WHITE PAPER

Intel Software

A Shared File System on SAS Grid Manager* in a Cloud Environment

Abstract

The performance of a shared file system is a critical component of implementing SAS Grid Manager* in a cloud environment. Intel® Cloud Edition for Lustre* (Intel® CE for Lustre*) software is purpose-built for use with the dynamic computing resources available from Amazon Web Services* (AWS*) and provides high-performance, software-defined storage on AWS using compute, storage, and I/O resources.[1,2] Based on the performance achieved on the Amazon Cloud Drive*[6], Intel CE for Lustre is being used by SAS customers to deploy SAS Grid Manager on AWS.

This paper will address:

- Selection of Amazon Elastic Compute Cloud* (EC2*) instances and configuration of the Intel CE for Lustre file systems on AWS for SAS Grid Computing applications
- Performance of SAS grid computing applications in a cloud environment with Intel CE for Lustre
- The impact of using Amazon EC2 instances with Intel CE for Lustre as a cluster file system and whether the solution provides a substantial benefit for SAS large-block, sequential IO patterns

"Based on the performance achieved on the Amazon Cloud Drive*[6], Intel® Cloud Edition for Lustre* software is being used by SAS customers to deploy SAS Grid Manager* on Amazon Web Services* (AWS*)."

Introduction

To address the performance of SAS grid computing applications in a cloud environment, testing was conducted with the Intel CE for Lustre shared file system on Amazon Web Services (AWS) using a mixed analytics SAS workload. We will begin by introducing the key components used in the test.

1.1 Introduction of Intel Cloud Edition for Lustre

Lustre is a scalable, high-performance, distributed parallel software for storage clusters. Today, the Lustre* file system is one of the most popular, powerful, and widely used shared file systems in supercomputing scenarios: 60% of the top 100 clusters in the world are currently running Lustre. A Lustre storage system consists of the following key components:[1,2]

- **Servers:** Including management servers (MGS), metadata servers (MDS), and object storage servers (OSS) to manage the names and directories in the file system and provide the file I/O service
- **Storage:** Including management target (MGT), metadata target (MDT), and object storage target (OST) to store data
- **File system clients:** Compute nodes to access the file system
- **Network:** Including Lustre Networking* (LNET*) and routers to provide the communications infrastructure required for the Lustre file system

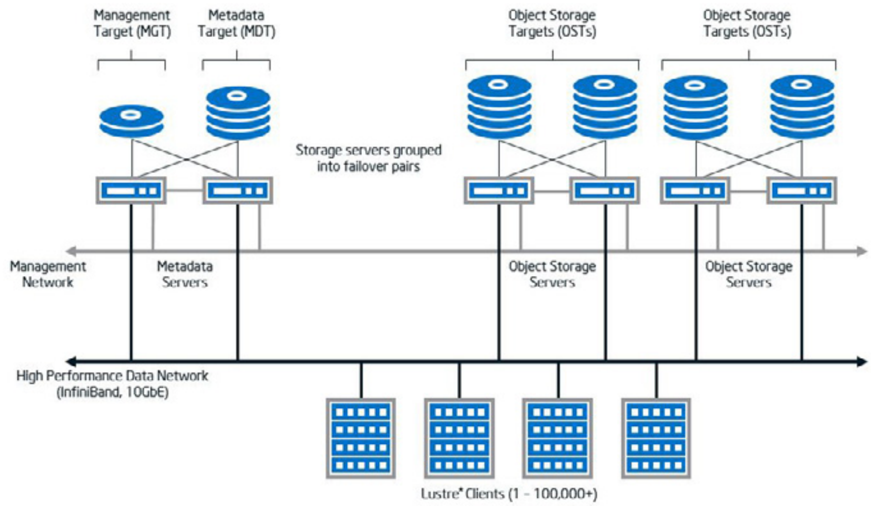

Figure 1.1. Typical Lustre* File System Configuration.

"Mixed analytics SAS workload (MASW) is a scaled workload that provides a mix of CPU, I/O, and memory-intensive SAS jobs for concurrent workload performance testing and is used to evaluate OS releases, chip technology, and storage subsystems across vendors."

