# INTEGRATED GRID SOLUTIONS FROM SAS, EMC ISILON AND GREENPLUM

sas THE POWER TO KNOW.

EMC PERSPECTIVE

EMC²

## INTRODUCTION

Intensifying competitive pressure and vast growth in the capabilities of analytic computing platforms are driving companies in all industries to strengthen their data analytics capabilities and improve the availability, performance, flexibility, and scalability of the infrastructure that supports it. Organizations looking to modernize their existing data analytics architectures and to deploy more robust IT infrastructure to support it are increasingly looking to SAS and EMC for best-in-class, integrated solutions.

To support the needs of large-scale analytics deployments, EMC and SAS have introduced a new solution that accelerates SAS® applications development, testing and deployment by integrating SAS Grid Computing with the EMC Greenplum Data Computing Appliance (DCA), Greenplum Database and EMC Isilon scale-out NAS.

## SAS GRID COMPUTING

SAS Grid Computing enables the creation of a managed, shared compute cluster for processing large volumes of data and analytic programs quickly using dynamic, resource-based load balancing. Individual SAS jobs are split into many smaller tasks that are then run in parallel across multiple CPUs sharing physical storage. This enables IT to build and manage a consolidated, high-performance yet flexible infrastructure at lower cost, while retaining the ability to easily scale to grow with changes in workloads.

## GREENPLUM DATABASE

Greenplum Database is the industry-leading massively parallel processing (MPP) database, regarded as the most scalable mission-critical analytical database available. The Greenplum Database architecture provides automatic parallelization of data loading tasks and query workloads. For SAS applications, it represents feature-rich analytical RDBMS optimized specifically for the types of workloads generated by analytical applications.

Unlike general-purpose databases, Greenplum Database delivers lightning-fast query processing for analytics and offers in-database processing that works with SAS In-Database technology to execute complex calculations within the database. By doing so, massive data can be analyzed quickly withou moving it, using a large number of processors can divide work into many smaller tasks to be accomplished in parallel. In-database analytics makes possible fast analytics on large volumes of data commonly encountered in statistics, predictive modeling and scoring operations.

SAS In-Database technology uses the power of Greenplum Database to run performance-intensive SAS procs directly in the database. Complex analytical and scoring operations run in-database improves performance, often by orders of magnitude, and eliminates very large data movements that would be required to move large data sets to the SAS Grid Computing nodes if the computation were run there. Greenplum Database and SAS Scoring Accelerator bring the ultimate in analytical processing performance to SAS applications transparently.

## GREENPLUM DATA COMPUTING APPLIANCE (DCA)

The Greenplum Data Computing Appliance (DCA) brings together SAS Grid Computing and Greenplum Database to offer the fastest loading and analytical performance while also eliminating the complexity, delays and administration costs of customer-assembled infrastructure.

Greenplum DCA can also be configured with additional modules beyond those used to run Greenplum Database. SAS Grid Computing is deployed directly on the DCA hardware itself. One or more 4-server modules are configured to run SAS Grid Computing and connected to Isilon for SAS data set storage. The result is in a uniquely capable, analytics-focused environment with all of the benefits of the DCA.

The flexibility of the DCA is such that in addition to deploying modules running Greenplum Database and SAS Grid Computing, customers can purchase modules for Greenplum HD, EMC's Hadoop distribution, and benefit from rapid time-to-market for challenges where the ingest, filtering and analysis of semi-structured and unstructured data are required.

Greenplum DCA can also host integrated ETL servers with ETL engines from Greenplum partner companies hosted on DCA-resident servers called Data Integration Accelerator modules.

The entire design of the DCA is built around flexibility, performance, scale and simplified administration. When delivering a DCA, all modules and required power, cooling and high-speed interconnects are pre-configured for rapid deployment and from the Greenplum Command Central framework for simplified administration.

Greenplum DCA is a key component of the Greenplum Universal Analytics Platform (UAP), a unified platform enabling powerful and agile analytics that can transform how organizations use data. The DCA's modular packaging permits customer-dependent deployment of integrated SAS Grid Computing, Greenplum Database, ETL and Hadoop components, all running on optimized EMC servers alongside Isilon scale-out Network Attached Storage (NAS).

Purpose-built for data co-processing, the Greenplum DCA shortens time to value by sidestepping the delays, risks and maintenance costs inherent in "build from scratch" approaches to MPP solutions, whether using name brand or "white-box" servers, storage and interconnects.

## EMC ISILON SCALE-OUT STORAGE

To address the data storage needs for SAS Grid Computing environments, EMC Isilon Scale-out NAS integrates easily with the EMC Greenplum DCA to deliver the storage I/O throughput needed to complete SAS jobs in the timeframe required by SAS Grid Computing users. Isilon scale-out storage also provides a number of important advantages for SAS Grid Computing environments:

**Massive scalability:** EMC Isilon NAS can easily scale to over 15 PB in capacity with a highly efficient, easy-to-manage, single file system/single volume storage solution.

