HP Operations Orchestration

Benchmark Performance Guide

Version 10.50

Document Release Date: November 2015
Software Release Date: September 2015

Legal Notices

Warranty

The only warranties for HP 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. HP 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 HP 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.

Copyright Notice

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation. UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: **http://h20230.www2.hp.com/selfsolve/manuals**

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to: **http://h20229.www2.hp.com/passport-registration.html**

Or click the **New users - please register** link 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 HP sales representative for details.

Support

Visit the HP Software Support Online web site at: https://softwaresupport.hp.com/
This web site provides contact information and details about the products, services, and support that HP Software offers.
HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:
http://h20229.www2.hp.com/passport-registration.html
To find more information about access levels, go to:
http://h20230.www2.hp.com/new_access_levels.jsp
HP Software Solutions Now accesses the HPSW Solution and Integration Portal Web site. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this Web site is http://h20230.www2.hp.com/sc/solutions/index.jsp

Contents

Executive Summary

This document provides an overview of the performance of HP Operations Orchestration version 10.50.

- Performance Improvements in HP OO 10.50
- Performance results for HP OO 10.50

Following are the results, which are described in more detail throughout this document:

- HP OO 10.50 shows a significant increase in the overall execution throughput in comparison with HP OO 9.07.0003.
- HP OO 10.50 shows a significant increase in the overall execution throughput in comparison with HP OO 10.20.

Objectives

This document details the performance improvements made in HP Operations Orchestration 10.50 and the performance tests and results using flow/step execution throughput (steps/time). The document includes the following:

- HP OO throughput in several environments:
  - Low cost FOSS (Free and Open-Source Software) operating system
  - High cost non-FOSS environments
  - Oracle-based environments
  - Clustered and stand-alone environments
- Single flow performance results of various scenarios comparing HP OO 10.50 to HP OO 10.20, and to HP OO 9.07.0003

Basic tuning was applied to the environments described in this document. These configurations are described in Recommendations for Environment Tuning.

Changes in the Performance Benchmark in OO 10.50

The HP Operations Orchestration 10.50 benchmark document introduces several changes related to the comparison methods and the distribution of the actual flows.

It also introduces newly designed performance flows, which were used for both the throughput and the single flow performance comparisons.

The original flow distribution has been kept and used as a reference but **will be removed in the future version(s)** of the benchmark document. As well, the original flows have been slightly redesigned in HP OO Studio to match the authoring best practices that are available in the HP OO Community.

Changes in Throughput Flows

Step Numbering Computation Changes

In the 10.20 benchmark and earlier, we relied on HP OO's database mechanism of step storage for step numbering computation. Starting with the 10.50 benchmark, we decided to align the step numbering computation process with the HP OO Central UI. This change enables us to be consistent with the benchmark of any future versions of the product.

The table below shows the number of steps at runtime for each flow that was used in this benchmark (Steps according to the HP OO Central UI).

Flow Type	Number of Steps per Flow Type at Runtime
Advanced Flow	120
Advanced Context Flow	120
Medium Flow	102
Parallel Flow	112
Multi-Instance Flow	603
Sub Flow	104
Short Flow	2
Large Context Flow	103
Long Flow	10002

Note: In future benchmarks, we will only use the new flow distribution.

Flow Distribution Changes

Starting with the 10.50 benchmark, we added two extra flows. The newly added flows (referred to below as Advanced Flow and Advanced Context Flow) have been designed to match:

- Customer-like cases of multiple sub-flows being called iteratively and data passed from parent to sub-flow, from step to step, and vice-versa.
- Customer-like cases as mentioned above but with large data being used.

The addition of the two flows triggered the need to change the original flow distribution, as we wanted to keep the 5000 runs being triggered at once. Thus, we ended up with the distribution shown below.

Flow Type	Number Per Flow Type
Advanced Flow	1800
Medium Flow	1000
Parallel Flow, Short Flow, Sub Flow, Multi-Instance Flow	480
Advanced Context Flow	200
Large Context Flow, Long Flow	40

Note: For the new flow distribution, all the results were achieved using the new step numbering computation. The number of steps at runtime for the new flow distribution is 1140280. These will become a reference for future benchmark documents.

Changes in Single Flow Performance Flows

We decided to create new flows with longer execution times. We will provide the results of the old flow and the new flows. In the next benchmark, we will use only the new flows.

The table below shows the number of steps at runtime for each flow that was used in this benchmark (Steps according to the HP OO Central UI – for more details, see page 5).

Flow Type	Number of Steps per Flow Type at Runtime
Large MI Flow	5080
Parallel Flow	2045
Subflows Level10	1030
Large Sequential Flow	20005
Large Context Flow	130
Long Flow	10002
Advanced Context Flow	120

Note: In future benchmarks, we will only use the new SFP flow.

Version-related Changes

STANDARD and EXTENDED persistence level:

In HP OO 10.50, we added the capability to handle database persistence level either as STANDARD or EXTENDED, which controls the level of detail saved to the run log. Because of this change, we ran benchmark tests for both persistence levels. The results and charts have been computed accordingly.

The STANDARD persistency level potentially means keeping less data in the database per-step (truncating large step inputs/outputs/results).

If flows are using small inputs/outputs/results (not reaching the truncation threshold), the difference between STANDARD and EXTENDED is minor in terms of database size.

Note: We recommend to use STANDARD persistence level for better performance.

Note: We will continue to keep STANDARD and EXTENDED results for the new flow distribution in future benchmark documents.

Setup

This section describes the different benchmark tests that are described in this document, including:

- Environment-related details
- The tools that were used
- The flows that were triggered and the flow distribution
- The results that were achieved, which showcase:
    - Throughput for:
        - The original flow distribution, which was used for previous benchmark documents, and which uses the original step computation mechanism
        - The new flow distribution, which uses the new step computation mechanism
    - Single flow performance results for:
        - Flows that were used in previous benchmark documents
        - New flows that were added in the HP OO 10.50 benchmark document

Environment

The following table describes the hardware and software components used for the benchmark tests:

	Model	Processors	Memory	Storage	Network	Notes
Server	ProLiant BL460c G7	12 core 2667Mhz	16 GB	Local	1 GB	Windows 2012/RHEL 6.3
Database	ProLiant DL380 G7	12 core 2933Mhz	32 GB	DAS	1 GB	RHEL 6.3 - Oracle

Tools

The following tools were used to produce this benchmark:

- HP LoadRunner 11.52
- HP SiteScope 11.20

Throughput Flows

This section describes the flows that were used during the benchmark throughput tests.

These flows were designed to stress the different functionality aspects of HP OO and to load the different resources of the system (CPU, memory, and so on). By running a mix of all of these flows, we tried to simulate a heterogeneous customer environment.

