Move your sensitive data to Azure securely

Astha Malik, Program Manager II,
Azure Global Financial Services

Contents

Introduction

Data is the foundation for financial services. Data breaches are an ever-increasing threat to every industry that needs to be protected against both internal and external threats. FSI has been a primary target for cyber security with 35% of all data breaches because of the sensitive information available to them. Average cost of data breach is predicted to exceed $150 million by 2020.

Data encryption is the most powerful tool that helps secure and prevent compromise of your data. For financial institutions (FI) that routinely deal with large volumes of highly confidential (financial records, trade records) , confidential (PII, IP, regulatory requirements), sensitive (emails, docs) and public (unrestricted) data, security is of utmost importance.

Moving to Azure or using a hybrid platform means you need to be sure that your data is secure during the cloud migration and once it is in the Azure cloud. Protecting data in transit, at rest and in use should be an essential part of your data protection strategy as data could be exposed to risk in all stages. With a constantly evolving economy, FIs are even more data-driven and data security is most instrumental for them. Today, many FIs in Banking and Capital markets and Insurance trust Azure platform to secure their sensitive data. Azure provides you with the right set of controls and tools, and helps you secure the data in all its stages – in transit, at rest and in use.

Data at rest

Data can be exposed while it is at rest, meaning when it is not actively moving from device to device or network to network, and rather persists in the cloud. Encryption at Rest protects against unauthorized data access. The Encryption at Rest designs in Azure use a symmetric encryption (Data Encryption key – DEK) to encrypt and decrypt the data and protect the data encryption key with a Key Encryption Key (KEK) for ease of management, access control, and auditing of encryption. Key Encryption Keys allow you, the customer, to meet regulatory requirements to bring your own keys, manage your own keys, and rotate them as mandated by regulations. You can manage the KEK in your own Azure Key Vault, including bringing your own keys (BYOK) from your hardware security module (HSM) to encrypt data with customer managed keys (CMK). By keeping the data encrypted on disk, it prevents compromise of data by the attacker.

Data in transit

Data can be accessed over the network while moving back and forth between locations – whether it's in motion between user and service, between or across datacenters or during end-to-end encryption. Azure recommends end-to-end encryption for data moving between cloud services with standard protocols such as the Transport Layer Security (TLS). Azure also uses Internet Protocol Security (IPsec), to provide authentication, integrity, and confidentiality of data at the IP

packet level as the data is transferred across the network. In doing so, Azure honors the encryption requirements customers bring in.

Data in use

Data can even be exposed while in use, example while being processed or in memory. An attacker with admin privileges, a malware or a hacker can have access to your data or code while it is being used. Use of technologies such as full memory encryption, enclave technologies, such as Intel's Secure Guard Extensions (SGX) and cryptographic techniques, such as homomorphic encryption (HE) can be used to create secure, trusted execution environments (TEE).
HE allows computations to be done on encrypted data, without requiring access to a secret (decryption) key. The results of the computations are encrypted and can be revealed only by the owner of the secret key. With Azure Confidential computing through TEE, you can now build, deploy, and run applications that protect data confidentiality and integrity in the cloud. TEE is the secure area that runs in an isolated environment protecting data being processed from access outside the TEE. Only authorized code is permitted to run and to access data, so code and data are protected against viewing and modification from outside of TEE. SQL Always Encrypted with Secure Enclaves protect sensitive data in use while preserving rich queries and providing in place encryption.

Data Governance

Azure ensures that your data is handled in a manner that meet customer data protection, regulatory, and sovereignty requirements. Security is built into the Azure platform, beginning with the development process, which is conducted in accordance with the Security Development Lifecycle (SDL), and includes technologies, controls and tools that address data management and governance. You can use built-in and custom Azure policies to set guardrails in your subscriptions. Azure Blueprints like PCI -DSS, ISO 27001 etc. help easily deploy fully governed environments throughout your organization. Azure Governance also helps you manage costs by gaining insights into your cloud spend so that you get the most from your cloud investments.

