Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

MAXIMIZE SAS ANALYTICS 9.4 PERFORMANCE WITH HPE ALLETRA 9000

CONTENTS

EXECUTIVE SUMMARY

In today's digital world, enterprises deploy business intelligence and analytical solutions to gain a competitive edge and make real-time decisions based on data that's generated. These decisions require data that is compute-intensive and uses a lot of storage. These solutions need compute and storage platforms that are highly scalable and reliable.

This technical paper highlights the key findings and demonstrates scalability of a solution running SAS 9.4 using the Mixed Analytics Workload suite on HPE ProLiant DL580 servers and HPE Alletra 9000 arrays.

SAS 9.4 customers require a hardware and software configuration that delivers data analysis results quickly and accurately. Enterprises demand reliable and fast storage that can scale to meet their business analytics requirements. This paper details a proof point for SAS 9.4 that demonstrates the performance that four HPE ProLiant DL580 servers and an HPE Alletra 9000 array can bring to the SAS workload.

HPE ProLiant DL580 servers help to manage workloads processing exponentially growing data sets, while also ensuring flexibility for all transaction, analytical, and data management workloads.

HPE Alletra 9000 is an intelligent, cloud-native, Tier-0 data infrastructure for mission-critical apps that delivers ultra-low latency and extreme resiliency in an as-a-service model via HPE GreenLake. HPE Alletra 9000 uses artificial intelligence (AI) and machine learning (ML) powered by HPE InfoSight for analytics and prediction. This end-to end, app-aware approach to resiliency maximizes uptime for mission-critical apps and offers a 100% Availability Guarantee. HPE Alletra 9000 can be fully managed and operated from the Data Services Cloud Console (DSCC).

Target audience

The target audience for this technical paper is chief information officers (CIOs), chief technology officers (CTOs), data center managers, enterprise architects, and deployment and implementation engineers. Others who are interested in implementing a SAS solution may find this report useful for its sample SAS configuration and demonstration of the HPE ProLiant DL580 and HPE Alletra 9000 solution.

INTRODUCTION

SAS specializes in advanced and predictive analytics with industry-specific product sets for banking, government, retail, telecommunications, and aerospace, and for marketing optimization or high-performance computing. The SAS 9.4 platform streamlines data access and preparation, enhances data quality, and ensures optimal processes to increase the productivity of analytics projects.

The SAS Mixed Analytics Workload suite uses the real-world data volumes and structures of a typical SAS customer. The scenario simulates jobs received from various SAS clients such as Display Manager, Batch, SAS Data Integration Studio, SAS Enterprise Miner, SAS Add-In for Microsoft Office, SAS Enterprise Guide, and SAS Studio. Many customer environments have several SAS users or jobs that use analytics in their company's day-to-day business activities.

The scenario described in this paper consists of typical analytic jobs designed to replicate a light-to-heavy workload. These jobs were launched via a script that includes time delays to simulate scheduled jobs and interactive users launching at different times.

The purpose of these tests was to characterize the HPE ProLiant DL580 server and the HPE Alletra 9000 array, help users understand how to configure HPE ProLiant DL580 servers and HPE Alletra 9000 arrays to use the ideal combination of storage and compute, and to demonstrate the level of performance this combination delivers in a SAS 9.4 environment. The number of cores and the clock speed of the processor are what ultimately matter for the compute. This paper is a guideline for the performance that can be achieved on this compute and storage combination. When designing a SAS solution using HPE Alletra 9000, you can choose from different compute offerings based on the number of cores needed and the cost sensitivity. These tests will still give a fair expectation of the performance that can be achieved in your chosen environment.

SOLUTION OVERVIEW

This solution is based on HPE ProLiant DL580 servers running on Red Hat® Enterprise Linux® 8.3 combined with an HPE Alletra 9080 array.

The hardware components that form the solution are shown in Figure 1.

