer's Guide</i> .
Execute	Specifies whether this App Server executes XQuery modules (for example, a WebDAV server does not). If true , the App Server executes XQuery modules; otherwise <code>false</code> .
File Log Level	The minimum log level for log messages sent to the MarkLogic Server log file (<code>ErrorLog.txt</code>). For a description of the log levels, see Understanding the Log Levels in the <i>Administrator's Guide</i> .
Group Name	The group to which this App Server belongs. For details on groups, see Groups in the <i>Administrator's Guide</i> .
Internal Security	Specifies whether the security database is used for authentication and authorization if the user is found in the security database. If true , use the security database; if false , external authentication/authorization is used. For details, see External Security in the <i>Security Guide</i> .
Keep Alive Timeout	The maximum number of seconds before a socket receives a timeout for subsequent requests over the same connection.
Log Errors	Specifies whether to log uncaught errors for this App Server to the <code>ErrorLog.txt</code> file. This is useful to log exceptions that might occur on an App Server for later debugging. If true , log errors; if false , do not log errors.
Max Inference Size	The maximum memory limit for semantic inference queries. For details, see Memory Available for Inference in the <i>Semantics Developer's Guide</i> .
Max Time Limit	The upper bound for any request's time limit. No request may set its time limit (for example with <code>xdmp:set-request-time-limit</code>) higher than this number. The time limit, in turn, is the maximum number of seconds allowed for servicing a query request. The App Server gives up on queries which take longer, and returns an error.

Field	Description
Multi Version Concurrency Control	<p>Specifies how the latest timestamp is chosen for lock-free queries. This only affects query statements, not update statements. The following are the possible values:</p> <ul style="list-style-type: none">• contemporaneous: App Server chooses the latest timestamp for which any transaction is known to have committed, even though there still may be other transactions for that timestamp that have not yet fully committed. Queries will see more timely results, but may block waiting for contemporaneous transactions to fully commit.• nonblocking: App Server chooses the latest timestamp for which all transactions are known to have committed, even though there may be a slightly later timestamp for which another transaction has committed. Queries won't block waiting for transactions, but they may see less timely results. <p>For details about queries and transactions in MarkLogic Server, see Understanding Transactions in MarkLogic Server in the <i>Application Developer's Guide</i>.</p>

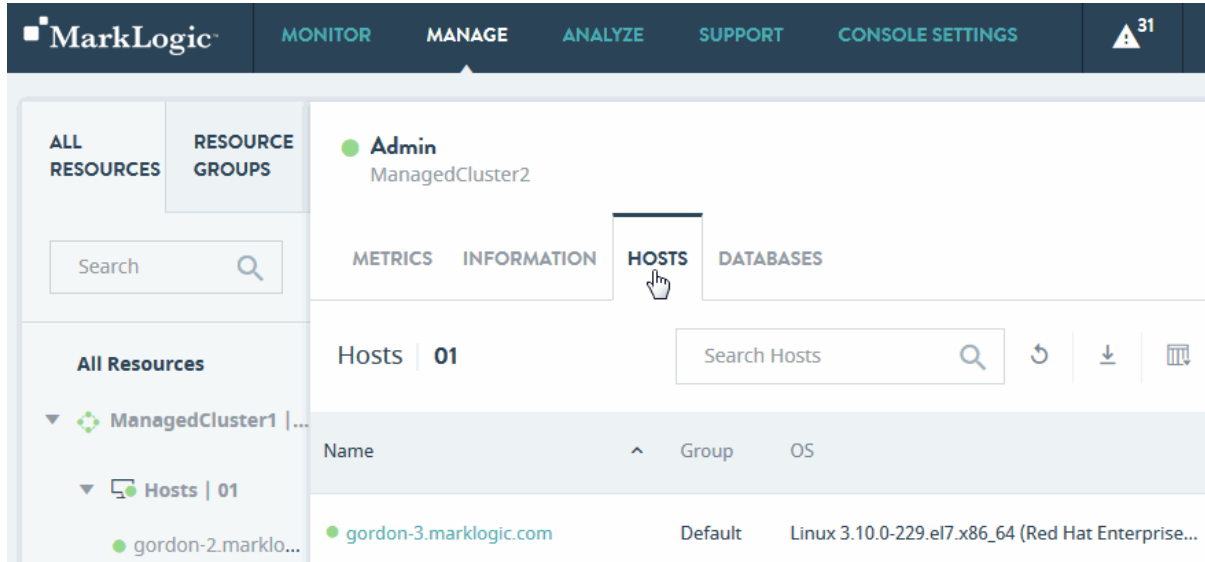
Field	Description
Output Byte Order Mark	The output options for the App Server.
Output Cdata Section Localname	<p>For details about controlling App Server output, see Controlling App Server Access, Output, and Errors in the <i>Application Developer's Guide</i> and Setting Output Options for an HTTP Server, Setting Output Options for an XDBC Server, Setting Output Options for an ODBC Server, and Setting Output Options for a WebDAV Server in the <i>Administrator's Guide</i>.</p> <p>For details on setting the serialization options in XQuery, see Declaring Options in the <i>XQuery and XSLT Reference Guide</i>. For XSLT output details, see the XSLT specification (http://www.w3.org/TR/xslt20#serialization).</p>
Output Cdata Section Namespace URI	
Output Doctype Public	
Output Doctype System	
Output Encoding	
Output Escape URI Attributes	
Output Include Content Type	
Output Include Default Attributes	
Output Indent	
Output Indent Tabs	
Output Indent Untyped	
Output Media Type	
Output Method	
Output Normalization Form	
Output Omit Xml Declaration	
Output Sgml Character Entities	
Output Standalone	
Output Undeclare Prefixes	
Output Version	

Field	Description
Port	The port number applications use to access this App Server.
Pre Commit Trigger Depth	The maximum depth (how many triggers can cause other triggers to fire, which in turn cause others to fire, and so on) for pre-commit triggers that are executed against this App Server. For more information on triggers, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .
Pre Commit Trigger Limit	The maximum number of pre-commit triggers this App Server can invoke for a single statement. For more information on triggers, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .
Privilege	The execute privilege required to access this App Server.
Profile Allow	Specifies whether to allow requests against this App Server to be profiled, using the MarkLogic Server profiling APIs. For details, see Profiling Requests to Evaluate Performance in the <i>Query Performance and Tuning</i> guide.
Request Timeout	The maximum number of seconds before a socket receives a timeout for the first request.
Rewrite Resolves Globally	Specifies whether to allow rewritten URLs to be resolved from the global <code>MarkLogic/Modules</code> directory or App Server root.
Root	<p>The directory in which programs executed against this App Server are stored. If the Modules field is set to a database, the root must be a directory URI in the specified modules database.</p> <p>If the Modules field is set to file system, the root directory is either a fully-qualified pathname or is relative to the directory in which MarkLogic Server is installed. For details on where the MarkLogic Server is installed by default on your platform, see Installing MarkLogic in the <i>Installation Guide</i>.</p>
Server Name	The name of this App Server.
Server Type	The App Server type (<code>http</code> , <code>odbc</code> , <code>xdbc</code> , or <code>WebDAV</code>).
Session Timeout	The maximum number of seconds before an inactive session times out.

Field	Description
SSL Allow SSLv3	Specifies whether to allow this App Server to communicate with applications by means of the SSLv3 security protocol.
SSL Allow TLS	Specifies whether to allow this App Server to communicate with applications by means of the TLSv1 security protocol.
SSL Ciphers	The SSL ciphers that may be used. For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
SSL Disable SSLv3	Specifies whether to prevent this App Server from communicating with applications by means of the SSLv3 security protocol.
SSL Disable TLSv1	Specifies whether to prevent this App Server from communicating with applications by means of the TLSv1 security protocol.
SSL Disable TLSv11	Specifies whether to prevent this App Server from communicating with applications by means of the TLSv11 security protocol.
SSL Disable TLSv12	Specifies whether to prevent this App Server from communicating with applications by means of the TLSv12 security protocol.
SSL Hostname	The hostname for the App Server SSL certificate. This is useful when many App Servers are running behind a load balancer. If not specified, each host will use a certificate for its own hostname. For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
SSL Require Client Certificate	Specifies whether to enable mutual authentication, where the client also holds a digital certificate that it sends to the server. Select which certificate authority is to be used to sign client certificates for the server. For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Static Expires	The number of seconds before an <code>expires</code> HTTP header is added for static content.
Threads	The maximum number of App Server threads.
URL Rewriter	The path to the script to run to rewrite URLs. For details, see Setting Up URL Rewriting for an HTTP App Server in the <i>Application Developer's Guide</i> .
WebDAV	Specifies whether this App Server is a WebDAV App Server.

5.4.3 App Server Hosts

Use the **HOSTS** tab to see the host that contains the selected App Server. This tab shows the host name as the first column, followed by other properties. Click on the host name to navigate to the manage view of the host.



The columns displayed in the App Server **HOSTS** tab are described in the following table.

Column	Description
Name	The hostname of the host.
Group	The name of the group that contains the host.
OS	The name and version of the operating system on which the host runs.
Server Version	The version of MarkLogic Server on the host.
Forests	The number of forests on the host.
Databases	The number of databases on the host.
App Servers	The number of App Servers on the host.
Disk Space (MB)	The amount of disk space (in MB) used on the host.
Uptime	The duration (Days Hrs:Min) the host has been available.
Maint. Mode	The host maintenance mode (normal or maintenance). For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .

Column	Description
Zone	The Amazon Web Services (AWS) zone in which the host resides, if applicable. For details, see <i>MarkLogic Server on Amazon Web Services (AWS) Guide</i> .

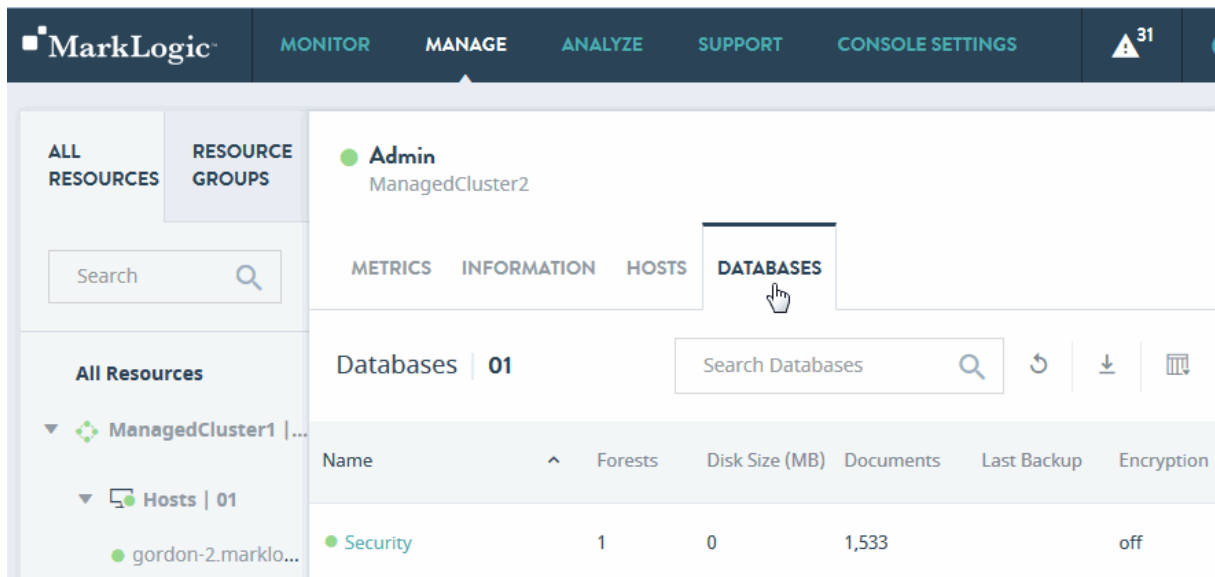
You can export data from the App Server Hosts tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the App Server Hosts table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

5.4.4 App Server Databases

Use the **DATABASES** tab to get a list of databases used by this App Server, with the database name as the first column, followed by other properties. Click on the database name to navigate to the manage view of the selected database.



The columns displayed in the App Server **DATABASES** tab are described in the following table.

Column	Description
Name	Name of the database.
Forests	The number of forests used by the database.
Disk Size (MB)	The amount of disk space used by the database forests, in megabytes.
Documents	The number of documents in the database.
Last Backup	The data-time of the last backup of the database. No value, if the database has never been backed up. For details on backing up a database, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Encryption	Specifies whether or not encryption at rest is enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .
HA	Specifies whether or not shared disk failover is enabled. For details, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled. For details, see the <i>Database Replication Guide</i> .
Replication Status	Specifies whether or not database replication is configured for the database.
Security DB	The name of the security database used by the database. For details, see Administering Security in the <i>Security Guide</i> .
Schemas DB	The name of the schemas database used by the database. For details, see Understanding and Defining Schemas in the <i>Administrator's Guide</i> .
Triggers DB	The name of the schemas database used by the database. For details, see Using Triggers to Spawn Actions in the <i>Application Developer's Guide</i> .

You can export data from the App Server **DATABASES** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the App Server DATABASES table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In the case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

6.0 ANALYZE View

The ANALYZE view presents a comprehensive set of detailed charts that allow to analyze utilization and performance of system resources in your enterprise, such as disks, CPU, memory, network, databases, and servers.

This chapter covers the following topics:

- [Configuring and Navigating the ANALYZE View](#)
- [Performance Charts by Resource](#)

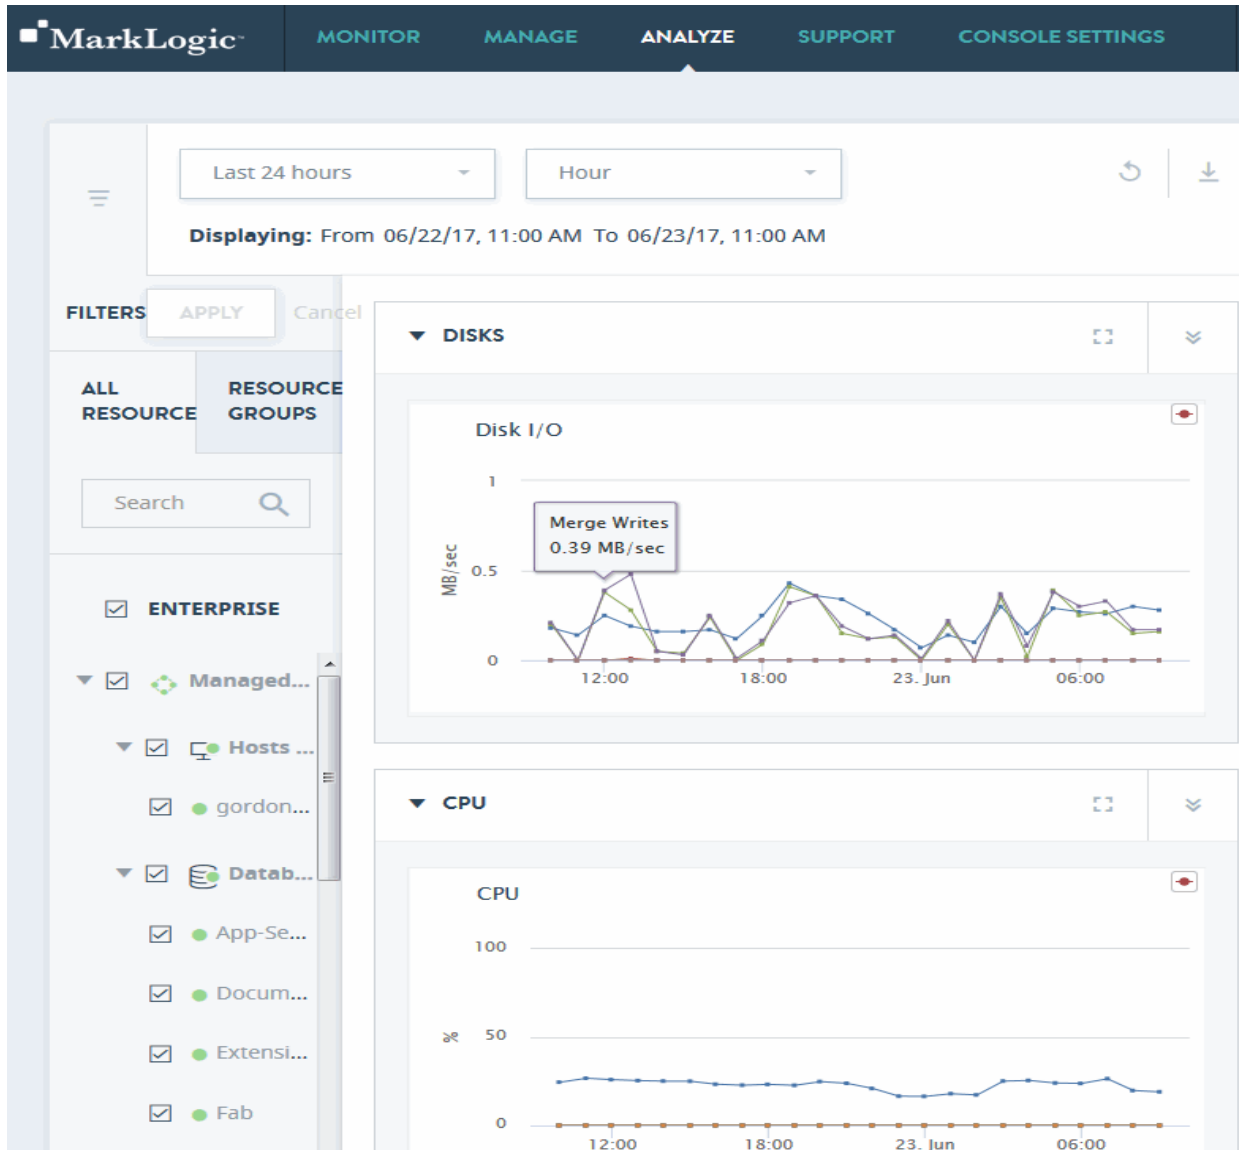
6.1 Configuring and Navigating the ANALYZE View

Use the ANALYZE view to examine key performance indicators of various system resources, such as disks, CPU, memory, network, databases, and servers.