With open-source Lustre software as a base, Intel CE for Lustre is optimized to orchestrate setup and configuration on AWS, which maximizes storage performance and enables rapid deployment. For example, it automatically creates metadata and object storage servers for the Lustre cluster atop EC2 using its own cloud controller. It also uses AWS CloudFormation* templates to launch Lustre clusters and adds specific features tailored for AWS, such as high-availability (HA) functionality, which automatically replaces failed instances. As a high-performance shared file system, Intel CE for Lustre has been used as the working file system for high performance computing (HPC), enterprise technical computing, and other I/O-intensive workloads.[5] Intel CE for Lustre is available through the AWS Marketplace*.[2]

1.2 Introduction of SAS

SAS is a collection of several hundred shared libraries that are dynamically loaded as needed by a separate process called an SAS job. Mixed analytics SAS workload (MASW) is a scaled workload that provides a mix of CPU, I/O, and memory-intensive SAS jobs for concurrent workload performance testing and is used to evaluate OS releases, chip technology, and storage subsystems across vendors. The three SAS file systems used in the workload exhibit a large-block, sequential access pattern as the norm.[1]

- **SASDATA:** A SAS permanent data file space. It houses persistent upstream data for SAS processes and permanently stored output data created by SAS Jobs.
- **SASWORK:** A SAS temporary workspace with 50/50 read-write (RW) pattern for all single-threaded operations. Since extreme data amplification can occur here, we find the heaviest use of all the file systems. The files do not persist if the processes terminate normally.
- **UTILLOC:** A SAS utility data file space. It behaves the same way as SASWORK, except it is for threaded procedure scratch space only.

The actual workload chosen for this test is composed of 19 different types of SAS tests: 10 computational, two memory-intensive, and seven I/O-intensive tests, designed to achieve the simultaneous sessions of 30 SAS users via launching 101 tests in total.^

1.3 Introduction of Amazon Web Services (AWS)

AWS is a comprehensive, evolving cloud computing platform provided by Amazon.com and managed by in-house infrastructure experts. It provides the flexibility for users to access technology resources on demand and launch applications regardless of the use case or industry. Selecting the right Amazon

^ Ephemeral storage is not recommended for permanent storage by systems that do not have a built-in replication scheme, such as HDFS. See Amazon EBS Volume Performance on Linux* Instances for more information.

EC2 instances for required testing is key to building a high-performance solution at a reasonable cost.

- **AWS Storage Options:** AWS has built different storage systems for various use cases. For example, ephemeral storage is direct-attached local storage. Even though it can be the fastest option for some of the Amazon EC2 instances, it has limitations on size and will disappear when the instance reboots or terminates. Amazon Elastic Block Storage* (Amazon EBS*) is networked storage with a maximum size of 16TB per EBS volume. It is persistent and can outlive the instance to which it is attached. EBS volumes also offer provisioned IOPS (100 to 10,000 IOPS). Amazon Simple Storage Service* (Amazon S3*) can help provide developers and IT teams with secure, durable, highly scalable object storage, which is recommended for long-term and durable data storage, and can be used to import data into the Intel CE for Lustre file system.
- **Amazon Elastic Compute Cloud (Amazon EC2) Instances:** Amazon EC2 is a service which can deliver scalable compute, memory, and network resources and help make Web-scale cloud computing easier for developers. AWS has optimized certain instances to meet the performance requirements of various workloads (see the architectural differences between Amazon EC2 instances in **Table 2.1**).

Experimental Framework

Next, we'll describe how to select Amazon EC2 instance types and design the system infrastructure to maximize I/O performance for mixed analytic SAS workload testing, based on the requirements of running applications.

2.1 Requirements of Mixed Analytic SAS Workload (MASW)

- **Data and I/O Throughput:** MASW has a high demand for I/O throughput and storage capacity. In general, the workload requires a minimum of 100MB/sec/core of I/O for temporary storage, and 50-75MB/sec/core for shared storage. As a scaled workload, the required storage capacity of MASW depends on the number of defined sessions of workload.

