White Paper

Scaling Objectivity Database Performance with Panasas® Scale-Out NAS Storage

A Benchmark Report

August 2011

Background

Objectivity/DB® uses a powerful distributed processing architecture to manage localized, centralized, or distributed object databases. Panasas storage appliances accelerate time-to-results with highly reliable, massively scalable storage solutions for performance-driven applications in Government R&D, Oil & Gas, Manufacturing, Higher Education & Research, Life Sciences, CAE manufacturing, and Financial markets.

This white paper discusses the results of a benchmark run to demonstrate the performance scalability of Objectivity/DB when accessing files stored on a Panasas® ActiveStor™ 12 scale-out NAS storage system. It covers the main features of Objectivity/DB and the Panasas equipment before explaining the benchmark software and the hardware configuration. It then presents and interprets the benchmark results. The benchmark demonstrated significant improvements in write and read speed over conventional disk storage.

Objectivity/DB Overview

Objectivity, Inc. is a leading provider of extremely large distributed database and data management technologies serving next-generation enterprise, government, web businesses, social networks, and data intensive research in cloud, non-cloud, and embedded environments. The company's flagship product, Objectivity/DB, is primarily used in Extremely Large Databases in the DoD, DHS and DoE communities as well as in complex real-time process control and in telecom hardware and systems.

No Mapping Layer - Objects and the relationships between them are stored directly, rather than in the column and row format of relational databases. This avoids the expensive and inefficient mapping layers needed between object oriented languages, such as C++, C# and Java, and the database storage mechanism. Objectivity/DB is extremely fast at navigating complex networks and tree structures as there is no relational join table processing.

Primary architectural features of Objectivity/DB include:

- A single logical view of distributed databases
- Simple, distributed server architecture
- Distributed Parallel Query Engine with replaceable components
- Predictable scalability

Single Logical View - Users see a single federated database that consists of local or distributed databases. The databases contain application controllable clusters of objects, relationships, collections, and indices.

Simple, Distributed Server Architecture - Objectivity/DB is directly linked with client applications. It accesses simple, distributed servers that manage transaction management, remote data access, and fragments of parallel queries. It can be configured in standalone, networked, peer-to-peer, Service Oriented Architectures (SOA), and grid or cloud computing environments and takes care of all heterogeneity issues that arise because of differences in languages, file systems, operating systems, and hardware.

Distributed Parallel Query Engine – Clients can service queries locally or use a parallel query engine that distributes fragments of queries to simple query servers located close to the databases to be accessed. There are replaceable components for controlling how queries are split, executed and serviced. Query servers can execute filters to refine results before they are returned to the client. The query servers can be replaced with agents for accessing external data sources such as RDBMSs or the Web.

Predictable Scalability - The address space available within a single logical view of a distributed federation of databases covers billions of terabytes. Objectivity/DB customers have built some of the world's largest databases, including a petabyte+ system at Stanford Linear Accelerator Center. Objectivity/DB has been used in data fusion systems that ingest and correlate over one terabyte of data per hour with simultaneous deletion and queries.

Panasas ActiveStor Overview

Panasas ActiveStor is the world's fastest parallel storage system, bringing plug-and-play simplicity to HPC storage deployments. Based on a fourth-generation storage blade architecture and the Panasas® PanFS™ storage operating system, ActiveStor delivers unmatched parallel file system performance in addition to the scalability, manageability, reliability, and value required by demanding technical computing organizations in the energy, government, finance, manufacturing, bioscience, and other research sectors.

Primary architectural features include:

- A single logical view of data for simple administration and access to data.
- Easy to use, tightly integrated storage appliance.
- Distributed parallel file system and high performance DirectFlow® protocol.
- Object RAID on a storage blade hardware architecture provides the basis for scalable performance.

Attributes for superior scalability: Panasas parallel storage utilizes a single global name space allowing seamless manageability at massive scale. Panasas customers have successfully deployed large capacity implementations, including a six petabyte shared file system at Los Alamos National Laboratory.

Reliability: Panasas storage is known for its superior reliability and data availability. The patented object RAID architecture and tiered parity technology ensures no single point of failure and superior data recovery on RAID protected data.

Benchmark Software

NetElem was selected to best test the scalability of Objectivity/DB using Panasas storage. This is because NetElem generates huge amounts of data in a relatively short amount of time. NetElem simulates typical telecommunications network entities (ports, cards, cables, etc.) which are efficiently represented and manipulated as objects. A telecommunications network has hundreds of thousands of such objects and therefore requires massive amounts of storage to store and manipulate these objects.

