Maintaining Data Protection in a Hybrid, Multi-Cloud World

AN ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) WHITE PAPER

PREPARED FOR IBM

BY PAULA MUSICH

JUNE 2020

EMA™
IT AND DATA MANAGEMENT
RESEARCH | INDUSTRY ANALYSIS | CONSULTING

Table of Contents

MAINTAINING DATA PROTECTION IN A HYBRID, MULTI-CLOUD WORLD

Digital collaboration and its inherent data sharing are a fact of life within the modern, cloud-connected enterprise. The traditional assumption of trust given to appropriately credentialed employees and contractors operating within the enterprise's hardened perimeter has been extended to modern cloud-based architectures. Teams working together to advance business objectives can freely share data across private hybrid cloud, multi-cloud, and on-premises-based applications. In most cases, that data is shared in cleartext. Research on data in the cloud shows that only 9.4% of cloud data is encrypted. If that data is exposed to the Internet or otherwise leaked, the organization has little to no means of recalling or deleting it. In the cases of theft, the overall lack of encryption in use means that once the data is stolen, it's game over.

All of these activities are founded upon sharing data across a tenuous ecosystem of trust that encryption could strengthen. Unfortunately, encryption is often not utilized due to the compartmentalization of technologies that protect data in its various states and the increased user friction at both ends of the transaction. To make matters worse, if a data recipient breaks confidence and shares the data with others (whether on purpose or by accident), the data owner/custodian has no knowledge of the breach of trust unless the offending party informs him or her. Enterprise data must be better and more broadly protected via a holistic, highly integrated set of technologies. This will aid the data owners/custodians in maintaining control and tracking their data throughout its lifecycle.

The McCumber Cube and Three Dimensions of Risk

In 1991, John McCumber released a cybersecurity risk model now known as the McCumber Cube. This model was revolutionary in the way it depicted cybersecurity risk factors as a three-dimensional cube. Each of the visible faces of the cube has three different aspects of cyber risk that need to be managed. Each of the three-dimensional intersections represents the union of three components, one from each face. The front-most mini cube, outlined in red, is the intersection of confidentiality, technology, and processing. This represents the idea of a technology control to protect the privacy of data in processing.

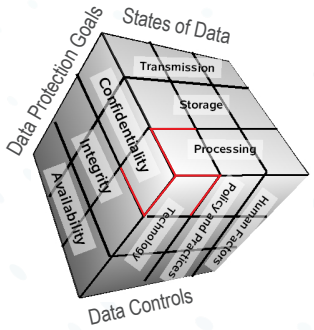

Figure 1: McCumber Cube

This paper will discuss how the proper application of encryption to each stage of the data's lifecycle, transmission, storage, and processing will create better controls that improve protection and privacy. In the age of digital data sharing and collaboration, these controls reduce the risk of data exposure, leakage, loss, and theft. This paper focuses on the concept of delivering confidentiality through new encryption controls and how the data controls affect and are affected by the use of encryption everywhere throughout the data's lifecycle.

ENCRYPTING DATA IN TRANSMISSION

The Intersection of Confidentiality, Technology, and Transmission

The proper application of encryption can significantly reduce the losses from data theft. In the McCumber Cube, the intersection of technology, confidentiality, and transmission has been addressed by the use of Transport Layer Security (TLS) protocol and its predecessor, Secure Sockets Layer (SSL). By the end of 2019, securing data in transmission over the Internet had greatly improved, with nearly 95% of website-focused traffic encrypted.

There are pros and cons to using transport-based encryption throughout the internal environment. The main benefit to those who own the data is the higher level of confidentiality it brings when they are moving data for legitimate business purposes. The main problem for those who own the data is the higher level of confidentiality encryption brings to black hats when they are moving stolen data. Security monitoring is often blind to what is in the traffic, or they must invest in gateway proxy tools to intercept traffic with an authorized man-in-the-middle attack to see what it is. This type of inspection can also add latency to applications, creating other problems and making the choice to use transport-based encryption difficult and often expensive.

