

# A new paradigm in data security and continuous operations through cryptographic data dispersal

Walter R. Lapinsky

White paper



**trusted**

This white paper presents a new method of storage security and virtualization that allows you to consolidate multiple storage networks, each dedicated to a single security level or community of interest, onto a single, virtualized storage infrastructure. It also provides a discussion of the SecureParser® technology and Unisys Stealth architecture, which in combination allow secure sharing, cost savings, and continuous operation in consolidated storage networks.

## **Table of Contents**

<b>Introduction</b>	5
<b>Securing the SAN</b>	6
<b>Consolidating parallel infrastructures</b>	7
Multilevel security	7
Communities of interest	8
<b>Achieving higher availability</b>	9
<b>Continuous operation equals replication</b>	9
<b>Measuring recovery: RTO and RPO</b>	10
<b>Unisys Stealth Solution: Cryptographic dispersal</b>	11
<b>Crypto-splitting: SecureParser overview</b>	11
<b>The Unisys Stealth Solution architecture</b>	13
<b>Cryptographic data dispersal</b>	14
<b>The Stealth SAN appliance</b>	15
<b>Conclusion</b>	16
<b>About the author</b>	16
<b>References</b>	17

This page is intentionally left blank.

## Introduction

With the multiple pressures of time and workflow, managing stored data can be difficult. You must ensure the confidentiality, integrity, and availability of your data.

- Confidentiality means ensuring that only those who are permitted can gain access to the data.
- Integrity means ensuring that the data cannot be modified except by authorized persons and processes.
- Availability means ensuring that the data is available when it needs to be available.

In short, the goal is to get the right data to the right people at the right time. In today's world, "the right time" is "right now" and "always." Increasingly, applications exist in government, finance, and other markets where not having data available for even seconds is unacceptable.

This white paper briefly reviews the established mechanisms to protect data stored in storage area networks (SANs) through the use of segregation of data into individual SAN infrastructures to physically separate data associated with different security levels or communities of interest as well as the use of cryptographic data dispersal to ensure data availability.

It also examines a solution to these shortcomings, called the Unisys Stealth Solution for SAN. The Unisys Stealth Solution including SecureParser® from Security First Corp. closes the gap in SAN security so that separate enclaves for different security levels or communities of interest need not be maintained, and you can securely intermixed all data within the same SAN infrastructure.

At the same time, you can use cryptographic data dispersal to provide secure continuous operation without replication and at significant cost savings. For point-of-the-spear applications, Stealth can create an environment in which you trust the security of the data, knowing that no one can recover the data even if the keys to the data are also captured.

## Securing the SAN

The majority of security threats come from insiders. Whether those threats are malicious or unintentional, storing all data within an enclave in clear text, although that enclave only supports one security level, represents a risk that is not always being addressed. Encrypting data for the long term has significant key management costs and risks. In lieu of a workable SAN security technology, enterprises often take extreme measures to physically secure the SAN infrastructure.

Physically securing the SAN infrastructure is what leads to the parallel storage area networks shown in Figure 1. Everyone at the same security level is using the same key (if the data is encrypted), Top Secret data cannot be present (and hence visible) on any storage area network to which Secret or Unclassified clients are connected. The same is true for Secret data on Unclassified networks. Thus, a strict physical segregation of storage area networks by classification level is implemented.




Figure 1. Multiple parallel storage networks

## Consolidating parallel infrastructures

Administratively, implementing multiple parallel storage area networks is fraught with problems. There is the obvious cost of obtaining, managing, and maintaining the necessary equipment, plus extra cost for power, cooling, space, and weight.

### Multilevel security

What is needed instead is a method of intermixing data classified at different security levels on the same storage area network in such a way that the data is protected

from being accessed by any client that is not authorized to do so. This storage scheme is the “holy grail” of multilevel security – a single SAN infrastructure, including the switches and all of the other pieces of equipment needed to implement a data-centric enterprise.

In essence, the storage network becomes virtualized on demand to support the storage of different classifications of data. Figure 2 shows how the networks shown in Figure 1 could be consolidated into a shared storage area network supporting multiple levels of security.




