Datadobi

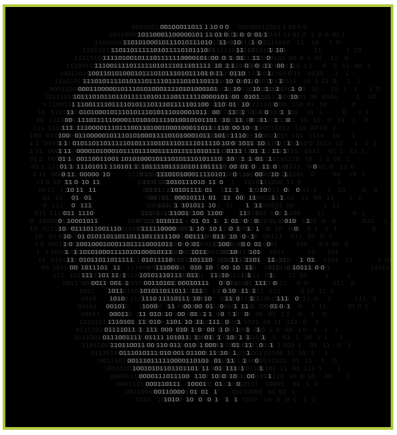

Protecting Business-Critical Data Against Cyber Threats

THE SCOURGE OF RANSOMWARE

Ransomware has grown to be the number one cyberthreat that organizations must protect against. Historically, data protection strategies have been put in place to provide an organization with the ability to recover from software and hardware failures within their data center. Unfortunately, cyber threats are forcing more organizations to consider malware (and ransomware in particular) as the new primary threat to protect against.

Once the attackers are inside, they observe your infrastructure for weeks or months. They detect what your business-critical data is, and they learn about your backup strategy, your snapshot policies, and your data replication setup.

A ransomware attack starts with an intrusion in your data center. Many mechanisms are possible, including malicious email attachments, attacks staged through an unprotected personal smartphone, viruses, network intrusions, USB sticks containing malicious data, or plain bribing of employees and IT personnel. Once the attackers are inside, they observe your infrastructure for weeks or months. They detect what your business-critical data is, and they learn about your backup strategy, your snapshot policies, and your data replication setup.

The attackers go undetected for months. They will first attempt to make your backups, snapshots, and other protection mechanisms unusable for data restoration. Next, they will encrypt all your important data. Finally, they will contact you to pay a ransom fee (often in Bitcoin) to get your data back. They will prove that they can decrypt your data and behave like a professional company with excellent customer service. They typically offer support over email or phone on how to get the ransom paid, and how to get your data back.

TRADITIONAL PROTECTION STRATEGIES DON'T PROTECT AGAINST RANSOMWARE

Traditional data protection involves the employment of many different technologies to create a layered defense. Filesystem snapshots provide the ability to quickly recover deleted or corrupted files. Filesystem backups to tape, VTLs (virtual tape libraries), or deduplication appliances provide longer retention and also serve to keep a copy of the data separate from the original storage system. In the case of VTL or deduplication appliances, the backup data can then be replicated to an off-site location to protect against the proverbial "smoking hole in the ground" should the primary data center be involved in some type of catastrophe.

Traditional data protection was conceived to protect against hardware and software failures and accidental misconfigurations. Ransomware, however was never a design criterion.

What are some challenges with this traditional approach when ransomware is introduced? Snapshots are negatively impacted because ransomware often succeeds in deleting snapshots. Certain types of ransomware will modify files on unstructured storage at a rapid rate. During an attack, snapshots grow very quickly; it is possible that a highly utilized array will become inoperable due to available capacity being completely consumed. Other types of ransomware attacks will stay undetected for months by slowly encrypting data. With the data and/or time of infection unknown, it is unlikely that snapshots can be leveraged for recovery since long-term retention is not their intended use.

Replication to a sibling system simply propagates the ransomware and replaces data with encrypted content, both in primary and secondary systems. Reverting to off-array backups made by Network Data Management Protocol (NDMP)-based software forces the restoration of data to either the original source system or to an identically configured system, which may not be available. Highly sophisticated ransomware will attack backup systems by destroying backups, backup catalogs, or backup configuration. This often goes unnoticed for a long time, since many enterprises do not continuously test the validity of their historical backups as they often lack the infrastructure and automation to regularly and proactively restore backups.

Snapshots, replication, and backup are all still critical components of an organization's data protection strategy; however, they cannot be relied upon in the same manner as when software and hardware failures were the only threat to be considered. Extending the protection capabilities with system-agnostic copies of data that are stored at difficult-to-reach locations provides the logical extension needed to protect against modern cyber threats.

WHY HETEROGENEOUS CAPABILITIES MATTER

NDMP is a protocol that has been in production for a considerable period of time. Rolling back in time to when we look into the data center of the past, we found servers fulfilling all types of roles; a main role was that of the file server. By today's standards, these servers managed small amounts of data. They were backed up using agents installed on the servers directing backup streams to backup servers, which then wrote the data to one or multiple tape systems. Roll the clock forward and large capacity purpose-built network-attached storage (NAS) systems began to appear. The large capacity of data stored by these systems posed new backup challenges because there was too much data to extract from the storage system and send through backup servers and then on to the tape systems. NDMP was developed as an open standard protocol allowing backup operations to be offloaded to the NAS device itself. With NDMP backups, the NAS device's binaries would write the backup streams directly to the backup device.

With source-agnostic, host-based software, the protection operations are largely offloaded from the array and the dataset is stored in an open format that can be restored to *any* desired system.

