CHAPTER 5

Application and Data Security Patterns

Organizations typically focus on building applications and data-driven software to harvest value from their data. Application security mainly attempts to prevent data or code from being stolen. Application security considerations include hardware, software, and procedures to minimize security vulnerabilities.

The previous chapter explored how to set up your cloud infrastructure and strategies to make it secure.

This chapter covers the following topics:

- Securing application access
- Data classification
- Securing data access
- Data encryption patterns

Securing Application Access

In traditional application development and environments, securing data and software applications involves the on-premise network perimeter and physical access to the data. Considering the current trend where software developers are able to work from home using the Bring Your Own Device (BYOD) concept and mobile and cloud applications, most of the workload happens outside the company's network.

Identity is considered the new security boundary for the enterprise organizations. So, providing a granular level of access and enabling only valid users to access the system is the key to controlling your data and applications.

S. Lad, *Azure Security For Critical Workloads*, https://doi.org/10.1007/978-1-4842-8936-5_5

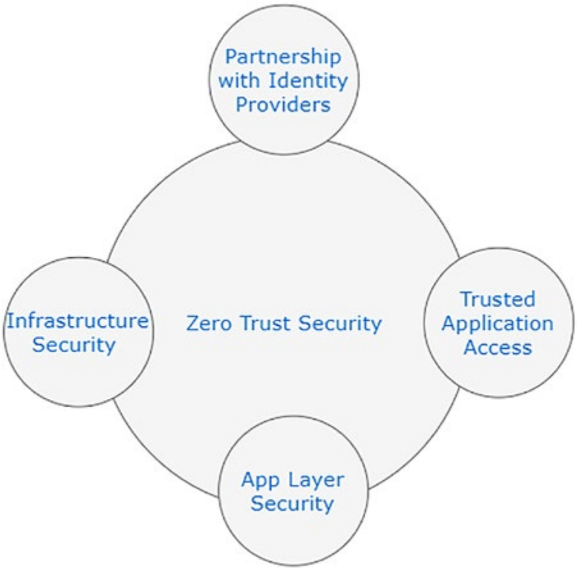
Partnership with Identity Providers
Infrastructure Security
Zero Trust Security
Trusted Application Access
App Layer Security

Figure 5-1. *Zero trust cloud security*

With zero trust cloud security (see Figure 5-1), cloud applications securely connect users to the applications. Authentication and authorization are required during the resource access in cloud infrastructure.

- The layered security approach with defense-in-depth: With layered security, you can easily detect attempts to access unauthorized access to the data. Since there are multiple layers, if one layer is breached, another layer is in place to prevent exposure.
- Regarding the PaaS and SaaS components of Microsoft Azure, Microsoft has a layered approach to security for physical data centers as well as across all Azure services. It protects information from intruders and prevents the data/information from being stolen.
- Confidentiality: Only authorized users can access the application. With the principle of least privilege, an entity or ID should be given only those privileges that are absolutely required to complete its tasks. Only individuals or non-personal entities who have explicit access can use the application. This includes securing user passwords, email content, or certificates.

- Integrity: In order to achieve integrity of an application, the goal is to prevent unauthorized changes to the data/information while the data is in rest or in transit. The standard approach while sharing data is that the sender creates unique fingerprint data using the hashing algorithm. The hash is then sent to the receiver along with the data. The receiver then checks the integrity of the data by calculating and comparing the hash to see if any data was lost or changed in transit.
- Availability: Cloud providers need to make sure that all services are available to authorized users. Denial of service attacks are the major reason to make applications unavailable to end users. Another major reason for the non-availability of the application is due to the natural disasters.

Now that you've read about the basics of application security, let's look at the security layers. You can think of layered security as a set of concentric rings, with the data being at the center, and it should be secured.

- Data: In a cloud environment, data normally resides as follows:
    1. Database
    2. Virtual machine disk
    3. SaaS application such as Microsoft 365
    4. Cloud storage
- Applications: You need to make sure that your applications are secure and free of vulnerabilities. Sensitive information should be stored in a secure storage mechanism. Integrating security into the application development lifecycle will lessen security vulnerabilities in the code. See Figure 5-2.