Figure 2. Consolidated storage network

## Communities of interest

An abstracted notion of a multilevel security storage area network beyond that of the current paradigm that supports separate Sensitive, Secret, and Top Secret designations would be a new paradigm. This paradigm compartmentalizes data by membership in flexible communities of interest (COIs), rather than by rigid security level classifications.

A COI is a group of people and/or servers that must share information and not allow anyone outside of their COI access to that information. The same individual could be part of more than one COI, either one at a time or multiple at the same time depending on the organization's security requirements.

Figure 3 compares the current rigid, hierarchical classification structure to a dynamic COI model.

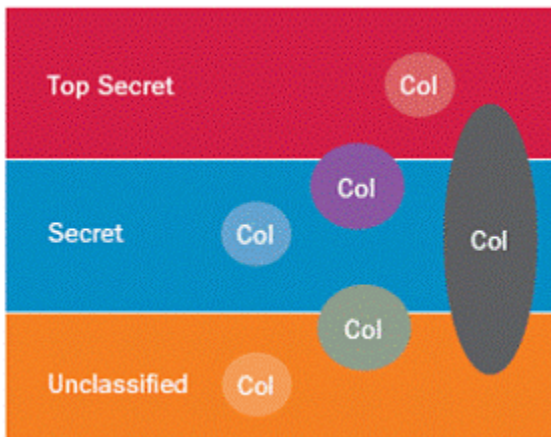


Figure 3. Communities of interest security model

The COI security model automates a true “need-to-know” capability for compartmented data. When necessary, you can authorize a particular user or server to participate in a COI for as long as needed. As a result, the COI-capable storage network supports not just multilevel security but also controlled data sharing/restricted access

- To support Multi-National Information Sharing (MNIS) in coalition operations
- To support unanticipated users during cooperative joint operations with law enforcement or first responders
- To control access to data within a financial institution

## Achieving higher availability

RAID – Redundant Array of Inexpensive Disks – is a technology that employs the simultaneous use of two or more disk drives to achieve greater levels of reliability. When several physical disks are set up to use RAID technology, they are said to be in a RAID array. The array is seen by the server and the operating system as one single disk drive.

Several different levels of RAID exist, but in each case, arrays of disks are redundant by writing extra data across the array such that the failure of any one disk in the array does not result in loss of data.

RAID levels 1 and 5 are the most commonly used to provide redundancy.

- RAID 1, also known as mirrored disk, could be described as a local backup solution. RAID 1 writes the same data to two or more drives so that data is not lost as long as one drive survives. The total capacity of the array is the capacity of a single disk drive, but it requires two or more disk drives. The failure of one drive does not increase the chance of a failure nor decrease the reliability of the remaining drives. In other words, it requires at least 20TB of physical disk to hold 10TB of data with RAID 1 (using 10TB drives).
- RAID 5, also known as striped disks with parity, combines three or more disks in a way that protects data against the loss of any one disk. The storage capacity of the array is reduced by one disk. It requires at least 15TB of physical disk to hold 10TB of data with RAID 5 (using 5TB drives).

## Continuous operation equals replication

To gain high availability of their data, most organizations replicate their data. They start by using technologies like RAID to replicate the data locally to protect against local failures, such as loss of a single storage unit. To protect against larger failures, such as power outages or natural disasters at one facility, organizations often have a remote site where they are periodically updating an additional copy of their critical data. That data is itself often replicated using RAID or similar technologies. Depending on the particular mechanism used to transport the data to the remote site, there might be an additional staging copy of the data might exist at the primary site that holds a snapshot of the data before it is sent to the remote location.

As terrorist attacks and natural disasters have demonstrated, you can lose a site not just for a few hours but for weeks or months. The impact area of an event might be larger than the separation of the prime and backup sites. Many organizations are creating an additional backup site hundreds or even thousands of miles away to protect their business or their country from a data outage.