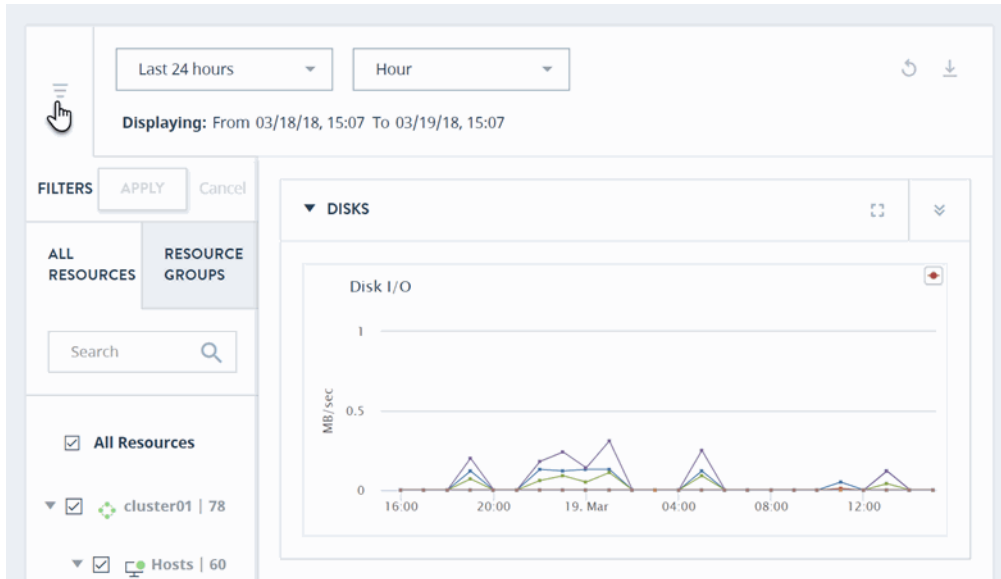
This section describes the general mechanisms for configuring and navigating the ANALYZE view.

Performance metrics are displayed in the central panel of the ANALYZE view.

Use the date picker to select the date/time range to inspect. Use the resources panel to select which resources to examine.




Toggle between a normal or expanded view of the data charts by selecting the chart-only icon beside the date filters.




Select which resources to examine by defining corresponding filters in the resources panel. You can use the search bar to restrict the list of resources to only those matching your specified keywords. These features are described in “Navigating and Filtering Ops Director Views” on page 81.

Drill down for greater detail by selecting the detail icon at the far right side of any top-level section of the Overview page.



The DISKS DETAIL page, for example, offers greater detail of the disk operational parameters.

To return to the Overview page from a DISKS DETAIL page, click on the detail icon at the upper right-hand section of the resource graph on the Detail page.



You may export data from the charts presented in the ANALYZE view by clicking the Export icon in the upper right corner.

As the result, a zip file with all exported metrics from the charts will be downloaded to your computer. The following rules apply:

- The zip file contains one CSV (Comma Separated Values) file for each plot (such as data series) across all charts in the view.
- The CSV files are grouped into folders by resource type: `Host`, `Database`, `Server`, and `Forest`.
- Each CSV file corresponds to one exported metric and has two columns: one for the timestamp and another for the chart data.

Metrics exported for the `Host` resource type are described in the following sections:

- [Disk Performance Data](#)
- [CPU Performance Data](#)
- [Memory Performance Data](#)
- [Network Performance Data](#)

Thus, the `Host` folder in the zip file will contain CSV files that correspond to these metrics.

Metrics exported for the `Database` resource type are described in the section [Database Performance Data](#).

Thus, the `Database` folder in the zip file will contain CSV files that correspond to these metrics.

Metrics exported for the `Server` resource type are described in the section [Server Performance Data](#). Thus, the `Server` folder in the zip file will contain CSV files that correspond to these metrics.

Metrics exported for the `Forest` resource type are described in the following sections:

- [Disk Performance Data](#),
- [CPU Performance Data](#),
- [Memory Performance Data](#),
- [Network Performance Data](#).

Thus, the `Forest` folder in the zip file will contain CSV files that correspond to these metrics.

You may export metrics also from the second-level pages in the ANALYZE view:

- If you click the Export icon from the DISKS DETAIL page, the resulting zip file will contain only `Host` and `Forest` folders.
- If you click the Export icon from the CPU DETAIL page, the resulting zip file will contain only `Host` and `Forest` folders.
- If you click the Export icon from the MEMORY DETAIL page, the resulting zip file will contain only `Host` and `Forest` folders.
- If you click the Export icon from the SERVERS DETAIL page, the resulting zip file will contain only `Server` folder.
- If you click the Export icon from the NETWORK DETAIL page, the resulting zip file will contain only `Host` and `Forest` folders.
- If you click the Export icon from the DATABASES DETAIL page, the resulting zip file will contain only `Database` folder.

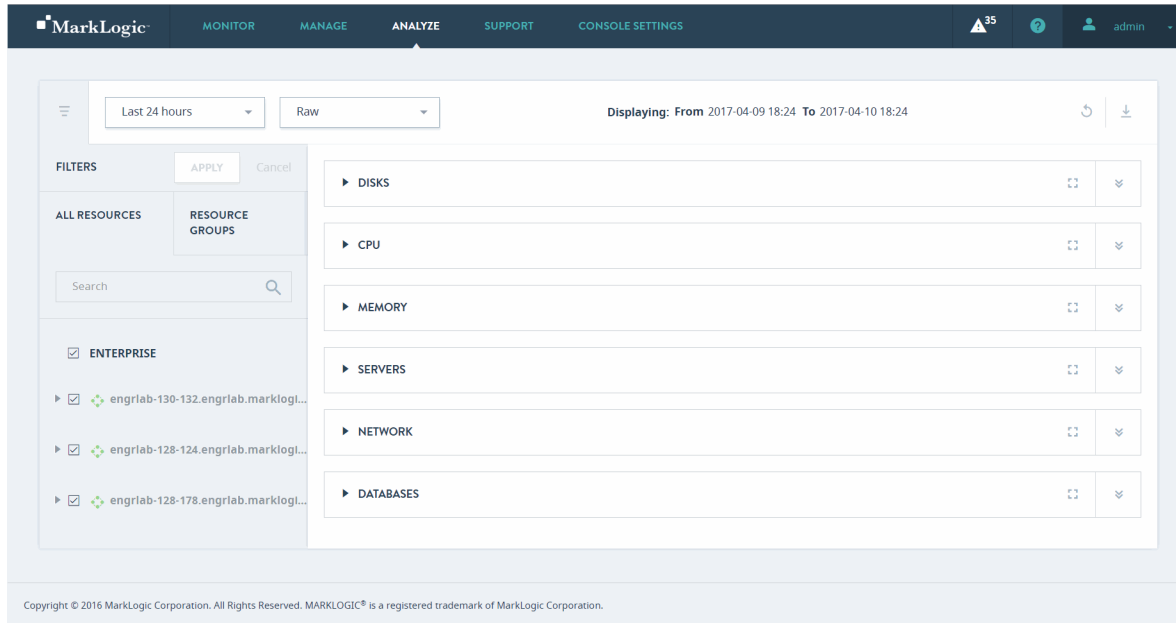
You may export metrics for selected resources and/or apply date filters:

- If you select some resources in the **ALL RESOURCES** tab of the left-hand navigation panel, click **Apply** to apply the new selection, and then click the Export icon, the resulting zip file will contain CSV files with content pertaining to the resources selected.
- If you select some resources in the **RESOURCE GROUPS** tab of the left-hand navigation panel, click **Apply** to apply the new selection, and then click the Export icon, the resulting zip file will contain CSV files with content pertaining to the resources selected.
- If you filter the ANALYZE view by selecting date/time range, click **Apply** to apply the new filtering, and then click the Export icon, the resulting zip file will contain CSV files with content pertaining to the resources that are visible with the applied filter.

You may then import the CSV files into other applications (such as Excel) for further processing or analysis.

6.2 Performance Charts by Resource

The ANALYZE view offers 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 Detail page, the lines represent the metric for each specific resource in the cluster.




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

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

6.2.1 Disk Performance Data

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




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 following table.

Metric	Description
Writes	The disk I/O performance (in MB/sec) during journal and save write operations. This is the sum of journal-write-rate, save-write-rate, and large-write-rate. For more information, see the <i>Query Performance and Tuning Guide</i> .
Query Traffic	The disk I/O performance (in MB/sec) during a query or queries. This is the sum of query-read-rate and large-read-rate. For more information, see the <i>Query Performance and Tuning Guide</i> .
Merge Reads	The disk I/O performance (in MB/sec) during a merge read operation. For more information on merging, see Understanding and Controlling Database Merges in the <i>Administrator's Guide</i> .
Merge Writes	The disk I/O performance (in MB/sec) during a merge write operation. For more information on merging, see Understanding and Controlling Database Merges in the <i>Administrator's Guide</i> .

Metric	Description
Backup Reads	The disk I/O read performance (in MB/sec) during a backup operation. For more information on database backup, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Backup Writes	The disk I/O write performance (in MB/sec) during a backup operation. For more information on database backup, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Restore Reads	The disk I/O read performance (in MB/sec) during a restore operation. For more information on database restore, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Restore Writes	The disk I/O read performance (in MB/sec) during a restore operation. For more information on database restore, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .

Click on the detail icon in the upper right-hand section of the DISKS section of the Overview page to view charts that present more detailed disk performance metrics.




The rate metrics displayed by the charts on the DISKS DETAIL page are described in the following table. For guidelines on how to interpret rate metrics, see “Assess MarkLogic Cluster Performance” on page 267.

Chart	Definition of Displayed Metric
Journal Write Rate	The moving average of data writes (in MB/sec) to the journal.
Save Write Rate	The moving average of data writes (in MB/sec) to in-memory stands.
Query Read Rate	The moving average of reading query data (in MB/sec) from disk
Merge Read Rate	The moving average of reading merge data (in MB/sec) from disk
Merge Write Rate	The moving average of writing data (in MB/sec) for merges
Backup Read Rate	The moving average of reading backup data (in MB/sec) to disk.
Backup Write Rate	The moving average of writing backup data (in MB/sec) to disk.
Restore Read Rate	The moving average of reading restore data (in MB/sec) from disk.
Restore Write Rate	The moving average of writing restore data (in MB/sec) from disk.
Large Binary Read Rate	The moving average of reading large documents (in MB/sec) from disk. For more information, see Working With Binary Documents in the <i>Application Developer's Guide</i> .
Large Binary Write Rate	The moving average of writing data for large documents (in MB/sec) to disk. For more information, see Working With Binary Documents in the <i>Application Developer's Guide</i> .

6.2.2 CPU Performance Data

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


Note: CPU metrics are not supported on the Mac OS X platform and are only partially supported on Windows.



Each performance metric in the CPU section of the Overview page is described in the following table.

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 detail icon to view graphs that present more detailed CPU performance metrics.



The charts on the CPU Detail page are described in the following table.

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.

6.2.3 Memory Performance Data

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

Note: CPU metrics are not supported on the Mac OS X platform and are only partially supported on Windows.



You can hover on a period point to view which CPU operation was taking place at that point in time. Each chart and associated performance metrics are described in the following table.

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


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 <code>/proc/vmstat</code>) for the hosts in pages/sec. • Page-Out Rate: The page-out rate (from Linux <code>/proc/vmstat</code>) for the hosts in pages/sec. • Swap-In Rate: The swap-in rate (from Linux <code>/proc/vmstat</code>) for the hosts in pages/sec. • Swap-Out Rate: The swap-out rate (from Linux <code>/proc/vmstat</code>) for the hosts in pages/sec.

Click on the detail icon to view graphs that present more detailed CPU performance metrics. The charts on the Memory Detail page are described in the following table. The displayed metrics are drawn from `/proc/vmstat`.

Chart	Description
RSS	The amount of MB of Process Resident Size (RSS) for each host in the cluster.
Anon	The amount of MB of Process Anonymous Memory for each host in the cluster.
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.

6.2.4 Server Performance Data

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




The Servers Overview page displays the charts described in the following table.

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.

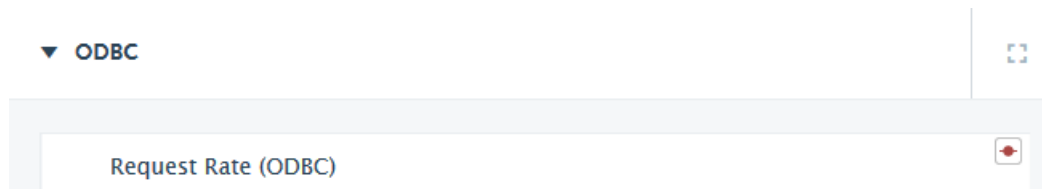
With the exception of the Task Server Queue Size chart, which only displays the queue size for the one task server, the color-coded metrics for the server charts are as shown in the following table.

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 detail icon 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 following table.



The server type (for example, ODBC) is shown in the upper right-hand section of each server type group.




The following repeating pattern of detailed charts are displayed for each of HTTP, XDBC, ODBC, Task, and WebDAV App Servers:

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 Hit Rate	The number of times queries could use the expanded tree cache on each App Server.
Expanded Tree Cache Miss Rate	The number of times queries could not use the expanded tree cache on each App Server.
Send Rate	The rate (in MB/sec) at which this App Server sends data.
Receive Rate	The rate (in MB/sec) at which this App Server receives data.
Queue Size (Task Server Only)	The number of tasks in the Task Server queue on each host.

6.2.5 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 Network section of 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 indicates excessively high during loads, running the MarkLogic Content Pump (m1cp) 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 Network section of the Overview page displays a chart with the metrics described in the following table.

Metric	Description
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> .

Click on the detail icon 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 following table.


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 <code>foreign-xdqp-client-receive-rate</code> and <code>foreign-xdqp-server-receive-rate</code> .
XDQP Write Rate	The amount of data (in MB/sec) written over XDQP by each host in the cluster. This is the sum of <code>foreign-xdqp-client-send-rate</code> and <code>foreign-xdqp-server-send-rate</code> .
XDQP Read Load	The execution time (in seconds) of read requests by each host in the cluster. This is the sum of <code>xdqp-client-receive-load</code> and <code>xdqp-server-receive-load</code> .
XDQP Write Load	The execution time (in seconds) of write requests by each host in the cluster. This is the sum of <code>xdqp-client-send-load</code> and <code>xdqp-server-send-load</code> .

Chart	Description
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 <code>foreign-xdqp-client-receive-rate</code> and <code>foreign-xdqp-server-receive-rate</code> .
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 <code>foreign-xdqp-client-send-rate</code> and <code>foreign-xdqp-server-send-rate</code> .
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 <code>foreign-xdqp-client-receive-load</code> and <code>foreign-xdqp-server-receive-load</code> .
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 <code>foreign-xdqp-client-send-load</code> and <code>foreign-xdqp-server-send-load</code> .

6.2.6 Database Performance Data

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.

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



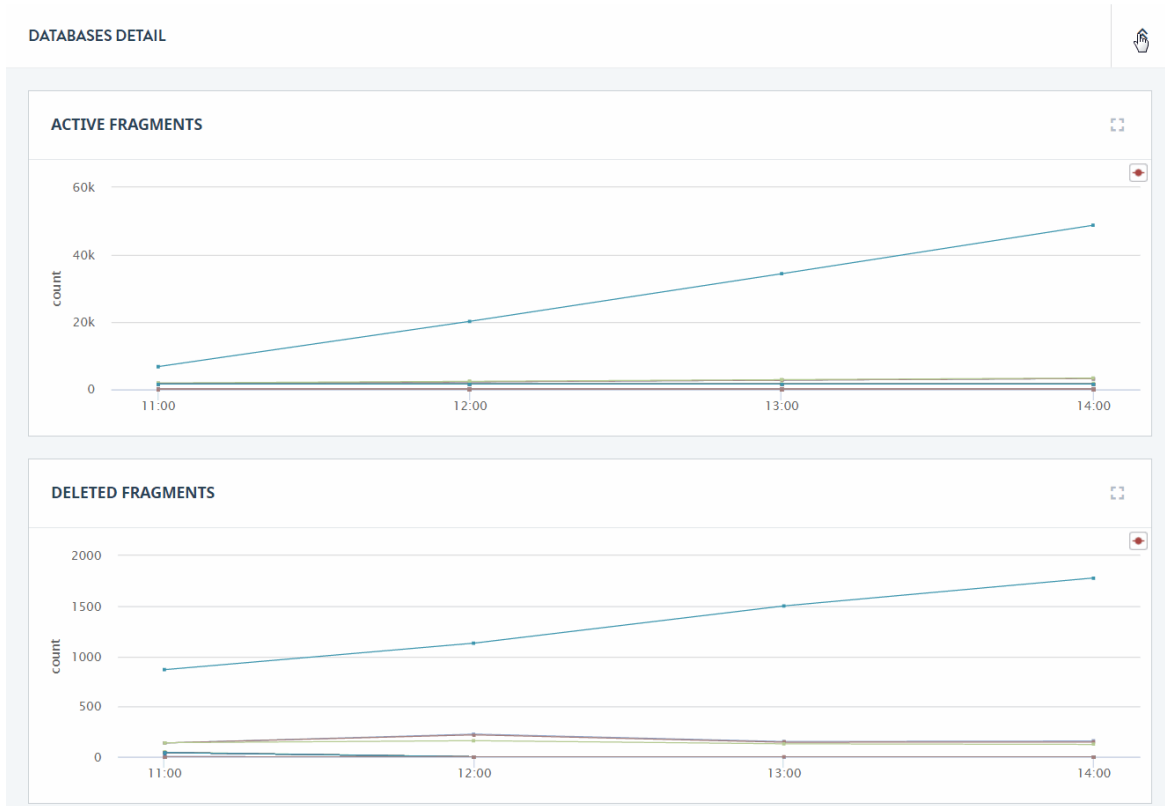
The following table describes the lines displayed in the DATABASES section of the Overview page.

