THE PURE STORAGE ROLE IN MALWARE PROTECTION

The important part Pure's products play in strategies for protecting data from attackers

Enterprises, governments, and institutions that rely on digital data have long used measures like regular backup and replication to remote locations to protect it against procedural errors and disasters. As reliance on data to conduct operations increases its value, it also becomes an attractive target for theft and malicious destruction.

THE MALWARE EXPLOSION

While storage media are occasionally stolen, the primary means of misappropriating enterprises' data is *malware*--software that attackers inject into systems to defeat security measures. A successful malware attack can copy data for illicit purposes, alter it, or destroy it. More recently, *ransomware* attacks encrypt data in place, making it unusable until a ransom is paid to the attacker, often in an untraceable cryptocurrency.

As data security techniques become more elaborate, attacks also become more sophisticated, although organizations remain susceptible to techniques like electronic mail "phishing" and social engineering. Many attacks begin by probing users and administrators to gain access to key systems. Attackers who succeed in injecting malware often use it sparingly at first to avoid detection. It may lie dormant for weeks or months, awaiting activity peaks or other times when organizations are especially dependent on access to their data.

State-directed attackers tend to seek strategic advantage, for example by corrupting, disabling, or destroying key IT infrastructure—a malfunctioning power grid or military network could severely restrict a nation's ability to function. Independent attackers are more likely to be motivated by potential financial gain, or occasionally, by revenge.

Whatever the motivation, the consequences of corrupted, misappropriated, and misused data can be severe. Ransomware attacks can make it impossible for organizations to function, even though their data is tantalizingly present. Faked money transfers can ruin financial institutions. Public disclosure of confidential health records can destroy confidence in providers and have severe consequences for individuals. Even after the immediate damage from an attack is repaired, the damage to the reputation of a formerly trusted institution can persist for a long time.

PROTECTING CRITICAL DATA FROM ATTACK

Thus, organizations that need their online data to operate must protect it, not only against procedural errors and disasters, but against deliberate attempts to corrupt, misappropriate, misuse, or destroy it. In today's highly interconnected digital world, most electronic attacks are from the outside via networks, but there are also potential threats from within. Suborned or disgruntled employees with access to the IT infrastructure can alter, pilfer, or destroy an organization's data if preventive measures are not sufficiently rigorous.

Attackers are becoming increasingly sophisticated, patiently acquiring multiple paths into systems and passively monitoring them for weeks or months to determine the optimal times and mechanisms for attacking. They often eradicate[1] snapshots, online backups, and administrator credentials shortly before attacking, making organizations powerless to recover.

Securing an IT infrastructure against data loss and exposure is therefore multi-faceted, requiring personnel policies, procedures, and rigorous enforcement as well as constantly evolving technology solutions that include analytics in addition to protecting the infrastructure itself.

THE DATA SECURITY SPECTRUM

Because protecting critical data against attackers requires continuous review and improvement of procedures and technology, *Chief Information Security Officers* (CISOs) responsible for the security of their organizations' digital assets are becoming a vital part of executive management teams. CISOs are responsible for:

Preventing Attacks

The best defense against attacks is prevention. Total prevention for all time is obviously impossible, but hardening the IT environment and implementing defensive measures promptly as vulnerabilities become known increases the barriers to attack.

Minimizing the Impact of Attacks

The ability to misuse misappropriated data depends on an attacker's ability to comprehend it. Encryption, both of "data at rest" (stored online) and of "data in transit" (while traversing networks) minimizes attackers' ability to use data for illicit purposes.

Recovering from Attacks

The most urgent priorities after a successful attack are determining the extent of damage to data and restoring IT services including valid data to legitimate users. Tamper-proof logs, snapshots, and backups of key data sets can speed recovery and minimize data loss and corruption, and help with post-attack forensic analysis.

[1] In Pure Storage parlance, *destroying* a data object makes it invisible to hosts, but preserves its contents for a specified period (usually 24 hours) during which it can be restored. *Eradicating* an object obliterates its contents, making restoration impossible.

THE FIVE PRINCIPLES FOR PROTECTING DATA AGAINST ATTACKS

It is a sad fact of digital information technology that as the value of digital data to an organization increases, so does the motivation to misappropriate and/or abuse it. A comprehensive program for protecting data from attacks has five principal components:

Understand the Environment

To secure an IT operation, one must know in detail what it consists of—the equipment, software, applications, administrators, and users that comprise it. Rigorous tracking of additions to and removals from IT infrastructures is key. Unpatched software, weak passwords, storage that isn't "scrubbed" before retirement, and so forth can all lead to data breaches. CISOs must ensure that all equipment and software in the environment has strong security features that inter-operate and don't clash with each other.