The bottom line is that all of the redundancy is expensive. Analysts indicate that organizations maintain about five copies of their data because of replication. If you are managing hundreds of terabytes of data, you probably own several petabytes of storage. Figure 4 shows how 10TB of data easily expands to more than 40TB.

First the local storage uses RAID technology. Periodically, the site suspends operations for a short time to have a consistent database and creates a business continuance volume (BCV) – a record of the changes since the last suspension. They use the BCV to update the remote copy of the data, which is also protected by RAID technology. The result is a requirement to have five times the original storage available in order to protect the data.

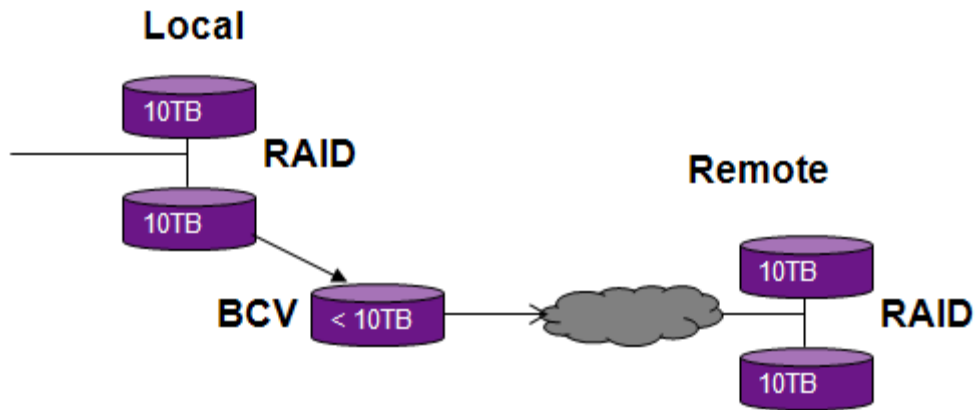


Figure 4. Typical replication of 10TB of data

## Measuring recovery: RTO and RPO

Having data replicated so it survives any event is the first step to having a highly available organization. But perhaps more important are the answers to these questions:

1. How long does it take to get at the data?
2. How far back in time did I just jump when I got the data?

Two measurements are used to answer these questions:

- **Recovery-Time Objective (RTO):** RTO is the time period after an event at which business functions need to be restored. An RTO of one hour says that you must be operational with your critical applications, and the data they need, within one hour of an event.
- **Recovery-Point Objective (RPO):** RPO is the age of the data after it is recovered from backup storage for normal operations to resume after an event. The RPO is expressed backward in time (that is, into the past) from the instant at which the failure occurs. You can specify it in seconds, minutes, hours, or days. For example, an RPO of an hour indicates that when you come up after an event, the data is no more than one hour old.

Clearly, the shorter the RTO and RPO values, the less impact an event has on your organization's ability to function. Shorter RTO and RPO values usually have significant cost implications. Different applications can have different RTO and RPO values. For example, an

airline's operations system, which must be operational to push an airplane away from a gate, needs RPO and RTO values in minutes. The airline's payroll system can have RTO and RPO values measured in days.

RTO values determine whether you need to have a populated infrastructure in your backup site – systems, networks connections, and so forth. Shorter RTO values indicate whether or not that infrastructure must actually be running and what kind of failover technology or other parallel processing techniques you are using.

RPO values determine the process and infrastructure necessary to transport the data between the sites. They determine everything from putting a tape on a truck every night to multiple high-speed network connections, and a synchronous or nearly-synchronous backup solution.

## Unisys Stealth Solution: Cryptographic dispersal

The Unisys Stealth Solution uses encryption plus cryptographic bit-level splitting and data dispersal to spread storage area network (SAN) data across an organization in a way that consolidates SANs, maintains high availability, reduces the total amount of storage required, and saves money. The consolidated SAN becomes a single infrastructure that supports multiple security-based COIs.