"Amazon Web Services (AWS) has built different storage systems for various use cases."

AWS INSTANCE TYPE	COMPUTE OPTIMIZED C4	STORAGE OPTIMIZED D2	GENERAL PURPOSE M4	MEMORY OPTIMIZED R3	IO OPTIMIZED I2	GRAPHICS OPTIMIZED G2	GRAPHICS OPTIMIZED G2
Intel® Procecessor	Custom Intel® Xeon® ES-2666 v3	Custom Intel® Xeon® ES-2676 v3	Custom Intel® Xeon® ES-2676 v3	Intel® Xeon® ES-2670 v2	Intel® Xeon® ES-2670 v2	Intel® Xeon® ES-2670	Intel® Xeon® Family
Intel® Process Technology	22nm Haswell	22nm Haswell	22nm Haswell	22nm Ivy Bridge	22nm Ivy Bridge	32nm Sandy Bridge	Various
Intel® AVX	AVX 2.0	AVX 2.0	AVX 2.0	Yes	Yes	Yes	Yes
Intel® AES-NI	Yes	Yes	Yes	Yes	Yes	Yes	No
Intel® Turbo Boost	Yes	Yes	Yes	Yes	Yes	No	Yes
SSD Storage	EBS-OPT by default	No	EBS-OPT by default	Yes	Yes	Yes	EBS Only

Table 2.1. Architecture overview of Amazon EC2* instances

- **Networking:** When processing enormous amounts of input, temporary, output, and cached data, SAS applications and solutions typically have a high demand for network bandwidth. When gigabytes of data are shared across the network, SAS data requires adequate bandwidth between grid nodes and file systems to maintain the necessary I/O throughput. Each of the fundamental tasks distributed across the grid must have access to all the required data regardless of where the tasks are executed. The compute tasks that require substantial data movement demand adequate bandwidth to write and fetch the data and achieve highest levels of efficiency.[1]

2.2 High-Performance Solution for MASW Testing on AWS

Key Components

Table 3.1 lists the selected Amazon EC2 instances for the test. These instances use Intel® Xeon® E5-2670 v2 or E5-2680 v2 processors with Red Hat Enterprise Linux* 6.6. (For details, see: aws.amazon.com/ec2/instance-types/.)

Amazon EC2 Instances for SAS Grid: The storage-optimized Amazon EC2 instances (I2 and D2) are the recommended choice for SAS grid nodes, because they are able to provide the best network performance to shared storage of any instance size, assuming minimal Amazon EBS traffic. In this test, i2.8xlarge instances are selected to configure the SAS grid cluster, because they can fully utilize the 10 gigabits of network bandwidth to access the shared Intel CE for Lustre file system.

Amazon EC2 Instances for Intel CE for Lustre: The compute-optimized Amazon EC2 instances are the recommended choice for Intel CE for Lustre servers, because they provide the performance and network bandwidth to maximize I/O performance on the Intel CE for Lustre servers. In this case, Amazon EC2 C3 instances are selected for MGS, MDS, and OSS.

SAS Grid Storage: The storage system is composed of ephemeral (attached local storage) and persistent storage.

- Each SAS grid node has 8 x 800GB solid-state drive (SSD) volumes of ephemeral storage.

FUNCTION	NAME	INSTANCE TYPE	MODEL	VCPU	MEM (GIB)
SAS Server	SAS grid node1	Storage Optimized	I2.8xlarge	32	244
	SAS grid node2	Storage Optimized	I2.8xlarge	32	244
	SAS grid node3	Storage Optimized	I2.8xlarge	32	244
	SAS grid node4	Storage Optimized	I2.8xlarge	32	244
Lustre* Nodes	MGS	Compute Optimized	c3.xlarge	4	7.5
	MDS	Compute Optimized	c3.4xlarge	16	30
	OSS00	Compute Optimized	c3.8xlarge	32	60
	OSS01	Compute Optimized	c3.8xlarge	32	60
	OSS02	Compute Optimized	c3.8xlarge	32	60
	NAT	General Purpose	m3.medimum	1	3.75