Azure solution for moving your Sensitive data

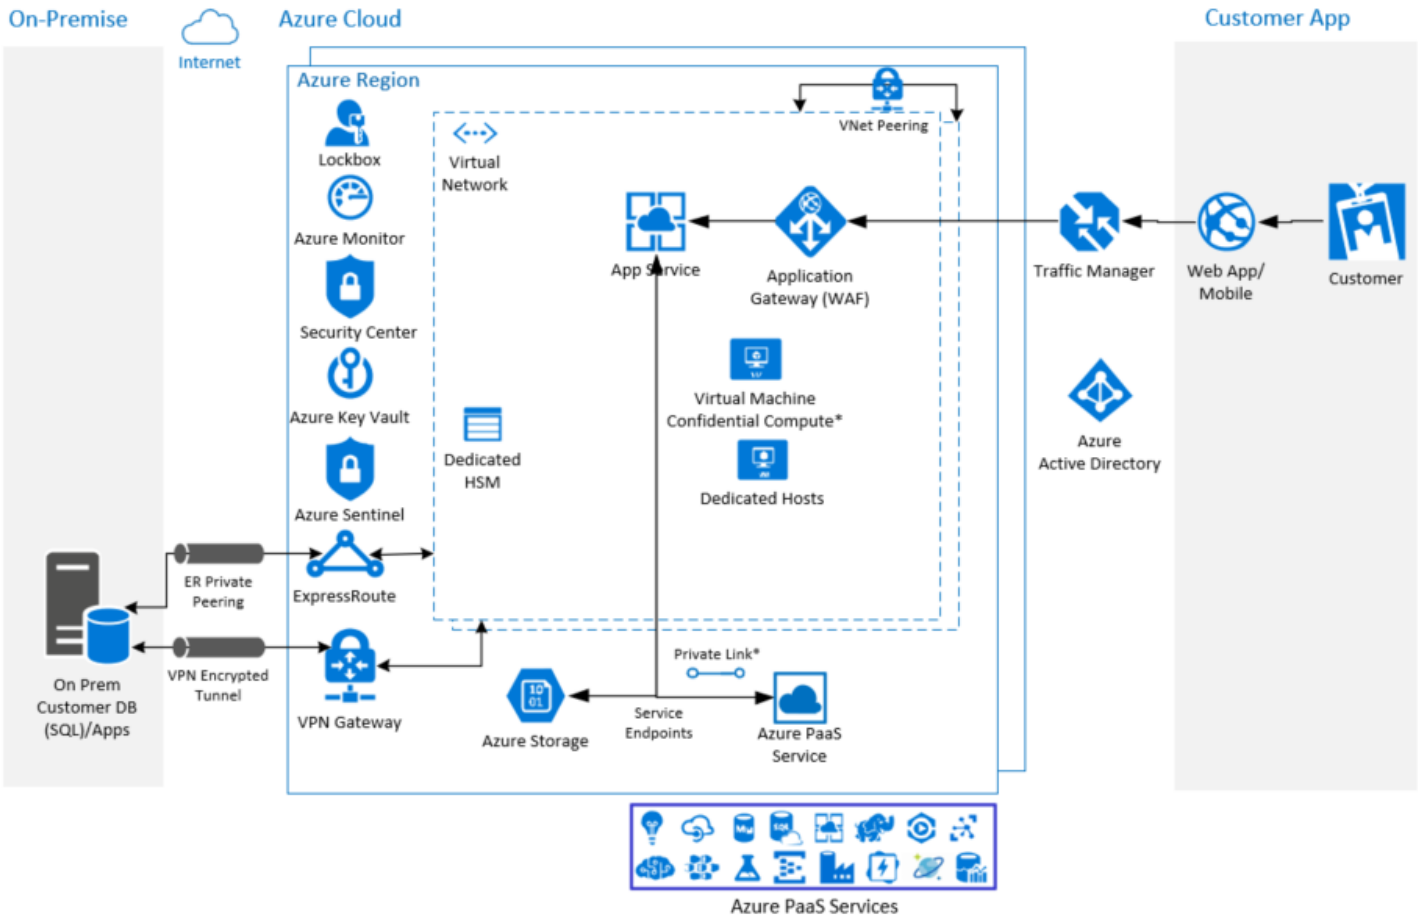

* Preview/Limited support

Storage and Key Management

Azure Storage and keys

Azure storage encryption is automatically enabled for all services and encrypts your data by default while persisting it to the cloud. FIPS 140-2 validated 256-bit Advanced Encryption Standard (AES) encryption is used to encrypt and decrypt the data in Azure Storage. By default, your storage account uses Microsoft managed encryption keys. Managing the KEK in your own Azure Key Vault, including bringing your own keys (BYOK) gives you more flexibility to create,

rotate, disable, manage access policies, network access policies and audit access to the KEK in Key vault.

