SGR® – October 2003
InfiniteStorage – CXFS™ / DMF

SGI: Alive And Well

For starters, yes, this is the same Silicon Graphics (SGI) that had the flashy workstations and servers 10 years ago. And while they are still in the business of building high performance workstations and servers, they also offer solid technology in the file system and data migration space. There have been a number of companies to come out of the closet these days ready to focus on the reference information space with a file system and / or volume management product. SGI's products complement this strategy as well fit in with their vertical industry-oriented server focus that includes media and entertainment, geosciences, the manufacturing industry (CAD/CAM) and government/defense. We initially wrote this brief in March, but find it appropriate to update some sections with the recent announcement of SGI's new branding of their storage business, InfiniteStorage.

The Problem

Reference information data is the largest growing segment of enterprise data (see "The Next Wave" at www.enterprisestoragegroup.com). Reference information is file-based data that is typically very large, and typically growing very rapidly. The data sets worked on by groups accessing this information may be up to several multiple terabytes in size. One example of how the media and entertainment business utilizes this type of information is outlined in Figure 1.

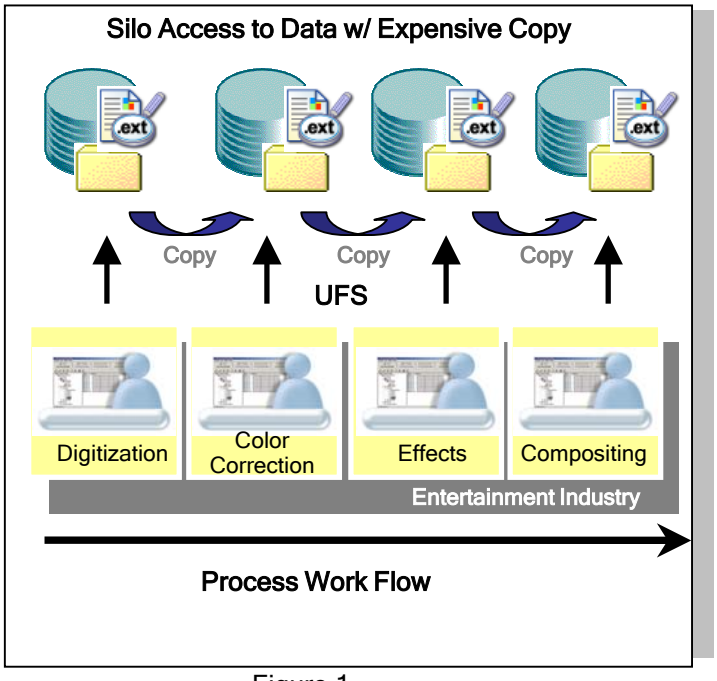

Figure 1

One group may access these large file sets to perform, for example, color correction. When this group is done with the data, the data is then copied to a second subsystem for another process such as adding effects. The copying of the data from one station to another is very expensive and time consuming--lasting anywhere from hours to days. Sobering fact: These users are not IT professionals and don't want to manage the data; they simply require immediate access to it. This traditional method of silo-based data access is very slow and expensive.

The Technology

SGI InfiniteStorage shared file system CXFS™ is a SAN-based shared file system technology based on SGI's original 64-bit XFS® file system. XFS has been around since 1994 and is widely trusted in the industry. Available since 1999, CXFS provides a single name space for clients connecting to the clustered file system and controls the locks on open files. As storage is added to the pool, clients automatically see the additional storage without any reconfiguration. CXFS also controls access to the files such that individual users cannot corrupt the data by writing to the same file at the same time. CXFS uses POSIX®, BSD, and SVR4 locks in a token system to lock open files. These standards ensure that no additional application file locking integration needs to take place.

When one user accesses a file for write purposes, CXFS checks out a token on that file so that other individuals only have read access to the file until the token is returned. CXFS supports IRIX, Linux® 32, and Linux 64 for Altix™, Solaris™, AIX® and Windows® operating systems as clients with support for Apple MAC® OS/X later this year. Additionally, Cray uses XFS on UNICOS® with each X1 box they ship. Environments also have the ability to configure two metadata servers for high availability. In addition, in the case of a failure, they have the ability to fail all CXFS functions to any other SGI® server in the cluster.

CXFS is one of the most scalable file systems available today. It supports file sizes of up to 9M TB and a full file