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 number of fragments available to queries. • Deleted Fragments: The number of fragments to be deleted during the next merge operation.

Chart	Description
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 disk space used by the data in the forest stands. This data is subject to periodic merges. • Fast Data Size: The amount of data in the forest Fast Data Directories. 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 <i>Query Performance and Tuning Guide</i>. • Large Data Size: The amount of data in the forest Large Data Directories. 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 <i>Application Developer's Guide</i>.
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.

Chart	Description
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.
List Cache Hits/Misses	<p>The displayed lines are:</p> <ul style="list-style-type: none"> • List Cache hits/sec • List Cache misses/sec
Compressed Tree Cache Hits/Misses	<p>The displayed lines are:</p> <ul style="list-style-type: none"> • Compressed Tree Cache hits/sec • Compressed Tree Cache misses/sec
Triple Cache Hits/Misses	<p>The displayed lines are:</p> <ul style="list-style-type: none"> • Triple Cache hits/sec • Triple Cache misses/sec
Triple Value Cache Hits/Misses	<p>The displayed lines are:</p> <ul style="list-style-type: none"> • Triple Value Cache hits/sec • Triple Value Cache misses/sec

Click on the detail icon to view graphs that present more detailed performance metrics for each database.



The charts displayed on the Databases Detail page are described in the following table. 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. For more information on Fast Data, see Fast Data Directory on Forests in the <i>Query Performance and Tuning Guide</i> .

Chart	Description
Large Data Size	The amount of data in the large data directories of the forests attached to each database. For more information on Large Data, see Working With Binary Documents in the <i>Application Developer's Guide</i> .
Read Lock Rate	The rate of read lock acquisitions, summed across all forests on the host.
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	Time threads spent waiting for read locks in proportion to the elapsed time, summed across all forests on the host.
Write Lock Wait Load	The aggregate 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.
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.
List Cache Hit Rate	The number of times per second that queries use (Hit) the expanded tree cache on each App Server.
List Cache Miss Rate	The number of times per second that queries could not use (Miss) the expanded tree cache on each App Server.
Compressed Tree Cache Hit Rate	The number of times per second that queries could use (Hit) the compressed tree cache on each App Server. For details, see Effect of External Binaries on E-node Compressed Tree Cache Size in the <i>Application Developer's Guide</i> .

Chart	Description
Compressed Tree Cache Miss Rate	The number of times per second that queries could not use (Miss) the compressed tree cache on each App Server. For details, see Effect of External Binaries on E-node Compressed Tree Cache Size in the <i>Application Developer's Guide</i> .
Triple Cache Hit Rate	The number of times per second that queries could use (Hit) the triple cache on each App Server. For details, see Triple Cache and Triple Value Cache in the <i>Semantics Developer's Guide</i> .
Triple Cache Miss Rate	The number of times per second that queries could not use (Miss) the triple cache on each App Server. For details, see Triple Cache and Triple Value Cache in the <i>Semantics Developer's Guide</i> .
Triple Value Cache Hit Rate	The number of times per second that queries could use (Hit) the triple value cache on each App Server. For details, see Triple Cache and Triple Value Cache in the <i>Semantics Developer's Guide</i> .
Triple Value Cache Miss Rate	The number of times per second that queries could not use (Miss) the triple value cache on each App Server. For details, see Triple Cache and Triple Value Cache in the <i>Semantics Developer's Guide</i> .
Reindex Refragment Rate	The average rate of the database reindex/refragment process. For more information, see Reindexing a Database in the <i>Administrator's Guide</i> .
Rebalance Rate	The average rate of the database rebalancing process. For details, see Database Rebalancing in the <i>Administrator's Guide</i> .

7.0 CONSOLE SETTINGS View

The Console Settings view allows you to configure role-based access control to resources, manage user accounts, manage licensing, and define telemetry settings.

This chapter covers the following topics:

- [Role Based Access Control \(RBAC\) Settings](#)
- [Resource Groups](#)
- [License Information](#)
- [Managed Clusters](#)
- [Configuring Email Notifications](#)

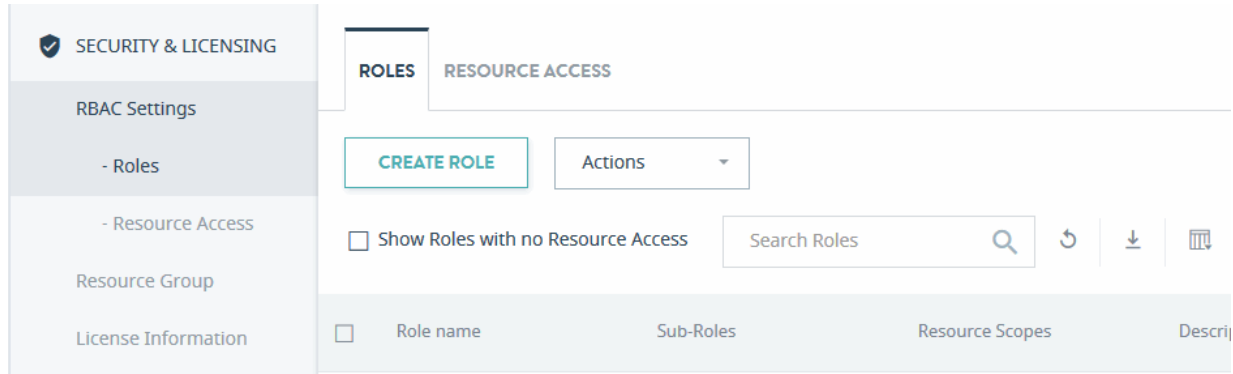
7.1 Role Based Access Control (RBAC) Settings

Use RBAC (Role Based Access Control) Settings to define new roles that assign sub-roles to Resource Groups to control which users have access to the resources defined by those Resource Groups (Resource Scope). The roles you create in this view will be accessible in the Admin Interface.

When assigning Resource Groups to a role, the resources in those groups and inherited resources will be accessible to users assigned that role. For details how resource groups define access to resources, see “Access Inheritance in Resource Groups” on page 14.

For example, to see the cluster resources, you must create a Resource Group for this cluster and assign it to a role. A practical configuration would be to restrict access of a particular user to one cluster, which would imply access to that cluster’s hosts, application servers, databases, and forests through the access inheritance mechanism in resource groups.

Note: If you do not have permission to see a resource, that resource will be displayed as blank or, if the resource type is presented as a count, it will be displayed as 0. Additionally, if you do not have permission to see a resource that is presented in chart form, you will see charts, but those charts will have no data (lines) for the prohibited resource.



This section covers the following topics:

- [Roles Tab](#)
- [Resource Access Tab](#)
- [Creating a Resource Group and Assigning it to a Role](#)
- [Editing or Deleting a Role](#)

7.1.1 Roles Tab

The **ROLES** tab lists the available roles. The columns displayed for a role are described in the following table.

Column	Description
Role name	The name of the Ops Director role.
Sub-Roles (optional)	The MarkLogic roles to be assigned to this role. For details, see Role-Based Security Model in the <i>Security Guide</i> and Appendix C: Pre-defined Roles in the <i>Administrator's Guide</i> . Note: Do not assign opsdir-admin as a sub-role, as opsdir-admin has access to view all of the resources in Ops Director, which defeats the purpose of RBAC.
Resource Scopes (optional)	The resource group(s) to which this role controls access.
Description (optional)	The description of this role.

You may export data from the ROLES tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Roles table in the UI.

- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such As Excel) for further processing or analysis.

7.1.2 Resource Access Tab

The **RESOURCE ACCESS** tab lists the resource groups and their assigned roles. The columns displayed are described in the following table.

Column	Description
Resource Scope	The name of the resource group.
Roles	The roles assigned to the resource group.

You may export data from the Resource Access tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Resource Access table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

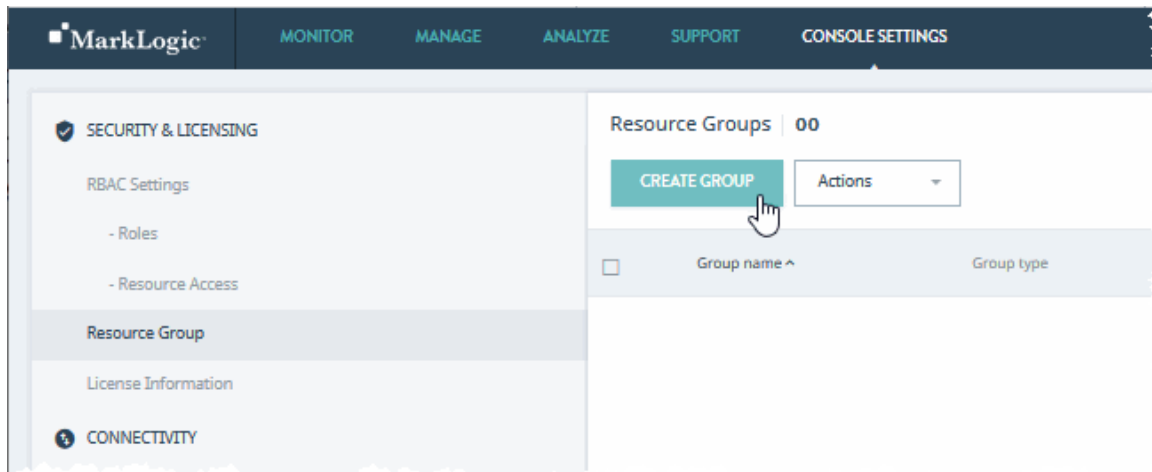
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.1.3 Creating a Resource Group and Assigning it to a Role

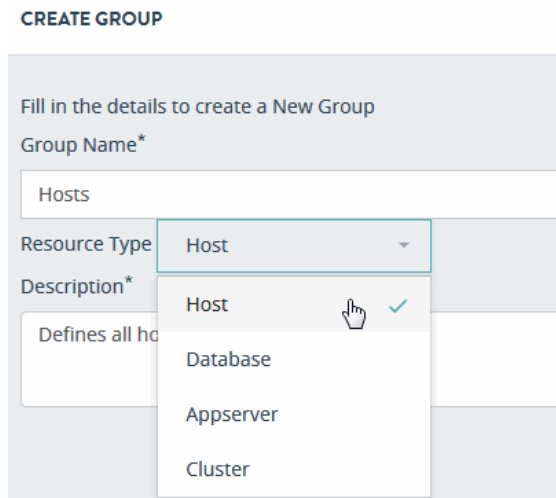
The following procedure creates a Resource Group that represents all of the hosts in the Managed Clusters and then restricts access to monitor those hosts to only users with the `opsdir-user` role.

1. In Ops Director, select the **CONSOLE SETTINGS** view.
2. Select **Resource Group** in the left menu.

3. In the right pane, click **CREATE GROUP**.



4. In the CREATE GROUP dialog box, enter the Group Name, Description, and select **Host** from the Resource Type menu.



5. After creating the group, select the group. In the Resource Groups page for the selected group, scroll down to the Total Resources section (which will likely be out of sight until

you scroll down) and select the check boxes to the left of each host in the cluster that you want to belong to this group.

Total Resources | 03

Search Hosts

<input type="checkbox"/>	Name	Cluster	Group	OS
<input checked="" type="checkbox"/>	gordon-1.marklogic.com	OpsDirectorCluster	Default	Linux
<input checked="" type="checkbox"/>	gordon-2.marklogic.com	ManagedCluster2	Default	Linux
<input type="checkbox"/>	gordon-3.marklogic.com	ManagedCluster3	Default	Linux

6. Click **ASSIGN** to assign the selected resources to the group.

Showing 1 to 3 of 3 entries

10 Per page

1

ASSIGN

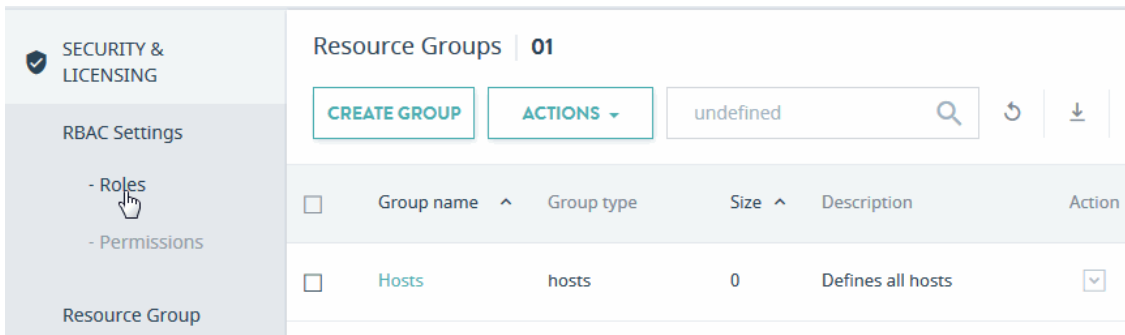
7. The Resource Group lists the assigned resources.

Assigned Resources | 03

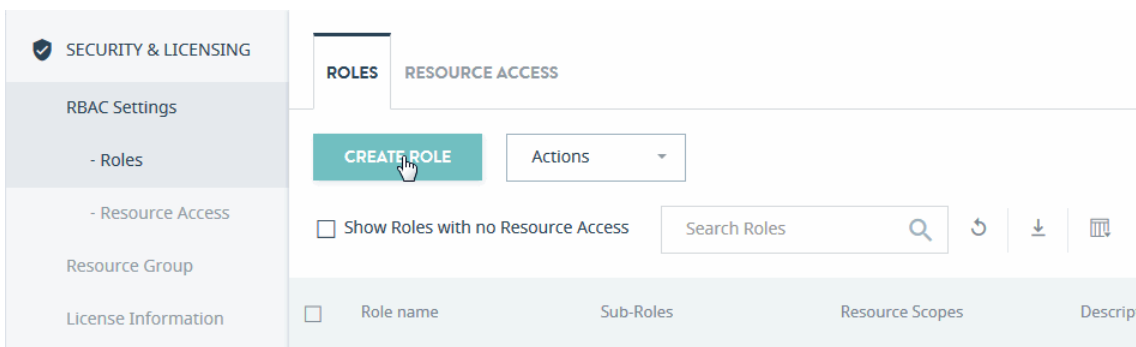
Search Hosts

<input type="checkbox"/>	Name	Cluster	Group	OS
<input type="checkbox"/>	gordon-1.marklogic.com	OpsDirectorCluster	Default	Linux
<input type="checkbox"/>	gordon-2.marklogic.com	ManagedCluster2	Default	Linux
<input type="checkbox"/>	gordon-3.marklogic.com	ManagedCluster3	Default	Linux

8. In the left menu, select **Roles** from under RBAC Settings.



9. In the ROLES tab, select **CREATE ROLE**.



10. In the Create Role dialog box, enter the Role Name (`HostAccess`, in this example) and Description. Select **opmdir-user** from the Sub-Roles menu.

CREATE ROLE

Role Name*

Description

Sub-Roles

Resource Scope

- opmdir-license-admin
- opmdir-user**

Select **Hosts** from the Resource Scope menu. Click **ok**.