Note: All the flows are the same as the previous benchmark accept except for Advanced Flow and Advanced Context Flow.

Note: The purpose of these flows was to load HP OO as a platform, and not to perform any actual work, as the goal of the benchmark is to verify the performance of HP OO as a platform and not to verify the performance of the HP OO content.

Large Context Flow

This flow receives a 4 MB context and has 103 steps.

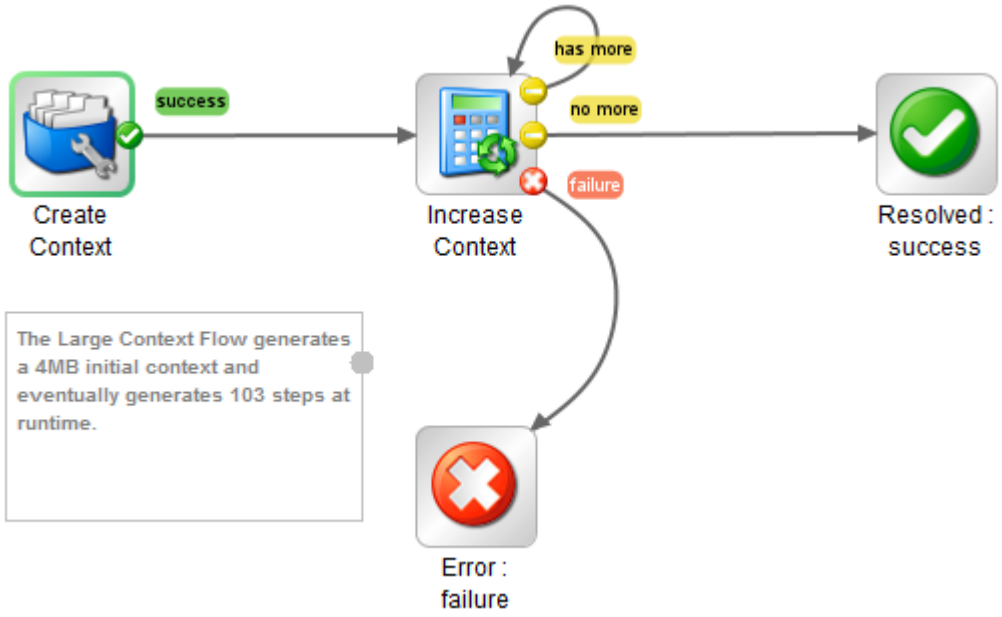

Short Flow

The following flow uses “Generate Data” operation and has 2 steps.

Medium Flow

The following flow uses the “Generate Data” operation and has 102 steps.

Long Flows

The following flow uses the “Generate Data” operation and has 10002 steps.

Multi-Instance Flow

This flow contains a multi-instance implementation of the UUID generator and runs with 300 lanes per flow.

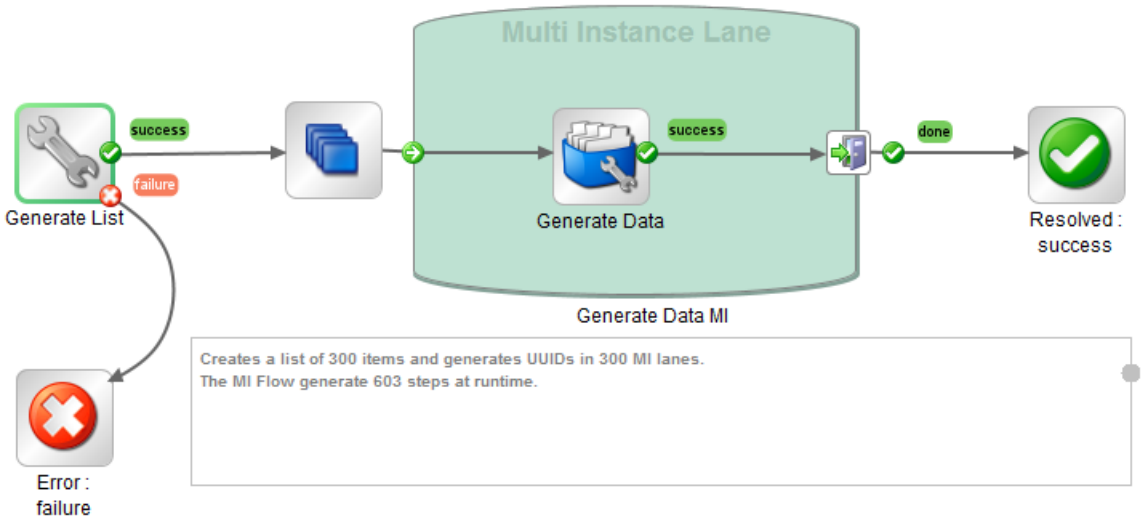

Subflow

This flow runs an instance of a medium sized flow as a sub-flow.

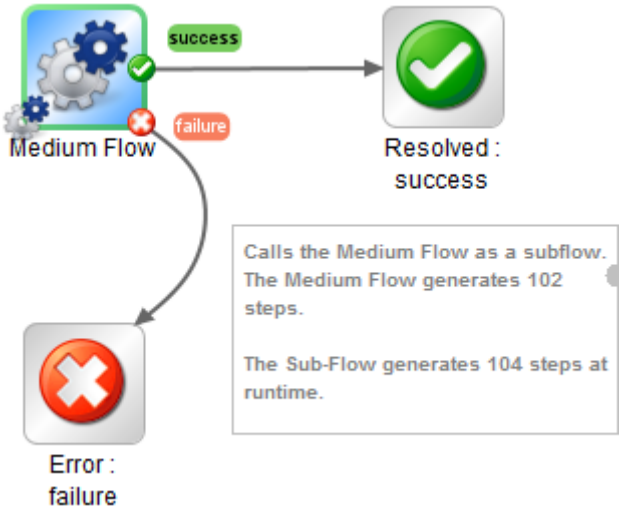

Parallel Flow

This flow runs 55 lanes of parallel split (only part of the flow can be seen in the following image).

