

# Storage Networking Times

## Issue 7

January  
2008

### Inside this issue:

Spotlight on: Storage Management Initiative 2

Analyst Briefing: Counting the Cost of Losing Data 3

Education: The Business Benefits of Fibre Channel over Ethernet (FCoE) for Server I/O Consolidation 4

Education: Only One Copy! 5

Myths Uncovered: Storage Security—Who Cares? 7

Industry Events 2008 10

## Letter From the Chair

Juergen Arnold, Chair, SNIA Europe, [eurochair@snia.org](mailto:eurochair@snia.org)



First and foremost I would like to wish you all a successful and prosperous 2008. This is the time of year when we all make new resolutions so why not make it one of your goals to ensure that

over the next 12 months your data storage infrastructure is going to be even more efficient than it is today? Efficiency is at the source of many benefits in the datacentre such as lower costs, increased productivity and of course reduced CO2 emissions.

In this issue of *Storage Networking Times* we bring you the latest news on a number of developments, namely around the Fibre Channel over Ethernet (FCoE) and the

Network Data Management Protocol (NDMP). Our contributors also discuss hot topics such as storage security, data deduplication and the real cost of losing data.

SNIA Europe will kick off its events calendar with the first Academy appointment for 2008 being held in Dubai on 5th February; this will be quickly followed by Zurich on 26th February. For 2008 we are also pleased to announce that we have added three key locations to the Academy itinerary: London and Paris, both taking place in May, and Milan, slated for September. You can find all the details at [www.storage-academy.com](http://www.storage-academy.com).

In addition to this programme we are actively working to bring you a number of webcasts following the success of last year's

*(Continued on page 2)*

## Standards Update: SNIA's New NDMP Technical Working Group

David Dale, Vice Chair, Storage Networking Industry Association, [info@snia.org](mailto:info@snia.org)



In December, the Storage Networking Industry Association announced the creation of a new Technical Working Group (TWG): "SNIA Creates Technical Work Group to Enhance NDMP

Standard Through Software Development" ([www.snia.org/about/news/newsroom](http://www.snia.org/about/news/newsroom).) This signalled an expansion of SNIA's portfolio of software development efforts, collectively targeted at expanding the universe of storage, storage management and data management interoperability tools.

For those of you who are not aware of NDMP, here's a little background...

NDMP stands for Network Data Management Protocol and was created by a multi-vendor taskforce; the specification was made available through [ndmp.org](http://ndmp.org) ([www.ndmp.org](http://www.ndmp.org)). Today, NDMP is the *de facto* standard protocol for NAS backup.

NDMP is fundamentally about data movement and its specification currently defines two services:

*(Continued on page 2)*

## Letter From the Chair

(Continued from page 1)

IP Storage Webinar, which attracted participants from throughout Europe. And of course the planning is already under way for SNW Europe 2008 taking place in Frankfurt from 27th–29th October. More details can be found at [www.snweurope.com](http://www.snweurope.com)

With so many storage and information management educational and peer networking opportunities on this year's cal-

endar, I sincerely hope to see many of you at our events. As for my new year's resolutions, I am looking forward to working with the rest of the Board, and all the committee members to make SNIA Europe an even better information resource for IT professionals and the industry overall.

Here's to a great 2008!

---

## Standards Update: SNIA's New NDMP Technical Working Group

(Continued from page 1)

- A *data* server—which either reads from disk and produces an NDMP data stream (in a specified format), or reads an NDMP data stream and writes to disk, depending upon whether a backup or restore is taking place;
- A *tape* server—which reads an NDMP data stream and writes it to tape or reads from tape and writes an NDMP data stream, depending upon whether a backup or restore is taking place. All tape-handling functions, such as split-image issues, are dealt with by this service.

Each service has a separate state diagram that dictates its behavior, e.g. the tape server can enter the pause state while tapes are being changed by the NDMP client. NDMP messages are categorised into distinct groups or NDMP interfaces, such as SCSI, *config* and *tape*, and can trigger state changes.

The NDMP standard has evolved over time and has continued to mature with version 4. However, there is currently no consistent software development kit (SDK) that developers can use to implement this version of the specification. The new SNIA NDMP TWG is initially chartered to create an SDK for NDMP v4.0.

