

# The Right Data, in the Right Place, at the Right Time

The Sun Storage and Archive Solution for HPC



## Highlights

- **Improved economics:** Reduce total cost of ownership (TCO) by up to 44% over three years
- **Open standards:** Leverage open source, published APIs, and industry standard file formats to take advantage of community-driven enhancements and knowledge sharing
- **Powerful scalability:** Scale capacity, performance, or both as cluster demands grow over time
- **Flexibility and choice:** Select from a variety of devices that best match your unique needs — and avoid the typical “one-size-fits-all” vendor approach

## Today's data challenge in the HPC world

Increasing computational power and advances in clustering technology have enabled a rapid expansion in the use of high-performance computing (HPC) for a wide variety of applications. And organizations are leveraging these technologies to gain a competitive advantage and achieve improved business results. For example, improved trading analysis can result in better trading decisions and increased trading profits. More rapid and accurate analysis of seismic data can help find oil and gas reserves more quickly and save costs by more accurately pinpointing reserve locations and increasing the hit rate for drilling. And in the manufacturing sector, more comprehensive crash analysis can lead to safer, more reliable motor vehicles, which should increase customer satisfaction and revenue.

To achieve these business gains, organizations are creating and applying more sophisticated analysis to their business problems. These efforts often result in two challenges related to data management. The first challenge is the need to get the data in and out of the HPC cluster at much higher rates than ever before (a computing cluster is idle while waiting for data). The second challenge is storing these large amounts of data in a way that is both simple to manage and economical to run.

The size of the data sets can range from tens of terabytes to petabytes of overall capacity; however, often only a small portion of that data is in use by the cluster where high I/O performance is required.

## The answer

The Sun Storage and Archive Solution for HPC enables you to manage the balance between high-performance access to your data when needed and cost effective, lower-performance storage for large amounts of data over the

long term. Built with an architecture that supports multiple storage tiers, the solution provides high-performance storage for current use, bulk disk storage for low-performance needs, and long-term storage on power-efficient tape for data requiring only occasional access. Data movement or migration is managed through an automated policy engine to help ensure that the right data is in the right place at the right time.

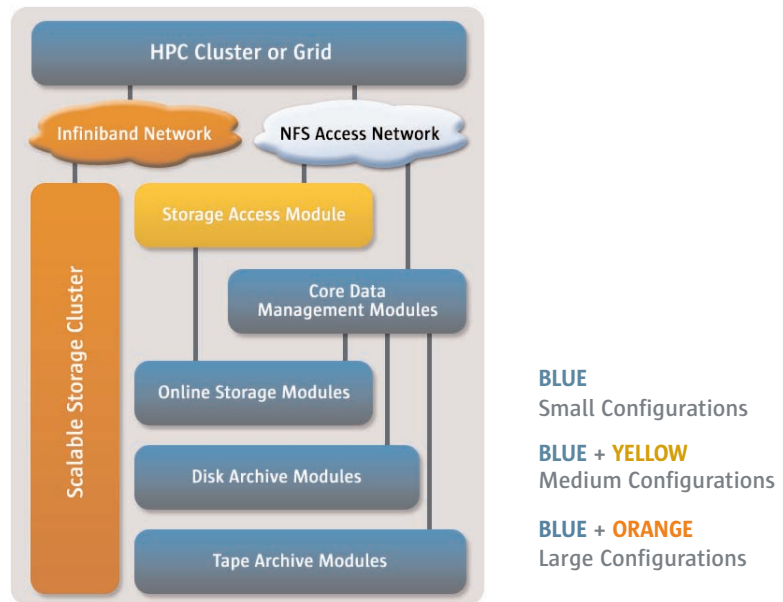
### Managing within a budget with limited staffing

Many HPC projects are funded with a fixed budget, which means organizations must work hard to maximize the amount of compute power available to run their applications. The Sun Storage and Archive Solution for HPC can help them achieve this goal by reducing the three-year total cost of ownership by up to 44% when compared to typical disk-based NFS appliances that serve HPC environments. By automatically migrating data that is not in current use to tape media, power consumption can be reduced by up to 24% (assuming 50% of the data is moved to tape).