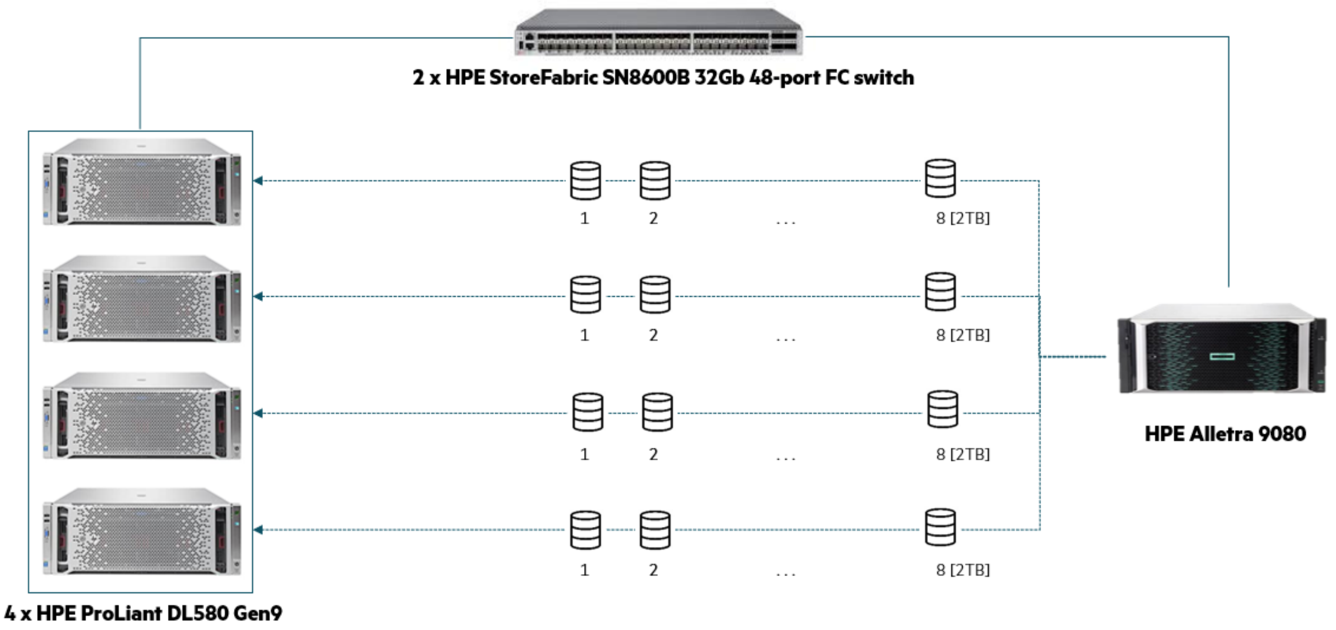

FIGURE 1. Front view of the solution featuring HPE ProLiant DL580 servers and HPE Alletra 9080 storage

SOLUTION COMPONENTS

The solution consists of the components described in this section.

HPE ProLiant DL580 server

The compute server used for this paper is the HPE ProLiant DL580 Gen9. The HPE ProLiant DL580 Gen9 server is a four-socket enterprise standard x86 server that offers reliability, commanding performance, and workload consolidation.

FIGURE 2. HPE ProLiant DL580 Gen9 server

Each of the four HPE ProLiant DL580 Gen9 servers used for this solution met the following specifications:

- Four sockets
- 18 cores per processor
- Intel® Xeon® CPU E7-8890 v3 @ 2.50GHz
- 1 TB of memory
- 45 MB of L3 cache
- HPE iLO (firmware HPE iLO4 2.0) 4 GB NAND
- Six dual-port QLogic Corp. ISP2722-based 16/32Gb Fibre Channel to PCIe Adapters

This solution uses HPE ProLiant DL580 Gen9 servers, but you can choose any of the other HPE compute models, such as the HPE ProLiant DL580 Gen10, HPE ProLiant DL380 Gen10, or HPE ProLiant DL380 Gen 10 plus, which ship with the latest Intel Xeon Scalable processors, support NVMe drives, and offer the latest in security, performance, and expandability.

HPE Alletra 9080 array

HPE Alletra 9000 is a cloud-native Tier-0 enterprise data infrastructure solution designed for simplicity, resiliency, and performance when running mission-critical applications and workloads. Built on proven resiliency and powered by the intelligence of DSCC and HPE InfoSight, HPE Alletra 9000 provides extreme latency and reliability and is delivered as a service. It ensures fast and always-on storage for all mission-critical applications. HPE Alletra 9000 is built on a massively parallel, multi-node, and all-active platform that helps consolidate traditional and next-generation mission-critical workloads at scale with predictable performance, ultra-low latency, and a 100% Availability Guarantee.