This software development effort was kick-started by significant code contributions, comprising the complete NDMP v3 reference implementation code base, donated to SNIA by NetApp and EMC. Other vendors who joined the NDMP TWG include Hitachi Data Systems, Pillar Data Systems, and Sun Microsystems.

The work product of the new NDMP TWG will help accelerate the already widespread adoption of the NDMP standard by storage vendors, and provide IT users in markets worldwide with more choices in data protection solutions.

## Spotlight On: Storage Management Initiative A Journey from the SMI Specification to the Management Framework

Frank Bunn, Chair, Storage Management Initiative (SMI), [smi-europe-chair@snia.org](mailto:smi-europe-chair@snia.org)



In the last few years, the SMI-S specification has established itself as THE comprehensive interface for the management of heterogeneous storage systems and is supported by all the leading providers of storage and storage management technologies. The ISO and ANSI have also certified SMI-S as a recognised standard. SMI-S provides the basis for interoperability and vendor-independence in business-critical storage environments.

But defining the standard is not the whole story. The time to market is a major challenge. It still takes too long for standardised implementations to be developed by the vendors and offered to the end users. This timeframe must be dramatically shortened to enable customers to experience a real added value from SMI-S. In order to address these demands, the SNIA formed the technical working group Management Framework in October 2006. The aim of the group is to define services and interfaces as core features

(Continued on page 6)

## Analyst Briefing: Counting the Cost of Losing Data

Hamish Macarthur, Macarthur Stroud International, [Hamish.macarthur@macarthurstroud.com](mailto:Hamish.macarthur@macarthurstroud.com)



### The Value of Information

The value of information is being realised and affecting everyone across the world. The impact of the sub-prime market shows that financial deals are done where buyers are not fully aware or informed of their liabilities. The packaging of debt through different financial vehicles has resulted in

the global financial markets having to re-evaluate their risks and liabilities, after the event.

Similarly in the UK, the government has been plagued with incident after incident of lost data; the loss of 25 million personal records relating to child benefit, the loss of a disk drive containing 3 million new drivers' records and the loss of 6,500 pension records from a pension company to the tax authorities. The care and attention with which personal and corporate data is maintained is under public, corporate and legal scrutiny.

Understandably, organisations do not wish to admit to any weaknesses in their systems and processes. But unless companies take a closer look at their systems and processes, the risks that are being run can be significant, with more examples of lost data coming to public attention.

From research carried out by Macarthur Stroud International, we know that CIOs and system administrators recognise that public awareness of data loss can affect their business. The evolving regulations mean that persons can now be prosecuted for poor information governance. After all, it is a reflection of the accepted standards of corporate governance. In addition, individuals are recognising that if they are seen to be responsible for such occurrences of data loss, they will lose the respect of their peers, they will have difficulty in progressing their careers and they will experience difficulties in finding a suitable role in a new company.

### Risks from Within

The greatest risk of fraud has always been from within an organisation, yet the focus of attention is consistently at keeping wrong-doers out. Having lost data on a disk or tape and announcing that there is no likelihood of the data falling into the wrong hands because it must be somewhere in the organisation, does not instil any confidence that there is no risk. With ever increasing volumes of digital data being stored on systems across the world, in different jurisdictions, the onus on businesses and the Chief

Information Officer is ever increasing to ensure that the information assets of an organisation and the associated personal data are absolutely secure and safe. This relates to both current application data and archive data.

The cost of implementing new procedures is invariably given as one reason for not reviewing the risks. "Are the risks real" and "it will not happen to me" are other such thoughts. But it is the law of the unexpected that will always catch each and every one of us out. Therefore, the issue is, what is the cost taking no action?

The challenges are increasing. With the growth of mobile computing and the reach of customers into online systems, the firewall boundary is becoming blurred. Virtualisation will deliver ever increasing benefits, but the way information is being processed and managed must be clearly understood.

### Need to Develop Information Security Policies and Practices

Users need to recognise that they must develop a total set of Information Security Policies and Practices. This must take into account firewalls, antivirus, identity management, data protection and system management. Following ITIL, COBIT or ISO 17799 guidelines can assist in the steps to be taken. But these have to be actioned and implemented.