Azure Key Vault is the recommended key storage solution and provides a common management experience across many Azure services to manage your keys and audit your key usage. It uses FIPS 140-2 Level 2 validated HSMs. You can either create your own keys in your own HSM and securely transfer them to Key Vault (Bring your own key – BYOK), or you can use the Azure Key Vault APIs to generate keys. For custom applications designed to use a traditional HSM, you could use Azure Dedicated HSM to store cryptographic keys. Azure Dedicated HSM devices are dedicated to just one customer, connected to your virtual network with full administrative control by you. Azure Dedicated HSM devices are FIPS 140-2 Level 3 compliant.

As mentioned, Azure uses Microsoft managed keys for data encryption by default. You could also bring in your own customer-managed key encryption keys and use them with your storage accounts. This provides a second layer of control over your data. In this scenario, the Account Encryption Key (AEK) is generated and wrapped or encrypted with the KEK in your Key Vault. Once you write data into the storage account, DEK is generated and the data in storage is encrypted using DEK and random salt. AEK version and the salt are stored as a metadata. If storage loses access to your KEK in Key Vault, the AEK cannot be unwrapped, data goes offline and the access to it is lost.

With more control, comes more responsibility**.** While using CMK, you should be careful of some of the responsibilities that come with it, including:

- Control access to key vault
- Use least privilege access principle to grant access
- Separation of duties
- Lock down access to subscription, resource group and key vaults (RBAC)
- Access policies for every vault
- Turn on Firewall and VNET Service Endpoints
- Use separate key vault for each application instance
- Reduces blast radius in case of a breach
- Take backups
- Take backup of your key on every add/edit
- Turn on logging
- Setup monitoring and alerts
- Use Azure Policy to enable on all Key Vaults, aggregating logs to one location
- Turn on recovery options
- Soft-delete – allows recovery from unintentional or malicious deletion
- Turn off purge – prevent permanent deletion of keys and avoids crypto lock-out

Access to a key vault requires proper authentication and authorization before a caller (user or application) can get access. Authentication establishes the identity of the caller, while authorization determines the operations that they are allowed

to perform. Authentication is done via Azure Active Directory. Authorization may be done via role-based access control (RBAC) or Key Vault access policy.

Azure Active Directory

Azure Active Directory accounts are used for permissions to use the keys stored in Azure Key Vault, either to manage or to access them for Encryption at Rest encryption and decryption. Azure Multi-Factor Authentication (MFA) helps admins protect their organization and users with additional methods. It provides an additional layer of security by safeguarding access to data and applications and delivers strong authentication visa range of easy to use authentication methods.

Conditional Access is a tool used by Azure Active Directory to bring identity signals (Ex- user, device, app, risk detection etc.) together, to make decisions (Ex- black access, grant access), and enforce organizational policies. Conditional Access policies at their simplest are if-then statements, if a user wants to access a resource, then they must complete an action.

Azure Key Vault provides a way to securely store credentials, secrets, and other keys, but your code has to authenticate to Key Vault to retrieve them. The managed identities for Azure resources feature in Azure Active Directory (Azure AD) solves this problem. The feature provides Azure services with an automatically managed identity in Azure AD. You can use the identity to authenticate to any service that supports Azure AD authentication, including Key Vault, without any credentials in your code.

RBAC is used when dealing with the management of the vaults. It is used for controlling management access to resources such as services, virtual machines, storage, and databases. It restricts access to users, groups and applications based on specific scope. Example, you could assign built-in roles like Key vault Contributor to grant access to a user to manage key vaults at a specific scope.

Azure documentation	Azure storage security guide
How to - Github	Blob and Key Vault encryption Azure AD Deployment plans

Azure Disk Encryption and Storage Service Encryption

All managed disks are encrypted at rest in storage with Storage Service Encryption. Azure Disk Encryption further protects with customer-managed keys and encryption within the Virtual Machine, protecting against even those with access to the storage account, meeting your organizational security and compliance requirements. It uses the BitLocker feature of Windows and the DM-Crypt feature of Linux to provide volume encryption.

Azure Storage Service Encryption (SSE) automatically encrypts data before it is stored at rest in Azure Blobs or Azure File share, and automatically decrypts the data when you want to retrieve it using 256-bit AES encryption.

