

# Storage Management and SAN Directions on the Cray X1

**Steve Johnson, Cray Inc., Mike Feuerstein,  
Cray Inc., and Paul Rutherford, ADIC (Advanced  
Digital Information Corporation)**

**ABSTRACT:** *The management of storage and data is a significant capability needed at all HPC sites. This talk will address the strategy and capability being developed for the X1 system and our relationship with ADIC (Advanced Digital Information Corporation) in implementing this strategy. Also to be discussed are challenges, trends, and various technologies existing in this area of HPC. ADIC product technology and its capabilities will be reviewed as it applies to HPC and the Cray storage management strategy, including DMF migration.*

## 1. Introduction

The Cray X1 system is capable of having almost an unlimited amount of direct attached storage. At the 2002 CUG in Manchester, the capabilities to manage that direct attach storage in terms of backup, archive, and migration was discussed. As a result of those discussions, Cray Inc. has enhanced its strategy towards storage management for the Cray X1 system.

The original strategy for data management on the Cray X1 relied heavily on the capabilities of IP networks and NFS, FTP, and UDP to explicitly move data from direct attach storage on the Cray X1 to a separate file server responsible for data management functions. The technological agreement negotiated between SGI and Cray Inc. placed limitations on the DMF capabilities available and the OS development required for the Cray X1 made supporting DMF a significant challenge. As a result, native data management functions, such as DMF, that had been available in previous Cray/SGI systems are not supported directly on the Cray X1. Likewise, support for direct attach tape has not been provided. Initial reasons were primarily complexity and availability of resources to develop and support such a capability. It can also be argued as to whether such a capability is an efficient use of the powerful Cray X1 processor resource.

There were several concerns and comments expressed by attendees of the 2002 Manchester CUG concerning this strategy. The deliverable performance of IP networks utilizing NFS, FTP, or UDP was one concern expressed. Secondly, users of HPC computing have become accustomed to a richer environment of capabilities and functionality for managing data such as had been provided by DMF

and similar data management packages. Finally, many attendees were concerned about the existing DMF based data archives they currently have on existing Cray Inc. and/or SGI equipment. These concerns were based upon the fact that DMF utilizes a proprietary tape format, and the expectation this data was going to need to be migrated to a different format and environment in the future.

## 2. Industry Trends

It was important for Cray Inc. to understand the trends of the storage industry and HPC computing as related to storage. There are a number of identified trends that we felt were significant to take into account as the data management strategy for the Cray X1 was reviewed.

The amount of storage and the size of the file systems continues to increase rapidly for HPC systems. A typical Cray X1 system currently ships with a minimum of eight Tbytes of storage per chassis in a system. The data management system needs to be able to manage hundreds of Tbytes of storage and tens of millions of files in a file system. Any capabilities provided for the Cray X1 system needs to be able to support this trend.

HPC data centers are increasingly become more heterogeneous in their makeup. Multiple compute systems from multiple vendors, along with a variety of storage and archive systems, are typically found at every HPC data center. As a result, data management capabilities that are capable of supporting truly heterogeneous environments are often a basic requirement.

In the search for improved efficiencies in the HPC datacenter, many locations are interested in consolidating or centralizing their storage and data management capabilities.

The support and utilization of industry standards for I/O and storage is almost a mandatory requirement for HPC systems. While standards still do not provide automatic assurance of interoperability, significant progress has been made by standards committees to improve interoperability of new products much faster than the past efforts. This potentially means faster adoption and utilization of new standards as they become available.

The increase in heterogeneous environments and the desire to consolidate or centralize storage and data management capabilities results in a strong need to find ways to share the data efficiently between various systems. Storage Area Networks with an efficient, high performance file system is one technology being used more frequently and successfully to allow the sharing of data and avoiding the performance constraints of network based sharing solutions.

Storage management software offerings are becoming increasing sophisticated and capable. As a result, customers expectations have increased along with the size of their data management problems.

### **3. Cray Inc. Data Management Strategy**

#### ***3.1 Developing the strategy***

Reviewing and developing an improved data management strategy for the Cray X1 required consideration of a number of factors in addition to the previously mentioned industry trends.