Table 3.1. List of Amazon EC2* instances used in the MASW test

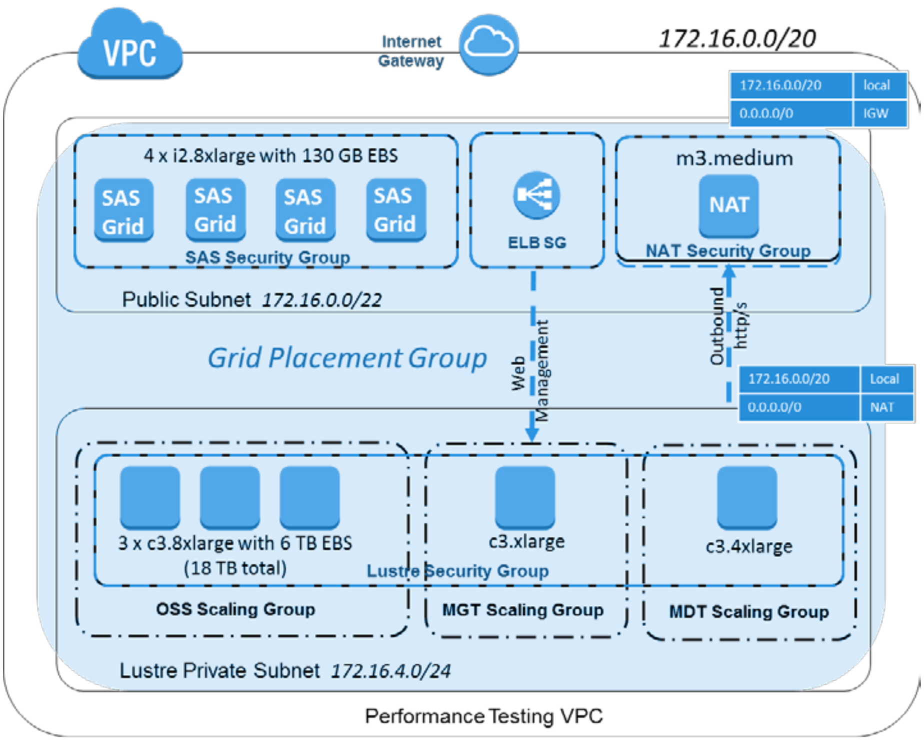

Figure 3.1. Designed system infrastructure of the test

- The Lustre file system consists of several volumes of persistent storage, including a management server target (MGT), a metadata server target (MDT), and three objective storage targets (OST) that serve the data for the file system. Each OST has 6TB of EBS general purpose SSD storage (also called gp2).

Network: A 10Gb/per/sec network connection was used for the traffic to the EBS volumes and SAS clients.

2.3 System Infrastructure

Figure 3.1 shows the system infrastructure used for testing, designed to maximize I/O throughput.

The SAS Grid Manager testing was run inside Amazon Virtual Private Cloud* (Amazon VPC*), which was configured with a public and private subnet. The VPC provides an isolation boundary fully controlled by customers using one subnet with a route to define the Internet (public subnet) and a second subnet to route within the internal network (Lustre private subnet). When the instances in the Lustre private subnet require outbound public access to interact with AWS APIs, the outbound http/https traffic is routed through the NAT instance in the public subnet, which isolates the private subnet from any public Internet traffic. Some capabilities available in Intel CE for Lustre, including high availability, require the use of these AWS API services. Even though the SAS client instances are deployed to the public subnet in this test, they could also be deployed to a private subnet in a production setting.

All instances of the test were deployed within a placement group to span subnets in the same availability zone within a VPC.[1]

SAS SG		ELB SG		NAT SG		LUSTRE* SG	
ALL	SASSG	80	0.0.0.0/0	22	0.0.0.0/0	ALL	Lustre SG
22	0.0.0.0/0			80	172.16.4.0/24	22	0.0.0.0/0
				443	172.16.4.0/24	80	0.0.0.0/0
						988	0.0.0.0/0

Figure 3.2. Port numbers and IP addresses of security groups[1].