The identity management, antivirus and firewall considerations must come further into the network, not just left on the periphery. Applications accessing data resources need to follow accepted routes. If the requests are coming from unknown routes within the firewall, does this mean that the requesting applications are all approved?

Logging activities and requests to produce a comprehensive audit trail becomes more important for all activities across the network. This relates to system changes, software updates and data movement as well as to application processes.

Data protection practices are designed to keep systems operational. The trend to using disk-based backups means that there are many images of data on disk drives as well as on tapes. Encrypting all these data images needs to be carefully considered. This spans the ever increasing volumes of archive data which must be maintained for legal, regulatory, contractual or corporate governance reasons.

(Continued on page 6)

## Education: The Business Benefits of Fibre Channel over Ethernet (FCoE) for Server I/O Consolidation

Gilles Chekroun, Technical Chair, SNIA Europe BeNeLux Committee, [beneluxt-chair@snia.org](mailto:beneluxt-chair@snia.org)



Data Centre I/O consolidation is quickly becoming a reality and the adoption of 10 Gigabit Ethernet is helping server administrators implement I/O convergence.

The 10 Gigabit Ethernet network connectivity is driven by the rapid evolution of server hardware with quad-core multi-socket CPUs and by the growth of server virtualisation software implementation allowing multiple virtual machines to run over the same hardware.

In turn, these virtual machines are driving the need for more access to storage and for more I/O bandwidth.

### FCoE in the Server

Fibre Channel over Ethernet (FCoE) is a technology that transports Fibre Channel protocol over a lossless Ethernet network. This technology allows Data Centre servers to natively connect to Fibre Channel SANs, thus protecting investments in storage equipment, management and administration.

The ability of FCoE to integrate seamlessly in the Data Centre allows an evolutionary approach to server administrators. The transport of IP and Fibre Channel over a single 10 Gigabit Ethernet cabling infrastructure reduces the number of server NICs and simplifies network topology with less need for power and cooling.

The business benefits are among others:

- Fewer interface cards per server, therefore smaller servers (IRU) and so fewer cables
- Power and cooling reduction
- One set of access FCoE switches with LAN and SAN connectivity
- Better performance with 10 Gigabit Ethernet connections
- Seamless integration with existing Fibre Channel infrastructure and management

### Simplifying Network Topology

FCoE simplifies the Data Centre network topology by reducing the number of server interfaces. Mission critical applications drive the need for multiple NIC and HBA interfaces on the server to create a redundant access to LAN and SAN switches. This number of adapters per server also drives the need to have large servers (3 or 4 U) and so reduce the number of servers per rack.




Figure 1: With 4 or 5 adapters per server, administrators are forced to use 3 or 4 U machines and 2 x 2 kind of Top of the Rack switches—LAN and SAN.

When FCoE access switches are used, a single pair of adapters and a single pair of cables are needed to connect every server to both Ethernet and Fibre Channel networks.




Figure 2: With one or 2 FCoE adapters per servers, Administrators can use 1 U server and simplify network topology

### Fewer Interfaces, Fewer Cables, Less Power, Less Cooling...And Greater Performance

FCoE technology is 10 Gigabit Ethernet-based and can run over copper cable like the new kind of twinax cable with SFP+ ends. IT managers can now fully realize the I/O consolidation and give every server access to centrally managed storage without having to invest in additional HBA adapters.

(Continued on page 8)

## Education: Only One Copy!

Glyn Bowden, Storage Engineer, UBS AG



Picture the scene, a large family gathered around a photograph album enjoying the many memories the pictures inspire.

Now one of the children leans over and points to a picture that they like, and three of the gathering all agree. With this news, the owner of the album promptly takes out the photo, makes four copies and places them all back in the album so

each person has their own copy. Sound ridiculous? Well every day, businesses are doing exactly this. The album is the shared resource of a storage pool and the photos are the files we send to each other every day, in e-mail, over instant messaging and in our home directories. The difference with the real life scenario is that our album is a far more expensive resource and it's usually a lot bigger than a family's collection of pictures. Imagine a family of thousands instead of tens and you can see the way in which the problem grows.