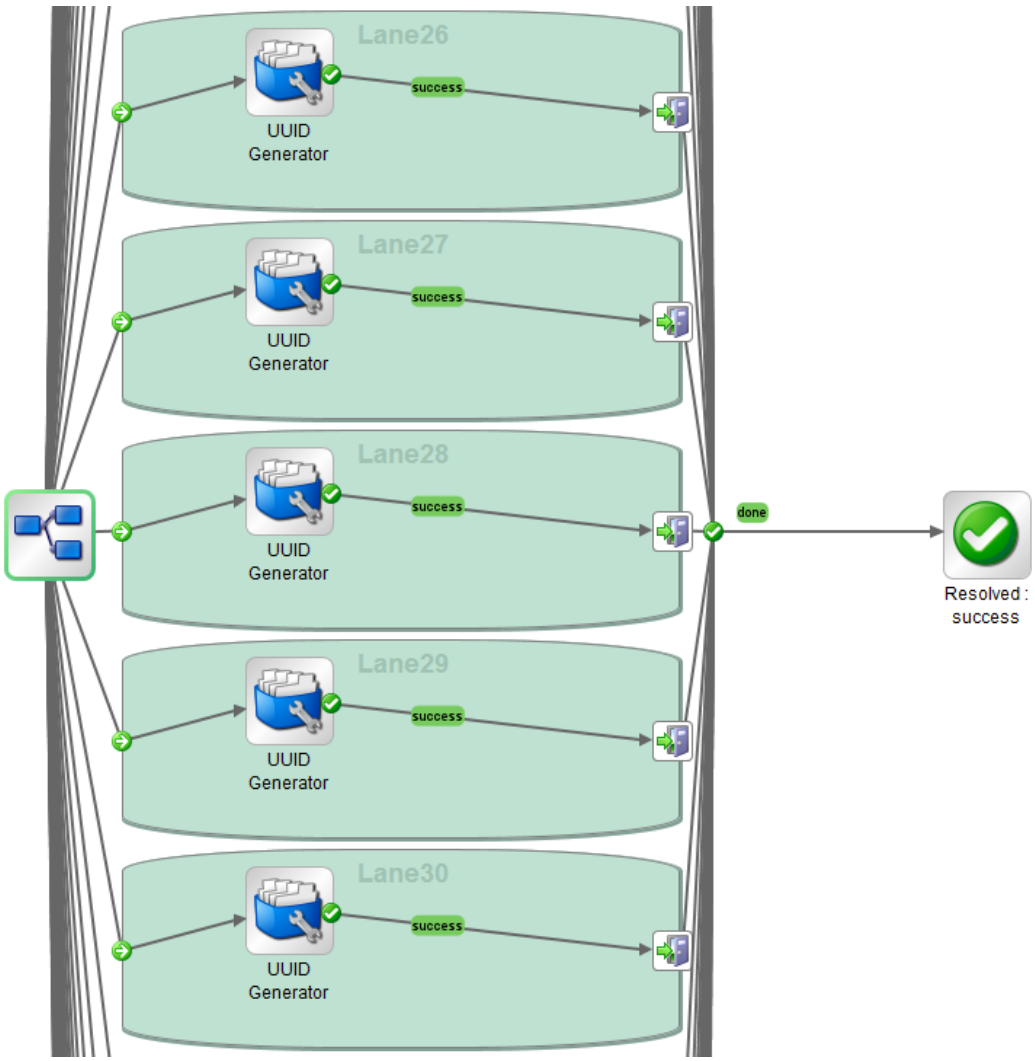

Advanced Flow

This flow was designed to mimic consecutive calls of multiple short and medium sized subflows from within a parent flow, while passing data from the parent flow to the subflow and vice-versa.

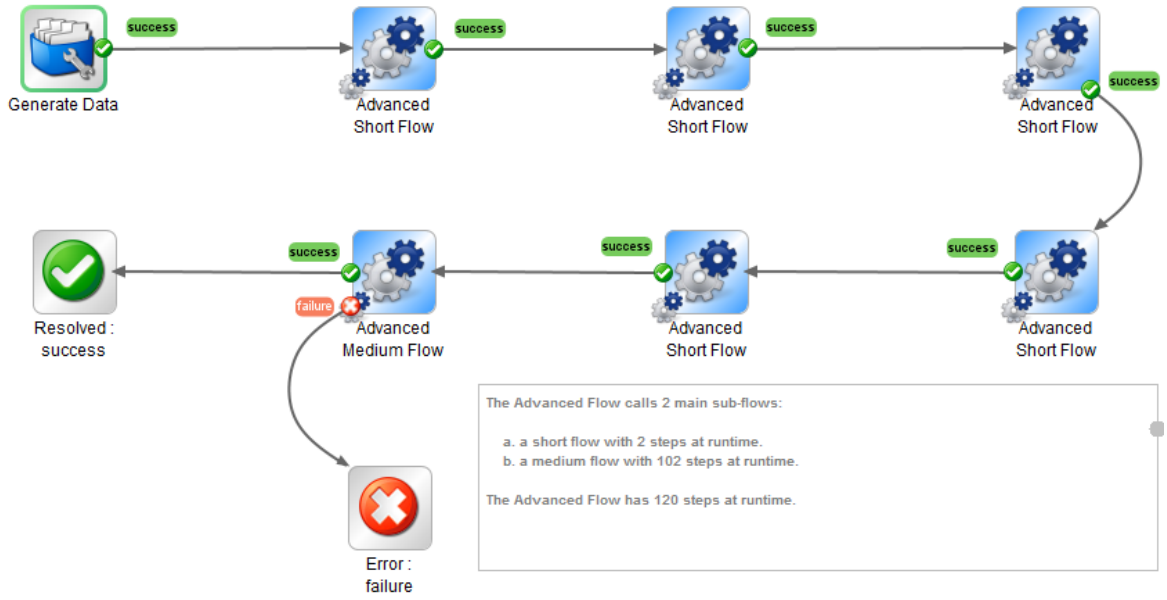

Advanced Context Flow

This flow was designed to mimic consecutive calls of multiple short and medium sized sub-flows from within a parent flow, while passing a large amount of data from the parent flow to the sub-flow and vice-versa. 4500 ASCII chars are used and passed from inputs to outputs, primary results and raw results. This flow uses scriptlets and filters.