"Security groups were used to restrict outbound traffic and inbound traffic, acting as stateful firewalls."

The Intel CE for Lustre cluster was configured with a management server (MGS) node, a metadata server (MDS) node, three object storage server (OSS) nodes, and a network address translation (NAT) node. Intel Cloud Edition for Lustre software, located on Amazon Web Services Marketplace, was subscribed, installed, and configured on it.

The security groups were divided into the SAS security group, ELB security group (ELB SG), NAT security (NAT SG) group, and the Lustre security group (Lustre SG). Security groups were used to restrict outbound traffic and inbound traffic, acting as stateful firewalls. SAS SG provides access rules to SAS instances. ELB SG defines rules for users to access the monitoring website served from the management server (MGS). NAT SG permits inbound traffic on port 80 and 443 from the Lustre private subnet IP address range, and port 22 from any address. Lustre SG protects all servers in the Intel CE for Lustre cluster.[1]

Experiments

Now we'll look at how to set up the testing environment and configure Intel CE for Lustre file systems, SAS Grid Manager, and mixed analytics SAS workload (MASW) on AWS to perform the experiments outlined.

3.1 Set up the testing environment[11]

- Open the Amazon EC2 console at console.aws.amazon.com/ec2/
- Create an AMI for the selected instances: i2.8xlarge, c3.xlarge, c3.4xlarge, c3.8xlarge, and m3.medium (see Launching Instances into a Placement Group)
  - From the Amazon EC2 dashboard, launch the instance
  - Connect to the instance. (see Connect to Your Linux Instance)
  - Install software and applications on the instance, copy data, or attach additional Amazon EBS volumes
  - Optional: Enable enhanced networking on Linux Instances in an Amazon Virtual private cloud (VPC)
  - Create image
- Create a grid placement group (a logical grouping of instances within a single availability zone)
- Launch the instances with the created AMI into the placement group
  - Launch 4 x i2.8xlarge instances for SAS grid cluster
- Configure the network of the placement group[1]
  - Create public subnet with IP: 172.16.0.0/22 for SAS grid cluster, ELB storage and NAT
  - Create private subnet with IP 172.16.4.0/24 for Intel CE for Lustre

3.2 Configure Intel CE for Lustre on AWS

Since Intel provides different options for launching Lustre using CloudFormation, the fully supported HVM version of Intel CE for Lustre was used for the SAS test on AWS.1 While the templates themselves are not region-specific, each template is required to be stored in the same region in which it will be used, so a corresponding CloudFormation template should be selected for the region in which it launched. Before creating instances for Intel CE for Lustre servers, you have to accept the terms on the Intel CE for Lustre Software AWS Marketplace page. During the initial process steps, you are able to input parameters such as the name of your secure key, placement group name, VPC, etc. Then, you have to modify the default template to include the use of placement groups and create c3.8xlarge OSS instances, etc.[1, 2]

For example, you can take the following steps to create c3.8xlarge OSS instances (see details at support.sas.com/rnd/scalability/grid/SGMonAWS.pdf).

- Locate "OSSLaunchConfig"
- Update Properties = InstanceType to c3.8xlarge
- Remove Properties = EbsOptimized (does not apply to c3.8xlarge)

At this stage, you can also specify the number of OSS nodes required and volume size to configure the storage system for a specific test. In this test, the template uses a default of 8 volumes per OSS, which sets the OST volume size to 750GB and will attach 6TB of storage per OSS. After completing changes to the template, you have to save the "new template" file and use it during the build process.