Azure Documentation and How to – Scripts	VMs Windows Linux VMs VM Scale sets
How to - Github	Azure Disk Encryption

Azure SQL Database

Azure SQL Database secures customer data by offering all three encryption types mentioned below:

1. **Data-in-transit with TLS.** TLS is enforced by the server to ensure that the data in transit between the client and server is always encrypted irrespective of Encrypt or TrustCertificate setting. However, we recommend that all applications always set the following parameters (Encrypt = On, TrustServerCertificate = Off), to ensure the client driver verifies the identity of the TLS certificate received from the server and reduce the risk of a "Man in the middle" attack. We also recommend that customers disable TLS 1.1 and 1.0 on the client especially if the application needs to comply with Payment Card Industry - Data Security Standard (PCI-DSS).
2. **Data-at-rest encryption with Transparent Data Encryption (TDE).** TDE is enabled by default for all databases and servers: TDE provides adds a layer of security for your data at rest to protect your data against offline access to raw files or theft. Transparent data encryption encrypts the storage of an entire database by using a symmetric key called the database encryption key (DEK). This database encryption key is protected by the transparent data encryption protector (TDE protector). The protector is either a service-managed certificate (service-managed transparent data encryption) or an asymmetric key stored in Azure Key Vault (Bring Your Own Key). In the BYOK scenario, the TDE Protector is stored in a customer-owned and managed Azure Key Vault.
3. **Data-in-use with Always Encrypted.** Introduced in SQL Server 2016, Always Encrypted is an industry first that protects the confidentiality of sensitive data from malware and high-privileged unauthorized users of SQL Server. High-privileged unauthorized users are DBAs, computer admins, cloud admins, or anyone else who has legitimate access to server instances, hardware, etc., but who should not have access to some or all of the actual data.

Always Encrypted

Always Encrypted with secure enclaves allows computations on plaintext data inside a secure enclave on the server side. A secure enclave is a protected region of memory within the SQL Server process, and acts as a trusted execution environment for processing sensitive data inside the SQL Server engine. There is no way to view any data or code inside the enclave from the outside, even with a debugger.

With secure enclaves, Always Encrypted protects the confidentiality of sensitive data while providing the following benefits not available with the first version:

- **In-place encryption.** Cryptographic operations on sensitive data, for example: initial data encryption or rotating a column encryption key, are performed inside the secure enclave and do not require moving the data outside of the database. Customers can issue in-place encryption using the ALTER TABLE Transact-SQL statement, and do not need to use tools, such as the Always Encrypted wizard in SSMS or the Set-SqlColumnEncryption PowerShell cmdlet.
- **Rich computations (preview).** Operations on encrypted columns, including pattern matching (the LIKE predicate) and range comparisons, are supported inside the secure enclave, which unlocks Always Encrypted to a broad range of applications and scenarios that require such computations to be performed inside the database system.

In Azure SQL IaaS (using SQL Server 2019), Always Encrypted with secure enclaves uses Virtualization-based Security (VBS) secure memory enclaves (also known as Virtual Secure Mode, or VSM enclaves) in Windows. In addition, Azure SQL DB will soon support INTEL's SGX enclaves. (refer to section on Azure Confidential Computing).

In addition to the above encryption features, Azure SQL Database offers a host of core security and advanced threat protection features such as:

- Row-level security
- Dynamic Data Masking
- Azure Active Directory Authentication and SQL Authentication
- Advanced Threat Protection
- Vulnerability Assessment
- Data Classification

Data Classification

Data discovery and classification allows you to discover, classify, label & protect the sensitive data in your databases. It helps you to determine and assign value to your organization's data and is a common starting point for governance. The data classification process categorizes data by sensitivity and business impact to identify risks. Once data is classified, it can be managed in ways that protect sensitive or important data from theft or loss.

Dynamic Data Masking

Dynamic data masking (DDM) limits sensitive data exposure by masking it to non-privileged users. Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to specify how much sensitive data to reveal with minimal impact on the application layer. DDM can be configured on designated database fields to hide sensitive data in the result sets of queries. With DDM the data in the database is not changed.