There has been a real explosion in the volume of stored data in recent years and this is largely due to the duplication of this data. So wouldn't it be remarkable if all of those duplicates could be identified and consolidated? Well, the technology is available to do this today; all that is missing, in most cases, is the understanding of how best to employ it. The first hurdle to overcome is the fear factor so let's deal with that first. In order to do that it is important to know how de-duplication works at the fundamental level.

### What is Data De-Duplication?

One definition of "what is a duplicate" is based upon the method used to identify duplication which may include meta-data, content, hashing, or other comparative techniques—the process of examining a data-set or I/O stream and removing duplicate data while maintaining data integrity and authenticity.

So, based on this definition, the end result is that when performed, patterns of data are only stored once and referenced by instances that would otherwise require a copy of that data. The granularity of the de-duplication depends on the method and technology used to achieve this. The process can in fact happen in one of two ways. Either before the data is written to the de-duplication system disks otherwise referred to as "inline" or "in band", or after the data is already stored on a media, also known as "out of band" or "post-process". Regardless of where the process occurs it needs to be transparent to the client accessing that data. Each of these methods has pros and cons, but both support the final goals of making more efficient use of storage media resources, enabling better use of network bandwidth for transmitting the processed data

set for disaster recovery or even migration and finally, to help build a greener data centre thanks to reducing the requirement for additional storage capacity.

### The Challenges with Data De-Duplication

Alongside these benefits lie some challenges. To start with, the system is essentially virtualising data sets and with any virtualisation comes another layer of abstraction. However, simplicity and keeping the de-duplication process transparent to the user should be design goals of any viable solution. Another difficulty faced by the early adopters was the complication of gathering metrics about the process. In order to deem the de-duplication exercise worthwhile, it must be possible to measure how much efficiency is gained. Solutions exist that report on what has taken place, and this feature should be considered when choosing a technology. This leads to the chargeback question. If multiple customers are affected by a single de-duplication process, who pays for the used blocks and how is that measured?

Finally, the process, by its very nature, can impact the performance of the storage device or network and could have a different use or impact than an existing service. It is important to understand the impact of the de-duplication process on the ability to deliver the storage service customers expect.

Some of the key questions that should be addressed by any investigation into data de-duplication are as follows:

The first is where the processing should occur. When inline de-duplication is used, only the de-duplicated data set is stored, making this method the most efficient in storage capacity, as the media does not need to be sized to cope with the full data set. This will also make any kind of snapshot technology more effective on the storage device, as it is impossible to capture the duplicate blocks in a snapshot. Inline de-duplication also happens in real time as the de-duplicated data is transmitted from client to storage or between storage tiers. However, this requires a considerable amount of processing power from the de-duplication device which will undoubtedly add latency to the transmission of the data. In applications that are latency sensitive, such as transactional or primary storage, this can cause issues. There is also the question of redundancy and how efficient the process can be if data is passing over more than one physical path. In order to be as efficient as possible, each redundant path would need to be aware of the data being processed on the other. It is unlikely that the database used for identifying duplicates in this case would be as efficient as that used on a device once the data is at rest. In this instance the metadata used to identify duplicates can be reproduced from simply scanning the data set. This

(Continued on page 8)

## Spotlight On: Storage Management Initiative

(Continued from page 2)

of a storage management application and make them available to all vendors as standard. A draft of the Management Framework Reference Architecture has been available for public review since November 2007.

Put simply, the SMI-S Management Framework is a collection of common components which can be leveraged as re-useable services in any storage management application. From a technical point of view, the framework is situated between the vendor-specific management application and the agents representing the storage elements to be managed. The *Management Framework Reference Architecture* itself is split into three layers, each offering different basic storage services. The vendors do not have to concern themselves with these anymore but can simply leverage the services provided. *The Infrastructure Layer* provides agent discovery, collation of management data and the receipt of alerts and events. Scheduling, data modelling and policy and security services are made available in the *Core Services Layer*. The *Storage Domain Specific Services Layer* defines extensions to the Core Services, which are fundamentally

for specific storage domains, but on the other hand are not needed by others.