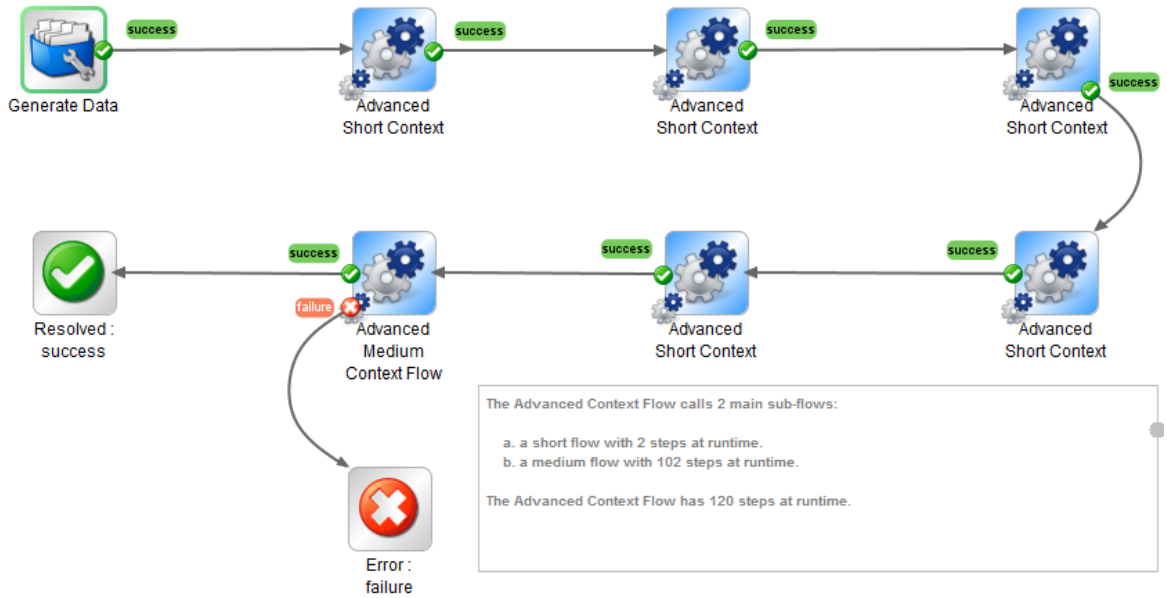

Single Flow Performance Flows

To be more exact with the results, we decided to create new flows with longer execution times.This section describes the new flows with longer execution times that were used during the benchmark for single flow performance tests.

Large MI Flow

This flow loops 25 times and uses two multi-instance steps, each of which has 50 instance steps.

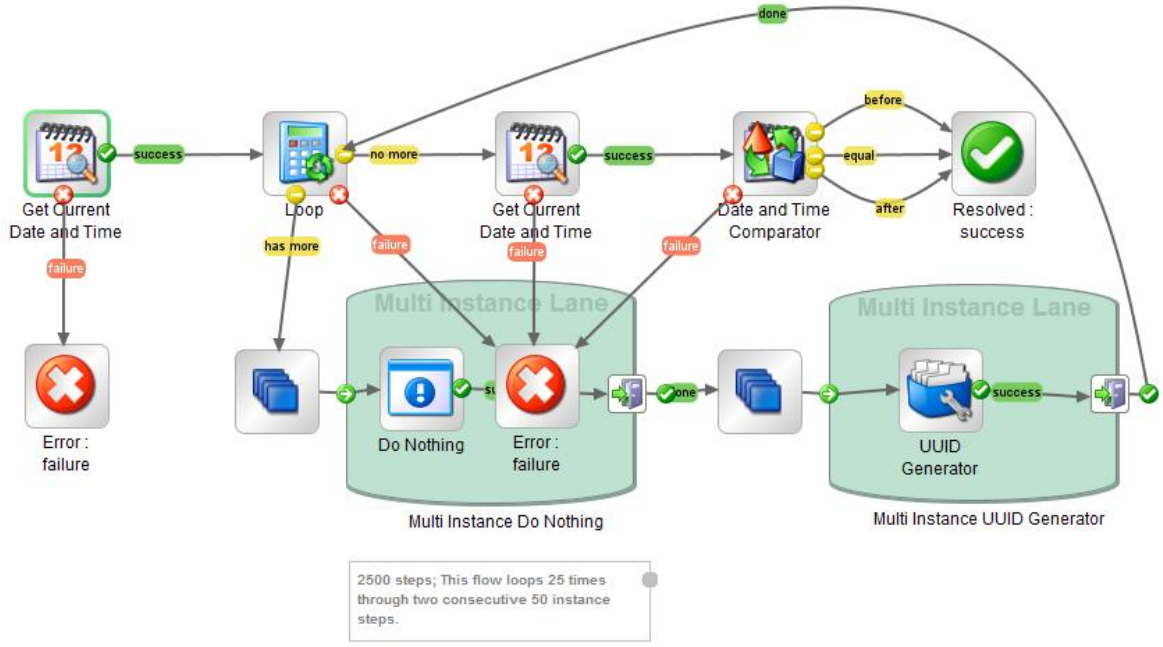

Parallel Flow

This flow loops 50 lanes of a parallel split step 20 times (only part of the flow can be seen in the following image).

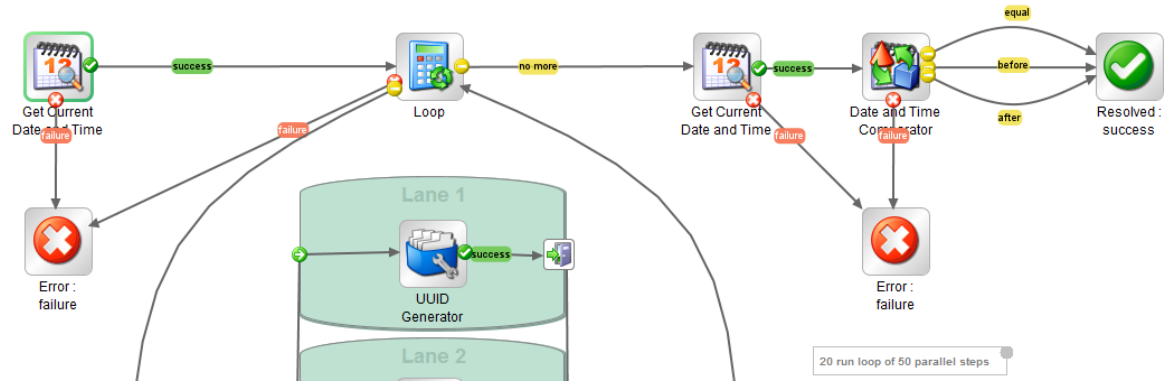

Subflows Level10

This flow loops 25 times, and each iteration contains 10 levels of subflows.