- Launch 3 x c3.8xlarge instances for OSS of Intel CE for Lustre
- Launch 1 x c3.4xlarge instance for MDS of Intel CE for Lustre
- Launch 1 x c3.xlarge instance for MGS of Intel CE for Lustre
- Launch 1 x m3.medium instance for NAT

3.3 Configure SAS Grid Manager

- Install Red Hat Enterprise Linux* OS (RHEL*) in each grid node of the SAS grid infrastructure, which consists of 4 i2.8xlarge EC2 instances as grid nodes (See **Figure 3.1**)
- Create "sas" user accounts at each grid node
- Preinstall SAS grid onto an instance
- Create an Amazon machine image (AMI)
- Replace the mapping for the AWSRHEL6AMI with the AMI ID of the created AMI
- Install SAS and reuse it for all of your instances

3.4 Configure MASW on AWS

- Upload the MASW on the shared Intel CE for Lustre file system
- The data of MASW has to be classified as "permanent" stored in SASDATA file system and "temporary" stored in SASWORK and UTILLOC file systems. Typically, temporary storage has the most demanding load. The file systems of MASW have been configured in two different ways for performance evaluation on AWS with Intel CE for Lustre.

Configuration 1): All data for MASW is located on the shared Intel CE for Lustre file system on Amazon EBS, which consists of 8 x 750GB General Purpose (SSD) volumes per OSS for the test.

- SASDATA – 2 TB
- SASWORK – 3 TB
- UTILLOC – 3 TB

Configuration 2): Permanent data for MASW is located on the shared Intel CE for Lustre file system, but temporary data is located on local file system of SAS nodes

Experimental Results

Table 4.1 shows the results of the

NO. OF TEST CASE	LOCATION OF TEMPORARY FILE SYSTEMS (SASWORK & UTILLOC)	SAS REAL TIME (MIN)	USER CPU TIME (MIN)	SYS CPU TIME (MIN)	USED MEMORY (GB)
1	Amazon Intel® CE Lustre* File System	1,493	1,094	434	192.6
2	Local EXT4 File Systems-Single SMP Host	1,481	1,750	275	92.4

Table 4.1. Results of the SAS workload test

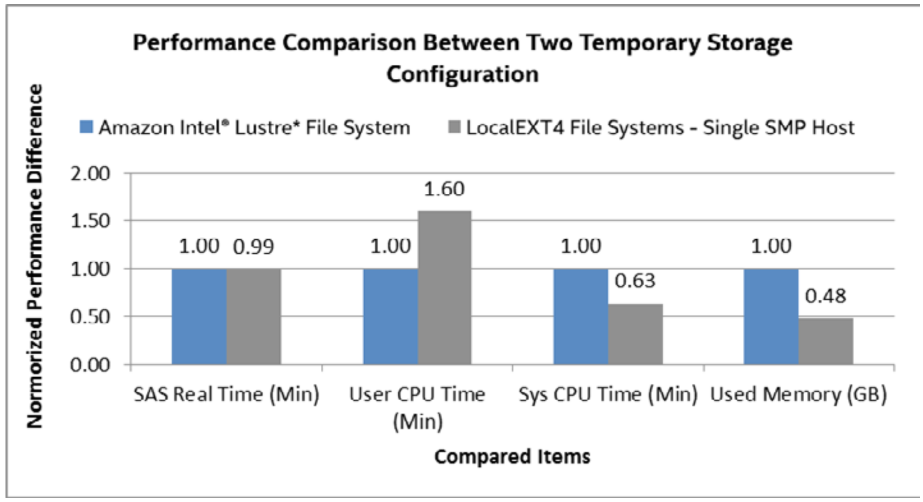

Figure 4.1. Performance comparisons between tests

MASW test with Intel CE for Lustre on AWS.

For each test, the ratio of (user CPU time + sys CPU time)/(SAS real time) can be used to determine how well SAS grid computing applications perform in a cloud environment with Intel CE for Lustre. For a single threading application, the ideal ratio is 1, as the CPU spends no time waiting on resources (usually I/O). In this test, the ratio is 0.95 for test case 1 and 1.15 for test case 2 because of threaded SAS procedures. For an I/O intensive application, the ratio of 0.95 or 1.15 indicates good efficiency for the CPU for service (because it spends little time waiting on resources).

This proves that AWS EC2 instances using Intel CE for Lustre as a cluster file system are able to yield substantial benefit for SAS large-block, sequential I/O patterns.