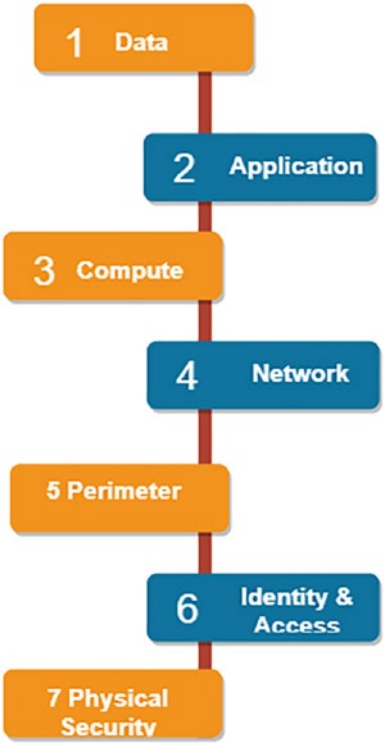

Figure 5-2. *Security layers*

- Compute: All compute resources should be kept secure by making sure that malware and patches are applied on time.
- Networking: Resource communication should be limited by using segmentation and access control. Inbound and outbound traffic should be denied by default and limited traffic should be open based on requirements.
- Perimeter: In order to avoid denial of service for end users, you can use the distributed denial of service (DDoS) protection. You can also use the perimeter firewalls to identify and alert about attacks against the network.
- Identity and access: The main goal of this security layer is to make sure that identities are secure and use a single sign-on and multifactor authentication for login and access management.

- Physical security: Physical security of the building and controlled access to the computer hardware in the data center is a top-most priority above all security layers.

Now that you have an understanding of application security access, the next section covers identity management in detail.

Identity Management

Digital identities are an integral part of enterprise organizations working on cloud or on-premise. Earlier identity and access services were restricted to operate only within the company's internal environment, so protocols like LDAP and Kerberos were designed and implemented.

Nowadays, mobile devices have become a primary way to interact with digital services. Organizations must evaluate the capabilities of the architecture in terms of its ability to bring such capabilities into the applications.

- Single sign-on: In a normal scenario, users have to manage multiple usernames and passwords to access applications/systems or services. More identities mean more usernames and passwords to remember. It becomes difficult for users to remember all that information. With single sign-on, users need to remember only one user Id and password. Access across the application is granted to one entity, which simplifies the security model. You can use Azure AD to enable single sign-on, which has a capability to combine multiple data sources into the security graph.
- Multifactor authentication (MFA): With multifactor authentication, you can add a layer of security to the application by requiring two or more elements for full authentication. Multifactor authentication enables security this way:
  - Something you know (for example, password)
  - Something you have (for example, security token)
  - Something you are (for example, biometrics such as fingerprints or face recognition)

- It increases the security of the identity by limiting the impact of confidential information being leaked. Azure AD has built-in multifactor authentication capability. Basic authentication features are available to Microsoft 365 and Azure AD administrators at no cost.

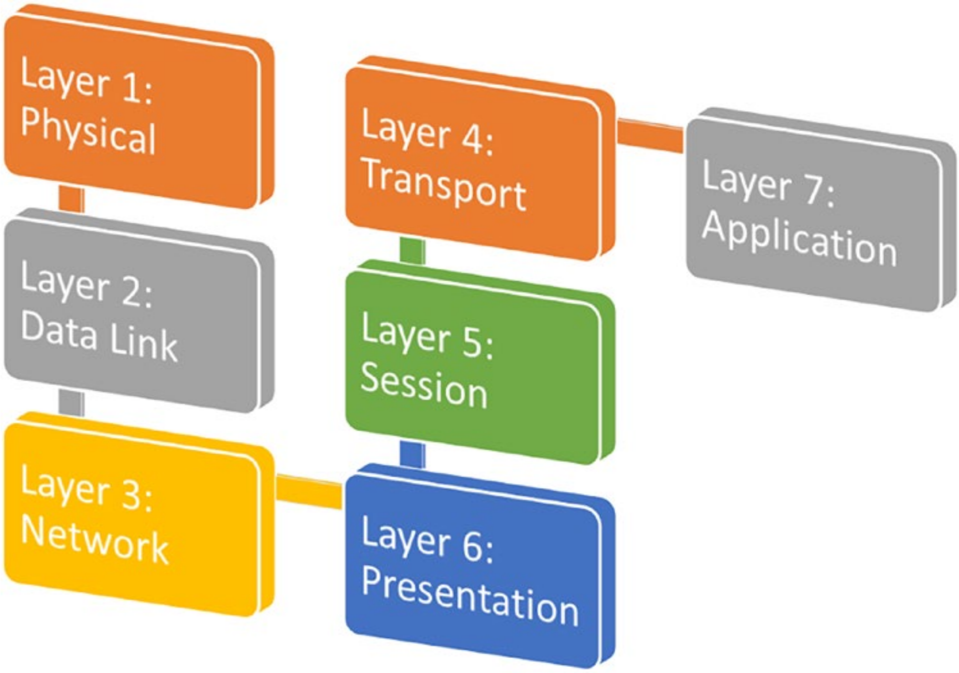

Figure 5-3. *Azure AD multifactor authentication*

The next section takes a quick tour of data classification.