Control the Environment

Attackers are opportunistic. They will probe thousands of users via phishing emails and login attempts, and try to exploit known application, operating system, and network vulnerabilities. Ensuring that equipment and software security features are enabled and are updated promptly, that access to key data is limited to applications and users with legitimate needs, and that strong user authentication is rigorously enforced all minimize the chances of successful intrusion. Attack prevention is about consistently maintaining good "hygiene" for the IT environment, including:

- Firewalling against external connections with tightly controlled, documented, with regular review and cleansing of firewall rulesets
- Using malware detection software throughout the environment and updating it promptly as new threats are discovered
- Minimizing administrative access to key equipment by vetting account holders, using role-based access control (RBAC) and multi-factor authentication wherever they are available, and keeping audit logs of all administrative actions
- Using artificial intelligence-based analytics to detect unusual access and usage patterns and acting promptly to determine their legitimacy.
- Regularly re-educating users about the types of attack and the risks they pose.

Minimize the "Surface Area" of Attacks

The more types of IT equipment and software an organization uses, the greater the number of potential vulnerabilities. The more applications and data sets a system or user has access to, the greater the damage an attacker who corrupts that system or user can do. Exposure to a successful attack can be limited by:

- Keeping numbers of equipment types, software versions, cloud providers, and so forth to a minimum and maintaining consistent update and patch levels

- Using storage network zones, VPNs, and VLANs to limit data access to applications and users with legitimate needs for it
- Maintaining rigorous control over IT user and administrator accounts and the access rights granted to them.

Make Attacking "Expensive"

The harder it is to penetrate an IT operation, the less likely an attacker is to invest the effort. The harder it is to make use of stolen data, the less likely it is to be stolen. The harder it is to prevent recovery from ransomware, the less likely an attacker is to use it to attempt extortion. In addition to controlling the IT environment as described above, practices that raise barriers for attackers include:

- Encrypting stored and transmitted data, both within and outside data centers
- Automatically detecting and verifying the legitimacy of successful network connections, especially for bulk data transfer, before allowing them to occur
- Wherever practical, using advanced techniques such as PPK pairs and multi-factor authentication to verify IT user and administrator credentials.

Respond, Recover, and Evolve

While the top priority after a successful attack is restoring services and data, it is also important to understand why the attack succeeded to prevent recurrences.

Once an attack is recognized, the most important immediate action is to prevent further damage, by some combination of blocking network access, stopping applications or entire systems, disconnecting storage systems, and disabling access for suspect users and administrators.

Without adequate before-the-fact recovery measures, paying ransom and trusting a ransomware attacker to provide encryption keys may be the only alternative. Some attackers "double-dip"—pilfer a copy of data before encrypting it, and after it is ransomed, threaten to publish it unless a second payment is made.

If snapshots or backups of data, applications, and operating system images that pre-date an attack are available, they can replace corrupted data and software. Updates made during or after an attack can be recovered and re-applied to restored snapshots, provided that malware and data affected by it can be identified and removed from them.

To avoid future recurrences, successful attacks must be analyzed to determine how they were made and preventive measures deployed.

ELECTRONIC ATTACK VECTORS

Virtually all processing, communication, and storage equipment and software in a modern data center is "intelligent," and can be administered via network connections. This is the only practical way to manage the scale and complexity of enterprise IT, but it increases the number of "vectors" through which an attacker can gain illicit access:

- Access to *hosts*[2] via operating system and application interfaces exposes data in the form in which it is processed
- Access to switches and routers can override VPN and VLAN configurations and expose data in the form in which it is transmitted
- Access to storage systems exposes data in the form in which it is stored and manipulated internally by those systems.

It is therefore vital to prevent unauthorized access to the entire data center "stack," of which storage is only a part. The remainder of the brief describes the Pure Storage product features that help organizations create and maintain comprehensive anti-malware programs.

HOW PURE CAN HELP

Securing an organization's IT infrastructure and data against malware attacks requires policies and procedures as well as technology. Policies and procedures are the responsibility of the organization, but Pure Storage products and services can make a significant contribution to the technology pillar of a comprehensive IT security program.

Three types of features designed into FlashArray™ and FlashBlade® systems since inception and continually enhanced are particularly well-suited to defending against malware and its impacts:

Encryption

Systems encrypt all user data and metadata all the time, using autonomously managed encryption keys. Encryption makes it extremely difficult for attackers who gain physical access to systems' storage media to exploit any data they might be able to retrieve.