Furthermore, comparing the results of test cases 1 and 2, you find that by changing the location of the SAS temporary files from the Lustre file system to Local EXT4 file systems (See **Figure 4.1**):

- User CPU time increased 60% (497 minutes)
- Sys CPU time decreased 37% (159 minutes)
- Used memory capacity reduced 52% (100.2GB)
- The aggregate real time across all 101 jobs was 12 minutes less

We conclude that SAS Grid Manager can benefit from a fast local storage system to store temporary SAS data.

Tips for Performance Tuning

Next, we'll demonstrate how to make SAS grid applications run faster with Intel CE for Lustre on AWS using the following tuning steps.

RHEL OS Tuning

- Use tuning tool with throughput-performance profile.
- Use striped LVM volumes.
- Use XFS file system.
- Change some parameters from their default values (see details at: support.sas.com/resources/papers/proceedings11/72480_RHEL6_Tuning_Tips.pdf).

SAS Grid Manager Tuning:

- Use direct-attached local SSD storage systems for temporary SAS data.
- Select the latest storage-optimized Amazon EC2 instances, such as c4.8xlarge, as SAS nodes.
- Install all SAS binaries and configuration files to a shared file system to minimize administrative tasks such as installation, SETINIT maintenance, etc.
- Install Load Sharing Facility

(LSF), which is used by SAS Grid Manager to schedule SAS tasks, in a persistent shared file system.

Intel CE for Lustre Tuning:

- Increasing the number of OSSs can significantly improve the I/O throughput because both read and write performance can near-linearly scale up with the number of OSS Amazon EC2 instances in the cluster (see the test results shown in **Figure 5.1**).[4]
- Intel CE for Lustre has been tuned using the following steps before running the tests[2] for performance.

1. Increase "readcache_max_filesize": This controls the maximum size of a file that both the read cache and write-through cache will try to keep in memory. If files are larger than readcache_max_filesize, they will not be kept in cache for either reads or writes. Because the SAS application has large-block, sequential I/O patterns, large readcache_max_filesize works well. In this test, readcache_max_filesize was set to 2M on each OSS via the command: lctl set_param obdfilter.*.readcache_max_filesize=2M
2. Increase "max_dirty_mb": This defines the amount of MBs of dirty data that can be written and queued up on the client. A good rule of thumb is 4x the value of max_rpcs_in_flight. The default value is 32MB. It was increased to 256 MB in this test via the command: lctl set_param osc.*.max_dirty_mb=256
3. Increase "max RPCs (remote procedure calls) in flight": This determines the maximum number of concurrent RPCs in flight from clients. The default value is 8. It was increased to 16 in this test with the command: lctl set_param osc.*.max_rpcs_in_flight=16

- Lustre performance may be further tuned[3] according to the experiments of Torben Kling-Peterson on HPC storage systems[3] as follows:

1. Disable Network Checksums: The default is turned on and it impacts performance
2. Disable LUR size (the number of client-side locks in an LRU queue), as well as enable write-through cache and read cache

Note: You need to mount the Lustre file systems with the –o flock option on all client nodes for SAS applications in order to activate the POSIX global locking system.

Conclusion

To build a high-performance solution with a reasonable cost on AWS, the first and most important step is to select the right Amazon EC2 instances. Understanding the requirements of the running application and knowing the properties of each Amazon EC2 instance provides a guideline for selecting the right approach.

The test results indicate that a properly configured Amazon AWS infrastructure with robust I/O and Intel CE for Lustre file system can work nicely with SAS Grid Manager. They prove that Amazon EC2 instances using Intel CE for Lustre as a clustered file system are able to yield substantial benefits for SAS large-block, sequential I/O patterns.

SAS grid applications can run faster with Intel CE for Lustre on AWS via OS tuning, SAS application tuning, and Intel CE for Lustre Tuning.