Creating and implementing the data transport encryption standards for SSL and TLS across browsers is logistically easier than international policy creation, enforcement, and key sharing. However, encryption creates a false sense of security for many Internet users. They believe their transported data is safe because it is encrypted, but once the data moves from transmission to processing or storage, the TLS transport security is dissolved and their data is more susceptible to attack.

It is also interesting to note that even before the transport encryption boom, the number of successful attacks on transport encryption was low, and the number of records stolen from data at rest was high. Stealing data on the fly is a more complex task that requires more technical setup and timeliness, especially to steal targeted data, than the technical setup for stealing data at rest.

ENCRYPTING DATA AT REST/IN STORAGE

The Intersection of Confidentiality, Storage, and Technology

Data at rest is by far the largest theft target. Though figures vary considerably depending on the reporting organization, the high-end numbers from Risk Based Security identify over four billion records in 2016, greater than seven billion records in 2017, more than five billion records in 2018, and just under eight billion records breached in 2019. The Thales Security 2019 Data Threat Report estimated that less than 30% of organizations deploy encryption within critical environments, and the amount of data encrypted falls in the single digits.

Historically, encryption systems have proved to be costly and difficult to install, configure, operate, and maintain. Business users who interact with the data criticize the friction it brings to their work and the negative performance impacts, even to the point of causing failed/lost requests. Thus, in the battle of usability versus security, usability is still winning.

For storage, conquering these issues requires addressing the underlying usability of the crypto system by focusing on crypto-key management. The usability of the data encryption tools to remove friction from the data customers, whether internal employees or external clients, is also important.

ENCRYPTING DATA IN PROCESSING AND USE

The Intersection of Confidentiality, Processing, and Technology

The McCumber Cube model identifies data processing, which can be applied to automated processing within a computer system or to manual data handling before or after being digitized.

Encrypting Data in Processing

Protecting data as it flows through an application is probably the most difficult aspect of data risk control. The first control gate is normal access provisioning. If someone does not have access to the front door of the application, then accessing the data is far more difficult. Beyond that, the focus of control is more often on physically protecting servers or hardening the electronic components that comprise the system. Attacks on data in processing require direct access to the system for probing, direct access to the data being operated on, or malware inserted into the application or system drivers to funnel data out while the application is operating.

Encrypting Data in Use

Depending on how processing is structured, controls may need to be added to defend against people siphoning data as they are handling/processing it. Caution should be taken to ensure that the right data goes to the right person. Traditionally, this is where encryption has been deployed the most. So long as only the proper people were included in the circle of trust, the data owner or custodian could be confident that the data was safe. However, until recently, the crucial limitation was that once the data left the owner or custodian's possession, it was still possible for someone to decrypt the information and forward it to another party unbeknownst to the data owner/custodian.

Defending Against Processing and Use Attacks

While antimalware is helpful for addressing malware used in attacks, locks and security guards are best used to keep people away from the processing systems. Encryption may be applied to processing and even some areas of use to limit data exposure during processing. There are research projects underway, like those with homomorphic encryption, that explore how relevant query information can be extracted from encrypted data without revealing the data itself. They are showing promise, but are most likely years away from commercial use.

Recently, encryption capabilities evolved to bind entitlements persistently with the data throughout its entire lifecycle. This is a tremendous leap forward, allowing the data owner to add, change, or revoke permissions for use to any user even after it has been shared.

REDUCE RISKS THROUGH ENCRYPTION EVERYWHERE

Protecting Data and Privacy Across All Locations and States of Data

In 2018, losses from identity theft were estimated at $1.7 billion USD in the United States alone. The Commission on the Theft of American Intellectual Property estimates that Chinese intellectual property thefts from U.S. companies cost as much as $600 billion USD annually. To avoid these losses and gain the promise of true protection throughout a data lifecycle requires a fundamental change in approach. Until recently, once data was delivered to the target, it was no longer under the original data custodian's control. Full trust was given to the next person in the data possession chain. If that person decided to further extend the circle of trust, he/she did not have to get permission from the original owner.

