TECHNICAL PAPER

Performance and Tuning Considerations for SAS® Viya® on Dell EMC® VxFlex™

Last update: January 2020

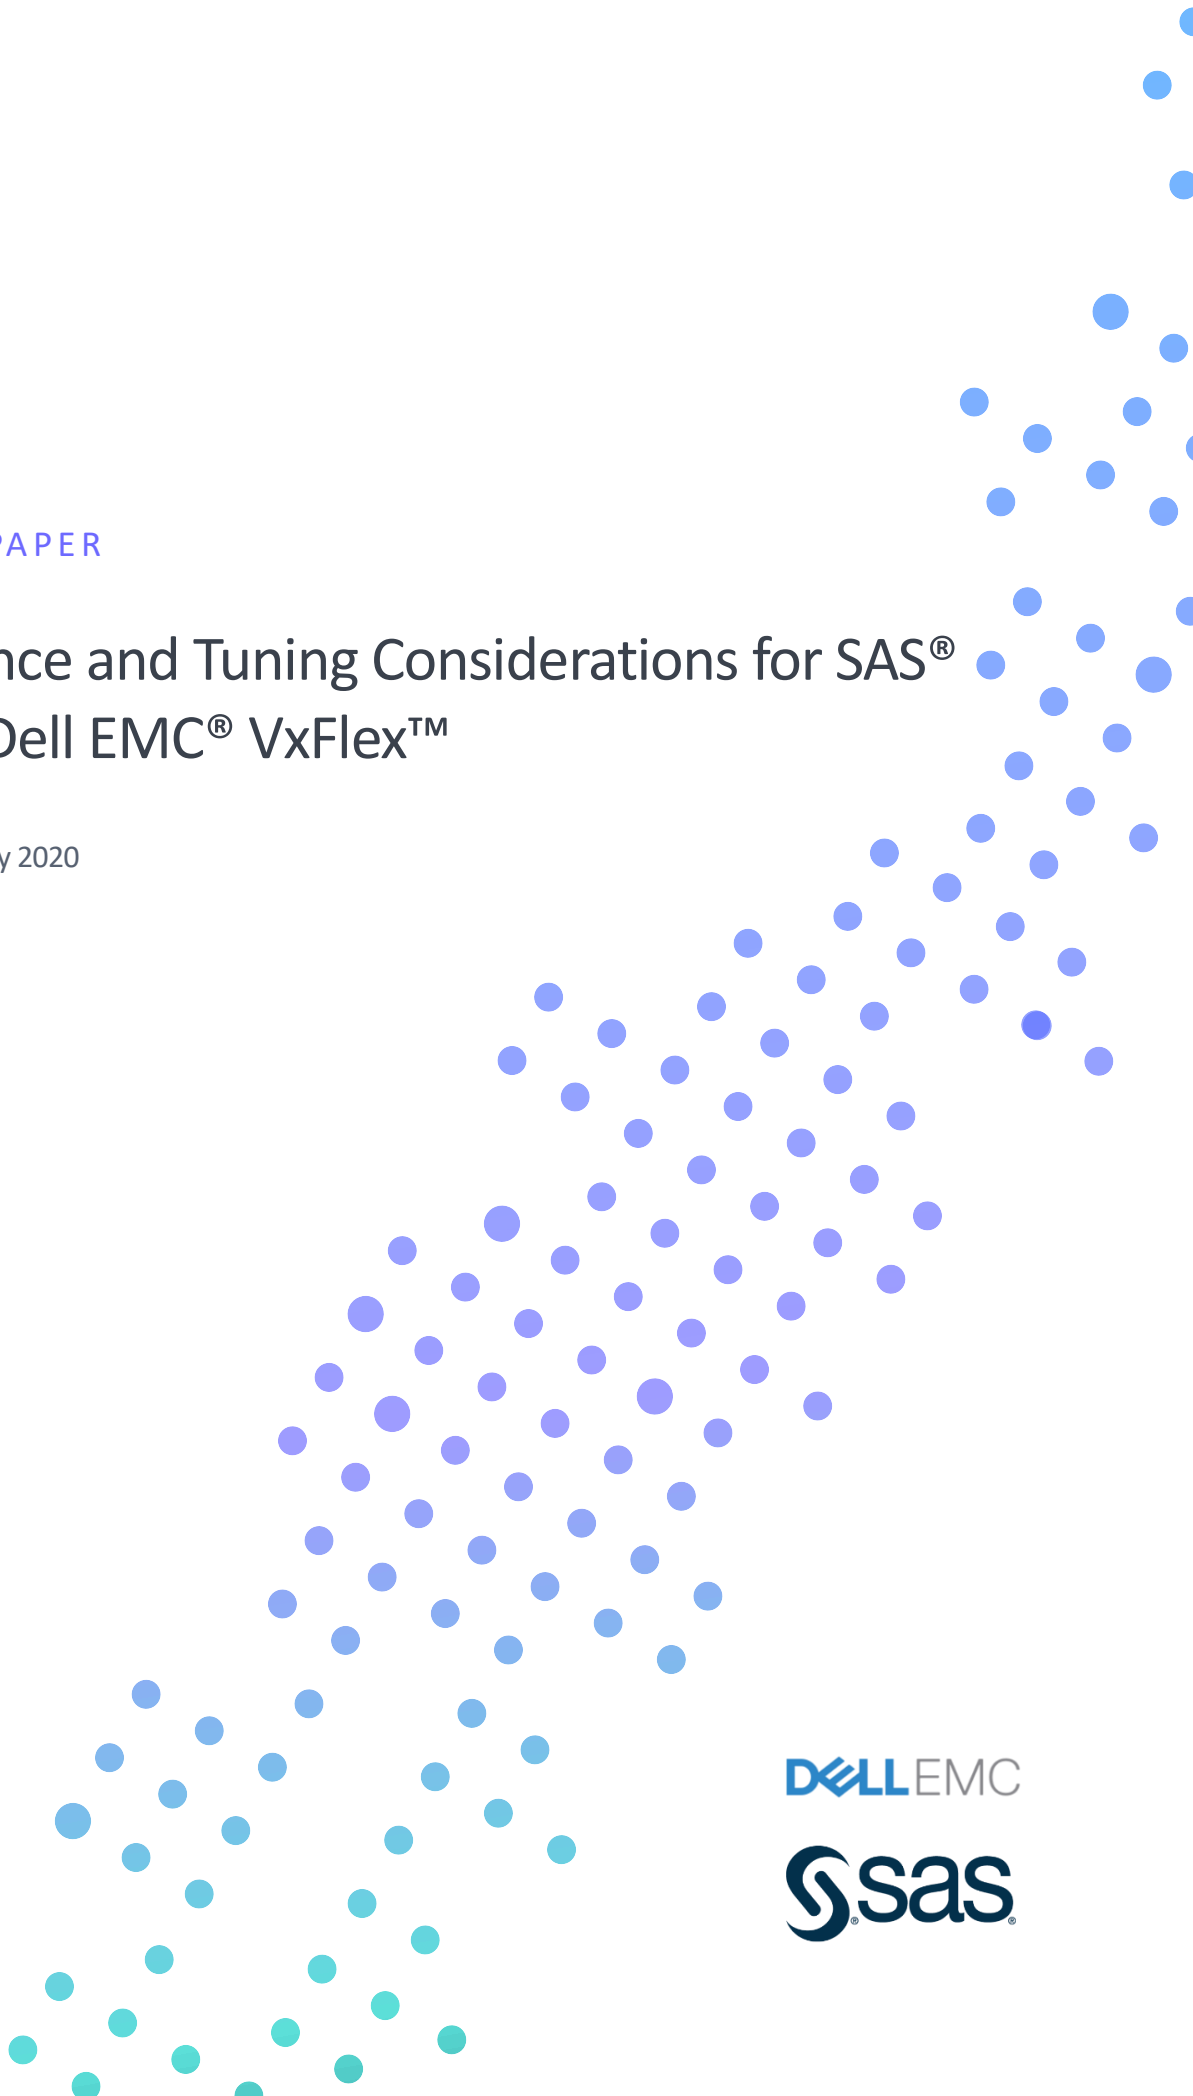

Contents

Introduction

This technical paper presents the testing results of a SAS® Viya® w orkload on the Dell EMC® VxFlex™ system. The Dell EMC VxFlex w as utilized as a massively parallel processing (MPP) architecture for the Viya system.