**Unmatched efficiency to reduce costs:** With Isilon storage systems, organizations can achieve 80 percent utilization with a single pool of storage. This industry-leading storage efficiency helps enterprises to dramatically improve the TCO and ROI of their storage infrastructure. This industry-leading storage efficiency, combined with a simple, easy-to-manage approach, helps enterprises to reduce capital expenditures as well as ongoing operating costs. In addition, with EMC Isilon SmartPools software, organizations can optimize resources with an automated tiered storage strategy that provides them with the right combination of performance and economy.

**Highly reliable storage:** Isilon data storage solutions leverage built-in enterprise data protection to provide the highest levels of reliability, availability, and serviceability in the industry. Isilon storage systems include dependable and efficient snapshot data protection with Isilon SnapshotIQ™ for backup and recovery, and for reliable disaster recovery protection, Isilon SyncIQ™ provides data replication of large, mission-critical data sets to multiple shared storage systems at multiple sites. Isilon Scale-out NAS storage is designed to withstand multiple simultaneous component failures while still affording unfettered access to the entire file system and dataset.

**Simplified data and storage management:** Isilon storage systems are simple to install, manage and scale, at virtually any size. And, unlike traditional enterprise storage, our solutions stay simple no matter how much storage is added. With powerful data and storage management software options such as Isilon SmartPools and InsightIQ, customers can also automate and further streamline management of their data assets.

## SYSTEM CONFIGURATION

**EMC Greenplum DCA Node Configuration (2 Master Servers and 4 Node Servers)**

Software

- Greenplum RDBMS v4.0.3
- Red Hat Enterprise Linux 5 Update 5 (RHEL 5.5)

Hardware

- Intel Xeon dual 6-core 2.93GHz Processor
- 48GB RAM
- Hard disk – 12 600 GB 15k SAS
- RAID controller – dual channel 6 Gb/s SA
- Dual-port converged network adapter – 2 x 10Gb/s

**SAS Grid Computing Node Configuration (1 SAS Grid Master and 3 SAS Grid Processing Nodes)**

Software

- SAS Grid Computing 9.3 64-bit Software
- Red Hat Enterprise Linux 5 Update 5 (RHEL 5.5)
- Single SAS Data Mount to each EMC Isilon S200 node

Hardware

- Intel Xeon dual 6-core 2.93GHz Processor
- 48GB RAM
- Hard disk – 12 600 GB 15k SAS
- RAID controller – dual channel 6 Gb/s SA
- Dual-port converged network adapter – 2 x 10Gb/s

**Isilon S200 Scale-out NAS Storage Configuration (8 S200 NAS Platform Nodes)**

- 22 SAS drives @ 10K and 600GB, 2 SSDs @ 400GB per S200
- 2 X 10Gb/s network connections

**Flexible multi-protocol support:** Isilon scale-out storage solutions help simplify workflows and enable enterprises to get more value from their business analytics initiatives by allowing users to interact with data using a wide range of standard protocols including iSCSI, NFS, SMB, HTTP, FTP, and HDFS. As a result, Isilon Scale-out NAS storage can serve as the central repository for data assets that can easily support data analytics initiatives as well as other enterprise applications and workloads while eliminating the need to manually move data around.

## INTEGRATED SOLUTION

The integration of SAS Grid Computing modules with the Greenplum DCA and Isilon Scale-out NAS delivers significantly improved availability, flexibility, performance, and scalability of analytic processing. Compared to traditional SAS infrastructures, this combination helps enterprises gain new business insights in a fraction of the time required with traditional architectures. Moreover, this integrated solution serves to bring previously unsolvable analytical problems within reach.

The increased computational capacity of SAS Grid Computing running on Greenplum DCA, in-database analytical computational performance of Greenplum Database, and the scale-out performance of Isilon Scale-out NAS permits SAS analysts to execute more iterations of development, testing, and validation processes on more detailed data than ever before, resulting in faster, more targeted and more accurate models and applications.

Greenplum DCAs are preconfigured to enable fast and easy installation and include comprehensive administration software that manages the DCA – including its grid of processors, memory, storage, and database software – as if it were a single platform. Together, EMC Isilon storage and Greenplum DCAs provide a highly reliable infrastructure for mission-critical SAS applications.

### VALIDATING PERFORMANCE AND SCALABILITY

To validate the performance and scalability benefits of the integrated grid solution, SAS and EMC performed two different tests designed to measure the performance of a SAS Grid Computing environment integrated into the EMC Greenplum DCA and with integrated EMC Isilon scale-out NAS storage.