Data custodians are mandated with sharing only the required information, often making decisions on what is required while aspects of the requirements are still fluid. This puts them in a precarious position. What was approved for internal sharing in the past can later be determined to be out of scope. In the current context of shared data, with or without encryption, protecting the organization by implementing a changed policy and retrieving the now out-of-scope data requires significant effort to undo. In most cases, verifying that all internal copies have been returned or destroyed is impossible. Even with no malicious intent, copies of shared data may have been captured in alternative repositories, such as backups, email, shared folders, and personal drives. The concept of encryption everywhere solves this problem by keeping data protected and private across the enterprise whether it is at rest, in flight, in storage, or in the cloud.

The Intersection of Confidentiality, Policies and Processes, and Human Factors

Policies and processes are the foundational elements of any solid cryptosystem. The policy dictates what can and cannot be shared and the parties that are part of the circle of trust for each protected data element. Unfortunately for traditional cryptographic systems, policies are more often based on a level of trust that humans will follow the prescribed policies to ensure data stays where and with whom it belongs. If the human factors make a mistake, then data leakage and exposure can occur.

No matter how unbreakable an encryption algorithm is and how well policies are documented, if a person decides to intentionally operate outside of them or makes a mistake, the protected data is at risk. In most cases this is merely a nuisance, but in others it can be devastating, causing huge reputational and financial impacts. Intellectual property theft at American Semiconductor is a prime example of the damage organizations can experience. American Semiconductor's stock price dropped by nearly 50% after the theft was discovered.

The first step in protecting data from human factors is to reduce the number of humans who can reveal or share the unprotected data. The data owner can always maintain control of the data entitlements and maintain that control separately from any delivery mechanisms or sharing environments. This creates two controls that can prevent an unintended data release.

The concept of encryption everywhere requires proactive data governance throughout the data's lifecycle. Data protection must be transformed into enforceable policies that create an embedded cryptosystem within the data. If policies are applied to the data prior to internal distribution and stick to that data as it moves, data owners can then rest assured that the entitlements they defined on that data remain with it in whatever state it is in. The owner must also define the length of time those entitlements are valid before a recheck of the control server is required, thus maintaining constant control over who can access shared data no matter where the data travels within the enterprise. Technology is then used to monitor and enforce policies throughout the data's lifecycle. This yields persistent protection that meets the changing demands of the business, from changes in personnel and business partners to other operationally driven requirements.

The assigned permissions are valid and are kept resident and protected within unstructured data. When data access is attempted, a request is sent to the key management system within the data owner's environment. If the requestor has the appropriate entitlements assigned, a temporary token is sent back to unlock the information and allow the entitlements to be exercised. In structured data, the data can be protected at the attribute or table levels. The data maintained at the recipient's location remains within the data owner/custodian's control the entire time. If at any time the data owner determines that access parameters need to change or be entirely revoked, all he/she needs do is change the policy and it is applied to the remote data copies.

Adaptive Policy Enforcement for Lifelong Data Protection

At whatever point the sharing relationship is terminated or the data owner determines that the entitlements require a change, he/she can update the entitlements in the policy engine. Upon the next request when the entitlements are checked, the updated permissions are enforced. The updated permissions are applied to any copies that were created before they were accessed. If full revocation is applied the keys are destroyed, rendering the encrypted data inert. Since the recipient does not have access to the keys, data remains safe from unauthorized use.

DATA PROTECTION WITH ENCRYPTION EVERYWHERE

A technology-enforced policy, persistent data entitlements, a robust key management system, and a strong encryption suite together create a solid foundation for in-depth defense. What makes encryption everywhere unique in concept and application is persistent policy enforcement by creating continuous protection across data centers and the cloud to keep control of eligible data, which can be accessed through a JDBC connection. To make the concept of encryption everywhere a functional reality, data owners must leverage a technology ecosystem, not just point solutions. Today, point solutions are good at what they do but are not designed for a comprehensive encryption everywhere ecosystem. Broad integrations are not where they need to be. Thus, achieving comprehensive data protection requires tight integrations with high interoperability as a design goal. In doing so, data is protected at each phase of its existence.

