The true story of data-at-rest encryption & the cloud

KAREN SCARFONE | **PRINCIPAL CONSULTANT | SCARFONE CYBERSECURITY**

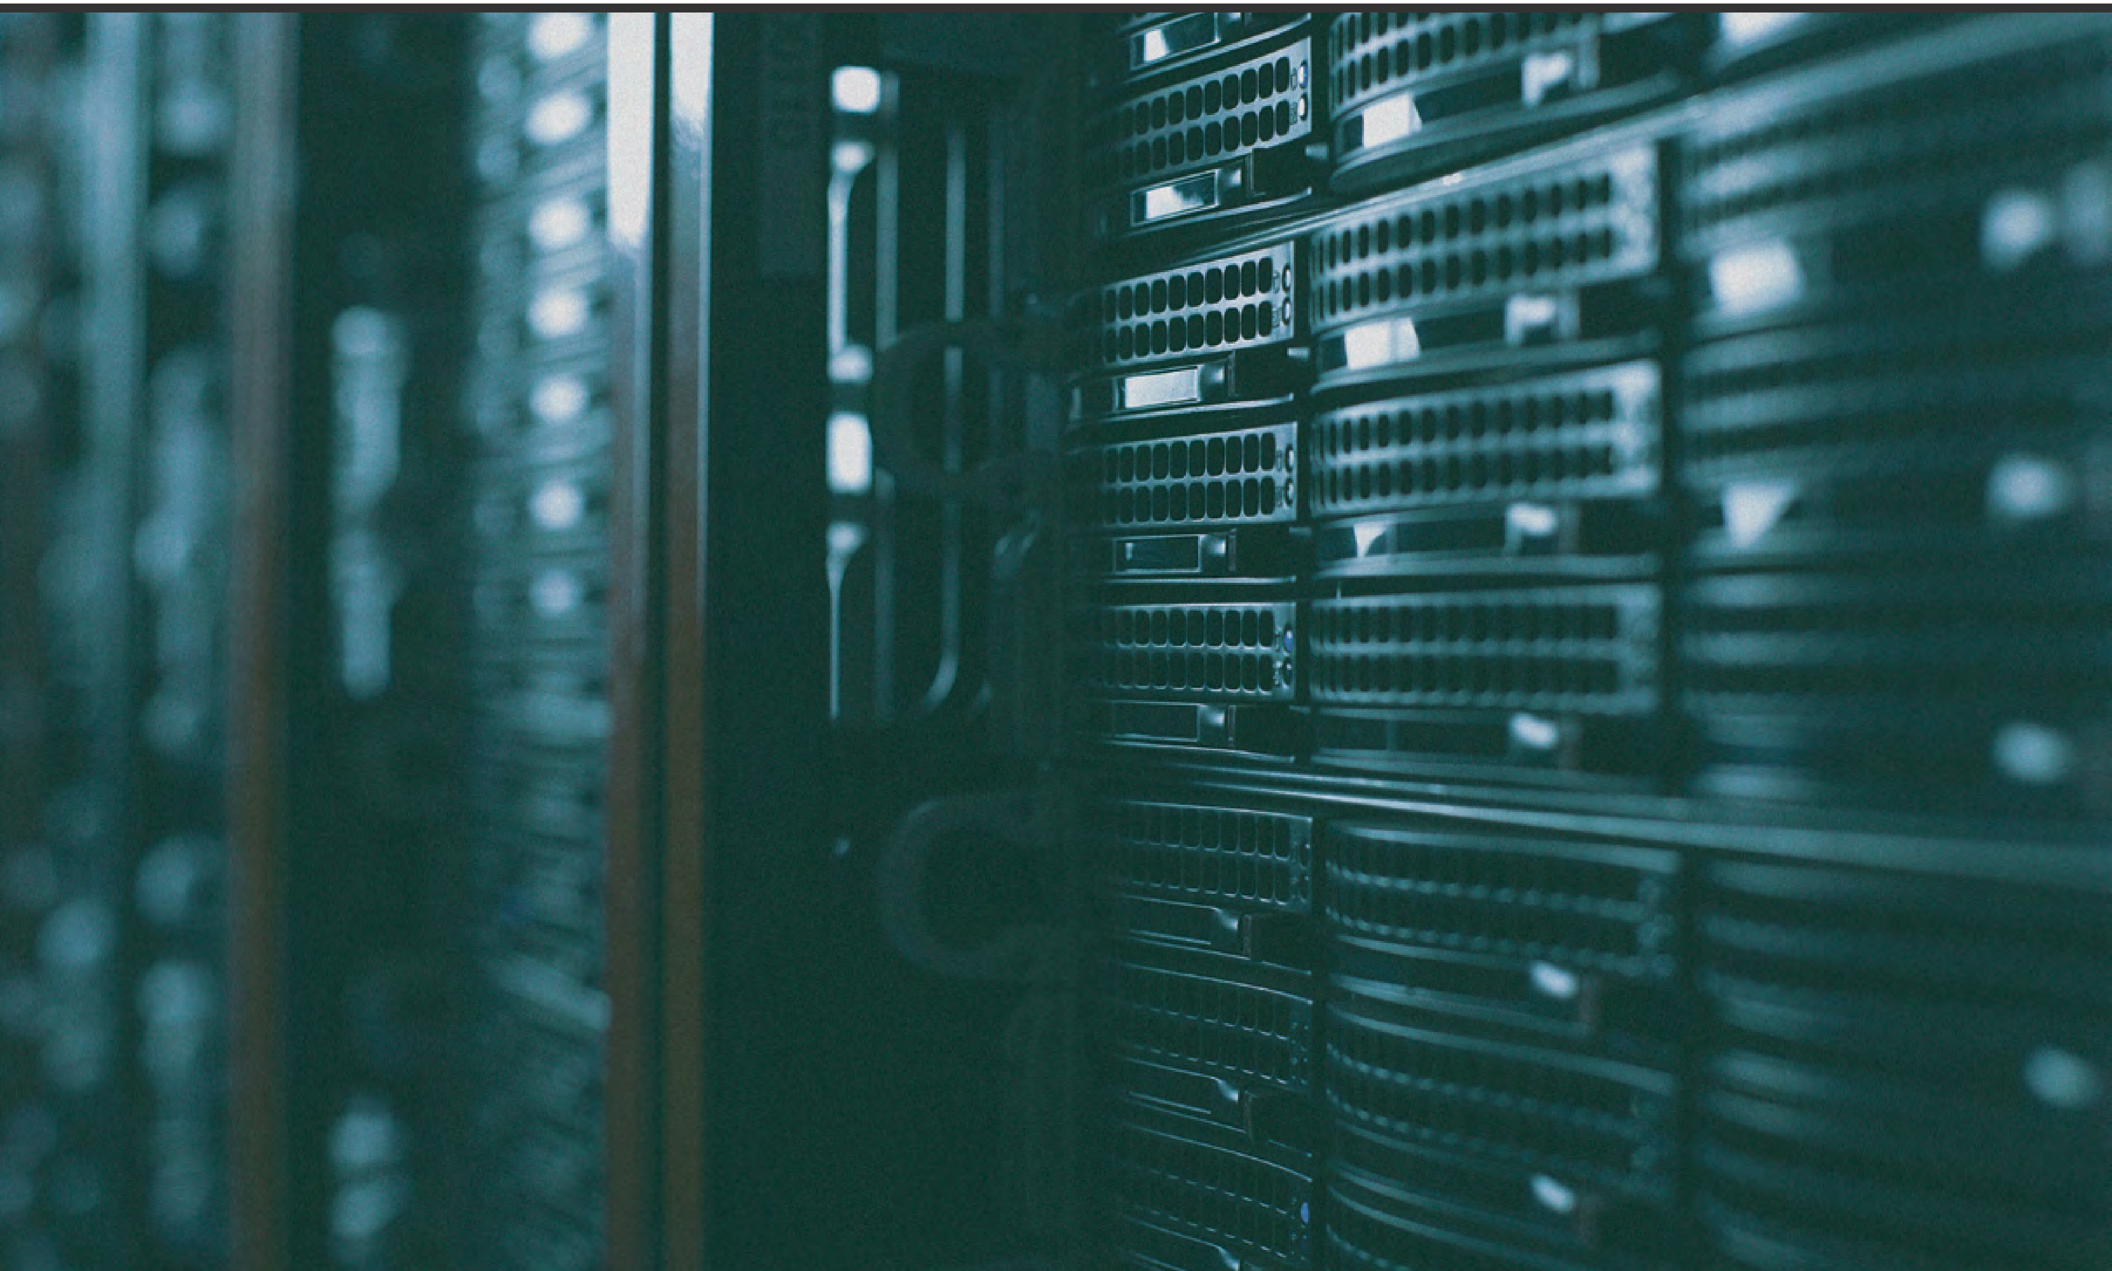

About Karen Scarfone

Karen Scarfone is the principal consultant for Scarfone Cybersecurity in Clifton, Va. She was formerly a senior computer scientist for the National Institute of Standards and Technology (NIST), where she oversaw the development of system and network security publications for federal civilian agencies and the public. She has co-authored more than 50 NIST Special Publications and Inter-agency Reports during the past 10 years, including NIST Special Publications 800-111, Guide to Storage Encryption Technologies for End User Devices, and 800-123, Guide to General Server Security.

