

PingDataMetrics™

Release 7.3

Server Administration Guide



PingDataMetrics Server™ Product Documentation

© Copyright 2004-2019 Ping Identity® Corporation. All rights reserved.

Trademarks

Ping Identity, the Ping Identity logo, PingFederate, PingAccess, and PingOne are registered trademarks of Ping Identity Corporation ("Ping Identity"). All other trademarks or registered trademarks are the property of their respective owners.

Disclaimer

The information provided in these documents is provided "as is" without warranty of any kind. Ping Identity disclaims all warranties, either express or implied, including the warranties of merchantability and fitness for a particular purpose. In no event shall Ping Identity or its suppliers be liable for any damages whatsoever including direct, indirect, incidental, consequential, loss of business profits or special damages, even if Ping Identity or its suppliers have been advised of the possibility of such damages. Some states do not allow the exclusion or limitation of liability for consequential or incidental damages so the foregoing limitation may not apply.

Support

<https://support.pingidentity.com/>

Table of Contents

- Chapter 1: Introduction** **1**
- PingDataMetrics Server Overview 2
- PingDataMetrics Server Components 2
- Data Collection 3
 - Performance Data 3
 - System and Status Data 3
- Charts and Dashboards 4
- PostgreSQL DBMS Details 4
- Chapter 2: Installing the PingDataMetrics Server** **6**
- Supported Platforms 7
- Install the JDK 7
- Configure a Non-Root User 7
- Optimize the Linux Operating System 8
 - Set the file descriptor limit 8
 - Set the filesystem flushes 9
 - Install sysstat and pstack on Red Hat 9
 - Install the dstat utility 10
 - Disable filesystem swapping 10
 - Manage system entropy 10
 - Set Filesystem Event Monitoring (inotify) 10
 - Tune IO scheduler 11
 - Enable the server to listen on privileged ports 11
- Configure Servers to be Monitored 12
 - Disk Space Requirements and Monitoring Intervals 12
 - Tracked Applications 13

Table of Contents

Ping license keys	14
Install the Server	14
Log into the Administrative Console	17
Server folders and files	18
Add Monitored Servers to the PingDataMetrics Server	19
Using the monitored-servers Tool	19
Remove Monitored Servers	20
Start and Stop the Server	20
Start the PingDataMetrics Server as a Background Process	20
Start the PingDataMetrics Server as a Foreground Process	20
Start the PingDataMetrics Server at Boot Time	21
Stop the PingDataMetrics Server	21
Restart the PingDataMetrics Server	21
Uninstall the server	21
Update servers in a topology	22
Update the server	23
Reverting an Update	24
Revert an Update	24
Reverting from Version 7.x to a Version Prior to 7.0	24
To Revert to the Most Recent Server Version	26
Administrative accounts	26
Change the administrative password	26
Chapter 3: Managing the PingDataMetrics Server	27
PingDataMetrics Server Error Logging	28
Logging Retention Policies	28

Logging Rotation Policies	28
Create Log Publishers	29
Error Log Publisher	29
Configure log file encryption	31
Backend Monitor Entries	32
Disk Space Usage Monitor	34
Notifications and Alerts	35
Configure alert handlers	36
The Alerts Backend	36
View Information in the Alerts Backend	36
Modify the Alert Retention Time	37
Configure Duplicate Alert Suppression	37
System alarms, alerts, and gauges	37
Test alerts and alarms	39
Back Up the PingDataMetrics Server Database	40
Historical Data Storage	41
Planning the Backup	41
Start the DBMS Backup	42
Restore a DBMS Backup	42
Management Tools	42
Available Command-Line Utilities	43
The tools.property File	44
Tool-Specific Properties	45
Specify Default Properties Files	46
HTTP Connection Handlers	46
Configure an HTTP Connection Handler	47

Table of Contents

HTTP Correlation IDs	49
Topology configuration	52
Topology master requirements and selection	53
Topology components	54
Monitor data for the topology	55
Updating the server instance listener certificate	56
Remove the self-signed certificate	57
Use an existing key-pair	57
Use the certificate associated with the original key-pair	58
Use the Configuration API	60
Authentication and authorization	60
Relationship between the Configuration API and the dsconfig tool	61
API paths	69
Sorting and filtering configuration objects	70
Update properties	71
Administrative actions	73
Update servers and server groups	73
Configuration API Responses	74
Domain Name Service (DNS) caching	75
IP address reverse name lookups	75
Configure traffic through a load balancer	76
Configure authentication with a SASL external certificate	76
Server SDK extensions	78
Chapter 4: Collecting Data and Metrics	80
Metrics Overview	81

Count Metrics	81
Continuous Metrics	81
Discrete Metrics	82
Dimensions	82
Query Overview	84
Select Query Data	84
Aggregate Query Results	85
Format Query Results	86
The query-metric Tool	86
Performance Data Collection	87
System Monitoring Data Collection	89
Stats Collector Plugin	89
System Utilization Monitors	90
External Collector Daemon	90
Server Clock Skew	91
Tune Data Collection	91
Reducing the Data Collected	92
Reducing the Frequency of Data Collection	92
Reducing the Frequency of Sample Block Creation	92
Reducing PingDataMetrics Server Impact on Performance	92
Data Processing	93
Importing Data	93
Aggregating Data	94
Monitoring for Service Level Agreements	95
SLA Thresholds	95
Threshold Time Line	97

Table of Contents

Configure an SLA Object	98
Chapter 5: Configuring Charts and Dashboards	100
Available Dashboards	101
Customize the LDAP Dashboard	103
Debug Dashboard Customization	104
Preserve Customized Files	104
The Chart Builder Tool	105
Chart Presentation Details	106
Chart Builder Parameters	107
Chart Properties File	108
Available Charts for PingData Servers	108
Charts for All Servers	109
PingDirectory Server Charts	109
PingDirectoryProxy Server Charts	109
PingDataSync Server Charts	109
PingDataMetrics Server Charts	110
PingDataGovernance Server Charts	110
Velocity Templates	110
Supporting Multiple Content Types	113
Velocity Context Providers	114
Velocity Tools Context Provider	114
Chapter 6: Troubleshooting	116
Use the collect-support-data tool	117
Slow Queries Based on Sample Cache Size	117
Insufficient memory errors	118
Unexpected Query Results	118

Conditions for automatic server shutdown	119
Problems with SSL communication	119
Installation and maintenance issues	119
The setup program will not run	120
The server will not start	121
The server has shutdown	123
The server will not accept client connections	124
The server is unresponsive	124
Problems with the Administrative Console	125
Chapter 7: PingDataMetrics Server API Reference	127
Connection and Security	128
Secure Error Messages	129
Response Codes	129
List Monitored Instances	130
EXAMPLES	130
Retrieve Monitored Instance	131
EXAMPLE:	132
List Available Metrics	132
EXAMPLES	134
Retrieve a Metric Definition	135
EXAMPLE	135
Perform a Metric Query	136
Data Set Structure	138
Google Chart Tools Datasource Protocol	140
Access Alerts	142
Retrieving Event Types	142

Table of Contents

Retrieving Events	142
LDAP SLA	143
Retrieving the SLA Object	144
EXAMPLE	145
Pagination	146
Index	147

Chapter 1: Introduction

The PingDataMetrics Server collects performance data from the Ping Data Platform.

Topics include:

[PingDataMetrics Server Overview](#)

[PingDataMetrics Server Components](#)

[Data Collection](#)

[Charts and Dashboards](#)

[PostgreSQL DBMS Details](#)

PingDataMetrics Server Overview

The PingDataMetrics Server provides insight into the transactions and performance of the Ping Data Platform. The PingDataMetrics Server collects data from configured instances and replicas of the PingDirectory Server, the PingDirectoryProxy Server, the PingDataSync Server, and the PingDataGovernance Server. Data collected from the PingDataMetrics Server enables:

- Measuring the performance of the identity infrastructure as a whole service, not a collection of individual servers.
- Identifying client applications that require the greatest amount of resources.
- Determining which servers have the most available resources to handle requests.
- Predicting the capacity and needs of the identity infrastructure to plan for increased traffic.
- Analyzing all aspects of the identity infrastructure for troubleshooting performance issues.

PingDataMetrics Server Components

The PingDataMetrics Server consists of the following components:

PingDataMetrics Server – A stand-alone server that relies on the PostgreSQL database for collected metrics. The PingDataMetrics Server gathers data for itself and configured PingData servers.

Metrics API – A REST API that provides access to collected metrics data. The API is accessible over HTTPS and supports multiple management parameters including filtering, averaging, and setting ranges for multiple data sets.

query-metric tool – The primary command-line tool for metric data access. This tool can also be used for scripted automation of extracting data from the PingDataMetrics Server. An explore option enables custom queries and additions to charts and dashboards.

SNMP access – System-level metrics can be accessed over SNMP.

Data Set – A proprietary data structure that is designed for interoperability with charting libraries such as Highcharts, or FusionCharts.

Charts, Chart Builder, and Dashboard Templates – Tools for customizable, web-based metrics charts and dashboards.

Data Collection

The PingDataMetrics Server collects data from all monitored servers through LDAP queries to the server's backend. Each monitored server collects and stores a limited history of data locally. Data includes system status and performance information. To collect data, the PingDataMetrics Server regularly polls all monitored servers for data that is stored in time-contiguous blocks, gathers the recent data, and stores data in a PostgreSQL database. Polling has minimal impact on the monitored servers.

Performance Data

The majority of information collected represents the performance of the monitored server. Each monitored server should be configured to enable the PingDataMetrics Server to adequately keep up with the flow. Performance data represents multiple dimensions of a metric. For example, a response time metric can represent the request type, time to respond, the application that made the request, and the action that was taken.

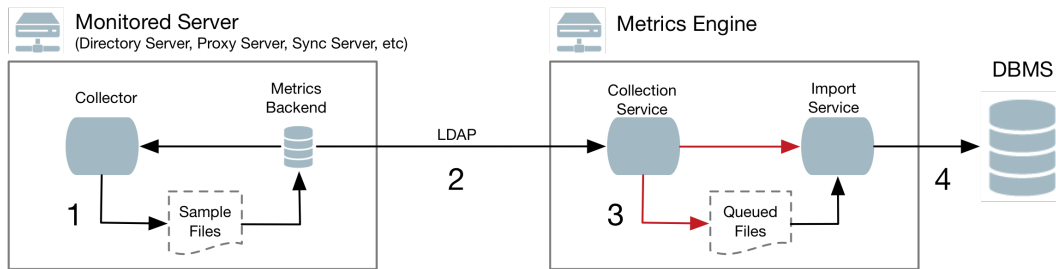
System and Status Data

All servers configured to be monitored by the PingDataMetrics Server store server and host system data. Server and machine metrics are retrieved from the `cn=monitor` backend of the monitored server.

The Stats Collector plugin is responsible for collecting performance data from the `cn=monitor` backend. Data includes server responses, replication activity, local database activity, and host system metrics. Stats Collector configuration defines:

- Data sample and collection intervals.
- The granularity of data collected (basic, extended, or verbose).
- The types of host system data collected such as CPU, disk, and network.
- The type of data aggregation that occurs for LDAP application statistics.

See [Tuning Data Collection](#) for more information. The following illustrates the data collection process:



Data Collection

Data collection occurs in the following steps:

1. Data samples are taken and stored in time-contiguous blocks on the disk of the monitored server.
2. The PingDataMetrics Server collection service polls for new sample blocks.
3. The new sample blocks are queued to disk on the PingDataMetrics Server.
4. The PingDataMetrics Server import service loads new blocks into the database.

Charts and Dashboards

The PingDataMetrics Server provides a number of charts and dashboards to display metrics information. A Chart Builder tool enables configuring charts on an HTML page and saving the properties for use in a dashboard. Several charts are provided for general system information and specific PingData server functions. All dashboards are viewed from the PingDataMetrics Server.

PostgreSQL DBMS Details

The PingDataMetrics Server uses a PostgreSQL DBMS to store data, which is included in the installation. This is a traditional table-based DBMS best suited for tabular data. The PingDataMetrics Server interacts with the DBMS in four ways:

- **Data import** – Import places steady write load on the DBMS and accounts for 80% of the writes. This single-threaded interaction puts a lock on the target table. A PingDataMetrics Server that monitors 20 servers keeps a single 10K RPM disk 70% busy with this single interaction.

- **Data aggregation** – Data aggregation places a less frequent read/write load on the DBMS. This interaction is responsible for the aggregation of the data samples from one time resolution to the next, reading from one set of tables and writing to another set. Sample aggregation uses no table-level locks and the ratio of records between read:write is between 60:1 and 24:1.
- **Data sample age-out** – Sample age-out occurs at regular intervals and results in a table being dropped and/or added. Age-out occurs every 30 minutes, though some intervals may drop and/or add more than one table.
- **Data query** – Sample queries occur when clients request metric samples from the public API. The API can aggregate multiple dimensions and multiple servers in a single request. A single request may fetch several million rows from the DBMS, though it only returns a few hundred data points to the client. Samples from previous queries are cached by the PingDataMetrics Server, but initial queries for a given metric may take several seconds and result in a large amount of disk read activity.

Over time, the storage of samples in the data tables is optimized to match the access patterns of the queries. However, the public API supports queries where the results are the aggregate of thousands of different dimension sets, and each dimension set may have thousands of samples within the time range of the query. For example, a query about the throughput of all PingDirectory Server and PingDirectoryProxy Servers for all applications and all LDAP operations over the last 72 hours might result in four to six million DBMS records being read into memory, aggregated, and finally reduced to 100 data values. The results from each query are cached so that a subsequent request for the same data results in less DBMS activity. Both disk seek time and rotational delay impact the performance of a first-time query, so disks with faster RPM speeds provide a measurable improvement for first-time queries.

Chapter 2: Installing the PingDataMetrics Server

This section describes how to install and run the PingDataMetrics Server. It includes pre-installation requirements and considerations.

Topics include:

[Supported Platforms](#)

[Installing Java](#)

[Creating a Non-Root user](#)

[Optimize the Linux Operating System](#)

[Configuring the PingData Servers to Gather Metrics](#)

[Ping License Keys](#)

[Installing the PingDataMetrics Server](#)

[Logging into the Administrative Console](#)

[Server Folders and Files](#)

[Adding Monitored Servers to the PingDataMetrics Server](#)

[Removing Monitored Servers](#)

[Start and Stop the PingDataMetrics Server Server](#)

[Uninstall the PingDataMetrics Server](#)

[Update Servers in a Topology](#)

[Revert an update](#)

[Administrative Accounts](#)

Supported Platforms

The following platforms and versions are supported for this release.

Operating systems	Virtualization platforms	Java versions
RedHat 6.6	VMWare vSphere & ESX 6.0	OpenJDK 8.x 64-bit
RedHat 6.8	KVM	OpenJDK 11.x 64-bit
RedHat 6.9	Amazon EC2	Oracle JDK 8.x 64-bit
RedHat 7.4	Microsoft Azure (Supported by Professional Services)	Oracle JDK 11.x 64-bit
RedHat 7.5		
CentOS 6.8		
CentOS 6.9		
CentOS 7.4		
CentOS 7.5		
SUSE Enterprise 11 SP4		
SUSE Enterprise 12 SP3		
Ubuntu 16.04 LTS		
Ubuntu 18.04 LTS		
Amazon Linux		
Windows Server 2012 R2		
Windows Server 2016		

Note

It is highly recommended that a Network Time Protocol (NTP) system be in place so that multi-server environments are synchronized and timestamps are accurate.

Install the JDK

The Java 64-bit JDK is required on the server. Even if Java is already installed, create a separate Java installation for use by the server to ensure that updates to the system-wide Java installation do not inadvertently impact the installation.

Configure a Non-Root User

The PingDataMetrics Server installer cannot be run as the root user, and generally the PingDataMetrics Server (and PostgreSQL) should not be run as root. As a non-root user, network port numbers below 1024 cannot be used.

In general, this account will need the ability to do the following:

- Listen on privileged network ports.
- Bypass restrictions on resource limits.

For security, the account should be restricted from the following:

- The ability to see processes owned by other users on the system.
- The ability to create hard links to files owned by other users on the system.

Optimize the Linux Operating System

Configure the Linux filesystem by making the following changes.

Note

The server explicitly overrides environment variables like `PATH`, `LD_LIBRARY_PATH`, and `LD_PRELOAD` to ensure that settings used to start the server do not inadvertently impact its behavior. If these variables must be edited, set values by editing the `set_environment_vars` function of the `lib/_script-util.sh` script. Stop and restart the server for the change to take effect.

Set the file descriptor limit

The server allows for an unlimited number of connections by default, but is restricted by the file descriptor limit on the operating system. If needed, increase the file descriptor limit on the operating system with the following procedure.

Note

If the operating system relies on `systemd`, refer to the Linux operating system documentation for instructions on setting the file descriptor limit.

1. Display the current hard limit of the system. The hard limit is the maximum server limit that can be set without tuning the kernel parameters in the `proc` filesystem.

```
ulimit -aH
```

2. Edit the `/etc/sysctl.conf` file. If the `fs.file-max` property is defined in the file, make sure its value is set to at least 65535. If the line does not exist, add the following to the end of the file:

```
fs.file-max = 65535
```

3. Edit the `/etc/security/limits.conf` file. If the file has lines that set the soft and hard limits for the number of file descriptors, make sure the values are set to 65535. If the lines are not present, add the following lines to the end of the file (before `#End of file`). Insert a tab between the columns.

Chapter 2: Installing the PingDataMetrics Server

```
* soft nofile 65535
* hard nofile 65535
```

4. Reboot the system, and then use the `ulimit` command to verify that the file descriptor limit is set to 65535 with the following command:

```
ulimit -n
```

Once the operating system limit is set, the number of file descriptors that the server will use can be configured by either using a `NUM_FILE_DESCRIPTOR` environment variable, or by creating a `config/num-file-descriptors` file with a single line such as, `NUM_FILE_DESCRIPTOR=12345`. If these are not set, the default of 65535 is used. This is strictly optional if wanting to ensure that the server shuts down safely prior to reaching the file descriptor limit.

Note

For RedHat 7 or later, modify the `20-nproc.conf` file to set both the open files and max user processes limits:

```
/etc/security/limits.d/20-nproc.conf
```

Add or edit the following lines if they do not already exist:

```
* soft nproc 65536
* soft nofile 65536
* hard nproc 65536
* hard nofile 65536
root soft nproc unlimited
```

Set the filesystem flushes

Linux systems running the ext3 filesystem only flush data to disk every five seconds. If the server is on a Linux system, edit the mount options to include the following:

```
commit=1
```

This variable changes the flush frequency from five seconds to one. Also, set the flush frequency in the `/etc/fstab` file to make sure the configuration remains after reboot.

Install sysstat and pstack on Red Hat

The server troubleshooting tool `collect-support-data` relies on the `iostat`, `mpstat`, and `pstack` utilities to collect monitoring, performance statistics, and stack trace information on the server's processes. For Red Hat systems, make sure that these packages are installed, for example:

```
$ sudo yum install sysstat gdb dstat -y
```

Install the dstat utility

The `dstat` utility is used by the `collect-support-data` tool.

Disable filesystem swapping

It is recommended that any performance tuning services like `tuned` be disabled. As root, change the current value in the operating system and by adding a line `vm.swappiness = 0` to `/etc/sysctl.conf` to ensure that the correct setting is applied when the system restarts.

If performance tuning is required, `vm.swappiness` can be set by cloning the existing performance profile then adding `vm.swappiness = 0` to the new profile's `tuned.conf` file in `/usr/lib/tuned/profile-name/tuned.conf`. The updated profile is then selected by running `tuned-adm profile customized_profile`.

Manage system entropy

Entropy is used to calculate random data that is used by the system in cryptographic operations. Some environments with low entropy may have intermittent performance issues with SSL-based communication. This is more typical on virtual machines, but can occur in physical instances as well. Monitor the `kernel.random.entropy_avail` in `sysctl` value for best results.

If necessary, update `$JAVA_HOME/jre/lib/security/java.security` to use `file:/dev/./urandom` for the `securerandom.source` property.

Set Filesystem Event Monitoring (inotify)

An event monitoring tool such as `inotify` can be configured for notifying processes about filesystem events (including file creation, deletion, and updates). The Linux system puts a limit on the number of `inotify` watches assigned to each user. To increase the limit, edit `etc/sysctl.conf` to add a line:

```
fs.inotify.max_user_watches = 524288
```

Run the command:

```
$ sudo sysctl -w fs.inotify.max_user_watches=524288
```

Tune IO scheduler

Using the correct IO scheduler can increase performance and reduce the possibility of database timeouts when the system is under extreme write load. For file systems running on an SSD, or in a virtualized environment, the `noop` scheduler is recommended. For all other systems, the `deadline` scheduler is recommended. To determine which scheduler is configured on your system, run this command:

```
$ cat /sys/block/<block-device>/queue/scheduler
```

For example:

```
$ cat /sys/block/sda/queue/scheduler
```

Changing the scheduler on a running system can be done with the following command:

```
$ echo 'deadline' > /sys/block/sda/queue/scheduler
```

The change will take effect after the system is restarted. The procedure for configuring a scheduler to use at startup depends on the version of Linux. See the Linux documentation for your specific version for the correct way to configure this setting.

Enable the server to listen on privileged ports

Linux provides 'capabilities' used to grant specific commands the ability to do things that are normally only allowed for a root account. Instead of granting the ability to a specific user, capabilities are granted to a specific command. It may be convenient to enable the server to listen on privileged ports while running as a non-root user.

The `setcap` command is used to assign capabilities to an application. The `cap_net_bind_service` capability enables a service to bind a socket to privileged ports (port numbers less than 1024). If Java is installed in `/ds/java` (and the Java command to run the server is `/ds/java/bin/java`), the Java binary can be granted the `cap_net_bind_service` capability with the following command:

```
$ sudo setcap cap_net_bind_service=+eip /ds/java/bin/java
```

The java binary needs an additional shared library (`libjli.so`) as part of the Java installation. More strict limitations are imposed on where the operating system will look for shared libraries to load for commands that have capabilities assigned. So it is also necessary to tell the operating system where to look for this library. This can be done by creating the file

`/etc/ld.so.conf.d/libjli.conf` with the path to the directory that contains the `libjli.so` file. For example, if the Java installation is in `/ds/java`, the contents of that file should be:

```
/ds/java/lib/amd64/jli
```

Run the following command for the change to take effect:

```
$ sudo ldconfig -v
```

Configure Servers to be Monitored

Before installing the PingDataMetrics Server, configure the servers to be monitored:

- PingDirectory Server
- PingDirectoryProxy Server
- PingDataSync Server
- PingDataGovernance Server

The monitored servers require sufficient disk space to store the monitoring data, and can be configured with Tracked Applications if there are specific application bind DNs that should be monitored.

Disk Space Requirements and Monitoring Intervals

The metrics backend on the monitored servers is responsible for the temporary storage of metric data, and is configured to keep a maximum amount of metric history based on log retention policies, which are configured with the `dsconfig` tool.

The default retention policies define a cap on disk space usage, which in turn determines the amount of metric history retained. If the PingDataMetrics Server is stopped for a period of time, the monitored servers should be configured to retain enough metrics history to prevent gaps in data when the PingDataMetrics Server restarts. The amount of disk space required for metrics history can also depend on the monitored server's Stats Collector Plugin settings. In general, 500MB is enough to retain an eight-hour span of metrics history.

The value of the `sample-flush-interval` property of the monitored server's metrics backend determines the maximum delay between when a metric is captured and when it can be picked up by the PingDataMetrics Server. The flush interval can be set between 15 and 60 seconds, with longer values resulting in less processing load on the PingDataMetrics Server. However,

this flush interval increases the latency between when the metric was captured and when it becomes visible in a chart or dashboard. Changing the `sample-flush-interval` attribute to 60 seconds, has the PingDataMetrics Server keep 2000 minutes of history.

The number of metrics produced per unit of time varies based on the configuration. No formula can be provided to compute exact storage required for each hour of history. In general, 60MB per hour is a standard estimate.

Tracked Applications

If the PingDataMetrics Server will monitor client applications associated with the monitored servers, the Tracked Applications feature should be configured for monitored servers as well. Activity performed by a particular LDAP Bind DN can be associated with a PingDataMetrics Server application-name, which in turn can be included in PingDataMetrics Server SLA definitions.

The Processing Time Histogram plugin is configured on each PingDirectory Server and PingDirectoryProxy Server as a set of histogram ranges. These ranges should be defined identically across all monitored servers. For each monitored server, set the `separate-monitor-entry-per-tracked-application` property of the processing time histogram plugin to `true`. Per-application monitoring information will appear under `cn=monitor`. The `per-application-ldap-stats` property must also be set to `per-application-only` in the Stats Collector Plugin. See the *Ping Identity PingDirectory Server Administration Guide* for Tracked Application configuration details.

The following sets the required property of the Processing Time Histogram plugin:

```
$ bin/dsconfig set-plugin-prop \  
  --plugin-name "Processing Time Histogram" \  
  --set separate-monitor-entry-per-tracked-application:true
```

The following example sets the required property of the Stats Collector plugin:

```
$ bin/dsconfig set-plugin-prop \  
  --plugin-name "Stats Collector" \  
  --set per-application-ldap-stats:per-application-only
```


Ping license keys

License keys are required to install all Ping products. Obtain licenses through Salesforce or from <https://www.pingidentity.com/en/account/request-license-key.html>.

- A license is always required for setting up a new single server instance and can be used site-wide for all servers in an environment. When cloning a server instance with a valid license, no new license is needed.
- A new license must be obtained when updating a server to a new major version, for example from 6.2 to 7.0. Licenses with no expiration date are valid until the server is upgraded to the next major version. A prompt for a new license is displayed during the update process.
- A license may expire on particular date. If a license does expire, obtain a new license and install it using `dsconfig` or the Administrative Console. The server will provide a notification as the expiration date approaches. License details are available using the server's `status` tool.

When installing the server, specify the license key file in one of the following ways:

- Copy the license key file to the server root directory before running `setup`. The interactive `setup` tool will discover the file and not require input. If the file is not in the server root, the `setup` tool will prompt for its location.
- If the license key is not in the server root directory, specify the `--licenseKeyFile` option for non-interactive `setup`, and the path to the file.

Install the Server

Use the `setup` tool to install the server. The server needs to be started and stopped by the user who installed it.

Note

A Windows installation requires that the Visual Studio 2010 runtime patch be installed prior to running the `setup` command.

1. Log in as a user, other than root.
2. Obtain the latest zip release bundle from Ping Identity and unpack it in a directory owned by this user.

```
$ unzip PingData<server>-<version>.zip
```

3. Change to the server root directory.

```
$ cd PingData<server>
```

Chapter 2: Installing the PingDataMetrics Server

4. Run the `setup` command.

```
$ ./setup
```

5. Type **yes** to accept the End-User License Agreement and press **Enter** to continue.
6. Read the installation process and prerequisites. Press **Enter** to continue.
7. Type the port number of for the PostgreSQL database instance to use to store monitoring data, or press **Enter** to accept the default port.
8. Enter the directory to be used for PostgreSQL data files, or press **Enter** to accept the default (`pgsql_data`). If the name entered is a relative path name, it will be created in the current working directory.
9. Enter a name for the database administrative account, or press **Enter** to accept the default. The `setup` tool will create a user (role) and database to be used by the PingDataMetrics Server. These credentials are strictly for use by this tool during this session and are not retained.
10. Enter and confirm a password for this account.
11. Specify the name of the PostgreSQL account to be associated with the PingDataMetrics Server historical monitoring data, or press **Enter** to accept the default (`metrics`). The tool will create this user account using the administrative account specified in step 9.
12. The password generated for this account is `metrics`, press **Enter** to accept the default, or type **yes** and provide and confirm a new password.
13. Enter the fully-qualified host name for the server, or press **Enter** to accept the default.
14. Create the initial root user DN for the server, or press **Enter** to accept (`cn=Directory Manager`)
15. Enter and confirm a password for this account.
16. Enter the port for HTTPS connection to the Platform (SCIM and the Configuration) APIs, or press **Enter** to accept the default.
17. Enter the port on which the PingDataMetrics Server will accept LDAP client connections, or press **Enter** to accept the default.
18. To enable LDAPS, type **yes**, or press **Enter** to accept the default **no**.
19. If LDAPS is enabled, enter the port on which the server will accept LDAPS client connections, or press **Enter** to accept the default (`2636`).
20. To enable StartTLS, type **yes**, or press **Enter** to accept the default **no**.
21. Choose a certificate option for the server.

```
Certificate server options:
```


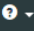
- 1) Generate self-signed certificate (recommended for testing purposes only)
- 2) Use an existing certificate located on a Java KeyStore (JKS)

- 3) Use an existing certificate located on a PKCS12 KeyStore
- 4) Use an existing certificate on a PKCS11 token

Depending on the option chosen, other information may be needed. If the Java or the PKCS#12 KeyStore is chosen, the KeyStore path and PIN is needed. If the PKCS#11 token is chosen, the key PIN is needed.

22. Choose the desired encryption for the directory data, backups, and log files from the choices provided:
 - Encrypt data with a key generated from an interactively provided passphrase. Using a passphrase (obtained interactively or read from a file) is the recommended approach for new deployments, and you should use the same encryption passphrase when setting up each server in the topology.
 - Encrypt data with a key generated from a passphrase read from a file.
 - Encrypt data with a randomly generated key. This option is primarily intended for testing purposes, especially when only testing with a single instance, or if you intend to import the resulting encryption settings definition into other instances in the topology.
 - Encrypt data with an imported encryption settings definition. This option is recommended if you are adding a new instance to an existing topology that has older server instances with data encryption enabled.
 - Do not encrypt server data.
23. Choose an option to assign the amount of memory that the server should allocate to the PingDataMetrics Server and press **Enter**.
24. Press **Enter** (*yes*) to start the server when configuration is complete.
25. Press **Enter** to install the PingDataMetrics Server with the defined parameters.

After the PingDataMetrics Server server is installed, access the Metrics landing page (<https://<host>:<HTTPS-port>/view/index>) for access to the default dashboards, chart builder tool, and online documentation.

 **Data Metrics Server** 

The Ping Identity Data Metrics Server includes these customizable dashboards for monitoring server operations and for analysis of the activity, performance, and capacity of a PingData deployment.

[Documentation](#) [List of Available Metrics](#) [REST API Documentation](#) [Try the Chart Builder](#)

Server Information
velocity/templates/server-info.vm
A simple template that displays basic server data.

SLA Viewer
velocity/templates/sla-viewer.vm
Displays the current operation response time and throughput, with drill-down to historical values. Use this template for monitoring adherence to the Service Level Agreements that are defined for critical applications.

Topology Dashboard
velocity/templates/dap-dashboard.vm
Use this template for monitoring the entire PingData deployment. Includes operational metrics at the individual server, data center, and aggregate level by client application and service.

OAuth2 Dashboard velocity/templates/oauth2-dashboard.vm	Consumer Dashboard velocity/templates/consumer.vm	Enforcement Dashboard velocity/templates/enforcement.vm
---	---	---

Three variations of an example dashboard—showing trends and metrics collected from the Data Governance Broker. These provide details about data use and policy enforcement, such as trends in consumer consent, applications accessing the most data, and the policies used most often—and to allow or deny data requests.

