

Broadcast Engineering®

THE JOURNAL OF DIGITAL TELEVISION

APPLIED TECHNOLOGY

NEW PRODUCTS & REVIEWS

Avid Unity ISIS

The system offers real-time access to multiple users.

BY BILL MOREN



Each Unity ISIS 4RU chassis, called the ISIS Engine, holds 16 storage blades, two Ethernet switch blades, and redundant power and cooling.

Within a decade, the need for multiple users to have common, unrestricted access to digital media files has gone from a convenience to a business necessity. What had been referred to as the digital media shared storage system is becoming understood as the media network — the backbone of the broadcast or post-production facility. This term better conveys the concept of connectivity and bandwidth needed to deliver digital media in real time.

The fact that storage is involved is a given. It's how the storage is organized to deliver reliable, digital media in real time that differentiates the media network from the more generic shared storage system or SAN.

The problem

In terms of performance and need for availability, there are few data applications more demanding than broadcast and post-production video editing. These are collaborative environments where real-time access is constantly needed. Every client requires simultaneous isochronous performance and data delivered with time constraints.

Also, broadcast environments often support 24/7 operations, further driving up the demands on the media network infrastructure that is the cornerstone of the facility. Add to this the growing need to accommodate higher bandwidth HD and an increasing number of users requiring access to media, and it becomes apparent that the system's ability to scale — without compromising access or availability — creates a difficult problem set.

A new approach

Building on the principles of the Avid Unity MediaNetwork, the Unity ISIS was developed to address these challenges. ISIS stands for infinitely scalable intelligent storage. It is a modular design, leveraging industry standard GigE throughout.

By combining a modular architecture with Ethernet, it is possible to offer a highly reliable, redundant storage infrastructure that is easy to configure, expand, service and administer, as well as provide the real-time isochronous performance required in multi-client media applications.

**The storage system
... has no intrinsic
upper bounds in ...
storage capacity.**

Fundamental to the system is a real-time distributed 64-bit file system with key attributes of distributed file system management among clients, a centralized management resource and intelligent storage elements. This type of data architecture enables efficient client storage access, increasing performance and ensuring real-time access.

A distributed approach also means that as storage and clients are added, the processing power to manage the system grows as well, eradicating potential CPU bottlenecks that would limit scalability. The storage system, in fact, has no intrinsic upper bounds

in terms of client counts or storage capacity. Any limitations are functions of available Ethernet switching hardware along with the practical boundaries set to ensure thorough testing of configurations that will guarantee real-time delivery.

Two independent and fully functional Ethernet switches are integrated into each storage chassis, called the ISIS Engine, eliminating the need for external switches for many configurations. This integrated Ethernet architecture simplifies the configuration of a shared storage system, reducing administration and service costs while increasing reliability.

Next-generation media protection

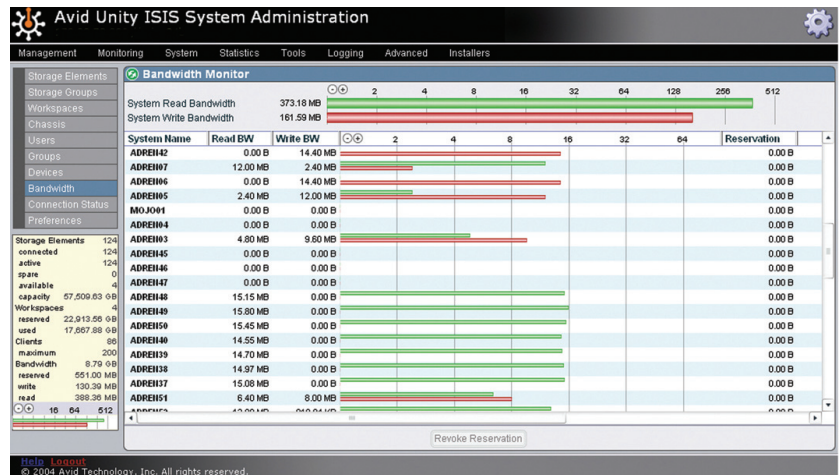
The Unity ISIS media network employs a distributed data striping plus mirroring approach. Media is divided into data chunks, and every time a data chunk is written on a storage blade, it is also written on another drive in the system. Unlike traditional RAID-1 (mirroring), however, the system does not pair the physical drives. The second copies of the data chunks are stored randomly across all of the other storage blades in the system.