NDMP helped dramatically with backup activities but this came at a cost. NDMP backup datasets, since they were created by the filer binaries, must be restored to either the same or compatibly configured system. In other words, backup data from an array running version 1.2.3 of its operating environment must be restored to that same type of array running a compatible operating environment. This means that cross-platform NDMP restores are not an option.

Additionally, NDMP backups are yet another workload on the array so most arrays are limited in the number of active streams that can be created at a single time. This fundamentally limits the amount of data that can effectively be written through the backup streams to the backup target. NDMP is showing its age here because NAS systems have grown significantly in terms of capacity and NDMP backups of over 100 TB result in backup times that are unsustainable. NDMP does not lend itself well to incremental forever backup strategies so a regular Level 0 (full) backup is required. If this full backup takes days to complete, there are large holes in the protection scheme. There is no mystery as to what happens when there are petabyte levels of data that many organizations now manage.

Breaking the handcuffs of NDMP-based backup systems requires a software-based approach that avoids the limitations of NDMP while providing source- and target-agnostic capabilities. With source-agnostic, host-based software, the protection operations are largely offloaded from the array and the dataset is stored in an open format that can be restored to *any* desired system. The type of source system no longer matters, which opens the door to massive flexibility when recalling protected data.

A protection solution that will maintain the structure and permissions associated with NAS data is also required. If only file copies are maintained, there will be access challenges if the permissions and associated metadata are not maintained or restored to the system of choice. Additionally, when protecting NAS data, being able to recreate complicated NAS SMB share and/or NFS export definitions is critical to maintain proper access by end users and applications. It is easy to overlook the impact of configuration settings such as share/export definitions, but they are critical to protect in addition to the content and the access permissions.

In regard to system-agnostic capabilities, remote offices or facilities that have storage systems differing from those residing in the central data center also need a solution that enables the protection of that data using a single solution. For example, scheduled copies could be made from the remote sites back to the central data center storage where the remote content can then be included in the main protection scheme – snapshots can be created, backup copies can be made, etc. This allows remote copies or edge data services to be treated in much the same way as the systems residing in your main data center.

AIR-GAPPED "GOLDEN COPY" OF DATA

As you adjust your data protection strategy to mitigate the risks associated with cyber threats such as ransomware, you want to consider the notion of maintaining a storage-agnostic "golden copy" of your data. A golden copy of data involves creating an additional copy of business-critical data and effectively serves as an insurance policy should malware infect your systems and corrupt both primary and backup copies of data. This copy should also be storage-agnostic, meaning it can be restored or recalled to any system while maintaining the original full directory structures, metadata, and permissions.

One of the myriad challenges faced by anti-malware solutions is that they tend to be fundamentally reactive to known threats. As new threats appear, the anti-malware solution must be updated.

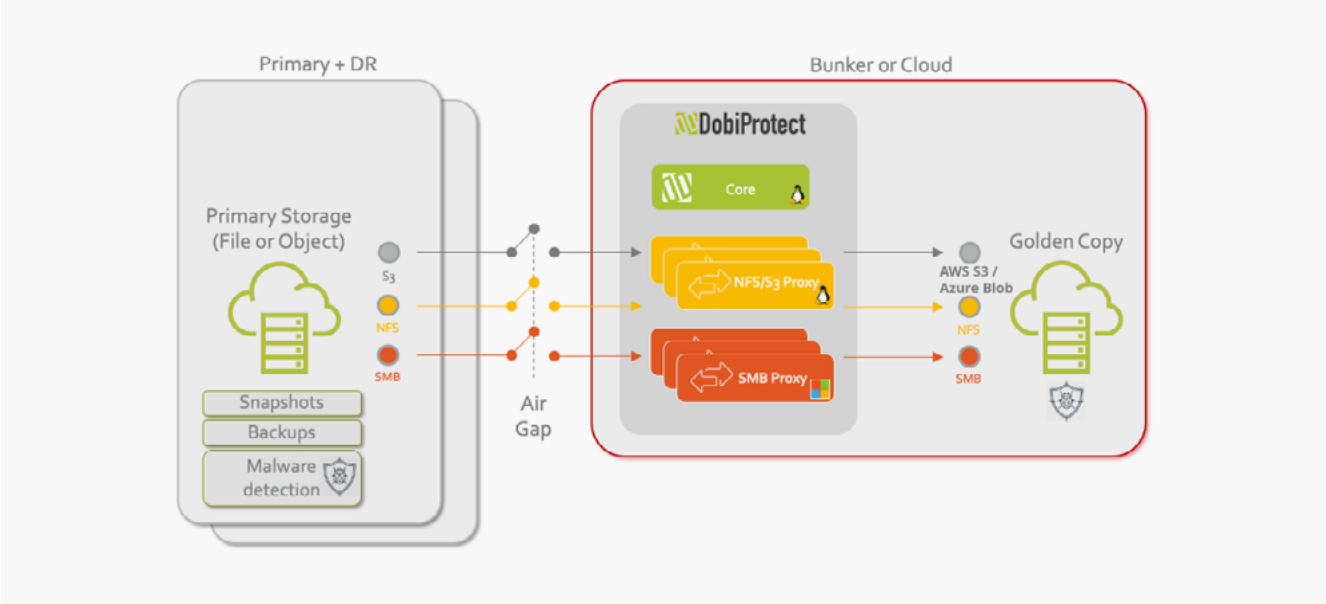

In the meantime, there is unknown vulnerability. In some cases, organizations employ the use of multiple anti-malware products from multiple vendors in an attempt to protect everything from end-point devices (laptops, cell phones, tablets, etc.) to the core file and application servers within the data center. Even with anti-malware products deployed and diligently maintained, it's possible for new malware to evade these systems.

