Blanket Encryption

Datrium™

385 Moffett Park Dr. Sunnyvale, CA 94089 844-478-8349 www.Datrium.com

Technical Report

Introducing Blanket Encryption

Datrium Blanket Encryption software offers a groundbreaking new alternative for private cloud operators.

Blanket Encryption provides comprehensive encryption over the lifecycle of VM data, from its genesis on a host, through in-use storage on host flash, across networks, and persisted at rest. It manages this while simultaneously supporting two additional challenges:

1] Supporting global deduplication and compression across the DVX system.

2] Being software-only, so not requiring self-encrypting drives or other hardware, thus further enabling DVX users to take advantage of emerging host flash technologies without sacrificing important data management features.

This combination of capabilities has not previously been available for virtualized, converged infrastructure. In a time of escalating data security issues and privacy requirements, Datrium believes private clouds can offer a haven of high trust infrastructure if the right tools are available.

This paper will review the requirements for encryption, some of the difficulties of prior models, and introduces the architecture used by Datrium in its Blanket Encryption software.

The need to encrypt data

Organizations increasingly have sensitive and private data that must be protected from unintended disclosure. Such data runs the gamut from corporate secrets to sensitive customer medical records. For many CIOs, government and industry regulations such as HIPAA and PCI enforce proper protection of sensitive data. There are significant costs associated with data breaches, including fines from regulatory bodies and reputation loss stemming from the negative publicity caused by data breach incidents. The risk of unintended disclosure is real for a variety of reasons.

First, cyber attacks happen. Second, hardware failures and upgrades happen. When hardware is replaced, any sensitive data that is on the decommissioned hardware is at risk for inadvertent disclosure. Third, human errors happen. We may have carefully crafted workflows for setting up access control and decommissioning storage devices, but mistakes and omissions can occur by human error.

Data encryption can mitigate the risk of inadvertent disclosure and the damage that results from data breaches. It offers peace of mind and a safe harbor in many cases.

Challenges to deploying encryption

Data encryption can help with securing sensitive data. But there are several challenges to deploying it in an effective way.

Protection Scope

Data processing systems have become ever more sophisticated. Modern storage hierarchy includes multiple storage tiers and caches made up of a variety of storage devices. The data on each of these tiers in the storage hierarchy and on each of the storage devices must be protected. Furthermore, storage systems are increasingly distributed, increasing the risk of leaks through the network.

The net result is that in the storage and network infrastructure alone, there are many elements that must be secured to properly protect the data. Furthermore, if each of these elements requires a different point-solution to secure, deploying and managing all these solutions appropriately will be an administrative nightmare.

Operational Flexibility

Another challenge with deploying encryption is that some solutions may complicate and impede routine IT operations. For example, as the risk of data spillage becomes widely appreciated, a common business need is to encrypt data for an existing application or to deploy a new application with encrypted data.

If an encryption solution relies on specific storage devices such as self-encrypting drives (SEDs) to perform encryption, the business need cannot be addressed until the encrypting storage devices are procured, acquired, deployed, and data is migrated onto them. Moreover, having encrypting versus non-encrypting storage devices adds to the number of distinct silos of storage that must be managed.

The capability to enable encryption on existing infrastructure is highly desirable because it provides operational flexibility. However, this is difficult to achieve in practice because non-overwriting file and storage systems are widely deployed at multiple levels in the storage hierarchy. If the encryption solution is not aware that data may not be overwritten when updated, it cannot ensure that data is not accessible in the unencrypted form. Such a solution will also have difficulty satisfying basic operational requirements such as the ability to rotate encryption keys.

Data Efficiency

As data continues to grow unabated, data reduction techniques such as data duplication and compression are critical to enable the data to be effectively managed. For most workloads, these techniques significantly reduce the data size for both storage and replication to a small fraction. Particularly as data centers adopt solid state technologies such as flash and NVMe, it is hard to imagine a modern storage system that does not employ these tested and proven techniques. However, if an encryption solution is not designed with these techniques in mind, these techniques may be rendered ineffective.