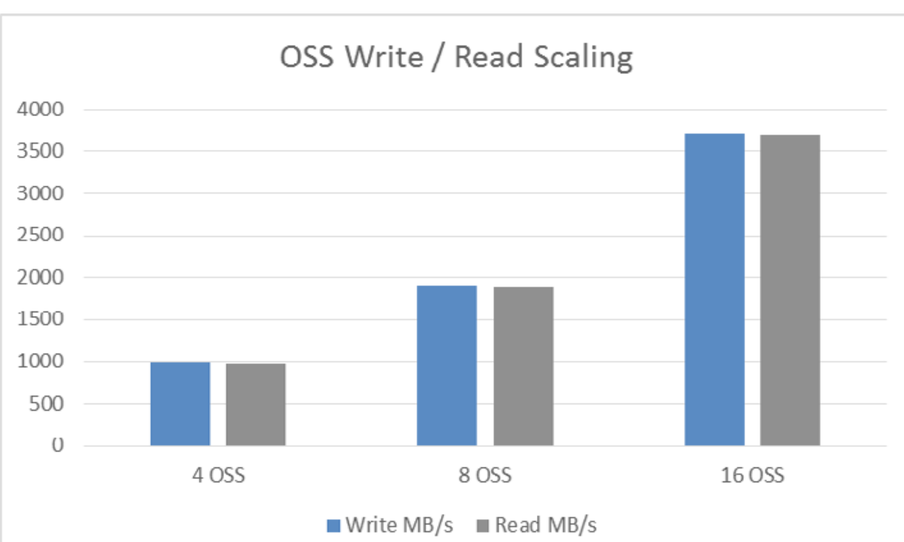

Figure 5.1. The I/O performance scalability of Intel® CE for Lustre* with OSSs

^^ The tuning values used in this testing are referenced from previous SAS tests with Intel® Enterprise Edition of Lustre* in a cluster environment. The configuration used in these tests is available in section 3.1.

References

1. "Performance and Tuning Considerations for SAS® Grid® Manager 9.4 on Amazon (AWS) Cloud Using Intel® Cloud Edition for Lustre File System," support.sas.com/rnd/scalability/grid/SGMonAWS.pdf
2. "Developing High-Performance, Scalable, Cost-Effective Storage Solutions with Intel® Cloud Edition for Lustre* Software and Amazon Web Services," intel.com/content/dam/www/public/us/en/documents/reference-architectures/ICEL-RA.pdf the number of client-side locks in an LRU queue
3. "Optimizing Performance of HPC Storage Systems," Torben Kling-Petersen, John Fragalla, ieee-hpec.org/2013/index_htm_files/17-Optimizing-Performance-KlingPetersen-2875157.pdf
4. "Developing High-Performance, Scalable, Cost-Effective Storage Solutions with Intel® Cloud Edition for Lustre* Software and Amazon Web Services," intel.com/content/dam/www/public/us/en/documents/reference-architectures/ICEL-RA.pdf
5. "Intel® Solutions for Lustre* Software," intel.com/content/www/us/en/software/intel-cloud-edition-for-lustre-software.html
6. "Lustre* Running Better Than Expected on AWS," Timothy Prickett Morgan, December 6, 2013, enterprisetech.com/2013/12/06/lustre-running-better-expected-aws
7. "Cloud AWS Infrastructure Vs. Physical Infrastructure," July 2010, highscalability.com/blog/2010/7/8/cloud-aws-infrastructure-vs-physical-infrastructure.html
8. "Amazon EC2 Instances," aws.amazon.com/ec2/instance-types/
9. "Running a High-Performance SAS Grid Manager Cluster on AWS with Intel® Cloud Edition for Lustre*, Chris Keyser, Margaret Crevar, April 2015, blogs.aws.amazon.com/bigdata/post/Tx1DSZMORS09PQE/Running-a-High-Performance-SAS-Grid-Manager-Cluster-on-AWS-with-Intel-Cloud-Edit
10. "Architecting a High-Performance Storage System," January 2014, intel.com/content/dam/www/public/us/en/documents/white-papers/architecting-lustre-storage-white-paper.pdf
11. "Placement Groups," docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

Your comments and questions are valued and encouraged. Contact us at: ying.m.zhang@intel.com and margaret.crevar@sas.com.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2®, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.