Log into the Administrative Console

After the server is installed, access the Administrative Console, `https://<host>/console/login`, to verify the configuration and manage the server. The root user DN or the common name of a root user DN is required to log into the Administrative Console. For example, if the DN created when the server was installed is `cn=Directory Manager, directory manager` can be used to log into the Administrative Console.

If the Administrative Console needs to run in an external container, such as Tomcat, a separate package can be installed according to that container's documentation. Contact Ping Identity Customer Support for the package location and instructions.

Server folders and files

After the distribution file is unzipped, the following folders and command-line utilities are available:

Directories/Files/Tools	Description
ldif	Stores any LDIF files that you may have created or imported.
import-tmp	Stores temporary imported items.
classes	Stores any external classes for server extensions.
db	For the PingDirectory Server, this is where its Berkeley DB files reside.
bak	Stores the physical backup files used with the backup command-line tool.
velocity	Stores Velocity templates that define the server's application pages.
update.bat, and update	The update tool for UNIX/Linux systems and Windows systems.
uninstall.bat, and uninstall	The uninstall tool for UNIX/Linux systems and Windows systems.
ping_logo.png	The image file for the Ping Identity logo.
setup.bat, and setup	The setup tool for UNIX/Linux systems and Windows systems.
revert-update.bat, and revert-update	The revert-update tool for UNIX/Linux systems and Windows systems.
README	README file that describes the steps to set up and start the server.
License.txt	Licensing agreement for the product.
legal-notice	Stores any legal notices for dependent software used with the product.
docs	Provides the release notes, Configuration Reference (HTML), API Reference, and all other product documentation.
metrics	Stores the metrics that can be gathered for this server and surfaced in the PingDataMetrics Server.
bin	Stores UNIX/Linux-based command-line tools.
bat	Stores Windows-based command-line tools.
lib	Stores any scripts, jar files, and library files needed for the server and its extensions.
collector	Used by the server to make monitored statistics available to the PingDataMetrics Server.
locks	Stores any lock files in the backends.
postgres	Stores PostgreSQL files.
tmp	Stores temporary files.
resource	Stores the MIB files for SNMP and can include ldif files, make-ldif templates, schema files, dsconfig batch files, and other items for configuring or managing the server.
config	Stores the configuration files for the backends (admin, config) as well as the directories for messages, schema, tools, and updates.
logs	Stores log files.

Add Monitored Servers to the PingDataMetrics Server

Configure the PingDataMetrics Server to monitor servers using the `monitored-servers` tool.

Using the `monitored-servers` Tool

The `monitored-servers` command-line tool configures communication between the servers and the PingDataMetrics Server, then adds external server definitions to the PingDataMetrics Server based on the server's administrative data. Before a server is added to the PingDataMetrics Server configuration, the system determines whether communication needs to be configured. If so, the `cn=Monitoring User` root user account is created on the external server.

Running the tool with the `add-servers` subcommand creates an external server based on the information discovered about the remote server. It also uses the information located in the `cn=admin` data entry to discover other servers in the topology, which are also added to the configuration.

The following examples use the `monitored-servers` tool:

- Run the `monitored-servers` tool with the `add-servers` subcommand. Specify connection information for the PingDataMetrics Server, as well as connection information for any remote servers in use.

```
$ bin/monitored-servers add-servers \  
  --bindDN uid=admin,dc=example,dc=com \  
  --bindPassword password \  
  --monitoringUserBindPassword password \  
  --remoteServerHostname localhost \  
  --remoteServerPort 1389 \  
  --remoteServerBindPassword password
```

- Use the `--dry-run` option to generate output detailing the work that would be done in a live session without actually making changes to the server configuration.

```
$ bin/monitored-servers add-servers \  
  --bindDN uid=admin,dc=example,dc=com \  
  --bindPassword password \  
  --monitoringUserBindPassword password \  
  --remoteServerHostname localhost \  
  --remoteServerPort 1389 \  
  --remoteServerBindPassword password \  
  --dry-run
```

Remove Monitored Servers

Use the `monitored-servers` tool to remove servers from the PingDataMetrics Server:

1. List the monitored servers, with the `/bin/status` tool.
2. Choose a server from the Monitored Servers list. Remove the server by name with the following command:

```
bin/monitored-servers remove-server "Austin Instance 1"
```

Start and Stop the Server

When the PingDataMetrics Server starts for the very first time, it downloads new samples from the monitored servers and adds data to the database. Until it has finished this first data collection, the PingDataMetrics Server will not be able to answer metric queries to the database. The PingDataMetrics Server processes samples from the oldest to the newest, so queries on more recent data may require more start-up time. If the monitored servers have been collecting samples for several days, there may be a significant backlog of data to collect.

To determine if the server is ready to respond to metric queries, run the `status` tool. If the `Sample Import Backlog` property is zero (0), the server is ready.

Note

The PingDataMetrics Server needs to be started and stopped by the user who installed it. If the PingDataMetrics Server is started or stopped by a different user, the following error is listed in the `postgres.log` file when Postgres starts:

```
FATAL: role "<username>" does not exist
```

Start the PingDataMetrics Server as a Background Process

Navigate to the server root directory, and run the following command:

```
$ bin/start-server
```

For Windows systems:

```
$ bat/start-server
```

Start the PingDataMetrics Server as a Foreground Process

Navigate to the server root directory, and run the following command:

```
$ bin/start-server --nodetach
```

Start the PingDataMetrics Server at Boot Time

By default, the PingDataMetrics Server does not start automatically when the system is booted. To configure the monitoring server to start automatically when the system boots, use the `create-rc-script` tool to create a run control script as follows:

1. Create the startup script as the non-root PingDataMetrics Server user. In this example `ds` is the user.

```
$ bin/create-rc-script --outputFile Ping-Identity-ME.sh \  
--userName ds
```

2. Log in as root, move the generated `Ping-Identity-ME.sh` script into the `/etc/init.d` directory, and create symlinks to it from the `/etc/rc3.d` (starting with an "S" to start the server) and `/etc/rc0.d` directory (starting with a "K" to stop the server).

```
# mv Ping-Identity-ME.sh /etc/init.d/  
# ln -s /etc/init.d/Ping-Identity-ME.sh /etc/rc3.d/S50-Ping-Identity-ME.sh  
# ln -s /etc/init.d/Ping-Identity-ME.sh /etc/rc0.d/K50-Ping-Identity-ME.sh
```

Stop the PingDataMetrics Server

Navigate to the server root directory, and run the following command:

```
$ bin/stop-server
```

Restart the PingDataMetrics Server

Restart the PingDataMetrics Server using the `--restart` or `-R` option. Running this command is equivalent to shutting down the server, exiting the JVM session, and then starting up again, which requires a re-priming of the JVM cache.

Navigate to the server root directory, and run the following command:

```
$ bin/stop-server --restart
```

Uninstall the server

Use the `uninstall` command-line utility to uninstall the server using either interactive or non-interactive modes. Interactive mode provides options, progress, and a list of the files and directories that must be manually deleted if necessary.

Non-interactive mode, invoked with the `--no-prompt` option, suppresses progress information, except for fatal errors. All options for the `uninstall` command are listed with the `--help` option.

The `uninstall` command must be run as either the root user or the user account that installed the server.

Perform the following steps to uninstall in interactive mode:

1. Navigate to the server root directory.

```
$ cd PingData<server>
```

2. Start the uninstall command:

```
$ ./uninstall
```

3. Select the components to be removed, or press **Enter** to remove all components.
4. If the server is running, press **Enter** to shutdown the server before continuing.
5. Manually remove any remaining files or directories, if required.

Update servers in a topology

An update to the current release includes significant changes, and the introduction of a topology registry, which will store information previously stored in the admin backend (server instances, instance and secret keys, server groups, and administrator user accounts). For the admin backend to be migrated, the `update` tool must be provided LDAP authentication options to the peer servers of the server being updated.

The LDAP connection security options requested (either plain, TLS, StartTLS, or SASL) must be configured on every server in the topology. The LDAP credentials must be present on every server in the topology, and must have permissions to read from the admin backend and the config backend of every server in the topology. For example, a root DN user with the `inherit-default-privileges` set to true (such as the `cn=Directory Manager` user) that exists on every server can be used.

After enabling or fixing the configuration of the LDAP connection handler(s) to support the desired connection security mechanism on each server, run the following `dsframework` command on the server being updated so that its admin backend has the most up-to-date information:

```
$ bin/dsframework set-server-properties \
--serverID serverID \
--set ldapport:port \
--set ldapsport:port \
--set startTLSEnabled:true
```

Chapter 2: Installing the PingDataMetrics Server

The `update` tool will verify that the following conditions are satisfied on every server in the topology before allowing the update:

- When the first server is being updated, all other servers in the topology must be online. When updating additional servers, all topology information will be obtained from one of the servers that has already been updated. The `update` tool will connect to the peer servers of the server being updated to obtain the necessary information to populate the topology registry. The provided LDAP credentials must have read permissions to the config and admin backends of the peer servers.
- The instance name is set on every server, and is unique across all servers in the topology. The instance name is a server's identifier in the topology. After all servers in the topology have been updated, each server will be uniquely identified by its instance name. Once set, the name cannot be changed. If needed, the following command can be used to set the instance name of a server prior to the update:

```
$ bin/dsconfig set-global-configuration-prop \  
  --set instance-name:uniqueName
```

- The cluster-wide configuration is synchronized on all servers in the topology. Older versions have some topology configuration under `cn=cluster,cn=config` (JSON attribute and field constraints). These items did not support mirrored cluster-wide configuration data. An update should avoid custom configuration changes on a server being overwritten with the configuration on the mirrored subtree master. To synchronize the cluster-wide configuration data across all servers in the topology, run the `config-diff` tool on each pair of servers to determine the differences, and use `dsconfig` to update each instance using the `config-diff` output. For example:

```
$ bin/config-diff --sourceHost hostName \  
  --sourcePort port \  
  --sourceBindDN bindDN \  
  --sourceBindPassword password \  
  --targetHost hostName \  
  --targetPort port \  
  --targetBindDN bindDN \  
  --targetBindPassword password
```

If any of these conditions are not satisfied, the `update` tool will list all of the errors encountered for each server, and provide instructions on how to fix them.

Update the server

This procedure assumes that an existing version of the server is stored at `PingData-server-old`. Make sure a complete, readable backup of the existing system is available before upgrading the server. Also, make sure there is a clear backout plan and schedule.

1. Download the latest version of the server software and unzip the file. For this example, the new server is located in the `PingData-server-new` directory.
2. Use the `update` tool of the newly unzipped build to update the server. Make sure to specify the server instance that is being upgrading with the `--serverRoot` option. The server must be stopped for the update to be applied.

Reverting an Update

If necessary, a server can be reverted to the previous version using the `revert-update` tool. The tool accesses a log of file actions taken by the `update` tool to put the filesystem back to its prior state. If multiple updates have been run, the `revert-update` tool can be used multiple times to revert to each prior update sequentially. For example, the `revert-update` command can be run to revert to the server's previous state, then run again to return to its original state. The server is stopped during the revert-update process.

Note

Reverting an update is not supported for upgrades to version 7.0, due to the topology backend changes.

Use the `revert-update` tool in the server root directory revert back to the most recent version of the server:

```
$ PingData-server-old/revert-update
```

Revert an Update

Once the server has been updated, you can revert to the most recent version (one level back) using the `revert-update` tool. The `revert-update` tool accesses a log of file actions taken by the updater to put the filesystem back to its prior state. If you have run multiple updates, you can run the `revert-update` tool multiple times to revert to each prior update sequentially. You can only revert back one level. For example, if you have run the update twice since first installing the server, you can run the `revert-update` command to revert to its previous state, then run the `revert-update` command again to return to its original state.

Reverting from Version 7.x to a Version Prior to 7.0

Reverting from version 7.0 or later to a pre-7.0 version can be done using the `revert-update` command with some extra steps. This is also the case when updating or reverting from a pre-6.2.0.2 version to 6.2.0.2 or later. These steps are listed when the `update` and `revert-update`

Chapter 2: Installing the PingDataMetrics Server

tool are run as well. You may need to perform one or more of the following tasks, depending on your installation and configuration:

- When updating or reverting from 6.2.0.2 or later to a pre-6.2.0.2 version, indexes may need to be rebuilt. Older versions of the server use an incompatible format for Local DB Composite Indexes. To update a server with composite indexes in the previous format, delete these indexes and re-run the update. After the update is complete, recreate the indexes and use the `rebuild-index` tool to rebuild the indexes. The command for recreating an index will be in the "Undo" portion of the `logs/config-audit.log` file. If you wish to later revert to an older version, delete and recreate those composite indexes again after the revert has completed.
- When updating to 7.x for the first time, instance names will need to be set for each server in the topology if they were not previously set. This is done with the following `dsconfig` command:

```
$ bin/dsconfig --bindDN "cn=Directory Manager" \  
  --bindPassword secret \  
  --no-prompt set-global-configuration-prop \  
  --set instance-name:<name>
```

- Topology information such as server instances, instance and secret keys, server groups, and administrator users have moved to the topology portion of the configuration from the `admin` backend. As long as new servers are not added to the topology after this update, the `revert-update` command can be used to return to the previous version. However, if new servers are added, then the restored `admin` backend of this server will not contain information about the new servers, and the local server will not be able to communicate with any other servers in the topology. New servers should not be added to the topology if reverting this update is a possibility.
- If new servers were added to the topology after the update, the new servers must be temporarily removed from the topology until all servers have been reverted to the previous version.
- When a server is reverted to a pre-7.x version, any servers in the topology using the topology portion of the configuration (rather than the `admin` backend) will need to know that the reverted server was downgraded to the `admin` backend. This is done by running the following `dsconfig` command on one of the servers that has not been reverted:

```
$ bin/dsconfig set-server-instance-prop \  
  --instance-name <Reverted server instance name> \  
  --set server-version:<Version to which server is reverted>
```

- If the topology does not have a master server when this command is run, it will not succeed. In this case, one of the remaining updated servers in the topology must be made master with the following command. This will enable the chosen instance to run the first command successfully.

```
$ bin/dsconfig set-global-configuration-prop \
  --set force-as-master-for-mirrored-data:true
```

- The 7.x server version includes database changes that are not compatible with previous server versions (6.x or older). If you wish to later revert to an older version, the data must be exported to LDIF before performing the reversion. Re-import the data after the revert process has completed. In addition, the `changelogDb/` and `db/changelog/` directories in the reverted server root must be deleted after the revert has completed.

When starting up the server for the first time after a revert has been run, and the necessary extra steps have been completed, the server will display warnings about "offline configuration changes," but they are not critical and will not appear on subsequent startups.

To Revert to the Most Recent Server Version

Use `revert-update` in the server root directory revert back to the most recent version of the server.

```
$ <PingServer>-old/revert-update
```

Administrative accounts

Users that authenticate to the Configuration API or the Administrative Console are stored in `cn=Root DNs,cn=config`. The `setup` tool automatically creates one administrative account when performing an installation. Accounts can be added or changed with the `dsconfig` tool.

Change the administrative password

Root users are governed by the Root Password Policy and by default, their passwords never expire. However, if a root user's password must be changed, use the `ldappasswordmodify` tool.

1. Open a text editor and create a text file containing the new password. In this example, name the file `rootuser.txt`.

```
$ echo password > rootuser.txt
```

2. Use `ldappasswordmodify` to change the root user's password.

```
$ bin/ldappasswordmodify --port 1389 --bindDN "cn=Directory Manager" \
  --bindPassword secret --newPasswordFile rootuser.txt
```

3. Remove the text file.

```
$ rm rootuser.txt
```

Chapter 3: Managing the PingDataMetrics Server

There are several ways to manage the PingDataMetrics Server status and performance. The PingDataMetrics Server Administrative Console enables browser-based server management, the `dsconfig` tool enables command line management, and the Configuration API enables management by third-party interfaces.

Topics include:

[PingDataMetrics Server Error Logging](#)

[Backend Monitor Entries](#)

[Configure Alert Handlers](#)

[The Alerts Backend](#)

[System Alarms and Gauges](#)

[Backup the PingDataMetrics Server Database](#)

[Management Tools](#)

[Configure HTTP Connection Handlers](#)

[Topology Configuration](#)

[Using the Configuration API](#)

[Domain Name Service \(DNS\) Caching](#)

[IP Address Reverse Name Lookup](#)

[Configuring Traffic Through a Load Balancer](#)

[Configuring Authentication with a SASL External Certificate](#)

[Server SDK Extensions](#)

PingDataMetrics Server Error Logging

The PingDataMetrics Server provides logging for warnings, errors, or significant events that occur within the server. Log publishers rely on log rotation and retention policies.

Customization options for log publishers are available with the `dsconfig` command, (`bin/dsconfig` or `bat/dsconfig` on Windows).

Each log publisher must have at least one log rotation policy and log retention policy configured. Configure the log rotation policy for each log publisher. When a rotation limit is reached, the server rotates the current log and starts a new log.

Logging Retention Policies

Select retention configuration from the following:

Time Limit Rotation Policy – Rotates the log based on the length of time since the last rotation. Default implementations are provided for rotation every 24 hours and every seven days.

Fixed Time Rotation Policy – Rotates the logs every day at a specified time (based on 24-hour time). The default time is 2359.

Size Limit Rotation Policy – Rotates the logs when the file reaches the maximum size for each log. The default size limit is 100MB.

Never Rotate Policy – Used in a rare event that does not require log rotation.

Logging Rotation Policies

Select rotation configuration from the following:

File Count Retention Policy – Sets the number of log files for the PingDataMetrics Server to retain. The default file count is 10 logs. If the file count is set to 1, then the log will continue to grow indefinitely without being rotated.

Free Disk Space Retention Policy – Sets the minimum amount of free disk space. The default free disk space is 500MB.

Size Limit Retention Policy – Sets the maximum size of the combined archived logs. The default size limit is 500MB.

Custom Retention Policy – Create a new retention policy. This requires developing custom code to implement the log retention policy.

Never Delete Retention Policy – Used in a rare event that does not require log deletion.

Create Log Publishers

Create a new log publisher with `dsconfig`, either from the command line or in interactive mode. Retention and rotation policies must be configured for the log publisher. For more information about policy options, see [Configure Logs](#).

Note

Compression cannot be disabled or turned off once configured for the logger. Determine logging requirements, prior to creating and configuring them.

Perform the following steps to create a log publisher:

1. The following creates a log publisher with the `dsconfig` command that logs disconnect operations.

```
$ bin/dsconfig create-log-publisher \
  --type file-based-access --publisher-name "Disconnect Logger" \
  --set enabled:true \
  --set "rotation-policy:24 Hours Time Limit Rotation Policy" \
  --set "rotation-policy:Size Limit Rotation Policy" \
  --set "retention-policy:File Count Retention Policy" \
  --set log-connects:false \
  --set log-requests:false --set log-results:false \
  --set log-file:logs/disconnect.log
```

To configure compression on the logger, add this option to the previous command:

```
--set compression-mechanism: gzip
```

2. To view log publishers, enter the following command:

```
$ bin/dsconfig list-log-publishers
```

Error Log Publisher

The Error Log reports errors, warnings, and informational messages about events that occur during the course of the server's operation. Each entry in the error log records the following properties (some are disabled by default and must be enabled):

Chapter 3: Managing the PingDataMetrics Server

Time Stamp – Displays the date and time of the operation in the format

DD/Month/YYYY:HH:MM:SS <offset from UTC time>.

Category– Specifies the message category that is loosely based on the server components.

Severity – Specifies the message severity of the event, which defines the importance of the message in terms of major errors that need to be quickly addressed. The default severity levels are `fatal-error`, `notice`, `severe-error`, and `severe-warning`.

Message ID – Specifies the numeric identifier of the message.

Message – Stores the error, warning, or informational message.

The following example displays an error log for the PingDataMetrics Server. The log is enabled by default and is accessible in the `<server-root>/logs/errors` file.

```
[21/Oct/2012:05:15:23.048 -0500] category=RUNTIME_INFORMATION severity=NOTICE
msgID=20381715 msg="JVM Arguments: '-Xmx8g', '-Xms8g', '-XX:MaxNewSize=1g',
'-XX:NewSize=1g', '-XX:+UseConcMarkSweepGC', '-XX:+CMSConcurrentMTEnabled',
'-XX:+CMSParallelRemarkEnabled', '-XX:+CMSParallelSurvivorRemarkEnabled',
'-XX:+CMSScavengeBeforeRemark', '-XX:RefDiscoveryPolicy=1',
'-XX:ParallelCMSThreads=4', '-XX:CMSMaxAbortablePrecleanTime=3600000',
'-XX:CMSInitiatingOccupancyFraction=80', '-XX:+UseParNewGC', '-XX:+UseMembar',
'-XX:+UseBiasedLocking', '-XX:+UseLargePages', '-XX:+UseCompressedOops',
'-XX:PermSize=128M', '-XX:+HeapDumpOnOutOfMemoryError',
'-Dcom.unboundid.directory.server.scriptName=setup'"
[21/Oct/2012:05:15:23.081 -0500] category=EXTENSIONS severity=NOTICE
msgID=1880555611 msg="Administrative alert type=server-starting
id=4178daee-ba3a-4be5-8e07-5ba17bf30b71
class=com.unboundid.directory.server.core.MetricsEngine
msg='The Metrics Server is starting'"
[21/Oct/2012:05:15:23.585 -0500] category=CORE severity=NOTICE
msgID=1879507338 msg="Starting group processing for backend api-users"
[21/Oct/2012:05:15:23.586 -0500] category=CORE severity=NOTICE
msgID=1879507339 msg="Completed group processing for backend api-users"
[21/Oct/2012:05:15:23.586 -0500] category=EXTENSIONS severity=NOTICE
msgID=1880555575 msg="'Group cache (2 static group(s) with 0 total
memberships and 0 unique members, 0 virtual static group(s),
1 dynamic group(s))' currently consumes 7968 bytes and can grow to a maximum
of an unknown number of bytes"
[21/Oct/2012:05:16:18.011 -0500] category=CORE severity=NOTICE
msgID=458887 msg="The Metrics Server (Metrics Server 4.5.1.0
build 20121021003738Z, R12799) has started successfully"
```

Use `dsconfig` to modify the default File-Based Error Log, as in the following command:

```
$ bin/dsconfig set-log-publisher-prop \
  --publisher-name "File-Based Error Logger" \
  --set include-product-name:true --set include-instance-name:true \
  --set include-startup-id:true
```

Configure log file encryption

The server supports the ability to encrypt log files as they are written. The `encrypt-log` configuration property controls whether encryption will be enabled for the logger. Enabling encryption causes the log file to have an `.encrypted` extension (and if both encryption and compression are enabled, the extension will be `.gz.encrypted`). Any change that affects the name used for the log file could prevent older files from getting properly cleaned up.

Like compression, encryption can only be enabled when the logger is created. Encryption cannot be turned on or off once the logger is configured. For any log file that is encrypted, enabling compression is also recommended to reduce the amount of data that needs to be encrypted. This will also reduce the overall size of the log file. The `encrypt-file` tool (or custom code, using the LDAP SDK's

`com.unboundid.util.PassphraseEncryptedInputStream`) is used to access the encrypted data.

To enable encryption, at least one encryption settings definition must be defined in the server. Use the one created during setup, or create a new one with the `encryption-settings create` command. By default, the encryption will be performed with the server's preferred encryption settings definition. To explicitly specify which definition should be used for the encryption, the `encryption-settings-definition-id` property can be set with the ID of that definition. It is recommended that the encryption settings definition is created from a passphrase so that the file can be decrypted by providing that passphrase, even if the original encryption settings definition is no longer available. A randomly generated encryption settings definition can also be created, but the log file can only be decrypted using a server instance that has that encryption settings definition.

When using encrypted logging, a small amount of data may remain in an in-memory buffer until the log file is closed. The encryption is performed using a block cipher, and it cannot write an incomplete block of data until the file is closed. This is not an issue for any log file that is not being actively written. To examine the contents of a log file that is being actively written, use the `rotate-log` tool to force the file to be rotated before attempting to examine it.

The following commands can be used to set log file encryption:

1. Use `dsconfig` to enable encryption for a Log Publisher. In this example, the File-basedAccess Log Publisher "Encrypted Access" is created, compression is set, and

rotation and retention policies are set.

```
$ bin/dsconfig create-log-publisher-prop --publisher-name "Encrypted
Access" \
  --type file-based-access \
  --set enabled:true \
  --set compression-mechanism:gzip \
  --set encryption-settings-definition-
id:332C846EF0DCD1D5187C1592E4C74CAD33FC1E5FC20B726CD301CDD2B3FFBC2B \
  --set encrypt-log:true \
  --set log-file:logs/encrypted-access \
  --set "rotation-policy:24 Hours Time Limit Rotation Policy" \
  --set "rotation-policy:Size Limit Rotation Policy" \
  --set "retention-policy:File Count Retention Policy" \
  --set "retention-policy:Free Disk Space Retention Policy" \
  --set "retention-policy:Size Limit Retention Policy"
```

2. To decrypt and decompress the file:

```
$ bin/encrypt-file --decrypt \
  --decompress-input \
  --input-file logs/encrypted-access.20180216040332Z.gz.encrypted \
  --output-file decrypted-access
Initializing the server's encryption framework...DoneWriting decrypted
data to file '/ds/PingDirectory/decrypted-access' using akey generated
from encryption settings definition
'332c846ef0dcd1d5187c1592e4c74cad33fc1e5fc20b726cd301cdd2b3ffbc2b'Success
fully wrote 123,456,789 bytes of decrypted data
```

Backend Monitor Entries

Each PingData server exposes its monitoring information under the `cn=monitor` entry.

Administrators can use various means to monitor the servers through SNMP, the Administrative Console, JConsole, LDAP command-line tools, and the Stats Logger.

The Monitor Backend contains an entry per component or activity being monitored. The list of all monitor entries can be seen using the `ldapsearch` command as follows:

```
$ bin/ldapsearch --hostname server1.example.com \
  --port 1389 \
  --bindDN "uid=admin,dc=example,dc=com" \
  --bindPassword secret \
  --baseDN "cn=monitor" "(objectclass=*)" cn
```

The following table lists a subset of monitor entries.

Monitoring Components

Component	Description
Active Operations	Provides information about the operations currently being processed by the server including the number of operations, information on each operation, and the

Monitoring Components

Component	Description
	number of active persistent searches.
Backends	Provides general information about the state of a server backend, including the entry count. If the backend is a local database, there is a corresponding database environment monitor entry with information on cache usage and on-disk size.
Client Connections	Provides information about all client connections to the server including a name followed by an equal sign and a quoted value, such as <code>connID="15"</code> , <code>connectTime="20100308223038Z"</code> .
Connection Handlers	Provides information about the available connection handlers on the server including the LDAP and LDIF connection handlers.
Disk Space Usage	Provides information about the disk space available to various components of the server.
General	Provides general information about the state of the server, including product name, vendor name, and server version.
Index	Provides information on each index including the number of preloaded keys and counters for read, write, remove, open-cursor, and read-for-search actions. These counters provide insight into how useful an index is for a given workload.
HTTP/HTTPS Connection Handler Statistics	Provides statistics about the interaction that the associated HTTP connection handler has had with its clients, including the number of connections accepted, average requests per connection, average connection duration, total bytes returned, and average processing time by status code.
JVM Stack Trace	Provides a stack trace of all threads processing within the JVM.
LDAP Connection Handler Statistics	Provides statistics about the interaction that the associated LDAP connection handler has had with its clients, including the number of connections established and closed, bytes read and written, LDAP messages read and written, and operations initiated, completed, and abandoned.
Processing Time Histogram	Categorizes operation processing times into a number of user-defined buckets of information, including the total number of operations processed, overall average response time (ms), and number of processing times between 0ms and 1ms.
System Information	Provides general information about the system and the JVM on which the server is running, including system host name, operation system, JVM architecture, Java home, and Java version.
Version	Provides information about the server version, including build ID, and revision number.
Work Queue	<p>Provides information about the state of the server work queue, which holds requests until they can be processed by a worker thread, including the requests rejected, current work queue size, number of worker threads, and number of busy worker threads.</p> <p>The work queue configuration has a <code>monitor-queue-time</code> property set to <code>true</code> by default. This logs messages for new operations with a <code>qtime</code> attribute included in the log messages. Its value is expressed in milliseconds and represents the length of time that operations are held in the work queue.</p>

Disk Space Usage Monitor

The disk space usage monitor provides information about the amount of usable disk space available for server components. It also provides the ability to generate administrative alerts, as well as take action if the amount of usable space drops below the defined thresholds.

The disk space usage monitor evaluates the free space at locations registered through the `DiskSpaceConsumer` interface. Disk space monitoring excludes disk locations that do not have server components registered. However, other disk locations may still impact server performance, such as the operating system disk, if it becomes full. When relevant to the server, these locations include the server root, the location of the `config` directory, the location of every log file, all JE backend directories, the location of the changelog, the location of the replication environment database, and the location of any server extension that registers itself with the `DiskSpaceConsumer` interface.

All values must be specified as absolute values or as percentages. A mix of absolute values and percentages cannot be used. The following thresholds are available:

- **Low space warning** – This threshold defines either a percentage or an absolute amount of usable space. If the amount of usable space drops below this threshold, the server generates an administrative alert. It generates alerts at regular intervals, based on configuration settings, until the amount of usable space is increased, or as the amount of usable space is further reduced.
- **Low space error** – This threshold is also defined as either a percentage or an absolute size. Once the amount of usable space drops below this threshold, the server will generate an alert notification and will begin rejecting all operations requested by non-root users with "UNAVAILABLE" results. Once the server enters this mode, some action must be taken before the server will resume normal operations. This threshold must be less than or equal to the low space warning threshold. If they are equal, the server will begin rejecting requests from non-root users immediately upon detecting low usable disk space.
- **Out of space error** – This threshold can also be defined as a percentage or an absolute size. Once the amount of usable space drops below this threshold, the server will generate a final administrative alert and will shut itself down. This threshold must be less than or equal to the low space error threshold. If they are equal, the server will shut itself down rather than rejecting requests from non-root users.

Notifications and Alerts

Each PingData sever provides delivery mechanisms for account status notifications, administrative alerts, and alarms using SMTP, JMX, or SNMP. Alerts, alarms, and events reflect state changes within the server that may be of interest to a user or monitoring service. Account status notifications are only delivered to the account owner.

Alert handler implementations include:

Error Log Alert Handler – Sends administrative alerts to the configured server error logger (s).

Exec Alert Handler – Executes a specified command on the local system if an administrative alert matching the criteria for this alert handler is generated by the server. Information about the administrative alert is made available to the executed application as arguments provided by the command.

Groovy Scripted Alert Handler – Provides alert handler implementations defined in a dynamically-loaded Groovy script that implements the `ScriptedAlertHandler` class defined in the Server SDK.

JMX Alert Handler – Sends administrative alerts to clients using the Java Management Extensions (JMX) protocol. PingData uses JMX for monitoring entries and requires that the JMX connection handler be enabled.

SMTP Alert Handler – Sends administrative alerts to clients via email using the SMTP. The server requires that one or more SMTP servers be defined in the global configuration.

SNMP Alert Handler – Sends administrative alerts to clients using the Simple Network Monitoring Protocol (SNMP). The server must have an SNMP agent capable of communicating through SNMP.

SNMP Subagent Alert Handler – Sends SNMP traps to a master agent in response to administrative alerts generated within the server.