USE CASES FOR ENCRYPTION EVERYWHERE

The use cases listed identify a combination of IBM and non-IBM components that can be used to deliver continuous data protection and privacy. Though all of the phases of the encryption everywhere concept can be achieved using other vendor solutions, IBM is the only vendor currently offering a tightly integrated ecosystem of solutions on IBM Z to provide continuous protection and privacy of eligible data. The following are components used across the use cases:

1. IBM z15 running z/OS or Linux on Z using pervasive encryption capabilities
2. IBM Data Privacy Passports
3. IBM Z Fibre Channel adapters
4. IBM DS8900F Storage
5. IBM Z Fibre Channel Endpoint Security
6. IBM Hyper Protect Virtual Servers
7. TLS or IPSec
8. Your choice of public and or private cloud(s)
9. IBM Data Privacy for Diagnostics (Vendors and Suppliers)
10. A hardware security module (HSM)
11. Commodity hardware for data processing and storage

Use Case 1: Data Protection and Privacy Within the Family of IBM On-Premises Solutions

For many demanding environments, such as high-volume retail, large banks, credit card processing, and other payments at scale, an IBM processing infrastructure is most likely already in place and the z15 is a foundational technology. Within the z15, pervasive encryption can be enabled to protect eligible data and processing within the system. Data protection and privacy for eligible data can be extended from IBM z15 environments to the rest of the enterprise with Data Privacy Passports' appropriate policy controls.[1]The Passport Controller for IBM Data Privacy Passports can be installed to maintain and manage entitlements and entitlement verifications for eligible data from data sources that can be accessed through a JDBC connection. With pervasive encryption protecting eligible on-box data, the focus can be moved to protecting data that needs to flow in and out of the system.

Once policies are enabled, data does not have to remain on the IBM Z to be granted protection. The data associated with the policies is encrypted before it leaves its host storage. When operating within the full IBM ecosystem, IBM Fibre Channel adapters and switches with ultra-high throughput can be used to move data very quickly within the data center. These are compatible with other Fibre Channel interfaced systems, but if used with IBM DS8900F storage, the protections can be increased by adding IBM Fibre Channel Endpoint Security, which protects data in transit at the hardware level. The Fibre Channel and DS8900F combination also adds data encryption and authentication for data in flight.

1 Disclaimer: Data Privacy Passports supports data sources that can be accessed through a JDBC connection.

Use Case 2: Data Protection and Privacy in Heterogeneous Enterprise Data Centers

Most organizations with other computing platforms already deployed are not in a position to wholly rip and replace their current compute infrastructure. Data Privacy Passports is designed to robustly protect critical and sensitive data residing on virtually any network-connected hardware in those data centers. Once the connection to the IBM z15 is made, protection for eligible data can be applied anywhere in the data center for the duration of the data's existence. The z15 is designed for security with FIPS 140-2 level 4 certified cryptographic HSM. It's also designed for speed and is capable of processing over 19 billion encrypted transactions per day.[2]

Use Case 3: Data Protection and Privacy in any Cloud and Across Shared Data

The IBM z15 with Linux on Z provides IBM Hyper Protect Virtual Servers to create a secure private cloud infrastructure. With Hyper Protect Virtual Servers, workload owners maintain complete control over their workloads and their data. Not even system or cloud administrators have access to the workloads unless permitted by the data owner. Data Privacy Passports can be applied to eligible data in even highly distributed hyperconverged, multi-cloud environments so long as that cloud has access to the Passport Controller that enforces policies. TLS tunnels can be added at the Internet gateways for additional transport security where the communicating endpoints need to be hidden from Internet view.

With data controls in place, any data owner can share data with anyone across the enterprise. Whatever the business needs, data access, distribution, and expiration within the enterprise are under the full control of the policy manager. Data owners can have full confidence that when the needs change, the policy can easily change to adapt to those needs. Once the need no longer exists, both protection and privacy can remain intact. The eligible data in any location within the enterprise can be rendered inert by destroying the local keys through the policy manager using Data Privacy Passports.