The HPE Alletra 9000 storage systems family includes HPE Alletra 9060 and HPE Alletra 9080 arrays, both available in two- and four-node configurations. The nodes in an HPE Alletra 9000 array are always active. For more information on the number of host ports and maximum storage capacities for each model, see the HPE Alletra 9000 data sheet. Figure 3 shows an HPE Alletra 9080.

FIGURE 3. HPE Alletra 9080 four-node array

The DSCC brings the cloud experience to where data lives, be it edge, core, or the cloud. It acts as the single destination to streamline data provisioning, access, monitoring, protection, and mobility from edge to cloud. The DSCC offers predefined blueprints that simplify the setup and provisioning of HPE Alletra 9080 arrays.

The HPE Alletra 9080 array that was used for performance benchmarking had the following configuration:

- OS version: 9.4.0
- 48 x 3.84 TB NVMe SSD
- Eight 32 Gb four-port Fibre Channel HBAs
- Eight 2 TB volumes per HPE ProLiant DL580 server with thin provisioning and data reduction
- LVM configured on each server with a strip size and read-ahead of 256 KB

Software

The software components used for this technical paper are:

- Red Hat Enterprise Linux 8.3
- SAS Foundation Software

Application software

The application software used for this solution includes the components described in detail in the following section.

SAS 9.4

SAS 9.4 is a highly flexible and extensible fourth-generation programming language (4GL) for data access, data transformation, analysis, and reporting. It is included with SAS Foundation. It also includes a powerful macro facility that reduces programming time and maintenance issues.

SAS 9.4 software optimizes hardware resources and scales across different infrastructure components. Numerous layered applications are available for additional functionality, including direct access to standardized data sources and advanced statistical analysis.

Prewritten SAS procedures handle many common tasks, including data manipulation and management, information storage and retrieval, descriptive statistics, and report writing. They deliver significant functionality with just a few simple commands so programmers can be more efficient and productive.[1]

SAS 9.4 workload has the following characteristics and requirements:

- Large-block I/O
- Sequential access I/O
- A minimum of 150 MB/sec/core to be supported for SASDATA, WORK, and UTILLOC file systems.[2]

SAS Mixed Analytics Workload

The SAS Mixed Analytics Workload used for this test had the following characteristics:

- 50% CPU intensive and 50% I/O intensive jobs
- SAS procedures including:
  - DATA step
  - PROC RISK, PROC LOGISTIC, PROC GLM (general linear model), PROC REQ, PROC SQL, PROC MEANS, PROC SUMMARY, PROC FREQ, and PROC SORT
- SAS program input sizes up to 50 GB per job
- Text, SAS data set, and SAS transport files input data types
- Memory use per job up to 1 GB
- Varied job run-times (short- and long-running tasks)

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

HPE Alletra 9000 volumes are provisioned to the four HPE ProLiant DL580 by servers using the HPE Alletra 9000 UI, the HPE Alletra 9000 command line, or the DSCC.

The multipathing configuration for HPE Alletra 9000 is described in a later section of this paper.

The tests found that it is optimal to have multiple virtual volumes (LUNs), typically eight, exported from the storage system to each of the servers. At each server, these eight LUNs are then aggregated into a volume group from which logical volumes are created for each of the storage areas needed for SAS. In total, there are 32 LUNs of 2 TB each created on the HPE Alletra 9080 array.

[1] Reprinted from Base SAS Software Fact Sheet: https://www.sas.com/content/dam/SAS/en_us/doc/factsheet/base-sas-101198.pdf
[2] Best Practices for Configuring Your I/O Subsystem for SAS 9 Applications: https://support.sas.com/resources/papers/proceedings16/SAS6761-2016.pdf

The SAS Mixed Analytics Workload is deployed in self-contained groups of 30 users each. This means a separate area on the disk was created for each 30-user grouping. The area on the disk is created from the same volume group as all the other groupings. By doing this, any number of users can be emulated, because they all share the same server, and they all use the same LUNs.