Third Party Alert Handler – Provides alert handler implementations created in third-party code using the Server SDK.

A complete listing of system alerts, alarms, and their severity is available in `<server-root>/docs/admin-alerts-list.csv`

Configure alert handlers

Alert handlers can be configured with the `dsconfig` tool. PingData servers support JMX, SMTP, and SNMP. Use the `--help` option for a list of configuration options. The following is a sample command to create and enable an SMTP Alert handler from the command line:

```
$ bin/dsconfig create-alert-handler \  
  --handler-name "SMTP Alert Handler" \  
  --type smtp \  
  --set enabled:true \  
  --set "sender-address:alerts@example.com" \  
  --set "recipient-address:administrators@example.com" \  
  --set "message-subject:Directory Admin Alert \%%alert-type\%%" \  
  --set "message-body:Administrative alert:\n\%%alert-message\%%"
```

The Alerts Backend

PingData servers generate administrative alerts under the `cn=alerts` branch. The backend makes it possible to obtain admin alert information over LDAP for use with remote monitoring. The backend's primary job is to process search operations for alerts. It does not support add, modify, or modify DN operations of entries.

The alerts persist on disk in the `config/alerts.ldif` file so that they can survive server restarts. By default, the alerts remain on disk for seven days before being removed. However, administrators can configure the number of days for alert retention using the `dsconfig` tool. The administrative alerts of Warning level or worse that have occurred in the last 48 hours are viewable from the output of the `status` command-line tool and in the Administrative Console.

View Information in the Alerts Backend

Use `ldapsearch` to view the administrative alerts:

```
$ bin/ldapsearch --port 1389 --bindDN "cn=Directory Manager" \  
  --bindPassword secret --baseDN cn=alerts "(objectclass=*)" \  
dn: cn=alerts \  
objectClass: top \  
objectClass: ds-alert-root \  
cn: alerts \  
  
dn: ds-alert-id=3d1857a2-e8cf-4e80-ac0e-ba933be59eca,cn=alerts \  
objectClass: top \  
objectClass: ds-admin-alert \  
ds-alert-id: 3d1857a2-e8cf-4e80-ac0e-ba933be59eca \  
ds-alert-type: server-started \  
ds-alert-severity: info
```

```
ds-alert-type-oid: 1.3.6.1.4.1.32473.2.11.33
ds-alert-time: 20110126041442.622Z
ds-alert-generator: com.unboundid.directory.server.core.metrics.engine
ds-alert-message: The server has started successfully
```

Modify the Alert Retention Time

Use `dsconfig` to change the maximum time information about generated alerts retained in the alerts backend. After this time, the information is purged from the server. The minimum retention time is 0 milliseconds, which immediately purges the alert information.

```
$ bin/dsconfig set-backend-prop --backend-name "alerts" \
  --set "alert-retention-time: 2 weeks"
```

View the property using `dsconfig`:

```
$ bin/dsconfig get-backend-prop --backend-name "alerts" \
  --property alert-retention-time
```

```
Property : Value(s)
```

```
-----:-----
alert-retention-time : 2 w
```

Configure Duplicate Alert Suppression

Use `dsconfig` to configure the maximum number of times an alert is generated within a particular time frame for the same condition. The `duplicate-alert-time-limit` property specifies the length of time that must pass before duplicate messages are sent over the administrative alert framework and the maximum number of messages should be sent.

```
$ bin/dsconfig set-global-configuration-prop \
  --set duplicate-alert-limit:2 \
  --set "duplicate-alert-time-limit:3 minutes"
```

System alarms, alerts, and gauges

An alarm represents a stateful condition of the server or a resource that may indicate a problem, such as low disk space or external server unavailability. A gauge defines a set of threshold values with a specified severity that, when crossed, cause the server to enter or exit an alarm state. Gauges are used for monitoring continuous values like CPU load or free disk space (Numeric Gauge), or an enumerated set of values such as 'server available' or 'server unavailable' (Indicator Gauge). Gauges generate alarms, when the gauge's severity changes due to changes in the monitored value. Like alerts, alarms have severity (NORMAL, WARNING,

Chapter 3: Managing the PingDataMetrics Server

MINOR, MAJOR, CRITICAL), name, and message. Alarms will always have a `Condition` property, and may have a `Specific Problem` or `Resource` property. If surfaced through SNMP, a `Probable Cause` property and `Alarm Type` property are also listed. Alarms can be configured to generate alerts when the alarm's severity changes.

There are two alert types supported by the server - standard and alarm-specific. The server constantly monitors for conditions that may need attention by administrators, such as low disk space. For this condition, the standard alert is `low-disk-space-warning`, and the alarm-specific alert is `alarm-warning`. The server can be configured to generate alarm-specific alerts instead of, or in addition to, standard alerts. By default, standard alerts are generated for conditions internally monitored by the server. However, gauges can only generate alarm-alerts.

The server installs a set of gauges that are specific to the product and that can be cloned or configured through the `dsconfig` tool. Existing gauges can be tailored to fit each environment by adjusting the update interval and threshold values. Configuration of system gauges determines the criteria by which alarms are triggered. The Stats Logger can be used to view historical information about the value and severity of all system gauges.

PingData servers are compliant with the International Telecommunication Union CCITT Recommendation X.733 (1992) standard for generating and clearing alarms. If configured, entering or exiting an alarm state can result in one or more alerts. An alarm state is exited when the condition no longer applies. An `alarm_cleared` alert type is generated by the system when an alarm's severity changes from a non-normal severity to any other severity. An `alarm_cleared` alert will correlate to a previous alarm when `Condition` and `Resource` property are the same. The Alarm Manager, which governs the actions performed when an alarm state is entered, is configurable through the `dsconfig` tool and Administrative Console.

Like the Alerts Backend, which stores information in `cn=alerts`, the Alarm Backend stores information within the `cn=alarms` backend. Unlike alerts, alarm thresholds have a state over time that can change in severity and be cleared when a monitored value returns to normal. Alarms can be viewed with the `status` tool. As with other alert types, alert handlers can be configured to manage the alerts generated by alarms. A complete listing of system alerts, alarms, and their severity is available in `<server-root>/docs/admin-alerts-list.csv`.

Test alerts and alarms

After alarms and alert handlers are configured, verify that the server takes the appropriate action when an alarm state changes by manually increasing the severity of a gauge. Alarms and alerts can be verified with the `status` tool.

1. Configure a gauge with `dsconfig` and set the `override-severity` property to `critical`. The following example uses the CPU Usage (Percent) gauge.

```
$ dsconfig set-gauge-prop \
  --gauge-name "CPU Usage (Percent)" \
  --set override-severity:critical
```

2. Run the `status` tool to verify that an alarm was generated with corresponding alerts. The `status` tool provides a summary of the server's current state with key metrics and a list of recent alerts and alarms. The sample output has been shortened to show just the alarms and alerts information.

```
$ bin/status
```

```

--- Administrative Alerts ---
Severity : Time           : Message
-----:-----:-----
Error    : 11/Aug/2016       : Alarm [CPU Usage (Percent). Gauge CPU Usage
(Percent)
          : 15:41:00 -0500       : for Host System has
          :                       : a current value of '18.583333333333332'.
          :                       : The severity is currently OVERRIDDEN in the
          :                       : Gauge's configuration to 'CRITICAL'.
          :                       : The actual severity is: The severity is
          :                       : currently 'NORMAL', having assumed this
severity
          :                       : Mon Aug 11 15:41:00 CDT 2016. If CPU use is
high,
          :                       : check the server's current workload and make
any
          :                       : needed adjustments. Reducing the load on the
system
          :                       : will lead to better response times.
          :                       : Resource='Host System']
          :                       : raised with critical severity
Shown are alerts of severity [Info,Warning,Error,Fatal] from the past 48
hours
Use the --maxAlerts and/or --alertSeverity options to filter this list

```

```

--- Alarms ---
Severity : Severity   : Condition : Resource   : Details
          : Start Time  :           :           :
-----:-----:-----:-----:-----
--
Critical : 11/Aug/2016: CPU Usage : Host System : Gauge CPU Usage

```

```
(Percent) for
      : 15:41:00 : (Percent) : : Host System
      : -0500    :           :           : has a current value of
      :         :           :           : '18.785714285714285'.
      :         :           :           : The severity is
currently
      :         :           :           : 'CRITICAL', having
assumed
      :         :           :           : this severity Mon Aug 11
      :         :           :           : 15:49:00 CDT 2016. If
CPU use
      :         :           :           : is high, check the
server's
      :         :           :           : current workload and
make any
      :         :           :           : needed adjustments.
Reducing
      :         :           :           : the load on the system
will
      :         :           :           : lead to better response
times
Shown are alarms of severity [Warning,Minor,Major,Critical
Use the --alarmSeverity option to filter this list
```

Back Up the PingDataMetrics Server Database

The PingDataMetrics Server stores all historical metric samples in the PostgreSQL DBMS, along with several other data tables that are used for bookkeeping and normalization of the sample data. Even a small PingDataMetrics Server installation, which monitors three to four servers, will use sample tables that occupy 95% of the total DBMS space. While a functional backup must capture a consistent view of several tables, the size of the sample tables dictates the desired approach to a regular backup strategy.

The historical samples enable:

- Diagnosing past performance problems.
- Capacity planning and historical reporting.
- Access to data needed for a revenue stream, such as data used for billing and charge back.

Defining data that is important to the infrastructure will help determine the right backup strategy. In the case of billing, the data needed is typically small compared to the total population of the DBMS. This may be all the data needed, and the planning and resources required to backup the DBMS will be minimal.

If it's not possible to determine what data will be important in the future, backing up all DBMS data is the safest approach.

Historical Data Storage

The PingDataMetrics Server DBMS stores all historical sample data. It can store time-aggregated data for up to twenty years. The data in the DBMS is continually changing as long as the PingDataMetrics Server is running.

The system that feeds data to the PingDataMetrics Server is designed to allow the PingDataMetrics Server to be offline for hours at a time without dropping any data. The collection points hold the data for hours, giving the PingDataMetrics Server time for maintenance tasks. The collection points do have a limit on how long they hold data, so the PingDataMetrics Server cannot be offline for an indeterminate time.

If the PingDataMetrics Server is offline so long that the collection points start to delete data that has not yet been captured, then there will be gaps in the data. Aggregation still works, even with these gaps. If the data gap is four hours, four time samples will be missing in the one hour aggregation level, and no data will be missing in the one day aggregation level. However, the one day aggregation level will use only 20 hours of data rather than 24. By default, the PingDataMetrics Server can be offline for about eight hours before any data is lost.

The PingDataMetrics Server responds to queries that result in data with time gaps. The resulting data differentiates between data with zero value and missing data.

Planning the Backup

Choose a time window during which the PingDataMetrics Server can be offline and ensure that there is enough disk space to hold the new image. The exact size of a DBMS table and its corresponding backup depends on the number of monitored servers, the number of tracked applications, the collected metrics, and the retention duration for each of the aggregation levels. The following table provides values from installations used during testing.

Data From Sample Deployments		
Data	25 Monitored Servers	50 Monitored Servers
Number of tracked applications	20	20
1 second data resolution	8 hours	8 hours

Data From Sample Deployments

Data	25 Monitored Servers	50 Monitored Servers
1 minute data retention	14 days	14 days
1 hour data retention	52 weeks	52 week
1 day data retention	20 years	20 years
1 second table size	22 G	42 G
1 minute table size	8 G	18 G
1 hour table size	4 G (estimated)	9 G (estimated)
1 day data retention	4 G (estimated)	7 G (estimated)
time to backup	15 minutes (estimated)	30 minutes (estimated)
time for import catchup	10 minutes	42 minutes
size of compressed backup image	3 G (estimated)	5.5 G (estimated)
time to restore	1 hour (estimated)	2 h (estimated)

If no backups are performed and the DBMS is completely lost, reinitialize the DBMS, restart the PingDataMetrics Server, and start collecting data again. All collected metric and event data are lost, but the configuration required to start collecting data again is retained.

Start the DBMS Backup

Shut down the PingDataMetrics Server before a backup or restore.

To backup the entire DBMS use the following command:

```
$ tar -cf backup.tar <path-to-postgres-data-directory>
```

Restore a DBMS Backup

To restore the full backup to a new database, use the following command:

```
$ tar -xvf backup.tar
```

Run the command from the base directory of the PostgreSQL data directory.

For more information, documentation is available on the PostgreSQL website.

Management Tools

The PingDataMetrics Server provides several command-line tools to administer the server. The command-line tools are available in the `bin` directory for UNIX or Linux systems and `bat` directory for Microsoft Windows systems.

Each command-line utility provides a description of the subcommands, arguments, and usage examples needed to run the tool. View detailed argument options and examples by typing `--help` with the command.

```
$ bin/dsconfig --help
```

To list the subcommands for each command:

```
$ bin/dsconfig --help-subcommands
```

To list more detailed subcommand information:

```
$ bin/dsconfig list-log-publishers --help
```

Available Command-Line Utilities

The following command-line utilities are available, which can be run in interactive, non-interactive, or script mode.

Command Line Tools

Command-Line Tool	Description
backup	Run full or incremental backups on one or more PingDataMetrics Server backends. This utility also supports the use of a properties file to pass predefined command-line arguments. See Managing the tools.properties File for more information.
base64	Encode raw data using the base64 algorithm or decode base64-encoded data back to its raw representation.
collect-support-data	Collect and package system information useful in troubleshooting problems. The information is packaged as a ZIP archive that can be sent to a technical support representative.
create-rc-script	Create an Run Control (RC) script that may be used to start, stop, and restart the server on UNIX-based systems.
config-diff	Generate a summary of the configuration changes in a local or remote server instance. The tool can be used to compare configuration settings when troubleshooting issues, or when verifying configuration settings on new servers.
dsconfig	View and edit the server configuration.
dsframework	Manage administrative server groups or the global administrative user accounts that are used to configure servers within server groups.
dsjavaproperties	Configure the JVM arguments used to run the server and associated tools. Before launching the command, edit the properties file located in <code>config/java.properties</code> to specify the desired JVM options and <code>JAVA_HOME</code> environment variable.
ldapmodify	Perform LDAP modify, add, delete, and modify DN operations.
ldappasswordmodify	Perform LDAP password modify operations.
ldapsearch	Perform LDAP search operations.
ldif-diff	Compare the contents of two LDIF files, the output being an LDIF file needed to bring

Command Line Tools

Command-Line Tool	Description
	the source file in sync with the target.
ldifmodify	Apply a set of modify, add, and delete operations against data in an LDIF file.
manage-extension	Install or update extension bundles. An extension bundle is a package of extension (s) that utilize the Server SDK to extend the functionality of the server. Extension bundles are installed from a zip archive or file system directory. The server is restarted to activate the extension(s).
metric-engine-schema	Show current and required DBMS schema version information.
monitored-servers	Configure the set of servers to be monitored by this PingDataMetrics Server and prepare external servers for monitoring.
query-metric	Explore collected monitoring data by forming queries for data.
queryrate	Execute metric queries.
restore	Restore a backup of the server backend.
revert-update	Returns a server to the version before the last update was performed.
review-license	Review or accept the product license.
server-state	View information about the current state of the server process.
setup	Perform the initial setup for the server instance.
start-server	Start the server.
status	Display basic server information.
stop-server	Stop or restart the server.
sum-file-sizes	Calculate the sum of the sizes for a set of files.
uninstall	Uninstall the server.
update	Update the server to a newer version by downloading and unzipping the new server install package on the same host as the server to update. Use the update tool from the new server package to update the older version of the server. During the update process, the server is stopped if running, then the update is performed. A check is performed to determine if the newly updated server starts without major errors. If it cannot start cleanly, the update is backed out and the server is returned to its prior state. See the <code>revert-update</code> tool for information on reverting an update.

The tools.property File

The `tools.properties` file simplifies command-line invocations by reading in a set of arguments for each tool from a text file. Each property consists of a name/value pair for a tool's arguments.

Two types of properties files are supported:

- Default properties files that can be applied to all command-line utilities.
- Tool-specific properties file that can be specified using the `--propertiesFilePath` option.

All of the server's command-line utilities can be over-written using the `config/tools.properties` file.

Create a properties file with a text editor or using the standard Java properties file format (`name=value`). For example, create a simple properties file that defines a set of LDAP connection parameters as follows:

```
hostname=server1.example.com
port=1389
bindDN=cn=Directory\ Manager
bindPassword=secret
```

Specify the location of the file using the `--propertiesFilePath` option. For example, specify the path to the properties file with `ldapsearch` as follows:

```
$ bin/ldapsearch --propertiesFilePath bin/mytools.properties "(objectclass=*)"
```

Properties files do not allow quotation marks around values. Any spaces or special characters should be escaped.

Tool-Specific Properties

The server also supports properties for specific tool options using the format:

`tool.option=value`. Tool-specific options have precedence over general options. For example, the following properties file uses `ldapsearch.port=2389` for `ldapsearch` requests by the client.

All other tools that use the properties file use `port=1389`.

```
hostname=server1.example.com
port=1389
ldapsearch.port=2389
bindDN=cn=Directory\ Manager
```

Another example using the `dsconfig` configuration tool is as follows:

```
hostname=server1.example.com
port=1389
bindDN=cn=Directory\ Manager
dsconfig.bindPasswordFile=/ds/config/password
```


Specify Default Properties Files

The server provides a default properties file, `tools.properties`, that applies to all command-line utilities used in client requests. The file is located in the `<server-root>/config` directory.

To use a file with a different filename in this default location, specify the path using the `--propertiesFilePath` option.

Evaluation Order

The following evaluation ordering is used to determine options for a given command-line utility:

- All options used with a utility on the command line take precedence over any options in any properties file.
- If the `--propertiesFilePath` option is used with no other options, the server takes its options from the specified properties file.
- If no options are used on the command line including the `--propertiesFilePath` option (and `--noPropertiesFile`), the server searches for the `tools.properties` file at `<server-root>`.
- If no default properties file is found and a required option is missing, the tool generates an error.
- Tool-specific properties (for example, `ldapsearch.port=3389`) have precedence over general properties (for example, `port=1389`).

HTTP Connection Handlers

HTTP Connection Handlers are responsible for managing the communication with HTTP clients and invoking servlets to process requests from those clients. They can also be used to host web applications on the server. Each HTTP connection handler must be configured with one or more HTTP servlet extensions and zero or more HTTP operation log publishers.

If the HTTP Connection Handler cannot be started (for example, if its associated HTTP Servlet Extension fails to initialize), then this will not prevent the entire Directory Proxy Server from starting. The server's `start-server` tool will output any errors to the error log. This allows the server to continue serving LDAP requests even with a bad servlet extension.

The configuration properties available for use with an HTTP connection handler include:

- **listen-address**. Specifies the address on which the connection handler will listen for requests from clients. If not specified, then requests will be accepted on all addresses bound to the system.
- **listen-port**. Specifies the port on which the connection handler will listen for requests from clients. Required.
- **use-ssl**. Indicates whether the connection handler will use SSL/TLS to secure communications with clients (whether it uses HTTPS rather than HTTP). If SSL is enabled, then `key-manager-provider` and `trust-manager-provider` values must also be specified.
- **http-servlet-extension**. Specifies the set of servlet extensions that will be enabled for use with the connection handler. You can have multiple HTTP connection handlers (listening on different address/port combinations) with identical or different sets of servlet extensions. At least one servlet extension must be configured.
- **http-operation-log-publisher**. Specifies the set of HTTP operation log publishers that should be used with the connection handler. By default, no HTTP operation log publishers will be used.
- **key-manager-provider**. Specifies the key manager provider that will be used to obtain the certificate presented to clients if SSL is enabled.
- **trust-manager-provider**. Specifies the trust manager provider that will be used to determine whether to accept any client certificates presented to the server.
- **num-request-handlers**. Specifies the number of threads that should be used to process requests from HTTP clients. These threads are separate from the worker threads used to process other kinds of requests. The default value of zero means the number of threads will be automatically selected based on the number of CPUs available to the JVM.
- **web-application-extension**. Specifies the Web applications to be hosted by the server.

Configure an HTTP Connection Handler

An HTTP connection handler has two dependent configuration objects: one or more HTTP servlet extensions and optionally, an HTTP log publisher. The HTTP servlet extension and log publisher must be configured prior to configuring the HTTP connection handler. The log publisher is optional but in most cases, you want to configure one or more logs to troubleshoot any issues with your HTTP connection.

1. The first step is to configure your HTTP servlet extensions. The following example uses the `ExampleHTTPServletExtension` in the Server SDK.

```
$ bin/dsconfig create-http-servlet-extension \  
  --extension-name "Hello World Servlet" \  
  \
```

Chapter 3: Managing the PingDataMetrics Server

```
--type third-party \  
--set  
"extensionclass:com.unboundid.directory.sdk.examples.ExampleHTTPServletEx  
tension" \  
--set "extension-argument:path=/" \  
--set "extension-argument:name=example-servlet"
```

2. Next, configure one or more HTTP log publishers. The following example configures two log publishers: one for common access; the other, detailed access. Both log publishers use the default configuration settings for log rotation and retention.

```
$ bin/dsconfig create-log-publisher \  
--publisher-name "HTTP Common Access Logger" \  
--type common-log-file-http-operation \  
--set enabled:true \  
--set log-file:logs/http-common-access \  
--set "rotation-policy:24 Hours Time Limit Rotation Policy" \  
--set "rotation-policy:Size Limit Rotation Policy" \  
--set "retention-policy:File Count Retention Policy" \  
--set "retention-policy:Free Disk Space Retention Policy"
```

```
$ bin/dsconfig create-log-publisher \  
--publisher-name "HTTP Detailed Access Logger" \  
--type detailed-http-operation \  
--set enabled:true \  
--set log-file:logs/http-detailed-access \  
--set "rotation-policy:24 Hours Time Limit Rotation Policy" \  
--set "rotation-policy:Size Limit Rotation Policy" \  
--set "retention-policy:File Count Retention Policy" \  
--set "retention-policy:Free Disk Space Retention Policy"
```

3. Configure the HTTP connection handler by specifying the HTTP servlet extension and log publishers. Note that some configuration properties can be later updated on the fly while others, like listen-port, require that the HTTP Connection Handler be disabled, then re-enabled for the change to take effect.

```
$ bin/dsconfig create-connection-handler \  
--handler-name "Hello World HTTP Connection Handler" \  
--type http \  
--set enabled:true \  
--set listen-port:8443 \  
--set use-ssl:true \  
--set "http-servlet-extension:Hello World Servlet" \  
--set "http-operation-log-publisher:HTTP Common Access Logger" \  
--set "http-operation-log-publisher:HTTP Detailed Access Logger" \  
--set "key-manager-provider:JKS" \  
--set "trust-manager-provider:JKS"
```

4. By default, the HTTP Connection Handler has an advanced monitor entry property, `keep-stats`, that is set to TRUE by default. You can monitor the connection handler using the `ldapsearch` tool.

```
$ bin/ldapsearch --baseDN "cn=monitor" \
  "(objectClass=ds-http-connection-handler-statistics-monitor-entry) "
```

HTTP Correlation IDs

A typical request to a software system is handled by multiple subsystems, many of which may be distinct servers residing on distinct hosts and locations. Tracing the request flow on distributed systems can be challenging, as log messages are scattered across various systems and intermingled with messages for other requests. To make this easier, a correlation ID can be assigned to a request, which is then added to every associated operation as the request flows through the larger system. The correlation ID allows related log messages to be easily located and grouped. The server supports correlation IDs for all HTTP requests received through its HTTP(S) Connection Handler.

When an HTTP request is received, it is automatically assigned a correlation ID. This ID can be used to correlate HTTP responses with messages recorded to the HTTP Detailed Operation log and the trace log. For specific web APIs, the correlation ID may also be passed to the LDAP subsystem. For the SCIM 1, Delegated Admin, Consent, and Directory REST APIs, the correlation ID will also appear with associated requests in LDAP logs in the `correlationID` key. The correlation ID is also used as the default client request ID value in Intermediate Client Request Controls used by the SCIM 2, Consent, and Directory REST APIs. Values related to the Intermediate Client Request Control appear in the LDAP logs in the `via` key, and are forwarded to downstream LDAP servers when received by the PingDirectoryProxy server. The correlation ID header is also added to requests forwarded by the PingDataGovernance gateway

For Server SDK extensions that have access to the current `HttpServletRequest`, the current correlation ID can be retrieved as a string through the `HttpServletRequest`'s `com.pingidentity.pingdata.correlation_id` attribute. For example:

```
(String) request.getAttribute("com.pingidentity.pingdata.correlation_id");
```

Configure HTTP Correlation ID Support

Correlation ID support is enabled by default for each HTTP Connection Handler.

- To enable correlation ID support for the HTTPS Connection Handler:

```
$ bin/dsconfig set-connection-handler-prop \
  --handler-name "HTTPS Connection Handler" \
  --set use-correlation-id-header:true
```

- To disable correlation ID support for the HTTPS Connection Handler:

```
$ bin/dsconfig set-connection-handler-prop \  
  --handler-name "HTTPS Connection Handler" \  
  --set use-correlation-id-header:false
```

Configuring the correlation ID response header

- The server will generate a correlation ID for every HTTP request and send it in the response through the `Correlation-Id` response header. This response header name can be customized. The following example changes the `correlation-id-response-header` property to "X-Request-Id."

```
$ bin/dsconfig set-connection-handler-prop \  
  --handler-name "HTTPS Connection Handler" \  
  --set correlation-id-response-header:X-Request-Id
```

Accepting an incoming correlation ID from the request

- By default, the server generates a new, unique correlation ID for each HTTP request, and ignores any correlation ID that may be set on the request. This can be changed by designating the names of one or more HTTP request headers that contain an existing correlation ID value. This enables the server to integrate with a larger system consisting of every servers using correlation IDs.

```
$ bin/dsconfig set-connection-handler-prop \  
  --handler-name "HTTPS Connection Handler" \  
  --set correlation-id-request-header:X-Request-Id \  
  --set correlation-id-request-header:X-Correlation-Id \  
  --set correlation-id-request-header:Correlation-Id \  
  --set correlation-id-request-header:X-Amzn-Trace-Id
```

HTTP Correlation ID Example Use

In this example, a request to the Directory REST API is made and the correlation ID enables finding HTTP-specific log messages with LDAP-specific log messages. The response to the API call includes a `Correlation-Id` header with the value `a54aee33-c6c6-4467-be25-efd1db7a8b76`.

```
GET /directory/v1/me?includeAttributes=mail HTTP/1.1  
Accept: */*  
Accept-Encoding: gzip, deflate  
Authorization: Bearer ...  
Connection: keep-alive  
Host: localhost:1443  
User-Agent: HTTPie/0.9.9  
  
HTTP/1.1 200 OK  
Content-Length: 266  
Content-Type: application/hal+json  
Correlation-Id: ee919049-6710-4594-9c66-28b4ada4b127  
Date: Fri, 02 Nov 2018 15:16:50 GMT
```

```
Request-Id: 369
{
  "_dn": "uid=user.86,ou=People,dc=example,dc=com",
  "_links": {
    "schemas": [
      {
        "href": "https://localhost:1443/directory/v1/schemas/
inetOrgPerson"
      }
    ],
    "self": {
      "href": "https://localhost:1443/directory/v1/
uid=user.86,ou=People,dc=example,dc=com"
    }
  },
  "mail": [
    "user.86@example.com"
  ]
}
```

This correlation ID can be used to search the HTTP trace log for matching log records, as follows:

```
$ grep 'correlationID="ee919049-6710-4594-9c66-28b4ada4b127"'
PingDirectory/logs/debug-trace
[02/Nov/2018:10:16:50.294 -0500] HTTP REQUEST requestID=369
correlationID="ee919049-6710-4594-9c66-28b4ada4b127" product="Ping Identity
Directory Server" instanceName="ds1" startupID="W9ikqA==" threadID=52358 from=
[0:0:0:0:0:0:0:1]:58918 method=GET
url="https://0:0:0:0:0:0:0:1:1443/directory/v1/me?includeAttributes=mail"
[02/Nov/2018:10:16:50.526 -0500] DEBUG ACCESS-TOKEN-VALIDATOR-PROCESSING
requestID=369 correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
msg="Identity Mapper with DN 'cn=User ID Identity Mapper,cn=Identity
Mappers,cn=config' mapped ID 'user.86' to entry DN
'uid=user.86,ou=people,dc=example,dc=com'"
[02/Nov/2018:10:16:50.526 -0500] DEBUG ACCESS-TOKEN-VALIDATOR-PROCESSING
requestID=369 correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
accessTokenId="201811020831" msg="Token Validator 'Mock Access Token
Validator' validated access token with active = 'true', sub = 'user.86', owner
= 'uid=user.86,ou=people,dc=example,dc=com', clientId = 'client1', scopes =
'ds', expiration = 'none', not-used-before = 'none', current time = 'Nov 2,
2018 10:16:50 AM CDT' "
[02/Nov/2018:10:16:50.531 -0500] HTTP RESPONSE requestID=369
correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
accessTokenId="201811020831" product="Ping Identity Directory Server"
instanceName="ds1" startupID="W9ikqA==" threadID=52358 statusCode=200
etime=236.932 responseContentLength=266
[02/Nov/2018:10:16:50.531 -0500] DEBUG HTTP-FULL-REQUEST-AND-RESPONSE
requestID=369 correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
accessTokenId="201811020831" product="Ping Identity Directory Server"
instanceName="ds1" startupID="W9ikqA==" threadID=52358 from=
[0:0:0:0:0:0:0:1]:58918 method=GET
```

Chapter 3: Managing the PingDataMetrics Server

```
url="https://0:0:0:0:0:0:0:1:1443/directory/v1/me?includeAttributes=mail"
statusCode=200 etime=236.932 responseContentLength=266 msg="
```

The LDAP log messages associated with this request can also be located:

```
$ grep 'correlationID="ee919049-6710-4594-9c66-28b4ada4b127"'
PingDirectory/logs/access
[02/Nov/2018:10:16:50.529 -0500] SEARCH RESULT instanceName="ds1"
threadID=52358 conn=-371045 op=1657393 msgID=1657394 origin="Directory REST
API" httpRequestID="369" correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
authDN="uid=user.86,ou=people,dc=example,dc=com" requesterIP="internal"
requesterDN="uid=user.86,ou=People,dc=example,dc=com"
requestControls="1.3.6.1.4.1.30221.2.5.2" via="app='PingDirectoryds1'
clientIP='0:0:0:0:0:0:0:1' sessionID='201811020831' requestID='ee919049-6710-
4594-9c66-28b4ada4b127'" base="uid=user.86,ou=people,dc=example,dc=com"
scope=0 filter="(&)" attrs="mail,objectClass" resultCode=0
resultCodeName="Success" etime=0.684 entriesReturned=1
[02/Nov/2018:10:16:50.530 -0500] EXTENDED RESULT instanceName="ds1"
threadID=52358 conn=-371046 op=1657394 msgID=1657395 origin="Directory REST
API" httpRequestID="369" correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
authDN="cn=Internal Client,cn=Internal,cn=Root DNs,cn=config"
requesterIP="internal" requesterDN="cn=Internal Client,cn=Internal,cn=Root
DNs,cn=config" requestControls="1.3.6.1.4.1.30221.2.5.2"
via="app='PingDirectory-ds1' clientIP='0:0:0:0:0:0:0:1'
sessionID='201811020831' requestID='ee919049-6710-4594-9c66-28b4ada4b127'"
requestOID="1.3.6.1.4.1.30221.1.6.1" requestType="Password Policy State"
resultCode=0 resultCodeName="Success" etime=0.542 usedPrivileges="bypass-
acl,password-reset" responseOID="1.3.6.1.4.1.30221.1.6.1"
responseType="Password Policy State"
dn="uid=user.86,ou=People,dc=example,dc=com"
[02/Nov/2018:10:16:50.530 -0500] SEARCH RESULT instanceName="ds1"
threadID=52358 conn=-371048 op=1657397 msgID=1657398 origin="Directory REST
API" httpRequestID="369" correlationID="ee919049-6710-4594-9c66-28b4ada4b127"
authDN="cn=Internal Client,cn=Internal,cn=Root DNs,cn=config"
requesterIP="internal" requesterDN="cn=Internal Client,cn=Internal,cn=Root
DNs,cn=config" requestControls="1.3.6.1.4.1.30221.2.5.2"
via="app='PingDirectoryds1' clientIP='0:0:0:0:0:0:0:1'
sessionID='201811020831' requestID='ee919049-6710-4594-9c66-28b4ada4b127'"
base="cn=Default Password Policy,cn=Password Policies,cn=config" scope=0
filter="(&)" attrs="dscfg- password-attribute" resultCode=0
resultCodeName="Success" etime=0.065 preAuthZUsedPrivileges="bypass-
acl,config-read" entriesReturned=1
```

Topology configuration

Topology configuration enables grouping servers and mirroring configuration changes automatically. It uses a master/slave architecture for mirroring shared data across the topology. All writes and updates are forwarded to the master, which forwards them to all other servers. Reads can be served by any server in the group.