Data recovery

Immutable snapshots of *data objects*[3], originally designed to enable recovery of data corrupted or lost due to application or administrator errors, are also useful to expedite recovery from data corruption and ransomware

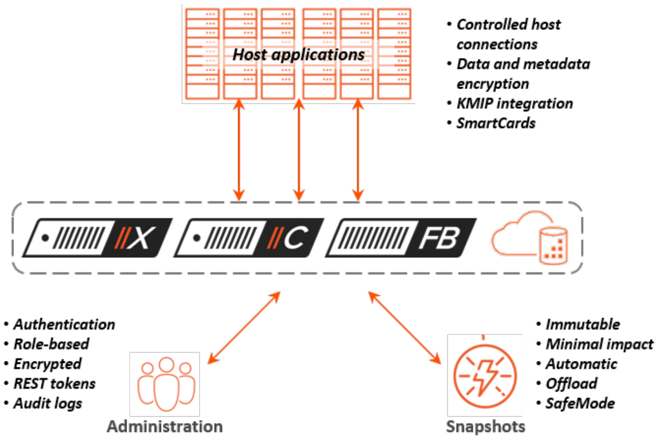

Figure 1: Attack Vectors and Protection Mechanisms

[2] This brief uses the term *host* to denote any computer that connects to a FlashArray, FlashBlade, or other storage system for the purpose of accessing and processing data.

[3] This brief uses the term *data object* to refer to volumes, file systems, and object stores collectively.

attacks. Snapshot contents cannot be changed;[4] they always represent the exact state of data when they were taken.

Administration

All host access to FlashArray and FlashBlade systems is explicitly controlled by administration. Systems do not respond to unauthorized hosts.

Administration itself is role-based, and offers authentication options including major identity managers (AD & LDAP). Systems connected to Pure1®, as most are,[5] use mutual TLS (mTLS) to authenticate every transmission, and RemoteAssist sessions require multi-factor authentication of Pure Storage Support engineers. Enabling *SafeMode™* (discussed on page 7) requires Pure Storage Support participation to eradicate snapshots or change schedules.

ENCRYPTION

FlashArray and FlashBlade systems both use the well-accepted AES-256 algorithm to encrypt *all* data and metadata stored in flash and staged in NVRAM. Encryption is always on—there is no way to disable it. The products manage encryption keys autonomously, refreshing them daily and on events such as device removal, and never expose them on any external interface. Keys themselves are encrypted and partitioned, with each partition stored on a different device. The partitioning algorithm requires more than half a system's devices to reconstruct a key.

Thus, even if an attacker should gain direct access to storage media, for example by acquiring obsolete devices removed from a system, decryption would require more than half the devices in addition to knowledge of the locations of key partitions and the key decryption algorithms.

To accommodate customers who centralize encryption key management, the products support certain *Key Management Internet Protocol* (KMIP) servers as an alternative to local key storage. For situations in which physical security of equipment cannot be guaranteed, FlashArray systems can be fitted with *SmartCards* whose removal renders stored data undecryptable.

The Appendix lists documents that describe FlashArray and FlashBlade data encryption, key management, and access control in further detail. These are available on purestorage.com or in hardcopy form from company representatives.

DATA RECOVERY

FlashArray snapshots are point-in-time images of administrator-defined sets of data objects that capture policies (schedules, replication targets, and retention periods) as well as data. FlashBlade snapshots capture point-in-time images of file systems. Both are *immutable*—administrators can destroy and eradicate them, but cannot alter their contents.

[4] Administrators can *clone* writable volumes from snapshots, for example for forensic analysis, but the snapshots themselves remain unchanged.

[5] With the exception of those deployed at "dark sites."

Snapshots are economical—they consume storage only when hosts make changes to the data objects from which they originate. Benefits of the space-saving design include:

- Creation is near-instantaneous, as is recovering lost or corrupted data objects by replacing them with a snapshot image.
- Impact on I/O performance and space consumption is minimal, so they can be taken frequently (for example, at 10-minute intervals) and retained for relatively long periods.

Administrators can take snapshots at any time. More typically, they are scheduled to occur automatically at regular intervals with specified retention periods, providing a range of recovery points. They are stored locally, but can be offloaded to secondary storage for longer retention. For example, a production FlashArray//X might retain a week of hourly snapshots, and thereafter offload them to bulk storage in a FlashArray//C or FlashBlade, or to a public cloud.