Scarfone Cybersecurity | NIST

Executive summary

Encrypting sensitive data stored in the cloud — data-at-rest — prevents attackers from gaining unauthorized access to that data. Standards such as the Health Insurance Portability and Accountability Act (HIPAA) already require this security control to be in place to reduce insider attacks, exfiltration through malware, and other threats.

Unfortunately, these requirements don't specify what type of storage encryption is to be used, and many cloud vendors are choosing the wrong solutions for fighting these threats.

Instead of relying on full disk encryption, which only protects data against unlikely physical theft — when servers are not running — cloud vendors must promote the use of logical/role-based encryption solutions. These solutions are effective whether a server is running or not because they limit access to data based on permissions/roles, providing much greater protection against today's threats than full disk or storage area network (SAN) based encryption.

What's more, encryption should be implemented and controlled by the customer, not the cloud provider, so that the cloud provider does not have insider access to the sensitive data being protected by encryption.

> "… encryption should be implemented and controlled by the customer, not the cloud provider, so that the cloud provider does not have insider access to the sensitive data being protected by encryption."

The core problem

Storing sensitive data in the cloud, also known as data-at-rest, subjects that data to certain risks inherent in any computing environment, as well as a few risks specific to cloud deployments. The main threats causing risks to cloud data-at-rest are as follows:

For illustrative purposes, healthcare-related examples are used in this white paper, with a focus on HIPAA compliance. However, the concepts and conclusions presented in this paper are equally applicable to other sectors and security compliance efforts, such as the Payment Card Industry Data Security Standard (PCI DSS).

Malware

If a cloud instance becomes infected with malware, this malware could then be used to access all the sensitive data from the cloud instance and exfiltrate it to an external site chosen by the attacker.

Malicious Insiders

One threat is malicious insiders from the organization gaining access to sensitive data and exfiltrating it, much like the malware above would do. However, in cloud environments there's also the threat of malicious insiders from the cloud provider itself.

Cloud Leakage

In cloud environments, particularly public clouds, there's often concern about cloud leakage, which refers to an attack coming through one virtual machine to compromise another cloud instance on the same physical server. Basically, a vulnerability in one virtual environment could be used to gain unauthorized access to another.

Loss of Physical Control

Although cloud servers are normally in an always-on state, cloud storage is sometimes taken offline — for example, when a hard drive is retired or transferred to a vendor for repair or replacement. Such a drive may contain sensitive data from one or more previous cloud instances.

The wrong solution: Full disk or SAN-based encryption

Full disk encryption (FDE), also known as whole disk encryption, is a form of storage encryption that involves protecting all the data on the entire hard drive from pre-boot threats.

This means that it protects the hard drive while the device it is installed in is powered off or first powered on, before the user or administrator provides authentication to enable the device to boot up.

FDE is widely used for laptop security, because laptops are typically powered off when transported and frequently lost or stolen during this transport. Without FDE, an attacker armed with widely available tools could take a stolen laptop and directly access its sensitive stored data, circumventing all operating system-level security controls.

This same type of encryption is also touted by many SAN vendors as being built in to their solutions as a suitable way to protect data. However, the data is only encrypted when the SAN is powered down or when a drive is removed from the SAN.

Once a device is booted successfully, FDE does nothing to protect its stored data until it returns to an "off" state. So, although FDE is a great solution for protecting laptops, it's not as useful for servers and other devices that are on most or all of the time.

It does help protect them in one way — if the SAN, server, or server storage is being transported from one facility to another.

In that circumstance, having FDE protects the data from being compromised if the server or its storage is lost or stolen. But FDE does absolutely nothing to protect a server or SAN that is powered on and running in a cloud data center from malware, insider threats, and other current threats.

The right solution: Logical/role-based encryption

In order to ensure that data on an always-running server is properly protected, using a solution that includes logical/role-based access to encrypted data is critical. The act of encrypting data is generally straightforward, but control over the encryption keys is the most difficult and most important consideration.

There are three primary types of encryption that can be used to encrypt data outside of FDE: application layer, file level and database level. Each of these needs to be mated with a logical/role-based encryption key management solution to provide a proper total solution.