Data Classification

Data classification sorts data to determine and assign value to it. It categorizes the data by sensitivity and business impact to identify the risks. Since the data is classified, you can manage the data to protect the sensitive information from loss. It is a process of linking metadata with every asset in a digital estate, which identifies the type of data.

The following are ways to classify data as per the Microsoft standard. Depending on your needs or security requirements, data classification standards may already exist in your organization. If no standard exists, you can use the following sample classification:

- Non-business: Personal data that doesn't belong to the organization
- Public: Data that is freely available and approved for public use
- General: Data that is available but not approved for public use
- Confidential: Data that can create issues if it is shared with other people
- Highly confidential: Business data that can create major issues for an organization if it is overshared

Microsoft Azure SQL DB has a built-in data discovery and classification feature that provides basic capabilities to discover, classify, label, and report sensitive data in an Azure SQL database. Sensitive data often consists of personal or financial information. Azure SQL DB can be considered a platform to fulfil the following data classification requirements:

- Achieve standards of data privacy and audit regulatory compliance
- Access sensitive data in a controlled manner and audit data access for security purposes
- Harden the security database that contains sensitive data

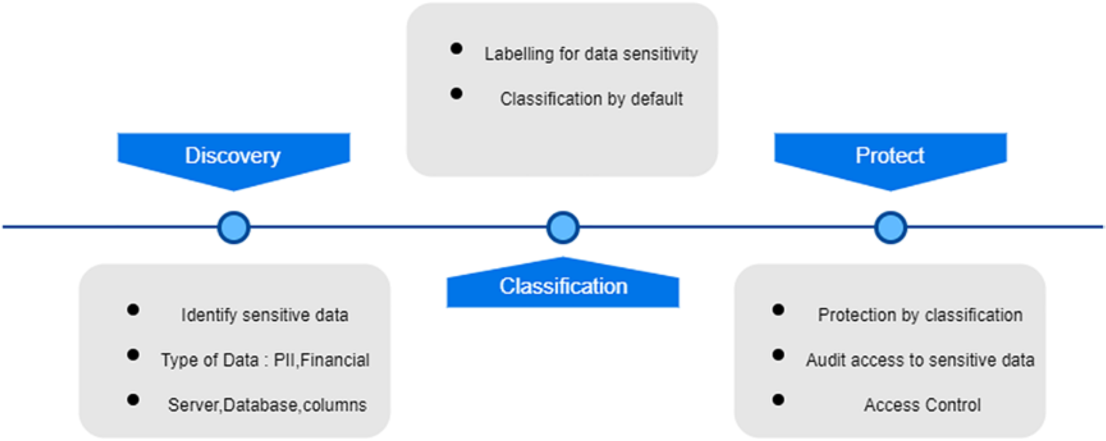
Labelling for data sensitivity
Classification by default
Discovery
Protect
Classification
Identify sensitive data
Type of Data : PII,Financial
Server,Database,columns
Protection by classification
Audit access to sensitive data
Access Control

Figure 5-4. *Data discovery and classification*

Data discovery and classification (see Figure 5-4) supports the following capabilities in Azure SQL Database:

- Discovery and recommendations: A classification engine scans the database and detects the columns in the database that might contain personal sensitive information. Based on the data, it provides a recommendation to review and apply classification using the Azure Portal. See Figure 5-5.

