



Hewlett Packard
Enterprise

Operations Orchestration

Software Version: 10.70

Windows and Linux Operating Systems

Tuning Guide

Document Release Date: November 2016

Software Release Date: November 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

(missing or bad snippet)

Copyright Notice

© 2005-2016 Hewlett Packard Enterprise Development LP

Trademark Notices

(missing or bad snippet)(missing or bad snippet)

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: <https://softwaresupport.hpe.com/>.

This site requires that you register for an HP Passport and to sign in. To register for an HP Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

(missing or bad snippet)

Contents

Introduction	4
Minimum Requirements	4
Parameters that May Affect your System Performance	4
Increasing the Number of Worker Threads	6
Increasing the JVM Heap Size	7
Purging the Database	8
Purging Flows	8
Purging APIs	9
Adjusting the Number of Database Connections	10
Analyzing Whether You Need to Adjust the Database Connections	10
Changing the Number of Database Connections	10
Configuring the Amount of Data Written to the Database	12
Scaling Out	13

Introduction

HPE OO is a generic platform that can be installed in a variety of environments and used in various use cases. Therefore, the platform performance and experience may be influenced by many parameters.

This document aims to help the administrator to understand the different parameters that can impact system performance and to provide a set of tools for tuning the system in cases of less than optimal performance.

Note: The appropriate tuning depends on how you use the system. If you change the parameters described in this document, you will need to monitor your system performance and re-tune if required.

Minimum Requirements

The minimal requirements for HPE OO are described in the *HPE OO System Requirements* document.

The requirements for database size are described in the *HPE OO Database Guide*.

Parameters that May Affect your System Performance

This document aims to help the administrator in the process of tuning the system.

If you face a performance issue, you need to identify the cause or causes. This section lists the different parameters that might be impacting your system's performance.

It is recommended to read the list below, and to consider whether the various parameters apply to your use case. If so, click each relevant link to see more information about how to tune that parameter.

Performance may be affected by the following parameters:

- **Load** - A heavy load can exhaust the available resources (threads). This may be caused by running flows with a large number of parallel or multi-instance lanes, or by triggering a large number of flows simultaneously.

In this case, the solution is to increase the number of threads. See ["Increasing the Number of Worker Threads" on the next page](#).

- **Memory consumption** - Your performance may be low because the JVM heap size is not appropriate and garbage collection is slowing down your system.

It is recommended to analyze the time and frequency of garbage collection. You may need to adjust the initial and maximum size of the Central/RAS heap so that it is in accordance with your memory needs.

See ["Increasing the JVM Heap Size" on page 7](#).

- **Database size** - HPE OO is database-intensive, and this may cause your database to become very large, and this will slow down performance.

To keep your database is running efficiently, you need to purge it regularly to keep the size down. See ["Purging the Database" on page 8](#).

If you have other issues with your database, contact your database administrator or see the *OO Database Guide*.

- **Run Log persistence level** - Your database may also be getting too large because of the run history that is persisted to the Run Log.

You can reduce the information that is saved by adjusting the persistence level in Central. See ["Configuring the Amount of Data Written to the Database" on page 12](#).

- **Number of database connections** - Your system may be running slowly because either the Central server or the database server is limiting the number of concurrent connections.

For more information about how to check whether the number of database connections needs tuning, and how to adjust this, see ["Adjusting the Number of Database Connections" on page 10](#).

- **Number of Centrals** - If you have tried all the methods listed above and are still having performance issues, you may need to scale out, by installing additional Central servers or by adding more workers. Our recommendation is to add Central servers. See ["Scaling Out " on page 13](#).

Increasing the Number of Worker Threads

By default, each HPE OO node has 20 worker threads. If your flows have a large number of parallel or multi-instance lanes, or if you trigger a large number of flows simultaneously, we recommend increasing this number. For example, you might increase this number to 200 threads per worker or Central.

Note: The number of threads that can be configured is dependent on the amount of memory available to the Central or worker.

Increasing the Number of Worker Threads in Central or RAS

1. Open the **central-wrapper.conf** or **ras-wrapper.conf** file (located under **<installation_folder>/central/conf** and **<installation_folder>/ras/conf**, respectively) in a text editor.
2. To configure the number of execution threads, edit the property
`-Dcloudslang.worker.numberOfExecutionThreads`
The default value is 20.
3. To configure the size of the incoming buffer, edit the property
`-Dcloudslang.worker.inBufferCapacity`
The default value is 200.
4. Restart the configured node.