Note

To remove a server from the topology, it must be uninstalled with the `uninstall` tool.

Topology master requirements and selection

A topology master server receives any configuration change from other servers in the topology, verifies the change, then makes the change available to all connected servers when they poll the master. The master always sends a digest of its subtree contents on each update. If the node has a different digest than the master, it knows it's not synchronized. The servers will pull the entire subtree from the master if they detect that they are not synchronized. A server may detect it is not synchronized with the master under the following conditions:

- At the end of its periodic polling interval, if a server's subtree digest differs from that of its master, then it knows it's not synchronized.
- If one or more servers have been added to or removed from the topology, the servers will not be synchronized.

The master of the topology is selected by prioritizing servers by minimum supported product version, most available, newest server version, earliest start time, and startup UUID (a smaller UUID is preferred).

After determining a master, the topology data is reviewed from all available servers (every five seconds by default) to determine if any new information makes a server better suited to being the master. If a new server can be the master, it will communicate that to the other servers, if no other server has advertised that it should be the master. This ensures that all servers accept the same master at approximately the same time (within a few milliseconds of each other). If there is no better master, the initial master maintains the role.

After the best master has been selected for the given interval, the following conditions are confirmed:

- A majority of servers is reachable from that master. (The master server itself is considered while determining this majority.)
- There is only a single master in the entire topology.

If either of these conditions is not met, the topology is without a master and the peer polling frequency is reduced to 100 milliseconds to find a new master as quickly as possible. If there is no master in the topology for more than one minute, a `mirrored-subtree-manager-no-master-found` alarm is raised. If one of the servers in the topology is forced as master with

the `force-as-master-for-mirrored-data` option in the Global Configuration configuration object, a `mirrored-subtree-manager-forced-as-master-warning` warning alarm is raised. If multiple servers have been forced as masters, then a `mirrored-subtree-manager-forced-as-master-error` critical alarm will be raised.

Topology components

When a server is installed, it can be added to an existing topology, which will clone the server's . Topology settings are designed to operate without additional configuration. If required, some settings can be adjusted to fit the needs of the environment.

Server configuration settings

Configuration settings for the topology are configured in the Global Configuration and in the Config File Handler Backend. Though they are topology settings, they are unique to each server and are not mirrored. Settings must be kept the same on all servers.

The Global Configuration object contains a single topology setting, `force-as-master-for-mirrored-data`. This should be set to `true` on only one of the servers in the topology, and is used only if a situation occurs where the topology cannot determine a master because a majority of servers is not available. A server with this setting enabled will be assigned the role of master, if no suitable master can be determined. See [Topology master requirements and selection](#) for details about how a master is selected for a topology.

The Config File Handler Backend defines three topology (`mirrored-subtree`) settings:

- `mirrored-subtree-peer-polling-interval` – Specifies the frequency at which the server polls its topology peers to determine if there are any changes that may warrant a new master selection. A lower value will ensure a faster failover, but it will also cause more traffic among the peers. The default value is five seconds. If no suitable master is found, the polling frequency is adjusted to 100 milliseconds until a new master is selected.
- `mirrored-subtree-entry-update-timeout` – Specifies the maximum length of time to wait for an update operation (add, delete, modify or modify-dn) on an entry to be applied by the master on all of the servers in the topology. The default is 10 seconds. In reality, updates can take up to twice as much time as this timeout value if master selection is in progress at the time the update operation was received.
- `mirrored-subtree-search-timeout` – Specifies the maximum length of time in milliseconds to wait for search operations to complete. The default is 10 seconds.

Topology settings

Topology meta-data is stored under the `cn=topology,cn=config` subtree and cluster data is stored under the `cn=cluster,cn=config` subtree. The only setting that can be changed is the cluster name.

Monitor data for the topology

Each server has a monitor that exposes that server's view of the topology in its monitor backend, so that peer servers can periodically read this information to determine if there are changes in the topology. Topology data includes the following:

- The server ID of the current master, if the master is not known.
- The instance name of the current master, or if a master is not set, a description stating why a master is not set.
- A flag indicating if this server thinks that it should be the master.
- A flag indicating if this server is the current master.
- A flag indicating if this server was forced as master.
- The total number of configured peers in the topology group.
- The peers connected to this server.
- The current availability of this server
- A flag indicating whether or not this server is not synchronized with its master, or another node in the topology if the master is unknown.
- The amount of time in milliseconds where multiple masters were detected by this server.
- The amount of time in milliseconds where no suitable server is found to act as master.
- A SHA-256 digest encoded as a base-64 string for the current subtree contents.

The following metrics are included if this server has processed any operations as master:

- The number of operations processed by this server as master.
- The number of operations processed by this server as master that were successful.
- The number of operations processed by this server as master that failed to validate.
- The number of operations processed by this server as master that failed to apply.
- The average amount of time taken (in milliseconds) by this server to process operations as the master.
- The maximum amount of time taken (in milliseconds) by this server to process an operation as the master.

Updating the server instance listener certificate

To change the SSL certificate for the server, update the keystore and truststore files with the new certificate. The certificate file must have the new certificate in PEM-encoded format, such as:

```
-----BEGIN CERTIFICATE-----
MIIDKTCCAhGgAwIBAgIEacgGrDANBgkqhkiG9w0BAQsFADBFMR4wHAYDVQQKEExVVbmJvdW5kSUQgQ2
VydG1maWNhdGUxIzAhBgNVBAMTGnZtLW1lZG11bS03My51bmJvdW5kaWQubGF1bG4XDTE1MTAxMjE1
MzU0OFoXDTM1MTAwNzE1MzU0FowRTEeMBWGA1UEChMvVW5ib3VuZElEIEEN1cnRpZmljYXR1MSMwIQ
YDVQQDExp2bS1tZWZlZD0tNzU0FowRTEeMBWGA1UEChMvVW5ib3VuZG1kLm1hYjYjCCAS1wDQYJKoZI
hvcNAQEBBQADggEPADCCAQoCggEBAKN4tAN3o9Yw6Cr9hivvVDxJqF6+aEi9Ir3WGFYLSrggRNXsiaOfWkSMWdIC5vyF5OJ9Dl
IgvHL4OuuqP/YNEGzKDKgr6MwtUeVSK14+dCixygJGC0nY7k+f0WSCjtIHzrmc4WWdrZXmgb+qv9Lup
S30JG0FXtcbGkYpjaKXIEqMg4ekz3B5cAve0SQUFYXEdN4rWOn96nVFkb2CstbiPzAgne2tu7paJ6S
GFOU0UF7v018XY1m2WHBIOd0WC8nOVLtG9zFUavaOxtlt1TlhClkI4HRMNg8n2EtSTdQRizKuw9DdT
XJBb6Kfvnp/nI73VHRyt47wUVueehEDfLDP8pMCAwEAAAMhMB8wHQYDVR0OBYYEFMrwjWxl2K+yd9
+Y65oKn0g5jITgMA0GCSqGSIb3DQEBCwUAA4IBAQBpsBYodblUGew+HewqtO2i8Wt+vAbt31zM5/kR
vo6/+iPEASTvZdCzIBcgletxKKGKeCQ0GPeHr42+erakiwmGD1UTYrU3LU5pTGTDLuR2I1lTT5xlEhC
WJGWipW4q3P13cX/9m2ffY/JLYdFTJaoJvnXrh7Sg719skkhjWZQgOHX1kPLx5TxFGhAovE1D4qLVR
WGohdpWDrIgFh0DVfoyan1Ws9ICCXdRayajFI4Lc6K1m6SA5+25Y9nno8BhVPf4q5OW6+UDc8MsLbB
sxpwwR6RJ5cv3ypfOriTehJsG+9ZDo7YeqVsTVGwAlW3PiSd9bYP/8yu9Cy+0MfcWcSeAE
-----END CERTIFICATE-----
```

If clients that already have a secure connection established with this server need to be maintained, information about both certificates can reside in the same file (each with their own begin and end headers and footers). If the listener certificate needs to be updated, it may be temporarily necessary for this property to have information about the old and new certificates. This can be done by including information about both certificates in the same file, each with their own begin and end headers and footers. Blank lines, and lines that start with the # character will be ignored.

After the keystore and truststore files are updated, run the following `dsconfig` command to update the server's certificate in the topology registry:

```
$ bin/dsconfig set-server-instance-listener-prop \
  --instance-name <server-instance-name> \
  --listener-name ldap-listener-mirrored-config \
  --set listener-certificate<path-to-new-certificate-file>
```

The `listener-certificate` in the topology registry is like a trust store. The public certificates that it has are automatically trusted by the local server. When the local server attempts a secure LDAP connection to a peer, and the peer presents it with its certificate, the local server

will check the `listener-certificate` property for that server in the topology registry. If the property contains the peer server's certificate, the local server will trust the peer.

Remove the self-signed certificate

The server is installed with a self-signed certificate and key (`ads-certificate`), which are used for internal purposes such as replication authentication, inter-server authentication in the topology registry, reversible password encryption, and encrypted backup/LDIF export. The `ads-certificate` lives in the keystore file called `ads-truststore` under the server's `/config` directory. If your deployment requires removing the self-signed certificate, it can be replaced.

The certificate is stored in the topology registry, which enables replacing it on one server and having it mirrored to all other servers in the topology. Any change is automatically mirrored on other servers in the topology. It is stored in human-readable PEM-encoded format and can be updated with `dsconfig`. The following general steps are required to replace the self-signed certificate:

1. Prepare a new keystore with the replacement key-pair.
2. Update the server configuration to use the new certificate by adding it to the server's list of certificates in the topology registry so that it is trusted by other servers.
3. Update the server's `ads-truststore` file to use the new key-pair.
4. Retire the old certificate by removing it from the topology registry.

Note

Replacing the entire key-pair instead of just the certificate associated with the original private key can make existing backups and LDIF exports invalid. This should be performed immediately after setup or before the key-pair is used. After the first time, only the certificate associated with the private key should have to be changed, for example, to extend its validity period or replace it with a certificate signed by a different CA.

Prepare a new keystore with the replacement key-pair

The self-signed certificate can be replaced with an existing key-pair, or the certificate associated with the original key-pair can be used.

Use an existing key-pair

If a private key and certificate(s) in PEM-encoded format already exist, both the original private key and self-signed certificate can be replaced in `ads-truststore` with the `manage-certificates` tool. The following command imports existing certificates into a new keystore file, `ads-truststore.new`:

Chapter 3: Managing the PingDataMetrics Server

```
$ bin/manage-certificates import-certificate \  
--keystore ads-truststore.new \  
--keystore-type JKS \  
--keystore-password-file ads-truststore.pin \  
--alias ads-certificate \  
--private-key-file existing.key \  
--certificate-file existing.crt \  
--certificate-file intermediate.crt \  
--certificate-file root-ca.crt
```

The certificates listed using the `--certificate-file` options must be ordered so that each subsequent certificate is the issuer for the previous one. So the server certificate comes first, the intermediate certificates next (if any), and the root CA certificate last.

Use the certificate associated with the original key-pair

The certificate associated with the original server-generated private key can be replaced with the following commands:

1. Create a CSR for the `ads-certificate`:

```
$ bin/manage-certificates generate-certificate-signing-request \  
--keystore ads-truststore \  
--keystore-type JKS \  
--keystore-password-file ads-truststore.pin \  
--alias ads-certificate \  
--use-existing-key-pair \  
--subject-dn "CN=ldap.example.com,O=Example Corporation,C=US" \  
--output-file ads.csr
```

2. Submit `ads.csr` to a CA for signing.
3. Export the server's private key into `ads.key`:

```
$ bin/manage-certificates export-private-key \  
--keystore ads-truststore \  
--keystore-password-file ads-truststore.pin \  
--alias ads-certificate \  
--output-file ads.key
```

4. Import the certificates obtained from the CA (the CA-signed server certificate, any intermediate certificates, and root CA certificate) into `ads-truststore.new`:

```
$ bin/manage-certificates import-certificate \  
--keystore ads-truststore.new \  
--keystore-type JKS \  
--keystore-password-file ads-truststore.pin \  
--alias ads-certificate \  
--private-key-file ads.key \  
--certificate-file new-ads.crt \  
--certificate-file intermediate.crt \  
--certificate-file root-ca.crt
```

Update the server configuration to use the new certificate

To update the server to use the desired key-pair, the `inter-server-certificate` property for the server instance must first be updated in the topology registry. The old and the new certificates may appear within their own begin and end headers in the `inter-server-certificate` property to support transitioning from the old certificate to the new one.

1. Export the server's old ads-certificate into `old-ads.crt`:

```
$ bin/manage-certificates export-certificate \
  --keystore ads-truststore \
  --keystore-password-file ads-truststore.pin \
  --alias ads-certificate \
  --export-certificate-chain \
  --output-file old-ads.crt
```

2. Concatenate the old, new certificate, and issuer certificates into one file. On Windows, an editor like notepad can be used. On Unix platforms, use the following command:

```
$ cat old-ads.crt new-ads.crt intermediate.crt root-ca.crt > chain.crt
```

3. Update the `inter-server-certificate` property for the server instance in the topology registry using `dsconfig`:

```
$ bin/dsconfig -n set-server-instance-prop \
  --instance-name <instance-name> \
  --set "inter-server-certificate<chain.crt"
```

Update the ads-truststore file to use the new key-pair

The server will still use the old `ads-certificate`. When the new `ads-certificate` needs to go into effect, the old `ads-truststore` file must be replaced with `ads-truststore.new` in the server's `config` directory.

```
$ mv ads-truststore.new ads-truststore
```

Retire the old certificate

The old certificate is retired by removing it from the topology registry when it has expired. All existing encrypted backups and LDIF exports are not affected because the public key in the old and new server certificates are the same, and the private key will be able to decrypt them.

```
$ cat new-ads.crt intermediate.crt root-ca.crt > chain.crt
```

```
$ bin/dsconfig -n set-server-instance-prop \
  --instance-name <instance-name> \
  --set "inter-server-certificate<chain.crt"
```

Use the Configuration API

PingData servers provide a Configuration API, which may be useful in situations where using LDAP to update the server configuration is not possible. The API is consistent with the System for Cross-domain Identity Management (SCIM) 2.0 protocol and uses JSON as a text exchange format, so all request headers should allow the `application/json` content type.

The server includes a servlet extension that provides read and write access to the server's configuration over HTTP. The extension is enabled by default for new installations, and can be enabled for existing deployments by simply adding the extension to one of the server's HTTP Connection Handlers, as follows:

```
$ bin/dsconfig set-connection-handler-prop \  
  --handler-name "HTTPS Connection Handler" \  
  --add http-servlet-extension:Configuration
```

The API is made available on the HTTPS Connection handler's `host:port` in the `/config` context. Due to the potentially sensitive nature of the server's configuration, the HTTPS Connection Handler should be used, for hosting the Configuration extension.

Authentication and authorization

Clients must use HTTP Basic authentication to authenticate to the Configuration API. If the username value is not a DN, then it will be resolved to a DN value using the identity mapper associated with the Configuration servlet. By default, the Configuration API uses an identity mapper that allows an entry's UID value to be used as a username. To customize this behavior, either customize the default identity mapper, or specify a different identity mapper using the Configuration servlet's `identity-mapper` property. For example:

```
$ bin/dsconfig set-http-servlet-extension-prop \  
  --extension-name Configuration \  
  --set "identity-mapper:Alternative Identity Mapper"
```

To access configuration information, users must have the appropriate privileges:

- To access the `cn=config` backend, users must have the `bypass-acl` privilege or be allowed access to the configuration using an ACI.
- To read configuration information, users must have the `config-read` privilege.
- To update the configuration, users must have the `config-write` privilege.

Relationship between the Configuration API and the dsconfig tool

The Configuration API is designed to mirror the `dsconfig` tool, using the same names for properties and object types. Property names are presented as hyphen case in `dsconfig` and as camel-case attributes in the API. In API requests that specify property names, case is not important. Therefore, `baseDN` is the same as `baseDn`. Object types are represented in hyphen case. API paths mirror what is in `dsconfig`. For example, the `dsconfig list-connection-handlers` command is analogous to the API's `/config/connection-handlers` path. Object types that appear in the schema URNs adhere to a `type:subtype` syntax. For example, a Local DB Backend's schema URN is `urn:unboundid:schemas:configuration:2.0:backend:local-db`. Like the `dsconfig` tool, all configuration updates made through the API are recorded in `logs/config-audit.log`.

The API includes the filter, sort, and pagination query parameters described by the SCIM specification. Specific attributes may be requested using the `attributes` query parameter, whose value must be a comma-delimited list of properties to be returned, for example `attributes=baseDN,description`. Likewise, attributes may be excluded from responses by specifying the `excludedAttributes` parameter. See [Sorting and Filtering with the Configuration API](#) for more information on query parameters.

Operations supported by the API are those typically found in REST APIs:

HTTP Method	Description	Related dsconfig Example
GET	Lists the attributes of an object when used with a path representing an object, such as <code>/config/global-configuration</code> or <code>/config/backends/userRoot</code> . Can also list objects when used with a path representing a parent relation, such as <code>/config/backends</code> .	<code>get-backend-prop</code> <code>list-backends</code> <code>get-global-configuration-prop</code>
POST	Creates a new instance of an object when used with a relation parent path, such as <code>config/backends</code> .	<code>create-backend</code>
PUT	Replaces the existing attributes of an object. A PUT operation is similar to a PATCH operation, except that the PATCH is determined by determining the difference between an existing target object and a supplied source object. Only those attributes in the source object are modified in the target object. The target object is specified using a path, such as <code>/config/backends/userRoot</code> .	<code>set-backend-prop</code> <code>set-global-configuration-prop</code>
PATCH	Updates the attributes of an existing object when used with a path representing an object, such as <code>/config/backends/userRoot</code> . See PATCH Example .	<code>set-backend-prop</code> <code>set-global-configuration-</code>

HTTP Method	Description	Related dsconfig Example
		prop
DELETE	Deletes an existing object when used with a path representing an object, such as <code>/config/backends/userRoot</code> .	delete-backend

The OPTIONS method can also be used to determine the operations permitted for a particular path.

Object names, such as `userRoot` in the Description column, must be URL-encoded in the path segment of a URL. For example, `%20` must be used in place of spaces, and `%25` is used in place of the percent (%) character. So the URL for accessing the HTTP Connection Handler object is:

```
/config/connection-handlers/http%20connection%20handler
```

GET Example

The following is a sample GET request for information about the `userRoot` backend:

```
GET /config/backends/userRoot
Host: example.com:5033
Accept: application/scim+json
```

The response:

```
{
  "schemas": [
    "urn:unboundid:schemas:configuration:2.0:backend:local-db"
  ],
  "id": "userRoot",
  "meta": {
    "resourceType": "Local DB Backend",
    "location": "http://localhost:5033/config/backends/userRoot"
  },
  "backendID": "userRoot2",
  "backgroundPrime": "false",
  "backupFilePermissions": "700",
  "baseDN": [
    "dc=example2,dc=com"
  ],
  "checkpointOnCloseCount": "2",
  "cleanerThreadWaitTime": "120000",
  "compressEntries": "false",
  "continuePrimeAfterCacheFull": "false",
  "dbBackgroundSyncInterval": "1 s",
  "dbCachePercent": "10",
  "dbCacheSize": "0 b",
  "dbCheckpointIntervalBytes": "20 mb",
  "dbCheckpointHighPriority": "false",
  "dbCheckpointWakeupInterval": "1 m",
  "dbCleanOnExplicitGC": "false",
  "dbCleanerMinUtilization": "75",
```

```

"dbCompactKeyPrefixes": "true",
"dbDirectory": "db",
"dbDirectoryPermissions": "700",
"dbEvictorCriticalPercentage": "0",
"dbEvictorLruOnly": "false",
"dbEvictorNodesPerScan": "10",
"dbFileCacheSize": "1000",
"dbImportCachePercent": "60",
"dbLogFileMax": "50 mb",
"dbLoggingFileHandlerOn": "true",
"dbLoggingLevel": "CONFIG",
"dbNumCleanerThreads": "0",
"dbNumLockTables": "0",
"dbRunCleaner": "true",
"dbTxnNoSync": "false",
"dbTxnWriteNoSync": "true",
"dbUseThreadLocalHandles": "true",
"deadlockRetryLimit": "10",
"defaultCacheMode": "cache-keys-and-values",
"defaultTxnMaxLockTimeout": "10 s",
"defaultTxnMinLockTimeout": "10 s",
"enabled": "false",
"explodedIndexEntryThreshold": "4000",
"exportThreadCount": "0",
"externalTxnDefaultBackendLockBehavior": "acquire-before-retries",
"externalTxnDefaultMaxLockTimeout": "100 ms",
"externalTxnDefaultMinLockTimeout": "100 ms",
"externalTxnDefaultRetryAttempts": "2",
"hashEntries": "false",
"id2childrenIndexEntryLimit": "66",
"importTempDirectory": "import-tmp",
"importThreadCount": "16",
"indexEntryLimit": "4000",
"isPrivateBackend": "false",
"javaClass": "com.unboundid.directory.server.backends.jeb.BackendImpl",
"jeProperty": [
  "je.cleaner.adjustUtilization=false",
  "je.nodeMaxEntries=32"
],
"numRecentChanges": "50000",
"offlineProcessDatabaseOpenTimeout": "1 h",
"primeAllIndexes": "true",
"primeMethod": [
  "none"
],
"primeThreadCount": "2",
"primeTimeLimit": "0 ms",
"processFiltersWithUndefinedAttributeTypes": "false",
"returnUnavailableForUntrustedIndex": "true",
"returnUnavailableWhenDisabled": "true",
"setDegradedAlertForUntrustedIndex": "true",
"setDegradedAlertWhenDisabled": "true",
"subtreeDeleteBatchSize": "5000",
"subtreeDeleteSizeLimit": "5000",
"uncachedId2entryCacheMode": "cache-keys-only",
"writabilityMode": "enabled"
}

```

GET list example

The following is a sample GET request for all local backends:

```
GET /config/backends
Host: example.com:5033
Accept: application/scim+json
```

The response (which has been shortened):

```
{
  "schemas": [
    "urn:ietf:params:scim:api:messages:2.0:ListResponse"
  ],
  "totalResults": 24,
  "Resources": [
    {
      "schemas": [
        "urn:unboundid:schemas:configuration:2.0:backend:ldif"
      ],
      "id": "adminRoot",
      "meta": {
        "resourceType": "LDIF Backend",
        "location": "http://localhost:5033/config/backends/adminRoot"
      },
      "backendID": "adminRoot",
      "backupFilePermissions": "700",
      "baseDN": [
        "cn=admin data"
      ],
      "enabled": "true",
      "isPrivateBackend": "true",
      "javaClass": "com.unboundid.directory.server.backends.LDIFBackend",
      "ldifFile": "config/admin-backend.ldif",
      "returnUnavailableWhenDisabled": "true",
      "setDegradedAlertWhenDisabled": "false",
      "writabilityMode": "enabled"
    },
    {
      "schemas": [
        "urn:unboundid:schemas:configuration:2.0:backend:trust-store"
      ],
      "id": "ads-truststore",
      "meta": {
        "resourceType": "Trust Store Backend",
        "location": "http://localhost:5033/config/backends/ads-truststore"
      },
      "backendID": "ads-truststore",
      "backupFilePermissions": "700",
      "baseDN": [
        "cn=ads-truststore"
      ],
      "enabled": "true",
      "javaClass":
```

```
"com.unboundid.directory.server.backends.TrustStoreBackend",
  "returnUnavailableWhenDisabled": "true",
  "setDegradedAlertWhenDisabled": "true",
  "trustStoreFile": "config/server.keystore",
  "trustStorePin": "*****",
  "trustStoreType": "JKS",
  "writabilityMode": "enabled"
},
{
  "schemas": [
    "urn:unboundid:schemas:configuration:2.0:backend:alarm"
  ],
  "id": "alarms",
  "meta": {
    "resourceType": "Alarm Backend",
    "location": "http://localhost:5033/config/backends/alarms"
  },
  ...
}
```

PATCH example

Configuration can be modified using the HTTP PATCH method. The PATCH request body is a JSON object formatted according to the SCIM patch request. The Configuration API, supports a subset of possible values for the `path` attribute, used to indicate the configuration attribute to modify.

The configuration object's attributes can be modified in the following ways. These operations are analogous to the `dsconfig modify-[object]` options.

- An operation to set the single-valued `description` attribute to a new value:

```
{
  "op" : "replace",
  "path" : "description",
  "value" : "A new backend."
}
```

is analogous to:

```
$ dsconfig set-backend-prop --backend-name userRoot \
  --set "description:A new backend"
```

- An operation to add a new value to the multi-valued `jeProperty` attribute:

```
{
  "op" : "add",
  "path" : "jeProperty",
  "value" : "je.env.backgroundReadLimit=0"
}
```

is analogous to:

Chapter 3: Managing the PingDataMetrics Server

```
$ dsconfig set-backend-prop --backend-name userRoot \  
  --add je-property:je.env.backgroundReadLimit=0
```

- An operation to remove a value from a multi-valued property. In this case, `path` specifies a SCIM filter identifying the value to remove:

```
{  
  "op" : "remove",  
  "path" : "[jeProperty eq \"je.cleaner.adjustUtilization=false\"]"}
```

is analogous to:

```
$ dsconfig set-backend-prop --backend-name userRoot \  
  --remove je-property:je.cleaner.adjustUtilization=false
```

- A second operation to remove a value from a multi-valued property, where the `path` specifies both an attribute to modify, and a SCIM filter whose attribute is `value`:

```
{  
  "op" : "remove",  
  "path" : "jeProperty[value eq \"je.nodeMaxEntries=32\"]"}
```

is analogous to:

```
$ dsconfig set-backend-prop --backend-name userRoot \  
  --remove je-property:je.nodeMaxEntries=32
```

- An option to remove one or more values of a multi-valued attribute. This has the effect of restoring the attribute's value to its default value:

```
{  
  "op" : "remove",  
  "path" : "id2childrenIndexEntryLimit"}
```

is analogous to:

```
$ dsconfig set-backend-prop --backend-name userRoot \  
  --reset id2childrenIndexEntryLimit
```

The following is the full example request. The API responds with the entire modified configuration object, which may include a SCIM extension attribute

`urn:unboundid:schemas:configuration:messages` containing additional instructions:

Example request:

```
PATCH /config/backends/userRoot  
Host: example.com:5033  
Accept: application/scim+json
```

```
{  
  "schemas" : [ "urn:ietf:params:scim:api:messages:2.0:PatchOp" ],  
  "Operations" : [ {
```

```

    "op" : "replace",
    "path" : "description",
    "value" : "A new backend."
  }, {
    "op" : "add",
    "path" : "jeProperty",
    "value" : "je.env.backgroundReadLimit=0"
  }, {
    "op" : "remove",
    "path" : "[jeProperty eq \"je.cleaner.adjustUtilization=false\"]"
  }, {
    "op" : "remove",
    "path" : "jeProperty[value eq \"je.nodeMaxEntries=32\"]"
  }, {
    "op" : "remove",
    "path" : "id2childrenIndexEntryLimit"
  } ]
}

```

Example response:

```

{
  "schemas": [
    "urn:unboundid:schemas:configuration:2.0:backend:local-db"
  ],
  "id": "userRoot2",
  "meta": {
    "resourceType": "Local DB Backend",
    "location": "http://example.com:5033/config/backends/userRoot2"
  },
  "backendID": "userRoot2",
  "backgroundPrime": "false",
  "backupFilePermissions": "700",
  "baseDN": [
    "dc=example2,dc=com"
  ],
  "checkpointOnCloseCount": "2",
  "cleanerThreadWaitTime": "120000",
  "compressEntries": "false",
  "continuePrimeAfterCacheFull": "false",
  "dbBackgroundSyncInterval": "1 s",
  "dbCachePercent": "10",
  "dbCacheSize": "0 b",
  "dbCheckpointIntervalBytes": "20 mb",
  "dbCheckpointIntervalHighPriority": "false",
  "dbCheckpointIntervalWakeup": "1 m",
  "dbCleanOnExplicitGC": "false",
  "dbCleanerMinUtilization": "75",
  "dbCompactKeyPrefixes": "true",
  "dbDirectory": "db",
  "dbDirectoryPermissions": "700",
  "dbEvictorCriticalPercentage": "0",
  "dbEvictorLruOnly": "false",
  "dbEvictorNodesPerScan": "10",

```

Chapter 3: Managing the PingDataMetrics Server

```
"dbFileCacheSize": "1000",
"dbImportCachePercent": "60",
"dbLogFileMax": "50 mb",
"dbLoggingFileHandlerOn": "true",
"dbLoggingLevel": "CONFIG",
"dbNumCleanerThreads": "0",
"dbNumLockTables": "0",
"dbRunCleaner": "true",
"dbTxnNoSync": "false",
"dbTxnWriteNoSync": "true",
"dbUseThreadLocalHandles": "true",
"deadlockRetryLimit": "10",
"defaultCacheMode": "cache-keys-and-values",
"defaultTxnMaxLockTimeout": "10 s",
"defaultTxnMinLockTimeout": "10 s",
"description": "123",
"enabled": "false",
"explodedIndexEntryThreshold": "4000",
"exportThreadCount": "0",
"externalTxnDefaultBackendLockBehavior": "acquire-before-retries",
"externalTxnDefaultMaxLockTimeout": "100 ms",
"externalTxnDefaultMinLockTimeout": "100 ms",
"externalTxnDefaultRetryAttempts": "2",
"hashEntries": "false",
"importTempDirectory": "import-tmp",
"importThreadCount": "16",
"indexEntryLimit": "4000",
"isPrivateBackend": "false",
"javaClass": "com.unboundid.directory.server.backends.jeb.BackendImpl",
"jeProperty": [
  "\"je.env.backgroundReadLimit=0\"",
],
"numRecentChanges": "50000",
"offlineProcessDatabaseOpenTimeout": "1 h",
"primeAllIndexes": "true",
"primeMethod": [
  "none"
],
"primeThreadCount": "2",
"primeTimeLimit": "0 ms",
"processFiltersWithUndefinedAttributeTypes": "false",
"returnUnavailableForUntrustedIndex": "true",
"returnUnavailableWhenDisabled": "true",
"setDegradedAlertForUntrustedIndex": "true",
"setDegradedAlertWhenDisabled": "true",
"subtreeDeleteBatchSize": "5000",
"subtreeDeleteSizeLimit": "5000",
"uncachedId2entryCacheMode": "cache-keys-only",
"writabilityMode": "enabled",
"urn:unboundid:schemas:configuration:messages:2.0": {
  "requiredActions": [
    {
      "property": "jeProperty",
```

```

    "type": "componentRestart",
    "synopsis": "In order for this modification to take effect,
    the component must be restarted, either by disabling and
    re-enabling it, or by restarting the server"
  },
  {
    "property": "id2childrenIndexEntryLimit",
    "type": "other",
    "synopsis": "If this limit is increased, then the contents
    of the backend must be exported to LDIF and re-imported to
    allow the new limit to be used for any id2children keys
    that had already hit the previous limit."
  }
]
}
}

```

API paths

The Configuration API is available under the `/config` path. A full listing of root sub-paths can be obtained from the `/config/ResourceTypes` endpoint:

```

GET /config/ResourceTypes
Host: example.com:5033
Accept: application/scim+json

```

Sample response (abbreviated):

```

{
  "schemas": [
    "urn:ietf:params:scim:api:messages:2.0:ListResponse"
  ],
  "totalResults": 520,
  "Resources": [
    {
      "schemas": [
        "urn:ietf:params:scim:schemas:core:2.0:ResourceType"
      ],
      "id": "dsee-compat-access-control-handler",
      "name": "DSEE Compat Access Control Handler",
      "description": "The DSEE Compat Access Control
        Handler provides an implementation that uses syntax
        compatible with the Sun Java System Directory Server
        Enterprise Edition access control handler.",
      "endpoint": "/access-control-handler",
      "meta": {
        "resourceType": "ResourceType",
        "location": "http://example.com:5033/config/ResourceTypes/dsee-compat-
        access-control-handler"
      }
    },
    {

```


Chapter 3: Managing the PingDataMetrics Server

```
"schemas": [
  "urn:ietf:params:scim:schemas:core:2.0:ResourceType"
],
"id": "access-control-handler",
"name": "Access Control Handler",
"description": "Access Control Handlers manage the
  application-wide access control. The server's access
  control handler is defined through an extensible
  interface, so that alternate implementations can be created.
  Only one access control handler may be active in the server
  at any given time.",
"endpoint": "/access-control-handler",
"meta": {
  "resourceType": "ResourceType",
  "location": "http://example.com:5033/config/ResourceTypes/access-
control-handler"
}
},
{
...

