WHITE PAPER

NEXSAN™
by imation

SECURE ARCHIVE

MAXIMIZE PRIMARY STORAGE ROI & DATA PROTECTION

imation

INTRODUCTION

"Do more with less" seems to be the dogmatic decree that confronts today's IT professionals as they struggle to reconcile constrained budgets with apparently limitless growth in unstructured and structured data. With costly primary storage rapidly filling up (and slowing down), the need to migrate less active data onto a more cost-effective storage tier is clear (see Figure 1 below).

Storage optimization—matching storage to data's performance, capacity, and accessibility needs—entails moving infrequently-accessed data off expensive, high-performance primary storage to less costly, lower-performance secondary storage. This maximizes primary storage ROI by freeing up more capacity, boosting performance, shrinking the backup window, and reducing its cost by ensuring all of the data that remains on it actually needs to be there. So far, so good . . .

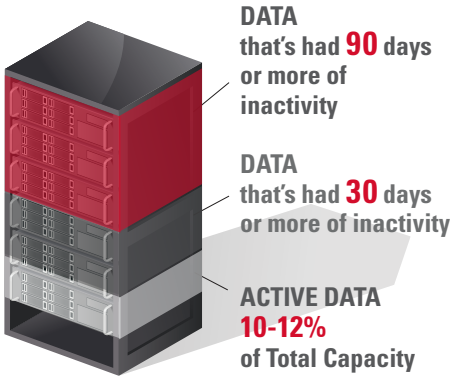

FIGURE 1
ONLY A SMALL PERCENTAGE OF DATA ON TYPICAL PRIMARY STORAGE SYSTEMS IS CURRENTLY ACTIVE

ARCHIVES, BACKUPS . . . AND DATA LOSS

If you don't plan well, moving data to the secondary tier presents new challenges; more administration headaches (how is that less-active data moved, and to where, and how do users find their data without contacting the IT department), and the additional expense of deploying and managing an extra backup solution . . . wait, why is more backup needed?

Because this secondary tier functions as an **archive** (storage employed for long-term retention and future reference of data no longer actively used for everyday operations), it contains **original** files that have been **moved** from their initial location (in this case, from primary storage) and placed elsewhere for secure preservation. Thus one or more **copies** of the archived data is needed, which can then be used to restore the original data if it's lost or damaged beyond repair.

Which brings us to the elephant in the room: What about data loss? Any storage optimization and archive initiative must address data loss.

Active files on primary storage are constantly being accessed and opened, thus any data corruption or missing files quickly become evident. Further, any issue in missing or corrupted data is made worse by the frequent snapshots and backup of your primary data. But what about files that are less frequently opened (whether on primary storage or in an archive)? It could be weeks, months or even years until an organization discovers a file is damaged...or just simply gone.

That's the fundamental problem with conventional storage solutions — IT pros only know there's a problem when they go to access a file and it doesn't open, or isn't there at all. And remember, if the backup was done after the corruption or data loss, it does nothing to repair or restore a file to its undamaged, original state.

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

NEXSAN™
by imation

HOPE IS NOT A STRATEGY

In its landmark 2007 study , the world-renowned CERN tested 3,000 servers attached to RAID subsystems[1]; in three weeks it found 500 instances of corrupted files in 17 percent of the RAID arrays. In short, the equivalent of one in every 1,500 files had become corrupt.

On a similar note, in February 2013 Oracle featured an article that discussed the dangers of silent data corruption[2], noting that "[It] can happen without warning and can be defined as the non-malicious loss of data resulting from component failure or inadvertent administrative action. Silent data corruption occurs when invalid data is read or written rather than resulting in a failed I/O operation. This type of corruption is by the far the most cataclysmic, and there are no effective ways to detect it without **end-to-end integrity checking.**

The sheer volume and variety of ways in which data can be breached is daunting, including: silent data corruption due to hardware failure and software failure, malicious attacks from cybercriminals or human error such as accidental file deletion or overwriting by employees, etc. Yet for all too many data center managers, their response to these threats seems to be a strategy based on . . . hope.

Unless they can answer the following four critical questions, their data protection strategy is tantamount to just hoping that nothing goes wrong:

1. **How do you know if all your files are in your backup or archive?**
2. **How do you know if there is a second copy of all your files at your remote site?**
3. **What is the health (integrity) of your files at each site?**
4. **If the files are different, which file is the correct one?**

For IT professionals that employ standard archive or backup systems, these questions simply cannot be answered because the information that's required to respond to such queries is not readily available. Conventional archive or backup solutions do not have the capability to monitor the availability and health of each file, and manually verifying the existence and integrity of each file (by opening millions, perhaps billions of files) is a practical impossibility.