CREATE ROLE

Role Name*

Description*

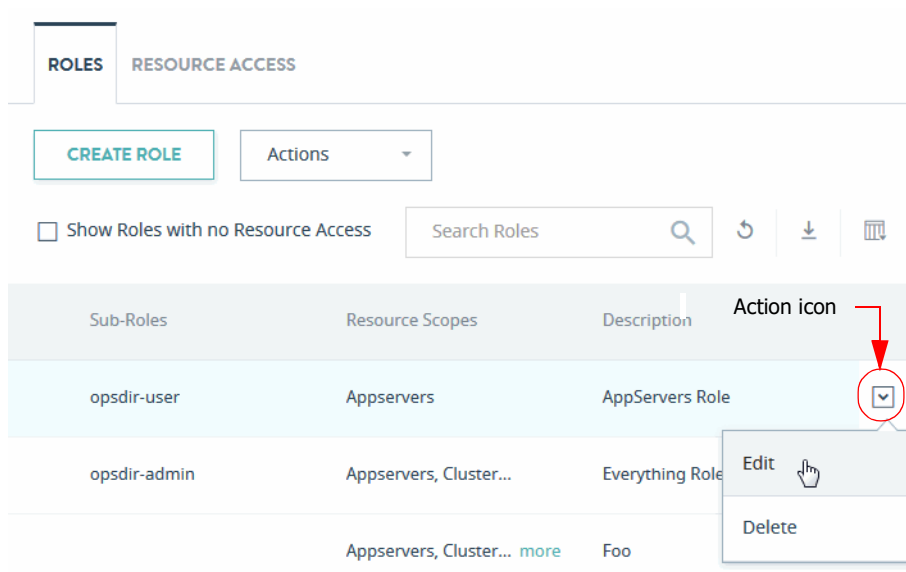
Sub-Roles

Resource Scope

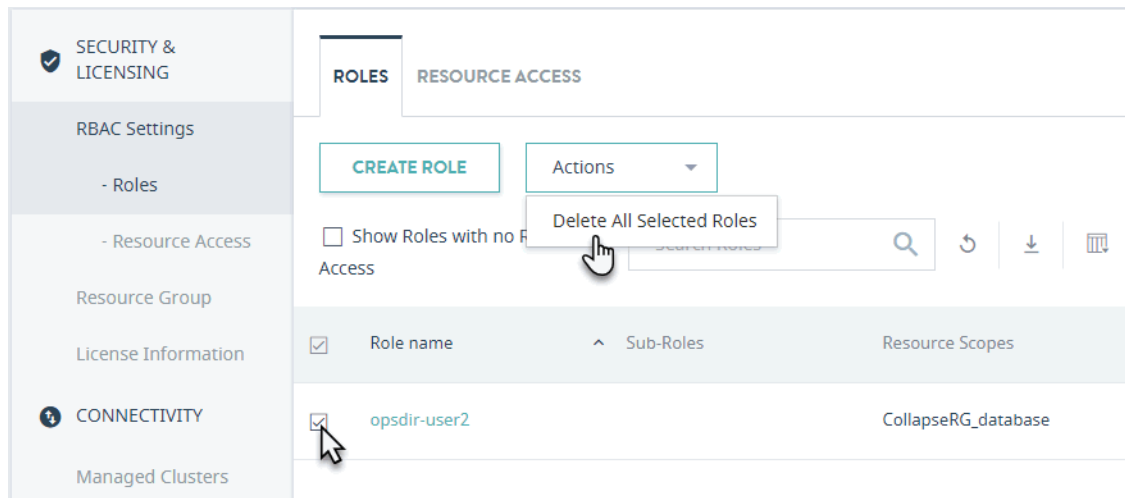
- Hosts**

7.1.4 Editing or Deleting a Role

To edit or delete a role, click the Action icon next to that role and select the desired action.



To delete multiple roles, select the checkboxes next to all roles you want to delete, then select **Delete All Selected Roles** in the ACTIONS menu.

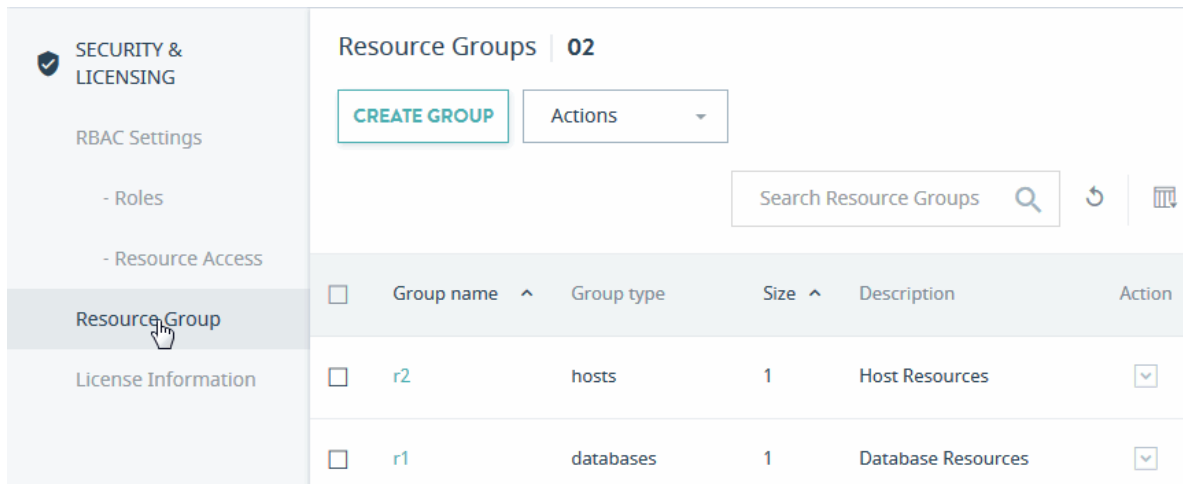


Note: The **Delete All Selected Roles** option is enabled when at least one role is selected.

7.2 Resource Groups

You may want to establish roles and privileges at a finer and more ad hoc granularity than is provided by the pre-defined MarkLogic roles. It is likely that roles defined within the enterprise are fairly coarse-grained and that changing roles (in an external LDAP server, for example), may be considered too “heavyweight” for ad hoc groupings.

Resource Groups define sets of resources to which you can assign specific roles to customize user access to those resources. For details on how resource groups define access to resources, see “Access Inheritance in Resource Groups” on page 14.



The columns displayed are described in the following table.

Column	Description
Group Name	The name of the resource group.
Group type	The type of resources in the group (Hosts, Databases, App Servers, Clusters).
Size	The number of resources in the resource group.
Description	The description of the resource group.
Action	The action to take on the resource group (Edit or Delete).

You may export data from the Resource Groups tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Resource Groups table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.

- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

Note: The Resource Groups CSV file, in addition to the columns from the Resource Groups table in the UI, has one additional column: `Resource Id`. This column contains comma-separated list of identifiers of all resources in this group.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

This section covers the following topics:

- [Creating a Resource Group](#)
- [Editing or Deleting a Resource Group](#)
- [Resource Group Views](#)

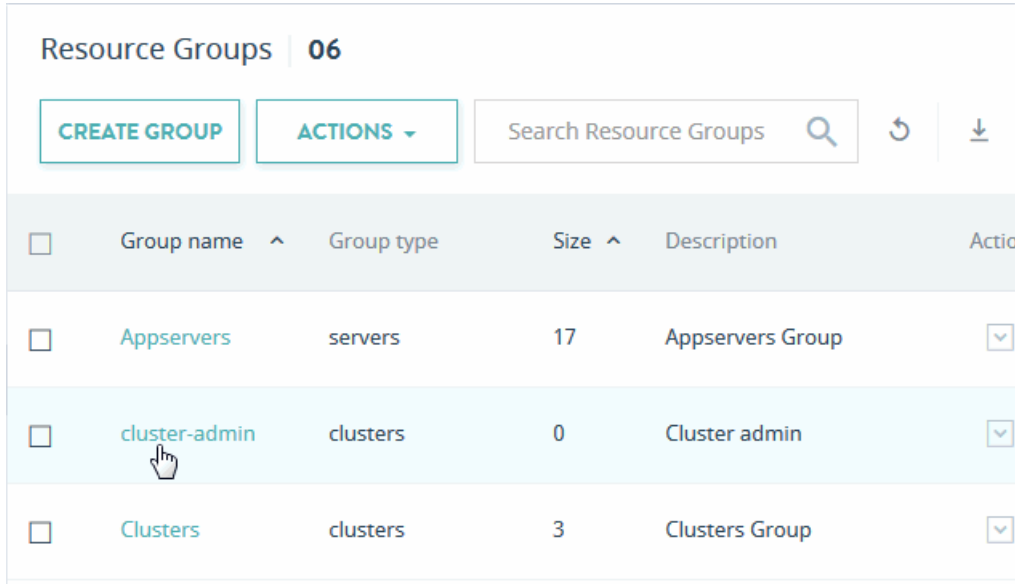
7.2.1 Creating a Resource Group

Do the following to create a Resource Group.

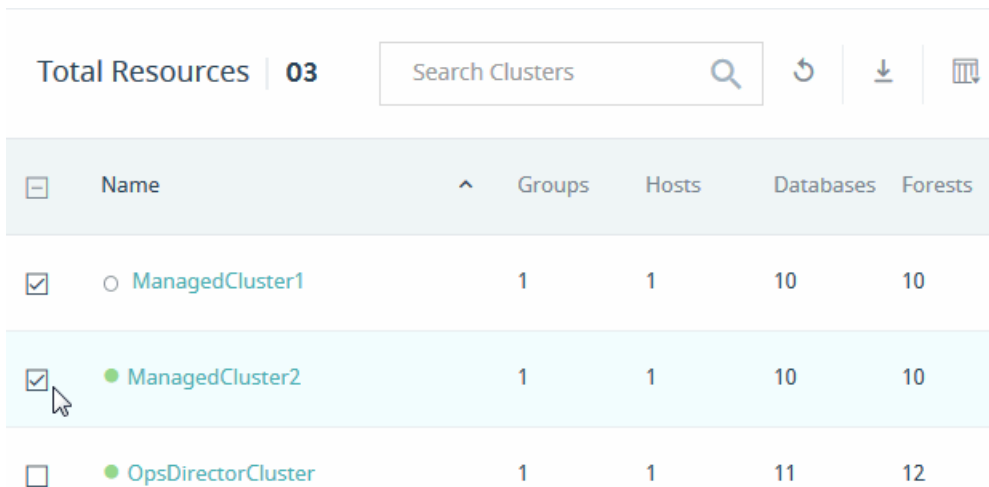
1. Click **CREATE GROUP**.
2. In the pop-up window, enter a Group Name, select a Resource Type (Host, Database, Appserver, or Cluster), and enter a description for the resource group.

The screenshot shows a 'CREATE GROUP' dialog box. The 'Group Name*' field contains 'cluster-admin'. The 'Resource Type' dropdown is open, showing options: Host, Database, Appserver, and Cluster. The 'Cluster' option is selected, indicated by a checkmark and a mouse cursor. The 'Description' field is empty. At the bottom, there are 'SAVE' and 'Cancel' buttons.

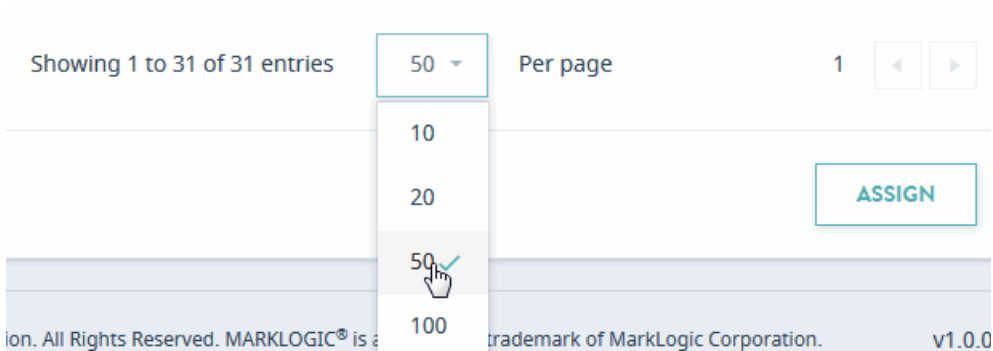
3. Click `save`. The new Resource Group is added to the list of Resource Groups.
4. By default, no resources are included in the Resource Group. To include resources, click on the name of the newly created Resource Group.



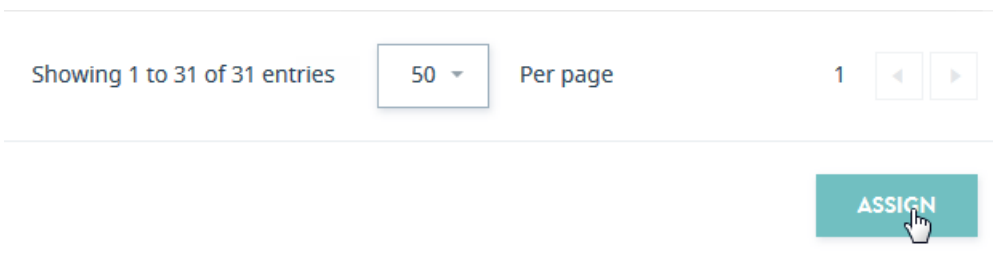
5. Scroll down to the Total Resources section and select the resource to be included in the Resource Group. This view will differ for each type of Resource Group as described in “Resource Group Views” on page 227.



By default, you can view 10 resources per page. You can adjust how many resources to view in the Resource Group page by changing the number in the menu at the bottom of the page.

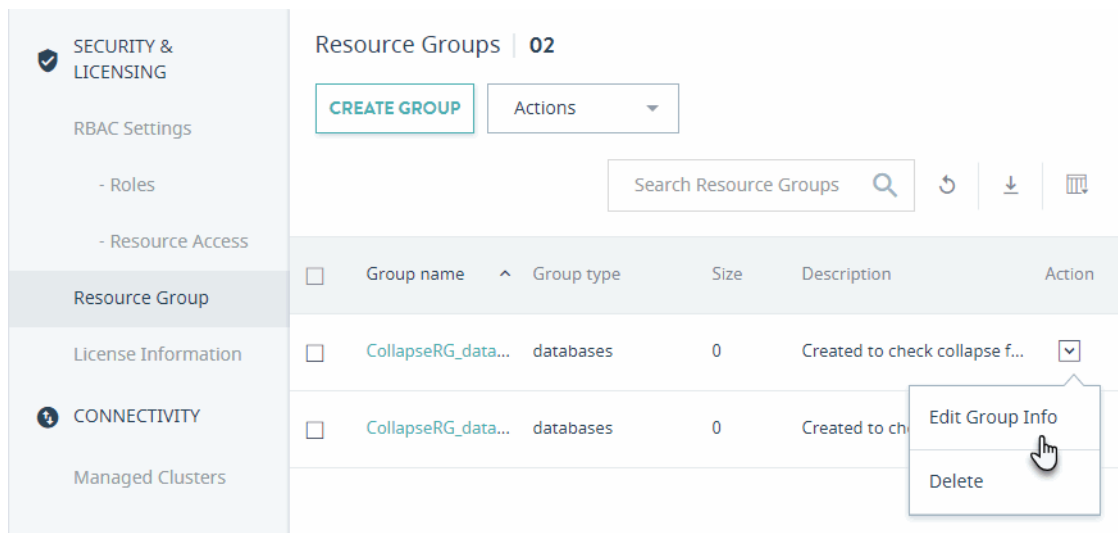


- When you have finished selecting resources for the Resource Group, click **ASSIGN** at the bottom of the page.

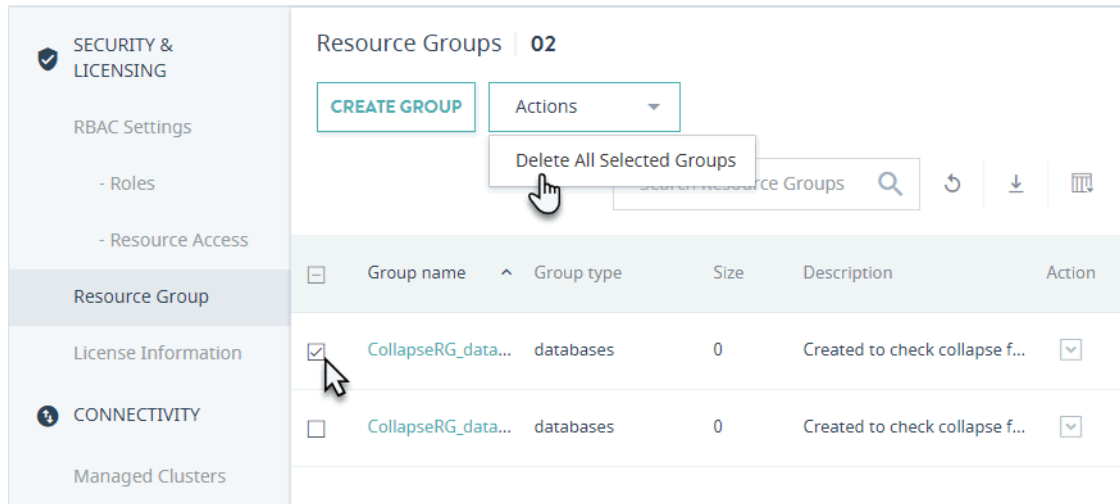


7.2.2 Editing or Deleting a Resource Group

To edit or delete a resource group, click the Action icon next to that resource group and select the desired action.



To delete multiple resource groups, mark checkboxes next to all resource groups you want to delete and select **Delete All Selected Groups** in the Actions menu.



Note: The **Delete All Selected Groups** option is enabled when at least one resource group is selected.

7.2.3 Resource Group Views

Click on a resource group to display the assigned and unassigned resources, as well as assign and deassign resources. The contents of each type of resource group are described in the following sections:

- [Host Groups](#)
- [Database Groups](#)
- [App Server Groups](#)
- [Cluster Groups](#)

7.2.3.1 Host Groups

The columns displayed for a host group are described in the following table. These settings are described in the [Hosts](#) chapter in the *Administrator's Guide*.

Column	Description
Name	The hostname of the host.
Cluster	The name of the cluster on which the host resides.
Group	The name of the group that contains the host.

Column	Description
OS	The name and version of the operating system on which the host runs.
Server Version	The version of MarkLogic Server running on the host.
Forests	The number of forests configured for the host.
Databases	The number of databases configured for the host.
App Servers	The number of App Servers configured for the host.
Disk Space	The amount of disk space (in MB) used on the host.
Uptime	The duration (Days Hrs:Min) the host has been available.
Maint. Mode	The host maintenance mode (normal or maintenance). For details, see Rolling Upgrades in the <i>Administrator's Guide</i> .
Zone	The Amazon Web Services (AWS) zone in which the host resides, if applicable.

You may export data from the Host Groups tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Host Groups table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.2.3.2 Database Groups

The columns displayed for a database group are described in the following table. These settings are described in the [Databases](#) chapter in the *Administrator's Guide*.