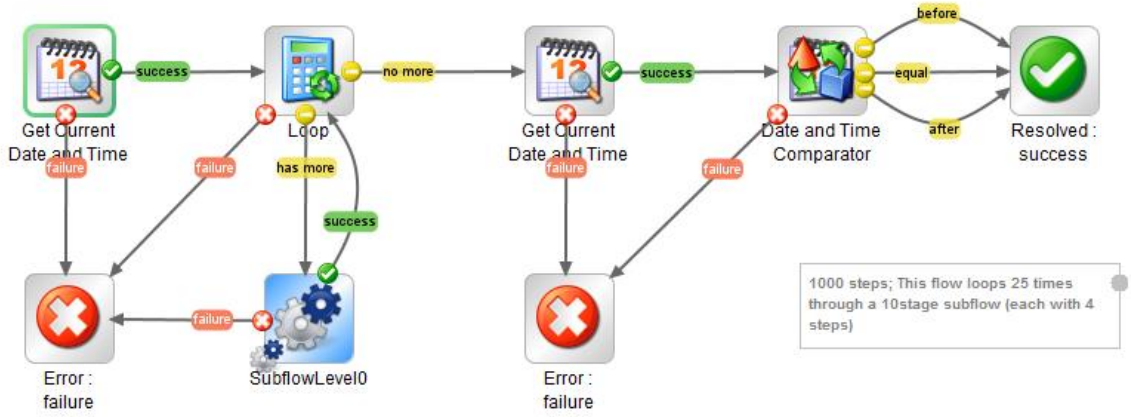

Large Sequential Flow

This flow loops 10,000 times and uses the “Do Nothing” operation.

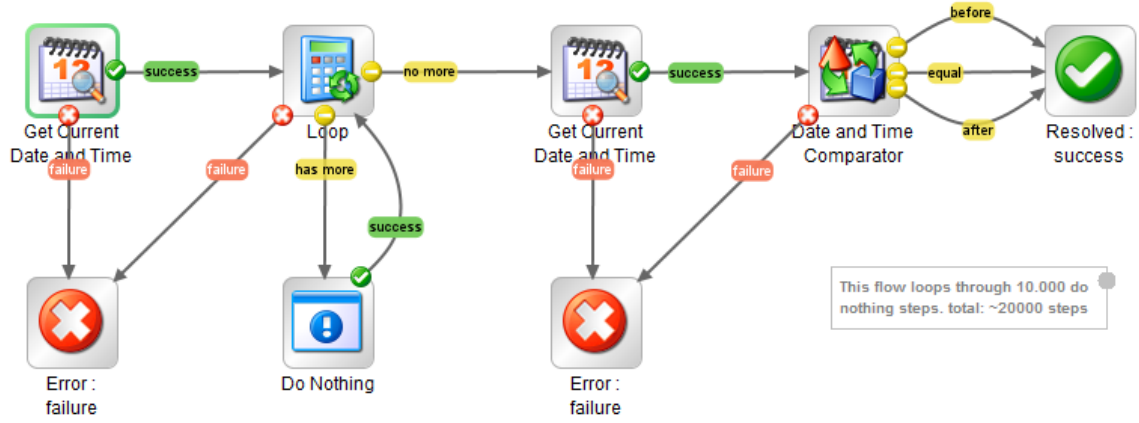

Large Context Flow

This flow loops 150 times, and uses the “List Appender” operation, which creates a large context variable.

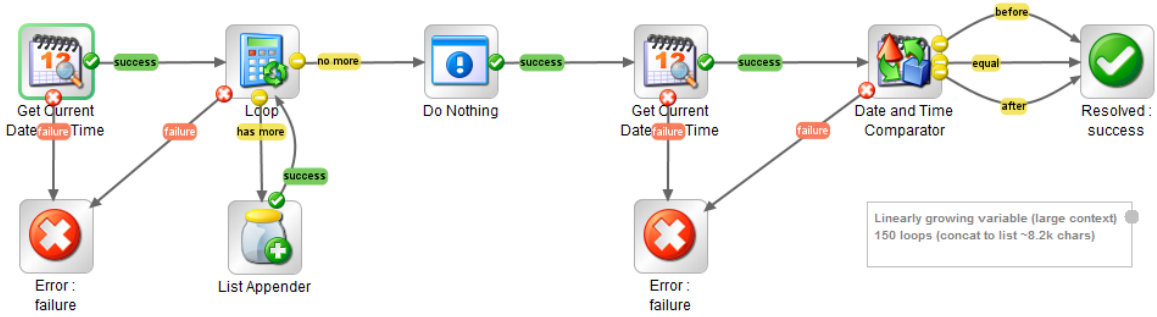

Scenario

This section describes the scenario used for the benchmark:

- The workload was generated using HP LoadRunner.
- The flow triggering was done using HP OO REST API calls.
- The number of flows ran in each benchmark is 5000, which amounts to:

  Original flow distribution: 1660640 steps according to the original model of step computation and 1310780 according to the new model of step computation. The results within the document were computed with the original model of step computation.

  New flow distribution: 1140280 steps according to the new model of step computation.

  See Step numbering computation changes.

The distribution of the flows was as follows:

Original Flow Distribution (10.20 Flow Distribution):

Flow	Number Per Flow Type
Medium Flow	1000
Parallel Flow, Short Flow, Sub Flow, Multi-Instance Flow	980
Large Context Flow, Long Flow	40

New Flow Distribution (10.50 Flow Distribution):

Flow	Number Per Flow Type
Advanced Flow	1800
Medium Flow	1000
Parallel Flow, Short Flow, Sub Flow, Multi-Instance Flow	480
Advanced Context Flow	200
Large Context Flow, Long Flow	40

- We used HP Sitescope integration with LoadRunner to monitor the different parts of the system during the tests, including JMX monitors for the JVM monitors (memory, garbage collection).

Comparison

Throughput – Original (10.20) Flow Distribution Results

The following chart compares the throughput between the different HP OO deployments used in this benchmark.