As malware sophistication has grown, so has its ability to target more than just the core data it finds. Backup data in the form of snapshots is at risk with newer threats targeting the corruption or outright deletion of any snapshots present. Even without directly attacking snapshots, malware making massive changes on the filesystem (such as when encrypting files) can generate change rates that can increase the size of the active snapshots to the point where the storage system becomes full and inaccessible. If the storage system is replicating to a sibling system, the replication capabilities are propagating the malware attack across systems. Malware can also reside "quietly" in a dormant state for a brief period before unleashing its full attack. This is where traditional backups can fall short because backups will be potentially capturing malware-infected content on a daily basis.

A data protection strategy involving a golden copy of your business-critical data means keeping an additional copy of data in another site or a public cloud provider. The golden copy provides additional insurance against malware that has managed to corrupt primary, secondary, and even backed-up datasets. The existence and configuration of the golden copy is known to few people in the organization.

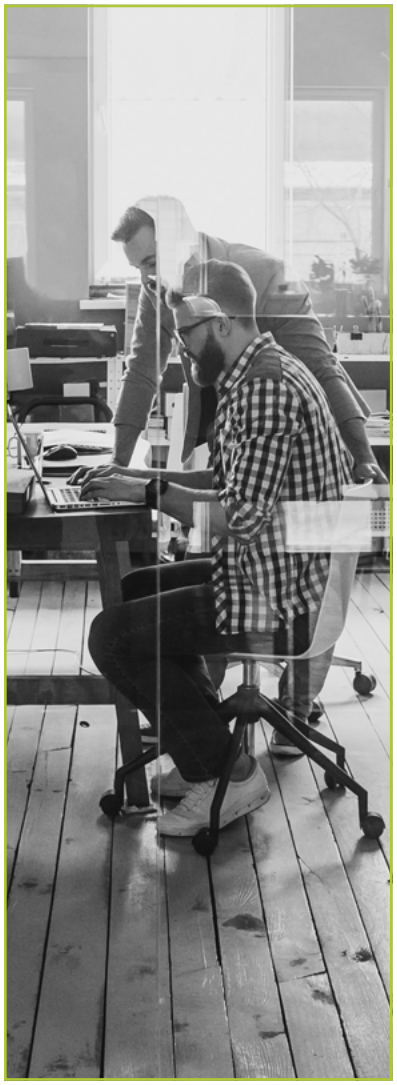

Additional protection offered by the golden copy is created by leveraging an air-gapped environment. The term *air gap* refers to limited network connectivity between the source and the target sites. Instead of constant network connectivity being available, the network connection is periodically activated for the purpose of pulling incremental updates from the source since that last transfer session was initiated. When the incremental updates have been transmitted and verified, the network connection is then terminated leaving the golden copy fully isolated from the rest of the organization's networks and sites. A protection solution such as DobiProtect® has the ability to create a schedule allowing copy sessions to run on a schedule that is coordinated with air-gapped network availability to the target site or public cloud provider.

ABOUT DOBIPROTECT

DobiProtect makes protecting your most business-critical NAS and object data against cyber threats, ransomware, accidental deletions, and software vulnerabilities simple, quick, and cost effective.

SUMMARY

Data availability is more likely to be negatively impacted by cyber threats such as ransomware than traditional software or hardware failures. Malware and ransomware are targeting NAS and backup datasets. To mitigate risk associated with cyber threats, it is recommended to extend your data protection strategy with new tools such as DobiProtect that can work with traditional or cloud-based resources.

Heterogeneous capabilities for unstructured data protection are key. It must be possible to restore data to any available system and to protect any type of source – even at remote locations. DobiProtect provides these critical storage-agnostic capabilities and complements existing data protection measures such as snapshots, site-to-site replication, and backup.

An air-gapped golden copy of data is a critical insurance policy. By definition, the pattern recognition capabilities of security monitoring products are reactive so despite diligent deployments of these solutions it is always possible for new malware to infect corporate systems without detection. This golden copy of data could be the critical piece in being able to resume business operations. An isolated golden copy of data accessible only over an air-gapped network provides even better protection. DobiProtect allows the creation of a schedule so that copy sessions can run in coordination with air-gapped network availability and can work with any source–target combination.

FOR MORE INFORMATION

To request a demo of DobiProtect or to learn more about Datadobi's range of products and services, visit www.datadobi.com.