The answers to these questions — and indeed the answer to this dilemma — comes from purpose-built secure archive solutions, which are specifically designed to provide maximum data security, integrity and privacy from the moment a file is ingested into the archive.

[1]Data integrity, Bernd Panzer-Steindel, CERN/IT Draft 1.3 8., April 2007

[2]How to Prevent Silent Data Corruption, Martin Petersen and Sonny Singh, published February 2013, http://www.oracle.com/technetwork/articles/servers-storage-admin/silent-data-corruption-1911480.html

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

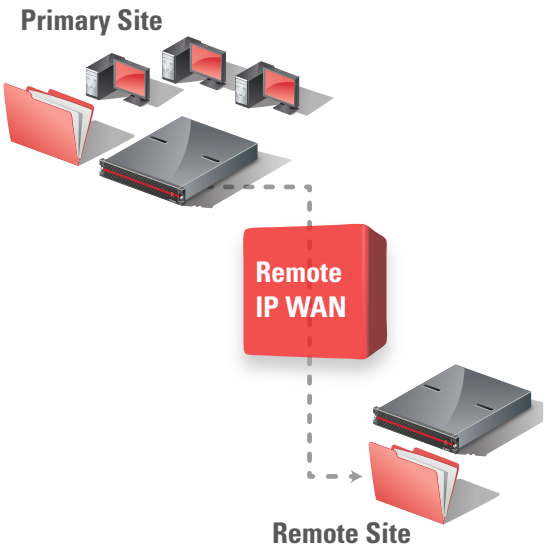

FIGURE 2
SECURE ARCHIVE AUTOMATICALLY KEEPS TWO COPIES OF EVERY FILE, EITHER LOCALLY OR REMOTELY

HOW SECURE ARCHIVE DELIVERS UNRIVALED DATA PROTECTION

As noted in the quote above, end-to-end integrity checking is the only way that silent data corruption can be detected (and thus corrected). Any archive solution that does not include this capability cannot credibly claim to be "secure" because it lacks the means to comprehensively ensure data protection.

The ingestion process of a secure archive solution should begin with a fingerprint of each file - a gold standard for original file integrity - and a duplication of each file; this provides a copy of every original file and its contents and metadata, stored either in a separate RAID disk set within the local secure archive system or on another secure archive system in a remote location—for example, an organization's main office, in the cloud, etc. (See Figure 2 left).

When files are moved from the primary storage system onto the secure archive, a small shortcut for each file should be left on the primary storage, enabling end users to immediately access any files they need (unlike a conventional backup). This not only makes the secure archive transparent to end users, it also maximizes data security by ensuring the archive is not directly accessed by end users for browsing or any other actions.

Maintaining a second, redundant copy of every original file ingested enables a secure archive solution to perform crucial comparative file analyses through utilization of these two powerful data protection technologies:

- File serialization
- File fingerprinting

What's more, the integration of redundant file copies within a secure archive system **eliminates the need for overt backup and restore operations to be performed.** Because these copies are subject to ongoing file serialization audits and file integrity audits (see below), they provide significantly greater data protection than is possible with a conventional backup/restore solution.

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

FILE SERIALIZATION AND AUDIT

Ideally, every file ingested into a secure archive will have a unique serial number assigned to it (the same serial number to be used for both copies of a file — the original and its redundant copy). This file serialization enables the secure archive to periodically verify the existence and location of every file in the archive, both at the archive's primary site and at its secondary site (often a remote location). Think of secure archive file serialization and audit as similar to the asset tags and tracking systems used by companies to identify and control their physical assets – PC's, servers, industrial tools, etc. – that are critical to a company's bottom line.

For example, every three months a secure archive might employ these serial numbers to perform a file audit, checking to make sure that the millions of original files ingested into the primary archive's disk set are still where they should be, and that their corresponding copies in the secondary archive's disk set are also in place. Should a missing file be detected, the archive can notify the administrator and automatically replace it using its serialized redundant copy (see Figure 3 below).

FIGURE 3
UNIQUE FILE SERIAL NUMBERS ENABLE EASY IDENTIFICATION OF MISSING FILES.

SERIALIZATION AUDIT SUMMARY:

- Files entering the secure archive are given unique serial numbers
- Every file is periodically checked to ensure it is still in the archive
- Checks are performed at both primary and secondary archive sites
- Reports any missing data, provides an audit of file availability

The end result of these serialization audits is that data availability is confirmed, enabling IT professionals to unequivocally answer the basic query embodied in questions 1 and 2 above: ***"Are all of my files there?"***

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

FILE FINGERPRINTING AND INTEGRITY AUDIT

To guarantee file-level integrity within the archive, a secure archive should generate a unique gold standard "fingerprint" of each file when it is ingested and when it is copied. Subsequent copies of the original file, for example, stored in a remote location, can be validated as a correct copy of the original file after the copy's fingerprint is compared to the original's fingerprint. Modern, best-in-class secure archive solutions perform file fingerprinting by combining two hashing algorithms, such as MD5 and SHA1, on the same file.