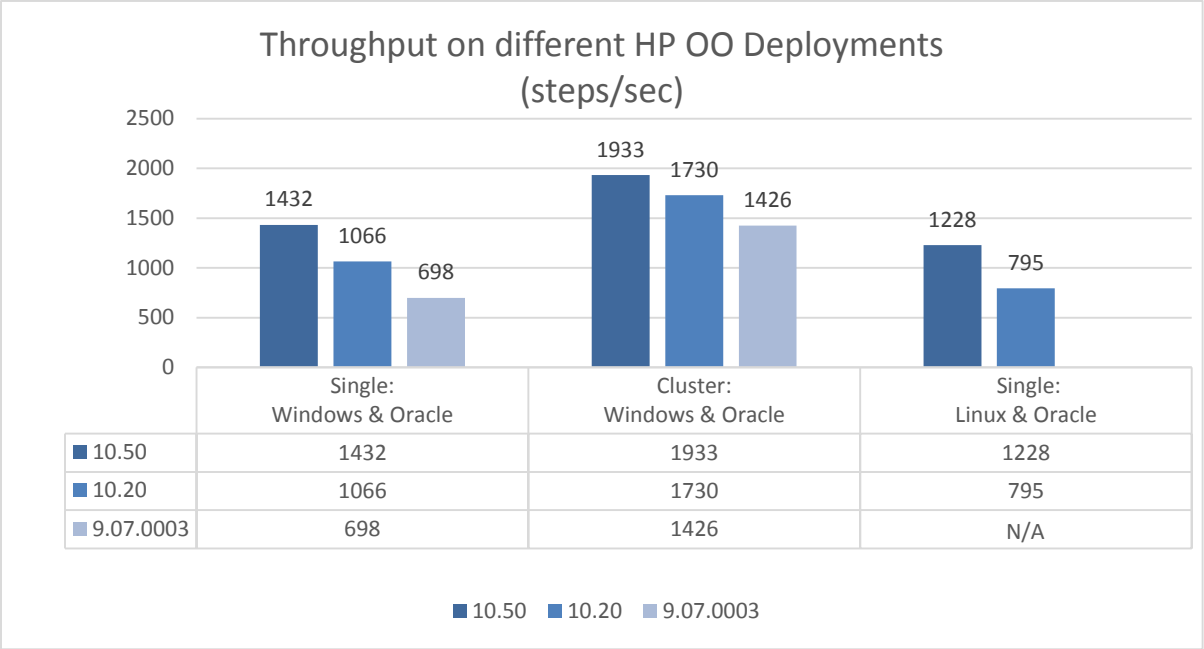

Note: A higher result shows better performance.

Note: This chart is only relevant for the current benchmark, and will be removed in future Benchmark Performance documents. The new results can be found here.

Note: HP Operations Orchestration was scaled out (in versions 9.07.0003, 10.20, and 10.50) by adding additional Central servers.

The following chart show the throughput and the performance increase of 10.20 and 10.50 over 9.07.0003 in either single or cluster mode configurations.

Note: The following charts use the original flow distribution, the original step numbering computation, and **EXTENDED** persistence level.

Note: The following charts relevant for the current benchmark, and will be replaced by new flow distribution in future Benchmark Performance documents.

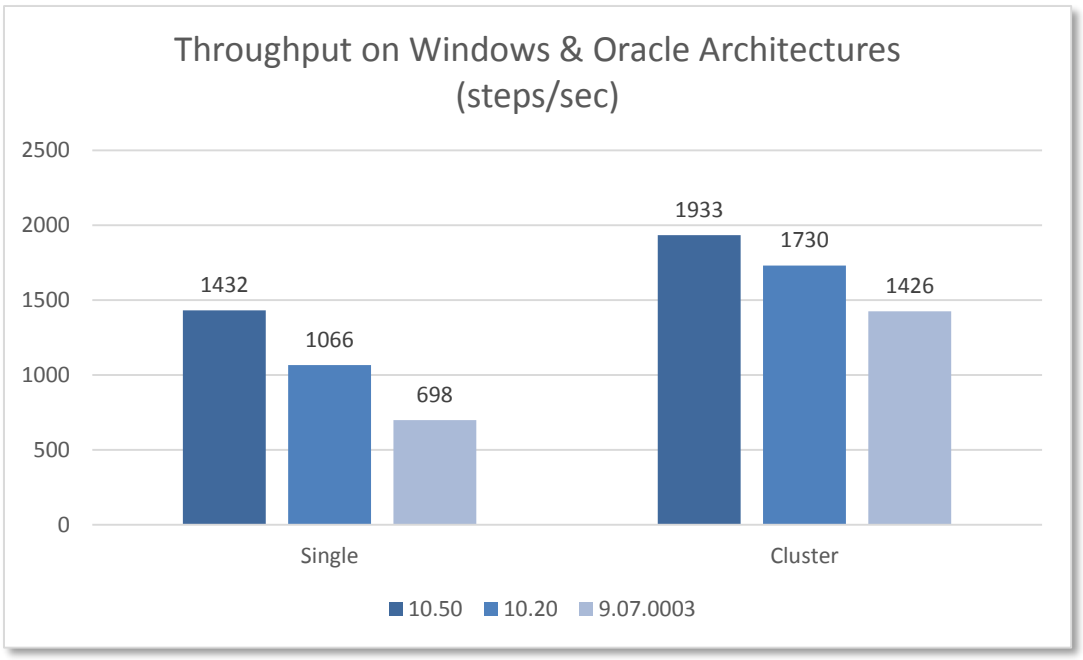

Note: A higher result shows better performance.

The following chart shows the performance increase in % of 10.20 and 10.50 over 9.07.0003.

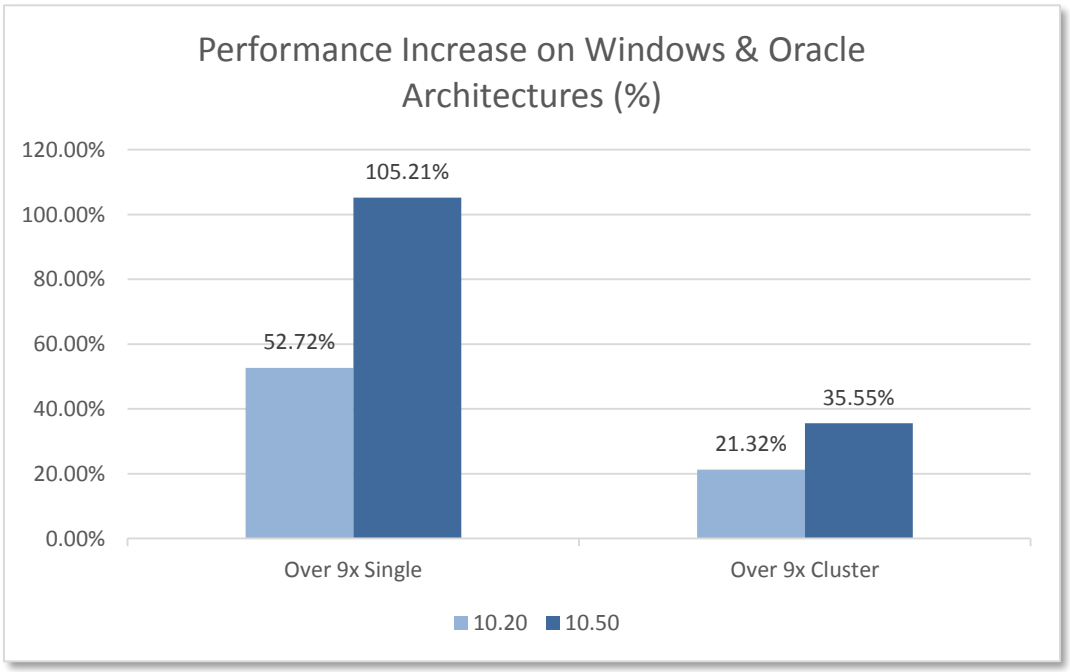

Throughput - New Flow (10.50) Distribution Results

The following charts show the throughput and the performance of 10.50 in either single or cluster mode configurations.