Azure Documentation	Always Encrypted Row-level Security Dynamic data Masking
How to - Github	SQL DB Always Encrypted code SQL DB – Always Encrypted Key Vault code SQL DB - TDE

Azure Cosmos DB

Azure Cosmos DB offers turnkey global distribution across any number of Azure regions by transparently elastically scaling and replicating your data wherever your users are. It is designed to provide low latency, elastic scalability of throughput, well-defined semantics for data consistency, and high availability. Cosmos DB provides native support for NoSQL and OSS APIs including MongoDB, Cassandra, Gremlin and SQL. It offers multiple well-defined consistency models and guarantees single-digit-millisecond read and write latencies at the 99th percentile, and guarantees 99.999 high availability with multi-homing anywhere in the world—all backed by industry-leading, comprehensive service level agreements (SLAs). Encryption at rest is on by default.

Azure Documentation	Consistency levels Security controls
How to - Github	Cosmos hub Create multi-region account

Networking

Azure lets you connect your on-premise infrastructure and services or colocation environment securely to Azure Virtual Networks.

Azure VPN Gateway

Azure VPN Gateway helps send the encrypted traffic between your Azure virtual network and on-premises location across a public connection or between virtual networks across Azure backbone. Once the VPN Gateway is created, an IPsec/IKE VPN tunnel connection could be created between that VPN gateway and another VPN gateway (VNet-to-VNet), or a cross-premises IPsec/IKE VPN tunnel connection between the VPN gateway and an on-premises VPN device (Site-to-Site) or a Point-to-Site VPN connection (VPN over OpenVPN, IKEv2 or Secure Socket Tunneling Protocol), which lets you connect to your virtual network from a remote location, such as from a conference or from home.

IPSec is used for Site-to-site encryption, and you could use a custom IPSec/IKE policy with specific cryptographic algorithms and key strengths.

Point-to-site (P2S) VPN lets you secure access from your on-premise individual workstation to an Azure network. P2S either uses OpenVPN protocol (SSL/TLS based protocol) or Secure Socket Tunneling Protocol (proprietary SSL-based protocol) or IKEv2 VPN (standards based IPSec solution).You can use your own internal public key infrastructure (PKI) root certificate authority (CA) for P2S connectivity. Site-to-site VPN lets you secure access from multiple workstations located on-premises to an Azure virtual network.

Azure documentation	Point-to-site Site-to-site
How to – Github	Point-to-site Site-to-site

Azure ExpressRoute

Azure ExpressRoute (ER) lets you create private connections between Azure datacenters and infrastructure on your premises or in a colocation environment. ExpressRoute connections don't go over the public Internet, and they offer more reliability, faster speeds, and lower latencies than typical Internet connections. In some cases, using ExpressRoute connections to transfer data between on-premises systems and Azure can give you significant cost benefits.

How to -Documentation	Configure MACsec - PowerShell
How to - Github	VNET to ExpressRoute Circuit ExpressRoute circuit

Virtual Network (VNet) Service Endpoints extends your virtual network private address space and the identity of your VNet to the Azure services, over a direct connection. Traffic from your VNet to the Azure service always remains on the Microsoft Azure backbone network. You no longer need reserved, public IP addresses in your VNet to secure Azure resources through IP firewall.

Azure Private Link

Azure Private Link (preview) provides private connectivity between services and VNet without exposing your data to the public internet. It helps you consume Azure services as well as 3rd party services privately from within your VNet. You could access services running in Azure from on-premises over ExpressRoute private peering / VPN tunnels and peered virtual networks using private endpoints. With Azure Private Link, the private endpoint in the VNet is mapped to a specific instance of the customer's PaaS resource as opposed to the entire service helping against data exfiltration risks.

How to - Documentation	*Create Private Endpoint - CLI* *Create Private Endpoint - PowerShell* *Create Private Link - CLI* *Create Private Link - PowerShell*