The test configuration included an EMC Greenplum DCA, SAS Grid Computing and EMC Isilon Scale-out NAS. The EMC Greenplum DCA used was a quarter-rack model equipped with 2 master servers and 4 segment servers, each capable of spawning multiple processing threads. The DCA was configured with a 4-node SAS Grid Computing module that included one SAS master node and 3 SAS grid processing nodes. An 8-node EMC Isilon S200 Scale-out NAS platform provided shared storage for the test configuration. All servers were internally connected via an interconnect bus using dual 10 Gigabit Ethernet ports.

The DCA segment servers that run Greenplum Database and SAS Grid Computing nodes in the test configuration are each hosted on 4 servers. Within each server are dual 6-core Intel Xeon processors with 48GB memory and 12 600GB SAS hard disk drives (HDDs). Both the segment servers and SAS Grid Computing nodes run Red Hat Enterprise Linux 5 Update 5 (RHEL 5.5).

Based on the testing as described below, SAS concluded that the EMC Isilon S200 Scale-out NAS component of the solution meets the storage I/O requirements of SAS Grid Computing by a wide margin, thereby validating sufficient performance plus scalability. The tests also determined that for optimal performance, SASWORK files should be stored on SAS Grid Computing local storage (inside the DCA), while the Isilon S200 Scale-out NAS should be utilized for input and output SAS data files.

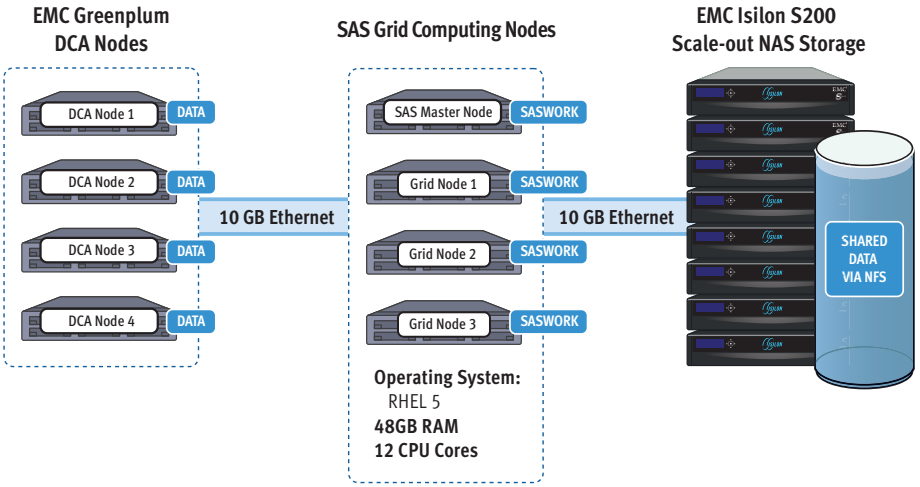


Figure 1. Test Configuration comprising SAS Grid Computing running on EMC Greenplum DCA with Isilon S200 Scale-out NAS

## SAS WORKLOAD CALIBRATION TESTING SCENARIO

The SAS calibration workload was engineered by SAS to represent a basic retail application. It uses sales history from comparable retail stores to calibrate demand model components and marketing mix effects, which are consumed by the forecast process to generate forecasts. It is typically run in batch during off-hours to prepare for forecasting.

The calibration workload consists of a total number of 750 jobs, which are spread across the SAS Grid Computing nodes by the SAS Grid Computing master node. The test runs a configurable number of SAS processes concurrently that execute the demand model calibration process. Specifically, the number of concurrent processes varied from 40 to 250 per SAS Grid Computing node (or 3.3 to 20.8 jobs per server CPU core) for the configuration-under-test.

All jobs in each "wave" of simultaneous jobs must finish before the next "wave" is launched. The workload emulates a SAS solution deployed in a distributed environment using shared input data and a shared output analytic data mart.

The first test used the SAS workload calibration benchmark to assess the performance of EMC Greenplum DCA running the 4-node SAS Grid Computing and 8-node EMC Isilon Scale-out NAS platform, as well as the optimal NFS client mount options for SAS Grid Computing nodes.

The workload calibration test used the Isilon S200 Scale-out NAS platform for storage of both input and output SAS data, while both SAS Grid Computing local storage and Isilon S200 shared storage were tested as options for storage of temporary SASWORK files. Furthermore, in addition to default NFS client mount options (including the soft mount option); the following specific mount settings were used for SAS Grid Computing nodes for different tests in order to determine the optimal settings for minimizing the test run-time.

- The noac NFS mount option disables attribute caching on the NFS client. This means that every reference to attributes is satisfied directly from the server though file data may still be cached, which impacts performance through additional latency, network load, and server load.
- The noatime NFS mount option disables updating of file access time metadata on NFS server systems reducing client, network, and server overhead.
- As with the noac option, setting actime0=0 disables attribute caching on the client, which adversely impacts performance and overhead.