Similarly to the file serialization audit process described above, these fingerprints enable the archive to periodically audit the integrity of each file against its original fingerprint, in order to confirm the data has not been changed (due to silent data corruption, disk error, virus, tampering or replication error). Should this process reveal a file has been altered, the audit reports the corruption and the archive automatically replaces the corrupted file with its undamaged copy (see Figure 4 below).

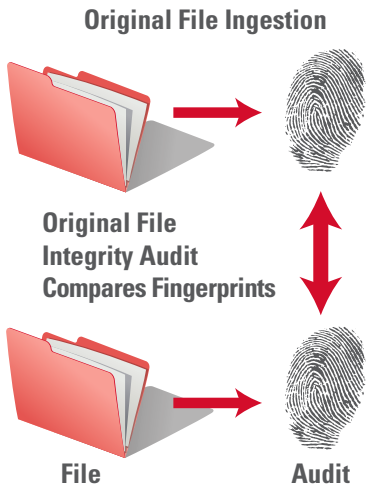

FIGURE 4
UNIQUE FINGERPRINT ENABLES RECURRING CHECKS FOR CORRUPTED FILES

INTEGRITY AUDIT SUMMARY:

- Files entering the secure archive are given unique fingerprints; archive generates a new fingerprint whenever a file is saved
- Keeps tracks of data to make sure the file has not changed (due to silent data corruption, disk error, virus, tampering or replication error)
- Reports and repairs any data corruption
- Provides an audit of file integrity

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

Archives Key Strategy for Implementing Cost-Effective Storage

Archives Key Strategy for Implementing Cost-Effective Storage

In 2014, active archives will increasingly become a top strategic purchasing intention of CIOs as they grapple with implementing cost-effective methods to store, retain and retrieve data.

Active Archive: Top Five Data Predictions for 2014; David Cerf, Active Archive Alliance.

Smarter Storage Will Double Administrator Productivity

By 2016, storage system functional enhancements will double storage administrator productivity on a petabytes-per-full-time-equivalent basis.

Gartner, May 2013, Market Share Analysis: Attached Storage and Unified Storage, Worldwide, 2012.

The end result of these file integrity audits is that data integrity is confirmed, enabling IT professionals to affirmatively answer the basic query embodied in questions 3 and 4 above: ***"Are all of my files still good?"***

While the above discussion details two vital technologies of a truly secure archive system, there are numerous other data protection features that a comprehensive, complete archive solution should include. While finding an archive solution that incorporates all of these capabilities can be difficult, it is not impossible . . .

NEXSAN ASSUREON™: SECURE ARCHIVE FOR COMPLETE DATA PROTECTION

Nexsan Assureon™ is a set of secure storage solutions that reduce storage costs by offloading and deduplicating data that is infrequently used or has aged by policy from primary storage. Assureon, through policy automation, can eliminate or greatly reduce the size, cost and complexity of backups for primary and infrequently- accessed data.

Assureon includes multi-tenancy, with secure copy creation, data movement and long term storage, and charge back capabilities for public and private cloud deployments. Data integrity features like file fingerprinting and automated self-healing integrity checks ensure that your high value data is protected. Assureon's security features comply with corporate and governmental regulatory requirements that make it well suited for medical, financial and government organizations.

Unlike alternative archive solutions, Assureon can guarantee the integrity of archived data through the use of file serialization, file fingerprinting, audit trails, self-auditing, and self-healing capabilities. Because every file is duplicated upon ingestion into Assureon, **data does not need to be backed up,** eliminating the significant hardware, IT management and capacity-metered (per TB) backup application costs associated with weekly full or daily incremental backups.

Despite the powerful technologies that underpin its data protection, Assureon is virtually invisible to end users. By utilizing shortcuts on primary storage to all files that have been migrated to its archives, Assureon ensures that users do not need to change how they access data or learn any new processes.

SIMPLE SECURITY FOR REMOTE BRANCH OFFICES

It can be very challenging to ensure data protection for files that reside in remote or branch offices, which often have few or no onsite IT personnel. A simple, cost-effective solution is to deploy an Assureon secure archive at the organization's headquarters, complemented by an Assureon Edge in each remote office (providing NFS and CIFS shares).

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

Shortfall in Data Protection

In 2013, while about 40 percent of the information in the digital universe required some type of data protection, less than 20 percent of the digital universe actually had these protections.

IDC, The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things, April 2014.

Massive Data Growth Projected

From 2013 to 2020, the digital universe will grow by a factor of 10 – from 4.4 trillion gigabytes to 44 trillion. It more than doubles every two years.