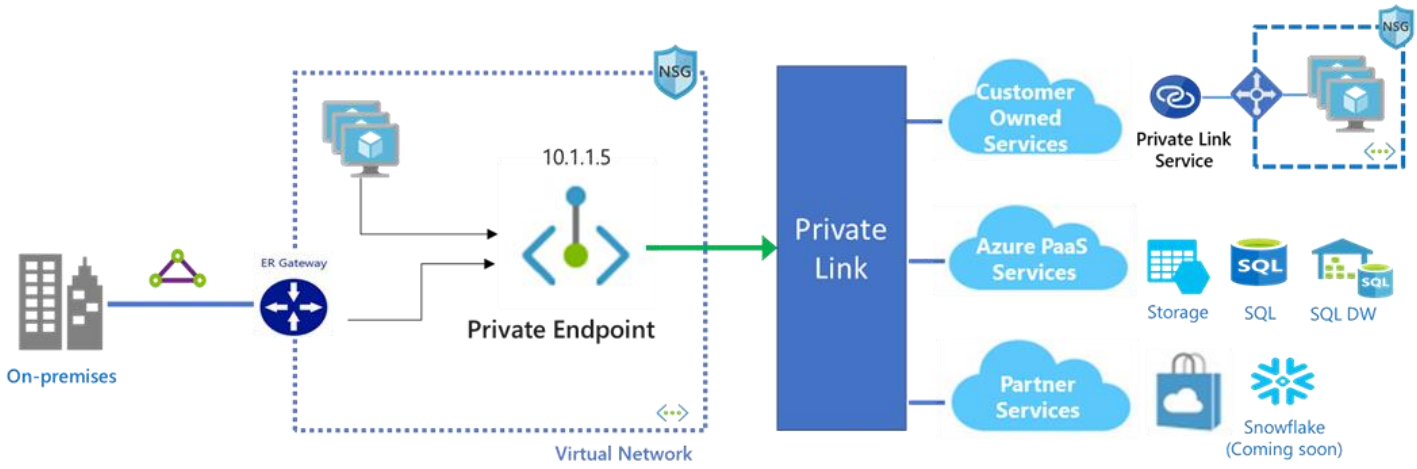

Security

Customer Lockbox

Customer Lockbox for Azure provides transparency and control over Microsoft's access to your content on the cloud. It provides an interface to review and approve or reject customer data access request. You would get explicit control in the very rare instance when a Microsoft Support Engineer may need direct access to your data to resolve an issue. In these rare instances, Microsoft engineer would use Just-in-time service which provides time-bound authorization with limited access to the service. The request would go out to you for review. Until the request is approved, Microsoft engineer will not be granted access. Lockbox maintains an audit of access requests for full visibility.

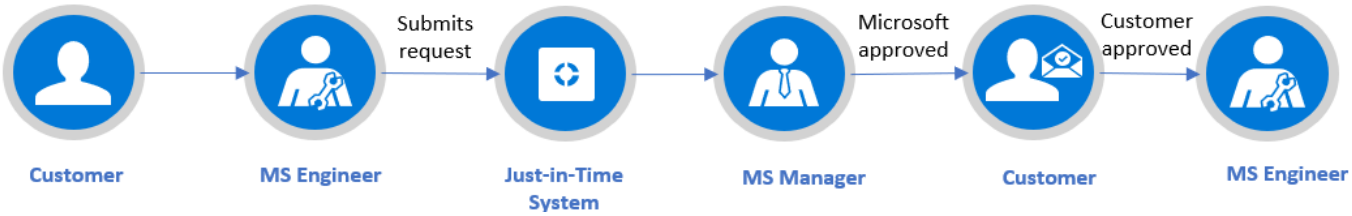

Azure Security Center

Azure Security Center is a unified infrastructure security management system that strengthens the security posture of your data centers and provides advanced threat protection across your hybrid workloads in the cloud, be it Azure, any other cloud, or on-premise.

Azure Sentinel

Azure Sentinel is a scalable, cloud-native, security information event management (SIEM) and security orchestration automated response (SOAR) solution. Azure Sentinel delivers intelligent security analytics and threat intelligence across the enterprise, providing a single solution for alert detection, threat visibility, proactive hunting, and threat response.

Compute

Azure Confidential Computing

Azure Confidential Compute (preview) provides you full control over your data, by securing your code and data while in 'use', beyond at rest and in transit and. It uses TEE to protect your data and safeguards it from malicious and insider threats while in use. Intel Secure Guard Extensions (SGX) and other Hypervisor enclave technologies like Virtualization Based Security allow developers to create secure and trusted execution environments. Enclaves act as a blackbox and provide an encrypted area for data and code that can only be processed by CPU-based security mechanisms in the process-embedded TEE. Only authorized code is permitted to run and to access data, which protects the data modification and viewing from outside of TEE.

SQL Server Always Encrypted

Azure SQL Always Encrypted with Secure Enclaves protects data in use while preserving rich queries and supporting initial data encryption and key rotation in-place without moving the data out of the database. An attestation protocol and attestation service are used to verify the code running inside the enclave is the genuine enclave code.