```

The response's `endpoint` elements enumerate all available sub-paths. The path `/config/access-control-handler` in the example can be used to get a list of existing access control handlers, and create new ones. A path containing an object name like `/config/backends/{backendName}`, where `{backendName}` corresponds to an existing backend (such as `userRoot`) can be used to obtain an object's properties, update the properties, or delete the object.

Some paths reflect hierarchical relationships between objects. For example, properties of a local DB VLV index for the `userRoot` backend are available using a path like `/config/backends/userRoot/local-db-indexes/uid`. Some paths represent singleton objects, which have properties but cannot be deleted nor created. These paths can be differentiated from others by their singular, rather than plural, relation name (for example `global-configuration`).

Sorting and filtering configuration objects

The Configuration API supports SCIM parameters for filter, sorting, and pagination. Search operations can specify a SCIM filter used to narrow the number of elements returned. See the SCIM specification for the full set of operations for SCIM filters. Clients may also specify sort parameters, or paging parameters. As previously mentioned, clients may specify attributes to include or exclude in both get and list operations.

GET Parameters for Sorting and Filtering

GET Parameter	Description
filter	Values can be simple SCIM filters such as <code>id eq "userRoot"</code> or compound filters like <code>meta.resourceType eq "Local DB Backend" and baseDn co "dc=example,dc=com"</code> .
sortBy	Specifies a property value by which to sort.
sortOrder	Specifies either <code>ascending</code> or <code>descending</code> alphabetical order.
startIndex	1-based index of the first result to return.
count	Indicates the number of results per page.

Update properties

The Configuration API supports the HTTP PUT method as an alternative to modifying objects with HTTP PATCH. With PUT, the server computes the differences between the object in the request with the current version in the server, and performs modifications where necessary. The server will never remove attributes that are not specified in the request. The API responds with the entire modified object.

Request:

```
PUT /config/backends/userRoot
Host: example.com:5033
Accept: application/scim+json
{
  "description" : "A new description."
}
```

Response:

```
{
  "schemas": [
    "urn:unboundid:schemas:configuration:2.0:backend:local-db"
  ],
  "id": "userRoot",
  "meta": {
    "resourceType": "Local DB Backend",
    "location": "http://example.com:5033/config/backends/userRoot"
  },
  "backendID": "userRoot",
  "backgroundPrime": "false",
  "backupFilePermissions": "700",
  "baseDN": [
    "dc=example,dc=com"
  ],
  "checkpointOnCloseCount": "2",
  "cleanerThreadWaitTime": "120000",
  "compressEntries": "false",
```

Chapter 3: Managing the PingDataMetrics Server

```
"continuePrimeAfterCacheFull": "false",
"dbBackgroundSyncInterval": "1 s",
"dbCachePercent": "25",
"dbCacheSize": "0 b",
"dbCheckpointIntervalBytes": "20 mb",
"dbCheckpointHighPriority": "false",
"dbCheckpointWakeupInterval": "30 s",
"dbCleanOnExplicitGC": "false",
"dbCleanerMinUtilization": "75",
"dbCompactKeyPrefixes": "true",
"dbDirectory": "db",
"dbDirectoryPermissions": "700",
"dbEvictorCriticalPercentage": "5",
"dbEvictorLruOnly": "false",
"dbEvictorNodesPerScan": "10",
"dbFileCacheSize": "1000",
"dbImportCachePercent": "60",
"dbLogFileMax": "50 mb",
"dbLoggingFileHandlerOn": "true",
"dbLoggingLevel": "CONFIG",
"dbNumCleanerThreads": "1",
"dbNumLockTables": "0",
"dbRunCleaner": "true",
"dbTxnNoSync": "false",
"dbTxnWriteNoSync": "true",
"dbUseThreadLocalHandles": "true",
"deadlockRetryLimit": "10",
"defaultCacheMode": "cache-keys-and-values",
"defaultTxnMaxLockTimeout": "10 s",
"defaultTxnMinLockTimeout": "10 s",
"description": "abc",
"enabled": "true",
"explodedIndexEntryThreshold": "4000",
"exportThreadCount": "0",
"externalTxnDefaultBackendLockBehavior": "acquire-before-retries",
"externalTxnDefaultMaxLockTimeout": "100 ms",
"externalTxnDefaultMinLockTimeout": "100 ms",
"externalTxnDefaultRetryAttempts": "2",
"hashEntries": "true",
"importTempDirectory": "import-tmp",
"importThreadCount": "16",
"indexEntryLimit": "4000",
"isPrivateBackend": "false",
"javaClass": "com.unboundid.directory.server.backends.jeb.BackendImpl",
"numRecentChanges": "50000",
"offlineProcessDatabaseOpenTimeout": "1 h",
"primeAllIndexes": "true",
"primeMethod": [
  "none"
],
"primeThreadCount": "2",
"primeTimeLimit": "0 ms",
"processFiltersWithUndefinedAttributeTypes": "false",
```

```

"returnUnavailableForUntrustedIndex": "true",
"returnUnavailableWhenDisabled": "true",
"setDegradedAlertForUntrustedIndex": "true",
"setDegradedAlertWhenDisabled": "true",
"subtreeDeleteBatchSize": "5000",
"subtreeDeleteSizeLimit": "100000",
"uncachedId2entryCacheMode": "cache-keys-only",
"writabilityMode": "enabled"
}

```

Administrative actions

Updating a property may require an administrative action before the change can take effect. If so, the server will return 200 Success, and any actions are returned in the `urn:unboundid:schemas:configuration:messages:2.0` section of the JSON response that represents the entire object that was created or modified.

For example, changing the `jeProperty` of a backend will result in the following:

```

"urn:unboundid:schemas:configuration:messages:2.0": {
  "requiredActions": [
    {
      "property": "baseContextPath",
      "type": "componentRestart",
      "synopsis": "In order for this modification to
        take effect, the component must be restarted,
        either by disabling and re-enabling it, or by
        restarting the server"
    },
    {
      "property": "id2childrenIndexEntryLimit",
      "type": "other",
      "synopsis": "If this limit is increased, then the
        contents of the backend must be exported to LDIF
        and re-imported to allow the new limit to be used
        for any id2children keys that had already hit the
        previous limit."
    }
  ]
}
...

```

Update servers and server groups

Servers can be configured as part of a server group, so that configuration changes that are applied to a single server, are then applied to all servers in a group. When managing a server that is a member of a server group, creating or updating objects using the Configuration API

requires the `applyChangeTo` query parameter. The behavior and acceptable values for this parameter are identical to the `dsconfig` parameter of the same name. A value of `singleServer` or `serverGroup` can be specified. For example:

```
https://example.com:5033/config/Backends/userRoot?applyChangeTo=singleServer
```

Note

This does not apply to mirrored subtree objects, which include Topology and Cluster level objects. Changes made to mirrored objects are applied to all objects in the subtree.

Configuration API Responses

Clients of the API should examine the HTTP response code in order to determine the success or failure of a request. The following are response codes and their meanings:

Response Code	Description	Response Body
200 Success	The requested operation succeeded, with the response body being the configuration object that was created or modified. If further actions are required, they are included in the <code>urn:unboundid:schemas:configuration:messages:2.0</code> object.	List of objects, or object properties, administrative actions.
204 No Content	The requested operation succeeded and no further information has been provided, such as in the case of a DELETE operation.	None.
400 Bad Request	The request contents are incorrectly formatted or a request is made for an invalid API version.	Error summary and optional message.
401 Unauthorized	User authentication is required. Some user agents such as browsers may respond by prompting for credentials. If the request had specified credentials in an Authorization header, they are invalid.	None.
403 Forbidden	The requested operation is forbidden either because the user does not have sufficient privileges or some other constraint such as an object is edit-only and cannot be deleted.	None.
404 Not Found	The requested path does not refer to an existing object or object relation.	Error summary and optional message.
409 Conflict	The requested operation could not be performed due to the current state of the configuration. For example, an attempt was made to create an object that already exists or an attempt was made to delete an object that is referred to by another object.	Error summary and optional message.
415 Unsupported Media Type	The request is such that the Accept header does not indicate that JSON is an acceptable format for a response.	None.
500 Server Error	The server encountered an unexpected error. Please report server errors to customer support.	Error summary and optional message.

An application that uses the Configuration API should limit dependencies on particular text appearing in error message content. These messages may change, and their presence may depend on server configuration. Use the HTTP return code and the context of the request to create a client error message. The following is an example encoded error message:

```
{
  "schemas": [
    "urn:ietf:params:scim:api:messages:2.0:Error"
  ],
  "status": 404,
  "scimType": null,
  "detail": "The Local DB Index does not exist."
}
```

Domain Name Service (DNS) caching

If needed, two global configuration properties can be used to control the caching of hostname-to-numeric IP address (DNS lookup) results returned from the name resolution services of the underlying operating system. Use the `dsconfig` tool to configure these properties.

network-address-cache-ttl – Sets the Java system property `networkaddress.cache.ttl`, and controls the length of time in seconds that a hostname-to-IP address mapping can be cached. The default behavior is to keep resolution results for one hour (3600 seconds). This setting applies to the server and all extensions loaded by the server.

network-address-outage-cache-enabled – Caches hostname-to-IP address results in the event of a DNS outage. This is set to `true` by default, meaning name resolution results are cached. Unexpected service interruptions may occur during planned or unplanned maintenance, network outages or an infrastructure attack. This cache may allow the server to function during a DNS outage with minimal impact. This cache is not available to server extensions.

IP address reverse name lookups

Ping Identity servers do not explicitly perform numeric IP address-to-hostname lookups. However, address masks configured in Access Control Lists (ACIs), Connection Handlers, Connection Criteria, and Certificate handshake processing may trigger implicit reverse name

lookups. For more information about how address masks are configured in the server, review the following information for each server:

- ACI dns: bind rules under *Managing Access Control* (PingDirectory Server and PingDirectoryProxy Servers)
- ds-auth-allowed-address: *Adding Operational Attributes that Restrict Authentication* (PingDirectory Server)
- Connection Criteria: *Restricting Server Access Based on Client IP Address* (PingDirectory Server and PingDirectoryProxy Servers)
- Connection Handlers: restrict server access using Connection Handlers (Configuration Reference Guide for all PingData servers)

Configure traffic through a load balancer

If a PingData server is sitting behind an intermediate HTTP server, such as a load balancer, a reverse proxy, or a cache, it will log incoming requests as originating with the intermediate HTTP server instead of the client that actually sent the request. If the actual client's IP address should be recorded to the trace log, enable `X-Forwarded-*` handling in both the intermediate HTTP server and the PingData server. See the product documentation for the device type. For PingData servers:

- Edit the appropriate Connection Handler object (HTTPS or HTTP) and set `use-forwarded-headers` to `true`.
- When `use-forwarded-headers` is set to `true`, the server will use the client IP address and port information in the `X-Forwarded-*` headers instead of the address and port of the entity that's actually sending the request, the load balancer. This client address information will show up in logs where one would normally expect it to show up, such as in the `from` field of the HTTP REQUEST and HTTP RESPONSE messages.

Configure authentication with a SASL external certificate

By default, the PingDataMetrics Server authenticates to the PingDirectory Server using LDAP simple authentication (with a bind DN and a password). However, the PingDataMetrics Server can be configured to use SASL EXTERNAL to authenticate to the PingDirectory Server with a client certificate.

Note

This procedure assumes that PingDataMetrics Server instances are installed and configured to communicate with the backend PingDirectory Server instances using either SSL or StartTLS.

After the servers are configured, perform the following steps to configure SASL EXTERNAL authentication:

1. Create a JKS keystore that includes a public and private key pair for a certificate that the PingDataMetrics Server instance(s) will use to authenticate to the PingDirectory Server instance(s). Run the following command in the instance root of one of the PingDataMetrics Server instances. When prompted for a keystore password, enter a strong password to protect the certificate. When prompted for the key password, press **ENTER** to use the keystore password to protect the private key:

```
$ keytool -genkeypair \  
-keystore config/metrics-user-keystore \  
-storetype JKS \  
-keyalg RSA \  
-keysize 2048 \  
-alias metrics-user-cert \  
-dname "cn=Metrics User,cn=Root DNs,cn=config" \  
-validity 7300
```

2. Create a `config/metrics-user-keystore.pin` file that contains a single line that is the keystore password provided in the previous step.
3. If there are other PingDataMetrics Server instances in the topology, copy the `metrics-user-keystore` and `metrics-user-keystore.pin` files into the `config` directory for all instances.
4. Use the following command to export the public component of the user certificate to a text file:

```
$ keytool -export \  
-keystore config/metrics-user-keystore \  
-alias metrics-user-cert \  
-file config/metrics-user-cert.txt
```

5. Copy the `metrics-user-cert.txt` file into the `config` directory of all PingDirectory Server instances. Import that certificate into each server's primary trust store by running the following command from the server root. When prompted for the keystore password, enter the password contained in the `config/truststore.pin` file. When prompted to trust the certificate, enter **yes**.

```
$ keytool -import \  
-keystore config/truststore \  
-alias metrics-user-cert \  
-file config/metrics-user-cert.txt
```


6. Update the configuration for each PingDataMetrics Server instance to create a new key manager provider that will obtain its certificate from the `config/metrics-user-keystore` file. Run the following `dsconfig` command from the server root:

```
$ dsconfig create-key-manager-provider \  
  --provider-name "Metrics User Certificate" \  
  --type file-based \  
  --set enabled:true \  
  --set key-store-file:config/metrics-user-keystore \  
  --set key-store-type:JKS \  
  --set key-store-pin-file:config/metrics-user-keystore.pin
```

7. Update the configuration for each LDAP external server in each PingDataMetrics Server instance to use the newly-created key manager provider, and also to use SASL EXTERNAL authentication instead of LDAP simple authentication. Run the following `dsconfig` command:

```
$ dsconfig set-external-server-prop \  
  --server-name ds1.example.com:636 \  
  --set authentication-method:external \  
  --set "key-manager-provider:Metrics User Certificate"
```

After these changes, the PingDataMetrics Server should re-establish connections to the LDAP external server and authenticate with SASL EXTERNAL. Verify that the PingDataMetrics Server is still able to communicate with all backend servers by running the `bin/status` command. All of the servers listed in the "--- LDAP External Servers ---" section should have a status of `Available`. Review the PingDirectory Server access log can to make sure that the BIND RESULT log messages used to authenticate the connections from the PingDataMetrics Server include `authType="SASL", saslMechanism="EXTERNAL", resultCode=0`, and `authDN="cn=Metrics User,cn=Root DNs,cn=config"`.

Server SDK extensions

Custom server extensions can be created with the Server SDK. Extension bundles are installed from a `.zip` archive or a file system directory. Use the `manage-extension` tool to install or update any extension that is packaged using the extension bundle format. It opens and loads the extension bundle, confirms the correct extension to install, stops the server if necessary, copies the bundle to the server install root, and then restarts the server.

Note

The `manage-extension` tool must be used with Java extensions packaged using the extension bundle format. For more information, see the "Building and Deploying Java-Based Extensions" section of the Server SDK documentation.

The Server SDK enables creating extensions for all PingData servers. Cross-product extensions include:

- Access Loggers
- Alert Handlers
- Error Loggers
- Key Manager Providers
- Monitor Providers
- Trust Manager Providers
- OAuth Token Handlers
- Manage Extension Plugins

Chapter 4: Collecting Data and Metrics

The PingDataMetrics Server polls all the monitored servers over LDAP to gather the following data:

- Status of each server.
- Alerts emitted by each server.
- Performance data exposed in the `cn=monitor` subtree of each server.

For a complete summary of the metrics and dimensions that can be exposed through the RESTful Metrics API, see the reference files located in the `docs/metrics-guide` directory. Most metrics have a count, minimum, maximum, and average.

Topics include:

[Metrics Overview](#)

[Query Overview](#)

[The query-metric Tool](#)

[Performance Data Collection](#)

[System Monitoring Data Collection](#)

[Server Clock Skew](#)

[Tune Data Collection](#)

[Data Processing](#)

[Monitoring for Service Level Agreements](#)

Metrics Overview

A metric corresponds to a single measurement made within the server. The PingDataMetrics Server collects three types of metrics:

- **Count metrics** – represent the number of times a specific event happens within the server. Examples of count metrics include the number of LDAP operations performed, network packets received, or new connections established.
- **Continuous-valued metrics** – measure things that always have a value. For example, these metrics include the amount of free disk space, the current number of connected clients, and the number of operations pending in the work queue.
- **Discrete metrics** – correspond to measurements that have both a value and a weight, such as the duration of an LDAP operation or the average duration of a checkpoint.

The statistics that can be applied to values depend on the metric type. Only count statistics are available for count metrics. Discrete metrics have count, average, and histogram statistics available, which expose a count of the values broken down into bucket ranges. Average, minimum, and maximum statistics are available for continuous-valued metrics.

Count Metrics

A count metric indicates the number of times a specific event happens within the server. For example, the number of packets received on a network interface during a measurement interval is a count metric. Each measurement returns the count of the number of packets received during that measurement interval only. The sample contains the number of occurrences, whether the measurement interval is five seconds or two minutes.

Another example of a count metric is the number of megabytes of data written to a disk device during a measurement interval. Using the COUNT statistic when querying for a count metric will return the sum of the counts. Count metrics can often be converted into a rate.

Continuous Metrics

A continuous metric is a measurement of a value where the thing being measured always has a valid value at each measurement point. For example, CPU percent busy is a continuous metric. For every sample CPU interval, a valid CPU percent busy measurement can be taken. A continuous metric differs from a count metric in that continuous metric samples cannot be

added across time in a meaningful way. Instead, continuous metric samples use average, minimum, and maximum statistics. To determine how busy the CPU has been since midnight, average, rather than sum, the samples since midnight.

Discrete Metrics

A discrete metric is a measurement that has both a value and a weight. Discrete metrics are different from continuous metrics because each measurement is weighted. A discrete metric is analogous to a weighted average and requires that multiple measurements be taken within a single sample interval. For example, LDAP operation response time is a discrete metric, where the actual response time of each operation is averaged, and the number of LDAP operations is provided as the weight. If no LDAP operations occur in a sample interval, the value would be zero and the weight would be zero.

Some continuous and discrete metrics may also report a minimum/maximum value if the measurement is composed of multiple sub-measurements. The minimum/maximum values are aggregated by averaging, so the values reflect the median.

Some discrete metrics may also convey histogram data. Histogram data represents an additional set of measurements that take individual measurements and place them into value ranges. The PingDataMetrics Server supports histograms with up to 15 value ranges. Histogram valued samples are unique because they give a picture of the distribution of the values, and because they more precisely answer the question of "How many samples are greater than X?"

Dimensions

Dimensions provide a means of aggregating and subdividing metric sample values in a way that logically follows what is actually measured. For example, metrics that measure disk activity have a `disk-device` dimension. Aggregating on the `disk-device` dimension shows the average disk activity for all disks, where pivoting (splitting) by the `disk-device` dimension shows the activity for specific disks.

Every metric has a logical instance dimension, which corresponds to the server on which the sample was created. Each metric may have up to three dimensions, which are defined in the metric definition.

Chapter 4: Collecting Data and Metrics

For example, the `sync-pipe-completed-ops` metric has two dimensions, the `pipe-name` and `pipe-result`. The `pipe-name` is the name of the sync pipe as configured for the PingDataSync Server. The `pipe-result` is one of the following values:

- `exception`
- `failed`
- `failed-at-resource`
- `failed-during-mapping`
- `match-multiple-at-dest`
- `no-match-at-dest`
- `already-exists-at-dest`
- `no-change-needed`
- `out-of-scope`
- `success`
- `aborted-by-plugin`
- `failed-in-plugin`

At each measurement interval for each sync pipe on each PingDataSync Server, there will be a value for each of the `pipe-result` values. So, for a single PingDataSync Server with two Sync Pipes, `pipe-one` and `pipe-two`, the samples generated for each sample period look like the following. The timestamp is constrained to time-only for brevity.

```
08:15:05, sync-pipe-completed-ops, pipe-one, exception, 1
08:15:05, sync-pipe-completed-ops, pipe-one, failed, 7
08:15:05, sync-pipe-completed-ops, pipe-one, failed-at-resource, 1
08:15:05, sync-pipe-completed-ops, pipe-one, failed-during-mapping, 1
08:15:05, sync-pipe-completed-ops, pipe-one, match-multiple-at-dest, 3
08:15:05, sync-pipe-completed-ops, pipe-one, no-match-at-dest, 0
08:15:05, sync-pipe-completed-ops, pipe-one, already-exists-at-dest, 0
08:15:05, sync-pipe-completed-ops, pipe-one, no-change-needed, 1
08:15:05, sync-pipe-completed-ops, pipe-one, out-of-scope, 1
08:15:05, sync-pipe-completed-ops, pipe-one, success, 125
08:15:05, sync-pipe-completed-ops, pipe-one, aborted-by-plugin, 1
08:15:05, sync-pipe-completed-ops, pipe-one, failed-in-plugin, 0
08:15:05, sync-pipe-completed-ops, pipe-two, exception, 3
08:15:05, sync-pipe-completed-ops, pipe-two, failed, 9
08:15:05, sync-pipe-completed-ops, pipe-two, failed-at-resource, 2
08:15:05, sync-pipe-completed-ops, pipe-two, failed-during-mapping, 1
08:15:05, sync-pipe-completed-ops, pipe-two, match-multiple-at-dest, 4
08:15:05, sync-pipe-completed-ops, pipe-two, no-match-at-dest, 0
08:15:05, sync-pipe-completed-ops, pipe-two, already-exists-at-dest, 0
08:15:05, sync-pipe-completed-ops, pipe-two, no-change-needed, 1
08:15:05, sync-pipe-completed-ops, pipe-two, out-of-scope, 1
```

```
08:15:05, sync-pipe-completed-ops, pipe-two, success, 217
08:15:05, sync-pipe-completed-ops, pipe-two, aborted-by-plugin, 1
08:15:05, sync-pipe-completed-ops, pipe-two, failed-in-plugin, 0
```

Compare how busy `pipe-one` is to `pipe-two` by pivoting on `pipe-name`. This results in the following:

```
pipe-one 141
pipe-two 239
```

Pivot by `pipe-result`, to get a set of counts that show the distribution of the counts of the specific error types, as well as the success and failure. This data provides a quick way of assessing the kinds of problems encountered by the Sync Pipes.

Dimensions provide a way to pivot or aggregate along a metric-specific axis. All metrics have the `instance` pivot and the `time` pivot. Metrics that support the histogram statistic can also have a `histogram` pivot.

Query Overview

A metric query consists of three components:

- The data used to calculate the query results.
- The aggregation method used on the data to calculate the query result.
- The format of the query result.

Select Query Data

The data used to generate the results of a metric query are driven by the following factors:

- Metric and statistic
- Time range
- Server instances included in the result (optional)
- Included dimension values (optional)
- Histogram range (optional)

Every query returns results for a single statistic and of a single metric. A query must include the time range used to generate the results. Time ranges can either be absolute dates (in ISO-8601 format) or relative dates (such as `-30m`). A relative start time offset is relative to the end time. A relative end time offset is relative to the current time. When no end time is specified, the server includes results up to the current time.

The time range and the desired number of points (for pivot by time) dictates the resolution of data used to process the query. For example, the finest granularity of data, one second resolution, is only kept for a few hours. It will not be used to satisfy a query spanning multiple days.

By default, all server instances that produce the metric are used to calculate the query results. However, the metric query can be restricted to one of the following:

- A specific list of servers
- Servers of a given type, such as PingDirectory Server
- Servers within a specific location

For metrics that include one or more dimensions, a query can be evaluated across a subset of dimension values. For example, the results returned for the `response-time` metric can be restricted to just the search and modify values of the `op-type` dimension.

For `discrete-valued` metrics that break their values down into histogram ranges, a query can count statistics applied to a subset of histogram buckets by specifying a minimum and/or maximum histogram value. For example, a query on the `response-time` metric could return a count of operations that took longer than 100 milliseconds.

Aggregate Query Results

A metric query can return the full, raw data that matches the query parameters, so that the server can aggregate metric results across time, server instance, dimension value, or histogram value. The server aggregates results, except when the query indicates not to, by using a pivot. The mechanism for aggregating the data depends on the type of metric. A pivot directs the query processor to not aggregate one component of the query data. A pivot can be based on time, server instance, a specific dimension, or histogram ranges.

- If no pivot is specified, the query returns a single number that represents the aggregation of all matched data. For example, a query with no pivot might return the total number of operations that have completed today.
- A single pivot results in one-dimensional data, such as a time-based chart with a single line or a simple bar chart.
- Two pivots results in two-dimensional data, such as a time-based chart with a separate line for each server instance, or a stacked bar chart that shows the number of completed operations broken down by server and operation type.

- Three pivots results in three-dimensional data, such as a stacked, grouped bar chart that shows completed operations broken down by server, operation type, and result.

Beyond aggregating multiple samples into one, the data returned by a metric query can be further manipulated. For example, queries can be scaled on the count statistic to return the count of events per second, per minute, or per hour. Counts of histogram values can be returned by a percentage of the total. For example, instead of returning the raw count of operations that took longer than 50 milliseconds to complete, the results could be returned as the percentage of all operations that took longer than 50 milliseconds to complete. A value of 0.02% is more meaningful than a value of 40.

Format Query Results

The query results can be converted into a format requested by the client, such as:

- CSV spreadsheet
- XML format
- JSON format

The query-metric Tool

The `query-metric` tool is a client application of the PingDataMetrics Server API that enables access to all the metrics gathered by the server. It includes subcommands that facilitate creating data queries for listing metrics, server instances, and dimension values. This tool runs in both interactive and non-interactive modes. Queries are formed using the following subcommands:

- `explore` – Creates a series of hyper-linked HTML files for a broad range of metrics. The tool generates these files by making a series of API queries for a set of servers and metrics. The tool highlights the breadth of available metrics and patterns or anomalies across multiple metrics. In interactive mode, the tool prompts for the servers and the metrics.
- `query` – Defines a query for specific data of interest. In interactive mode, the tool prompts for the server, metrics, dimensions, statistics, and pivot values. The tool can be used to request a data formatted in XML, JSON, or CSV.

To start the tool in interactive mode, enter the following command:

```
$ query-metric
```

Chapter 4: Collecting Data and Metrics

Or, specify a subcommand in interactive mode:

```
$ query-metric explore
```

In non-interactive mode, the tool generates a data table based on command-line input. For example, the following command requests information from the local PingDataMetrics Server listening on port 8080 and generates response-time and throughput data tables for Proxy Server instances in Austin for the previous two weeks:

```
$ query-metric explore \  
  --httpPort 8080 \  
  --instanceType proxy \  
  --instanceLocation Austin \  
  --metric response-time \  
  --metric throughput \  
  --startTime -2w
```

The following command line obtains a JSON formatted data table that shows average throughput for all PingDirectoryProxy Server instances, over time with 100 data points. Each line in the table represents either an application's search or modification throughput.

Throughput values are represented as operations per second:

```
$ query-metric query \  
  --hostname localhost \  
  --httpPort 8080 \  
  --username cn=user1,cn=api-users \  
  --password secret \  
  --table json \  
  --metric throughput \  
  --instanceType proxy \  
  --statistic average \  
  --pivot op-type \  
  --pivot application-name \  
  --dimension op-type:search,modify \  
  --rateScaling second \  
  --maxIntervals 100 \  
  --startTime 2012-09-01T17:41Z \  
  --endTime 2012-09-30T17:41Z
```

To see a list of all supported options, run the help option for the `query-metric` tool:

```
$ query-metric -?
```

Performance Data Collection

Performance data represents a majority of the data collected by the PingDataMetrics Server. Each server may produce hundreds of kilobytes of performance data per minute, though the

amount of data captured has little to no impact on the performance of the monitored system. By default, the PingDataMetrics Server stores performance data for 20 years. Configure the volume of performance data collected by each monitored server so that the PingDataMetrics Server can keep up with the flow.

The performance data model is a dimensional data model. Measurements can be taken on multiple simultaneous values that are distinguished by dimension values. For example, a response time metric provides the time in milliseconds it took a server to respond to an LDAP request. This `response-time` metric has two dimensions:

Application name – reflects the connection criteria of the request.

Operation type – corresponds to the LDAP operation, such as add, bind, or search.

If a server has 20 different connection criteria, each response-time sample may have 140 different values, one for each of the applications multiplied by the number of operation types.

The performance data captured on the monitored server has a record with the following fields.