This effort involved testing a SAS Viya w orkload in MPP mode against six Dell R640 host nodes. The tests w ere performed to show the appliance scalability as w ell as uniformity of performance per MPP node. This paper outlines performance test results performed by SAS, and general configuration and tuning considerations for the Dell EMC VxFlex system for SAS Viya deployments.

A testing overview is discussed first. Next, detailed descriptions of the actual test bed and w orkload are provided along w ith a description of the test hardw are. Test results follow , w ith a list of tuning recommendations. Finally, general considerations, recommendations for implementation w ith SAS Viya, and conclusions are discussed.

Dell EMC VxFlex SAS Viya Performance Testing

Performance testing w as conducted w ith six Dell R640 physical host nodes running the Red Hat Linux (RHEL) 7.6 operating system. For the full system description, see “Hardw are Description” below . The purpose of the testing w as to determine w hether the VxFlex could perform w ell w ith SAS Viya in-memory heavy w orkloads. In this section of the paper, w e describe the performance tests, the hardw are used for testing and comparison, and the test results.

Test Bed Description

The test bed chosen for the flash testing w as a SAS Viya w orkload. The scaled w orkload consisted of SAS® Visual Statistics, and SAS® Visual Data Mining and Machine Learning tests to measure concurrent, mixed job performance.

The actual w orkload chosen w as composed of 21 individual Viya tests: 18 Visual Statistics, and 3 Visual Data Mining and Machine Learning tests. Each test w as composed of multiple steps, relying on existing *.SASHDAT data stores. The tests w ere chosen as a matrix of long-running and shorter-running tests (ranging in duration from approximately 9 seconds to over 13 minutes). Multiple instances of some tests w ere run to balance the w orkload flow . In all, 29 test instances w ere launched. These tests generally place a heavy CPU and memory stress on the system.

Data Used

Test programs and scripts w ere launched from Node 1. The nodes w ere used for different roles in the SAS ® Cloud Analytic Services (CAS) cluster and are identified in “Hardw are Description” below .

An aggregate of approximately 670 gigabytes of data w as used by the tests for the SAS Viya w orkload across the CAS controller node and four CAS w orker nodes. The test data w as already prepared in the *.SASHDAT file format at the start of the test run. It w as stored in a permanent file system on the non-volatile memory express (NVMe)

storage cluster, accessed by the CAS controller (Node 2) to distribute to the CAS w orkers.

The CAS_Disk_Cache w as allocated on the NVMe cluster for each CAS w orker node (Nodes 3–6).

Hardware Description

This test bed w as run against six Dell R640 host nodes using the SAS Viya w orkload w ith 29 simultaneous test instances.

VxFlex Configuration

Physical Configuration

The VxFlex infrastructure test w as composed of the follow ing:

- Machines: Six physical RHEL hosts
- Machine Configuration
  - o Node 1 – Programming & Microservices Node
  - o Node 2 – CAS controller
  - o Nodes 3–6 – CAS w orkers
- Host: Dell R640 Server
- OS: RHEL 7.6
- CPU: Tw o Intel Skylake 8168 24 cores, 2.7 gigahertz (GHz), hyper-threading-enabled Intel Xeon Platinum 8168 Processor
- Memory: 384 GB
- Storage: Eight 1.5-TB NVMe Flash devices
- Fabric: Four 25-Gigabit Ethernet (75 Gb for VxFlex OS traffic and 25 Gb for the SAS Viya netw ork usage)

SAS Software Used

The SAS softw are used w as Viya 3.4 (19w 40).

SAS Viya Test Results on Six Dell EMC VxFlex System Nodes

SAS Viya w as run in a quiet setting (no competing activity on the system) using six nodes of the Dell EMC VxFlex system. Multiple runs w ere executed to standardize the results. Table 1 show s the performance of the Dell EMC VxFlex six -node test.