For each of the 30-user runs, logical volumes for SAS Data, SAS Work, and UTILLOC are created. Each of the logical volumes are striped across all eight LUNs with a stripe size of 256 KB. The Logical Volume Manager (LVM) read-ahead is also set to 256 KB by using the `lvchange` command. The LVM stripe size and read-ahead values differ for each type of storage system. This value is 256 KB for the HPE Alletra 9000 family. Hewlett Packard Enterprise also found that the best performing file system for SAS is XFS. Two scenarios were tested: first with thin provisioned volumes and second with data reduction volumes that compress and deduplicate data.

PERFORMANCE RESULTS

With SAS 9.4 running the Mixed Analytics Workload, a server is rated based on the amount of time it takes to complete the entire test, the combined read and write throughput, and the amount of time an average job takes to execute. A SAS 9.4 job demands a certain amount of processing power and I/O during its execution. The tests found that the ability to read and write data is just as important to the execution of a SAS 9.4 job as the ability of the CPU to process the data. The inability to move data effectively and rapidly to and from the processor means that the cores will have idle time, even though there are SAS 9.4 jobs that are ready to run.

The four servers underwent a series of tests. Each SAS test bundle was representative of 30 concurrent SAS 9.4 users. Each test bundle ran a total of 102 SAS jobs on behalf of those 30 representative users. With an HPE Alletra 9080 array and four HPE ProLiant DL580 servers, up to three such SAS test bundles could be run on each server (12x30 = 360 users running a total of 1224 concurrent jobs). There was a near-linear increase in the combined read and write throughput.

The environment was queried every five seconds during the tests. The results were then aggregated for the best 10 intervals (50 seconds), 24 intervals (two minutes), 60 intervals (five minutes), 120 intervals (10 minutes) and 240 intervals (20 minutes). This was done because the I/O demand starts high and trails off steadily during the duration of the tests. In fact, for several intervals during the later stages of the tests, only CPU-intensive jobs remained active and there was no I/O being performed.

For thin provisioned volumes

Figure 4 shows the average total read + write throughput for each of the job counts along with the total number of jobs for the four HPE ProLiant DL580 servers for eight 2 TB thin provisioned LUNs on the HPE Alletra 9080 array.