Encryption securely scrambles the data so that the encrypted data (cipher-text) is for all practical purposes incompressible. Broadly speaking, the cipher-text is a function of the original data, the data encryption key and the initial vector or 'tweak.' With a properly chosen tweak value, even multiple copies of the same block of data, encrypted with the same encryption key will create vastly different cipher-texts. Hence, the various copies will not be duplicates. If an encryption solution encrypts the data prior to de-duplication, it will destroy any savings from the downstream de-duplication.

Performance Impact

Encryption tends to be a compute-intensive process. Depending on how and where encryption is implemented, deploying encryption may have a significant impact on overall workload performance. For example, encryption has previously been associated with poor performance when implemented in software. Performance is especially a concern for encryption solutions that rely on limited controller resources within an array to perform encryption in software.

Problems with current approaches

The current approaches to encrypting data can be broadly classified into application-level, device-level and storage-array-level.

Application or VM-level Data Encryption

In this approach, data is encrypted before it is committed to durable storage through standard storage interfaces. For example, the encryption may be performed by an application such as a database management system or, in the case of virtualized environments, the hypervisor. This approach ensures that data is protected starting from the hosts running the application. However, it renders ineffective any downstream data efficiency techniques. This means that the storage system has to unnecessarily store and replicate several times more data than without encryption. Encrypting data upstream of the storage system also suffers from the issue that the encryption system is unaware of and has no control over any non-overwriting capability in the storage system. As discussed earlier, basic operations such as key rotation may not be properly performed in such cases.

Device-level Data Encryption

In device-level data encryption (i.e. SEDs), each of the storage devices performs encryption on the data it receives. Such an approach offers high performance through the use of embedded hardware as well as scaling of this hardware resource with the number of storage devices. However, this approach does not protect data until it arrives at each individual device. Also, for this approach to be effective, each and every

device type and model must be capable of performing encryption. Deploying encryption using such an approach is difficult because it does not extend to existing storage devices that lack the built-in capability to encrypt data. This typically means that new storage devices must be procured, acquired and deployed into discrete pools of storage that have to be separately managed.

Storage-array-level Data Encryption

The third approach to data encryption is to rely on the storage array to perform controller-based software encryption. In this approach, the data is not protected until it arrives at the storage array. This means that the storage devices on the hosts are not covered by this approach nor is the network between the hosts and the storage array. Performance is also a major challenge with this approach, because typically just two array controllers are tasked with all I/O operations as well as data encryption, which can create a performance chokepoint (even without being tasked to perform encryption).

DVX Components *figure 1*

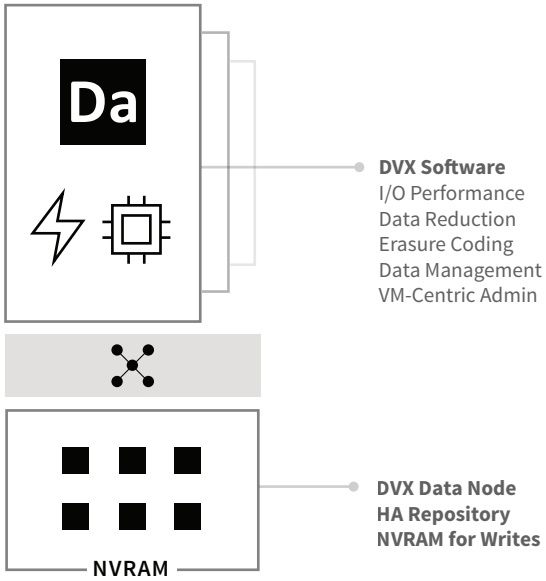

Datrium Blanket Encryption Software

The Datrium DVX system consists of:

1] **DVX Software on each host manages all active data for the Virtual Machines (VMs) within that compute node (host). It provides scalable IO performance, availability and data management capabilities by leveraging compute node CPU and memory for I/O processing, and flash-based storage devices on the compute node for data caching.**

