

SAS® on HP MSA 2040 Storage benchmark report



Up to 1.2 GB per second of I/O throughput with 64 concurrent workflows

Table of contents

Executive summary	2
Introduction	2
SAS job characteristics	2
Test topology	3
Test methodology	4
Test results	4
Test analysis summary	4
Recommendation	4
Appendix A – Volume Group create script	5
Appendix B – Logical Volume create script	5
Appendix C – File System create script	5
Appendix D – Device File manipulation script	5
For more information	6

Executive summary

We live and work in a new era of extreme business speed with heightened customer, partner, and employee expectations. To better compete and grow, businesses demand more innovation, speed, and flexibility from their data centers. Additionally, Foundation SAS customers require a hardware/software configuration that can deliver data analysis results quickly and accurately. To meet that need, enterprise customers demand reliable and fast storage that can scale to meet their business analytics requirements. This report details a proof point for Foundation SAS that demonstrates the performance available with the HP MSA 2040 Storage entry-level offering.

This report highlights the key findings from running Foundation SAS using a SAS workload running on an HP DL980 G7 Server and HP MSA 2040 Storage. This test showcased the performance of the MSA 2040 Storage with a heavy SAS workload.

SAS ran the tests and has certified that the test suite was run correctly, and that it represents the needs of a typical user.

Target audience: The target audience for this performance report is the IT community studying solutions for their environments. Business users and IT professionals who are interested in implementing a SAS solution may find this report useful for a sample SAS configuration. IT professionals that either already have a SAS implementation on MSA storage or are considering one may find suggestions for improving the performance of their systems, as well as a proof-point for running SAS with the MSA 2040.

This white paper describes testing performed in August 2013.

Introduction

SAS creates test kits to provide representative workloads for performance testing and tuning. This workload scenario utilizes real-world data volumes and structures of a typical SAS customer.

The scenario simulates activities that are typical of any data manipulation or data preparation process. The goal may be extract, transform, and load (ETL) from operational systems to a data mart, or it may be preparing and summarizing data in preparation for analysis, reporting or exploitation via other SAS tools.

SAS job characteristics

The Foundation SAS workload used for this test had the following characteristics:

- DATA Step
- Iterate 1,000³ calculations
- PROC Format
- PROC SQL to create indexes
- PROC SORT by Character Variables
- PROC SUMMARY

Important

This test scenario can provide a benchmark for comparing hardware and/or software products; it is not intended to be used as a sizing guideline. In the real world, server performance is highly dependent upon the application design and workload profiling.

As with any laboratory testing, the performance metrics quoted in this paper are idealized. In a production environment, these metrics may be impacted by a variety of factors.

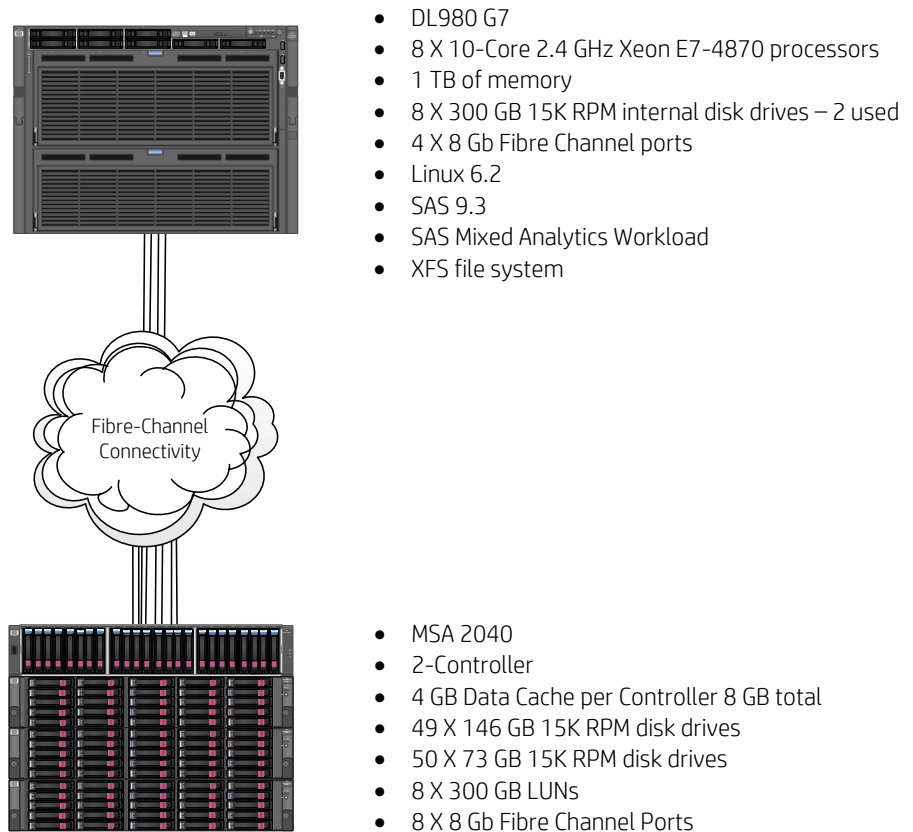
As a matter of best practice for all deployments, HP recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HP Solution Design Services (SDS) representative at sastech@hp.com.