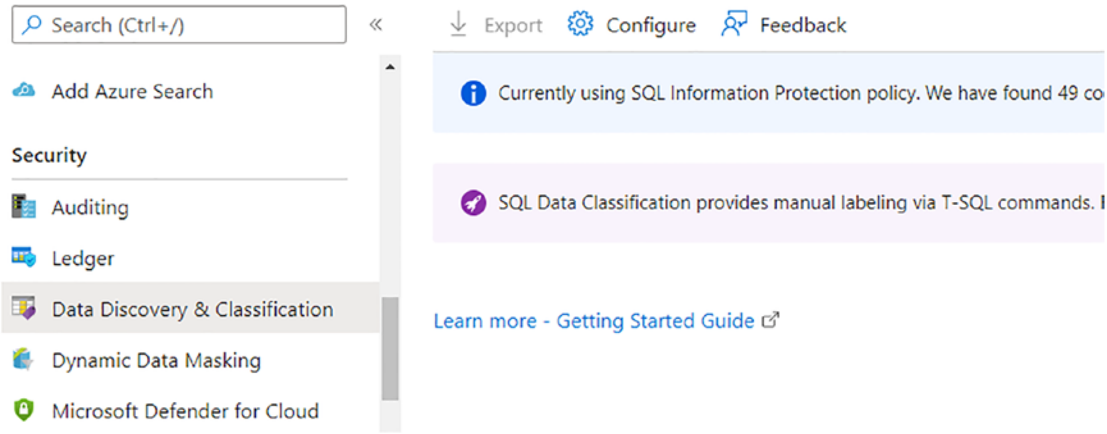

Currently using SQL Information Protection policy. We have found 49 co

SQL Data Classification provides manual labeling via T-SQL commands. I

Figure 5-5. *Data discovery and classification in Azure SQL DB*

- Labels: Using this feature, you can apply labels to the columns based on the metadata attribute available in the SQL database. You can take into account the metadata attributes in the SQL database based on the compliance and auditing requirements. See Figure 5-6.

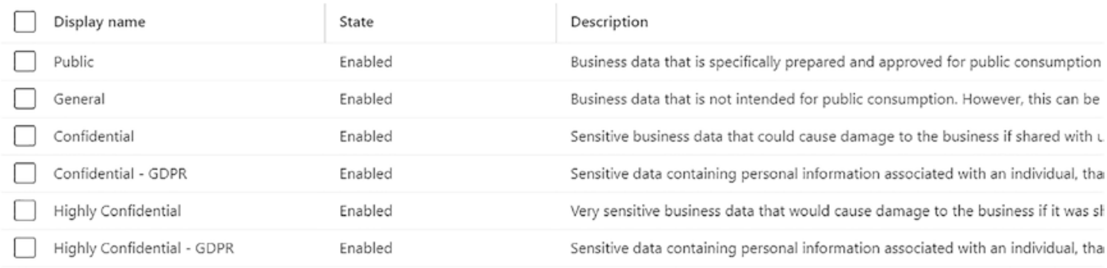

□ Display name	State	Description
□ Public	Enabled	Business data that is specifically prepared and approved for public consumption
□ General	Enabled	Business data that is not intended for public consumption. However, this can be
□ Confidential	Enabled	Sensitive business data that could cause damage to the business if shared with u
□ Confidential - GDPR	Enabled	Sensitive data containing personal information associated with an individual, tha
□ Highly Confidential	Enabled	Very sensitive business data that would cause damage to the business if it was sh
□ Highly Confidential - GDPR	Enabled	Sensitive data containing personal information associated with an individual, tha

Figure 5-6. *Labelling in Azure SQL DB*

- Query execution sensitivity: Once you execute the query, sensitivity of the query result is determined for auditing purposes.
- Dashboards: The SQL user can view the classification of the underlying data in visualized dashboards. You can also download reports in Excel format for auditing and governance purposes. See Figure 5-7.