Performance Data Fields

Name	Data Type	Description
Timestamp	Date	Time of measurement, using clock on the monitored server
Metric	String	Name of metric
Dimension	String	Values of dimensions 1 - 3
Count	Int	Number of measurements represented by this sample
Average	Double	Average value of this sample
Minimum	Double	Optional minimum value of this sample
Maximum	Double	Optional maximum value of this sample
Buckets	Int	Optional histogram data associated with this sample

When a performance record is imported into the PingDataMetrics Server, it is normalized to reduce the size of the record. The normalized record contains the following information.

Normalized Record in the PingDataMetrics Server

Name	Data Type	Description
batchID	Int	The ID of the batch of data to which this record belongs
sampleTime	Timestamp	The time the sample was captured or equivalent information after aggregation
metric_qual	Int	The ID of a structure that reflects the metric and all dimension values
definitionID	Int	ID of the histogram definition, if the data belong to a histogram-valued sample
count	Int	Number of measurements represented by this sample

Normalized Record in the PingDataMetrics Server

Name	Data Type	Description
avg_val	Real	Average value for this sample
min_val	Real	Minimum value for this sample
max_val	Real	Maximum value for this sample
val1-15	Long	Histogram bucket values

System Monitoring Data Collection

All servers have the ability to monitor their health and that of the host system. Servers do not collect any performance data until they are prepared by the PingDataMetrics Server. All of the important server and machine metrics are stored in the `cn=monitor` backend.

Stats Collector Plugin

The Stats Collector plugin is the primary driver of performance data collection for LDAP, server response, replication, local JE databases, and host system machine metrics. Stats Collector configuration determines the sample and collection intervals, granularity of data (basic, extended, verbose), types of host system collection (cpu, disk, network) and the type of data aggregation that occurs for LDAP application statistics. The Stats Collector plugin is configured with the `dsconfig` tool and collects data using LDAP queries. For example, the `--server-info:extended` option includes collection for the following:

- CPU
- JVM memory
- Memory
- Disk information
- Network information

The following are all options for the Stats Collector plugin:

```
>>>> Configure the properties of the Stats Collector Plugin
Property                               Value(s)
-----
1)  description                          -
2)  enabled                               false
3)  local-db-backend-info                 basic
4)  replication-info                     basic
5)  entry-cache-info                     basic
6)  host-info                             cpu, disk, network
7)  included-ldap-application             If per-application LDAP stats is enabled,
```

```

then stats will be included for all
applications.
8)  sample-interval      1 s
9)  collection-interval  500 ms
10) ldap-info            extended
11) server-info          basic
12) per-application-ldap-stats aggregate-only

```

System Utilization Monitors

The System Utilization Monitors interface directly with the host operating system to gather statistics about CPU utilization and idle states, memory consumption, disk input and output rates, and queue depths, as well as network packet transmit and receive activity.

Utilization metrics are gathered with externally invoked operating system commands, such as `iostat` and `netstat`, using platform-specific arguments and version-specific output parsing.

Enabling the Host System monitor provider automatically gathers CPU and memory utilization, but only optionally gathers disk and network information. Disk and network interfaces are enumerated in the configuration by device names (such as `eth0` or `lo`), and by disk device names (such as `sd1`, `sdab`, `sda2`, `scsi0`).

External Collector Daemon

The System Utilization monitor contains an embedded collector daemon that runs on systems affected by a Java process fork memory issue (RFE 5049299). When a process attempts to fork a child process, the system attempts to allocate the same amount of memory for the child process, which will likely fail when the parent process consumes a large amount of memory.

The embedded collector daemon is started automatically for the server and inspects the Host System Monitor provider configuration to conditionally determine whether the external daemon process is required.

The external collector daemon operates by having an internal table of repeatable commands that run on a schedule. The collector creates a simulated filesystem in the `<server-root>/logs` directory for each command type so that the Host System Monitor Provider can find the output of the most recently collected data.

Repeating commands use a subdirectory for each command type to keep results isolated from other command types and to help organize file cleanup. The filename of the output contains the

sample timestamp, such as `iostats-[sampltimestamp]`. If the collector daemon fails for any reason, the Host System Monitor provider is not left reading stale system data because the expected timestamp files is missing. To handle clock-edge timing, the monitor sampler will also look for data in a filename of the previous second. Timestamp files are deleted once their data have been collected.

The collector daemon runs with no inter-process communication and can be stopped if no longer necessary.

Server Clock Skew

Correlating metric samples from multiple servers requires that the timestamp associated with each sample from each monitored server is synchronized. The PingDataMetrics Server tracks system time information and makes it visible in the `cn=Monitored Server <servername>`, `cn=monitor` entry.

The `system-clock-skew-seconds` attribute indicates the difference between the PingDataMetrics Server system clock and the monitored server clock, in seconds. The larger this skew value, the less precision there is when comparing changes in data across servers.

While it is not necessary to keep the PingDataMetrics Server clock synchronized with all of the monitored servers, it can be convenient when issuing metric queries with time ranges specified by offsets. Because the offset is computed using the PingDataMetrics Server system clock, if this clock is very different from the monitored servers' system clocks, the start and end time of a metric query will not match the expected boundaries.

Tune Data Collection

Collecting all of the performance data at the most granular level from all of the servers may not be possible without a significant investment in hardware for the PingDataMetrics Server. Instead, tune data collection to fit within the limits of the existing PingDataMetrics Server hardware.

The remainder of this section describes several strategies for tuning data collection.

Reducing the Data Collected

If not all information collected by the PingDataMetrics Server is required, the Stats Collector Plugin's `entry-cache` property can be tuned using the `dsconfig` command-line tool. For example, to omit all metrics related to the entry cache set the `entry-cache-info` group on the monitored server:

```
$ bin/dsconfig set-plugin-prop --plugin-name "Stats Collector" \
  --set entry-cache-info:none
```

The server collects information for eight different information groups. Limit data collection to the devices of actual interest.

Reducing the Frequency of Data Collection

Monitored servers can produce metric samples every second, which is useful for short-duration changes. These samples are less useful hours later, after the per-second data is aggregated to per-minute data. Use the `dsconfig` tool to change the base sample production rate from the default of 1 second to 10 seconds:

```
$ bin/dsconfig set-plugin-prop --plugin-name "Stats Collector" \
  --set "sample-interval:10 seconds"
```

This change reduces the total data volume by about 90 percent.

Reducing the Frequency of Sample Block Creation

The number of sample blocks processed by the PingDataMetrics Server can also be reduced in a given time. By default, the monitored servers produce a new block of samples every 30 seconds. Increasing this to 60 seconds, while reducing the PingDataMetrics Server's polling rate to 60 seconds, reduces the sample processing overhead. Change the frequency at which the monitored servers create sample blocks using the following `dsconfig` command:

```
$ bin/dsconfig set-backend-prop --backend-name metrics \
  --set sample-flush-interval:60s
```

Reducing PingDataMetrics Server Impact on Performance

All Ping Identity servers expose performance data through the `cn=monitor` DN. Performance issues occur when data is read, either directly by an LDAP client, or by enabling either the Stats Logger or Stats Collector plugins.

Chapter 4: Collecting Data and Metrics

The Stats Logger plugin reads the configured monitors and writes the resulting values to a CSV file. The Stats Collector plugin also reads the configured monitors and writes the resulting values to a CSV file, but this file is made available for LDAP clients in `cn=metrics` DN. The Stats Collector CSV files are suitable for use by the PingDataMetrics Server, and contain one metric value per line.

Both the Stats Logger and the Stats Collector plugins are disabled by default. When enabled, each of these plugins adds an approximate 3% CPU utilization penalty, plus a negligible amount of disk I/O and JVM heap usage.

To enable the Stats Collector plugin, use `dsconfig` as follows:

```
$ bin/dsconfig set-plugin-prop --plugin-name "Stats Collector" \  
  --set enabled:true
```

The `monitored-servers` tool will enable the Stats Collector plugin on the monitored server.

Data Processing

When blocks of samples arrive in the PingDataMetrics Server, they are queued on disk and loaded into the database. Samples from a single server are processed in time-order, so that sample blocks with older data are always processed before a sample block containing newer data. The PingDataMetrics Server does not do time-correlation between blocks coming from different servers. So, server A samples from two hours ago may be loaded immediately after server B samples from two minutes ago. This flexibility enables servers to be unavailable to the PingDataMetrics Server, without affecting the overall system monitoring. Also, a query for data from server A and B may return data for server B but not server A, until the data queued for server A has been collected and imported. Samples collected from the PingDataMetrics Server itself are processed ahead of all other servers.

Importing Data

The PingDataMetrics Server polls all of the monitored servers at a regular interval. When new samples are available, the PingDataMetrics Server fetches them through LDAP. The PingDataMetrics Server has one dedicated thread taking sample blocks and converting them to the normalized form stored in the DBMS. The import queue's size is normally near zero, but under certain conditions it may become large. When the PingDataMetrics Server starts, it will

queue (for import) all sample blocks still on disk. Blocks that are older than two hours are discarded.

For example, if a monitored server becomes unavailable for an extended period of time, it will continue to queue blocks of samples locally. When it becomes available again, the PingDataMetrics Server collection poll of that server will capture hundreds or thousands of sample blocks. The PingDataMetrics Server captures the sample blocks at a much faster speed than it can import them, causing the queue to grow for a period of time. If the PingDataMetrics Server is stopped, this problem is compounded because all monitored servers will then have a backlog of sample blocks to be imported.

Aggregating Data

To maintain a size-limited DBMS while accumulating data over a period of years, the PingDataMetrics Server aggregates data into four different levels. Each level contains data with less time granularity, but covering a larger period of time. Data is aggregated from a lower (greater time granularity) to a higher level as soon as enough data for aggregation is available. For example, the level 0 data has one second granularity, and the level 1 data has one minute granularity. After level 0 has collected one minute's worth of data, the data from that minute can be aggregated to level 1.

To keep the data tables for each aggregation level at a constrained size, each aggregation level has a maximum age for the samples. When the samples are older than this age, they are deleted from the level. While aggregation occurs soon after the samples arrive in the level, pruning occurs only after all samples in a block have passed their age limit.

The PingDataMetrics Server attempts to collect data from all configured servers as efficiently as possible. However, Monitored Server availability, DBMS backlog, and PingDataMetrics Server load can all cause the data pipeline to slow down. The data aggregation system is designed to correctly handle gaps in the data.

The resolution of the aggregation levels cannot be changed, but the maximum age of each level can be configured. The following table lists the aggregation levels.

Aggregation Levels

Level	Resolution	Default Maximum Age	Maximum Age
0	1 second	2 hours	48 hours

Aggregation Levels

Level	Resolution	Default Maximum Age	Maximum Age
1	1 minute	7 days	34 days
2	1 hour	12 months	5 years
3	1 day	20 years	20 years

Monitoring for Service Level Agreements

The PingDataMetrics Server provides the ability to aggregate and track performance data for one or more service level agreements (SLAs). The server aggregates the data using an SLA object that tracks the current and historical performance of LDAP operations (throughput and response times) that are tied to specifically monitored applications. The SLA object consists of a tracked application name, one or more LDAP operations to be considered, a set of servers that contribute performance data to the SLA, and optionally, thresholds to generate alerts should the server exceed these limits.

Thresholds are optional configuration settings that enable the monitoring of performance data. Each threshold sets a limit that indicates a warning condition where the server's performance is nearing a limit and/or a critical condition. When the monitored server enters or ends a warning or critical state, the PingDataMetrics Server generates an alert. The generated alerts are the same as those created by the PingDirectory Server and PingDirectoryProxy Servers and can be routed through the Alert Handler to a monitoring console or administrator.


The SLA object can report the aggregate performance of all configured servers. The SLA object is configured with the following:

- **Designate Servers that Contribute to SLA Tracking** – The SLA object includes a Server Query component that is used to designate the servers that contribute to the SLA measurements.
- **REST API** – A REST API enables listing configured SLA objects and their current status. The PingDataMetrics Server REST API also enables listing alerts generated by SLA thresholds, and blending the alert information with the threshold information to provide a more contextual view of the tracked applications performance.

SLA Thresholds

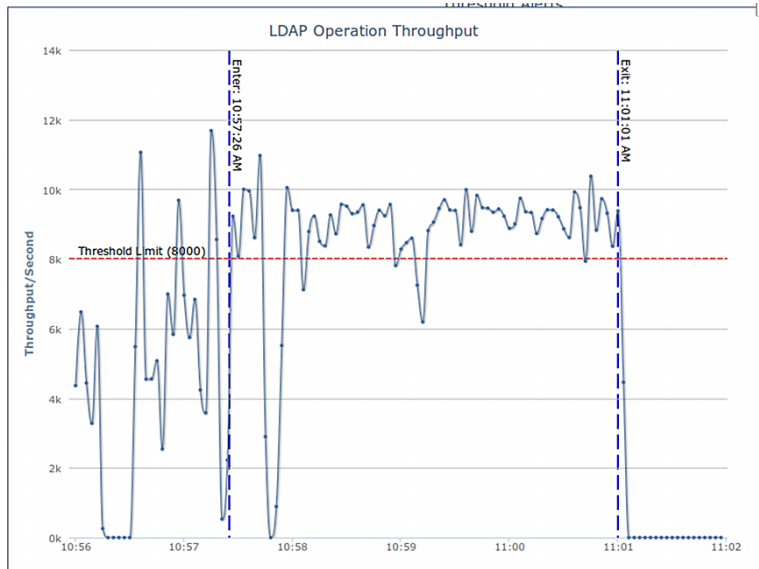
The PingDataMetrics Server uses a Monitoring Threshold mechanism that has two components:

- Spike Monitoring Threshold** – Used to configure a set of operational performance limits on a specific measurement, where the limit is specified as a percent change from the most recent measurement average value. A Spike Monitoring Threshold has warning and critical limits, and will enter or leave an alerted state when the monitored value exceeds either of the limits. This threshold is useful when the valid range of the measurement is not known in advance. This type of limit is useful in detecting short-term changes in a measurement that fluctuates broadly over time. The limit is applied in both positive and negative directions, so that this type of threshold can detect an upward or downward spike in the value. The following chart shows a spike monitoring threshold, where the red line is the average throughput/second and the green lines are the limits, showing the average window for the throughput.




- Static Level Monitoring Threshold** – Used to configure performance limits on a specific measurement, where the specified limits are fixed values that do not change over time. A Static Level Monitoring Threshold has warning and critical limits, and will enter or leave an alerted state when the monitored value exceeds any of the limits. The Static Level Monitoring Threshold is configured with static numeric limits, and is useful when the expected valid range of the measurement is known in advance.

Chapter 4: Collecting Data and Metrics



Threshold Time Line

The PingDataMetrics Server periodically evaluates each threshold, computing current value, current average, and alerted state. The default evaluation period is 30 seconds. Using the figure below, a threshold is evaluated at time 'Now.' The most recent data that threshold uses is one minute old (Newest Data). Each threshold evaluation requires at least one minute of new data (Minimum Data). At time 'Now,' the threshold is working with data that is between one and two minutes old (between Minimum Data and Newest Data).



The one minute delay between 'Now' and 'Newest Data' is not configurable. This delay ensures that the PingDataMetrics Server has had enough time to poll the monitored servers and get the most recent data. The one minute delay between 'Newest Data' and 'Minimum Data' is

configurable on a per-threshold basis for Spike-valued thresholds, but one minute is the minimum window. Generally, time between when a monitored performance anomaly occurs on a monitored system, and when an alert is created will be between two and three minutes.

Because the PingDataMetrics Server can capture metrics data with a very fine time resolution (one second data is the default), the data is often very "noisy." By default, the data is time-averaged (using five consecutive one-second samples to produce a single five-second value), and time-averaging will ultimately reduce the noise. However, "noisy" data can make it harder to choose an appropriate threshold limit value. If the limit value is too close to the noise levels, the threshold will alert due to values that have a very short time duration.

Each threshold is configured with a `minimum-time-to-trigger` property, which determines the minimum time allowed to exceed the threshold before an alert is generated, and a `minimum-time-to-exist` property that determines the time required for the threshold to exit an alerted state.

Configure an SLA Object

The SLA object relies on existing performance metrics and only aggregates the data for specific SLAs. Configure any number of SLA objects for monitored servers. For more information, see the *Ping Identity PingDataMetrics Server Configuration Reference* (HTML) documentation in the `<server-root>/docs` directory.

The following steps use the `dsconfig` command-line tool:

1. Create a server query that specifies which servers will contribute to SLA monitoring. In this example, the command specifies the PingDirectoryProxy Servers located in Austin.

```
$ bin/dsconfig create-server-query \
  --query-name "Austin Proxy Servers" \
  --set server-instance-type:proxy \
  --set server-instance-location:Austin
```

2. Create a `static-level` monitoring threshold. In this example, the alert condition is set to `entry`, which means that the server will generate an alert if the server enters a warning state (`alert-on-warn:true` and `warn-if-above:12`) or critical state (`critical-if-above:15`). When the server leaves its alerted state, an alert is generated (`alertcondition: exit`). The minimum amount of time that the threshold can be exceeded before an alert is generated is set to 15 seconds (`min-time-for-trigger:15s`).

```
$ bin/dsconfig create-monitoring-threshold \
  --threshold-name "15ms response time" \
```

Chapter 4: Collecting Data and Metrics

```
--type static-level \  
--set alert-condition:entry \  
--set alert-condition:exit \  
--set alert-on-warn:true \  
--set min-time-for-trigger:15s \  
--set min-time-for-exit:15s \  
--set warn-if-above:12 \  
--set critical-if-above:15
```

3. Create another `static-level` monitoring threshold. In this example, the alert condition is set to `entry`. The server generates an alert if it enters a warning state (`alert-on-warn:true` and `warn-if-above:4000`) or critical state (`critical-if-above:5000`). When the server leaves its alerted state, an alert is generated (`alertcondition:exit`). The minimum amount of time that the threshold can be exceeded before an alert is generated is set to 15 seconds.

```
$ bin/dsconfig create-monitoring-threshold \  
--threshold-name "5k ops/sec" \  
--type static-level \  
--set alert-condition:entry \  
--set alert-condition:exit \  
--set alert-on-warn:true \  
--set min-time-for-trigger:15s \  
--set min-time-for-exit:15s \  
--set warn-if-above:4000 \  
--set critical-if-above:5000
```

4. Create an SLA object that targets an SSO application and monitors the response and throughput times for LDAP bind operations. The response time threshold is set to 15ms. The throughput threshold is set to 5k operations per second. The targeted servers are the set of PingDirectoryProxy Servers, located in Austin.

```
$ bin/dsconfig create-ldap-sla \  
--sla-name "SSO Application" \  
--set enabled:true \  
--set "application-name:SSO Application" \  
--set "response-time-threshold-ms:15ms response time" \  
--set "throughput-threshold-ops-per-second:5k ops/sec" \  
--set ldap-op:bind \  
--set "sla-server-query:Austin Proxy Servers"
```

Chapter 5: Configuring Charts and Dashboards

The PingDataMetrics Server provides a set of dashboards with series of charts for each configured PingData server.

Charts can be built and customized with the PingDataMetrics Server Chart Builder tool. Dashboards and charts can be modified with Velocity templates.

Topics include:

[Available Dashboards](#)

[Available Charts for Identity Servers](#)

[The Chart Builder Tool](#)

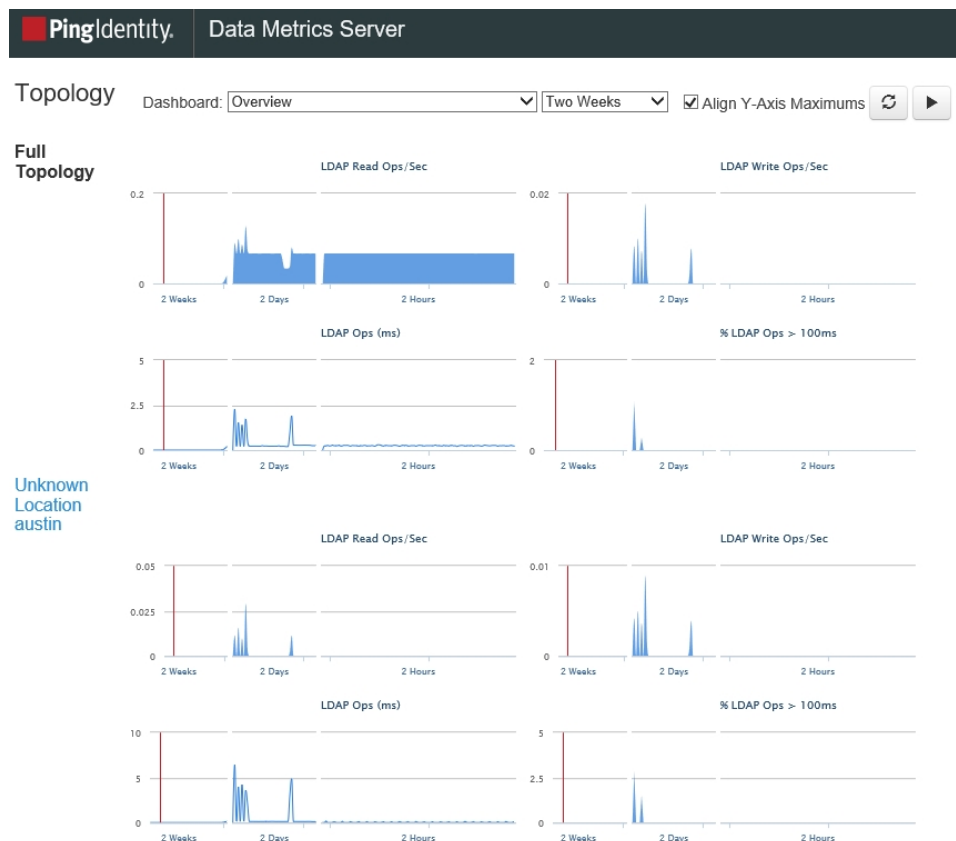
[Velocity Templates](#)

Available Dashboards

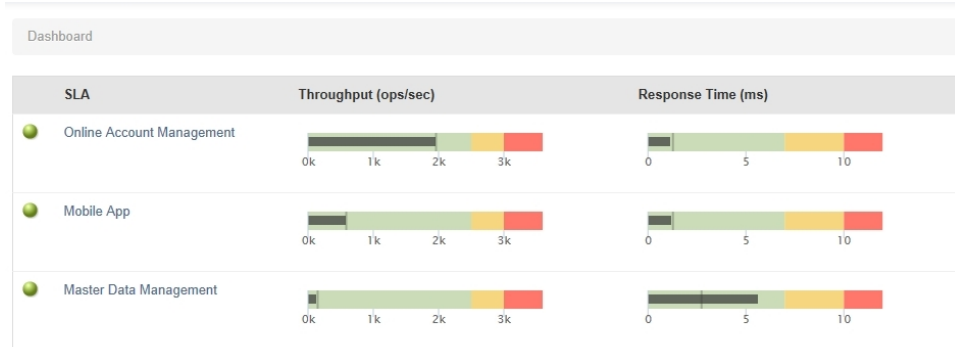
The PingDataMetrics Server includes several dashboards that can be used to display information for all servers in a data center, specific applications, or SLA specifics. The following dashboards are available:

ldap-dashboard – Displays charts for PingDirectory Server, PingDirectoryProxy Server, and PingDataSync Servers configured with the `monitored-servers` command. Charts are also displayed for the PingDataMetrics Server server. This dashboard is viewed from a browser at `http://<metrics-host>:<port>/view/ldap-dashboard`, and is easily customized. See [Customize the LDAP Dashboard](#). The charts can display information by:

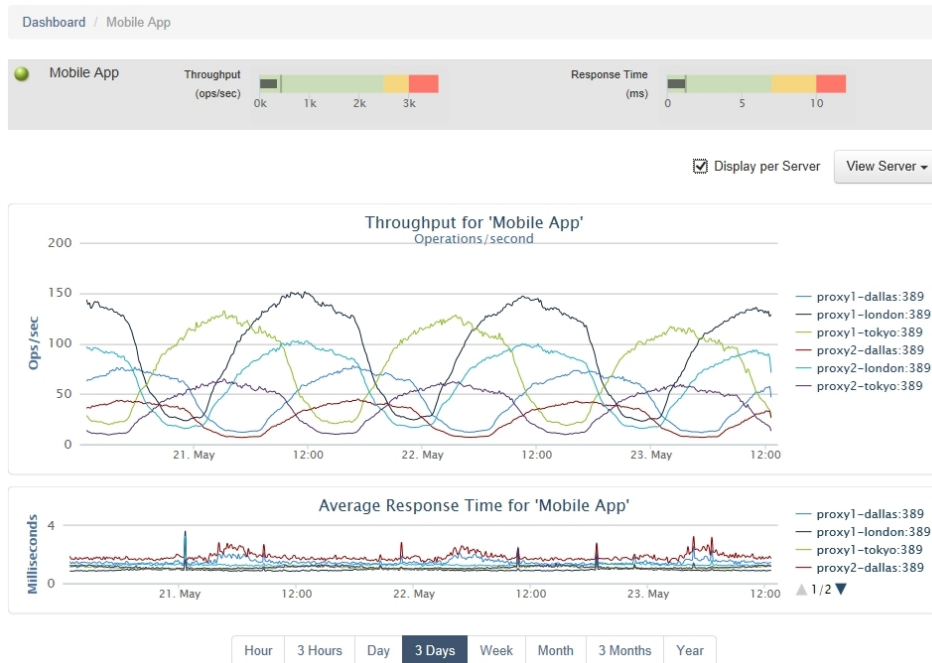
- Individual server, server location, or server type.
- Varying level of detail adjusted by server type.
- Time scale, providing either a recent or more historical data view.



sla-viewer – Displays throughput and response time graphs, and status for configured SLAs. This dashboard is viewed from a browser at `http://<metrics-host>:<port>/view/sla-viewer`. See [Monitoring for Service Level Agreements](#) for information about configuring SLAs.



sla-viewer-details – Displays SLA Viewer data and additional charts for response time and time ranges from the sla-viewer dashboard. Data can be viewed per server and includes server details.



demo-dashboard – Demonstrates how to display a set of charts for multiple servers and how to vary that set of charts per server type. This dashboard is viewed from a browser at

Chapter 5: Configuring Charts and Dashboards

`http://<metrics-host>:<port>/view/demo-dashboard`. This dashboard can be used as a starting point for custom dashboards.



A dashboard readme file provides general instructions for customizing any dashboard, and is located in:

```
<server-root>/config/dashboard/dashboard.README
```

Custom style sheets can be created and referenced in the dashboard template or styles can be configured for all charts. See [Chart Presentation Details](#) for information. The PingDataMetrics Server default style sheets should not be modified.

Customize the LDAP Dashboard

Dashboards are defined by Velocity templates. After servers are configured, the LDAP dashboard displays all metrics from monitored servers. See [Velocity Templates](#) for information about templates and template components.

Perform the following to configure the LDAP dashboard:

- The configuration file for the LDAP dashboard is `<server-root>/config/velocity/templates/_ldap-dashboard-config.vm`. This file should not be changed, but can be used as a guide for customization.

Note

Files within this directory that begin with an underscore (`_`) are templates that are referenced by each of the

dashboards. The `_ldap-dashboard-config.vm` template is the only file that contains all of the dashboard configuration inside the file. Configuration of other templates requires configuration of a corresponding dashboard file as well.

- The `_ldap-dashboard-config.vm` file references a template file that can be customized in `<server-root>/velocity/templates/_ldap-dashboard-config-overrides.vm`. This is the file that should be customized.
- Both files contain configuration instructions. The following can be customized in the LDAP dashboard overrides file:
 - The charts that display for each server type and their styles. See [Available Server Charts](#).
 - The charts that display for a data center and their styles.
 - The charts that display for an application type and their styles.
 - The default time resolution (two weeks, is the default for data displayed).
 - The size of the charts.

Debug Dashboard Customization

A debug option can be used in any Velocity template for exploring available information in the Velocity Context. This information includes the servers that are monitored and the metrics that are available. This option is included in the `ldap-dashboard` and `demo-dashboard` files:

```
## Uncomment this to have a window popup with detail of what's in the Velocity
Context.
##parse("_debug.vm")
##debug()
```

See [Velocity Templates](#) for more information.

Preserve Customized Files

Any files that are customized should be copied from the `config/velocity` subdirectories to the same subdirectory of the velocity directory under the server root (`<server-root>/velocity`). The files in `config/velocity` should not be modified. They are updated when the product is updated.

By default, any file of the same name under `<server-root>/velocity` will be loaded in place of `<server-root>/config/velocity`. This enables the preservation of customized files after a product upgrade.

After a product upgrade, review the files in `config/velocity` to determine if any changes should be incorporated into customized templates.

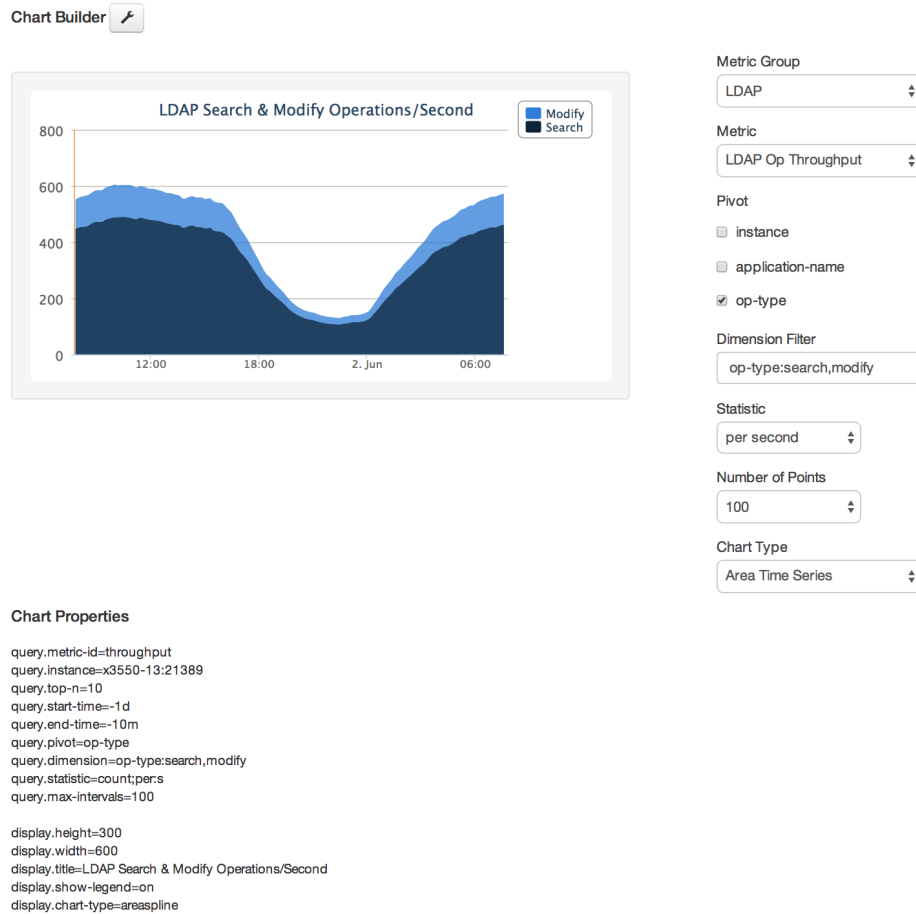
The Chart Builder Tool

The Chart Builder tool is used to create performance charts for all configured servers. As the settings in the Chart Builder are changed, the builder gathers the data from the PingDataMetrics Server using the PingDataMetrics Server REST API. Once configured, the dashboard page asynchronously fetches metric data for all charts, with each chart rendering when its data is returned. While most metric queries respond quickly (50-100ms), some queries may take longer. If the lag seems too long, consider making changes to the query to reduce the amount of data gathered.