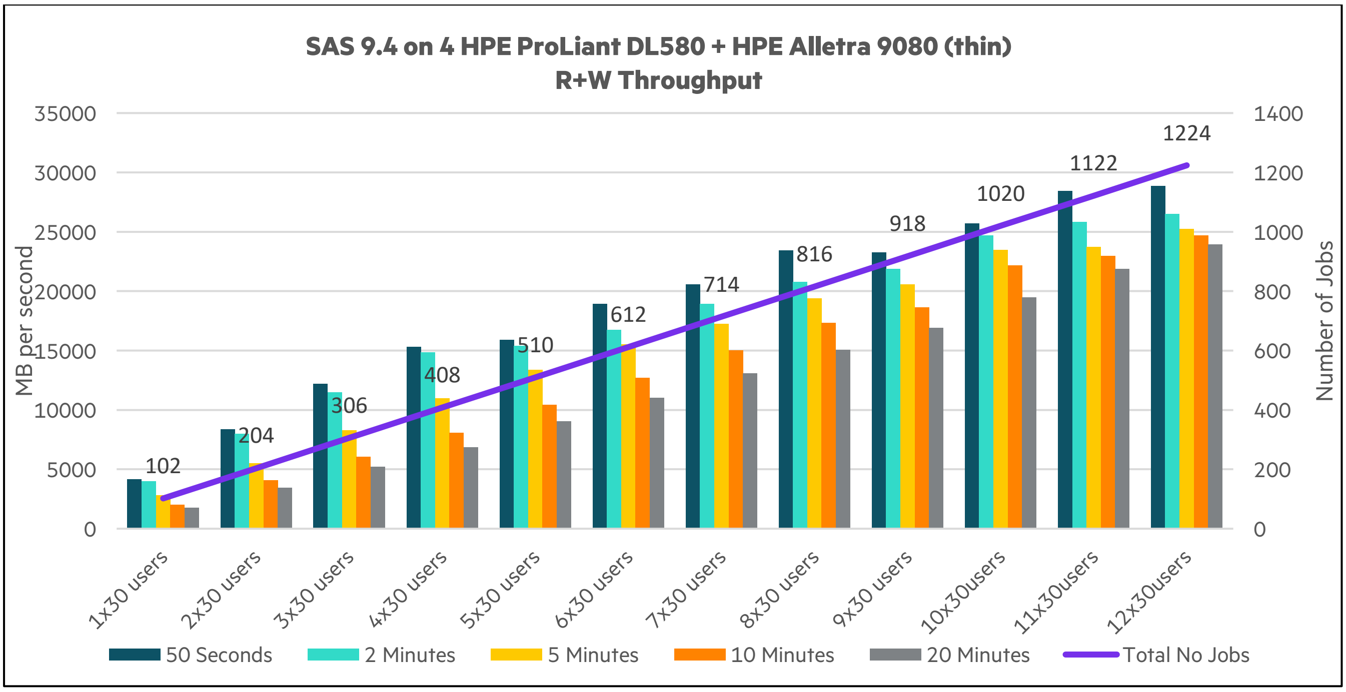

FIGURE 4. Read+write throughput at various time intervals and total number of SAS jobs for 4 x HPE ProLiant DL580 servers on HPE Alletra 9080 thin provisioned volumes

The legend tells the number of 30-user runs being executed on four HPE ProLiant DL580 servers. A legend of 8x30 users means eight sets of 30 users were executed and consisted of 8 x 102 = 816 SAS jobs.

A way to read this chart is to consider that if 300 (10x30) users are going to run up to 1020 SAS 9.4 analytics jobs for up to 10 minutes on this hardware stack, they can achieve a maximum read+write throughput of 22 GB/s. If the jobs run for up to 50 seconds, the maximum read+write throughput would be 25 GB/s and so on. The maximum read + write throughput that was achieved was 29 GB/s for 1224 concurrent jobs of duration up to 50 seconds. The maximum throughput possible is based on the number of concurrent users or concurrent jobs requirement from a customer and the duration of these jobs. Subsequently, for that throughput value, using SAS 9.4 performance guidelines, the required number of CPU cores for HPE ProLiant DL580 servers or any of their equivalent compute models and the sizing for the HPE Alletra 9080 can be determined.

Another important aspect of SAS 9.4 software is the amount of time it takes for a series of SAS jobs to complete in total and the time it takes for a single SAS job to complete. Figure 5 shows the time taken for all the jobs along with the average time it took a single SAS job to run during this scenario.

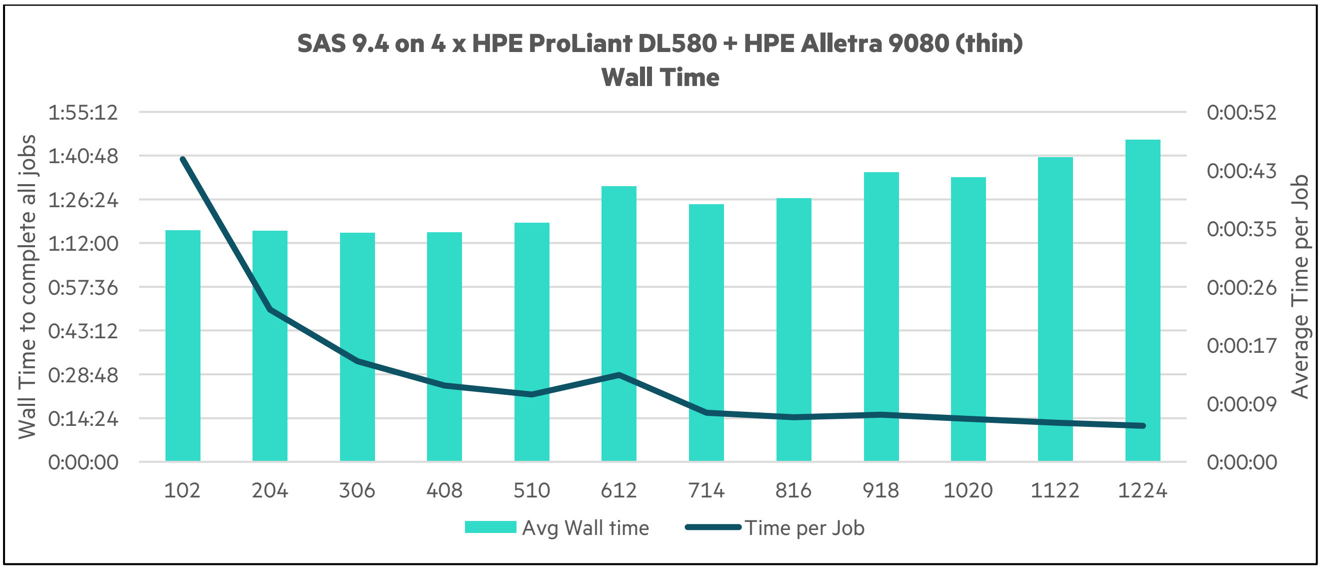

FIGURE 5. Average time recorded to complete all jobs along with the average time per job for thin provisioned volumes