Column	Description
Name	The name of the database.
Cluster	The name of the cluster on which the database resides.

Column	Description
Forests	The number of forests configured for the database.
Disk Size (MB)	The amount of disk space used by the database forests, in megabytes.
Documents	The number of documents in the database.
Last Backup	The data-time of the last backup of the database. No value, if the database has never been backed up. For details on backing up a database, see Backing Up and Restoring a Database in the <i>Administrator's Guide</i> .
Encryption	Specifies whether or not encryption at rest should be enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .
HA	Specifies whether or not shared disk failover is enabled. For details, see High Availability of Data Nodes With Failover in the <i>Scalability, Availability, and Failover Guide</i> .
Replication	Specifies whether or not database replication is enabled (On/Off). For details, see the <i>Database Replication Guide</i> .
Security DB	The name of the security database used by the database.
Schemas DB	The name of the schema database used by the database.
Triggers DB	The name of the triggers database used by the database.

You may export data from the Database Groups tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Database Groups table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.2.3.3 App Server Groups

The columns displayed for an Appserver group are described in the following table. These settings are described in the [HTTP Servers](#), [ODBC Servers](#), [XDBC Servers](#), and [WebDAV Servers](#) chapters in the *Administrator's Guide*.

Column	Description
Name	The name of the App Server.
Cluster	The name of the cluster on which the App Server resides.
Type	The App Server Type (HTTP, ODBC, XDBC, WebDAV).
Database	The content database used by the App Server.
Port	The App Server port number.
SSL	Whether the App Server has SSL enabled (Yes) or disabled (No). For details, see Configuring SSL on App Servers in the <i>Security Guide</i> .
Group	The name of the group that contains the App Server.
Modules DB+Root	The name of the modules database, or if file system, the root directory.
Security	The type of security (internal or external).

You may export data from the App Server Groups tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the App Server Groups table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.2.3.4 Cluster Groups

The columns displayed for a cluster group are described in the following table. These settings are described in the [Clusters](#) chapter in the *Administrator's Guide*.

Column	Description
Name	The name of the cluster.
Groups	The number of groups in the cluster.
Hosts	The number of hosts in the cluster.
Databases	The number of databases in the cluster.
Forests	The number of forests in the cluster.
App Server	The number of App Servers in the cluster.
Server Version	The version of MarkLogic Server running on the hosts in the cluster.
OS	The name and version of the operating system on which the host runs.
Uptime	The duration (Days Hrs:Min) the cluster has been available.
Encryption	Specifies whether or not encryption at rest should be enabled for the database. For details, see Encryption at Rest in the <i>Security Guide</i> .

You may export data from the Cluster Groups tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Cluster Groups table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.3 License Information

Click **License Information** for a summary of managed hosts running under one or more MarkLogic license editions, with a breakdown of licensed cores, used cores, and the operating system platforms on which MarkLogic is running.

License Edition	Licensed Cores	Used Cores	Platform
Essential Enterprise	30912	7020	linux
Enterprise Edition	256	16	linux

The displayed columns are described in the following table.

Column	Description
License Edition	The type of MarkLogic License. For details, see Pricing and Licensing on the MarkLogic website.
Licensed Cores	The number of licensed cores. For more information, see Scalability Considerations in MarkLogic Server in <i>Scalability, Availability, and Forest-Level Failover</i> .
Used Cores	The number of used cores. For more information, see Scalability Considerations in MarkLogic Server in <i>Scalability, Availability, and Forest-Level Failover</i> .
Platform	The host operating system. See Supported Platforms in the <i>Release Notes</i> .
Environment	The type of environment, such as production or test.

You may export data from the License Information tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the License Information table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

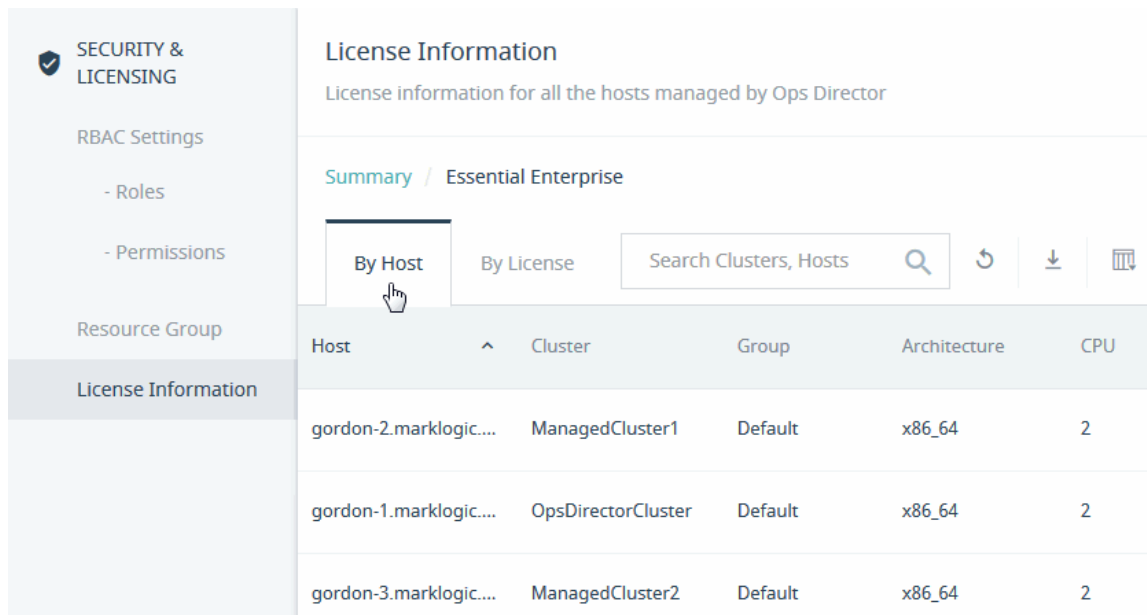
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

This section covers the following topics:

- [License Information By Host](#)
- [License Information By License](#)

7.3.1 License Information By Host

Select a specific MarkLogic License Edition to view details, broken down by host or by license edition, such as cluster name, group membership, processor architecture, and the number of CPUs, cores, and running threads.



The displayed columns are described in the following table.

Column	Description
Host	The list of licensed hosts in your enterprise.
Cluster	The host cluster.
Environment	The MarkLogic environment. Typically, Development or Production.
Group	The host group.
Architecture	The type of CPU hardware on which the host is running.

Column	Description
CPU	The number of CPUs configured on the host hardware.
Cores	The number of cores configured on the host hardware.
Threads	The number of threads used by the host.
Licensed CPUs	The number of licensed CPUs for the host.
Licensed Cores	The number of licensed cores for the host.
Options	Your licensed options. For details, see Displaying License Options in the <i>Administrator's Guide</i> and Pricing and Licensing on the MarkLogic website.
Expiration	The license expiration date.
License Key	The license key. For details, see Entering a License Key in the <i>Installation Guide</i> .
Licensee	The name of the person or organization that holds the license.

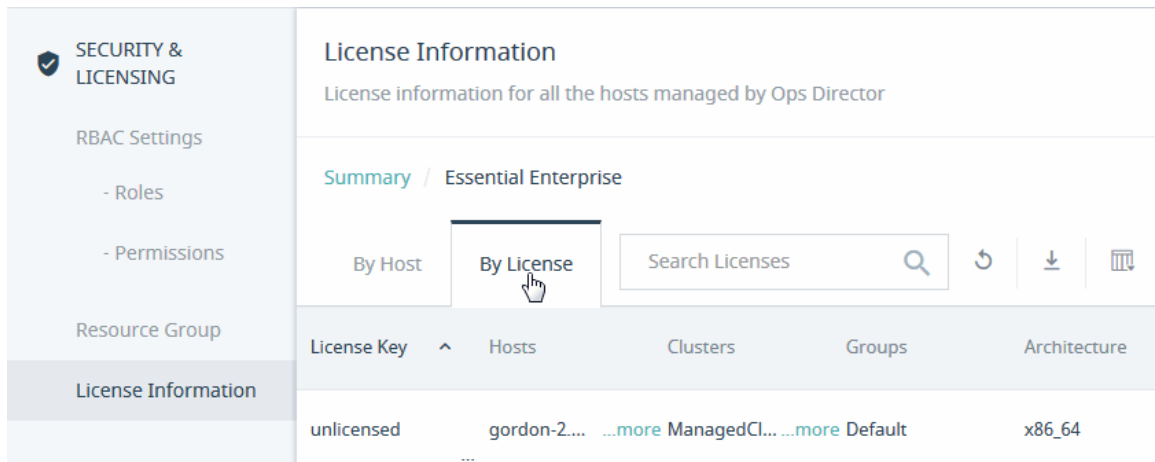
You may export data from the License Information by Host tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the License Information by Host table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.3.2 License Information By License

Click the **By License** tab to view the license information by license key.



The displayed columns are described in the following table.

Column	Description
Licensee	The name of the person or organization that holds the license.
Hosts	The hosts in your enterprise.
Clusters	The clusters in your enterprise.
Environment	The MarkLogic environment. Typically, Development or Production.
Groups	The groups in your enterprise.
Architecture	The type(s) of CPU hardware used by your enterprise.
CPU	The number of CPUs in your enterprise.
Cores	The number of cores in your enterprise.
Threads	The number of threads used by the enterprise.
Licensed CPUs	The number of licensed CPUs for the enterprise.
Licensed Cores	The number of licensed cores for the enterprise.
Options	Your licensed options. For details, see Pricing and Licensing on the MarkLogic website.
Expiration	The license expiration date.

Column	Description
License Key	The license key. For details, see Entering a License Key in the <i>Installation Guide</i> .

You may export data from the License Information by License tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the License Information by License table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.4 Managed Clusters

Under CONNECTIVITY, select **Managed Clusters** to view the list of clusters managed by Ops Director and remove clusters that are currently in the Unknown state.

This section covers the following topics:

- [Viewing and Filtering the List of Managed Clusters](#)
- [Removing Unknown Managed Clusters from the List](#)
- [Reconnecting a Managed Cluster to Ops Director](#)

7.4.1 Viewing and Filtering the List of Managed Clusters

Use the Managed Clusters page to view the list of all clusters currently managed by Ops Director, along with their health status.

The screenshot shows the MarkLogic Console Settings View. The top navigation bar includes 'MarkLogic', 'MONITOR', 'MANAGE', 'ANALYZE', 'SUPPORT', and 'CONSOLE SETTINGS'. The left sidebar has sections for 'SECURITY & LICENSING', 'CONNECTIVITY', and 'NOTIFICATIONS'. The 'Managed Clusters' tab is selected, showing a list of clusters managed by Ops Director. The table has columns for Cluster, Most Recent Update, Hosts, and OS. Three clusters are listed, each with a 'REMOVE' button.

Cluster	Most Recent Update	Hosts	OS
engr1ab-128-059.engr1ab.marklogi...	01/30/18, 05:13	2	Linux 3.10.0-514.16.1.el7.x86_...
engr1ab-128-124.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...
engr1ab-128-178.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...

The columns displayed in the Managed Clusters tab are described in the following table.

Column	Description
Name	The name of the cluster.
Most Recent Update	The most recent date and time the cluster status was updated.
Hosts	The number of hosts in the cluster.
OS	The name and version of the operating system of the hosts on the cluster.

You may export data from the Managed Clusters tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Managed Clusters table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

A managed cluster may become disconnected from the Ops Director, either due to temporary network unavailability or because MarkLogic Server had stopped on the hosts comprising the cluster. In this case, the cluster's state cannot be determined by Ops Director, and hence the cluster is assigned the Unknown state.

For additional reasons why a managed cluster might be assigned the Unknown state, see “Security and Database Dependencies of Managed Clusters” on page 13.

You may filter the list of the managed clusters to view only those clusters that are currently in the Unknown state by selecting the **Show only unknown** checkbox.

The screenshot shows the MarkLogic Console Settings page. The left sidebar contains navigation options: SECURITY & LICENSING (RBAC Settings, Roles, Resource Access, Resource Group, License Information) and CONNECTIVITY (Managed Clusters). The main content area is titled 'Managed Clusters' and shows a list of clusters. The 'Show only 2 unknown' checkbox is checked, and a mouse cursor is pointing at it. The table below shows two clusters in the Unknown state.

Cluster	Most Recent Update	Hosts	OS	
engr1ab-128-124.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...	REMOVE
engr1ab-128-178.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...	REMOVE

7.4.2 Removing Unknown Managed Clusters from the List

You may remove a managed cluster that is currently in the Unknown state from the list of clusters managed by Ops Director. Perform the following steps:

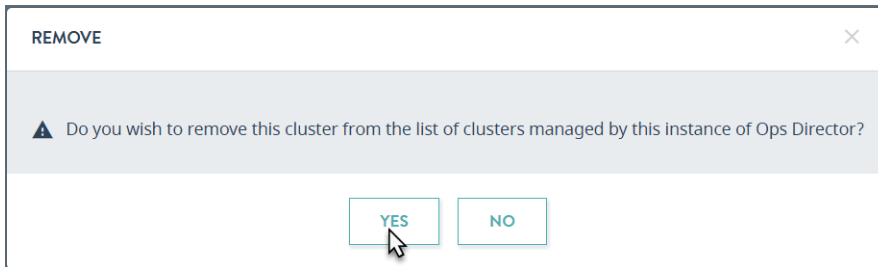
1. Click **REMOVE** to the right of the cluster you want to remove from the list.

The screenshot shows the MarkLogic Console Settings page. The left sidebar is the same as in the previous screenshot. The main content area is titled 'Managed Clusters' and shows a list of clusters. The 'Show only 2 unknown' checkbox is unchecked. A mouse cursor is clicking the 'REMOVE' button for the cluster 'engr1ab-128-124.engr1ab.marklogi...'. The table below shows three clusters.

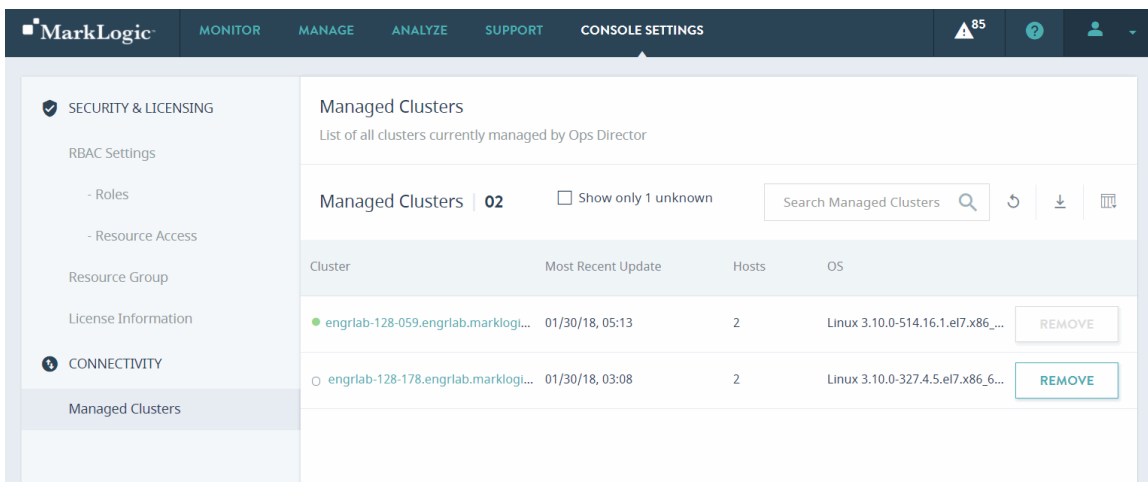
Cluster	Most Recent Update	Hosts	OS	
engr1ab-128-059.engr1ab.marklogi...	01/30/18, 05:13	2	Linux 3.10.0-514.16.1.el7.x86_...	REMOVE
engr1ab-128-124.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...	REMOVE
engr1ab-128-178.engr1ab.marklogi...	01/30/18, 03:08	2	Linux 3.10.0-327.4.5.el7.x86_6...	REMOVE

Note: The **Remove** button is enabled only for clusters that are currently in the Unknown state.

2. A confirmation box appears asking **Do you wish to remove this cluster from the list of clusters managed by this instance of Ops Director?** Click **YES**.



3. The cluster is removed, and the updated list of managed clusters is displayed by Ops Director.



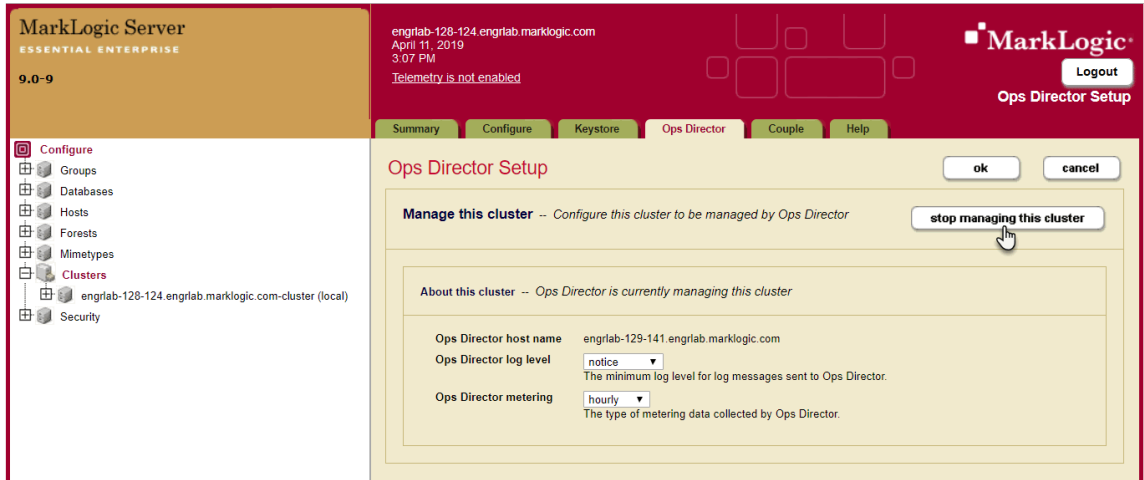
7.4.3 Reconnecting a Managed Cluster to Ops Director

If you removed a cluster from the list of clusters managed by Ops Director, you may want to reconnect this cluster later on, such as when the issue that caused the Unknown state is resolved.