## SAS WORKLOAD CALIBRATION TESTING RESULTS

The SAS workload calibration test measured the elapsed run-time for test completion for 8 different benchmarks using a numbered range of configurable of jobs per SAS Grid Computing node, varying location of SASWORK files between SAS Grid Computing local storage and Isilon Scale-out S200, and different NFS client options as discussed above. The two benchmarks with the lowest run-times had the following characteristics:

- 72 simultaneous jobs per SAS Grid Computing node, which translated into 6 simultaneous jobs per server CPU core
- Use of SAS Grid Computing local storage for SASWORK files
- Use of soft NFS client mount option

## THROUGHPUT TESTING SCENARIO

The SAS benchmarking staff also ran throughput tests on a single node (12 core) and 4-node (48 core) SAS Grid Computing environments to simulate the workload of a typical SAS Grid Computing job on EMC Greenplum DCA environments. The goal of the throughput test scenario was to evaluate the I/O performance of EMC Isilon Scale-out NAS Storage (specifically, the Isilon S200 platform) when it provides shared storage access for the SAS Grid Computing running on EMC Greenplum DCA. The SAS application benchmark test places an emphasis on I/O, which SAS software utilizes heavily during typical SAS Grid Computing computing program execution.

## THROUGHPUT TESTING SPECIFICS

An I/O-intensive SAS application benchmark was used to demonstrate high throughput and scalable I/O. This test consists of running multiple instances of file transfer of 10GB files between SAS Grid Computing node local storage and Isilon S200 Scale-out shared storage. Throughput is increased by running up to 24 concurrent instances each on each of the 4 SAS Grid Computing node server systems or a total of up to 96 concurrent jobs. The file transfer benchmark is primarily doing streaming (sequential) read, write or read-write I/O with SAS Grid Computing nodes using default NFS mount options (including the soft mount option which minimizes impact on clients of network or server overhead) and a block size of 128KB. The throughput test also comprised a range of benchmarks with a goal of determining best practices for locations of SASWORK and SAS data files for the configuration-under-test.

SAS does temporary file I/O in the SASWORK directory, which is a disk directory configurable as a SAS option. SAS I/O activity in SASWORK can be quite substantial both in terms of intensity and disk space, depending on the SAS job. A critical aspect of planning a SAS grid is allocating adequate storage as well as using the optimal location (i.e. local disk or shared network storage) for SASWORK.

## THROUGHPUT TEST RESULTS AND I/O TUNING

Throughput tests were conducted on single node (12 core) and 4-node (48 core) SAS Grid Computing configurations. Testing of the single node SAS configuration showed sustained throughput of 712-792 MB/s (or 59-66 Mb/s per core) for read-write tests using local storage for SASWORK and Isilon Scale-out NAS storage for SAS data files. SAS standard best practices recommendation for I/O throughput for external storage for SAS Grid Computing environments is 50 MB/s per core and the DCA solution exceeded that that by a wide margin demonstrating performance scalability.

Similarly, testing of the four node SAS configuration showed sustained throughput of 2450-3200 MB/s (or 51-66 Mb/s per core) for read-write tests using local storage for SASWORK and Isilon Scale-out NAS storage for SAS data files, again fully meeting the requirements for SAS grid storage. The 8-node EMC Isilon S200 scale-out NAS cluster delivered sustained aggregate throughput of 300 MBps per S200 node.

The throughput test shows that the following best practices are required for optimal performance of the DCA solution comprising SAS Grid Computing and EMC Isilon S200 Scale-out NAS solution.

- Deploy SASWORK on the SAS Grid Computing local storage
- Store all SAS data files on the EMC Isilon Scale-out NAS shared storage

## CONCLUSION

In-depth testing by SAS and EMC confirms that a single-rack EMC Greenplum DCA combining SAS Grid Computing with EMC Isilon S200 Scale-out NAS delivers the demanding compute, I/O, and throughput required to support customers' needs to modernize their SAS infrastructure for flexibility, high performance, and scalability. Performance tests showed that the combination of utilizing EMC Isilon S200 Scale-out NAS for input/output SAS data and SAS Grid Computing local storage for SASWORK temporary file data delivers outstanding storage I/O performance for the analytics solution in order to take advantage of the I/O throughput capabilities of SAS Grid Computing infrastructure.

## CONTACT US

To learn more about how EMC products, services, and solutions help solve your business and IT challenges contact your local representative or authorized reseller —or visit us at www.EMC.com/isilon

For more information about how SAS® software and services can help you make better decisions faster, contact your local SAS office at sas.com/offices



**EMC Corporation**
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

**EMC Isilon Corporation**
Seattle, Washington 98104
1-206-315-7500
www.EMC.com/isilon