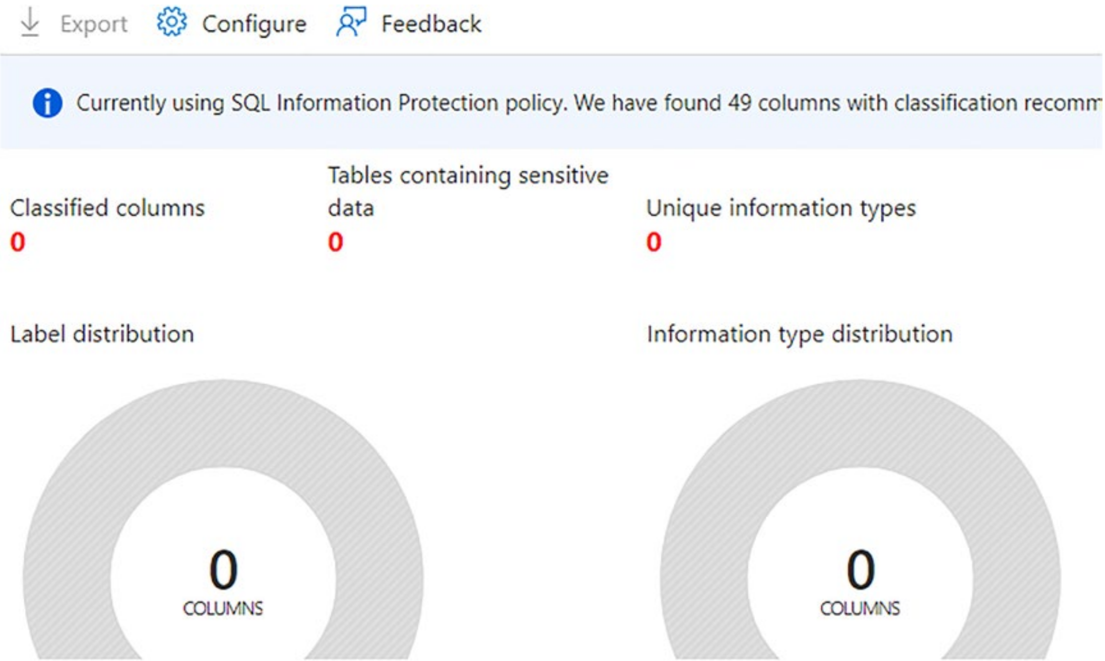

Currently using SQL Information Protection policy. We have found 49 columns with classification recomm

Classified columns 0

Tables containing sensitive data 0

Unique information types 0

Label distribution 0 COLUMNS

Information type distribution 0 COLUMNS

Figure 5-7. *Azure SQL DB data classification dashboards*

- Azure SQL offers the SQL Information Protection Policy as well as the Microsoft Information Protection policy to classify data. You can use either of them to discover and classify your data.
- SQL Information Protection Policy: Data discovery and classification has an built-in feature of sensitivity labels and information types that have discovery logic. You can define and customize the classification inside the central place using the Azure organizations. This central place is the Microsoft defender from the security policy. Only administrators of the management group can execute this activity. See Figure 5-8.

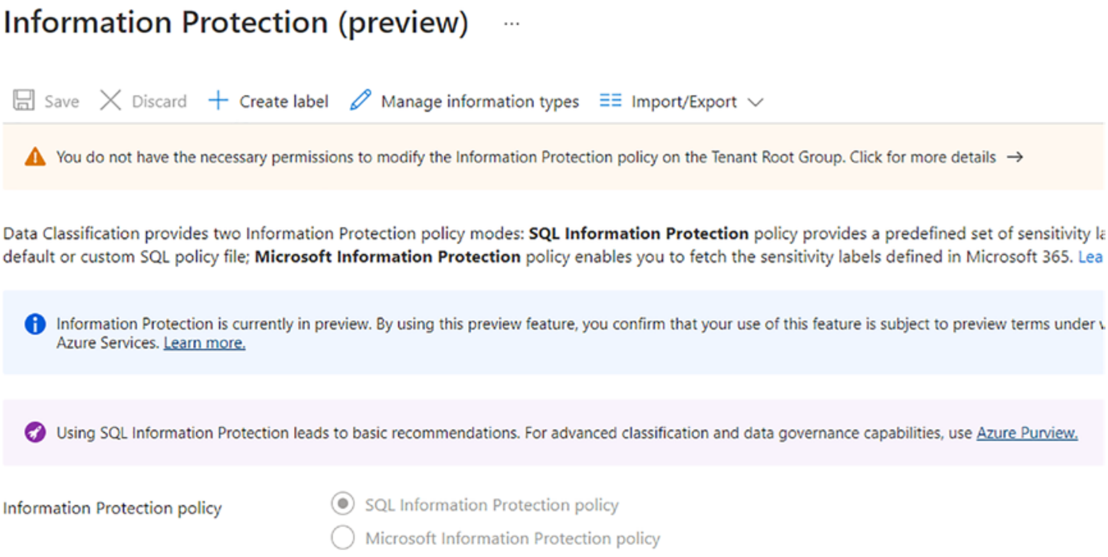

You do not have the necessary permissions to modify the Information Protection policy on the Tenant Root Group. Click for more details →

Data Classification provides two Information Protection policy modes: **SQL Information Protection** policy provides a predefined set of sensitivity la
default or custom SQL policy file; **Microsoft Information Protection** policy enables you to fetch the sensitivity labels defined in Microsoft 365. Lea

Information Protection is currently in preview. By using this preview feature, you confirm that your use of this feature is subject to preview terms under w
Azure Services. Learn more.

Using SQL Information Protection leads to basic recommendations. For advanced classification and data governance capabilities, use Azure Purview.

Information Protection policy
● SQL Information Protection policy
○ Microsoft Information Protection policy

Figure 5-8. *Azure SQL DB Information Protection*

During policy management, you can define custom labels, rank them, and associate them with information types. You can add custom information types and configure them with string patterns. See Figure 5-9.