SafeMode

FlashArray and FlashBlade snapshots are not visible to hosts,[6] so host-based attacks do not affect them. Attackers who gain administrative access to the products, however, could eradicate snapshots that might otherwise be used to restore corrupted or encrypted data. To avoid this, customers can, in cooperation with Pure Storage Support, enable *SafeMode* which:

- Allows support engineers to extend the normal 24-hour interval between destroying a snapshot and automatic eradication of its image to as much as a month
- Prevents administrators from explicitly eradicating destroyed snapshots or making changes to snapshot schedules
- Requires verbal contact between Pure Storage Support and two or more customer contacts (designated when SafeMode is enabled) to alter snapshot schedules or disable SafeMode.

SafeMode effectively prevents attackers from eliminating snapshots that could be used to recover data corrupted by malware or rendered unreadable by ransomware attacks.

Administration

Administrators must explicitly enable host connections to products. FlashArray and FlashBlade systems do not respond to commands from hosts that have not been connected. Connected hosts can only access data in the form of LBA or file contents. They have no access to data in the products' internal formats. Protecting against attacks on data via hosts is the responsibility of host security management.

[6] Product administrators can create host-accessible volumes and file systems from snapshots, but these have no effect on the snapshots themselves.

Attackers who gain administrative access to FlashArray and FlashBlade systems have no mechanism for altering data object contents. However, in certain administrator roles they can:

- Destroy and eradicate data objects (volumes, file systems, and object stores)
- Connect data objects to unauthorized hosts which if infected, can corrupt object contents.

It is therefore vital to protect administrative interfaces against use by unauthorized parties.

FlashArray and FlashBlade administration is *role-based*. Each administrative account is associated with a role that defines permissible actions. In both products, the ARRAY role has access to all administrative functions, so it should obviously be closely held, and individuals to whom it is assigned should be verifiably trustworthy.

The products are administered via *command line interfaces* (CLIs), *Graphical User Interfaces* (GUIs), and by commands from programs or scripts via *Representational State Transfer* (REST) APIs. CLI and GUI access are validated by an account name and password[7] and all interactions are encrypted. The products may validate passwords, or alternatively they may be configured to use Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) authentication services. Each administrator account is associated with a token generated when the account is created. REST API commands are encrypted, and authenticated by embedding the token of the account under which each one is issued.

Products log all administrative actions, including login attempts. They retain logs locally, typically for several weeks. If they are connected to the Pure1 *virtual private cloud* (VPC), as most are, the logs are also uploaded and retained by Pure for longer periods, primarily for troubleshooting by support engineers. All communication between products and Pure1 is authenticated by TLS Mutual Authentication (mTLS) and encrypted.

The net effect of these measures is that if FlashArray and FlashBlade users limit administrative access to the products to individuals and systems with legitimate needs and enforce strong authentication requirements, it is difficult for an attacker to gain administrative access.

SUMMARY

Protecting data against malware attacks requires trustworthy people and strictly enforced policies as well as technology. With transparent always-on encryption, immutable snapshots, SafeMode, and rigorous administrative access controls, Pure Storage products provide the industry’s best tools to help protect their critical data against destruction, loss of access, misappropriation, and misuse.

[7] FlashArray also supports public/private key (PPK) pair authentication for CLI access.

APPENDIX: RELATED MATERIAL

Technical Brief TB-160201: Securing FlashArray® "Data At Rest"

https://support.purestorage.com/FlashArray/FlashArray_Security/FlashArray_Security_Reference/The_FlashArray_Data_Security_Model_TB-160201

Technical Brief TB-160202: FlashArray® Access Security

https://support.purestorage.com/FlashArray/FlashArray_Security/FlashArray_Security_Reference/Pure_Storage_FlashArray_Access_Security_TB-160202

Technical Brief TB-190701: FlashBlade® Security

https://support.purestorage.com/FlashBlade/FlashBlade_Security/FlashBlade_Security_Reference/FlashBlade_Data_Security_TB-190701

Technical Brief TB-160501: Security in Pure1®

https://support.purestorage.com/Pure1/Pure1_Security/Pure1_Security_Reference/Pure1_Security%3A_Technical_Report_TB-160501

Technical Brief TB-171101: FlashArrays and GDPR Compliance

https://support.purestorage.com/FlashArray/PurityFA/FlashArray_Technical_Reports/FlashArray_Technical_Papers/FlashArrays_and_GDPR_Compliance_TB-171101

Technical Brief TB-180202: FlashBlade and GDPR Compliance

https://support.purestorage.com/FlashBlade/FlashBlade_Security/FlashBlade_Certifications_and_Compliance/FlashBlade_and_GDPR_Compliance_TB-180202

THIS DOCUMENTATION IS PROVIDED “AS IS” AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.