Use Case 4: Minimizing the Impacts of Shadow IT

Shadow IT occurs when someone in the organization decides to move or copy data to an unsanctioned location. Moving or copying it without permission creates security gaps and increases business risks. If a data leak or exposure occurs, even by accident, heavy reputational and financial consequences can be incurred. Implementing Data Privacy Passports on all critical or sensitive structured data greatly minimizes the impact of shadow IT. Even if data is copied and moved, it is still useless without permissions. If someone with access to a controlled database but no data access permissions moves the data to an unauthorized location(s), the data remains encrypted, minimizing the company exposure.

2 Disclaimer: This transaction rate is based on internal measurements of a z15 configuration consisting of two 8-way LPARs and a 4-way ICF running with dataset encryption and CF encryption enabled. Using these results, full-size z15 transaction rates were projected using standard LSPR MIPS. The performance that any user will experience may vary.

SUMMARY: PROTECT DATA THROUGHOUT ITS LIFECYCLE WITH ENCRYPTION EVERYWHERE

Maintaining data confidentiality provides the owning parties with business and operational advantages. Despite this fact, virtually every organization underutilizes encryption to protect their data and many are falling prey to malicious threat actors and careless individuals.

Within an organization, the primary issue for protecting data is that the most popular and common tools require different interfaces and separate policies to defend the data in its different states. The tools and management interfaces operate independently, with only loose integrations. The independence makes total cohesion and testing of policies and enforcement difficult and often leaves gaps in protection.

In collaboration settings, user friction and maintaining data control are two of the most difficult aspects. Increased user friction drives people away from traditional encryption platforms. The lack of flexibility in policy control and enforcement for data in the field make data owners and custodians reluctant to deploy protections.

While technologies to protect data at each lifecycle stage are common, the regular changes data experiences throughout its lifecycle make certain aspects of managing data encryption difficult. Regardless of the impediments, companies must define their business requirements to protect all their confidential data in transmission, processing, and storage. Cost/benefit analyses must always take place, but a realistic assessment will almost always show that there is benefit in expanding the use of encryption for sensitive and confidential data.

Be prepared. In some cases, the transition to encrypted data stores can take years. Though quantity of data is a factor, it is not the most impactful. The most difficult requirements to accommodate involve cataloguing and defining the data type and location diversity, user and application entitlements, and the application interfaces for data interaction and sharing. Legacy applications need an upgrade or replacement to perform with encryption, but if the data provides true business and/or operational advantage and is therefore worth keeping, it is worth protecting.

For organizations with highly sensitive data or high-volume transactional systems, running IBM z15 with z/OS or Linux on Z with pervasive encryption and Data Privacy Passports should be a prime consideration. The IBM z15 ecosystem provides unparalleled performance for in-house applications or as the foundation for any type of cloud environment being built. Its native security architecture includes built-in cryptographic accelerator chips, embedded hardware security module chips and services, crypto-key creation and lifecycle management services, encrypted multi-Gbps interfaces, and encryption-compatible high-speed storage. The platform delivers persistent data confidentiality, policy management, and enforcement that support any encryption requirement. Currently, there is no more comprehensive and performant mass production system available.

Regardless of the solution(s) chosen, implementing an encryption everywhere strategy significantly reduces costs for privacy and compliance-related breaches. If the data owner/custodian can provide reasonable proof that any data is leaked, stolen, or otherwise compromised, then notification, forensic, victim restoration, and fines are significantly reduced and sometimes eliminated. Negative brand impact can also be significantly reduced/eliminated. Reducing these factors helps decrease bottom line profit recognition.

To learn more about how your organization can benefit from the IBM approach to a comprehensive data encryption ecosystem using IBM z15 with pervasive encryption and Data Privacy Passports, please visit: https://www.ibm.com/it-infrastructure/z/capabilities/enterprise-security.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates® (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA's clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals, and IT vendors at www.enterprisemanagement.com or blog.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook, or LinkedIn.

Corporate Headquarters:
1995 North 57th Court, Suite 120
Boulder, CO 80301
Phone: +1 303.543.9500
www.enterprisemanagement.com

3933.03022020-06032020.revision9