The main difference between these three methods and FDE is that they all encrypt the data before it is written to the disk and require authenticated access to encryption keys to view decrypted data, thus ensuring that the data is encrypted on the disk while the server is operating and only available to authorized accounts.

Due to the highly sensitive nature of encryption keys, their management should always be under the full control of the customer. See the "Best Practices for Encryption Key Management" section on page 7 for more information on the reasoning behind this.

Encryption Guidance

The U.S. Department of Health & Human Services (HHS) Office for Civil Rights (OCR) has issued guidance involving protecting stored health information. The guidance includes the following statement:

"Valid encryption processes for data at rest are consistent with NIST Special Publication 800-111, Guide to Storage Encryption Technologies for End User Devices."

Unfortunately, this statement has caused a great deal of confusion. NIST Special Publication 800-111, as its title implies, is meant to be applicable only to end-user devices — desktops, laptops, etc. — and does not provide any recommendations for encryption on servers or within cloud environments.

That being said, the NIST publication does contain general information on several data-at-rest techniques, and these techniques are certainly relevant for servers.

So while the recommendations of NIST Special Publication 800-111 shouldn't be directly applied to servers, the principles of 800-111 should be, and these principles strongly indicate that full disk encryption is not an appropriate primary protection mechanism for servers.

Types of encryption

Encryption can take many forms, including:

Application Layer Encryption

In this form, the actual encryption of the data is carried out in software within the application. Developers can choose from a variety of strong encryption libraries and implement encryption and decryption routines within the application. It is critical that encryption key management be done securely outside of the application. Most encryption key management solutions involve a form of hardware security module (HSM) that provides for the creation and management of encryption keys via application programming interface (API).

File Encryption

In file encryption, files stored on a cloud server are individually encrypted as they are written to the disk. Accessing the contents of any of these files requires decrypting that particular file only, leaving all the other files encrypted. The main advantage of this type of encryption is that it is transparent to both applications and users, so it does not require any changes to application code or special process to grant user access.

Database Encryption

Database encryption can provide the most granular cloud encryption solution. Individual database records, or even particular fields within a record, can be maintained in an encrypted state and only decrypted individually when proper authorization is granted. Database encryption is widely used to protect electronic health information and other sensitive information stored in databases. Major database platforms including Microsoft SQL Server and Oracle include encryption solutions within specific versions and also provide key management solutions.

All of these forms of encryption have something in common: they can provide protection for sensitive data even when the cloud server is active.

This makes them invaluable for mitigating a wide range of threats. In the past there have been concerns about encryption causing additional overhead; generally, though, this overhead is negligible, and the use of file encryption and database encryption supports the principle of least privilege — granting the minimum access necessary.

Enforcing least privilege through logical/role-based encryption mitigates threats involving malware and malicious insiders accessing and exfiltrating sensitive data outside of the organization's control. Because the data is encrypted at all times except when it is specifically being used, the window of opportunity for the data to be stolen is minimized.

Even malware that gains administrator privileges or an insider with administrative privileges may be stymied by the use of logical/role-based encryption, assuming that best practices for encryption key management are being followed (see "Best Practices for Encryption Key Management" on the following pages for more information).

> "It is strongly recommended that organizations storing sensitive data in the cloud use file or database encryption, as appropriate, to protect that data from unauthorized access."

Best practices for encryption key management

This section will highlight selected best practices for encryption key management. An example is storing the encryption key separately from the encrypted data, so that unauthorized access to the data does not also grant access to the key. Another example is configuring the encryption so that it is transparent, not requiring user or administrator intervention in order to decrypt and re-encrypt data.

Organizations storing sensitive data in the cloud should follow best practices for encryption key management. These practices are intended to safeguard encryption keys and in general to make encryption usable while still secure. Best practices of note include the following:

Maintain control of all private/secret encryption keys.

A common mistake is to allow the cloud provider to control the encryption keys. This creates a new risk, because a malicious insider from the cloud provider could use those keys to gain unauthorized access to the customers' sensitive data stored in the cloud.

Allowing another party to have access to the encryption keys raises issues of accountability. It's fine to use encryption services offered by the cloud provider or a reputable third party, as long as the party offering the services doesn't get access to the encryption keys.

Store encryption keys separately from encrypted data.

Suppose that encrypted medical records are stored in the same logical volume as the keys used to encrypt those records. This may be convenient, but unfortunately it also makes it much easier to gain unauthorized access to the encrypted data.

A single compromise can allow an attacker to access both the keys and the data they protect, effectively circumventing the encryption. Encryption keys should also not be stored within application configuration files or compiled into the application itself. A best practice here

Configure encryption to be transparent to users.

