



## Isilon Accelerates Delivery of Digital Content

Analyst: Anne MacFarland

### Management Summary

With all the hype about technology architectures, it is easy to become distracted from one of the basic tenets of business: **optimize revenue-producing processes early and often.** Revenue production means profitability, so when those processes have special needs, you build special environments, like clean rooms. You use special tools. And, these days, you may want to use optimized IT technology.

Computing has transformed many businesses, but use in business has also transformed computing. While database transactions are at the heart of business and process *management*, the digitization of communication and all manner of once-physical assets has broadened and fragmented the definition of IT optimization. **Optimizing for protracted sequential reads is not the same as optimizing for a high volume of small writes.**

Digitized content, in the form of images, videos, and large data sets, can be the basis of key business processes – or in some cases it can be the key business deliverable. **Digitized content is stored as files in vast, active repositories, whose comprehensiveness and accessibility are key to its business value.** These values often are inadequately served by traditional optimization in IT environments like SAN and NAS, which have been designed for transactional data and text files. When providing access to digital content is the business, and quality of service translates as the ability to deliver throughput at high bit rates, the following concerns must be easily – even effortlessly – addressed:

- **Physical Scalability:** Not only are the files big, but they should be managed as assets in a single pool - by a single file system. Multiple file systems complicate management and searching the collection as a whole. In these environments, growth comes in BIG chunks, but revenue models do not allow for high-cost infrastructure incrementation. Adding more must be easy and relatively inexpensive.
- **Scalable Quality of Service:** In content delivery situations, qualities of service include both streaming (bit rate) and access (time to start). To support a mature business model, multiple qualities of service to the same file must be provided, as well as acceptable access times for a variety of file sizes. Caching and data placement strategies are needed to deliver concurrent, high-throughput sequential-read access as a service. For the service to scale, these intricacies have to be built into the system, not just bolted onto an administrative layer.
- **Security:** Beyond the traditional concerns of data and system integrity, file-level authentication and file-specific security treatments may be required. Policies must be implementable at a file level – and must follow the asset wherever it resides.

Isilon Systems' *Isilon IQ* storage addresses these needs. **Isilon IQ's extensive metadata structure and integrated utility processes are optimized for concurrent, read-intensive, sequential data access, where files are not changed but are, instead, versioned.** If this sounds like the core of your revenue stream, either directly as a content provider or indirectly as a heavy user of large-file content, you may benefit from an optimized approach. Read on.

### IN THIS ISSUE

> Isilon IQ.....	2
> The Research Channel.....	3
> Conclusion .....	4

## Isilon IQ

Isilon Systems, Inc., of Seattle, WA, provides commodity-based storage hardware, its *OneFS* clustered file system, and its software utilities (see below). Each element is optimized to the task of serving digital content.

### *The Hardware*

The *Isilon IQ* product is a 2U-high clusterable node consisting of Intel-based I/O architecture, memory, and 4 ATA 160 GB disks (1.44 TB). The entry model is 4.3 TB of storage for \$49,950, of which 2 TB is immediately accessible, with a software license to upgrade to the full capacity. Each node is a peer – there is no master node<sup>1</sup>. At launch, a 32-node cluster is supported (about 50 TB). Larger clusters will be supported in the future. Maintenance is provided by IBM Global Services.

The nodes in a cluster are connected by a low-latency interconnect (switch). Latency is not tolerated because of the need for cache coherency, (discussed below). IsilonIQ nodes are based on gigabit Ethernet and use standard networking and reporting protocols. No Fibre Channel skills are needed.

### *A Larger Scale of RAID*

The cluster of nodes acts as a single pool of storage. **Instead of doing RAID across a single disk, IsilonIQ does RAID 5 and mirroring across up to 9 nodes.** In clusters of more than 9 nodes, IsilonIQ's *Flex Protect* utility (described below) will select the nodes to be striped across according to parameters set by the administrator

### *Coherent Cache*

Each node contributes 4 GB of RAM to the cluster's coherent cache. The larger the cluster, the larger the coherent cache. Performance of the system increases linearly as the system grows, unlike traditional NAS products. Nodes are synchronized by a distributed lock manager, the coherent cache, and a remote block manager. Ninety percent of the large cache is dedicated to reads, 10% to writes, making IsilonIQ transactionally-safe, but too slow, for a general purpose transaction environment.

<sup>1</sup> This makes deployment simpler, but requires cache that is coherent across the nodes. In large streaming repositories, the need for a single name space justifies the complexities of cache coherence (which are transparent to administrators).

## *Isilon Software*

Other distributed file systems<sup>2</sup> differ in where they keep their metadata (centralized, striped, locally with each file, etc.) **Isilon OneFS differs from other approaches by enriching the metadata with administratively relevant elements, which allows granular approaches to security and easier administration in a system that still can scale.** It also allows Isilon to reorganize how data services are delivered.

In an environment where the files do not change but are versioned, the classical multiple layers of administration (volume, file, security), designed to serve multi-purpose environments are not necessary. Isilon's four integrated utility software processes ensure data and system integrity, and enable quick and easy administration across clusters as large as 50 TB<sup>3</sup>.