To add the cluster back to the list of clusters managed by Ops Director from the Admin Interface, perform the following steps:

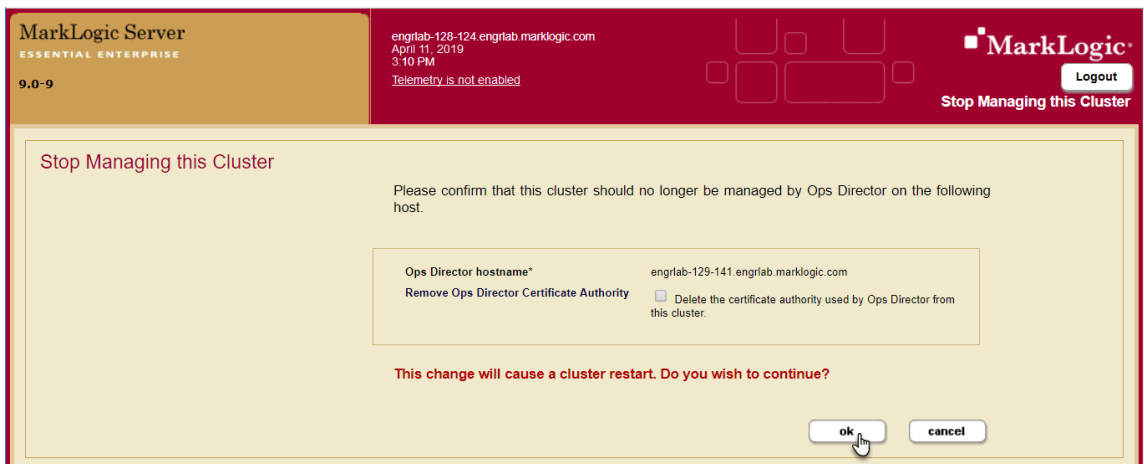
1. Log into the Admin Interface on the cluster to be managed by Ops Director.
2. In the left menu, select **Configure > Clusters**.
3. Select the local cluster. The Edit Local Cluster Configuration page appears.

4. Select the **Ops Director** tab at the top of the page.
5. The Ops Director Setup page is displayed. Select **stop managing this cluster**.



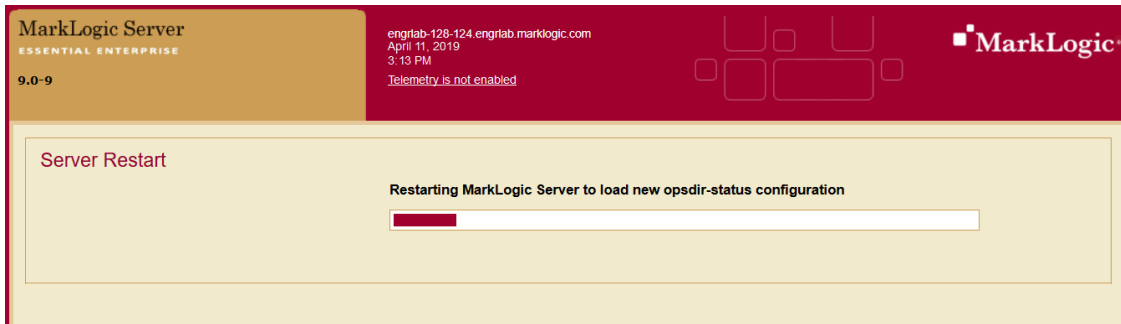
Note: When you remove the cluster from the list of managed clusters in Ops Director, the cluster is not notified that it is no longer managed, because the connection between the cluster and the Ops Director is down at that point. Therefore, you must first update the cluster state by selecting **stop managing this cluster** from the Admin Interface.

6. On the Stop Managing this Cluster page, make sure the box to the left of **Remove Ops Director Certificate Authority** box is unchecked, then click **ok**.

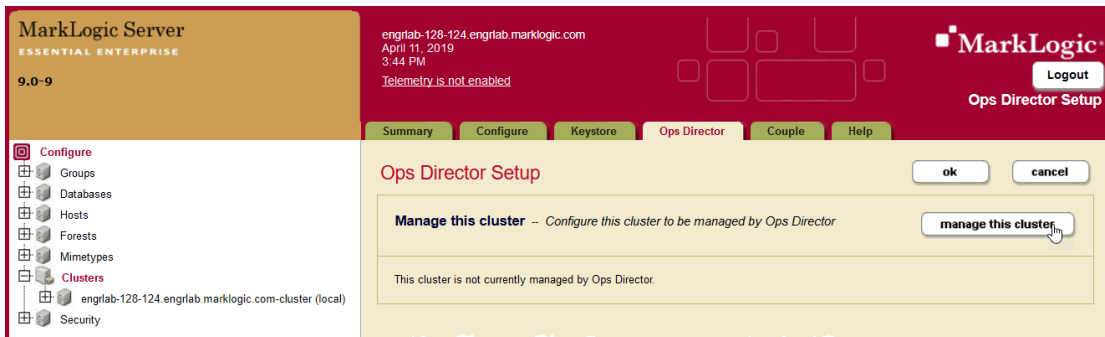


Note: You may stop managing and start managing a cluster without re-importing a certificate to it, in case the cluster will be managed by the same instance of Ops Director.

7. The Server Restart page is displayed. The page informs you that MarkLogic Server is being restarted to load new Ops Director status configuration.



8. When the restart is completed, the Ops Director Setup page is displayed again. Select **manage this cluster**.



9. The Configure as a Managed Cluster page is displayed.
 - Enter the name of the host where your Ops Director application runs.
 - From the Ops Director Certificate Authority menu, select **MarkLogic Ops Director Certificate Authority**.

You may optionally modify the level for log messages sent to Ops Director, as well as the frequency at which the metering data is collected.

Click **ok**.

MarkLogic Server
ESSENTIAL ENTERPRISE
9.0-9

engr1ab-128-124.engr1ab.marklogic.com
April 11, 2019
3:56 PM
Telemetry is not enabled

MarkLogic
Logout
Configure as a Managed Cluster

Configure as a Managed Cluster

Use the fields below to identify the Ops Director server already installed within your enterprise. This will allow Ops Director to manage this cluster.

Ops Director hostname	engr1ab-129-141.engr1ab.marklogic.com The (resolvable from this host) Ops Director hostname.
Ops Director system port	8009 The Ops Director port for system communication with managed clusters.
Ops Director log level	notice The minimum log level for log messages sent to Ops Director.
Ops Director metering	hourly The type of metering data collected by Ops Director.
Ops Director Certificate Authority	MarkLogic Ops Director Certificate Authority The certificate authority shared with the Ops Director system.
This hostname	engr1ab-128-124.engr1ab.marklogic.com The (resolvable from the Ops Director host) name of this host.

ok cancel

- In Ops Director, select the **CONSOLE SETTINGS** tab. In the left menu under **CONNECTIVITY**, select **Managed Clusters**.

Refresh the page. The updated list of managed clusters is displayed, with the newly reconnected cluster among them.

7.5 Configuring Email Notifications

You can configure Ops Director to notify you by email when a specified type of event occurs, or when an alert is enabled or disabled.

Under **NOTIFICATIONS**, select **Email Configuration** to set up emailed notifications.

This section covers the following topics:

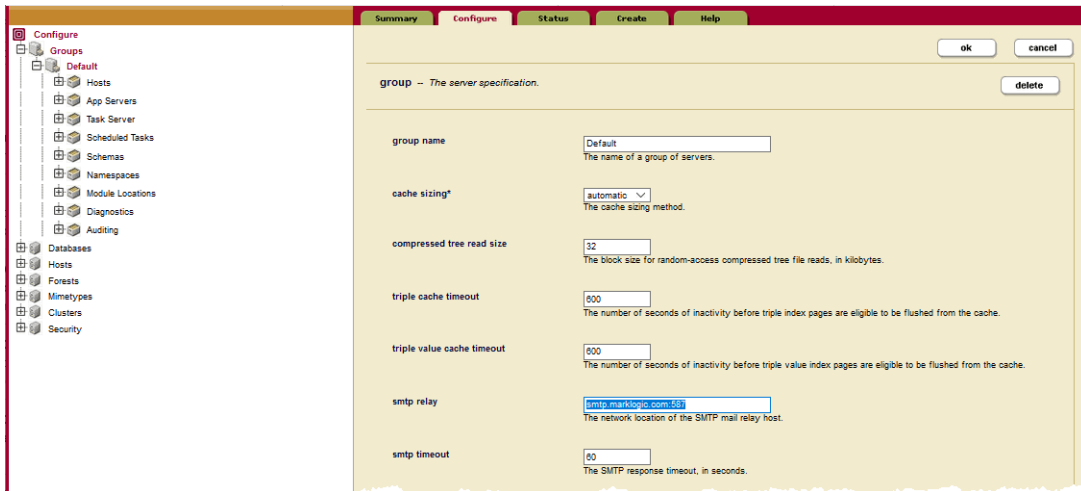
- [Setting Up Email Notifications](#)
- [Setting up and Managing Alerts](#)

7.5.1 Setting Up Email Notifications

To set up Ops Director to notify you by email when a specific type of event occurs:

- Make sure you have set up your SMTP relay. In the Admin Interface, navigate to **Configure > Groups > Default**. In the right pane, select the **Configure** tab. Make sure

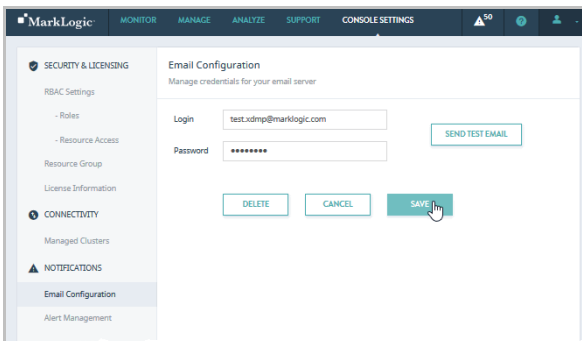
the **smtp relay** field contains the correct name and port of the SMTP relay host. If it needs to be changed, make sure to click **ok** at the bottom of the pane to save your changes.



2. Log into the XQuery console (port 8000) and submit the following query to update the “From” email address against the database **OpsDirector** and the server **OpsDirectorApplication**, replacing **Test User** and **test-xdmp-email@marklogic.com** with the name and email address from which the alerts are to come:


```
xquery version "1.0-ml";
import module namespace cfg="http://marklogic.com/v1/opsdirector/config"
  at "/common/config.xqy";
cfg:set-atomic-property("email-return-name", "Test User"),
cfg:set-atomic-property("email-return-address", "test-xdmp-email@marklogic.com")
```

3. In the Ops Director interface, select **CONSOLE SETTINGS**. In the left menu, select **Email Configuration**. The Email Configuration page displays in the right pane.

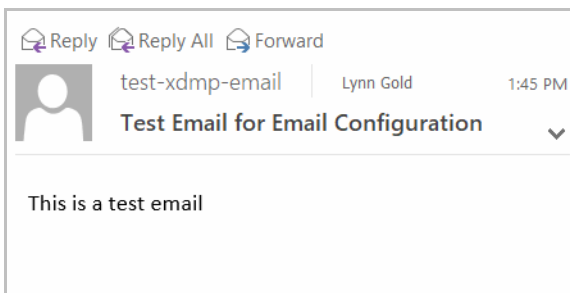


4. Enter the Login and Password credentials for the the SMTP relay server configured in the Admin UI in step [1](#). Click **SAVE** to save the credentials.

- To make sure the Ops Director email server is working correctly, click **SEND TEST EMAIL**. The SEND TEST EMAIL box displays.



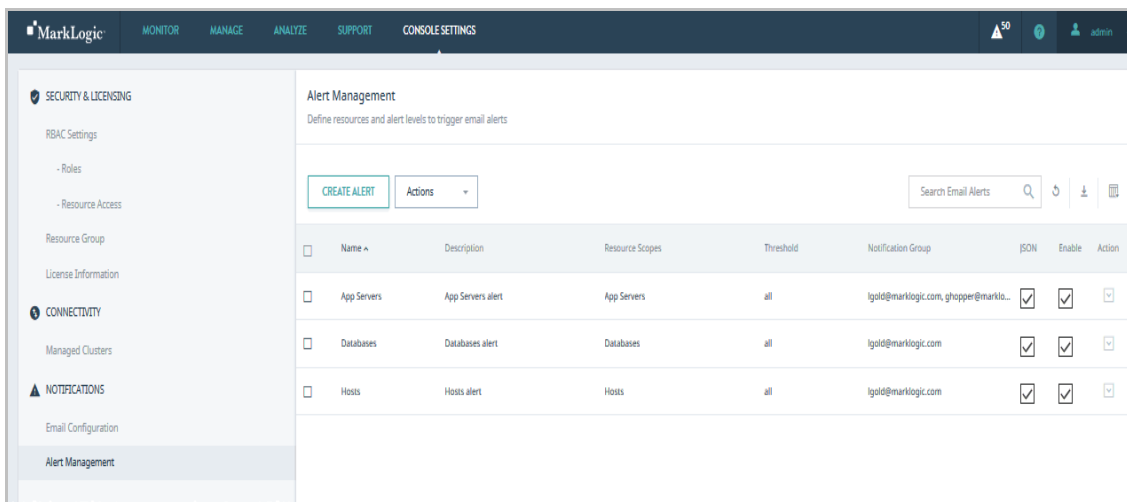
- Enter an email address to receive the test message and click **OK**. You will receive a test email message similar to the following:



To remove the account from Ops Director, click **DELETE**.

7.5.2 Setting up and Managing Alerts

The **Alert Management** page enables you to determine when email alerts are sent. To display the Alert Management page in the right pane, in the left menu, select **Alert Management**.

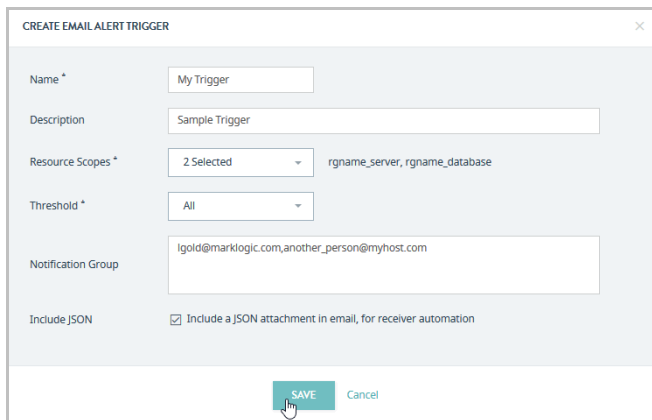


7.5.2.1 Creating an Email Alert

When you create or edit an email alert, make sure you consider the email recipient's RBAC (Role Based Access Control) settings. Make sure that email notifications with full alert details about MarkLogic components only be sent to users with full view of these components.