The Unisys Stealth Solution for SAN is a data security architecture that allows the intermixing of data for different COIs (or security levels) on the same SAN infrastructure in such a way that the separation of data for different COIs is maintained. The data is logically separated by COI as defined by a workgroup key for each COI, although it physically resides on the same SAN. You can physically disperse data within a single COI through bit-level data dispersal at different locations. You can perform this dispersal such that a different location has sufficient information to recover the original data or such that not enough data exists at a specific location to recover the data no matter how much other information is available, including the original cryptographic keys.

Stealth segregates the data by encrypting it and storing cryptographically-split pieces in different locations as specified by the site. Data is not physically segregated by COI or security level, but rather logically segregated by a COI-specific secret (workgroup) key, then by the random distribution of bits that make up the data.

The Unisys Stealth Solution accomplishes this cryptographic splitting of data by utilizing the SecureParser® from Security First Corp.

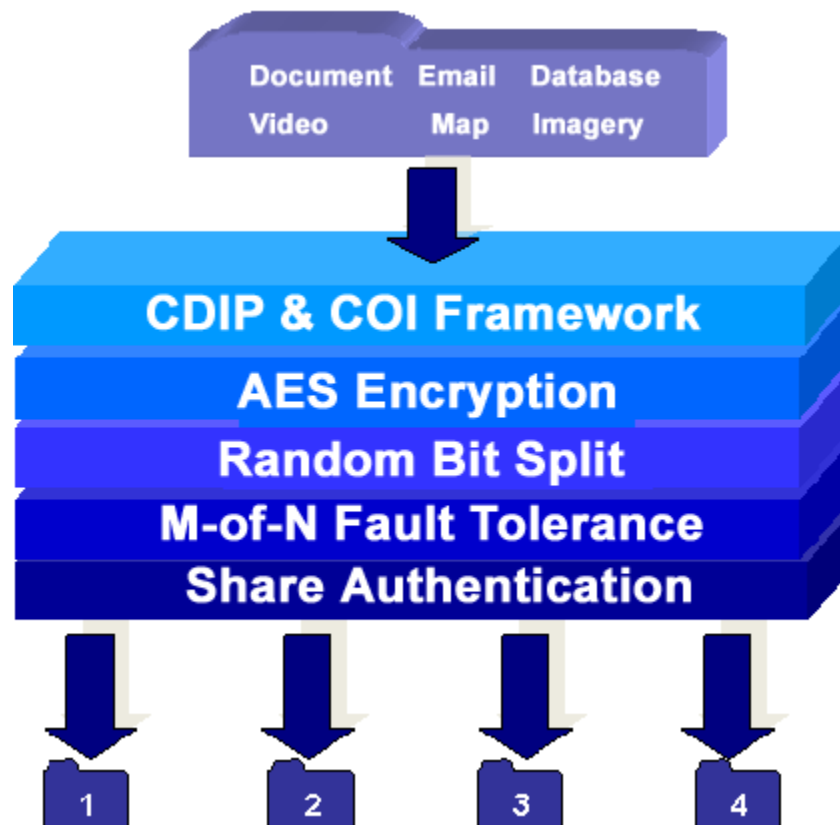
## Crypto-splitting: SecureParser overview

SecureParser is not an encryption method, but it works in conjunction with standard encryption techniques like DES and AES to add a layer of physical security. SecureParser takes an input buffer, shreds or “parses” the data at the bit level, then randomly assigns each bit to one or more output shares. A cryptographically secure pseudo-random number controls the distribution of the bits. The resultant shares have the characteristic that a minimum subset of them is required to restore the original data.

SecureParser operates on in-memory data segments of variable size. The SecureParser cryptographic split is made up of five modular components:

- The CDIP (Cross Domain Information Protection) and COI Framework Module allows an external workgroup key to be used to encrypt the internal keys generated by the SecureParser to encrypt, split, and authenticate the data. Use of these keys not only strengthens the protection of the shares but restricts access to the data to those in possession of, or having access to, the workgroup key.
- The AES Encryption Module uses provably strong AES 128, 192, or 256 to encrypt the data prior to splitting it into shares.
- The Random Bit Split Module uses a keyed IDA (Information Dispersal Algorithm) that randomizes and shreds the encrypted data at the bit level, and each of those pieces is randomly distributed to one or more of the output shares. This highly efficient algorithm provides two-factor secret sharing that requires not only a sufficient number of the underlying shares but also a split key to restore the data. This patented process is unique to SecureParser.
- When fault tolerance is specified (see the M-of-N discussion in a subsequent topic), the M-of-N Fault Tolerance Module writes each bit of shredded data to more than one output share. This redundancy allows the restoration of the original data with a minimum subset as opposed to all of the shares.
- The Share Authentication Module adds integrity information to each share to allow the detection of corrupted shares.

Figure 5 shows a schematic of the SecureParser processing.



**Figure 5. SecureParser processing steps**

In the simplest mode of operation, the Encryption Session Key and Split Session Key used are generated internally.

These keys are included with the data. These session keys are themselves split, and the key shares are stored together with the data shares. As a result, no external key management is needed. Rather, access to enough of the separate shares is, in essence, the “key.” This option is a big advantage over using just straight encryption to protect the data, especially when the data must persist for a long period of time, such as backup/archive data. If an external key is used, that key must persist and be available as long as the data it is protecting exists. Thus, for situations where physical separation of the data shares is sufficient protection, key management is not a concern.

Each piece of split data is parsed into one or more shares. The reason the data might be put into more than one share is to allow for resiliency in the case where one or more of the data shares are lost or corrupted.

You can configure SecureParser to support M-of-N redundancy: N shares are generated, but only M of them are required to restore the original data. Thus, in a 2-of-4 scenario, the original data is parsed into four shares such that any two of them can reconstruct the original.

In many situations, this type of redundancy is a big advantage. For disaster recovery purposes, mission-critical data must be duplicated, often to a remote site. Without splitting, all of the data would need to be recovered before processing could continue. Using split redundancy, processing can continue on the remaining shares (which can still restore the original data) while a new set of redundant shares is created.

Although the individual shares are smaller than the original data, there are no savings in the total amount of data. If the configuration allows one share to be lost, the data storage or bandwidth needed multiplies. For example, if M-of-N is 2-of-4, each bit must be in three shares as two can be lost. This scheme triples the data.

As noted, the Internal Encryption Session and Split Session keys are each split and stored in the output shares. This separation allows the data to be restored without using any external keys once a minimum subset of shares is located. For situations in which additional security beyond the physical separation of shares is required, you can encrypt the Split Session Key with an External Workgroup Key. The Workgroup Key is a symmetric key that is also required during restoration.

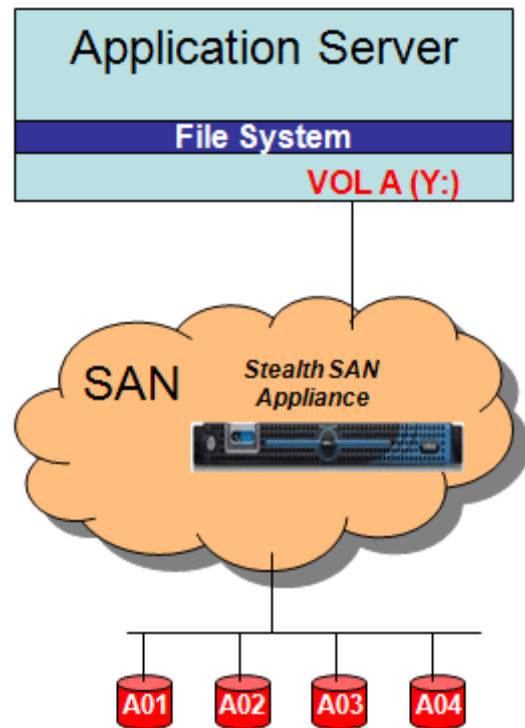
## The Unisys Stealth Solution architecture

All of the Stealth technology is implemented in a Stealth SAN appliance. The Stealth SAN appliance sits in the SAN infrastructure. A server sees it as a disk array, and a disk sees it as a server. In a normal SAN environment, a disk drive (for example, y:) is mapped to a network drive (for example, VOL A). In some SAN environments, that network drive, often called a Logical Unit Number (LUN), can be virtualized in multiple ways:

- Spread across multiple physical disk units
- Combined with other network drives in a single disk device
- Included in RAID technology in a variety of ways

The Stealth SAN appliance performs similar virtualization except with additional functionality that enables enhanced security, controls storage virtualization, and offers disruptive continuous operations capabilities.