- **Auto Balance** rebalances the data in a cluster according to business rules, in real time, non-disruptively. Business rules can be tiered so that an event at one layer can trigger a cascade of changes. This automation deals with the flow of events that are inherent to providing data access as a pay-for-view service.
- **Flex Protect** sets up security and replication policies. The administrator can change parameters, and immediately all data is written to the new policies. Additionally, as time permits, auto-balance goes out and adjusts how existing data is stored to meet the new policy. Data may be mirrored up to 8 times if the demand for access warrants it.
- **Media Scanner** runs in the background, continuously verifying the integrity of all files and repairing when necessary via error correcting algorithms.
- **Smart Read** algorithms translate the policies generated by the utilities above into the particulars of data placement. The placement is optimized for high concurrent throughput, but can also accommodate special needs such as security.

<sup>2</sup> See *Consolidated File Systems – Relieving the Pains of Scale in Data Storage* in **The Clipper Group Explorer** dated January 30, 2003, at [www.clipper.com/research/TCG2003003.pdf](http://www.clipper.com/research/TCG2003003.pdf).

<sup>3</sup> On the roadmap for next year is the ability to scale, not just to a bigger cluster, but to clusters across distances as their customer base demands.

As an integrated set, these utilities cover the administrative tasks needed, while automating the exacting but routine data and storage services. Should a disk fail, the system would stripe data elsewhere and use the existing RAID redundancy to rebuild files as needed. The journaling file system can rebuild itself as well.

The Isilon solution, taken as a whole, allows customers to scale both the size of their repositories and the performance (quality of service) characteristics of access time and throughput. Flex Protect and Media Scanner, in particular, address the security needs peculiar to these data environments.

### The Research Channel

Let's take a closer look at how one customer uses IsilonIQ. The Research Channel (ResearchChannel.org) is funded by a consortium of university, private and corporate research organizations to distribute content, both by broadcast and by on-demand over-the-Web access. It is hosted by the University of Washington, which also hosts a radio and a TV station, as well as the University's IT environment.

Science and research programming tends to be rich in graphics, not talking heads. And, according to Jim DeRoest, Assistant Director of University Computing Services, University of Washington, the higher the quality of the broadcast, the larger the audience, the longer they stay tuned, and the more successful the operation. While the Research Channel must be able to cope with lesser connectivity (down to 56 KB/s), their focus is on the high end of the bit-rate spectrum.

This focus on high throughput made traditional approaches like SAN less than attractive to the Research Channel as they surveyed their options. To achieve the 5.6 MB/s of today's MPEG2, and the 19 MB/s of tomorrow's HD, and, eventually, the studio quality bit rate of 270 MB/s in a SAN architecture, rampant replication of large files would be required. This would be costly and would constrain performance of the system as a whole. A NAS, file-based approach was more appropriate – but the large and growing repository had to be searchable as a whole, and storage administrators were already stretched thin – scalability had to be administratively

easy.

IsilonIQ's ability to stripe as needed, both for fault tolerance and for throughput scalability, was very attractive, according to DeRoest. The ability to add a node and push one button to rebalance the system<sup>4</sup>, made repository growth something to celebrate, not to dread. The ability to set digital policies once, and then to have them implemented across the integrated utilities makes Isilon an effective and comfortable solution in a complex environment. The ATA drives have performed well in this application.

The Research Channel started with a 3-node cluster in Isilon's beta program and grew it to 10 nodes during the first year. They are now ready for three more nodes. They started with a broadcast feed for cable TV organizations and are now offering more on-demand options for schools, universities and individuals. The quality they provide has led to inquiries from studios who are interested in sharing content for editing over the Internet.

The Research Channel uses StorageTek *Powderhorn* tape libraries as archives, and DeRoest does see a future need, as the collection grows even more massive, to stage content to tape and back. Additional projects, such as Digital Well and a project with the Library of Congress for a national video archive, indicate a need to think expansively about lifecycle criteria and how best to address different immediacies of on-demand on a massive scale.

### Conclusion

Using IT systems to access information streams is a growing part of computing. Presuming massive scale and optimizing for high throughput at a supportable cost is very different from the traditional optimizations of transaction processing, and requires very different storage characteristics and tight integration of data services. Commodity general-purpose storage is not the answer. In such a case, think smart. You probably need a purpose-built solution like IsilonIQ.



<sup>4</sup> It would automatically rebalance thereafter.

### ***About The Clipper Group, Inc.***

***The Clipper Group, Inc.***, is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at [www.clipper.com](http://www.clipper.com).***

### ***About the Author***

***Anne MacFarland is Director of Enterprise Systems Research with The Clipper Group.*** Ms. MacFarland specializes in strategic solutions offered by enterprise systems and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at [AnneM@clipper.com](mailto:AnneM@clipper.com) or at 781-235-0085 Ext. 28. (When you hear the automated attendant, please dial "1-28".)***

### ***Regarding Trademarks and Service Marks***

***The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and "clipper.com"*** are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

### ***Disclosure***

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

### ***Regarding the Information in this Issue***

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.