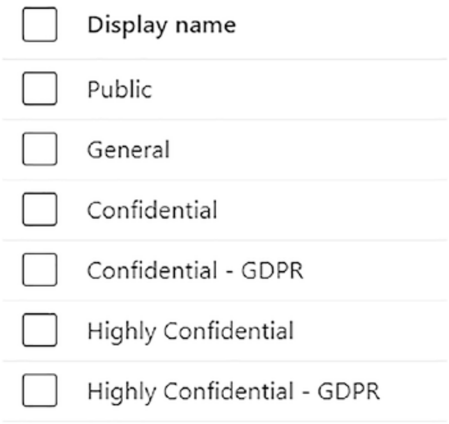

□ Display name
□ Public
□ General
□ Confidential
□ Confidential - GDPR
□ Highly Confidential
□ Highly Confidential - GDPR

Figure 5-9. *Azure SQL DB Information Protection Policy Labels*

If you want to classify your database using the SQL information protection policy mode, follow these steps:

1. First go to the Azure Portal using the link: `https://portal.azure.com/`
2. Go to the Azure database and open the Data Discovery and Classification option from the Security section. The Overview tab has a summary of classification from the SQL database.
3. It also has an option to download the report in Excel. Click the Export button in the top pane to do this. See Figure 5-10.

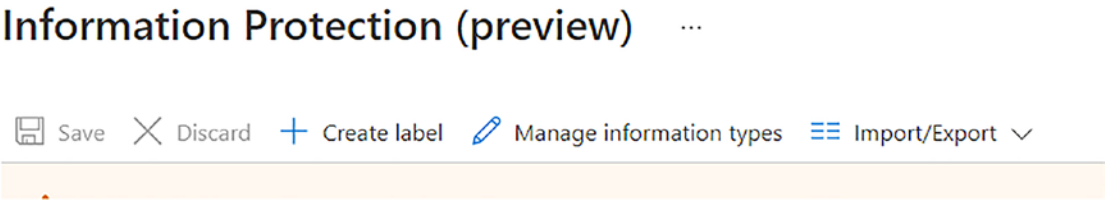

Figure 5-10. *Azure SQL DB Data Discovery an Classification: Export option*

4. In order to classify your data, select the Classification tab from the Data Discovery and Classification page. This classification engine scans the data and identifies any columns that contain personal and sensitive information.
5. Next, apply the classification recommendations as follows. First select the recommendation panel from the bottom pane.

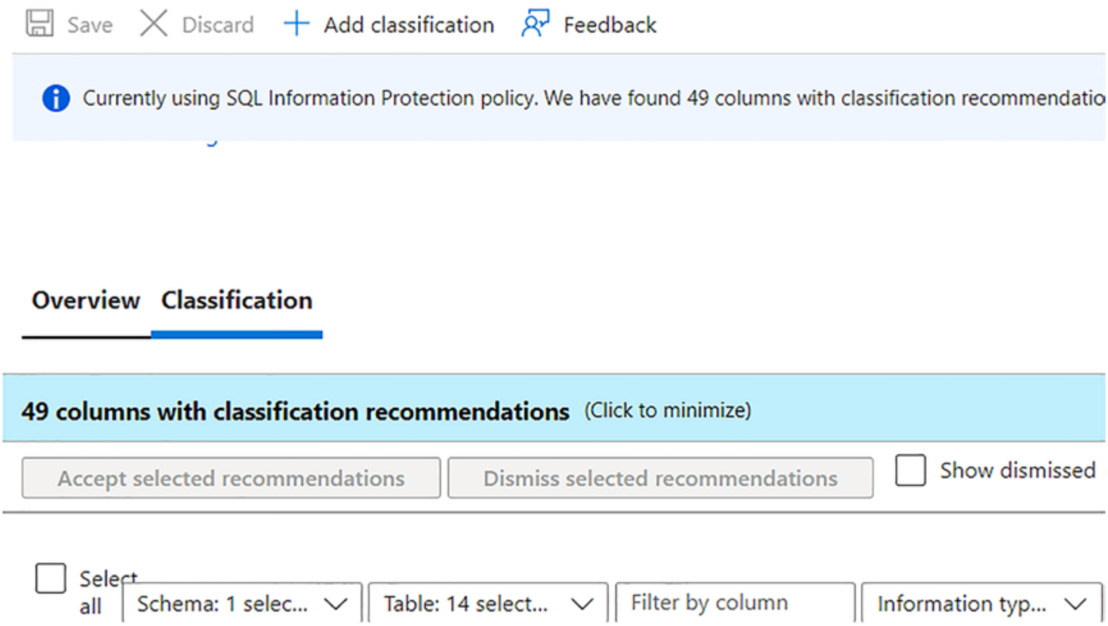

Currently using SQL Information Protection policy. We have found 49 columns with classification recommendatio

49 columns with classification recommendations (Click to minimize)

Figure 5-11. *Azure SQL DB classification recommendations*

You can accept the recommendations for a specific column or ignore them based on your choices. In order to apply the recommendations, choose the Accept Selected Recommendations option.