To create an email alert:

1. In the Ops Director Alert Management page, click **CREATE ALERT**. The **CREATE EMAIL TRIGGER ALERT** window appears.

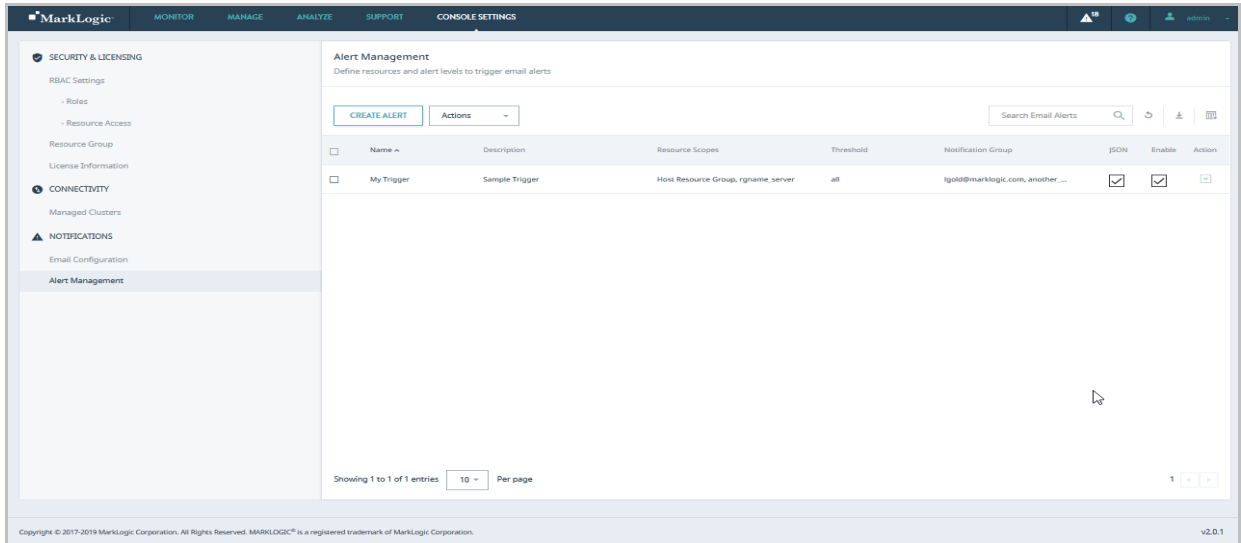


The screenshot shows a dialog box titled "CREATE EMAIL ALERT TRIGGER". It contains the following fields and controls:

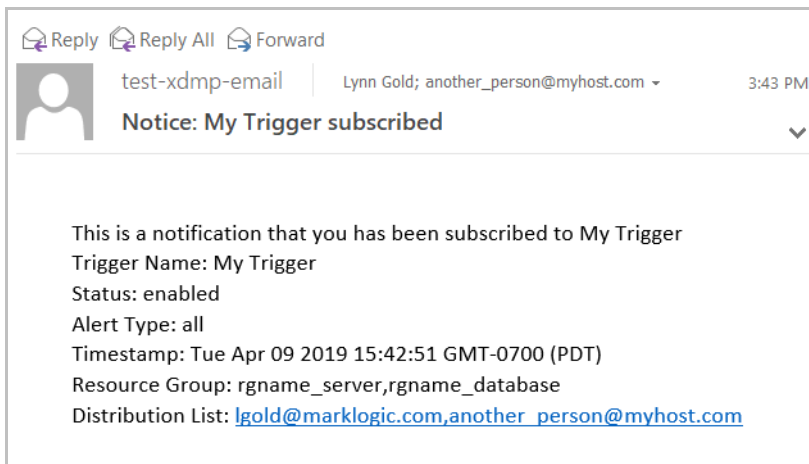
- Name ***: Text input field containing "My Trigger".
- Description**: Text input field containing "Sample Trigger".
- Resource Scopes ***: A dropdown menu showing "2 Selected" and a list of selected scopes: "rgname_server, rgname_database".
- Threshold ***: A dropdown menu showing "All".
- Notification Group**: Text input field containing "lgold@marklogic.com, another_person@myhost.com".
- Include JSON**: A checked checkbox with the label "Include a JSON attachment in email, for receiver automation".
- At the bottom, there are two buttons: "SAVE" (highlighted in green) and "Cancel".

2. In the **CREATE EMAIL TRIGGER ALERT** window:
 - Enter the **Name** of the trigger. You may also add a **Description**.
 - From the **Resource Scopes** menu, select one or more resource groups.
 - From the **Threshold** menu, select one of **Critical**, **At Risk**, or **All**, depending upon the level of alerts you want to monitor.
 - Check the **Include JSON** box if you want to include a JSON attachment in the email alert.
 - Enter one or more email addresses in the **Notification Group** box.

Click **SAVE**. The new alert shows up in the Alert Management pane and is automatically enabled.



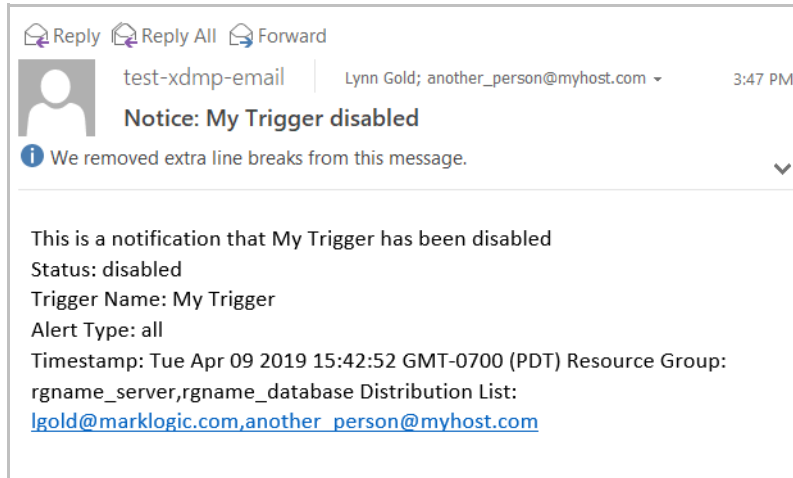
Anyone in the Notification Group will receive an email message similar to the following:



The fields in your email message are:

- **Trigger Name:** The name of the Ops Director alert.
- **Status:** Whether the alert is enabled or disabled.
- **Alert Type:** The alert threshold: one of critical, at risk, or all.
- **Timestamp:** When the email was sent.
- **Resource Group:** The resource groups involved in the alert
- **Distribution List:** The email addresses that receives this alert.

To disable a trigger, select the box to the right of the trigger in the **Enable** column so it is unchecked. If you are in the Notification Group, you will receive an email message similar to the following:



The fields that appear in this message are:

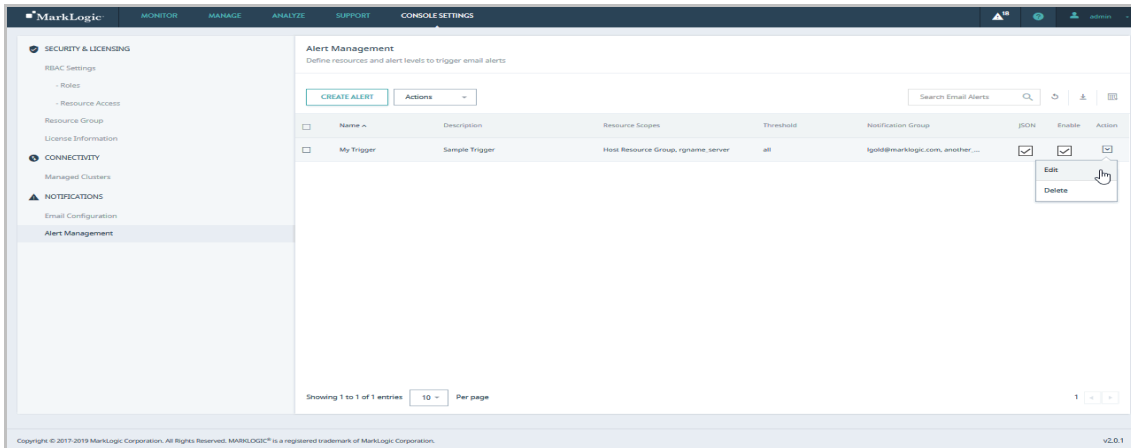
- **Status:** Whether the alert is enabled or disabled.
- **Trigger Name:** The name of the Ops Director alert.
- **Alert Type:** The alert threshold: one of critical, at risk, or all.
- **Timestamp:** When the email was sent.
- **Resource Group:** The resource groups involved in the alert
- **Distribution List:** The email addresses that received this alert.

7.5.2.2 Editing an Email Alert

When you create or edit an email alert, make sure you consider the email recipient's RBAC (Role Based Access Control) settings. Make sure that email notifications with full alert details about MarkLogic components only be sent to users with full view of these components.

To edit an email alert:

1. To the right of the alert you want to edit, select the **Action** menu.



2. From the **Action** menu, select **Edit**. The **EDIT EMAIL ALERT TRIGGER** window displays.

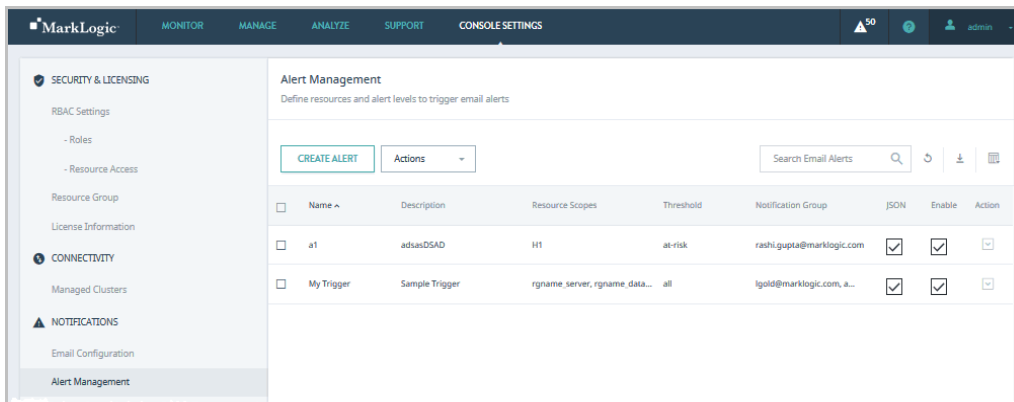
3. Edit one or more of the following fields:
 - The **Name** of the trigger.
 - The **Description** of the trigger.
 - From the **Resource Scopes** menu, you may select or deselect one or more resource groups.

- From the **Threshold** menu, you may select one of **Critical**, **At Risk**, or **All**, depending upon the level of alerts you want to monitor.
- Click the **Include JSON** box to include a JSON attachment in the alert email.
- Add or remove one or more email addresses in the **Notification Group** box.

Click **SAVE** to save your changes.

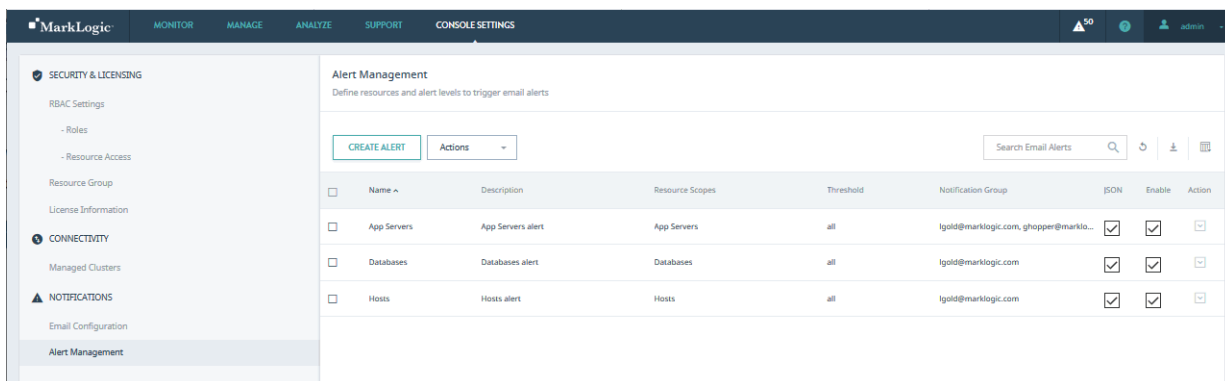
Note: The **SAVE** button does not become operational until you “edit” either the **Name**, **Description**, or **Notification Group** field.

Your changes appear in the Alert Management page.



7.5.2.3 Viewing and Filtering the List of Alerts

Use the Alert Management page to view the list of all alerts currently managed by Ops Director.



The columns displayed in the Alert Notifications page are described in the following table.

Column	Description
Name	The name of the alert.
Description	An optional description of the alert.
Resource Scopes	The resource groups to be monitored by this alert.
Threshold	The error threshold at which the alert is triggered.
Notification Group	The email addresses of users to be notified when an alert is triggered.
JSON	Whether to send a JSON notification when the alert is triggered.
Enable	Whether or not the alert is enabled.

You may export data from the Alert Management page as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Managed Clusters table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

7.5.2.4 Deleting an Email Alert

You can delete an email alert in one of the following ways:

- To the right of the alert you want to delete, select the **Action** menu. From the **Action** menu, select **Delete**.
- Select one or more alerts in the box column to the left of the **Name** column. To the right of **CREATE ALERT**, select **Actions**, then select **Delete All Selected Triggers**.

A DELETE EMAIL ALERT TRIGGER confirmation box displays. Click **YES** to delete all selected alerts.

8.0 SUPPORT view

The **SUPPORT** view shows alerts for application servers in your enterprise, allows you to view and filter server logs, and displays tasks that run on your managed clusters.

This chapter covers the following topics:

- [System Alerts](#)
- [Event Logs](#)
- [Task Console](#)

8.1 System Alerts

Use the **SYSTEM ALERTS** tab to review alerts from all monitored servers throughout the enterprise.

Note: Alerts are associated with all types of resources. For example, if you have a resource group of databases containing db1 and db2, then only alerts associated with db1 and db2 and their forests are displayed.

<input type="checkbox"/>	Alert Name	Severity	Code	Start Time	Stop Time	Cluster
<input type="checkbox"/>	Schemas	Info	HEALTH-DATABASE-NO-BA...	07/17/17, 9:58...		ManagedCluster2
<input type="checkbox"/>	Schemas	Info	HEALTH-DATABASE-NO-BA...	07/17/17, 9:58...		ManagedCluster1
<input type="checkbox"/>	Extensions	Info	HEALTH-DATABASE-NO-BA...	07/17/17, 9:58...		OpsDirectorCluste
<input type="checkbox"/>	Triggers	Info	HEALTH-DATABASE-NO-BA...	07/17/17, 9:58...		ManagedCluster1
<input type="checkbox"/>	Security	Info	HEALTH-DATABASE-NO-BA...	07/17/17, 9:58...		ManagedCluster2

The columns displayed for **SYSTEM ALERTS** are described in the following table.

Column	Description
Resource Name	The name of the resource related to the alert.
Severity	Color-coded alert severity level. For details, see “Alert Levels” on page 255.
Code	The alert code.
Start Time	Alert timespan starting date and time
Stop Time	Alert timespan ending date and time
Cluster	Name of the cluster that generated the alert
Alert Details	Detailed description of alert
Status	Alert status (Open/Closed/Acknowledged)
Action	Select Acknowledge to remove the alert from the list or Analyze to navigate to the ANALYZE view of the resource. For details, see the “ANALYZE View” on page 189.

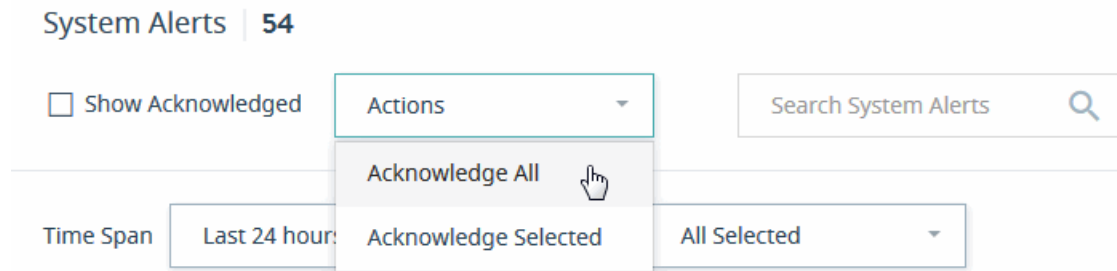
You may export data from the System Alerts tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the System Alerts table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

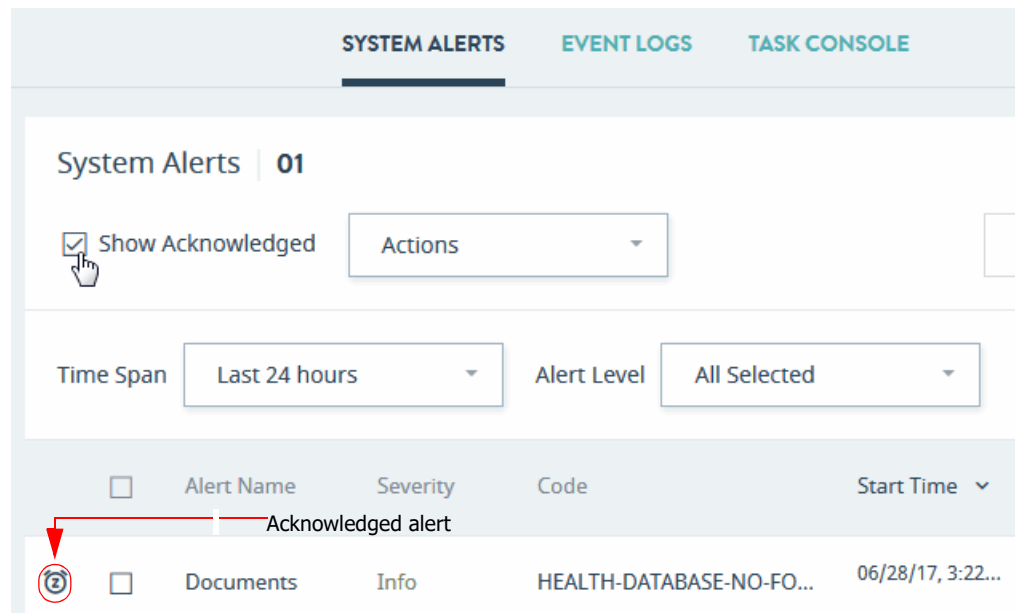
You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

Once you have reviewed an alert, you can mark specific or all alerts as acknowledged to remove them from the list. Once acknowledged, the alert is suppressed from further notifications. The acknowledge action toggles depending on the state of an alert. If the alert is suppressed then the option changes to Unacknowledge, which unsuppresses the alert.

Select **Acknowledge All** to acknowledge all alerts, whether or not they are currently displayed. Acknowledge Selected is enabled if one or more alerts is selected via checkboxes. Select **Acknowledge Selected** to acknowledge those selected alerts.



Acknowledged alerts are indicated by an icon to the left of each acknowledged alert. The Acknowledged state is a property of the alert that is shared for all users. You can include acknowledged alerts in the table by selecting the **Show Acknowledged** checkbox. Uncheck to exclude acknowledged alerts.



You can narrow the results by specifying a time period, resource name, and alert severity.

Time Span — You can select from the list of predefined time spans or select a custom date and time range by choosing a start date and time and end date and time, as described in “Date and Time Filters” on page 86.