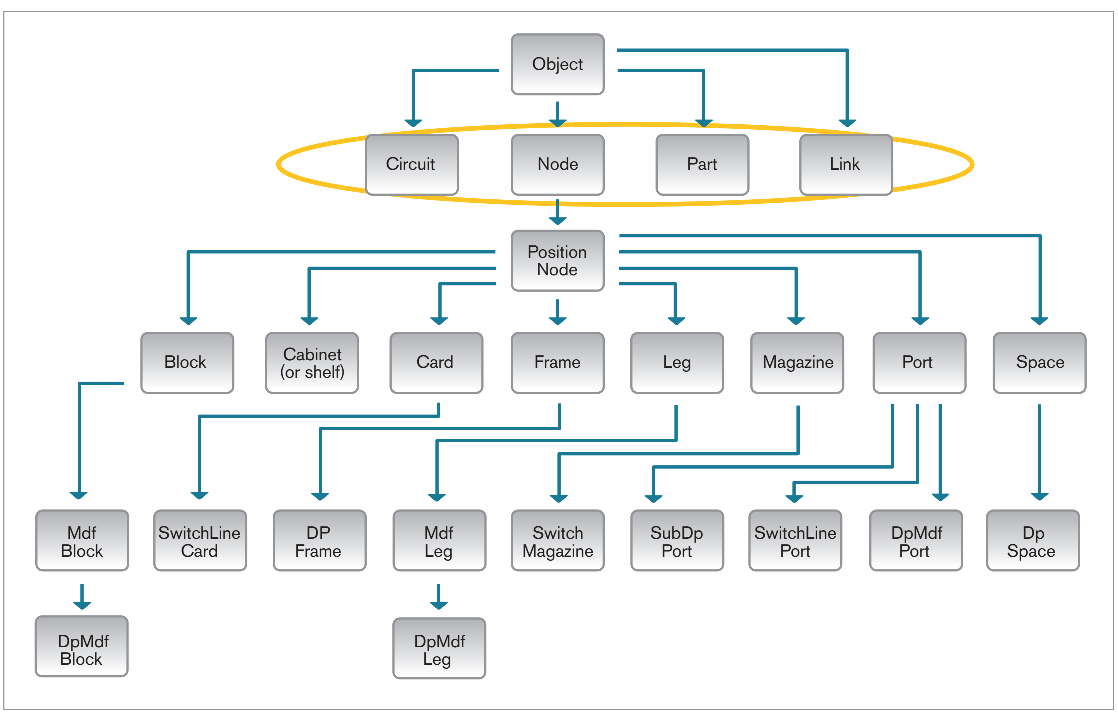

Network Element Management System object class model.

The NetElem Suite

The object class model for the Network Element Management System application is shown above. Only the circled types are used in the benchmark.

The benchmark has three phases:

- **Create -** Create the tree of network elements
- **Traversal -** Navigate the tree
- **Scan -** Scan one class of objects sequentially

NetElem Create Phase

In the first phase, NetElem creates 16 databases of 2GB each. Each one contains 86,482,160 objects of the following four types:

Object Type	Count
Circuits	6,144,000
Nodes	43,396,032
Parts	78,128
Links	36,864,000

The object instances are linked together using Objectivity/DB 1-to-many or many-to many associations. The Links are a type of object, not a database relationship type.

NetElem Navigate Phase

The navigation function starts from a root node and navigates to all node type objects, using Objectivity/DB persistent relationships (associations). The number of objects visited during navigation is 43,396,032.

NetElem Scan Phase

The scan function scans the 16 databases for Circuit and Port objects. Scanning in this case is mostly sequential read, since the container is created once and never modified. The number of Circuit and Port objects read during the scan is 6,144,000 and 18,432,000, respectively.

Benchmark Hardware

Objectivity/DB (version 10.1) was configured as a standalone application, running on a multi-core, dual socket server and accessing its database stored on Panasas ActiveStor 12.

The Panasas storage equipment was configured with a single ActiveStor 12 storage shelf with one director blade and 10 storage blades running PanFS 4.0.1. Eight client nodes were used for the benchmarking, each with dual 2.4 GHz quad core Intel Xeon processors and 24 GB of memory (reduced to 1GB to avoid local caching). The client OS was running Red Hat Enterprise Linux 5, kernel 2.6.18.