2] **DVX Data Node: The DVX Data Node provides persistence and resiliency for a durable copy of all data in the cluster. In normal operation it is write-only, but it also provides streaming read performance for flash uploads as well as cluster coordination for simple management.**

With Datrium Blanket Encryption, the DVX Software running on each compute node encrypts incoming data within compute node RAM. The encryption is performed before the data is sent on any of the multiple paths through the system:

1] **Data is sent encrypted over the network and stored encrypted in the non-volatile RAM (NVRAM) on the Data Node.**

2] **Data is sent encrypted over the network and stored encrypted in the persistent high capacity storage devices on the Data Node.**

3] **Data is stored encrypted in the Flash-based Cache on the host.**

4] **Data is sent encrypted over the network (LAN or WAN) to another Datrium DVX.**

Data Encryption *figure 2*

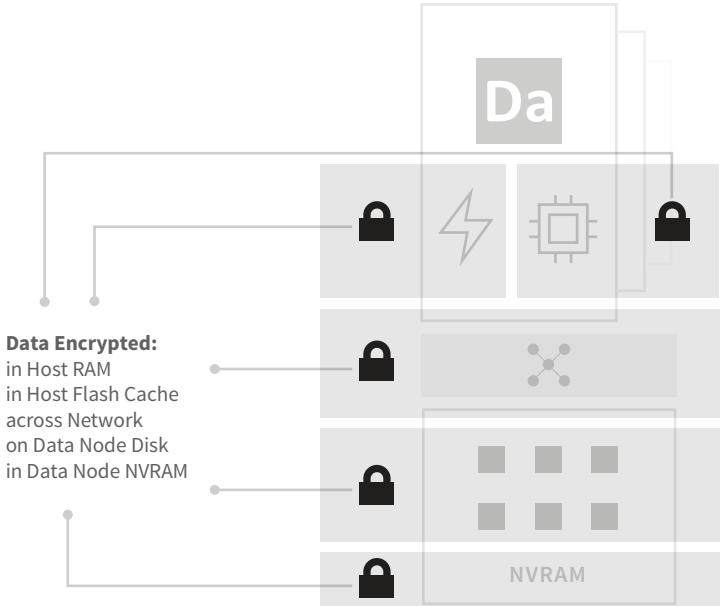

In other words, the DVX Software is the point of entry into the protection domain. DVX Software also decrypts outgoing data so that for the entire duration that the data is within the DVX, it is protected with encryption.

By performing encryption and decryption in the DVX Software, Datrium Blanket Encryption extends protection from the compute nodes, across the wire (internal and WAN), and down to the individual storage devices on both the compute nodes and DVX Data Node. We call this end-to-end storage protection.

Before encrypting the data, the DVX Software computes fingerprints (cryptographic-strength identifiers) on the data. The fingerprints are later used to perform deduplication on the data. By computing the fingerprints on the pre-encrypted (or clear-text) data, the DVX retains the ability to find and eliminate duplicate variably-sized groupings of blocks of data regardless of where the blocks of data are located.

DVX Software also performs compression on the clear-text data before the data is randomized by encryption and rendered incompressible. Thus, with Datrium Blanket Encryption, customers continue to enjoy the considerable benefits of de-duplication and compression. And the savings do not just apply to the durable copy of the data in the Data Node. They also apply to the active copy in the flash-based cache on each compute node, and to the network transmission between the compute nodes and Data Node, and across LAN and WAN to a remote DVX.