Test topology

The application environment under test included several software components:

- Foundation SAS 9.3
- Red Hat® Enterprise Linux (RHEL) 6.2
- XFS File System

Figure 1. Architectural diagram of test bed



The test environment was comprised of Foundation SAS 9.3 on a single HP DL980 G7 server running Red Hat Enterprise Linux 6.2. This server was connected to a Fibre Channel switch by 4 x 8 Gb Fibre Channel ports, which was connected to the HP MSA 2040 Storage with 8 x 8 Gb Fibre Channel ports. The MSA 2040 was configured with 8 X 300 GB LUNs in RAID-10 configuration. No shared or clustered file system was used. Logical storage for the test environment consisted of one volume group. The single volume group was divided into a SASData and SASWork file systems via Logical Volume Manager. Those file system sizes were 1 TB each¹.

Changing storage needs for server customers now requires greater levels of performance and functionality at a lower cost of ownership. The next generation HP MSA 2040 Storage arrays raises the bar and uniquely addresses HP customers' storage and data protection needs, reduces total cost of ownership while dramatically increasing performance using technologies such as SSDs, Snap and Volume copy and Replication.

The new HP MSA 2040 Storage arrays are positioned to provide an excellent value for customers needing increased performance to support initiatives such as deploying SAS. The MSA 2040 delivers this performance by leveraging new controller architecture with a new processor, four 8 Gb or 16 Gb FC ports with 4 GB cache per controller and using drive technologies such as SSDs.

¹ Volume Group and Logical Volume creation scripts are included in Appendix A and B respectively.

The tested MSA 2040 configuration had four 8 Gbps Fibre Channel connections and 4 GB of data cache per controller, with two controllers, for a total of 8 GB of data cache and eight 8 Gbps Fibre Channel connections. With HP's variety of storage systems there is a cost-effective solution that would meet the needs of any size SAS implementation.

HP ProLiant DL980 – an HP scale-up, resilient, x86 server based on the PREMA Architecture – is designed to take full advantage of the 10-core Intel® Xeon® processor E7-4800/2800 product families with Intel QuickPath Interconnect (QPI) technology. Working in concert, they form the foundation for unparalleled transactional performance, scalability, and energy efficiency, plus significantly lower TCO.

The Node Controller chipset design of the HP PREMA Architecture features Smart CPU Caching that reduces overhead, frees up memory bandwidth, and eliminates performance bottlenecks. With its Resilient System Fabric, PREMA delivers improved system availability, minimized downtime, improved system reliability, and performance – in addition to the HP ProLiant innovations, such as HP Integrated Lights-Out 3 (iLO 3) and Thermal Logic, which customers have come to expect.

Test methodology

The SAS workload was used for the testing. The workload consists of workflows. Each workflow is iterated twice. Performance differences are noted as the number of concurrent workflows is increased. We start with one workflow and then work with 2, 4, 8, etc. through a total of 64 workflows. The iteration average is then noted, as well as the total real time for the workflow.