For data reduction volumes

Figure 6 shows the average total read + write throughput for each of the job counts along with the total number of jobs for the four HPE ProLiant DL580 servers, each with eight 2 TB thin data reduction volumes on the HPE Alletra 9080. The maximum total read + write throughput achieved is 22 GB/s for more than 1200 jobs of up to 50 seconds duration.

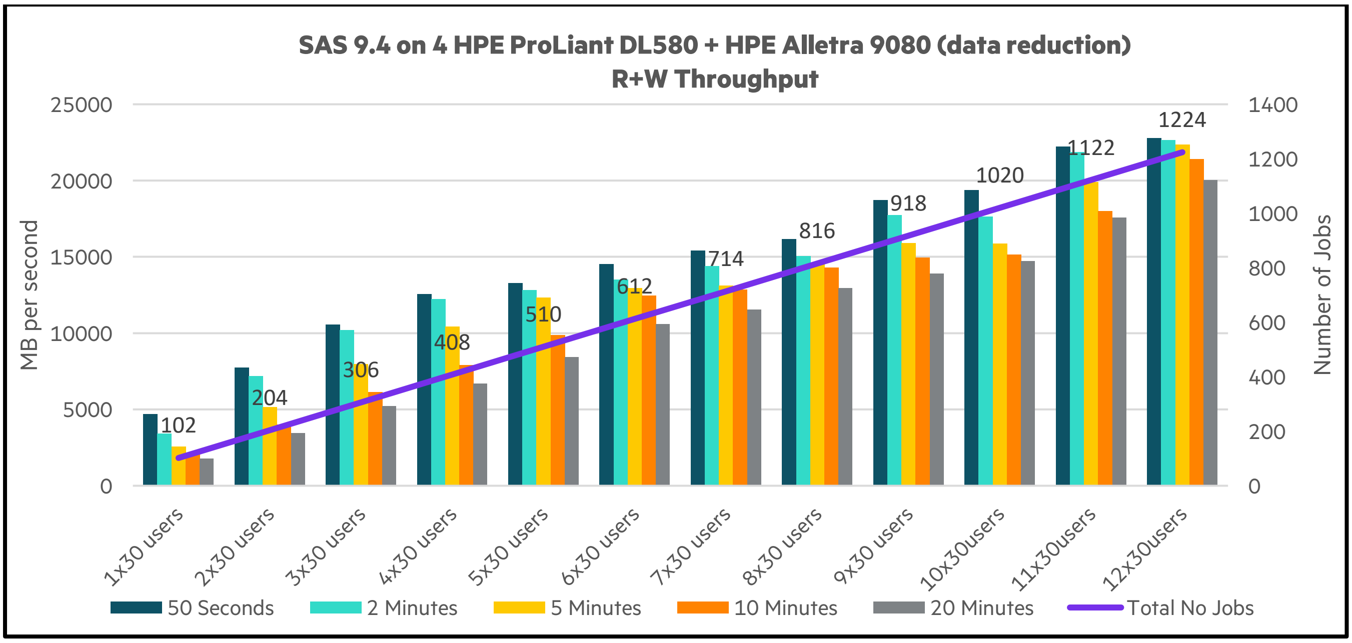

FIGURE 6. Read+write throughput at various time intervals and total number of SAS jobs for 4 x HPE ProLiant DL580 servers on HPE Alletra 9080 data reduction volumes

The data reduction volumes on an HPE Alletra 9080 array offer space saving of around 3:1 for SAS 9.4 Mixed Analytics Workload. However, the actual space savings in a customer environment depends on how compressible their data is. Figure 7 shows the time taken for all the jobs along with the average time it took a single SAS job to run during this scenario.