With the Sun Storage and Archive Solution for HPC, data protection is incorporated into the architecture, greatly reducing the time, effort, and cost of deploying and administering separate backup applications for HPC data. Through a simple policy interface, the Sun Storage and Archive Solution for HPC can automatically maintain one or more complete copies of an organization's data, or just the desired portions, on tape. This eliminates the need to purchase a separate environment to protect data and can result in additional savings.

### Scaling for explosive growth

A primary benefit of using clustered architectures for HPC is the ability to easily and rapidly scale compute power as needed. Growing the cluster also means increasing the storage I/O load and the overall amount of data that needs to be



*Solution Architecture*

managed. The Sun Storage and Archive Solution for HPC provides for the ability to grow by scaling performance, scaling capacity, or both. The architecture provides the ability to add only what is needed to avoid unnecessary expenditures.

### Leveraging open standards

Because many HPC environments are pushing the technology envelope, benefits can often be derived by customizing the workflow, modifying or tuning the source code, or reaching out to a community for specific expertise and advice. With proprietary source code or closed appliance solutions, users are often prevented from achieving these benefits and must instead lobby to vendors to achieve their goals. But with published APIs, interfaces, and open-source code, the Sun Storage and Archive Solution for HPC enables users to attain all of these advantages. In addition, data is stored in industry-standard file formats, freeing users from being locked into a single-vendor proprietary solution and ensuring that the data stored today can be read tomorrow by a variety of standards-compliant tools.

### Preserving flexibility

The flexible architecture of the Sun Storage and Archive Solution for HPC enables organizations to select the building blocks that best fit their needs. Modules can be added for growth or retired without having to “rip and replace” the entire environment. For new deployments, the solution can be integrated with existing assets to ensure the most economical solution possible. Availability of tuning parameters and control of data layouts lets users optimize the solution to fit their specific HPC workloads and needs.

### Speeding deployment

To help you get your solution up and running quickly and smoothly, Sun offers a suite of services. We'll install and implement the components, as well as design, install, configure, and test your servers so they are application-ready. Our service experts can assist you in avoiding pitfalls and streamlining your deployment so you can achieve your business goals.

### Solution operation and architecture

The Sun Storage and Archive Solution for HPC performs four main functions for an HPC environment. First, it serves data to the compute cluster or grid. Second, it stores and protects data on various storage tiers. Third, it provides a policy based hierarchical storage manager to automate the movement of data across the storage tiers. And fourth, it provides a continuous copy mechanism to protect critical data and alleviate the need for traditional backup applications.

To provide flexibility and scaling, the solution is architected from a series of modules or building blocks that allow selection of those components that are best matched to a particular set of needs. (see figure)

#### Core Data Management Modules

Consisting of Sun Fire™ servers running Sun StorageTek™ Storage Archive Manager and the Solaris™ 10 Operating System (OS), the Core Data Management Modules control the overall operation of the solution. The Core Data Management Modules manage the location and movement of data based on policies you set and manage metadata related to the environment. These modules also provide “data ingest” or loading of data into the environment, access to home directories for HPC users as well as connectivity between the deployed storage modules.

For smaller compute cluster environments, the Core Data Management Modules provide data access and I/O services between the compute cluster and Online Storage Modules which store the cluster working or scratch space. As needs evolve, performance can be scaled by adding more of these modules.

#### Storage Access Modules

Consisting of Sun Fire servers running Sun StorageTek QFS software and the Solaris™ 10 Operating System (OS), the optional Storage Access Modules can be deployed to achieve higher levels of I/O performance to the HPC cluster. When Storage Access Modules are used, they provide the compute cluster I/O services while the Core Data Management modules continue to serve user home directories and manage the overall environment. This results in a higher level of I/O performance to the compute cluster or grid.