The infrastructure within Cray Inc. did not exist to develop a Cray X1 proprietary solution for data management that would meet the expectations of our customer base. It also was felt that, strategically, Cray Inc. development resources would be better utilized focusing on the introduction of the Cray X1 to the marketplace and optimizing its performance capabilities than developing internal data management capabilities. Therefore, it was decided to pursue a strategy of partnering in order to acquire access to third party technologies that would enhance the data management capabilities of the Cray X1. Specifically, Cray Inc. wanted to identify Storage Area Network technology that would allow a majority of the data management functions to be offloaded to other participants of the network. This was desirable to ensure efficient use of the powerful Cray X1 processors and to minimize the technical and business challenges of porting and supporting 3rd

party data management technology directly on the Cray X1 system.

The requirements for a 3rd party data management solution for the Cray X1 were as follows:

- The solution needs to have minimal impact on overall Cray X1 I/O performance and functionality.
- The solution needs to support heterogeneous data center environments, i.e., have support for a broad distribution of equipment from multiple vendors. Client software support for Solaris, IRIX, LINUX, AIX, and HP/UX is desired along with UNICOS/mp.
- The solution needs to be flexible in terms of the types of environments and other technologies it can interact with. This is important in order to support and preserve investments in a data management infrastructure that a customer may already have made.
- The solution needs to be based upon proven, reliable technology currently being deployed in major HPC data centers from an established and stable provider.
- The 3rd party partner needs to have compatible business interests and vision with Cray Inc., regarding the HPC marketplace and data management requirements.

Cray Inc. investigated a number of potential technologies that it felt could meet its requirements for a data management solution for the Cray X1. The key element of the technology needed to implement a solution was a SAN based file system. Technology offerings from ADIC, SUN, IBM Tivoli, and SGI were investigated. The capabilities of the technology, the compatibility with UNICOS/mp of the component needed to port to the Cray X1, and the compatibility as a partner was evaluated for each case. The result of this activity was the signing of an agreement between Cray Inc. and ADIC (Advanced Digital Information Corporation) in early October of 2002.

#### ***3.2 Why ADIC?***

The agreement between Cray Inc. and ADIC includes aspects for source licenses and reselling and support of ADIC products by Cray Inc. as part of the data management strategy. The source license agreement results in ADIC's StorNext File System being ported to the Cray X1. This software is the key element in providing SAN based file system capabilities. The reseller agreement allows Cray Inc.

to provide to our customers any or all of ADIC's products as part of a data management solution supplied by Cray Inc. Finally, the service agreement defines the aspects of how the ADIC products are supported at a customer site when provided by Cray Inc. Normally, the customer will receive first and second levels of support directly from Cray Inc. with escalation support from ADIC being available.

ADIC was clearly the supplier of 3rd party data management solutions most compatible with Cray Inc.'s requirements. The product technology, especially the StorNext FileSystem and StorNext Management Suite, are mature products, well suited for the Cray X1 system and successfully deployed in a number of HPC environments. The technical team at ADIC is very competent (multiple ex-Cray and SGI personnel), and the management team has developed a strong vision for the future. ADIC's business is based solely on data management products which was considered as a positive to Cray Inc. Unlike other partners considered by Cray Inc., ADIC does not offer the server component of the solution (implying they truly design and support their products to operate in a heterogeneous environment; and they must maintain a strong focus on the performance, capability, and reliability of their products). The ability of ADIC's technology to interact with a variety of other data management solutions (hardware and software), including tape storage systems from suppliers such as STK, provides significant flexibility in responding to individual customer requirements.

With the ADIC StorNext technology, Cray Inc. acquires the ability to create high performance, multi-cabinet X1 systems and the ability to closely integrate with other systems in a heterogeneous environment. The StorNext File System allows for the creation of SANs which will provide a simplified data movement architecture and avoidance of the necessity of using IP network traffic to move and manage data. The StorNext Management Suite will provide an integrated data management capability. This capability includes data protection and disaster recovery, the potential integration of high performance and low cost disk, HSM, and data life cycle management.

Finally, the topic of DMF migration is addressed in ADIC's StorNext File System release 2.2 which will provide a utility for conversion of DMF files for UNICOS and IRIX systems.

#### **4. Status and Plans**

Activities directed towards implementing the storage management strategy are well underway. Engineers from ADIC and Cray Inc. have been

working since last December on the porting of the StorNext File System to the Cray X1 platform. The initial software port has been completed and the software is being run on one of the in-house development systems. Basic functionality has been verified and initial performance evaluations have been executed. Unicos/mp release 2.3 is targeted to be the first release supporting the StorNext File System. Fiber Channel switch support in the Cray X1 I/O driver is in development and targeted for availability in 4Q03.