Benchmark Results

Single Server Scaling

The first part of the benchmark measured the performance of an increasing number of processes running locally on one application server accessing data stored on a Panasas ActiveStor 12 storage system. The results showed that the time taken to create, traverse and scan the DB instances with ActiveStor 12 was significantly less than what is required when storing the data on conventional / local storage (ActiveStor 12 was 7.8x faster for Ingest, 14x faster for Traversal and up to 39x faster for the Scan phases). This increased throughput is achieved by processing the data streams through PanFS, the Panasas parallel file system.

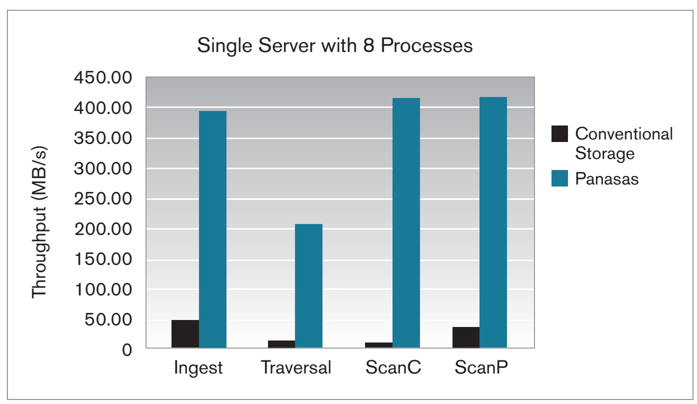

Next, we examined the throughput (in MB/s) for each phase while increasing the number of processes from 1 to 8 with the large databases stored on the ActiveStor 12 scale out NAS system. Note that the throughput increased almost linearly with the number of object databases proving the benefit of using ActiveStor 12 when processing multiple concurrent streams of data.

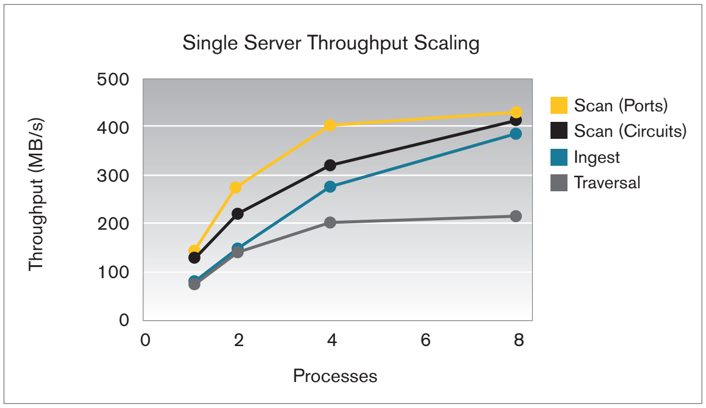

Multi-Server Scaling

The second part of the benchmark measured the effect of scaling the number of client nodes (each with two processes) accessing data on ActiveStor 12. What is particularly significant is that the throughput, especially the create (write) rate, increased linearly and did not reach a plateau with the number of processes and threads available.

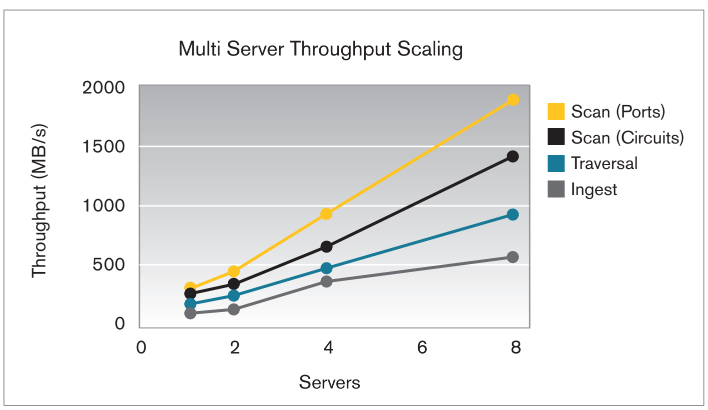

Summary

The benchmark convincingly demonstrated the superior performance that can be achieved with Objectivity/DB using a Panasas scale-out NAS system, versus local storage. Data ingest was 7.8 times faster, the traversal stage was 6.8x faster, and sequential scans were up to 39x faster with increasing numbers of processes. Furthermore, database I/O performance was demonstrated to scale linearly as the number of clients increased when using the Panasas PanFS file system.

Clearly, the combination of Objectivity/DB and Panasas ActiveStor 12 offers breakthrough performance for applications that manipulate or query large amounts of complex data. The benchmarked configuration is particularly attractive for NoSQL and other applications with extremely large storage requirements.

08172011 1065