As needs evolve, performance can be scaled by adding more Storage Access Modules however, for large scale HPC environments or moderate sized ones that require very high I/O performance needs, the Sun Customer Ready Scalable Storage Cluster can be seamlessly integrated with the solution to provide very high performance cluster working space instead of this module.

#### Online Storage Modules

Consisting of Sun Modular Storage, the Online Storage Modules provide high performance, reliable storage for user home directories, metadata related to the environment and in the case of a small or medium sized configuration provide cluster scratch space. Since the I/O profile of home directories, metadata structures and cluster scratch space are typically very different, the Online Storage Modules are set up specifically for these different needs. Both performance and capacity can be scaled by deploying additional Online Storage Modules.

#### Disk Archive Modules

Consisting of Sun Fire storage servers, the Disk Archive Modules provide high-density, low-cost, online capacity for storing large volumes of data that require regular access but are not under heavy usage by the cluster environment. Either via policy or explicit command, the Core Data Management Modules move data from the Online Storage Modules to the Disk Archive Modules to free up higher-performing space for data that requires it. Data can also be moved from the Tape Archive Modules to the Disk Archive Modules if there is a need for reprocessing data that was previously archived. Since all of the storage modules are contained within the same global name space, the movement of data is a fast, efficient operation that is transparent to users.

Additional Disk Archive Modules may be added to scale capacity or archiving bandwidth. These modules are optional, in some cases it may be preferable to archive directly to tape.

#### Tape Archive Modules

The Tape Archive Modules provide the same function as the Disk Archive Modules with the obvious advantages of lowest possible cost and power consumption for stored data. The Tape Archive Modules, like the Disk Archive Modules are contained within the same name space as the Online Storage Modules, so data access is seamless, regardless of media type. Unlike traditional tape backup, with StorageTek Storage Archive Manager software, files that are moved to tape archive are still visible and accessible to users without requiring a backup operator to initiate a restore operation from tape. The Sun StorageTek Storage Archive Manager software can also be configured to create additional copies of data as part of your data protection plans.

The Tape Archive Modules can consist of any of the leading Sun StorageTek Library portfolio. These modules are optional but are required if you want a tape protection copy of your data. Additional modules (libraries and/or drives) can be added to scale capacity and archive bandwidth.

#### **Sun Customer Ready Scalable Storage Cluster**

For clusters that require very high I/O bandwidth, the Sun Customer Ready Scalable Storage Cluster can be added to provide the utmost in I/O performance. The Scalable Storage Cluster is deployed in place of the Storage Access Modules. The Scalable Storage Cluster consists of the Lustre parallel file system, Sun Fire storage servers, and InfiniBand infrastructure that has been optimized to provide very high data rates. The Scalable Storage Cluster contains the storage capacity used for the cluster scratch space so in this instance the Online Disk Modules are sized only for user directories and the metadata requirements of the Core Data Management Modules.

Sun Fire servers are used as data movers that connect between the InfiniBand high performance compute or grid environment and the archive environment. Performance of the Sun Customer Ready Scalable Storage Cluster can be easily be scaled through the addition of Sun Fire storage servers.

#### **Policy-based management**

An integral part of the Sun Storage and Archive Solution for HPC, the StorageTek Storage Archive Manager software allows data management policies to be set automatically based on file attributes. Data can be managed according to the storage and access requirements of each user on the system and can be grouped, copied, and accessed based on the needs of the application and the users. By controlling the location of project data, the StorageTek Storage Archive Manager software helps maximize return on investments by storing data on the most appropriate media type.

#### **About Sun**

A singular vision, The Network is the Computer™, drives Sun in delivering industry-leading technologies that focus on the whole system — where hardware, software, and services combine. With a proven history of sharing, building communities, and innovation, Sun creates opportunities, both social and economic, around the world. You can learn more about Sun at sun.com.

Learn More  
For more information on the Sun Storage and Archive Solution for HPC, visit [sun.com/hpcdata](http://sun.com/hpcdata)