Selecting specific instances and using dimension filters can decrease query time. The Chart Builder tool and the underlying libraries constrain a chart to a single metric. The size of each chart is determined by the library default size (300x300) and can be overridden in the chart properties file. There are times when the legends and labeling of a chart dictate the minimum size for a chart.

The Chart Builder tool (`chart-builder.vm`) is shipped with the PingDataMetrics Server and is enabled after installation at the following URL:

```
https://<metricshost>:<port>/view/chart-builder
```



The metrics parameters used to build the chart can be saved to a properties file and added to a dashboard. If not using the Chart Builder tool, the `_chart-definition.template` file in `<server-root>/config/dashboard/charts` provides instructions about manually creating charts and adding them to a dashboard.

Chart Presentation Details

Chart presentation details can be configured per chart or for all charts in the `_chart-definition.template` file. A properties file can be created for common styles and referenced in this file. Instructions for adding custom styles are included in the file.

The following is a sample of the chart parameters that are available:

- Colors used in the data series.
- Enable and disable a legend.

Chapter 5: Configuring Charts and Dashboards

- Location (top/bottom/left/right) of the legend.
- Background color.
- Thickness of the time-series lines (absolute or as a function of the # of plotted series).
- Macro expansion in the specified title.
- Sub-title (with macro expansion that includes metric-name and current date/time).

Chart Builder Parameters

Use the Chart Builder tool to build or adjust system and performance charts. When the configuration is set, copy the parameters into a properties file, and add the chart to a dashboard.

Chart Builder Parameters

Parameters	Description
Metric Group	Selects a specific group of metrics to be considered for charting.
Metric	Displays the specific metric. Open the drop-down list and hover over a metric to view a description of the particular metric.
Pivot	Splits the chart result into multiple series based on the pivot dimension chosen.
Dimension Filter	Filters the data based on the dimension(s) entered, such as the type of operations that can be viewed for an LDAP operations metric.
Statistic	<p>Displays the type of "measurement" that may exist for each metric. For example, each response-time sample contains:</p> <ul style="list-style-type: none">• number of operations (count)• average time-per-op (average)• histogram-of-operation-time (histogram) <p>On a per-sample basis, the PingDataMetrics Server stores the following: count, average, minimum, maximum, and histogram. Any metric can have one or more of five statistics, but not all statistics are equally valuable. Note the following points:</p> <ul style="list-style-type: none">• The minimum and maximum statistics may be of limited value, because as they are time-averaged, they go to extremes (min of minimums and max of maximums).• The count and histogram statistics have high fidelity over time because they time-aggregate perfectly.• The average statistic loses fidelity over time, because as the time-window for averaging gets larger, the highs and lows are lost.
Number of Points	If the number of points is set to 1, all chart types, except time series, may be used. If the number of points is > 1, then only time series charts may be used.
Chart Type	<p>Displays the chart based on the type:</p> <ul style="list-style-type: none">• Area Time Series

Chart Builder Parameters

Parameters	Description
	<ul style="list-style-type: none"> • Bar Chart • Column Chart • Pie Chart • Stacked Bar Chart • Stacked Column Chart • Time Series
Chart Properties	Displays the generated chart properties. Copy the query properties into a properties file.
Data URL	The API for getting the data. This can be used to call the chart into a third-party client application.

Chart Properties File

Each dashboard uses a Velocity template (`<name>.vm`) and a set of chart properties files to render the charts. As charts are configured with the Chart Builder tool, the tool generates the corresponding properties for each customized item. The metrics configuration can be copied into a properties file and added to a dashboard. If no values are specified for a given property, the property will use a default value from the `<server-root>/config/dashboard/charts/_chart-definition.template` file. All properties and their descriptions are listed in this file.

The properties in the chart definition file are broken into two groups: properties that start with `display` affect the display of the data, and properties that start with `query` affect the metric query. When building a new chart, just copy the query parameters into a properties file. In general, display options should be referenced from common styles defined in the `_chart-definition.template` file, or the styles defined for a dashboard.

Available Charts for PingData Servers

The following are the default charts that display on the LDAP Dashboard for each configured server. These and additional charts for server and system metrics reside in `<server-root>/config/dashboard/charts`. They can be modified or used to create new charts.

Charts for All Servers

The following charts are displayed on the LDAP Dashboard for PingDirectory Server and PingDirectoryProxy Servers:

- LDAP Read Operations Per Second
- LDAP Write Operations Per Second
- LDAP Response Time
- LDAP Response Time Outliers
- LDAP Worker Thread Percent Busy
- LDAP Average Operations in Progress
- LDAP Average Queue Size
- LDAP Open Connections
- LDAP New Connections
- System CPU
- System Memory Percent Free
- System Network Read MB
- System Network Write MB
- System Disk Busy
- System Disk Service Wait
- System Disk Read MB
- System Disk Write MB

PingDirectory Server Charts

The following charts are displayed on the dashboard for the PingDirectory Server:

- Replication Backlog
- Replication Oldest Change
- Replication Unresolved Naming Conflicts
- Backend Entry Count
- Backend Cache Percent Full
- Backend Size on Disk
- Backend Cleaner Backlog

PingDirectoryProxy Server Charts

The following charts are displayed on the dashboard for the PingDirectoryProxy Server:

- External Server Total Operations
- External Server Failed Operations
- External Server Health

PingDataSync Server Charts

The following charts are displayed on the dashboard for the PingDataSync Server:

- Sync Pipe Unretrieved Changes
- Sync Pipe Percent Busy
- Sync Pipe Completed Operations Success
- Sync Pipe Completed Operations Failed
- Sync Pipe Completed Operations By Type

PingDataMetrics Server Charts

The following charts are displayed on the dashboard for the PingDataMetrics Server:

- Metrics Queries Per Minute
- Metrics Query Time
- Metrics Query Time Max
- Metrics Query Time Histogram
- Metrics Cache Entry Count
- Metrics Cache Hit Count
- Metrics Cache Miss Count
- Metrics Cache Expired Count
- Metrics Cache Evicted Count
- Metrics Import Delay
- Metrics Load Time
- Metrics DBMS Cluster Time

PingDataGovernance Server Charts

The following charts are displayed on the dashboard for the PingDataGovernance Server:

- Policy Evaluation
- Policy Evaluation Time

Velocity Templates

The PingDataMetrics Server exposes Velocity pages through an HTTP Servlet Extension. If the HTTP Connection Handler is enabled, the Velocity extension is enabled.

```
$ bin/dsconfig set-connection-handler-prop \
  --handler-name "HTTPS Connection Handler" \
  --add http-servlet-extension:Velocity
```

Velocity template files contain presentation content and variables that are replaced when the content is requested. Variables are expressed using a `$` followed by an identifier that refers to an object put into a context (VelocityContext) by the server.

Velocity extensions can be configured to expose a number of objects in the context using the `expose-*` properties:

- **expose-request-attributes** – Indicates whether HTTP request attributes are accessible to templates using the `$subid_request` variable. In general, request attributes are added by server components processing the HTTP request. Also the HTTP request parameters map is available as `$subid_request.parameters`. Request parameters are supplied by the requester, usually in the request URL query string or in the body of the request itself.
- **expose-session-attributes** – Indicates whether HTTP session attributes are accessible to templates using the `$subid_session` variable. Like request attributes, session

attributes are also added by server components processing the HTTP request. The lifetime of these attributes persists until the user's session has ended.

- **expose-server-context** – Indicates whether a Server SDK server context is accessible to templates using the `$ubid_server` variable. The server context provides access to properties and additional information about the server. See the Server SDK documentation for details.

The following are other properties of the Velocity HTTP Servlet Extension:

- **description** – A description of the extension.
- **cross-origin-policy** – Defines a cross origin policy for this extension.
- **base-context-path** – URL base context for the Velocity Servlet.
- **static-content-directory** – In addition to templates, the Velocity Servlet will serve miscellaneous static content related to the templates. By default this is `config/velocity/statics`.
- **require-authentication** – Requires credentials to access Velocity content.
- **identity-mapper** – Maps user credentials to backend entries. If the `require-authentication` property is set, use this property to map bind credentials from a users backend. This is set to `Exact Match` by default. PingDataMetrics Server Velocity template authentication should share the `api-users` LDIF backend used by the REST API. Details are available in the PingDataMetrics Server REST API servlet configuration, and in the [Connection and Security](#) section of the PingDataMetrics Server REST API Reference chapter.
- **static-custom-directory** – If static content is customized, it resides in `velocity/statics` by default.
- **template-directory** – The template directory from which templates are read. By default this is `config/velocity/templates`. This directory also serves as a default for Template Loaders that do not have a template directory specified.
- **static-context-path** – URL path beneath the base context where static content can be accessed.
- **allow-context-override** – Indicates whether context providers may override existing context objects with new values.
- **mime-types-file** – Specifies a file that is used to map file extensions of static content to a Content Type to be returned with requests.
- **default-mime-type** – The default Content Type for HTTP responses. Additional content types are supported by defining one or more additional Velocity Template Loaders.

The VelocityContext object can be further customized by configuring additional Velocity context providers. The dot notation used for context references can be extended to access properties and methods of objects in context using Java Bean semantics. For example, if the HTTP request URL includes a `name` query string parameter like:

```
http://example.com:8080/view/hello?name=Joe
```

An HTML template like the following could be used to generate a page containing a friendly greeting to the requestor:

```
<html>
  <body>
    Hello $subid_request.parameters.name
  </body>
</html>
```

A pop-up window displays a table on the page that lists all variables that are in the Velocity Context. References like `$subid_request` can appear in the template file and be replaced when the template is rendered. This information can be used to check which variables are permitted to be in the template along with the variable values.

Note

For security, all template substitutions are HTML escaped by default. To substitute unescaped content, a variable name ending with "WithHtml" must be used. For example, `$addressWithHtml`, would substitute the contents of the `$addressWithHtml` variable into the page generated from the HTML template without escaping it.

A debug option can be used in any Velocity template for verifying available information in the Velocity Context:

```
parse ("_debug.vm")
debug ()
```

If a variable is added to a template for something that does not exist, the rendered page will contain a literal string of the unfulfilled variable (for example `$undefined_variable`).

By default, the Velocity Servlet Extension expects to access content in subdirectories of the server's `config/velocity` directory:

- **templates** – This directory contains Velocity template files that are used to generate pages in response to client requests.
- **statics** – This directory contains static content such as cascading style sheets, HTML, and Javascript files as well as images and third-party libraries.

Supporting Multiple Content Types

By default, the Velocity Servlet Extension is configured to respond to HTTP requests with a content type `text/html`. Change this request type by setting the default MIME type using `dsconfig`. For example, the following can be used to set the default type to XML:

```
$ bin/dsconfig set-http-servlet-extension-prop \  
  --extension-name Velocity \  
  --set default-mime-type:application/xml
```

HTML requests can be supported as well as clients that seek content in other formats. Create one or more Velocity template loaders to load templates for other content types like XML or JSON.

The ability to serve multiple formats of a document to clients at the same URL is typically called *content negotiation*. HTTP clients indicate the type of content desired using the `Accept` header. A client may use a header like the following to indicate that they prefer content in XML but will fallback to HTML if necessary:

```
Accept: application/xml,text/html;q=0.9
```

The following can be used to create a Velocity template loader for XML content:

```
$ bin/dsconfig create-velocity-template-loader \  
  --extension-name Velocity \  
  --loader-name XML \  
  --set evaluation-order-index:502 \  
  --set mime-type-matcher:application/xml \  
  --set mime-type:application/xml \  
  --set template-suffix:.vm.xml
```

Upon receiving a request, the Velocity Servlet first creates an ordered list of requested media types from most desired to least based on the value of the `Accept` header. Starting from the most desired type, it will then iterate over the defined template loaders according to the `evaluation-order-index` property from lowest value to highest.

A template loader can indicate that it can handle content for requested media type by comparing the requested type to its `mime-type-matcher` property. A loader can be configured to load templates from a specific directory or load template files having a particular suffix. For example, XML templates are expected to be named using a `.vm.xml` suffix. If a loader indicates it handles the requested content type and a template exists for the requested view, the template is loaded and used to generate a response to the client. If no loaders are found

for the requested media type, the next most preferred media type (if any) is tried. If no loaders indicated that they could satisfy the requested view, the client is sent an HTTP 404 (not found) error. If no loaders could provide acceptable media but the requested view exists in some other format, the client is sent an HTTP 406 (not acceptable) error.

In this example, a template file called `hello.vm.xml` can be used to generate a response in XML:

```
<hello name="$ubid_request.parameters.name"/>
```

In this case, the response will contain an HTTP Content-Type header with the value of the `mime-type` property of the Velocity template loader.

Velocity Context Providers

The previous examples use a value supplied as an HTTP request query string parameter to form a response. The templates contain a variable `$ubid_request.parameters.name` that was replaced at runtime with a value from the Velocity Context.

The Velocity Extension can be configured to make some information available in the Velocity Context such as the HTTP request, session, and Server SDK Server Context. Velocity Context Providers provide more flexibility in populating the Velocity Context for template use.

Here are some of the properties of a Velocity Context Provider:

- **enabled** – Indicates whether the provider will contribute content for any requests.
- **object-scope** – Indicates to the provider how often objects contributed to the Velocity Context should be re-initialized. Possible values are: `request`, `session`, or `application`.
- **included-view/excluded-view** – These properties can be used to restrict the views for which a provider contributes content. A view name is the request URL's path to the resource without the Velocity Servlet's context or a leading forward slash. If one or more views are included, the provider will service requests for just the specified views. If one or more views are excluded, the provider will service requests for all but the excluded views.

Velocity Tools Context Provider

Apache's Velocity Tools project is focused on providing utility classes useful in template development. The Velocity Context can be configured by specifying Velocity Tool classes to be

Chapter 5: Configuring Charts and Dashboards

automatically added to the Velocity Context for template development. For more information about the Velocity Tools project, see the Velocity website.

The following command can be used to list the set of Velocity Tools that are included in the Velocity Context for general use by templates:

```
$ bin/dsconfig get-velocity-context-provider-prop \  
  --extension-name Velocity \  
  --provider-name "Velocity Tools" \  
  --property request-tool \  
  --property session-tool \  
  --property application-tool \  
  \
```

Chapter 6: Troubleshooting

There are several ways to troubleshoot issues with data gathering or with the PingDataMetrics Server itself.

Topics include:

[Collect Support Data Tool](#)

[Enable JVM Debugging](#)

[Slow Queries Based on Sample Cache Size](#)

[Insufficient Memory Errors](#)

[Unexpected Query Results](#)

[Conditions for Automatic Server Shutdown](#)

[Problems with SSL Communication](#)

[Installation and Maintenance Issues](#)

Use the collect-support-data tool

PingData servers provide information about their current state and any problems encountered. If a problem occurs, run the `collect-support-data` tool in the `/bin` directory. The tool aggregates all relevant support files into a zip file that can be sent to a support provider for analysis. The tool also runs data collector utilities, such as `jps`, `jstack`, and `jstat` plus other diagnostic tools for the operating system.

The tool may only archive portions of certain log files to conserve space, so that the resulting support archive does not exceed the typical size limits associated with e-mail attachments.

The data collected by the `collect-support-data` tool may vary between systems. The data collected includes the configuration directory, summaries and snippets from the `logs` directory, an LDIF of the monitor and RootDSE entries, and a list of all files in the server root.

Perform the following steps to run this tool:

1. Navigate to the server root directory.
2. Run the `collect-support-data` tool. Include the host, port number, bind DN, and bind password.

```
$ bin/collect-support-data \  
--hostname 100.0.0.1 --port 389 \  
--bindDN "cn=Directory Manager" \  
--bindPassword secret \  
--serverRoot /opt/PingData<server> \  
--pid 1234
```

3. Email the zip file to a support provider.

Slow Queries Based on Sample Cache Size

The `evicted-count` attribute of the sample cache sets the number of entries that have been evicted from the cache due to a lack of space. The cache may not be large enough for the query load placed on the server. Increase the size of the sample cache with the following command, which sets the maximum size to 200000:

```
$ dsconfig set-monitoring-configuration-prop \  
--set sample-cache-max-cached-series:200000
```


Some queries are so infrequent that the cached data expires due to age. The default age is 10 minutes, but this can be increased up to one hour. If the `expired-count` monitor attribute is increasing between queries, consider increasing the idle timeout as follows:

```
$ dsconfig set-monitoring-configuration-prop \
  --set sample-cache-idle-series-timeout:20m
```

Insufficient memory errors

If the server shuts down due to insufficient memory errors, it is possible that the allocated heap size is not enough for the amount of data being returned. Consider increasing the heap size, or reducing the number of request handler threads using the following `dsconfig` command:

```
$ bin/dsconfig set-connection-handler-prop \
  --handler-name "HTTP Connection Handler" \
  --set num-request-handlers:<num-of-threads>
```

Unexpected Query Results

The query API aggregates data samples across servers and dimension values. The samples for different servers, or even different dimension values, are imported into the PingDataMetrics Server at different times. All metric data is imported in time-order for each server. The ordering cannot be set across servers, and samples for a specific time may arrive in stages. Therefore, a metric query that aggregates across servers or dimensions may get partial data when the query time range ends. This problem can be compounded when the monitored servers clocks are not synchronized (samples have the monitored server timestamp). The query looks at a single time range. The more [clock skew](#) between the monitored servers, the higher the probability of the results not being accurate for the range.

With the query API, the data can be pivoted (split) by server and dimension. The API enables formatting the results as an HTML table. The following sequence of API URLs return the last three minutes of data in 10-second increments:

```
http://<metrics-server-host:port>/api/v1/metrics/throughput/datatable?
maxIntervals=30&startTime=-3m&tqx=out:html&tz=US/Central

http://<metrics-server-host:port>/api/v1/metrics/throughput/datatable?
maxIntervals=30&startTime=-3m&tqx=out:html&tz=US/Central&pivot=instance

http://<metrics-server-host:port>/api/v1/metrics/throughput/datatable?
```

```
maxIntervals=30&startTime=-
3m&tqx=out:html&tz=US/Central&pivot=instance&pivot=op-type
```

- The first URL aggregates all servers and LDAP operations into a single number split across time.
- The second URL splits out the data by server and time.
- The third URL splits out the data by server, LDAP operation, and time.

As dimension pivots (splits) are added, the results display more aggregations of partial data.

Conditions for automatic server shutdown

All PingData servers will shutdown in an out of memory condition, a low disk space error state, or for running out of file descriptors. The PingDirectory Server will enter lockdown mode on unrecoverable database environment errors, but can be configured to shutdown instead with this setting:

```
$ dsconfig set-global-configuration-prop \
--set unrecoverable-database-error-mode:initiate-server-shutdown
```

Problems with SSL communication

Enable TLS debugging in the server to troubleshoot SSL communication issues:

```
$ dsconfig create-debug-target \
--publisher-name "File-Based Debug Logger" \
--target-name
com.unboundid.directory.server.extensions.TLSConnectionSecurityProvider \
--set debug-level:verbose \
--set include-throwable-cause:true
```

```
$ dsconfig set-log-publisher-prop \
--publisher-name "File-Based Debug Logger" \
--set enabled:true \
--set default-debug-level:disabled
```

In the `java.properties` file, add `-Djavax.net.debug=ssl` to the `start-ds` line, and run `bin/dsjavaproperties` to make the option take effect on a scheduled server restart.

Installation and maintenance issues

The following are common installation and maintenance issues and possible solutions.

The setup program will not run

If the `setup` tool does not run properly, some of the most common reasons include the following:

A Java Environment Is Not Available – The server requires that Java be installed on the system prior to running the `setup` tool.

If there are multiple instances of Java on the server, run the `setup` tool with an explicitly-defined value for the `JAVA_HOME` environment variable that specifies the path to the Java installation. For example:

```
$ env JAVA_HOME=/ds/java ./setup
```

Another issue may be that the value specified in the provided `JAVA_HOME` environment variable can be overridden by another environment variable. If that occurs, use the following command to override any other environment variables:

```
$ env UNBOUNDID_JAVA_HOME="/ds/java" UNBOUNDID_JAVA_BIN="" ./setup
```

Unexpected Arguments Provided to the JVM – If the `setup` tool attempts to launch the `java` command with an invalid set of arguments, it may prevent the JVM from starting. By default, no special options are provided to the JVM when running `setup`, but this might not be the case if either the `JAVA_ARGS` or `UNBOUNDID_JAVA_ARGS` environment variable is set. If the `setup` tool displays an error message that indicates that the Java environment could not be started with the provided set of arguments, run the following command:

```
$ unset JAVA_ARGS UNBOUNDID_JAVA_ARGS
```

The Server Has Already Been Configured or Started – The `setup` tool is only intended to provide the initial configuration for the server. It will not run if it detects that it has already been run.

A previous installation should be removed before installing a new one. However, if there is nothing of value in the existing installation, the following steps can be used to run the `setup` program:

- Remove the `config/config.ldif` file and replace it with the `config/update/config.ldif.{revision}` file containing the initial configuration.
- If there are any files or subdirectories in the `db` directory, then remove them.
- If a `config/java.properties` file exists, then remove it.

- If a `lib/setup-java-home` script (or `lib\set-java-home.bat` file on Microsoft Windows) exists, then remove it.

The server will not start

If the server does not start, then there are a number of potential causes.

The Server or Other Administrative Tool Is Already Running – Only a single instance of the server can run at any time from the same installation root. Other administrative operations can prevent the server from being started. In such cases, the attempt to start the server should fail with a message like:

```
The <server> could not acquire an exclusive lock on file
/ds/PingData<server>/locks/server.lock:
The exclusive lock requested for file
/ds/PingData<server>/locks/ server.lock
was not granted, which indicates that another
process already holds a shared or exclusive lock on
that file. This generally means that another instance
of this server is already running.
```

If the server is not running (and is not in the process of starting up or shutting down), and there are no other tools running that could prevent the server from being started, it is possible that a previously-held lock was not properly released. Try removing all of the files in the locks directory before attempting to start the server.

There Is Not Enough Memory Available – When the server is started, the JVM attempts to allocate all memory that it has been configured to use. If there is not enough free memory available on the system, the server generates an error message indicating that it could not be started.

There are a number of potential causes for this:

- If the amount of memory in the underlying system has changed, the server might need to be re-configured to use a smaller amount of memory.
- Another process on the system is consuming memory and there is not enough memory to start the server. Either terminate the other process, or reconfigure the server to use a smaller amount of memory.
- The server just shut down and an attempt was made to immediately restart it. If the server is configured to use a significant amount of memory, it can take a few seconds for all of the memory to be released back to the operating system. Run the `vmstat`

command and wait until the amount of free memory stops growing before restarting the server.

- If the system is configured with one or more memory-backed filesystems (such as `/tmp`), determine if any large files are consuming a significant amount of memory. If so, remove them or relocate them to a disk-based filesystem.

An Invalid Java Environment or JVM Option Was Used – If an attempt to start the server fails with 'no valid Java environment could be found,' or 'the Java environment could not be started,' and memory is not the cause, other causes may include the following:

- The Java installation that was previously used to run the server no longer exists. Update the `config/java.properties` file to reference the new Java installation and run the `bin/dsjavaproperties` command to apply that change.
- The Java installation has been updated, and one or more of the options that had worked with the previous Java version no longer work. Re-configure the server to use the previous Java version, and investigate which options should be used with the new installation.
- If an `UNBOUNDID_JAVA_HOME` or `UNBOUNDID_JAVA_BIN` environment variable is set, its value may override the path to the Java installation used to run the server (defined in the `config/java.properties` file). Similarly, if an `UNBOUNDID_JAVA_ARGS` environment variable is set, then its value might override the arguments provided to the JVM. If this is the case, explicitly unset the `UNBOUNDID_JAVA_HOME`, `UNBOUNDID_JAVA_BIN`, and `UNBOUNDID_JAVA_ARGS` environment variables before starting the server.

Any time the `config/java.properties` file is updated, the `bin/dsjavaproperties` tool must be run to apply the new configuration. If a problem with the previous Java configuration prevents the `bin/dsjavaproperties` tool from running properly, remove the `lib/set-java-home` script (or `lib\set-java-home.bat` file on Microsoft Windows) and invoke the `bin/dsjavaproperties` tool with an explicitly-defined path to the Java environment, such as:

```
$ env UNBOUNDID_JAVA_HOME=/ds/java bin/dsjavaproperties
```

An Invalid Command-Line Option was Used – There are a small number of arguments that can be provided when running the `bin/start-server` command. If arguments were provided and are not valid, the server displays an error message. Correct or remove the invalid argument and try to start the server again.

The Server Has an Invalid Configuration – If a change is made to the server configuration using `dsconfig` or the Administrative Console, the server will validate the change before

Chapter 6: Troubleshooting

applying it. However, it is possible that a configuration change can appear to be valid, but does not work as expected when the server is restarted.

In most cases, the server displays (and writes to the error log) a message that explains the problem. If the message does not provide enough information to identify the problem, the `logs/config-audit.log` file provides recent configuration changes, or the `config/archived-configs` directory contains configuration changes not made through a supported configuration interface. The server can be started with the last valid configuration using the `--useLastKnownGoodConfig` option:

```
$ bin/start-server --useLastKnownGoodConfig
```

To determine the set of configuration changes made to the server since the installation, use the `config-diff` tool with the arguments `--sourceLocal --targetLocal --sourceBaseline`.

The `dsconfig --offline` command can be used to make configuration changes.

Proper Permissions are Missing – The server should only be started by the user or role used to initially install the server. However, if the server was initially installed as a non-root user and then started by the root account, the server can no longer be started as a non-root user. Any new files that are created are owned by root.

If the user account used to run the server needs to change, change ownership of all files in the installation to that new user. For example, if the server should be run as the "ds" user in the "other" group, run the following command as root:

```
$ chown -R ds:other /ds/PingData<server>
```

The server has shutdown

Check the current server state by using the `bin/server-state` command. If the server was previously running but is no longer active, potential reasons may include:

- Shut down by an administrator – Unless the server was forcefully terminated, then messages are written to the error and server logs stating the reason.
- Shut down when the underlying system crashed or was rebooted – Run the `uptime` command on the underlying system to determine what was recently started or stopped.
- Process terminated by the underlying operating system – If this happens, a message is written to the system error log.

- Shut down in response to a serious problem – This can occur if the server has detected that the amount of usable disk space is critically low, or if errors have been encountered during processing that left the server without worker threads. Messages are written to the error and server logs (if disk space is available).
- JVM has crashed – If this happens, then the JVM should provide a fatal error log (a `hs_err_pid<processID>.log` file), and potentially a core file.

The server will not accept client connections

Check the current server state by using the `bin/server-state` command. If the server does not appear to be accepting connections from clients, reasons can include the following:

- The server is not running.
- The underlying system on which the server is installed is not running.
- The server is running, but is not reachable as a result of a network or firewall configuration problem. If that is the case, connection attempts should time out rather than be rejected.
- If the server is configured to allow secure communication through SSL or StartTLS, a problem with the key manager or trust manager configuration can cause connection rejections. Messages are written to the server access log for each failed connection attempt.
- The server may have reached its maximum number of allowed connections. Messages should be written to the server access log for each rejected connection attempt.
- If the server is configured to restrict access based on the address of the client, messages should be written to the server access log for each rejected connection attempt.
- If a connection handler encounters a significant error, it can stop listening for new requests. A message should be written to the server error log with information about the problem. Restarting the server can also solve the issue. Another option is to create an LDIF file that disables and then re-enables the connection handler, create the `config/auto-process-ldif` directory if it does not already exist, and then copy the LDIF file into it.

The server is unresponsive

Check the current server state by using the `bin/server-state` command. If the server process is running and appears to be accepting connections but does not respond to requests received on those connections, potential reasons for this include:

Chapter 6: Troubleshooting

- If all worker threads are busy processing other client requests, new requests are forced to wait until a worker thread becomes available. A stack trace can be obtained using the `jstack` command to show the state of the worker threads and the waiting requests.

If all worker threads are processing the same requests for a long time, the server sends an alert that it might be deadlocked. All threads might be tied up processing unindexed searches.

- If a request handler is busy with a client connection, other requests sent through that request handler are forced to wait until it is able to read data. If there is only one request handler, all connections are impacted. Stack traces obtained using the `jstack` command will show that a request handler thread is continuously blocked.
- If the JVM in which the server is running is not properly configured, it can spend too much time performing garbage collection. The effect on the server is similar to that of a network or firewall configuration problem. A stack trace obtained with the `pstack` utility will show that most threads are idle except the one performing garbage collection. It is also likely that a small number of CPUs is 100% busy while all other CPUs are idle. The server will also issue an alert after detecting a long JVM pause that will include details.
- If the JVM in which the server is running has hung, the `pstack` utility should show that one or more threads are blocked and unable to make progress. In such cases, the system CPUs should be mostly idle.
- If there is a network or firewall configuration problem, communication attempts with the server will fail. A network sniffer will show that packets sent to the system are not receiving TCP acknowledgment.
- If the host system is hung or lost power with a graceful shutdown, the server will be unresponsive.

If it appears that the problem is with the server software or the JVM, work with a support provider to diagnose the problem and potential solutions.

Problems with the Administrative Console

If a problem occurs when trying to use the Administrative Console, reasons may include one of the following:

- The web application container that hosts the console is not running. If an error occurs while trying to start it, consult the logs for the web application container.
- If a problem occurs while trying to authenticate, make sure that the target server is online. If it is, the access log may provide information about the authentication failure.

- If a problem occurs while interacting with the server instance using the Administrative Console, the access and error logs for that instance may provide additional information.

Chapter 7: PingDataMetrics Server API Reference

The PingDataMetrics Server REST API can be used to build custom dashboards and other applications for processing and viewing data. The API interface can be accessed using standard tools and charting packages. The PingDataMetrics Server API is also easily accessed from a Web browser.

Topics include:

[Connection and Security](#)

[List Monitored Instances](#)

[Retrieve Monitored Instance](#)

[List Available Metrics](#)

[Retrieve a Metric Definition](#)

[Perform a Metric Query](#)

[Data Set Structure](#)

[Google Chart Tools Datasource Protocol](#)

[Access Alerts](#)

[LDAP SLA](#)

[Pagination](#)

Connection and Security

No sensitive user data is collected by the PingDataMetrics Server and stored in the DBMS. If secure access to the PingDataMetrics Server REST API is required, enable secure HTTPS connections and require authentication. A secure HTTPS Connection Handler and authentication can be enabled using `dsconfig`, if not configured during setup.

Note By default, the PingDataMetrics Server can open up to 20 simultaneous database connections. The HTTP Connection handler that runs the REST API servlet has a default value of 15 connections. If the PingDataMetrics Server receives requests through multiple HTTP Connection Handlers, make sure that the total number of request handlers does not exceed the maximum number of database connections.

When authentication is enabled, the REST API service requires HTTP basic authentication.