IDC, The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things, April 2014.

All data stored on each branch office Assureon Edge is securely transmitted to the Assureon archive storage system at headquarters, where it will be archived. Alternatively, an Assureon Client can be installed on branch office Windows servers; the client archives selected directories and files by transmitting them to the Assureon system at the organization's headquarters.

PURPOSE-BUILT FOR CLOUD STORAGE

Assureon archive storage systems are designed with multi-tenancy in mind. Online archive storage providers can utilize security features such as certificate-based authentication and separate AES-256 encryption for each cloud services customer. Standard reports keep track of storage usage for each cloud services customer, enabling easy import of this data into the provider's billing systems.

For private clouds, an Assureon can be configured as a Virtual Archive; this makes it possible for multiple secure applications, departments or even separate companies to operate with complete physical, logical and encryption separation.

OPTIMUM STORAGE SOLUTION FOR HIGH-VALUE DATA

From CT and PET scans to MRIs, EKGs to lab reports, ongoing patient electronic records to collaborative care, healthcare providers must manage an enormous, ever-expanding quantity of data. Assureon is specifically designed to safeguard vital, irreplaceable information. Unique in the storage industry, Assureon blends the privacy, integrity and longevity of a secure archive with the high-speed access of online disk drive technology—all while meeting the rigorous regulatory compliance requirements that characterize healthcare environments.

As such, it should come as no surprise that Assureon archive systems are widely used in the healthcare industry; however, Assureon is also frequently deployed in federal, state and local government organizations, as well as a wide variety of business settings (for example, call centers, video surveillance, firms with multiple remote offices, etc.) where robust data protection is required.

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

NEXSAN™
by imation

ASSUREON SECURE ARCHIVE: PROTECTED, EFFICIENT AND CLOUD-READY

PROTECTED

- **File Integrity:** Each time a file is saved, a unique fingerprint is generated using both an MD5 and SHA1 hash of its contents and metadata to ensure history and contents cannot be altered after the fact; every 90 days the integrity of every file is audited against the original fingerprint
- **Data Availability:** Each file is assigned a unique serial number which is used to ensure no files are missing or inappropriately added; every 90 days every file is checked to make sure it is still in the archive
- **File Redundancy:** Each file and its fingerprint are stored twice by Assureon; the second copy is stored in a separate RAID disk set within the same Assureon unit or on a remote Assureon

EFFICIENT:

- **Improved Primary Storage ROI:** Migrating less active files to Assureon frees up capacity and boosts performance on the primary tier while shrinking backup windows
- **No Backups Needed:** Redundant file copies within the secure archive system eliminate the need for costly backup and restore operations
- **Fast Restores:** Restores only need to replace tiny shortcuts rather than actual file contents, enabling IT pros to meet even the toughest Recovery Point Objective (RPO) and Recovery Time Objective (RTO) targets

CLOUD-READY:

- **Multi-Tenancy:** Multi-tenancy features enable cloud service providers to offer highly secure Archive-as-a-Service with logical, physical and encryption data separation
- **Online Archive:** An on-premise Assureon can replicate to a cloud-based Assureon-powered archive service; for private clouds, Assureon supports one-to-one and many-to-one replication

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)

NEXSAN™
by imation

CONCLUSION

Explosive data growth is highlighting the inefficiencies of housing huge quantities of unstructured and structured data on costly primary storage. The vast majority of this data is seldom accessed, and doesn't require the high performance (and high costs) that primary storage entails. Not all data is equal, and IT managers are applying storage optimization principles to better match storage to data's performance, capacity, and connectivity needs.

But as these storage optimization initiatives migrate less-active data from primary to archive storage, a crucial issue is frequently overlooked—the risk of data loss. Nexsan Assureon secure archive solutions are purpose-built to deliver unrivaled data protection for archive data, whether it's stored for days or decades. In short, Assureon makes it possible for IT professionals to cut primary storage costs without cutting corners on data protection.

ABOUT IMATION

Imation is a global data storage and information security company. Imation's Nexsan portfolio features solid-state optimized unified hybrid storage systems, secure automated archive solutions and high-density enterprise storage arrays. Nexsan solutions are ideal for mission-critical IT applications such as virtualization, cloud, databases, and collaboration; and energy efficient, high-density storage for backup and archiving. There are more than 11,000 customers of Nexsan solutions worldwide with more than 33,000 systems deployed since 1999. Nexsan systems are delivered through a worldwide network of cloud service providers, value-added resellers and solutions integrators. For more information, visit www.imation.com/nexsan.

Imation Corp. 1 Imation Way, Oakdale, MN 55128-3414 | p. 651.704.4000 f. 651.537.4675 | www.imation.com/nexsan

(Rev. 09/11/14)