A Storage Area Network test bed has been assembled in Chippewa Falls consisting of multiple Cray X1 nodes, RS200 and RS100 storage systems, heterogeneous server components (Linux, Solaris, and Irix based), and Qlogic Fiber Channel switches. Expansion plans for this test bed are evolving, but do include the addition of, a tape storage subsystem from ADIC. The test bed is being used to verify the interoperability of typical heterogeneous components of a HPC SAN environment, develop Cray Inc. expertise in designing, integrating and deploying SAN based storage management solutions, tune and optimize SAN solutions for HPC applications, and demonstrate the capabilities and progress to our sponsors and customers.

Cray Inc. has added additional Product Line and Professional Services resources experienced in the area of storage and associated technologies to develop and deploy the storage management strategy. The storage practices leader position within Professional Services is key to eventual deployment of the strategy to Cray Inc.'s customers. This individual is directly engaged with the Product Line developers in designing, integrating, and implementing the technologies that are part of the storage management strategy. The storage practice leader will also interact with the Cray Inc. sales teams and support teams to identify the specific storage management needs of existing and new Cray X1 customers. He will be prepared to offer support and services such as:

- Legacy Migrations
- SNSM Integration
- Ad hoc Training
- Custom Configurations
- X1/SNFS, RS200 Extensions

#### **5. Summary**

##### ***Cray Inc. - Storage Strategy***

Cray Inc. has always integrated high performance storage components into supercomputer systems. While technology continues to advance and storage vendors continue to come and go, Cray Inc.'s basic approach remains the same – offer a best-of-

breed storage subsystem capable of sustaining supercomputer I/O.

Given the recent growth of standards-based Storage Area Networks (SAN), this now means that Cray Inc. storage solutions must not only leverage the best HPC storage systems, but they must also support a wider range of storage functionality, including heterogeneous resource sharing and file sharing. Meeting this challenge requires Cray Inc. storage to interoperate with the most common SAN products. To this end, Cray Inc. has and will continue to establish working partnerships with many of the leading storage and SAN vendors, including:

- Advanced Digital Information Corporation (ADIC)
- Cisco Systems, Inc.
- Emulex Corporation
- LSI Logic Storage Systems, Inc.
- Myricom, Inc.
- QLogic Corporation
- Quadrics, Ltd.

The ADIC relationship is especially notable since the Cray X1 leverages ADIC's StorNext File System to provide a shared SAN file system. Cray Inc. is also authorized to provide the complete ADIC StorNext Management Suite to its customers. The solution can be software only, or it can optionally include any of the ADIC tape libraries. In either case, it can provide what ADIC calls Total Data Life Management. Cray Inc.'s standard (SAN-capable) X1 systems include a core set of SAN components, all selected specifically to support the high performance X1 environment. For existing customer SANs or storage environments with components not qualified by Cray Inc., Professional Storage Services are available to implement customer-specific solutions. Storage Services are also provided for extensions to the base Cray RS200 functionality, for exceptional SNFS implementations, SNMS integrations, and legacy system data migrations. The storage strategy for future Cray Inc. products or HPC solutions will be similar to that for the Cray X1 in that it will leverage and depend on standards-based SAN products from industry leaders.

The Cray Inc. storage strategy reflects a seasoned understanding of what our core business is (supercomputing) and who our customers are. Cray Inc. storage subsystems and related services are specifically tailored for our unique market segment: a relatively small number of customers with very specialized requirements which demand a combination of extreme performance, extreme capacity, extreme reliability, and compelling value.

### ***Acknowledgements***

The authors would like to thank colleagues, users, and vendor staff for their assistance and input in generating this paper.

### ***About the Authors***

Steve Johnson is I/O Product Manager at Cray Inc. He has spent the majority of his career in HPC at the various forms of Cray involved in product development and I/O systems. He can be reached at Cray Inc., 1050 Lowater Road, Chippewa Falls, Wisconsin 54729. Mike Feuerstein is the storage practices leader for Cray Inc. Professional Services and has many years of experience in the storage and services industry for HPC. Paul Rutherford is VP of Technology at ADIC and has attended many CUGs in the past as a Cray Research employee. He can be reached at Advanced Digital Information Corporation, 11431 Willows Road NE, Redmond, Washington 98052