If users are prompted to enter keys and passphrases every time they want to access protected information, they're quickly going to try to circumvent those protections.

To make security usable, it's critical that encryption be as transparent as possible, ideally so that users aren't even aware it's in use. Many solutions, including both Microsoft SQL Server and Oracle, offer transparent encryption options.

Conclusion

Sensitive data-at-rest in the cloud is subject to major threats that can lead to data breaches. Although full disk encryption is often used to protect this data, it is ineffective against nearly all the major categories of threats because it only works when the cloud servers are powered off.

Instead of or in addition to full disk encryption, cloud customers should use strong logical/role based encryption technologies, such as file encryption or database encryption, to protect their sensitive data from unauthorized access. These technologies protect data while the cloud server is in operation.

Organizations with sensitive data stored in clouds should encrypt this data in such a way that they maintain control over the encryption keys. These keys should be stored separately from the encrypted data to prevent a single compromise from granting access to both the keys and the data they protect.

And encryption should be configured to be transparent to users so that it does not affect usability.

If you are unsure what form of encryption is protecting your sensitive cloud data, don't hesitate to contact your cloud provider and ask what forms of encryption they provide, if any, for cloud storage. You may find that they aren't providing sufficient protection for your data.

Don't panic — there are a variety of third-party encryption services available that will protect your data while still giving you full and exclusive control over it. But act quickly to get a solution in place before your organization becomes the subject of the next data breach headline.

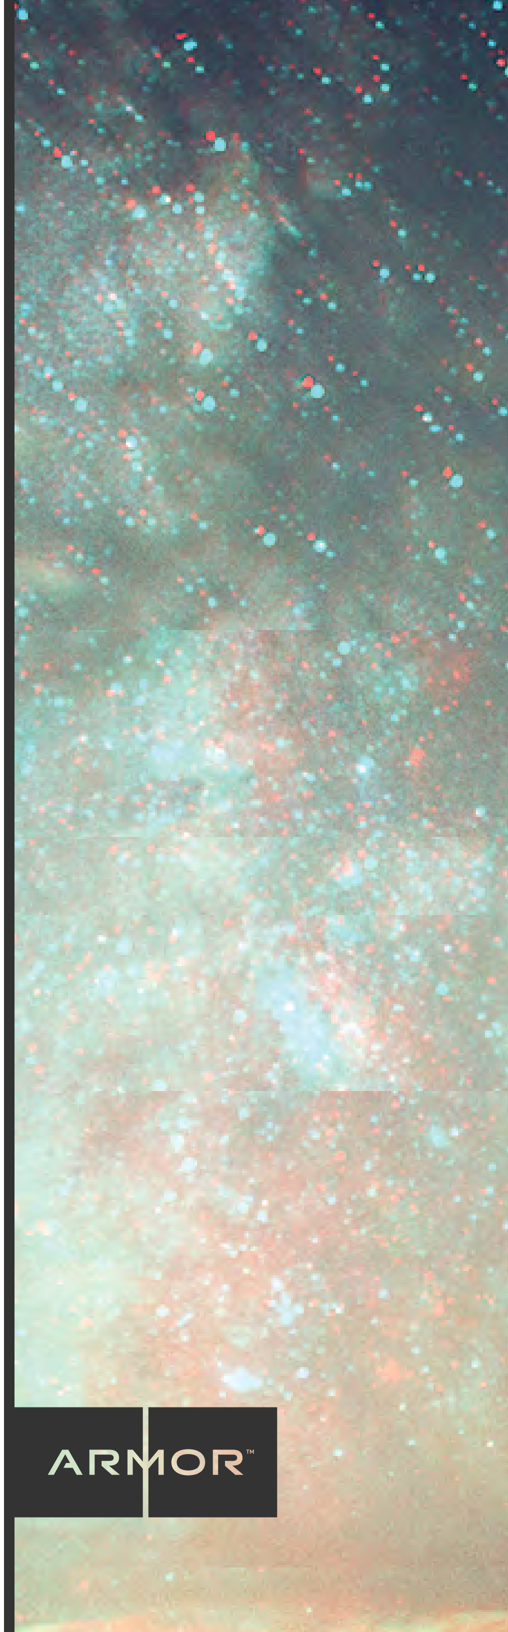

US 2360 Campbell Creek Boulevard, Suite 525, Richardson, Texas 75082 | Phone: +1 877 262 3473
UK 268 Bath Road, Slough, Berkshire SL1 4AX | Phone: +44 800 500 3167

ARMOR™