Data Master & Agents *figure 3*

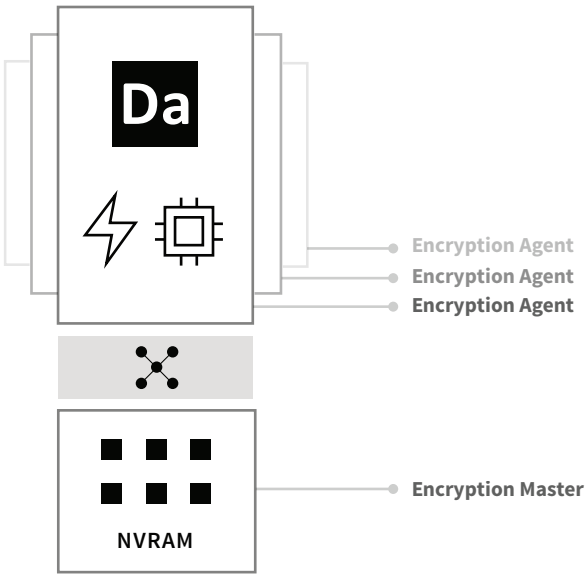

With Datrium Blanket Encryption, the scope of protection extends from the DVX software on each of the compute nodes, through the network to the DVX Data Node, and down to the individual storage devices on both the compute nodes and DVX Data Node. The encryption is consistent across the entire protection scope so that there is only a single encryption system to manage. Each of the DVX Software instances within each compute node includes an Encryptor Agent that operates on behalf of an Encryptor Master to provide the simplicity of a single encryption system while enabling robust distributed protection beginning on the compute nodes running the applications.

Datrium Blanket Encryption is software-based so it works across all types and models of supported storage devices. This is an especially important consideration because flash-based storage devices are undergoing rapid technological change. Customers will invariably wish to deploy a variety of such devices over time, and they can be assured that Datrium Blanket Encryption will protect their data on any storage device that they deploy within the system. A software-based encryption system also provides operational flexibility to enable encryption on existing hardware when sensitive workloads are added or when new encryption requirements arise.

A common concern with using software-based encryption is that it burns CPU cycles and results in poor performance. Datrium Blanket Encryption, however, imposes virtually no performance impact. This is achieved through a combination of techniques. First, Datrium Blanket Encryption uses extensively the AES-NI (Advanced Encryption Standard – New Instructions) engine that is available and otherwise mostly idle on modern processors. So it offers the flexibility of a software-based solution and the performance of a hardware accelerated implementation.

Second, Datrium Blanket Encryption uses a cryptographic algorithm (FIPS-approved AES-XTS-256 with 512-bit keys) that has been specially optimized to leverage the AES-NI engine. Third, Datrium Blanket Encryption performs encryption and decryption in the DVX Software on each of the compute nodes, so its performance scales as workloads expand and additional compute nodes are provisioned. The net result is that there is virtually no performance impact to enabling encryption.

To simplify deployment, Datrium Blanket Encryption includes a built-in key manager integrated with the DVX high-availability capability. The key manager supports two modes of operation.

1] In unlocked mode, the system utilizes an embedded system store to remember a key derived from the encryption password, and uses the derived key to unlock the system when started. In this mode, the system resumes service after a restart without requiring administrator action to provide the encryption password.

2] In shipping mode, the derived key is not stored persistently anywhere in the system. In this mode, the DVX and/or ESX compute nodes can be transported without risk of a data breach while in transit. When the system is powered up in this mode, the administrator must provide the encryption password before the system will serve data.

Blanket Encryption Conclusion

Datrium Blanket Encryption offers:

1] **End-to-end storage protection** that extends from the compute nodes running the applications, through the network, and down to the individual storage devices on both the compute nodes and DVX Data Node.

2] **Simplicity** of a single encryption system and coverage of multiple point solutions.

3] **Flexibility** of a software-based solution with virtually no impact on performance.

4] **De-duplication and compression savings** in storage capacity for Flash-based caches and Netshelf, and network bandwidth including WAN bandwidth between DVXs.

	Application-Level	Device-Level	Storage-array-Level	Datrium Blanket Encryption
Protection Scope	✓			✓
Operational Flexibility			✓	✓
Data Efficiency		✓	✓	✓
Performance Impact		✓		✓