System Alerts | 54

Show Acknowledged Actions Search System Alerts

Time Span: Last 24 hours Alert Level: All Selected

	Code	Start Time	Stop Time
<input type="checkbox"/>	HEALTH-HOST-RECENT-RES...	06/23/17, 9:52...	
<input type="checkbox"/>	HEALTH-HOST-RECENT-RES...	06/23/17, 9:48...	

Alert Level — You can select one or more alert levels to narrow the list of alerts that have the specified levels.


System Alerts | 54

Show Acknowledged Actions Search System Alerts

Time Span: Last 24 hours Alert Level: All Selected

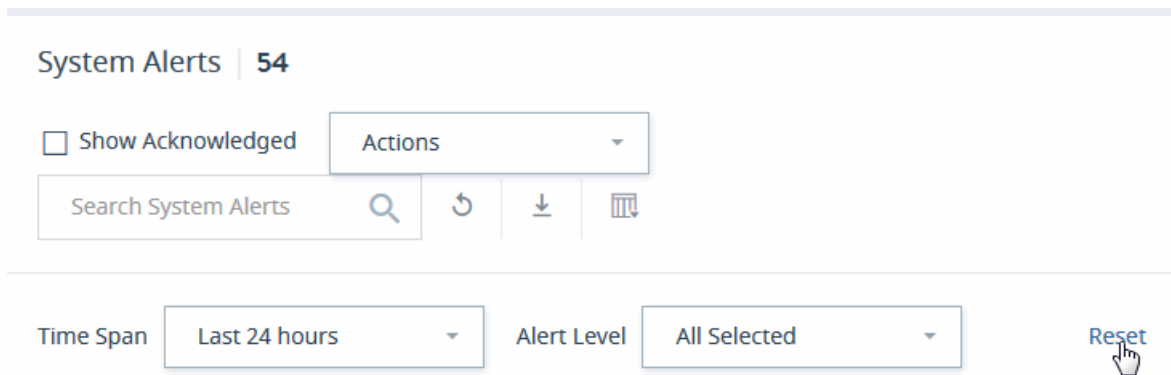
	Alert Name	Severity	Code	Stop Time
<input type="checkbox"/>	engrlab-130-1...	Info	HEALTH-HOS	
<input type="checkbox"/>	engrlab-130-2...	Info	HEALTH-HOS	
<input type="checkbox"/>	Fab	Critical	HEALTH-DAT	

Search System Alerts — You can search for specific system alerts. The search text must be more than two characters.



Note: The spaces in the string that you pass in the Search System Alerts box are trimmed off. For example, entering HEALTH SERVER with a space at the end will search for all alerts starting with HEALTH SERVER, such as HEALTH SERVER-DISABLED and so on.

Click **Reset** to reset the applied filters to their default values. The Reset button is enabled when you make changes to any of the currently applied filters.



8.1.1 Alert Levels

The color-coded alert levels are described in the following table.

Level	Description
Critical (red)	Resource is in a critical state - there is no hard definition of this alert as the resource can also be offline or non operational.

Level	Description
At Risk (yellow)	<p>Resource that is dependent on another resource that is critical generates an At Risk alert.</p> <p>Specifically:</p> <ul style="list-style-type: none"> • A cluster is At Risk if any of its resources is in a Critical state • A group is At Risk <ul style="list-style-type: none"> • if a host is Critical or At Risk • if an App Server is Critical or At Risk • A host is At Risk <ul style="list-style-type: none"> • if a forest is Critical or At Risk • if an App Server is Critical or At Risk (note the App Server At Risk alert will show up on every host) • An App Server is At Risk <ul style="list-style-type: none"> • if the database is Critical or At Risk • A database is At Risk <ul style="list-style-type: none"> • if a forest is Critical or At Risk • if a replica is Critical or At Risk • if a foreign replica is Critical • if a schema or trigger database is Critical or At Risk • if a sub-database is Critical or At Risk <p>There are other At Risk scenarios, such as:</p> <ul style="list-style-type: none"> • A forest has low storage space • A foreign replica is unavailable • 30% of hosts in a cluster are failing would place cluster At Risk
Healthy (green)	Resource has no cluster-health-reports associated it is considered healthy.
Maintenance (dark grey)	Resource status has been intentionally disabled.

Level	Description
Offline (light grey)	Resource state cannot be determined or resource is in an Offline state (for example, result of setting enabled property to false for the resource).
Information (dark green)	Optional information and advice.
Unknown (white / hollow)	State of resource cannot be determined or an Ops Director application error has been thrown.

8.1.2 Point-in-Time Alerts

Point-in-time alerts are derived from logs.

Certain log messages in the server logs make Ops Director generate alerts. Generation of the alerts is based on rules configured in the `debug.xml` file in the Ops Director database. Each rule specifies a log level and a regular expression to which a log message of that level has to match. This log message is translated into an alert of the severity specified in the rules.

8.2 Event Logs

Use the **EVENT LOGS** tab to review events from all monitored App Servers throughout the enterprise. Narrow the results by specifying a time period, log type, and log file.

Note: Access to log data is based on your access to hosts and clusters. For example, if you have a resource group of hosts containing `host1` and `host2`, only logs associated only with `host1` and `host2` are displayed.

The columns displayed for **EVENT LOGS** are described in the following table.

Column	Description
Date & Time	Datetime of the logged event.
Level	The log level of the event. For a description of the log levels, see Understanding the Log Levels in the <i>Administrator's Guide</i> .
Cluster Name	The name of the cluster on which the logged event occurred.
Host Name	The name of the host on which the logged event occurred.
Message	The logged error message.
Log File	The full text of the logged event.

You may export data from the **EVENT LOGS** tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Event Logs table in the UI.

- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such As Excel) for further processing or analysis.

By default, all event logs are displayed. Use the Log Level filter to select which types of event logs to view. For a description of the log levels, see [Understanding the Log Levels](#) in the *Administrator’s Guide*.

Event Logs | 08 Sea

Time Span: Last 24 hours | Log Level: 1 Selected

Date & Time	Level	Cluster Name	Host
06/23/17, 12:39 ...	Warning	engrlab-130-132.engrl...	en...
06/23/17, 12:39 ...	Warning	engrlab-130-132.engrl...	en...
06/23/17, 12:38 ...	Warning	engrlab-128-124.engrl...	en...
06/23/17, 12:36 ...	Warning	engrlab-130-132.engrl...	en...
06/23/17, 12:36 ...	Warning	engrlab-130-132.engrl...	en...
06/23/17, 12:33 ...	Warning	engrlab-130-132.engrl...	en...
06/23/17, 11:59 ...	Warning	engrlab-128-178.engrl...	en...
06/23/17, 11:56 ...	Warning	engrlab-128-124.engrl...	engrlab-128-124.engrl... Pooled

All

Finest

Finer

Fine

Debug

Config

Info

Notice

Warning ✓

Error

Critical

Alert

Emergency

Select any log file to access its full contents, which may be downloaded for further analysis, or for archival and audit purposes.

...	engr1ab-128-124.en...	engr1ab-130-141.en...	Uploaded 1 records, 1 MB of Error Logs to O...	ErrorLog.txt
...	engr1ab-128-178.en...	engr1ab-130-153.en...	Uploaded 1 records, 1 MB of Error Logs to O...	ErrorLog.txt

The text for the event log is displayed in a popup window. The selected log entry is boldfaced. The window includes entries up to 10 minutes before and 5 minutes after the selected entry.

ERROR.TXT

(from rh7v-intel64-90-opsdir-3.marklogic.com)

```

2017-08-30 11:13:29.338 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/236410446819226879/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:29.339 Debug: Retrying xdm:invoke update-configs.xqy 2183132511081466225 Update 1 because XDMP-
DEADLOCK: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179/236410446819226879
/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:29.340 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/236410446819226879/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:29.342 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/236410446819226879/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:29.343 Debug: Retrying xdm:invoke update-configs.xqy 4737871800559412130 Update 1 because XDMP-
DEADLOCK: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179/236410446819226879
/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:30.325 Debug: Uploaded 19 records, 1 MB of Error Logs to Ops Director
2017-08-30 11:13:32.528 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/13347275998269105556/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:32.528 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/13347275998269105556/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:32.528 Debug: Retrying xdm:invoke update-configs.xqy 17238586939899957666 Update 1 because XDMP-
DEADLOCK: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179/13347275998269105556
/2017-08-30T11:10:47.502-07:00.json
2017-08-30 11:13:32.528 Info: Deadlock detected locking OpsDirector-4 /resource/databases/15555203222799652179
/13347275998269105556/2017-08-30T11:10:47.502-07:00.json

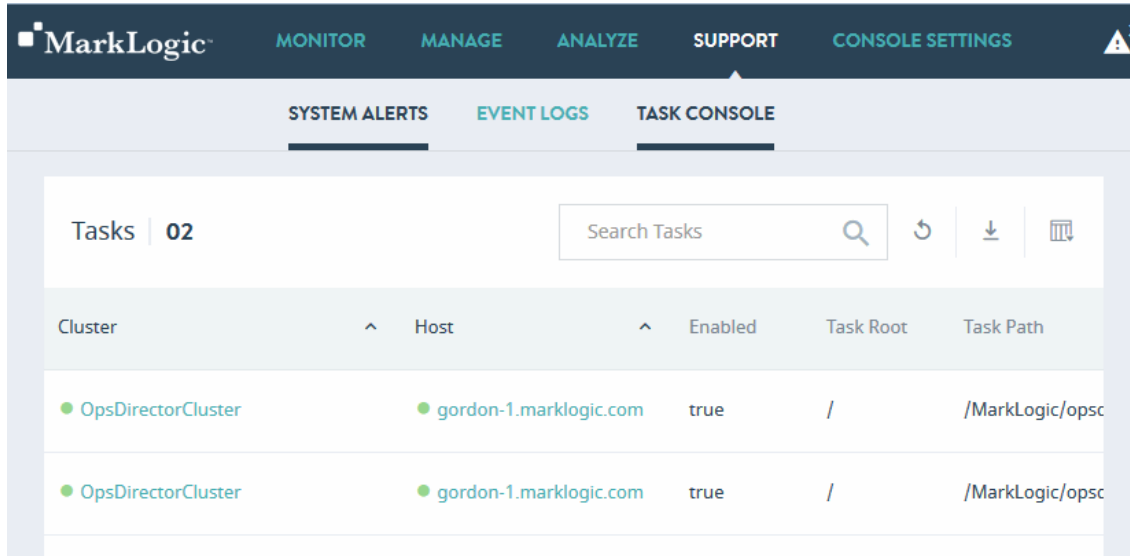
```

Note: Log times are expressed in GMT -07:00.

8.3 Task Console

Use the **TASK CONSOLE** tab for an overview of all tasks throughout the enterprise. The data table presents the following information:

- the task name
- the cluster, host, and database on which it is running
- the user who initiated the task
- the task priority
- the current state of the task



The columns displayed for **TASK CONSOLE** are described in the following table.

Column	Description
Cluster	Cluster on which the task host is located.
Host	The hostname of the host computer on which the scheduled module is to be invoked.
Enabled	Whether the task is enabled (true) or disabled (false).
Task Root	The root directory (filesystem) or URI root (database) that contains the module.
Task Path	The module the task is to invoke.

Column	Description
Task Type	<p>The task type:</p> <ul style="list-style-type: none"> • minutely specifies how many minutes between each invocation of the module. • hourly specifies how many hours and minutes between each invocation of the module. • daily specifies how many days between each invocation of the module and the time of day (in 24:00 notation). • weekly specifies how many weeks between each invocation of the module, check one or more days of the week, and the time of day (in 24:00 notation) for the task to start. • monthly specifies how many months between each invocation of the module, select one day of the month (1-31), and the time of day (in 24:00 notation) for the task to start. • one-time specifies the start date in MM/DD/YYYY notation (for example, 07/29/2009 for July 29, 2009) and time in 24:00 notation.
Task Period	How often the module is to be invoked (every n months, weeks, days, hours, or minutes).
Created on	The datetime the task was created.
Database	The database to which the scheduled module connects for query execution.
Task Modules	The name of the database in which the scheduled module locates the application code. If set to (file system), then any files in the specified task root directory are executable (given the proper permissions). If set to a database, then any documents in the database whose URI begins with the specified task root directory are executable.
User	The user with permission to invoke the module.
Priority	<p>The priority of the task:</p> <ul style="list-style-type: none"> • normal specifies the task is queued with normal priority. • higher specifies the task is queued with higher priority.

You may export data from the Task Console tab as a CSV (Comma Separated Values) file by clicking the Export icon in the upper right corner. The following rules apply:

- The resulting CSV file will have the same columns as the Task Console table in the UI.
- All available columns are exported, regardless of whether they are visible or hidden in the UI.
- All records that satisfy the set filter parameters are exported.
- Data types of the fields in the CSV file correspond to those in the UI. In case of boolean values, Yes/No in the UI corresponds to TRUE/FALSE in the CSV file.

You may then import the CSV file into other applications (such as Excel) for further processing or analysis.

9.0 Troubleshooting with Ops Director

Ops Director may facilitate troubleshooting of MarkLogic clusters in your enterprise.

Ops Director has set thresholds on specific metrics to alert you when a metric exceeds a pre-specified value. Many metrics that can help in alerting and troubleshooting are meaningful only in relation to 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.

This chapter describes major use cases of troubleshooting with Ops Director. It provides a set of guiding questions to help you understand and identify the metrics that are of interest under various circumstances.

This chapter covers the following topics:

- [Assess Whether MarkLogic Has Adequate Resources](#)
- [Assess the Overall State of the System](#)
- [Assess MarkLogic Cluster Performance](#)
- [Assess Severity of Problems in the System](#)

9.1 Assess Whether MarkLogic Has 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.

Aspect	Analysis
Sufficient resources for-MarkLogic Server 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. The metrics on the MONITOR View will generally give you a high-level view of the loads on your cluster resources. For more detail, use the ANALYZE View .
Sufficient 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 <i>Installation Guide</i>). If a merge runs out of disk space, it will fail. The metrics described in “Database Performance Data” on page 207 will help isolate disk space problems.

Aspect	Analysis
Sufficient 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. The metrics described in “Database Performance Data” on page 207 will help isolate disk space problems.
Sufficient 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. The metrics described in “Memory Performance Data” on page 200 will help isolate memory problems.
Correctness of swap space configuration	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. The metrics described in “Memory Performance Data” on page 200 will help isolate swap space problems. For details on configuring swap memory, see Tuning Query Performance in MarkLogic Server in the <i>Query Performance and Tuning Guide</i> .
Number of hosts in a cluster and their configuration	How many hosts are in a 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 <i>Scalability, Availability, and Failover Guide</i> .
Applications with resource-intensive features	Applications with resource-intensive features include CPF, replication, and point-in-time recovery. Are the hardware, software, and network resources available and configured to most efficiently support such applications?

9.2 Assess the Overall State of the System

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

Aspect	Analysis
Efficiency of CPU usage	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?

Aspect	Analysis
Efficiency of I/O usage	What amount of data is currently being read from or written to disk? Are there any I/O bottlenecks?
Free disk space per filesystem	Is there enough free disk space for each filesystem?
Network state	What is the current state of the network?
Errors messages in application logs	Are there any errors or warnings appearing in the logs of MarkLogic Server or applications?
Errors messages in system logs	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.

9.3 Assess MarkLogic Cluster Performance

When you suspect an error or performance problem originates from MarkLogic Server, some questions to ask are as follows. Most of these metrics can be viewed on the [MONITOR View](#) and [ANALYZE View](#).

Aspect	Analysis
Whether all resources in the cluster are utilized	Are all of the hosts in the cluster online? Are all of the App Servers enabled? In what states are the forests?
Query optimization and load balancing	What are the patterns of queries and updates? Do they appear to be evenly distributed across the hosts in the cluster?
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.
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.

Aspect	Analysis
I/O rates and loads pattern	<p>In this context, <i>rates</i> refers to amount of data applications are currently reading from or writing to MarkLogic Server databases (throughput) and <i>loads</i> 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.</p> <p>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.</p> <p>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.</p>
Journal and save write rates and loads pattern	<p>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.</p>
XDQP rates and loads pattern	<p>In this context, <i>rates</i> refers to amount of data hosts are currently reading from or writing to other hosts and <i>loads</i> 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.</p>
Cache hit/miss rates	<p>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.</p>
Concurrent updates and reads are in progress	<p>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.</p>

Aspect	Analysis
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 <i>Administrator's Guide</i> .
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 <i>Administrator's Guide</i> .
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 <i>Administrator's Guide</i> .

9.4 Assess Severity of Problems in the System

The [MONITOR View](#) alerts you to the more serious problems in your MarkLogic clusters. If you are encountering a serious problem in which MarkLogic Server is unable to effectively service your applications, use the following problem analysis for the system troubleshooting.

Problem	Analysis and Workaround/Solution
MarkLogic Server aborts or fails 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.
An application is 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.
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 <i>Query Performance and Tuning Guide</i> .

Problem	Analysis and Workaround/Solution
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 <i>Scalability, Availability, and Failover Guide</i> for possible causes.
Messages of the error level and higher in the log files	The various log levels are described in Understanding the Log Levels in the <i>Administrator's Guide</i> . 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 are listed in the following table.

Error message	Root Cause Analysis
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 should not happen. Look at the error text for details and the logs for context. If you have an active maintenance contract, you can contact MarkLogic Technical Support.

10.0 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](#).

11.0 Copyright

MarkLogic Server 9.0 and supporting products.
Last updated: March 25, 2019

COPYRIGHT

© 2019 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.