Figure 6 shows the basic architecture of the Stealth SAN appliance.

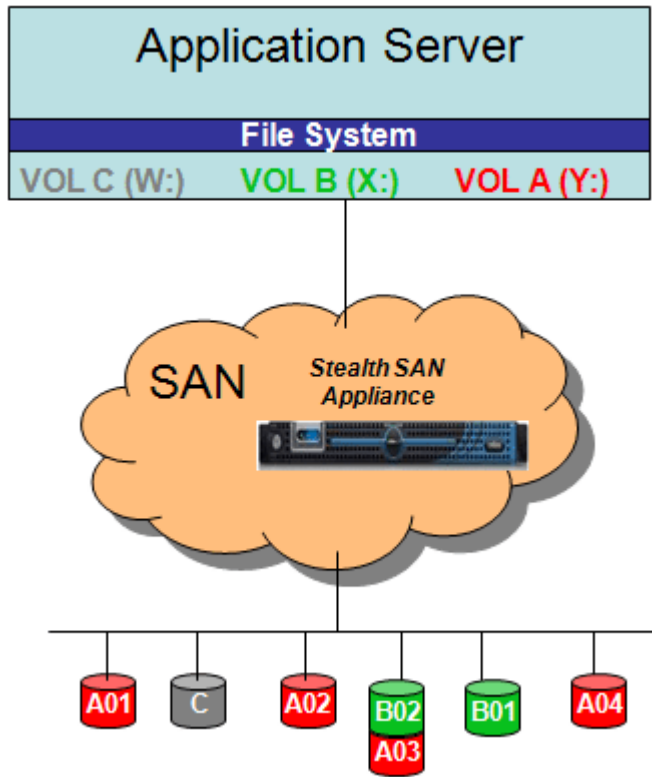


**Figure 6. Stealth SAN appliance**

The site administrator sets up each LUN in a similar way to a traditional SAN virtualization environment, except that each LUN has N shares based on the M-of-N shares configured for that LUN. In figure 6, LUN A is virtualized into 4 shares (M-of-4 configuration). Each bit of data written to LUN A will be encrypted and shredded into four blocks or shares – one written to each of the logical units A01, A02, A03, and A04. These multiple writes are performed in parallel, and the data written to each share is smaller than the original data. Each LUN belongs to a single COI, and each COI has a unique set of encryption and splitting keys. Multiple LUNs can belong to the same COI as appropriate.

The Stealth SAN appliance slips into the SAN fabric and is agnostic to the vendor and type of servers or the vendor and type of storage. The same Stealth SAN appliance can support multiple LUNs belonging to multiple COIs. Also, additional legacy LUNs can exist that are not managed by the Stealth SAN appliance.

Figure 7 shows LUN A virtualized into 4 shares (M1-of-4 configuration), LUN B virtualized into 2 shares (M2-of-2 configuration), and LUN C not managed by Stealth at all.