At the end of the day, the Management Framework holds advantages for vendors and for end users alike. The vendors benefit from the SNIA's investment in standardised infrastructure services. Routine features to manage new storage systems based on SMI-S can now be implemented considerably faster and technologies are less prone to error. As a result, storage vendors save time, money and resources which they can invest in the continued developments of their management applications in order to offer the customers a new added value and set themselves apart from the competition.

Customers benefit from the faster provisioning of new standardised management functionality. They receive a much wider support from the providers of interoperable applications. This leads ultimately to considerably lower purchasing and management costs for heterogeneous storage environments.

---

## Analyst Briefing: Counting the Cost of Losing Data

(Continued from page 3)

Understanding how the system resources are being used, where the information is resident and what elements of the system are redundant are important steps in the process. This contributes to better information management as well as contributing to better utilisation of power, cooling, office space and cost containment.

And when it comes to archived information, can it clearly be proven that it has not changed or been tampered with since it was created. Such forensic considerations become important if put to the test by courts of law or regulators some time in the future.

### Loss of Information Could Cost You Dearly

Embarking on this journey to secure the system operations and information assets will mitigate risks for each and every organisation. Recognising that there are risks is the first step, followed by identifying possible weaknesses, completing an appropriate risk assessment and identifying the necessary actions. Implementing appropriate solutions will help organisations to minimise embarrassing information management practices being exposed to the gaze of television and the press.

Information has a value. Respect this and protect the data. Otherwise, loss of information could cost you dearly.

## Myths Uncovered: Storage Security—Who Cares?

Bjarne Madsen, Chair, SNIA Europe Nordics Committee, [nordicschair@snia.org](mailto:nordicschair@snia.org)



One of the great advantages being an 'oldie' in the IT industry is the pleasure of being witness to the changes of focus and priority for various topics over the years. In the past, storage security was not high on the IT department's agenda mainly because the storage was strictly attached to a single host and associated applications. Later on, storage became shared and accessed through networks such as SANs and LANs, and that was when IT administrators first realised the importance of a strong security strategy to protect their storage infrastructures. But did they follow through? Not really. Storage security was mostly still a 'tick in the box' feature—customers didn't understand it or didn't care about it much. All that mattered was that the solution could provide data storage to the enterprise while minimising the risk of data loss and corruption. Today's solutions usually involve the geographic distribution of data for business continuity for example; but multiple points of entry also mean multiple opportunities for security breaches. After many meetings with end users I strongly believe that the need for security around the storage infrastructure is highly recognised, but mostly not implemented.

So what will change the status quo? The major drivers that will take storage security from buzz word to reality at the moment seem to be the increasing information management regulations with which companies must comply together with the rapid growth in the rate of security incidents throughout the industry. Regulatory mandates such as the Sarbanes-Oxley Act of 2002, the California Database Protection Act of 2001, the Gramm-Leach-Bliley Act, the Health Insurance Portability and Accountability Act (HIPAA), the Basel II accords, Markets in Financial Instruments Directive (MiFID) and EuroSox are an additional catalyst for applying due diligence in the security decision and implementation process. These laws impose strict requirements on enterprises to establish or identify, document, test and monitor necessary internal control processes. Because information technology supports most, if not all, of these processes, these laws significantly affect companies' security strategies. As a result these new regulations force security designers and architects to impose and maintain suitable security controls throughout their enterprises.

(Continued on page 9)

