WHITE PAPER

BETTER DATA BREACH PROTECTION WHILE REDUCING COSTS WITH VSPHERE VM ENCRYPTION AND INFINIBOX

vmware®

INFINIDAT

Contents

The Increased Business Risk from a Data Breach

With organizations undergoing digital transformation, data, and the business's ability to leverage it, has become a key component of the business's strategic advantage. A lot of that data falls under various privacy regulations, from the European GDPR to localized state-wide privacy regulations to market-segment-specific regulations such as HIPAA (healthcare) and NYDFS Cybersecurity (financial).

Government regulators across the world have made it clear that by profiting from user's private data, businesses also take on the responsibility of guarding that data from unauthorized access, whether internal or external. At the same time, businesses face the reality that in the long run, a data breach is inevitable, and they must design for failure, planning for the day after the data breach.

The risk of a data breach is even higher in light of the massive credentials theft from insecure websites, and the realization that many users reuse their corporate passwords, effectively poking a hole in the perimeter defenses of the data infrastructure. The 2017 Verizon Data Breach Incident Report places the number of stolen credentials at 1 Billion, while the McAfee Cloud Adoption and Risk Report shows that 92% of organizations surveyed had their credentials for sale on the dark web.

> ***"You should always have data encrypted […] there's still lots of breaches, but this dramatically reduces the attack surface." (Pat Gelsinger, VMware CEO)***

The only way to design for the day after the data breach is to prevent attackers from accessing the stolen data through anonymization or encryption.

This is the only way to show the regulators that your organization has acted in good faith in safeguarding data privacy:

> ***The communication to the data subject referred to in paragraph 1 shall not be required if any of the following conditions are met:***
>
> *The controller has implemented appropriate technical and organizational protection measures, and those measures were applied to the personal data affected by the personal data breach, in particular those that render the personal data unintelligible to any person who is not authorized to access it, such as encryption*
>
> ***Source: GDPR, Article 34***

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Anonymization

More often than not, anonymization of the data requires redesigning database structures which impacts the application design and creates a high cost for the business. Anonymizing new applications from day one, so that only a subset of the data needs to be encrypted, is a good practice for new applications. For existing applications, transparent encryptions offer the simplest migration path.

Transparent Encryption Alternatives

Many organizations have a false sense of security by utilizing encryption at rest, which only encrypts the data inside the storage media. This is easy to implement, as it requires zero changes throughout the IT stack. However, that is also its biggest weakness – it exposes the data to anyone who was able to gain access to any layer in the IT stack.
Moving from encryption-at-rest to encryption-in-flight means the data gets encrypted at the source, as early as possible. It is then encrypted as it traverses the network, remains encrypted as it's stored (at rest) as well as when it's replicated to a remote DR site and stored there. As attackers try to exfiltrate the data, they will fail to recognize the patterns they are looking for (social security numbers, credit card numbers, etc.) and therefore are less likely to keep exfiltrated data - an operation that is likely to expose their presence. If they do exfiltrate the encrypted data, it will be useless to them (it can't be sold in the dark net) and your organization will be able to show the regulators that the data couldn't have been abused by the attackers.

Where to Encrypt Data in the IT Stack

The main question for any organization planning to protect data using encryption is where in the IT stack, and there are multiple alternatives.

Application Level Encryption

While considered the most secure option, as data starts its lifecycle already encrypted, it is also the hardest to implement, especially in existing applications. This also poses the challenge of scale, as many existing applications will require rewriting to incorporate data encryption. This is a good alternative for new applications.

Database encryption

Offers the highest level of granularity, allowing DBAs to encrypt only specific columns, tables, table spaces Etc. yet this capability also requires complex data classification processes, so that each column in the database that may contain sensitive data gets encrypted, and those classifications are not easy to maintain over time. There's also a financial consideration here as licensing another CPU core for the DB server is usually more expensive than licensing another core as part of the hypervisor.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Volume / Filesystem level Encryption

This works within the physical or virtual machine to encrypt data as soon as it's stored to persistent storage, and offers a very secure alternative, as even the VM administrator can't see the data. However, there are some weaknesses – the challenge of implementing such a solution at scale, supporting multiple operating systems, some very old and no longer supported (legacy applications).

Hypervisor Encryption

Encrypts the data as soon as it leaves the VM, and keeps it encrypted throughout its lifecycle. Requires no changes to the Guest OS or applications and has a single integration point (hypervisor), hence it is the easiest to implement and the most scalable. vSphere VM encryption offers separate keys for each virtual machine. This prevents the use of that virtual machine elsewhere as connectivity and trust with the Key Management System is required.

Network / data fabric encryption

Provides encryption over the wire only, doesn't solve the need to encrypt the data early at the source as well as when stored & replicated to the remote site.

Operational Implications of vSphere VM Encryption

Key Generation

Encryption requires Data Encryption Keys (DEKs). These are generated by the ESXi host using FIPS-140-2 validated cryptography. Key Encryption Keys (KEK) are generated by the Key Manager using dedicated products that yield the highest levels of entropy (e.g. Hardware Security Module, HSM). vSphere allows a simple integration with external HSMs using the Key Management Interoperability Protocol (KMIP), which oversees securely getting the Key Encryption Keys to the ESXi host when they are needed, among other functions.

Key lifecycle management

The DEK for each VM is encrypted using the KEK from the key manager. Once KEK's are used, they need to be managed, protected and backed up to prevent downtime and data loss.