**Figure 7. Stealth SAN appliance with multiple LUNs**

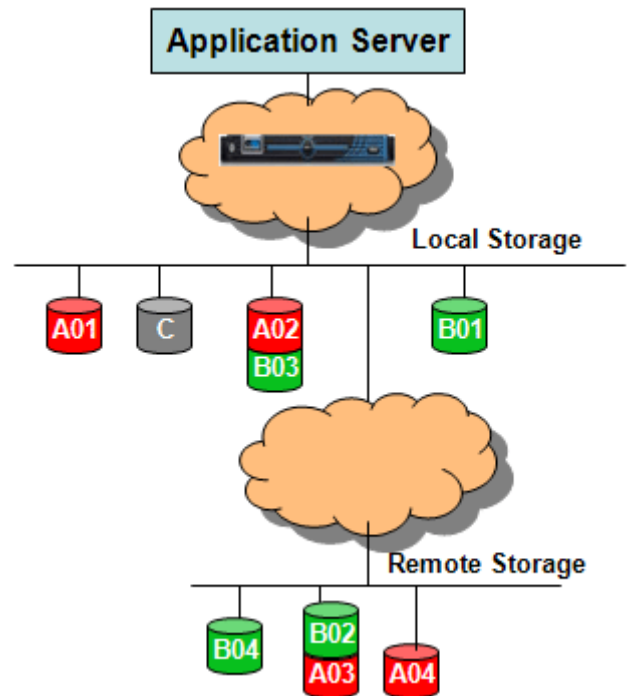
In this example, LUN B belongs to a different COI than LUN A and therefore has a different set of encryption and splitting keys. You can place shares belonging to different COIs on the same physical drive without fear that someone (or an application) that does not belong to the same COI could access the data.

An important characteristic of this design is that it can fit into an existing SAN infrastructure without change. Based on the individual needs of applications and associated LUNs, you can integrate the Stealth Solution into the SAN infrastructure at the pace that makes sense for the organization

## Cryptographic data dispersal

In a previous topic, you learned about the ability of Stealth to parse data into multiple shares such that Stealth does not need all of the shares to recover the data. In Figure 8, Stealth is parsing LUN A into four shares. If all four shares are needed to recover the data, then each share is approximately a quarter of the original data size. The M-of-N in this case is 4-of-4.

A more interesting scenario is an M-of-N of 2-of-4. In this case, only two of the four shares are necessary to recover the original data. Each bit of the original data is placed in three of the shares, so Stealth writes approximately three times as much data as in the original block. Each share is approximately 75 percent of the size of the original. If two of the shares are physically local to the application server and two are remote, the result appears similar to the diagram in Figure 8. In this case, both LUN A and LUN B are written 2 of 4.



**Figure 8. Local and remote share**

For LUN A in the normal case, an application running in the application server reads from the two shares in the local storage, A1 and A2. If a failure occurs on one of those drives, for example on A1, then Stealth automatically reads the data from A2 and one of A3 or A4. The application is not aware that process has occurred. If a larger failure took out both A1 and A2, the application would still be protected with no application failover required.

Depending on the connection to the remote site, some degradation in response time often occurs after a drive failure. However, Stealth allows the application to continue without the interruption necessary to perform a failover.

Once a failed drive is repaired, Stealth can recover the share that was on that drive by rebuilding it from the remaining shares, without decrypting the data.

If the connection to the remote site is temporarily lost, Stealth caches the writes to remote shares locally until the connection is restored, and then it copies those writes to the remote site.

If a backup application server and Stealth SAN appliance are at the remote site, the loss of the entire facility that holds the application server and its local storage would require the failover of the applications to the remote site. The data is already at the remote site.

In this 2-of-4 configuration, the total data written is approximately three times the size of the original data. Thus, if LUN A represents 10 TB of data, the combined size of the four shares would be approximately 30 TB – at least a 25 percent savings over the 40 to 50 TB shown in Figure 4 in a standard data replication model.

## The Stealth SAN appliance

The Unisys Stealth SAN appliance provides highly available and provably secure disk virtualization. Your administrator can create a virtual disk, allocate storage to it, and assign the virtual disk to a specific COI. Multiple virtual disks using different COIs can share the same storage areas, and virtual disks representing different COIs can securely share the same physical disk drive.

The Stealth SAN appliance provides configuration management of the storage. An administrator can monitor status, check for events, and perform security functions from a single console. The administrator can create, modify, and destroy virtual disks, and make the virtual disks visible or invisible to the appropriate COIs without requiring the rezoning of the individual switches in the SAN network. The appliance software also provides a management interface that follows SNIA (Storage Networking Industry Association) standards.