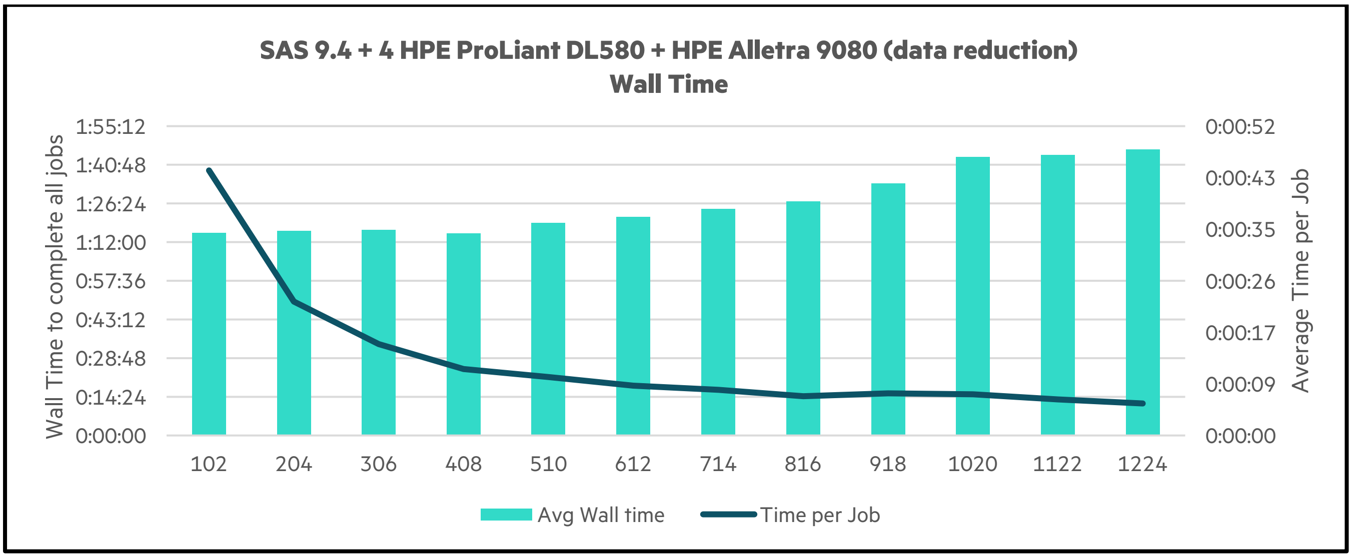

FIGURE 7. Average time recorded to complete all jobs along with the average time per job for data reduction volumes

IMPLEMENTING A PROOF-OF-CONCEPT

As a best practice for all deployments, Hewlett Packard Enterprise 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 HPE Services representative (hpe.com/us/en/services/consulting.html) or your HPE partner.

Configuration for the HPE Alletra 9000 array

Parameters used in the /etc/multipath.conf for specific HPE Alletra 9000 device configurations are as follows:

```
# The defaults section

defaults {
        polling_interval 10
}
# For Alletra 9000 device
devices {
     device {
         vendor                 "3PARdata"
         product                "VV"
         no_path_retry          0
         path_grouping_policy group_by_prio
         prio                   "alua"
         hardware_handler       "1 alua"
         path_selector          "service-time 0"
         rr_weight              uniform
         rr_min_io_rq           1
         path_checker           tur
         failback               immediate
     }
}
```

Creating volume groups

The script used to create the volume group is as follows:

```
#!/bin/bash
#Volume group creation#
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
```

Creating logical volumes

The script used to create the logical volumes is as follows. You can create multiple asuites using these commands.

```
#!/bin/bash
#logical volume creation#
lvcreate -y -L 50G -i 8 -I 262144 -n sas vgsas
lvcreate -y -L 2000G -i 8 -I 262144 -n saswork1 vgsas
lvcreate -y -L 1000G -i 8 -I 262144 -n utilloc1 vgsas
lvcreate -y -L 1000G -i 8 -I 262144 -n asuite1 vgsas

lvchange -readahead 262144 /dev/vgsas/utilloc1
lvchange -readahead 262144 /dev/vgsas/asuite1
lvchange -readahead 262144 /dev/vgsas/saswork1
```