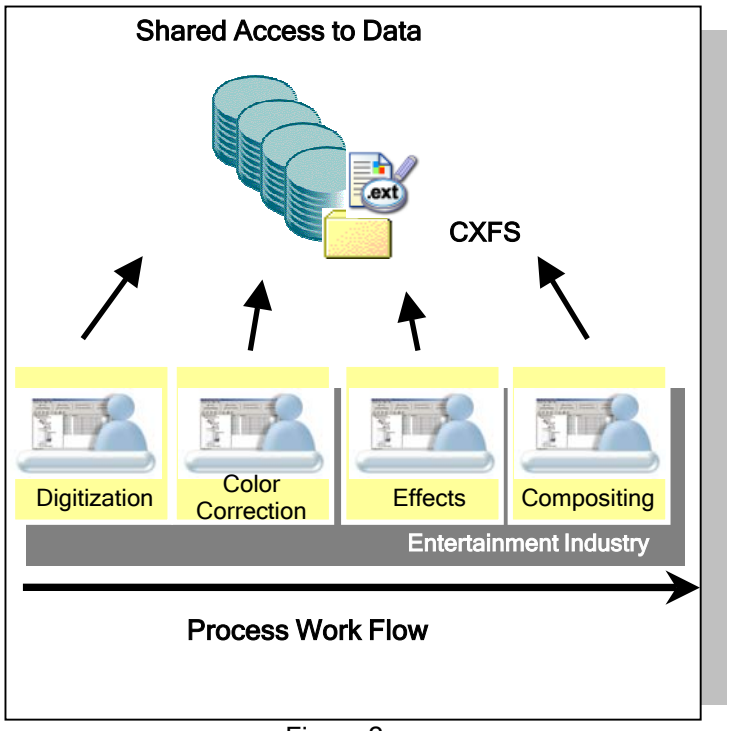

Figure 2

system of 18M TB. Such are claims not made by any other cluster file system vendor. By having the ability to scale to file sizes and systems this large, CXFS becomes a good fit for industry segments requiring reference information data. Media and entertainment, geosciences, weather and others that employ reference information can benefit from scalability of this magnitude. Figure 2 shows that having simultaneous; shared data access can save a great deal of time and money.

SGI also has a product called Data Migration Facility (DMF). DMF is an automated data migration product. The package–which works in conjunction with CXFS allows users to set policies that enable them to migrate data from primary storage to some form of nearline storage. As a file system fills up, policies governing data movement can be set. For example, as capacity is met, older data can be automatically migrated to near-line storage, keeping more expensive storage resource free for the most frequently accessed data. DMF supports and manages over 25 million files across single or multiple file systems. When a user makes a request for data that may have been migrated to near line storage, CXFS calls DMF to retrieve the file and is completely transparent to the applications. The retrieval rate of near line devices tends to be slower than online storage, but it is much faster than having IT perform a data recovery from tape.

SGI has developed a bundled SAN server and software solution called the SGI InfiniteStorage SAN 3000 for high-bandwidth SANs, and the SAN 2000 targeted at SAN/NAS convergence environments. The product is made up of LSI storage, 2GB Fibre Channel Brocade® switches and dual metadata servers for SAN and File System high availability. The SAN Server is available in configurations from 1 TB to 224 TB of raw capacity, support for 2 to 26 CXFS systems, and 800 MB/s to 2.4 GB/s of performance in single or multi-rack implementations.

Information Lifecycle Management (ILM) is a process that is becoming more prevalent in the storage industry. Implementing ILM for use with reference information is difficult due to the need to move large files. SGI offers some ILM components with their CXFS and DMF products. These solutions make it easier to manage reference data.

Partnerships

Recently SGI signed an agreement with AppIQ, where by AppIQ will provide SAN based monitoring services for the SGI platform. The agreement allows SGI to market and resell AppIQ's SRM technology. This will help to enable SGI to play in the ILM arena where a number of vendors are putting marketing dollars today. AppIQ allows IT to manage information from application creation throughout its lifecycle and ensure that it ends up on the right device throughout its lifecycle. This bodes well for SGI and coupled with their comprehensive shared file system and data migration tool have a very comprehensive, data management solution for the markets they play in.

Real World Example

One organization we spoke with is responsible for the acquisition, processing and storage of brain image data. Continuing advancements in imaging technology have produced increasingly higher resolution data, which has in turn increased file sizes and created the need for higher performance storage networking infrastructures for practical analyses. Prior to using SGI's CXFS, this organization used NFS servers over a 100BaseT network to analyze multigigabyte data sets. As file sizes grew and more people needed access to more data, the workflow became excruciatingly slow. As a result, bogged-down users were required to wait for the data they needed to transfer and load from the NFS servers. The IT group then deployed a SAN using CXFS with a DMF back end. The SAN has a couple of mainframes connected with 100MB HBAs through a switch. There are also a couple of

RAID 5 arrays. The organization now gets 400 MB/s sustainable throughput on just one of their systems and has greatly increased productivity.

CXFS can take advantage of DMF to monitor free space, migrating files based on age or size criteria off to near-line disk or to tape when free space thresholds are surpassed. Migrated data can access time generally ranges from sub-second to 20 seconds. Now both near-line disk and tape can be attached directly to the CXFS SAN giving users a single shared view of all data regardless of it's physical location.

The Bottom Line

Over the past year SGI has put a great deal of emphasis on facilitating storage in their most successful vertical markets. In the process, they have increased their storage attach rate from 20% to 44%. SGI has over 300 clients using its CXFS and DMF technologies, and these clients are managing hundreds and hundreds of terabytes. SGI has over 50 customers running Multi OS CXFS clusters – such as IRIX, Solaris and Windows at the same time.

The market segment CXFS focuses on, reference information, is growing at a CAGR of 92% as opposed to 61% for traditional information. ESG predicts that by 2005, reference information will represent 54% of all new information compared to 31% at the end of 2001. So if you do the math, what does all this add up to? Seems like SGI has a pretty large opportunity to make a lot of cash.

While heterogeneous clustered file systems are finally becoming a reality, SGI has very quietly attained more real-life runtime than anybody else and users should be advised to consider them as the requirement nears.

Author: Steve Kenniston