These are newly-supported properties. If this is the first time that you have configured them, you will need to add them manually as follows:

```
wrapper.java.additional.<next available number>=  
-Dcloudslang.worker.numberOfExecutionThreads=<new value>  
wrapper.java.additional.<next available number>=  
-Dcloudslang.worker.inBufferCapacity=<new value>
```

Increasing the JVM Heap Size

You can adjust the initial and maximum size of the Central/RAS heap, so that it is in accordance with your memory needs and garbage collection is faster.

1. Open the **central-wrapper.conf** and **ras-wrapper.conf** files (located under **<installation_folder>/<central or ras>/conf/**).
2. Edit the following properties:

 `wrapper.java.initmemory=<value in MB>`

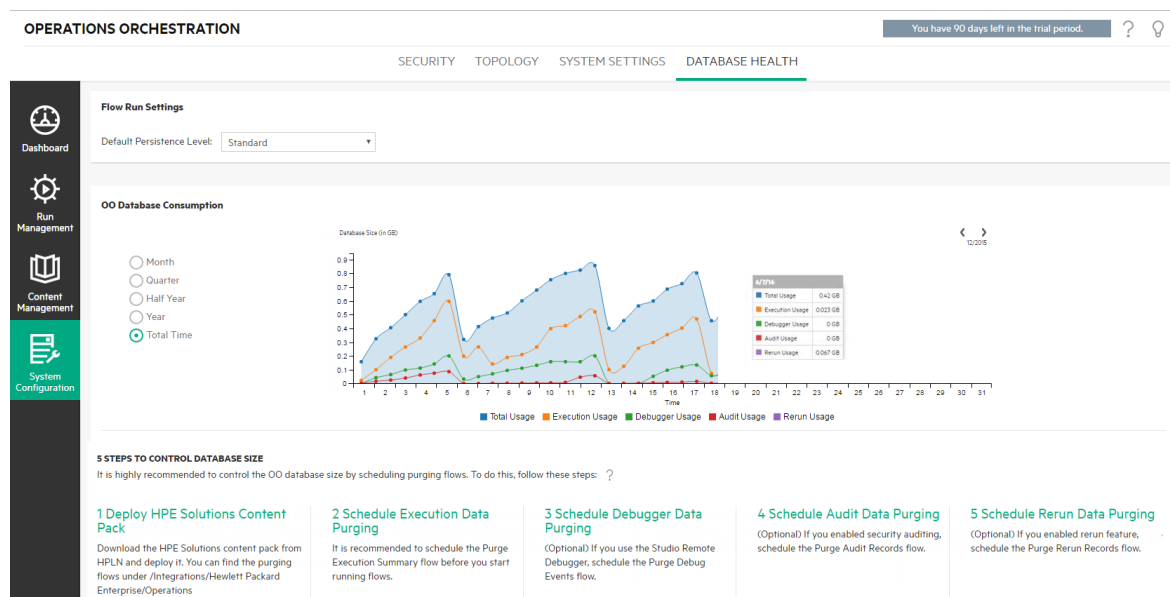
 `wrapper.java.maxmemory=<value in MB>`
3. Restart the configured node.

Purging the Database

By default, HPE OO 10.x saves all flow and step execution data in the database, in order to support debugging of flow runs. As a result, the database size will grow according to system throughput and flow complexity.

It is highly recommended to purge the database regularly, in order to control its size and enhance performance.

You can monitor the size of the database in Central, under the **Database Health** tab in the System Configuration workspace.



Purging Flows

The purging flows are available in the **HPE Solutions** content pack (available on HPE Live Network at <https://hpln.hpe.com/product/operations-orchestration/content>). It is recommended to deploy this content pack, configure the flows with your required settings, and schedule them in Central.

The following purging flows are located under **Library > Integrations > Hewlett-Packard Enterprise > Operations Orchestration > 10.x > Database**:

- **Purge Execution Summary** - Purges run data.

It is recommended to schedule this flow as soon as you start to run flows.

- **Purge Debug Events** - Purges Studio Remote Debugger event data.

If you use the Studio Remote Debugger, schedule this flow.

- **Purge Audit Records** - Purges old auditing records if auditing is enabled.