Metrics monitored during the tests are:

- Workflow Iteration Average – Amount of CPU time it took the workflows to complete divided by the number of workflows run.
- Workflow Total Real Time – Total amount of CPU time all of the workflows consumed during the test.
- Percent Mean User Time – Amount of user time awarded to all SAS processes.
- Percent Mean System Time – Amount of system time consumed during the test
- Percent Mean Idle Time – Amount of idle time the system had during the test.
- Percent Mean Wait Time – Amount of time spent waiting for I/O to complete
- Mean Blocks In – This is the number of blocks read from disk summarized in 2 second intervals. Each block is 4096 bytes.
- Mean Blocks Out – This is the number of blocks written to disk, summarized in 2 second intervals.

SAS I/O is predominantly serial in nature, and because of that we configure storage differently for SAS than we would for more traditional OLTP environments. SAS I/O is large block serial in nature and typically has a 50-50% read-write mix. As a result, we configure for larger blocks of data to be transferred with each I/O than if we were performing the more random type of I/O typical in an OLTP environment. You can see from the logical volume creation scripts in Appendix B, that we used a 64K stripe size to enable these larger transfers.

Test results

The overall test results demonstrate the MSA 2040, as configured, will deliver 1.2 GB/sec of throughput to the SAS processes.

Test analysis summary

This workload demonstrated a high-performance, entry level MSA 2040 storage array. The MSA 2040 supported the SAS users with excellent scalability.

The SAS workload clearly showed that the MSA 2040 storage was able to support the loads necessary for a typical customer solution with outstanding throughput.

Recommendation

HP recommends HP MSA 2040 storage for entry level SAS data storage implementations. From the input/output area to the SAS work area, HP MSA 2040 is able to meet the demanding I/O footprint that SAS requires to provide users with good response times and shorter elapsed job runtimes.

While this report demonstrates a sample Foundation SAS implementation, SAS solutions will vary greatly between implementations. HP recommends that all customers get a sizing that is specific to their environment. HP provides this as a complimentary service at sastech@hp.com.

Appendix A – Volume Group create script

DISCLAIMER OF WARRANTY

This document contains the following HP scripts. These are provided as a courtesy, free of charge, “AS-IS” by Hewlett-Packard Company (“HP”). HP shall have no obligation to maintain or support this software. **HP MAKES NO EXPRESS OR IMPLIED WARRANTY OF ANY KIND REGARDING THIS SOFTWARE INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE OR NON-INFRINGEMENT. HP SHALL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES, WHETHER BASED ON CONTRACT, TORT OR ANY OTHER LEGAL THEORY, IN CONNECTION WITH OR ARISING OUT OF THE FURNISHING, PERFORMANCE OR USE OF THIS SOFTWARE.**

```
#!/bin/bash

vgcreate vgsas /dev/mapper/mpathb \
/dev/mapper/mpathc \
/dev/mapper/mpathd \
/dev/mapper/mpathe \
/dev/mapper/mpathf \
/dev/mapper/mpathg \
/dev/mapper/mpathh \
/dev/mapper/mpathi
```

Appendix B – Logical Volume create script

```
#!/bin/bash

lvcreate -n sasdata -i 8 -I 64 -L 1T vgsas
lvcreate -n saswork -i 8 -I 64 -L 1T vgsas
```

Appendix C – File System create script

```
#!/bin/bash

mkfs.xfs -f -b size=4096 /dev/vgsas/sasdata
mkfs.xfs -f -b size=4096 /dev/vgsas/saswork
```

Appendix D – Device File manipulation script

```
#!/bin/bash

lvchange --readahead 16K /dev/vgsas/sasdata
lvchange --readahead 16K /dev/vgsas/saswork
```

For more information

HP and SAS, hp.com/go/sas

HP ProLiant Servers and Blades, hp.com/go/proliant

HP MSA Storage, hp.com/go/msa

HP Complimentary Customized Sizing, sastech@hp.com

SAS Partner Directory, sas.com/partners/directory/hp

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.

Sign up for updates

hp.com/go/getupdated

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. 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.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries.
® indicates USA registration.

Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