On the other hand, Microsoft Information Policy labels provide a simple way for end users to classify sensitive data across various Microsoft applications. These sensitivity labels are created and maintained in the Microsoft 365 compliance center.

Let's now go through the process of securing data access.

Securing Data Access

Public cloud services offer various in-built features to secure access to the data stored in cloud storage services. Configuring these features enables enterprise organizations to secure and protect sensitive information. However, many companies fail to protect their data due to misconfigurations or lack of awareness. Securing data is not only the responsibility of cloud providers but it is also the responsibility of enterprise organizations to properly configure security settings. See Figure 5-12.

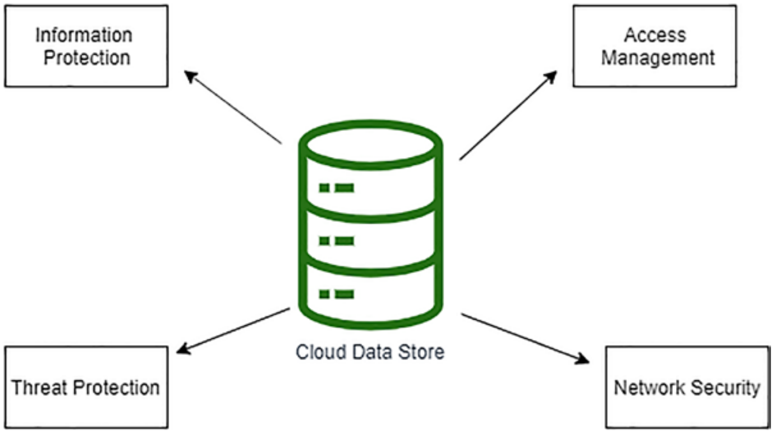

Figure 5-12. *Cloud data security*

Major challenges faced by enterprise organizations to secure and monitor data usage are as follows:

- Centrally monitoring security events in the logs
- Encrypting data at rest and in transit
- Guaranteeing authenticated and authorized access to data
- Handling centralized identity management for securely accessing stored data

Data Protection

The first step to protecting information is to identify and determine which data should be protected and where it resides. There are three types of data: data at rest, data in transit, and data that's in use (see Figure 5-13). Once you identify the location, the next step is to create clear, simple guidelines to secure the data. Once the data is available in the cloud storage, you need to protect the data at rest and in transit.

- Data at rest: You need to protect the data that exists statically or physically inside the cloud storage.
- Data in transit: When data is being transferred to components or locations over the network or across the service bus, it must be secured.

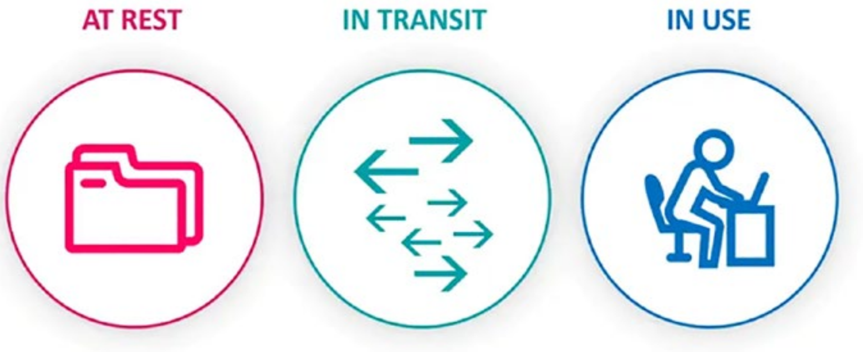

Figure 5-13. *Data at rest or in transit*

Access Control

Securing and administrating data stored in the cloud is a combination of identity access management and control.

Here is the communication flow for single sign-on (SSO) implementation.

- Centralized identity management. Centralized identity management enables much easier authentication and authorization to manage enterprise applications.