Data Classification Process

New data being created in existing or new applications need to be assigned the relevant sensitivity classification, so it gets the right level of protection. By encrypting in the application, that classification is made easier by leveraging vSphere

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

policies for the entire VM Since virtual disks include large datasets, they enable “blanket encryption” on the entire dataset, avoiding the need to manage this per database column, filesystem, Etc.

Implications of Encryption on IT infrastructure

Enabling encryption effects IT infrastructure in multiple ways. Understanding them in advance and designing for them is instrumental to a successful encryption implementation.

Application Development

Encryption is a compute-intensive process that can impact application performance. Fortunately, there are optimized instruction sets both in Intel and AMD CPUs that enable developers to reduce this impact. vSphere encryption solutions leverage FIPS 140-2 validated cryptography libraries for VM and vSAN encryption. These are the same libraries for both solutions. In addition, these libraries are used to secure TLS connections and certificate generation by VMCA, the purpose-built certificate authority embedded in vCenter.

Compute Layer

Since the compute layer will need to handle the computational effort of encrypting and decrypting data, it’s important that this layer have some flexibility to add compute power (scale out). Databases and hypervisors are a good example of encryption locations that offer good scalability with different benefits.

Storage

Storage is the most affected tier when it comes to encryption in flight due to two components:

- Storage is already among the top two IT budget items
- Many storage arrays rely heavily on data reduction technologies (compression, deduplication, pattern removal) to overcome the high cost of the storage media they use

Since encrypted data has high entropy levels, **current data reduction techniques do not work on encrypted data**. Storage costs, dependent on the space savings from compression, deduplication, and pattern removal techniques, are at risk of increasing by 5x or more if a scalable storage solution is not implemented correctly.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Backup Layer

Backup is another layer that heavily relies on data reduction, which means it also requires attention when deploying encryption.

There are two types of backup scenarios for encrypted data:

Format in which the data is read by the backup client	Data stays encrypted	Data is decrypted (cleartext)
Security implications	More secure, data stays encrypted through the backup lifecycle too	Requires the backup team to handle data encryption while in flight (often over a WAN) and at rest (inside the backup target)
Data reduction on target	No data reduction	No change
Cost	Higher (more capacity)	Lower (same capacity)
CPU implications	Lower	Higher, as data is decrypted and then encrypted again
WAN optimization	Not applicable as data is encrypted	Applicable, data is in cleartext

VMware's Virtual Disk Development Kit (VDDK) allows backup clients integrating with vSphere to see the data in cleartext to maintain as much of the existing operational model and cost structure, however customers who wish to send data encrypted for increased security can leverage storage snapshots to do so.

VM Encryption

What is vSphere VM Encryption

VMware vSphere® virtual machine encryption (VM encryption) is a feature introduced in vSphere 6.5 to enable the encryption of virtual machines within the vSphere hypervisor layer. VM encryption provides security to VMDK data by encrypting I/Os from a virtual machine (which has the VM encryption feature enabled) before it gets stored in the VMDK.

Integration with Backup Through VADP

Use of the "Hot Add" backup method is supported to back up a VM Encrypted virtual machine. The backup vendor's virtual machine will need to be encrypted and the backup user in vCenter will need the Cryptography.DirectAccess permission enabled. At this point, the backup will complete as normal. The backup vendor's virtual machine will snapshot the encrypted VM and mount the parent virtual disk. It will see unencrypted data and will back up the data normally.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Storing Encrypted VMs on InfiniBox

What is InfiniBox

INFINIDAT InfiniBox® provides multi-petabyte enterprise storage with scalability exceeding 8 PB in a single 42U rack, high I/O performance and throughput, seven nines (99.99999%) availability, and multi-protocol support with incredible ease of use. Altogether, the InfiniBox provides unprecedented value for modern enterprise storage. With a disruptive price point, available in multiple configurations, InfiniBox enables customers to acquire, store and analyse the most data to achieve a competitive advantage.

Benefits of Running Encrypted VMs on InfiniBox

The INFINIDAT Software Defined Storage (SDS) approach implemented in InfiniBox leverages Neural Cache, a learning algorithm to optimize performance in real time. InfiniBox leverages a cache-optimized architecture to reduce the cost of storage while at the same time providing the majority of writes and reads from DRAM to accelerate performance beyond what all-flash can achieve. This approach avoids using expensive media to yield high performance, which typically requires the use of data reduction to satisfy the requirements for a cost-effective solution.
The innovation in the InfiniBox software allows customers to gain the highest performance and significantly reduce storage costs, all the while enjoying seven nines (99.99999%) of availability for multi-petabyte systems in a single rack.

Best Practices

When running encrypted VMs, several best practices should be followed for optimal experience:

- Disable data reduction on the target storage array if all the VMs are encrypted, reducing unnecessary CPU load on the storage array.
- Leverage thin provisioning to maintain some data reduction.
- Use the ability of InfiniBox Pools to grow automatically (“Emergency Buffer”) when nearing full capacity, to prevent thin provisioned VMs from filling up the space.
- Enable UNMAP within VMs to clear storage space when data is deleted within a VM (requires vSphere version 6.5 and VMFS6 filesystems).

Additional Resources

- vSphere Virtual Machine Encryption Performance
- vSphere 6.5/6.7: VM and vSAN Encryption FAQ
- Key Manager Concepts and Topology Basics for VM and vSAN Encryption

WP-VMENCRYPT-190313-US

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.