If you have enabled security auditing, schedule this flow.

- **Purge Rerun Info** - Purges rerun data.

If you have defined rerun points in your flows, schedule this flow.

For more information about these purging flows, see the flow descriptions in Central.

Purging APIs

As an alternative to using the purging flows, you can perform purging functions via API.

The following APIs are available:

- **DELETE /audit/records**: Purges old auditing records if auditing is enabled.
- **DELETE /debugger-events**: Purges Studio Remote Debugger event data.
- **DELETE /executions**: Purges run data such as bound inputs, outputs and step log events. This run data needs to be purged regularly, because running large numbers of runs can cause the database to reach the maximum table size.

Note: This only affects the data of completed runs.

- **DELETE /executions/rerun**: Purges the rerun data from the database.
- **DELETE /steps-log**: Purges step data according to time and number of executions to purge.

By using the purge APIs, you can purge the data manually as required, or by scheduling recurrent flows that incorporate these APIs.

For more information about purging APIs, see the Operations Orchestration *API Guide*.

For more information about purge operations, see the Operations Orchestration *Database Guide*.

Adjusting the Number of Database Connections

Your system may be performing badly because the minimum and maximum database connection pool size is not set correctly.

Analyzing Whether You Need to Adjust the Database Connections

To identify if the number of database connections is an issue in your environment:

1. Open the **database.properties** file (located under **<installation_folder>/central/conf/**), and register the value of the `db.pool.minPoolSize` and `db.pool.mxPoolSize` properties.

The Central server aims to keep the number of database connections at the defined minimum value. If necessary the Central server will add connections but will not exceed the maximal value.

2. Review your database server configuration and check the current limitation on the number of database connections.

Note that in some cases, this limitation is global (the sum of all connections to all database schemas) and in other cases, resource-usage profiles may apply. Consult your DBA, in that case.

3. Connect to the database server and track the number of connections from this Central server to the database throughout busy hours. It is important to count only the connections that originated from this specific Central server.

If you are using an HPE OO cluster, you must configure the database server to allow connections from a number of Central servers, as well as connections from other clients and consumers.

Changing the Number of Database Connections

In order to change the maximum number of database connections on the Central server's end:

1. Open the **database.properties** file (located under **<installation_folder>/central/conf/**).
2. Edit the `db.pool.maxPoolSize` property.
3. Restart the configured node.
4. Repeat for every Central node.

Configuring the Amount of Data Written to the Database

In Central, a very detailed run history is persisted. This makes it easy to troubleshoot, as all the information is available in the Run Log. However, if your database size has increased to the limit, you may need to reduce the information that is saved to the Run Log.

Two persistence levels are available, and each one saves a different set of data:

- **Standard** - large input/output values are truncated at approximately 4,000 bytes when logged
- **Extended** - large input/output values are not truncated when logged



The screenshot shows a configuration window titled "Run Log Persistence Level". Inside, there is a label "Default persistence level:" followed by a dropdown menu. The dropdown menu is open, showing three options: "Extended" (selected and highlighted in blue), "Standard", and "Extended". A mouse cursor is pointing at the bottom of the dropdown menu.

Under the **Database Health** tab in the System Configuration workspace, select the default log level. This will be applied, by default, to all flows that are run.

You can override this default for individual flows in the flow library, or when triggering or scheduling a flow run.

You can also set the log level in API runs (REST/SOAP). For more information, see the *HPE OO API Guide*.

Scaling Out

Scalability is the ability of the HPE OO system to be enlarged to accommodate a growing amount of work, so that it increases its total throughput under an increased load.

Adding More Central Servers

One way to scale out is to install additional Central servers in an HPE OO cluster.

Clustering provides high availability and scalability to enhance throughput. To create a cluster, you run the Installation wizard to create the first Central. Then, you run it again on the other machine to create the next node and, during this second installation, make it point to the same database schema.

For more information, see "Installing an HPE OO Central Cluster" in the *HPE OO Installation, Upgrade, and Configuration Guide*.

Adding More Workers

Another method of scaling out is to add more workers to the existing HPE OO Central server.

Workers are responsible for executing flows. An external worker connects to Central to obtain tasks (flow execution messages) to process.

To create a new worker, install a new RAS. For more information, see "Installing a RAS" in the *HPE OO Installation, Upgrade, and Configuration Guide*.