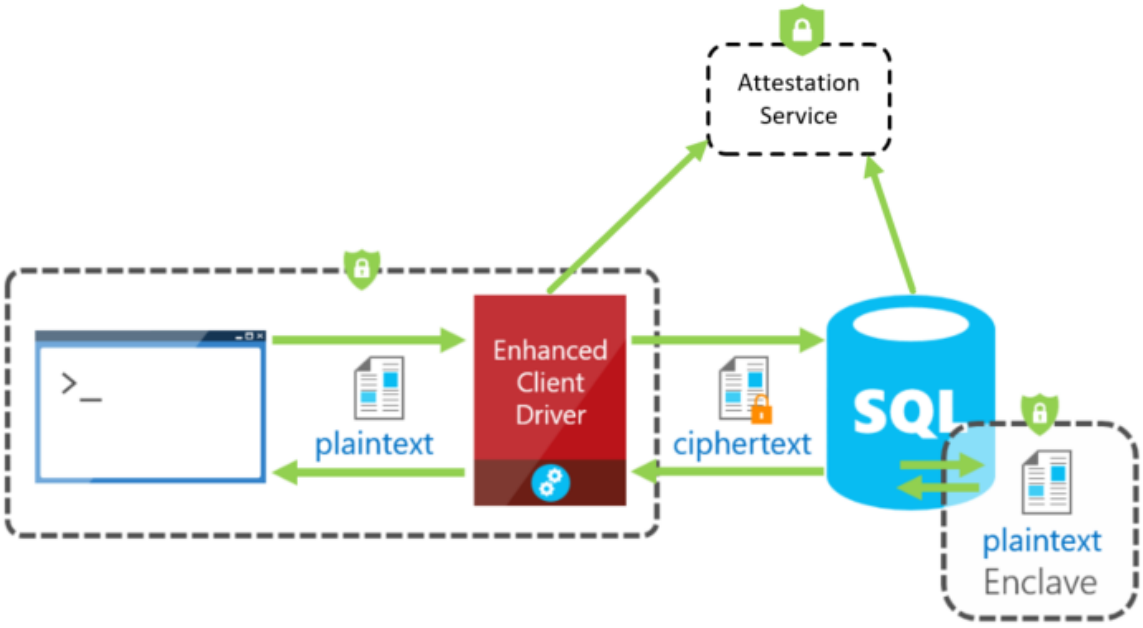

Multi-party collaboration

Privacy-preserving multi-party machine learning allows multiple organizations to perform collaborative data analytics while guaranteeing the privacy of their individual datasets. Data sources from different organizations could be combined using machine learning to better train models, without revealing private data to participants or the cloud platform.

How to - Github	Install Open Enclave SDK Clone Open Enclave Repo
Deployment through Azure Marketplace	Deploy DC Series VM in Azure

Azure Dedicated Hosts

Azure Dedicated Hosts (ADH) lets you have one or more dedicated virtual machines on an Azure physical server. This helps avoid side channel attacks by having workloads run on dedicated hosts. Also, for performance sensitive workloads, where you do not trust other VMs around you, ADH helps avoid noisy neighbors.

How to - Github	Deploy Dedicated Hosts

Application

Application Gateway

Azure Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications. It supports end-to-end SSL/TLS encryption as well as SSL termination. SSL certificate needs to be added to the application gateway listener for SSL termination. However, in case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway. Web application firewall (WAF) feature of Application gateway provides centralized protection of your web applications from common exploits and vulnerabilities, example SQL injection, bots, crawlers, cross-scripting etc.

How to - Github	Deploy Application Gateway in a VNet With WAF, end-to-end SSL, HTTP to HTTPS redirect on IIS servers

Conclusion

Data security and privacy are among the topmost concerns for financial institutions. Microsoft Azure provides confidentiality, integrity, and availability of customer data, while also enabling transparent accountability. As shown above, Azure provides financial services customers with features that build confidence that their sensitive data will stay secure and private. Whether a customer's data is at-rest, in-use, or in-transit these features can be used to protect it and the workloads that they support.

For more information on Azure Security, please visit Azure Security page and Azure Security Team Blog.

For Azure security recommendations and mapping of Azure services to security capabilities, download Azure Cloud Security Benchmark (draft).

Useful Links

Service Trust Portal

Azure Governance – Azure Policy, Azure Blueprint, Azure Cost Management

Azure Blueprints – PCI-DSS, ISO 27001, NIST, SWIFT CSP

Data Residency and Security