Table 1. *Aggregated Jobs Run Time, Total CPU Time, Absolute Value for CPU/Real-Time Ratio Performance Using Dell EMC VxFlex*

Dell EMC VxFlex Node	CAS Actions Real Time (*hh:mm:ss.ms*) Workload Aggregate	CAS Actions Total CPU Time (*hh:mm:ss.ms*) Workload Aggregate	CPU/Real-Time Ratio
Six CAS w orker nodes	2:10:14.83	45:54:22.23	21.15

Deploying the SAS Viya w orkload across four CAS w orker nodes, 192 total cores plus hyper-threads, resulted in an aggregated w all clock run time of 2 hours, 10 minutes, and 14 seconds. This w orkload is parallel processed and is incredibly fast as compared to serial processing it w ith one serial job per CPU. The Dell EMC VxFlex system performance w ith this instantaneous parallel system demand yielded a CPU/Real-Time ratio of 21.5. This illustrates the pow er of parallel processing on very capable hardw are.

The SAS Viya w orkload is composed of statistical models that require fast CPU speed, floating point performance, memory utilization, and IO to service data coming into the parallel job chain, and service the CAS_Disk_Cache memory map backing store. Your performance w ill vary according to w orkload differences.

Performance Graphs

The follow ing graphs show RHEL performance for the cluster w ith reports on read and w rite bandw idth, and read and w rite latency. The w orkload read bandw idth peaks at 2.5 GB/sec, and the w rites at almost 11 GB/sec. This is a high bandw idth w orkload for w rite activity. The read latency is generally less than 1 ms, and at peak hits a little over 15 ms. This is incredibly low latency for the block IO pushed through the system. Write latency is generally at 3 ms or less, w ith spikes up to 16 ms. The latency figures are referentially very good for large block IO running in the 11 GB/sec bandw idth range.