A benefit to this approach is the increased bandwidth gained in striping data across the storage blades. Intelligence in the ISIS file system enables the clients to determine which storage blade is able to optimally service each request. Since clients have direct access to each intelligent storage blade, clients do not access a central resource to balance the I/O load across the system.

Another benefit to this media protection method is a much faster reconstruction of a failed blade's data. When a blade failure is sensed, the system makes a new copy of

mirror. This happens in parallel across all the storage blades in the system. Because every storage blade is involved in the reconstruction process, it takes a fraction of the time required by the parity RAID subsystems that are typically used in Fibre Channel SANs.

A second System Director, which can be added for redundancy, is constantly updated with the same metadata and monitors the primary unit via a regular heartbeat inquiry through dual redundant private Ethernet connections. Each switch blade



In the collaborative environments of broadcast and post-production video editing, workflows benefit from the Unity ISIS' ability to provide real-time isochronous performance and accommodate higher bandwidth HD and multiple user access.

Rethinking storage hardware

A network of intelligent, Ethernet-connected storage blades couple processing power and memory with a pair of high-performance SATA II disks. Each storage blade connects to an integrated Ethernet network that is also accessible to the clients.

Both storage blades and integrated Ethernet switches are housed in a chassis. They can be configured into the system to provide the desired bandwidth, capacity and redundancy.

Also connected to the network is an ISIS System Director, a central resource server that facilitates access to the storage by clients. When a client or storage blade wants to access a file, it serves as an index to the file system from which all access to a particular file can be derived. With the index in hand, a client or storage blade can then algorithmically find and access any other part of the desired file without iteratively communicating with a central metadata controller.

The 4RU chassis holds 16 storage blades, two Ethernet switch blades, and redundant power and cooling. Each storage blade contains two 250GB or 500GB drives for a total chassis capacity of 8TB or 16TB. Each switch blade provides eight user ports or 16GigE client ports per chassis. The rack-mount System Director has dual connections to Ethernet ports on two integrated switch blades.

includes a high-speed 12Gb port that interconnects switches from multiple ISIS Engines. Up to 12 chassis can be interconnected, providing a total storage capacity of 192TB per system.

With heat and vibration as the major causes of disk drive failure, the storage chassis uses highly efficient front-drawn laminar airflow for cooling. Redundant power and cooling are key capabilities of the system. There are load-sharing, hot-swappable power supplies in each chassis, and each power supply has its own AC input, allowing a failed unit to be pulled easily. Through the use of rigid enclosures, vibration has been reduced to low levels.

The system's architecture leverages file system, networking, storage and processing technology to meet user demands. For high-pressure, deadline-driven broadcast, post and film production, this means uninterrupted availability, reliability and the ability to scale and adapt intelligently to changing business requirements. Its layers of intelligence and redundancy reflect a media network design philosophy. The result is the foundation for an enterprise-level workflow.

BE

Bill Moren is a senior product manager for storage at Avid Technology.

Clients have direct access to each intelligent storage blade.

the mirrored data, randomly distributing it across the storage blades. This process is called reconstruction.

When a blade reconstruction is initiated, each storage blade determines whether it contains any data that was on the failed blade. If so, the blade itself initiates the copying of the data to a new location in the available pool of storage blades, creating a new

Reprinted with permission from the January 2007 issue of *Broadcast Engineering*® (www.broadcastengineering.com)

Copyright 2007, Prism Business Media. All rights reserved.

BE-61-CUU