The following chart uses the new flow distribution, the new step numbering computation, and **STANDARD** persistence level.

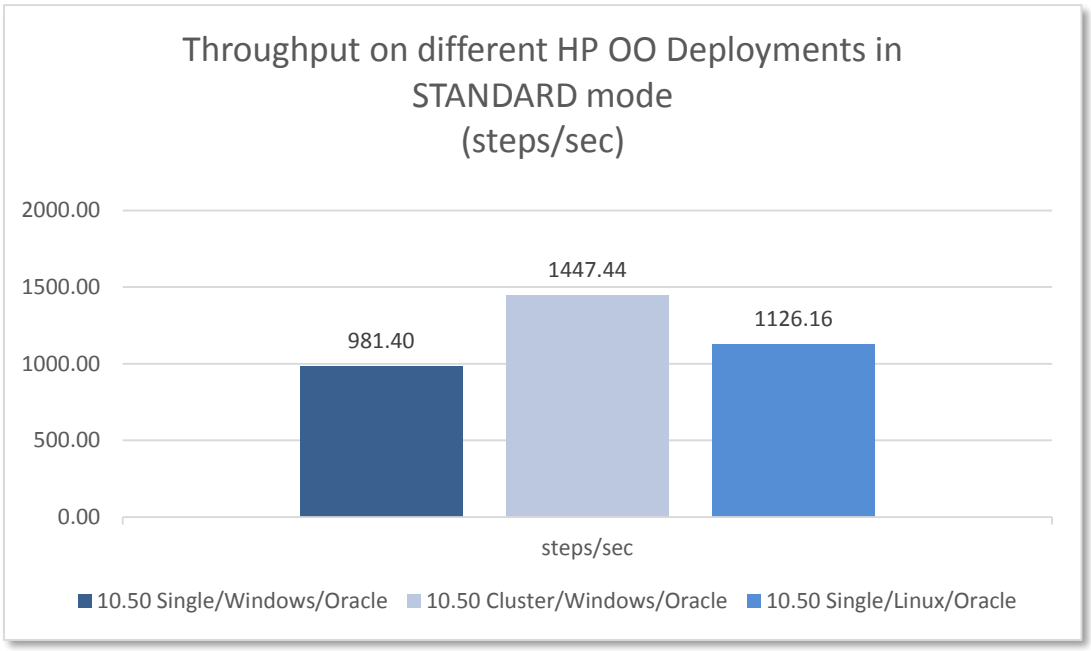

The following charts use the new flow distribution, the new step numbering computation, and **EXTENDED** persistence level.

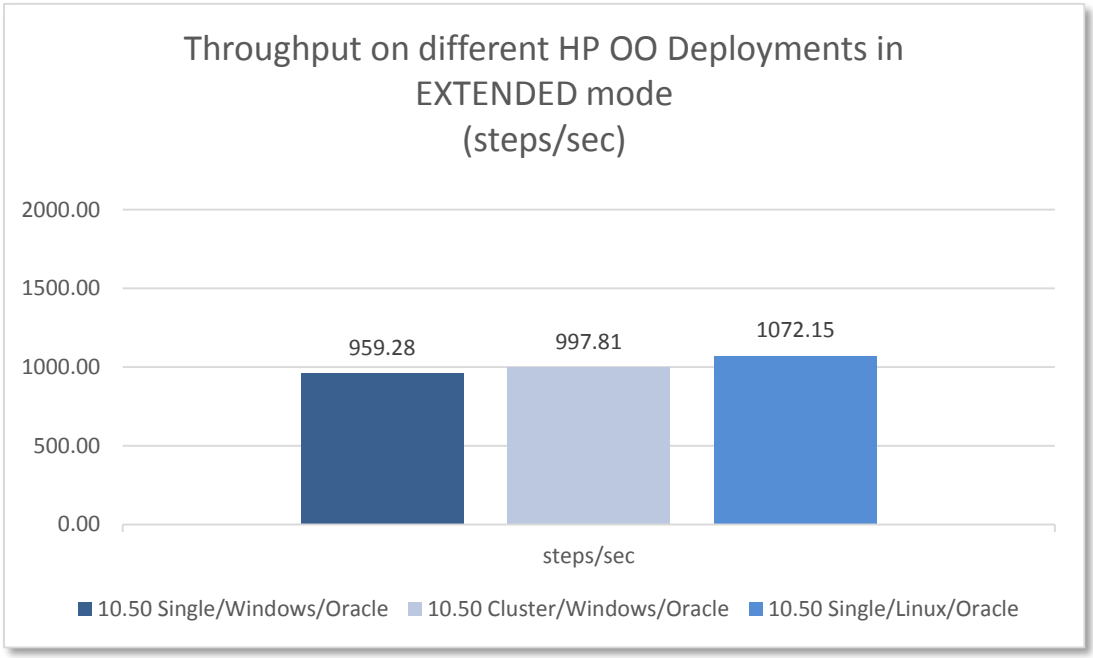

Single Flow Performance – Old Flows

The following chart compares the single flow performance of HP Operations Orchestration 10.50 to 10.20 and 9.07.0003.

Note: This section uses the 10.20 single flow performance benchmark flows (old flows)