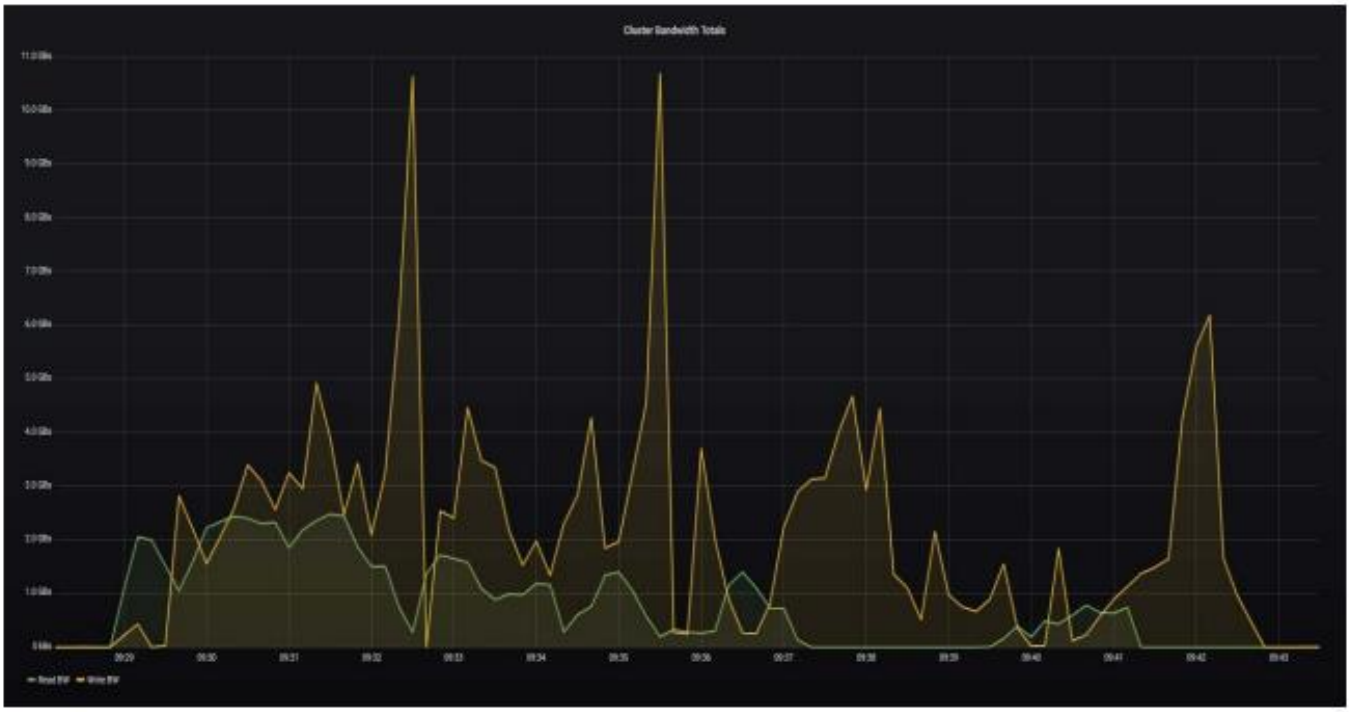

Graph 1: Read and Write Bandwidth

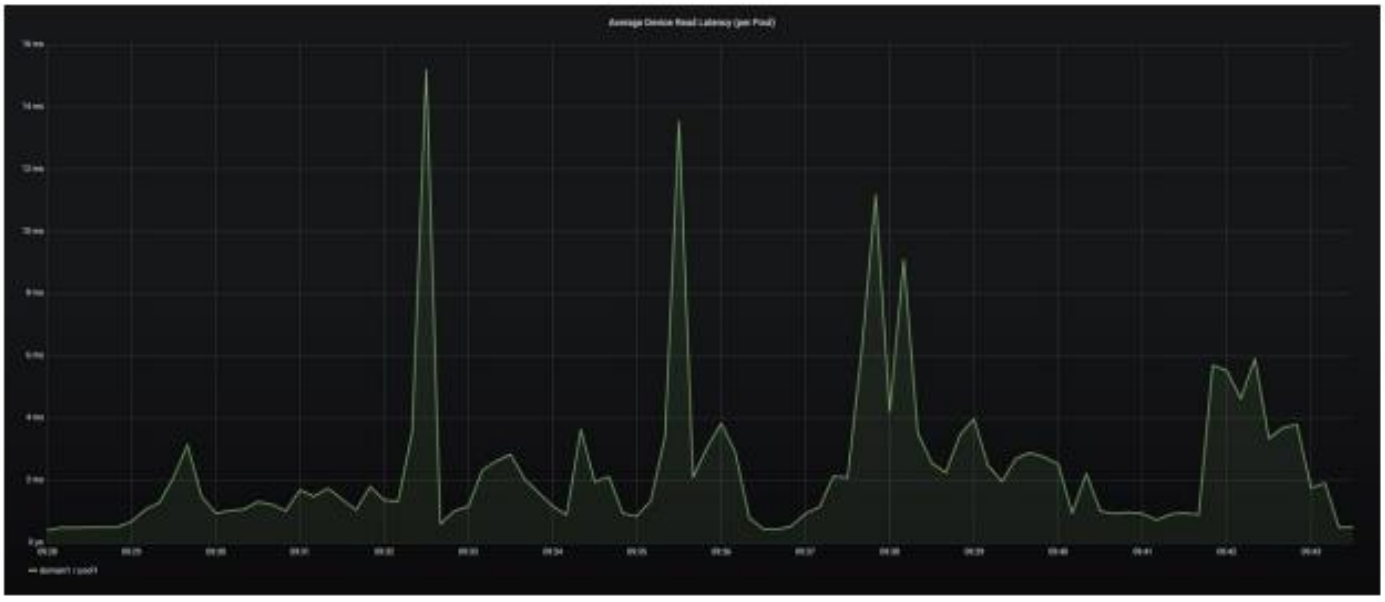

Graph 2: Read Latency

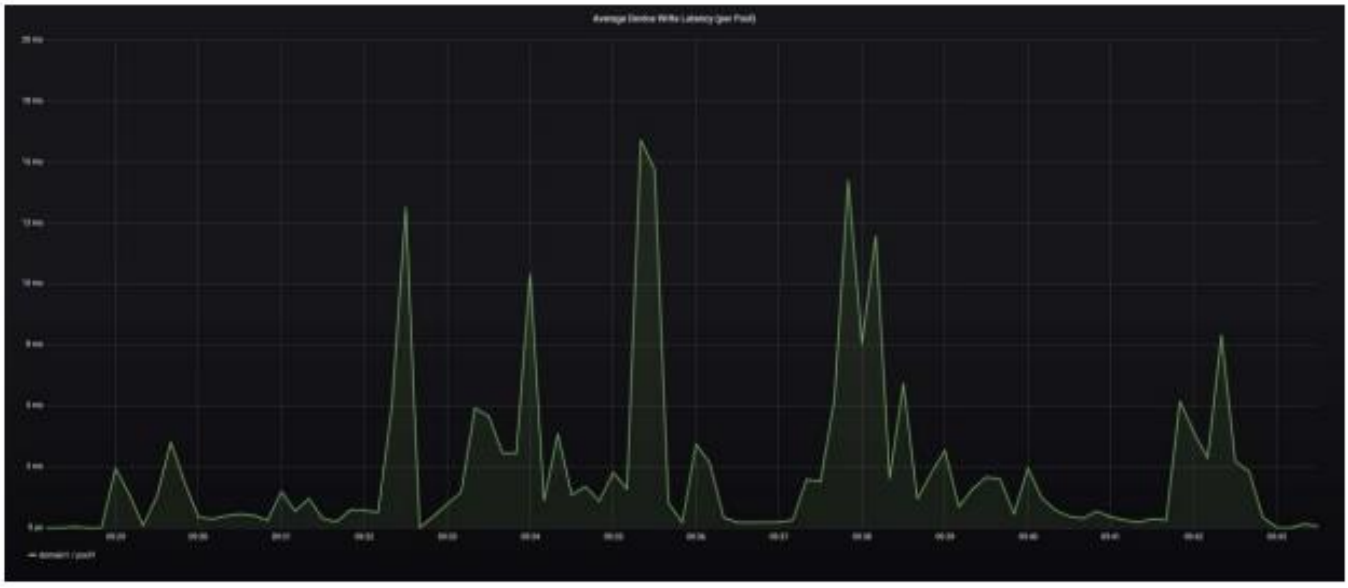

Graph 3: Write Latency

General Considerations

The Dell EMC VxFlex system has show n to deliver significant performance for an intensive SAS Viya w orkload. It is very important to use the SAS Viya performance and architectural tuning guides for your system to optimize server-side performance w ith Dell EMC VxFlex, as w ell as any additional suggestions provided by Dell EMC VxFlex engineers.

Dell EMC VxFlexTuning and Provisioning Recommendations

The follow ing settings w ere used for the Dell EMC VxFlex OS testing.

- Jumbo frames usage
    - Refer to Dell EMC VxFlex OS Networking Best Practices and Design Considerations .
- SDC/SDS/MDM set to a High-Performance Profile
    - Refer to the “VxFlex OS performance fine-tuning tasks” section of Dell EMC VxFlex OS Performance Fine-Tuning Technical Notes.
- NVMe devices w ere passed through to SVM using DirectPath IO.
- Three 25-GbE connections w ere used for data netw orks using VxFlex IP Roles.
    - Refer to Dell EMC VxFlex OS Networking Best Practices and Design Considerations .

Conclusion

The Dell EMC VxFlex system has been proven to be extremely beneficial for parallel SAS Viya w orkloads. In summary, the Intel Skylake processors enable the compute layer to perform more operations per second, w ith excellent memory handling, thus increasing the potential performance of the solution. Consistently low response times from the underlying NVMe storage layer, w ith the Flex OS softw are-defined storage system, enable CPU and memory potential to be more fully realized.

Using the Dell EMC VxFlex system is designed to be as straightforw ard as possible. It is crucial to w ork w ith your Dell EMC storage engineer to plan, install, and configure the hosts for the environment to attain maximum performance.

The guidelines listed in this paper are beneficial and recommended. Your individual experience might require additional guidance by Dell EMC and SAS engineers, depending on your host system and w orkload characteristics.

References

SAS Institute Inc. (2019). SAS Note 42917. “A list of papers useful for troubleshooting system performance problems.” Available at https://support.sas.com/kb/42/197.html

Release Information Content Version: 1.0 January 2020