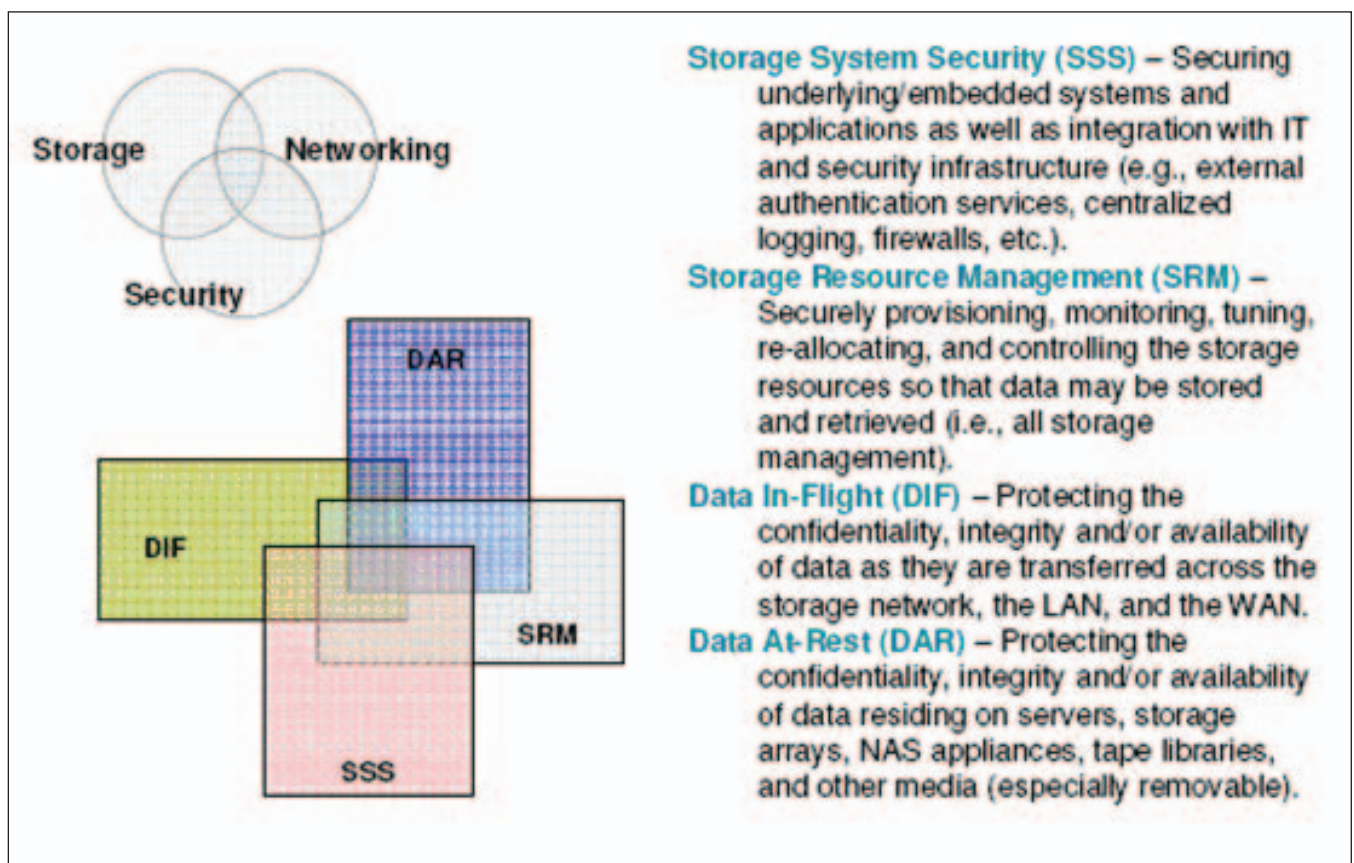


Figure 1. SNIA's view of storage security

## Education: The Business Benefits of Fibre Channel over Ethernet (FCoE) for Server I/O Consolidation

(Continued from page 4)

Inside this 10 Gigabit Ethernet pipe, administrators can use up to 8 Virtual Lanes, based on the User Priority in the VLAN 802.1q header, with VLANs for security, QoS and 802.1p Priority Flow Control (PFC) to provide the lossless Ethernet needed for FCoE.

The ability to transport Fibre Channel and IP over the same cable reduces the number of cables in the rack by a factor of two and having fewer adapters means consuming less power.

### Integration in Existing Environments

FCoE doesn't replace Fibre Channel but allows extension of native Fibre Channel networks up to the servers with 10 Gigabit Ethernet transport. FCoE switches in the top of the rack allow aggregation of 10 Gigabit Ethernet connections from the servers and split them to classical Ethernet LAN and Fibre Channel SAN thus preserving existing investments that have been made in the company.

### Why 10 Gigabit Ethernet?

With the rapid growth in CPU technology and in particular, the multi-socket, multi-core CPUs, servers are becoming very powerful machines.