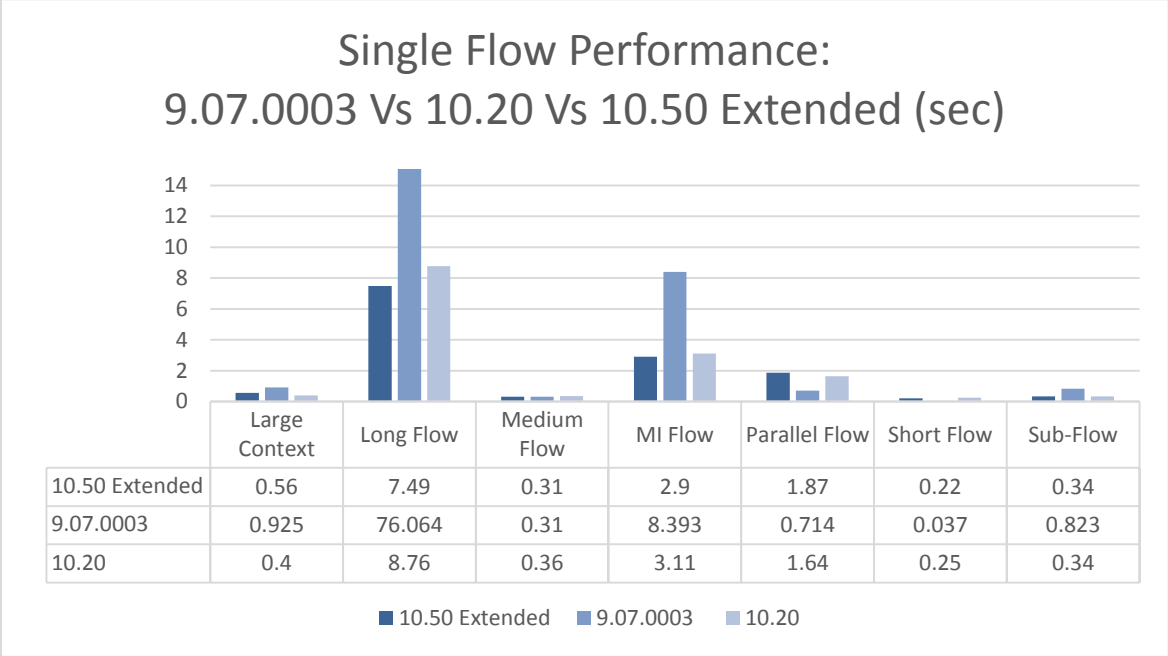

	Large Context	Long Flow	Medium Flow	MI Flow	Parallel Flow	Short Flow	Sub-Flow
10.50 Extended	0.56	7.49	0.31	2.9	1.87	0.22	0.34
9.07.0003	0.925	76.064	0.31	8.393	0.714	0.037	0.823
10.20	0.4	8.76	0.36	3.11	1.64	0.25	0.34

Note: A lower result shows better performance.

Single Flow Performance – New Flows

The following charts show the performance of single flows used for the throughput benchmark runs of the new flows.

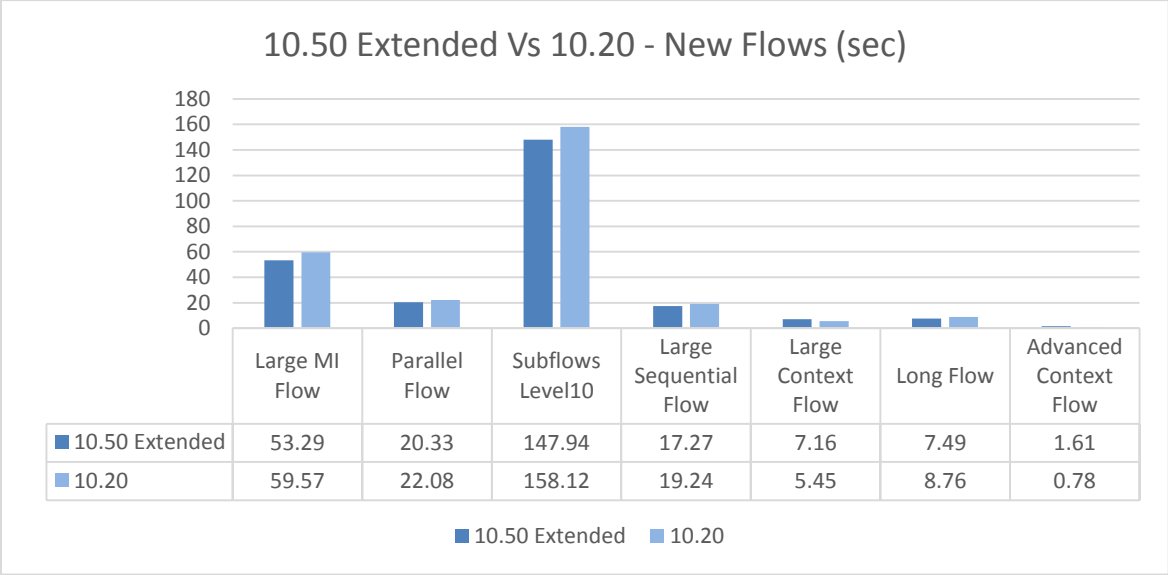

	Large MI Flow	Parallel Flow	Subflows Level10	Large Sequential Flow	Large Context Flow	Long Flow	Advanced Context Flow
10.50 Extended	53.29	20.33	147.94	17.27	7.16	7.49	1.61
10.20	59.57	22.08	158.12	19.24	5.45	8.76	0.78

Note: A lower result shows better performance.

Analysis of Results

HP Operations Orchestration 10.50 introduced significant architectural refactoring, which includes, among other changes, the design goals for the new execution queue:

- Use bulk operations
- Utilize a task and worker design pattern
- Minimize database locking
- Statelessness

Due to the fact that the system now accumulates bulks of steps for execution, the overall throughput is increased.

On the other hand, leveraging the task and worker design, this overhead does not exist when we compare flows that have a large number of steps or flows that have multiple lanes (parallel/multi-instance, and so on).

Why Scale-out was Done by Adding Central Servers and not RAS Servers?

During this benchmark test, we scaled our system by adding additional Centrals and not RAS servers. This is currently the recommendation due to the following reasons:

- RAS servers have an additional overhead, because they have to receive work and send results from Central remotely via REST.
- HP OO balances the work between all the workers in the same group equally. Remote workers will decrease the average throughput of the system, due to their inherent overhead.

Therefore, for these reasons, it is recommended to scale out using RAS only when it is needed due to a functional requirement.

Note: See the O*perations Orchestration 10.x Concepts Guide* for more information. You can download this document from HPLN or the SSO Portal. It is also available in Central and Studio in the online help or in PDF format located in the HP OO Documentation folder.

Recommendations for Environment Tuning

The following configurations were made during the benchmark tests.

OO 9.07.0003

- The *heap size* was increased to 1 GB - 4 GB.
- Persistency for flow steps was turned off. This was done by setting the following parameters under **Administration** > **System Configuration**:

OO 10.50

- The *heap size* was increased to 1 GB - 4 GB.

  This can be configured in **<OO Installation>\oo\central\conf\central-wrapper.conf**.

  # Initial Java Heap Size (in MB)

  wrapper.java.initmemory=1024

  # Maximum Java Heap Size (in MB)

  wrapper.java.maxmemory=4096

- The *number of execution threads* was increased to 300 while the i*nBuffer capacity* was increased to 500.

  These can both be configured in **<OO Installation>\oo\central\conf\central-wrapper.conf** starting from HP OO 10.50.

  wrapper.java.additional.25=-Dcloudslang.worker.numberOfExecutionThreads=300

  wrapper.java.additional.26=-Dcloudslang.worker.inBufferCapacity=500

- The number of *database connections* was increased to 20 - 100.

  This can be configured in **<OO Installation>\oo\central\conf\database.properties**.

  db.pool.maxPoolSize=100

  db.pool.minPoolSize=20

- In HP OO 10.50, the recommended way to scale out is by adding additional Central servers. This was done as part of these benchmark tests.