Creating file system

The script used to create the file systems is as follows:

```
#!/bin/bash
#file system creation#
mkfs.xfs /dev/vgsa/sas
mkfs.xfs /dev/vgsas/asuite1
mkfs.xfs /dev/vgsas/utilloc1
mkfs.xfs /dev/vgsas/saswork1
```

Creating mount points

The script used to create the mount points is as follows:

```
#!/bin/bash
# Mountpoint creations #
mkdir /opt/sas9.4
mount /dev/vgsas/sas /opt/sas9.4/
mkdir /sasdata
mkdir /sasdata/asuite1
mount /dev/vgsas/asuite1 /sasdata/asuite1/
mkdir /sasdata/asuite1/saswork
mount /dev/vgsas/saswork1 /sasdata/asuite1/saswork/
mkdir /sasdata/asuite1/saswork/utilloc/
mount /dev/vgsas/utilloc1 /sasdata/asuite1/saswork/utilloc/
chmod -R 775 /sasdata
chmod 755 /opt/sas
chown sas:sas /opt/sas9.4
chown -R sas:sas /sasdata
```

/etc/fstab script for remounting the file systems after a system boot

The following code resembles the /etc/fstab file on the HPE ProLiant DL580 servers after all the logical volumes, file systems, and mount points have been created. This allows the logical volumes to automatically mount after the system is rebooted.

```
# /etc/fstab
# created by anaconda
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
/dev/mapper/rhel-root   /                       xfs     defaults        0 0
UUID=211d0774-f5ff-4bca-8ca5-279e12ed262c /boot                   xfs     defaults        0 0
UUID=CAB7-56B0          /boot/efi               vfat    umask=0077,shortname=winnt 0 0
/dev/mapper/rhel-home   /home                   xfs     defaults        0 0
/dev/mapper/rhel-swap   swap                    swap    defaults        0 0
/dev/mapper/vgsas/asuite1  /sasdata/asuite1                  xfs    defaults      0 0
/dev/mapper/vgsas/saswork1 /sasdata/asuite1/saswork          xfs    defaults      0 0
/dev/mapper/vgsas/utilloc1 /sasdata/asuite1/saswork/utilloc  xfs    defaults      0 0
/dev/mapper/vgsas/sas            /opt/sas9.4                        xfs    defaults      0 0
```

Each asuite here is a group of 30 users. To create more users, repeat these commands with asuite2, asuite3, and so forth. These asuites can be on the same server if it has enough cores or can be distributed across multiple servers.

SUMMARY

The HPE Alletra 9080 is designed to provide extreme resiliency and ultra-low latency performance for analytics workloads like SAS 9.4. Combined with the power of HPE ProLiant DL580 servers, HPE delivers a linearly scaling performance for the read + write throughput for SAS 9.4 with a maximum of 29 GB/s for thin provisioned volumes for 1224 jobs of up to 50 seconds duration. On data reduction volumes, space savings of around 3:1 and a maximum throughput of 22 GB/s for 1224 jobs of up to 50 seconds duration can be observed. This combination is a sample solution and HPE Alletra 9000 can be used with any other HPE Compute server option for a favorable and predictable performance with cloud-native management in an as-a-service model via HPE GreenLake.

Resources, contacts, or additional links

HPE servers
hpe.com/servers

HPE Technology Consulting Services
hpe.com/us/en/services/consulting.html

SAS 9.4
https://www.sas.com/en_us/software/sas9.html

Programming documentation for SAS 9.4
https://documentation.sas.com/doc/en/pgmsascdc/9.4_3.5/pgmsaswlcm/home.htm?homeOnFail

SAS 9.4 Administration
https://documentation.sas.com/doc/en/bicdc/9.4/biwlcm/home.htm

HPE ProLiant DL580 Gen9 Server QuickSpecs
https://www.hpe.com/psnow/doc/c04601208?jumpid=in_lit-psnow-red

HPE Alletra 9000
https://www.hpe.com/psnow/doc/PSN1013540069IEEN.pdf

Data Services Cloud Console
https://www.hpe.com/in/en/storage/data-services-cloud-console.html

HPE Alletra 9000 Simplified Setup with Data Services Cloud Console
https://www.youtube.com/watch?v=vjEQg-oyl5w

LEARN MORE AT

hpe.com/storage

Hewlett Packard Enterprise

a00121094ENW