Next to this, there is a big Data Centre trend in server virtuali-

sation so we see multiple virtual machines running in each server. All these machines want to access storage as well and 10 Gigabit Ethernet is the underlying technology to enable I/O convergence and multiple I/O streams on the same cable.

In the Ethernet world, there is nothing between 1 Gigabit Ethernet and 10 Gigabit Ethernet like we have on Fibre Channel 2-4-8 Gbps and 1 Gigabit Ethernet is definitely not enough these days for the growth we see in server performance.

### Unified Fabric

Until now, Data Centre networks were primarily two separate and very different networks: LAN and SAN. The first phase adoption of FCoE technology by server administrators will allow a unified fabric in the Data Centre server racks with lower costs and higher performance while reducing power consumption and enabling seamless integration with existing LAN and SAN networks.

The second phase will be adoption of FCoE by storage vendors and the ability to provide 10 Gigabit Ethernet FCoE connections to storage arrays. This may already happen during the second part of 2008, since some demonstrations have already been made at SNW US and more recently at SNW Europe 2007.

---

## Education: Only One Copy!

(Continued from page 5)

is only possible when the data is already at rest on the device. Finally, inline de-duplication will only affect new data being transmitted, so for existing data-sets, unless a migration or retransmission occurs, this will not be effective.

Post-processing on the other hand offers the advantage of working with the full data set and the ability to re-examine any part of it at any time. This approach is also a batch process, meaning that it happens at regular intervals or on-demand, but not in real time. This allows the administrator to choose when the best time to perform the process is, for example prior to taking a snapshot or backup, or when the data is in less demand and the processing power of the device can be used for de-duplication. The disadvantage of this approach is that the planning of capacity to handle the unprocessed data set can lead to wasted storage capacity. Also this post-processing must be carefully timed so that snapshots and backups only occur once the de-duplication process has completed. This can be quite tricky to estimate until some meaningful metrics have been gathered.

### What to Consider?

The questions that span both implementations are those of positioning, interoperability and risk. If a tiered storage model is in place, there is somewhat of a dilemma. One might expect that the top tier of storage that uses expensive media would benefit the most from de-duplication. However, performing this process might degrade performance significantly enough to impact delivery of service levels expected for this tier. The disaster recovery or near-line locations might be a better choice from a performance perspective, but then the disks or tape might be more inexpensive here and so the savings perhaps not as great. However, the near-line or archive devices usually are at a location where data is being consolidated. In this instance, data de-duplication can reap huge benefits in extending retention capabilities yet reducing the capacity requirement and creating efficiencies in the amount of disk or tape and even storage enclosures that are required. And in addition, less capacity means less power and cooling consumption which helps the data centre costs and the environment!

(Continued on page 9)

## Education: Only One Copy!

(Continued from page 8)

Vendor interoperability is also important. In order to protect the data from a vendor lock-in scenario it is important that all related vendor equipment can read the de-duplicated data or that the data is exported or migrated in a fashion that is open and able to be read by third party devices.

Some of these questions are still being addressed by the industry, but exciting de-duplication products already exist that offer an easily deployable solution to a very pertinent problem. Data de-duplication is good for the bottom line, good for administrators and good for the environment!

### About SNIA and Data De-Duplication

The SNIA's Data Management Forum (DMF) has formed a new special interest group to focus specifically on data de-duplication. This new group is called Data De-Duplication & Space

Reduction (DDSR) Special Interest Group (SIG).

The SNIA DDSR SIG is dedicated to advancing space reduction in all networked storage technologies. This mission addresses the continued exponential growth of data and the increasing need for storage technologies utilising data de-duplication and other space reduction techniques. By defining and promoting efficient networked storage solutions and common implementations, the DDSR SIG is enabling sustainable data storage operations that reduce both storage costs and the environmental impact of data centre infrastructures.