Requests are authenticated against entries in the `api-users` LDIF backend, or entries in `cn=Root` DNs, `cn=config`. Root DN users have many privileges by default. To restrict access, authenticate with users in the `api-users` backend instead, to prevent the unnecessary use of more privileged account credentials.

Enable REST API authentication by setting the `require-api-authentication` property of the Metrics HTTP Servlet Extension Configuration object as follows:

```
$ bin/dsconfig set-http-servlet-extension-prop \
  --extension-name "PingDataMetrics Server REST API Servlet" \
  --set require-api-authentication:true
```

Perform the following steps to add a REST API user:

1. Create a file name `api-user1.ldif` containing one or more user entries with no privileges. Below is a sample user entry.

```
dn: cn=app-user1,cn=api-users
changeType: add
objectClass: inetOrgPerson
objectClass: person
objectClass: top
cn: app-user1
uid: app-user1
sn: User1
userpassword: api1
ds-pwp-password-policy-dn: cn=Default Password Policy,cn=Password
Policies,cn=config
```

The password is in clear text. It will be encrypted in the next step.

2. As a privileged user that can add API users, load the entry using `ldapmodify`:

```
$ bin/ldapmodify --filename api-user1.ldif
```

3. Authenticate using either the full DN of the user added (`cn=app-user1,cn=api-users`) or the UID (`app-user1`). The user name to DN map is governed by the `identity-mapper` setting of the Metrics REST HTTP Servlet Extension configuration object.
4. Enable Velocity Template authentication as well:

```
$ bin/dsconfig set-http-servlet-extension-prop \
  --extension-name Velocity \
  --set require-authentication:true
```

Secure Error Messages

When developing an application that uses the PingDataMetrics Server API, error messages should not be delivered from the API directly to a user. Also, the application should not depend on error messages or reason text. These messages may change over time, and their presence may depend on server configuration. Use the HTTP return code and the context of the request to create a client error message that is displayed to the user.

The PingDataMetrics Server API has an `omit-error-message-details` Metrics HTTP Servlet Extension Configuration object, that when enabled, restricts error messages to the typical reason phrase associated with the HTTP return code (such as, 'Not Found' for an HTTP 404 error). This prevents the server from inadvertently revealing information about itself or its data. Set this property as follows:

```
$ bin/dsconfig set-http-servlet-extension-prop \
  --extension-name "PingDataMetrics Server REST API Servlet" \
  --set omit-error-message-details:true
```

Response Codes

The following response codes are available:

Response Code	Description
200 OK	The request was processed successfully and the requested data returned.
400 Bad Request	The request contained an error. Refer to the error message to resolve the issue.
404 Not Found	The requested resource is not found or no samples are collected for the metric.
500 Internal Server Error	An unexpected server error occurred. Refer to the error message for more info.
503 Service Not Available	The metric query service is temporary offline. Refer to the error message for more info.

The following is a sample response:

```
Response Body <?xml version="1.0" encoding="UTF-8"?>
  <errorResponse xmlns="com.unboundid.directory.mon.api.v1.models">
```

```
<errors reason="unknown_data_source_id" message="There are  
no metrics defined with id connections"/>  
</errorResponse>
```

List Monitored Instances

Get a list of all monitored instances along with their current status. The default format is JSON. The servlet will use the HTTP Accept header as a hint if no specific format is specified. Results are filtered using the various `instance` query parameters.

URL	<code>/api/v1/instances</code>
Method	GET
Formats	JSON, XML
Query Parameters	<p><code>instanceHostname</code> – Hostname(s) of the servers from which data is gathered. Multiple values are evaluated as logical ORs.</p> <p><code>instanceLocation</code> – Location(s) of the servers from which data is gathered. Multiple values are evaluated as logical ORs.</p> <p><code>instanceType</code> – Types of server(s) to get data from. Possible values are:</p> <ul style="list-style-type: none">• <code>ds</code>• <code>proxy</code>• <code>sync</code>• <code>metrics-server</code> <p><code>instanceVersion</code> – Version(s) of the servers to get data from. Multiple values are evaluated as logical ORs.</p>

EXAMPLES

All instances in JSON format.

```
curl \  
-X GET \  
https://<metricsHost>:8080/api/v1/instances.json
```

All PingDirectory Server and PingDirectoryProxy Server instances in XML format:

```
curl \  
-X GET \  

```

```
https://<metricsHost>:8080/api/v1/instances.xml?
instanceType=data-store&instanceType=proxy
```

Response Code 200 OK

Response Body

```
{
  "found" : 2
  "offset" : 0,#
  "instances" : [ {
    "type" : "ds",
    "id" : "pingidentity4510",
    "hostname": "pingidentity5200.example.com",
    "displayName" : "pingidentity4510",
    "version": "Directory Server 7.3.0",
    "operatingSystem": "Linux",
    "status": {
      "state": "ONLINE"
    }
  }, {
    "type" : "ds",
    "id" : "unboundid3500",
    "hostname": "unboundid3500.example.com",
    "displayName" : "directory3500",
    "version": "UnboundID Directory Server 3.5.0.0",
    "operatingSystem": "Linux",
    "status": {
      "state": "DEGRADED",
      "unavailableAlerts": [
        "replication-backlogged"
      ]
    }
  }
] }
```

Retrieve Monitored Instance

Get a specific monitored instance along with its status. The default format is JSON. The servlet will use the HTTP Accept header as a hint if no specific format is specified.

URL /api/v1/instances/{instance}{.format}

Method GET

Formats JSON, XML

Query Parameters N/A

Server State

The PingDataMetrics Server returns the server state status of the monitored instance, which is displayed by the `status` parameter. The `status` parameter can have one of the following values:

OFFLINE – Server cannot be contacted.

STARTING_UP – Server is starting.

ONLINE – Server is available.

DEAD_LOCKED – Server is deadlocked and not able process more operations.

UNAVAILABLE – Server is unavailable, but not offline. The server may be in lock-down mode, but may be online.

DEGRADED – Server is available but is incapable of providing services.

CONNECTION_ERROR – Server could not connect or has lost connection to the host.

EXAMPLE:

Instance with ID `metrics-server` in JSON format.

```
curl \
-X GET \
https://<metricsServerHost>:8080/api/v1/instances/metrics-server.json
```

Response Code 200 OK

Response Body

```
{
  "displayName": "metrics-server",
  "hostname": "metrics-server.example.com",
  "id" : "metrics-server",
  "operatingSystem": "Linux",
  "status" : {
    "state" : "ONLINE"
  },
  "type" : "metrics-server",
  "version": "PingData PingDataMetrics Server 7.3"
}
```

List Available Metrics

Get a list of metric definitions with their units, dimensions, names, and other values. The default format is JSON. The servlet will use the HTTP Accept header if no specific format is specified.

URL /api/v1/metrics{.format}

Method GET

Formats JSON, XML

Query Parameters `name` – Limits the results to metrics whose names contain a matching substring. The search is not case-sensitive.

`type` – Limits the results to the metrics of the specified type. Possible values are:

- `discreteValued`
- `continuousValued`
- `count`

`group` – Limits the results to the metrics within the specified group.

Possible values are:

- `PingDirectory Server Backend`
- `Monitoring Data Cache`
- `Java Virtual Machine`
- `LDAP`
- `Entry Balancing`
- `PingDirectory Server Entry Cache`
- `External Server`
- `Host System`
- `Metric Query`
- `Monitoring DBMS`
- `Monitoring Data Processing`
- `Replication`
- `Sync Pipe`

`instanceType` – Limits the result to metrics that uses the specified instance types as sources. Possible values are:

- `ds`
- `proxy`
- `sync`
- `metrics-server`

`statistic` – Limits the results to metrics that provides the specified statistics. Possible values are:

- `count`
- `average`
- `maximum`

- minimum
- histogram

EXAMPLES

All metrics in JSON format.

```
curl \
  -X GET \
  https://<metricsServerHost>:8080/api/v1/metrics.json
```

All count type metrics in the "PingDirectory Server Backend" group providing either count or average statistics:

```
curl \
  -X GET \
  https://<metricsServerHost>:8080/api/v1/metrics.json?type=count&group=ds
%20backend&statistic=count&statistic=average
```

Note

Spaces in parameter values may be encoded as %20 or t.

Response Code 200 OK

Response Body

```
{
  "found": 7,
  "metrics": [
    {
      "countUnit": {
        "abbreviatedName": "Chkpt",
        "pluralName": "Checkpoints",
        "singularName": "Checkpoint"
      },
      "description": "Number of database checkpoints
        performed by the backend",
      "dimensions": [
        {
          "id": "backend",
          "values": [
            "userroot"
          ]
        }
      ],
      "group": "Directory Server Backend",
      "id": "backend-checkpoints",
      "instanceTypes": [
        "ds"
      ],
      "name": "Backend Checkpoints",
      "shortName": "Checkpoints",
      "statistics": [
        "count"
      ]
    }
  ]
}
```

```

    ],
    "type": "count"
  },
  ...

```

Retrieve a Metric Definition

Get a specific metric definition. The default format will be JSON if none is specified. The servlet will use the HTTP Accept header as a hint if no specific format is specified.

URL	/api/v1/metrics/{metricId}.{format}
Method	GET
Formats	JSON, XML
Query	Parameters N/A

EXAMPLE

Metric with ID backend-sequential-writes in XML format.

```

curl \
  -X GET \
  https://<metricsServerHost>:8080/api/v1/metrics/backend-sequential-
writes.xml

```

Response Code 200 OK

Response Body

```

<?xml version="1.0" encoding="utf-8" standalone="yes"?>
<countMetric xmlns="com.unboundid.directory.mon.api.v1"
  id="backend-sequential-writes" name="Sequential Disk Writes"
  shortName="Sequential Writes" group="Directory Server Backend">
  <description>Number of Sequential I/O Disk writes
    made by backend</description>
  <instanceTypes>
    <instanceType>ds</instanceType>
  </instanceTypes>
  <statistics>
    <statistic>count</statistic>
  </statistics>
  <dimensions>
    <dimension id="backend">
      <values>
        <value>userroot</value>
      </values>
    </dimension>
  </dimensions>
  <countUnit singularName="Sequential Write"
    pluralName="Sequential Writes" abbreviatedName="Seq Wr" />
</countMetric>

```

Perform a Metric Query

A metric query returns the collected sample data from the various monitored instances. The data returned can be presented many ways, depending on client requirements.

Common Query Parameters

`instanceType` – Type(s) of instances to get data from. Possible values are:

- `ds`
- `proxy`
- `sync`
- `metrics-server`

`instanceLocation` – Location(s) of the instances from which data is collected.

`instanceHostname` – Names of the machines hosting the instances.

`instanceVersion` – Version(s) of the instances providing the data.

`instance` (multi-valued) – ID(s) of the instances from which data is collected. The instance ID is the `cn` of the external server and the same name as listed by the `status` command.

`startTime` – Include samples on or after the specified time. The time is either an absolute time in ISO 8601 format (such as 2016-08-13T19:36:00Z) or a time relative to the `endTime` (such as -5m or -4h). By default, the start time is -5m.

`endTime` – Include samples on or before this time. The end time is either an absolute time in ISO 8601 format or a time relative to now (such as -5m or -4h). The default end time is now. Offset time values are relative to the current system clock time on the PingDataMetrics Server.

`maxIntervals` – The number of separate intervals, between the start and end times, returned. This is considered the “resolution” of the data over time. By default, the maximum number of intervals is 1, which means all samples collected between the start and end times will be

aggregated into one result according to the statistic selected.

`statistic` – Retrieve and apply this statistic to the data. Default for count based metrics is count and average for other metric types.

Possible values are:

- count
- average
- minimum
- maximum
- histogram

`dimension` – Include only these dimension values. A colon separates the dimension name and values, which are separated by commas (for example, `op-type:add,delete`).

`pivot` – Pivot by these dimensions. A pivot keeps the data separated along different dimensional values. The value "instance" may be used to keep the data separate between different instances. For metrics that have the histogram statistic, the histogram pivot may also be used to keep the values of each histogram bucket separate.

`tz` – Specifies the timezone to be used when displaying dates. By default, it is GMT. The timezone is specified in Java Time Zone format, so "US/Central" is CST in the United States.

Sub-parameters for the count and average statistics

Both the count and average statistics of count type metrics may have a rate scale applied to occurrences over a period of time using the per sub-parameter. The valid rate scaling values are:

- s OR second
- m OR minute
- h OR hour

Sub-parameters for the histogram statistic

The histogram statistic includes all buckets and keeps the raw value for each bucket. Graphs can be configured to show the percentage of all operations above a given threshold, such as 50 ms. These graphs are useful for looking at a small percentage of operations in a given category. If the value falls between histogram bucket boundaries, the

buckets where it falls will be included in the data. The possible values are:

- `min` - Includes in the calculation only the histogram data above the given threshold.
- `max` - Provides an upper bound on the histogram value
- `percent` - Allows the histogram values to be reported as a percentage of the overall values. Instead of returning raw counts, the value is a fraction of the total. This percentage is calculated within a pivot.

If both `min` and `max` are specified, the returned value is the sum of all buckets between and including `min` and `max`.

Data Set Structure

The data set structure is a proprietary data structure that is space-optimized and designed to work with charting libraries like Highcharts or FusionCharts. The default format is JSON. The servlet will use the HTTP Accept header as a hint if no format is specified.

URL	<code>/api/v1/metrics/{metricId}/dataset{.format}</code>
Method	GET
Formats	JSON, XML

Note

All of the Common Query parameters apply to this resource.

Get the average response time metric for add and delete operations from 7/7/2015 for all PingDirectory Server and PingDirectoryProxy Server in Austin and Houston:

```
curl \
-X GET \
https://<metricsServerHost>:8080/api/v1/metrics/response-time/dataset?
instanceType=directory-server
&instanceType=proxy&instanceLocation=austin&instanceLocation=houston&startTi
me=-1d
&endTime=2015-07-07&pivot=instance&dimension=op-type:add,delete
```

Get the new connections metric and scale the value per hour in the last 5 minutes:

```
curl \
-X GET \
https://<metricsServerHost>:8080/api/v1/metrics/new-connections/dataset?
statistic=count;per:hour
```

Get the percentage of all occurrences in the last hour where the response-time metric has a value above 50ms:

```
curl \
  -X GET \
  https://<metricsserverHost>:8080/api/v1/metrics/response-time/dataset?
  statistic=histogram;min:50;percent&startTime=-1h
```

Response Code 200 OK

Response Body When one time interval is requested, a category dataset is returned where the first pivoted dimension values are listed as categories and each data point corresponds to a category. Subsequent pivots and histogram buckets are included as a series and subseries. This example is the result of two pivots, op-type and instance:

```
{
  "type" : "category",
  "firstSampleTime" : 1344090300000,
  "lastSampleTime" : 1344090600000,
  "metric" : {
    "type" : "discreteValued",
    "id" : "response-time",
    "name" : "Response Time",
    "shortName" : "Response Time",
    "description" : "Time for server to process an LDAP
      operation and send a response to the client",
    "group" : "LDAP",
    "instanceTypes" : [ "ds", "proxy" ],
    "statistics" : [ "average", "count", "histogram" ],
    "dimensions" : [ {
      "id" : "application-name"
    }, {
      "id" : "op-type",
      "values" : [ "Search", "ModifyDN", "Add", "Delete",
        "Compare", "Bind", "Modify" ]
    } ],
    "countUnit" : {
      "singularName" : "Operation Response Time",
      "pluralName" : "Operation Response Time",
      "abbreviatedName" : "Response Time"
    },
    "valueUnit" : {
      "singularName" : "Millisecond",
      "pluralName" : "Milliseconds",
      "abbreviatedName" : "Msec"
    }
  },
  "series" : [ {
    "label" : "unboundid35",
```

```

    "data" : [ "0", "0", "0", "0", "0", "0", "0" ]
  }, {
    "label" : "unboundid3",
    "data" : [ "0", "0", "0", "0", "0", "0", "0" ]
  } ],
  "label" : "op-type",
  "categories" : [ "Search", "Delete", "Bind", "Modify",
    "Add", "ModifyDN", "Compare" ]
}

```

Google Chart Tools Datasource Protocol

Metrics data can be presented with Google's Chart Tools Datasource protocol. The Google Visualization API query language (the `tq` request parameter) is not supported. The PingDataMetrics Server supports JSON, HTML, CSV, and TSV data formats as outlined by the Datasource protocol.

URL `/api/v1/metrics/{metricId}/datatable`

Method GET

Formats JSON, HTML, CSV, and TSV

Query Parameters `tqx=out:html` – HTML formatted output.

`tqx=out:csv` – CSV formatted output.

`tqx=out:tsv-excel` – TSV formatted output.

`tz` – Specifies the timezone to be used when displaying dates. The Google Visualization API assumes that the times returned are in local time. The PingDataMetrics Server stores and returns all timestamps in GMT. This parameter specifies how the PingDataMetrics Server presents the time. Usually, the client will pass the user's local timezone in IANA Time Zone Database format, such as "US/Central."

Note

All Common Query parameters apply to this resource.

The following example gets the average response time metric for the last 5 minutes with 30 second ($5 * 60 / 10$) resolution and pivoted by op-type and then instance in CSV format:

```

curl \
  -X GET \
  https://<metricsServerHost>:8080/api/v1/metrics/response-time/datatable?
tqx=out:csv&maxIntervals=10
&pivot=op-type&pivot=instance&tz=US/Central

```

Response Code 200 OK

Response Body When only one time interval is requested, the first pivoted dimension values form the first column. For queries that request more than one time interval, the start of each time interval forms the first column. Combinations of subsequent pivoted dimension values and/or histogram buckets are included as additional columns. All date and time values are under the GMT time zone.

```
"Time","server35 AVERAGE Milliseconds","server3 AVERAGE Milliseconds"
"2012-08-04T14:38:00Z","0","0"
"2012-08-04T14:39:00Z","0","0"
"2012-08-04T14:40:00Z","0","0"
"2012-08-04T14:41:00Z","0","0"
"2012-08-04T14:42:00Z","0","0"
```

The following sample illustrates using Google chart tools:

```
<html>
  <head>
    <!--Load the AJAX API-->
    <script type="text/javascript"
src="https://www.google.com/jsapi"></script>
    <script type="text/javascript">

      // Load the Visualization API and the line chart package.
      google.load('visualization', '1.0', {'packages':['corechart']});
      // Set a callback to run when the Google Visualization API is loaded.
      google.setOnLoadCallback(drawChart);

      function drawChart() {
        var query = new google.visualization.Query
          ('https://<MetricsHost>:8080/
          api/v1/metrics/response-time/datatable?maxIntervals=10
          &pivot=optype&pivot=instance');
        query.send(handleQueryResponse);
      }
      function handleQueryResponse(response) {
        if (response.isError()) {
          alert('Error in query: ' + response.getMessage() + ' '
            + response.getDetailedMessage());
          return;
        }
        var data = response.getDataTable();

        var visualization = new
          google.visualization.LineChart(document.getElementById('chart_div'));
        visualization.draw(data, null);
      }
    </script>
  </head>
```



```
<body>
  <!--Div that will hold the chart-->
  <div id="chart_div"></div>
</body>
</html>
```

Access Alerts

The eventTypes and event APIs can be used to retrieve information and alerts from monitored servers. The eventTypes API provides the range of alert types that have occurred. The events API provides detail about individual alerts.

Retrieving Event Types

The range of alerts that have been generated by monitored servers can be retrieved, with optional filtering, based on the following API definition.

URL /api/v1/eventTypes/[?query-parameters] - gets a list of event types
Method GET
Formats JSON, XML
Query Parameters instance, instanceType, startTime, and endTime. See [Performing a Metric Query](#) for a description of each parameter.

Response Code 200 OK

Response Body ["health-check-available-to-degraded", "health-check-degraded-toavailable"]

Retrieving Events

The detailed information for one or more events can be retrieved, with optional filtering, based on the following API definition.

URL /api/v1/events/[?query-parameters] - gets a list of events
/api/v1/events/{eventId} - gets a single event
Method GET
Formats JSON, XML
Query Parameters type - Limits the result to include only events of the specified types. See the HTML API Reference for event types.
severity - Limits the result to include only events that have the matching severity. Valid "severity" values are: INFO, WARNING,

ERROR, and FATAL.

instance, instanceType, instanceLocation, instanceHostname, instanceVersion, startTime, and endTime. See [Performing a Metric Query](#) for a description of each parameter.

limit, offset. See [Pagination](#) for a description of each parameter.

Response Code
Response Body

200 OK

```
{
  "found" : 2,
  "offset" : 0,
  "events" : [

    { "id": "9bdfd1b8-3811-4a84-b779-93553ff35f83",
      "creationDate": 1351274815559,
      "eventType": "server-starting",
      "eventSeverity": "INFO",
      "sourceProductInstance": "lockdown-test",
      "summary": "Server Starting",
      "detail": "The Directory Server is starting" },
    { "id": "9bdfd1b8-3811-4a84-b779-93553ff35f83",
      "creationDate": 1351274815559,
      "eventType": "server-starting",
      "eventSeverity": "INFO",
      "sourceProductInstance": "directory-3",
      "summary": "Server Starting",
      "detail": "The Directory Server is starting" }
  ]
}
```

LDAP SLA

The LDAP SLA API lists the LDAP SLA objects (configuration data) and queries any single LDAP SLA object. The query of an LDAP SLA object results in the aggregated LDAP SLA configuration, scalar data containing current values for the LDAP SLA, and time-series data. Current data comes from the Threshold object. Historical data comes from a metric query. Historical data is more expensive to fetch and is only included if the client requests it. This allows an LDAP SLA query to get the configuration and current data very efficiently for clients that only need the current data. A client that needs both current and historical data can include the appropriate query parameter and get all the data in a single call.

Retrieving the SLA Object

List the LDAP SLA objects (configuration data) and query any single LDAP SLA object. The default format will be JSON if none is specified. The servlet will use the HTTP Accept header as a hint if no specific format is specified.

URL `/api/v1/sla/ldap` – Returns a list of all LDAP SLA configuration objects in name-order. This includes current values and status as held by the Threshold objects, but will only include any historical data.

`/api/v1/sla/ldap/{sla-name}` – Returns a single LDAP SLA configuration object plus optional historical data.

Method GET

Formats JSON, XML

Query Parameters **For the 1st URL:**

`instance` – Returns LDAP SLA's that reference the specified instance.

`application-name` – Returns LDAP SLA's that reference this application name.

`ldap-op` – Returns LDAP SLA's that reference this LDAP operation.

For the 2nd URL:

`historical` (multi-valued, optional):

- `time` – Includes time series data.
- `limits` – Includes the percent of time thresholds limits that have been exceeded. Requires Threshold.
- `alerts` – Includes all Threshold alerts. Requires Thresholding.
- `histogram` – includes response-time histogram as column data)
- `nines` – Includes response time values that correlate to 99%, 99.9%, 99.99%, and 99.999% response-time measurements)

`startTime` – (optional). The time at which the historical data starts. The default is 1hr.

`endTime` – (optional). The time at which the historical data ends. The default is 5m.

`pivot` – (optional). Historical time-series pivots by this dimension.

- instance – pivot by producing server.
- ldap-op – pivot by LDAP operation.
- histogram – pivot response-time series by histogram buckets.

maxIntervals – (optional). Number of points to include in the historical time series. The default is 100.

EXAMPLE

Retrieving an SLA object.

```
curl -X GET http://x3550-09:8080/api/v1/sla/ldap/Acme+Identity+Portal?historical=time
&historical=nines&pivot=instance&startTime=-15m
```

Response Code 200 OK

Response Body

(JSON, sample is abbreviated)

```
{
  "name":"Acme Identity Portal",
  "applicationName":"Application 5",
  "ldapOps":["search"],
  "servers":["x2270-08.pingidentity.lab:1389"],
  "enabled":true,
  "responseTimeState":"NORMAL",
  "throughputState":"normal",
  "currentResponseTime":6.002752,
  "currentThroughput":7032.794,
  "averageResponseTime":6.212055,
  "averageThroughput":5517.1323,
  "responseTimeWarnLimit":8.0,
  "responseTimeCriticalLimit":10.0,
  "throughputWarnLimit":8000.0,
  "throughputCriticalLimit":10000.0,
  "responseTimeSeries":{
    "type":"timeInterval",
    "firstSampleTime":1359045070000,
    "lastSampleTime":1359045970000,
    "rateScaling":"NONE",
    "statistic":"AVERAGE",
    "metric":{
      "type":"discreteValued",
      "id":"response-time",
      "name":"Response Time",
      "shortName":"Response Time",
      "description":"Time for server to process
an LDAP operation and send a response to the client.",
      "group":"LDAP",
      "instanceTypes":["directory-server","proxy"],
      "statistics":["average","count","histogram"],
      "dimensions":[{"id":"application-name",
```

```
"values":["unidentified directory application",
"unidentified proxy application","application 9",
"application 5","root user","admin user",
"application 6"]},{ "id":"op-type","values"
  ["search","modifydn","add","delete","compare",
  "bind","modify"]}],
"countUnit":{"singularName":"Operation Response Time",
"pluralName":"Operation Response Time",
"abbreviatedName":"Response Time"},
"valueUnit":{"singularName":"Millisecond",
"pluralName":"Milliseconds","abbreviatedName":"Msec"}
},
...
```

Pagination

Pagination is supported for both the metrics and instances listing URLs.

Query Parameters

`limit` – Specifies the maximum number of results to return. The default is to return all results.

`offset` – Specifies how many results to skip for the first results to return.

Response Parameters

`found` – The number of results that satisfied the query parameters.

`offset` – The index into the total result set where the current response begins.

Index

A

- administrative account
 - adding a root user account 26
- Administrative Console
 - URL 17
- administrative password 26
- aggregating data 94
- alarms 37
 - testing setup 39
- alerts
 - alarm_cleared alert type 38
 - configure alert handlers 36
 - list of system alerts 35, 38
 - notifications and alerts 35
 - overview 38
 - testing setup 39
- alerts backend
 - alert retention time 37
 - duplicate alert suppression 37
 - overview 36
 - view information 36
- API
 - access alerts 142

- add a REST user 128
 - connection and security 128
 - data set structure 138
 - LDAP SLA 143
 - list available metrics 132
 - list monitored instances 130
 - pagination 146
 - perform a metrics query 136
 - response codes 129
 - retrieve a metric definition 135
 - retrieve monitored instance 131
- authentication
- server authentication with a SASL External Certificate 76
- ## B
- backend monitors
 - disk space usage 34
 - entries 32
 - backup command 43
 - base64 command 43
- ## C
- chart builder tool
 - overview 105
 - charts
 - available server charts 108
 - chart builder parameters 107

- chart builder tool 105
 - chart properties file 108
 - presentation details 106
 - cn=monitor backend 89
 - collect-support-data tool 9, 43, 117
 - command-line
 - available tools 42
 - default properties file 46
 - tools.properties file 44
 - config-diff tool 43
 - Config File Handler Backend 54
 - create-rc-script command 43
 - D**
 - dashboards
 - available dashboards 101
 - configure dashboards 103
 - debug template files 104
 - metrics landing page 16
 - save custom files 104
 - data collection
 - aggregating data 94
 - importing data 93
 - overview 3
 - performance data 87
 - reduce data collected 92
 - reduce frequency of collection 92
 - reduce sample blocks 92
 - system monitoring data 89
 - Data Governance charts 110
 - Data Metrics Server
 - components 2
 - overview 2
 - start server 20
 - stop server 20
 - Data Sync Server charts 109
 - demo-dashboard 102
 - Directory Proxy Server charts 109
 - Directory Server charts 109
 - disk space usage monitor 34
 - DNS caching 75
 - dsconfig command 43
 - dsframework command 43
 - dsjavaproperties command 43
 - E**
 - error log publisher 29
 - external collector daemon 90
 - G**
 - gauges 37
 - testing related alarms and alerts 39
 - Global Configuration object 54
-

H

host system monitor provider 90
HTTP Correlation ID 49
HTTP Servlet Extension object 128-129

I

inter-server-certificate property 59
IO scheduler 11
IP address reverse name lookup 75

J

Java
 installing the JDK 7
JVM debugging
 during setup 120
 invalid options 122

L

landing page 16
ldap-dashboard 101
ldapmodify command 43
ldappasswordmodify command 43
ldapsearch command 43
ldif-diff command 43
ldifmodify command 44
license key 14
Linux configuration
 filesystem swapping 10

filesystem variables 8
install dstat 10
install sysstat and pstack 9
set file descriptor limit 8
set filesystem flushes 9

load balancers 76

logging

 encrypt log files 31

logs

 create log publisher 29
 error log publisher 29
 overview 28
 retention policies 28
 rotation policies 28

M

manage-certificates tool 57
manage-extension command 44
manage-extension tool 78
memory errors 118
metric-engine-schema command 44
metrics
 continuous metrics 81
 count metrics 81
 dimensions 82
 discrete metrics 82

- list available metrics 132
 - overview of metric types 81
 - performance impact 92
 - query overview 84
 - Metrics Server charts 110
 - monitored-servers command 44
 - remove a server 20
 - monitored-servers tool 19
 - monitored servers
 - add servers 19
 - configure servers to monitor 12
 - monitored-servers tool 19
 - processing time histogram 13
 - stats collector 13
 - tracked applications 13
 - monitoring entries 32
 - N**
 - non-root user 7
 - normalized records 88
 - P**
 - performance data
 - overview 3
 - performance data fields 88
 - pivots 85
 - PostgreSQL
 - backup database 40
 - data storage 41
 - install 15
 - plan the backup 41
 - restore the backup 42
 - start the backup 42
 - processing time histogram plugin 13
 - pstack utility 125
 - Q**
 - query-metric command 44
 - access metrics 86
 - query data
 - aggregate query results 85
 - pivots 85
 - select query data 84
 - unexpected results 118
 - query overview 84
 - queryrate command 44
 - R**
 - REST API
 - overview 2
 - restore command 44
 - revert-update command 44
 - review-license command 44
-

root user DN 17

S

sample-flush-interval property 12

self-signed certificate

 replacing 57

server-state command 44

server clock skew 91, 118

server install 14

server SDK

 extension types 78

server shutdown 119

server status 131

service level agreements 143

 monitoring overview 95

 SLA dashboard 102

 SLA object 95

 spike monitoring threshold 96

setup command 44

 troubleshooting 120

setup tool 120

sla-viewer-details dashboard 102

sla-viewer dashboard 102

SLA object 95

 configure object 98

SSL certificate 56

start-server command 44

start Data Metrics Server 20

stats collector plugin 13, 89, 93

 cn=monitor backend 89

status command 44

stop-server command 44

stop Data Metrics Server 21

sum-file-sizes command 44

system data

 overview 3

system entropy 10

system utilization monitors 90

T

tools.property file 44

topology

 force master setting 54

 inter-server-certificate property 59

 master selection 53

 monitor data 55

 overview 52

 replace self-signed certificate 59

 server configuration settings 54

 subtree polling interval 53

 update servers 22

 update SSL certificate 56

topology configuration

X

 update SSL Certificate 56

X-Forwarded values 76

tracked applications 13

troubleshooting

 client connections 124

 collect support data 117

 console 125

 installation 119

 memory errors 118

 server shutdown 119, 123

 server unresponsive 124

 slow queries 117

 SSL 119

 unexpected query results 118

U

uninstall command 44

uninstall server 21

update command 44

update tool 22

V

Velocity templates

 multiple content types 113

 overview 110

 save custom files 104

 tools context provider 114