The Stealth SAN appliance provides caching of input/output (I/O) operations. The administrator can select which LUNs are cached and the amount of memory to be used for caching.

Stealth provides for virtual storage capacity on demand. A virtual storage enabler key identifies the amount of virtual storage available for use. The dispersal strategy determines the amount of physical storage space that is required to support the virtual storage space. If additional virtual storage space is required, a new virtual storage enabler key is electronically shipped to the administrator and is entered into the Stealth appliance administrator console.

Stealth encrypts data at the AES-256 level then shreds the encrypted data into shares at the bit-level. The separate shares are stored on separate storage media in dispersed locations, possibly both local and geographically remote locations. Because the data is encrypted and shredded, you can transmit the data without the use of private networks.

This integrated resiliency capability means that remote mirror data replication is not required. A complete copy of the data does not exist on any storage media volume. This patented process is provably secure, while at the same time it provides the high availability characteristics that usually require separate investment in RAID and

remote mirror data-replication technology. This always-on data capability improves RTOs and RPOs. The shredding process requires only a subset of the shares to reconstitute original data resulting in a reduction in the storage capacity required because fewer copies of the original data exist.

## Conclusion

The new paradigm in data security described in this white paper uses a unique bit-level splitting technology to cryptographically disperse data through a SAN. The Unisys Stealth SAN Solution can also provide interesting continuous operations opportunities to reduce RTO and RPO. At the same time, Stealth uses less redundant storage than traditional replication methodologies for data protection while simultaneously offering easier storage management by allowing virtualization of storage and, in some environments, the consolidation of multiple SANs without sacrificing data access controls.

## About the author

Walter R. Lapinsky  
Product Marketing Director  
Systems and Technology, Unisys Corp.

Mr. Lapinsky earned a BS and MS in Mathematics from the University of Delaware and completed graduate work in Computer Science at the University of California, Berkeley. In his 25+ year career with Unisys he has held many positions – primarily in software engineering development and management. Many of those projects were within the U.S. federal market including missile guidance, communications systems and logistics systems.

Prior to rejoining Unisys, he was the Vice President of a small software development and consulting firm for 12 years, and the Manager of Software Engineering for a communications software company for 4 years.

Currently, he is responsible for product marketing for the Unisys Stealth Solution.

## References

Several of the references below are internal whitepapers that have been distributed to prospects at various conferences. You can obtain them from the author upon request. Contact the author at [Walter.Lapinsky@unisys.com](mailto:Walter.Lapinsky@unisys.com).

[JOHN07-1] R Johnson, "The Unisys Stealth Solution and SecureParser: A New Method for Securing and Segregating Network Data," Unisys Corporation white paper (2007).

[SCHN05] S. Schnitzer, R. Johnson, and H. Hoyt, "Secured Storage Using SecureParser®," Proceedings of the 2005 ACM Workshop on Storage Security and Survivability (2005).

[SFC05-1] Security First Corp., "SecureParser® Beyond Encryption v3.5," white paper (2005).

[SFC05-2] Security First Corp., "SecureParser® Design Specification v4.1," white paper (2005).

[SFC05-3] Security First Corp., "SecureParser® Storage Overview v4.1," white paper (2005).

[SFC06-1] Security First Corp., "SecureParser® Cryptographic Core Design v4.1," white paper (2006).

[SFC06-2] Security First Corp., "SecureParser® Workgroup Key Usage Notes v 4.1," white paper (2006).

[SFC06-3] M. Bellare, P. Rogaway, "Robust Computational Secret Sharing and a Unified Account of Classical Secret-Sharing Goals," unpublished manuscript (2006)

## Notes

## Notes

---

For more information, please visit our web site at [UnisysStealthSolution.com](http://UnisysStealthSolution.com).

©2009 Unisys Corporation.

All rights reserved.

Unisys and the Unisys logo are registered trademarks of Unisys Corporation. All other brands and products referenced herein are acknowledged to be trademarks or registered trademarks of their respective holders.

Printed in the United States of America 03/09.



3839 3732-000