For more information about the DDSR SIG—please visit [http://www.snia.org/forums/dmf/programs/data\\_protect\\_init/ddrsig](http://www.snia.org/forums/dmf/programs/data_protect_init/ddrsig) and for more information on the Data Management Forum contact [snia-dmf@advisorygroups.snia.org](mailto:snia-dmf@advisorygroups.snia.org)

---

## Myths Uncovered: Storage Security—Who Cares?

(Continued from page 7)

### What is Storage Security?

Storage security represents a major component of the overall information security plan for a data centre and a business. Consequently, business policies and practices must augment any hardware- or software-level security model, including network and system security. Security however, is not a simple commodity that you can order by weight and bolt onto an IT infrastructure. Security considerations permeate every aspect of your IT Infrastructure—from application to the management of technology and of people.

Another perception is that when security has been implemented we are done. Sorry—not true! Storage security requires specialised maintained knowledge, careful attention to detail, and ongoing reviews to ensure that the storage infrastructure continues to meet the organisation's evolving needs. Measuring security is difficult—how safe are we at any point? Unlike processor speed or storage capacity, we do not measure security in simple units—except after an incident when we can objectively demonstrate that the deployed security mechanisms were inadequate. As a result, enterprise security has traditionally been handled reactively in a fashion which is somewhat reminiscent of the old saying 'they shut the stable door after the horse had bolted'.

An exhaustive storage security strategy involves several areas; even the simple movement of data from point to point either through a network or to different media such as tapes and CDs, requires specific processes and procedures along with the appropriate encryption of the information. In fact, data should

be protected both as Data In-Flight (DIF) and Data At-Rest (DAR); see Figure 1. for SNIA's view of storage security.

Generally speaking storage security includes the following elements:

- Authentication—validates user, system and/or application
- Access control—determines what can be seen
- Integrity—validates that data is in the original form it was stored in
- Confidentiality—use of encryption to protect content
- Secure key management—keys must be available *whenever* and *wherever* data is accessed

In 2007 we noticed that data protection and ILM were among the most popular projects undertaken by user organisations. To continue and complete these projects thereby fulfilling regulatory and specific SLAs you need to integrate storage security into the overall company strategy for information management. Due to the tight integration of existing IT challenges such as data protection, information growth, and compliance, and their associated and increased costs, 2008 might be the year where we finally see well-developed and documented IT strategy plans across the IT community.

### For More Information:

SNIA Security Technical Work Group (TWG)  
[http://www.snia.org/tech\\_activities/workgroups/security/](http://www.snia.org/tech_activities/workgroups/security/)  
Storage Security Industry Forum (SSIF)  
<http://www.snia.org/forums/ssif/>

**Industry Events 2008****February 5**

SNIA Europe Academy  
Dubai

**February 5-7**

Storage Expo Italy  
Milan

**February 26**

SNIA Europe Academy  
Zurich

**February 26-27**

Data Centre World  
London

**March 3-4**

Information Management  
Taastrup, Denmark

**March 4-9**

CeBIT  
Hannover

**March 19-20**

Storage Expo Belgium  
Brussels

**March 25**

Disaster Recovery Planning  
London

**April 1**

SNIA Europe Academy  
Copenhagen

**April 3**

SNIA Europe Academy  
Stockholm

**April 15-June 12**

IDC Storage Roadshows  
Various EMEA Cities

**May (date tbc)**

SNIA Europe Academy  
Paris

**May 20**

SNIA Europe Academy  
London

**June 4**

SNIA Europe Academy  
Moscow

**June 11-12**

Storage Expo Spain  
Madrid

**September 18-November 25**

IDC Business Intelligence Roadshows  
Various EMEA Cities

**October 16-17**

Storage Expo UK  
London

**October 27-29**

Storage Networking World Europe  
Frankfurt

**November 13-14**

Storage Expo Netherlands  
Utrecht

**Full event calendar with weblinks  
on: <http://www.snia-europe.org>**



**STORAGE NETWORKING  
INDUSTRY ASSOCIATION EUROPE**

Erico House, Suite 406  
93-99 Upper Richmond Road  
London SW15 2TG  
United Kingdom

Phone: + 44 (0) 20 8785 5555  
Fax: + 44 (0) 20 8785 5624  
Email: [euroinfo@snia.org](mailto:euroinfo@snia.org)  
Web: <http://www.snia-europe.org>

**Subscribe to Storage Networking Times**

**[www.snia-europe.org/subscribe](http://www.snia-europe.org/subscribe)**