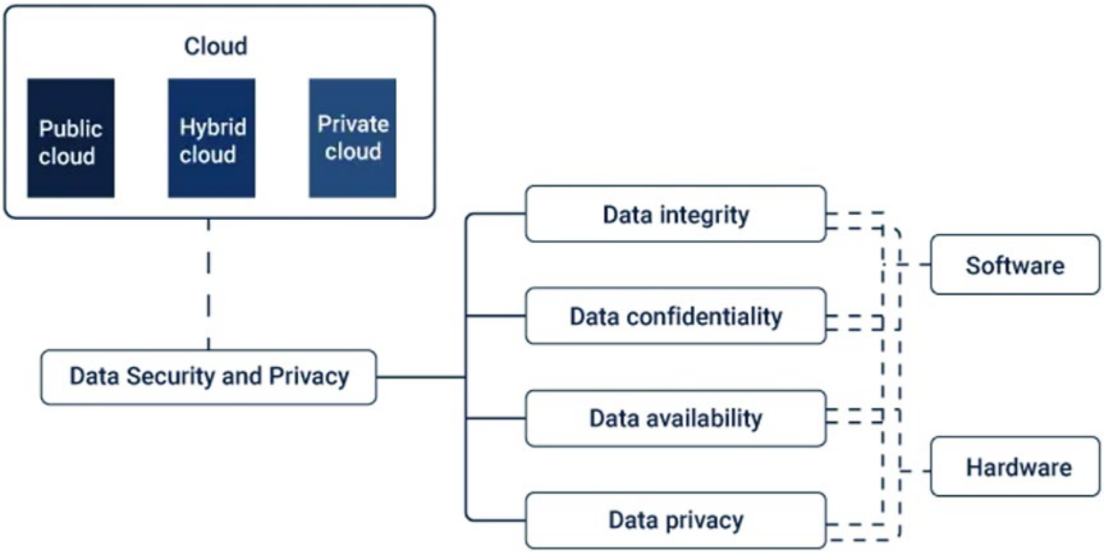

Figure 5-14. *Identity and Access Management (IAM) in Azure*

- Password management
- Multifactor Authentication(MFA): Multifactor authentication is an electronic authentication method in which a user can access the application or data only after providing two or more pieces of information—what the user knows, such as a password, what the user has, such as a token, and what the user is, using biometric verification.
- Role-based access control (RBAC): Role-based access control is a method for controlling what users can do within a company's IT application. RBAC enables access control by assigning roles to end users. See Figure 5-15.

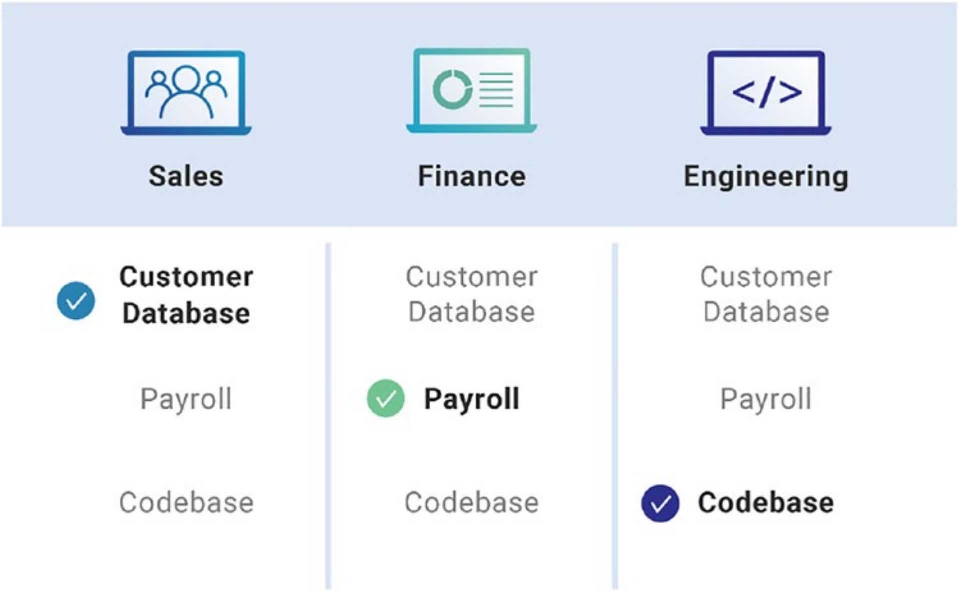

Figure 5-15. *Role-based access control*

- Conditional access policies
- Monitor suspicious activities

Another important consideration for protecting data is to select the right key management solution. Azure Key Vault (AKV) safeguards keys, certificates, and secrets that can be used by cloud applications and services. Azure Key Vault is designed to store keys and secrets. It is not intended to store the passwords.

Best practices for using Azure Key Vault to store data are as follows:

- Grant access to the users, groups, or application for specific purposes.
- Azure has many in-built roles as well and you can create custom roles as per the requirements. You can assign those predefined roles to users or groups as needed.
- Control users who have access to the data or application.
- Access to Key Vault is managed through two separate interfaces: management plane and data plane. For example, if you want to grant data plane access permissions using the Key Vault access policies then no management plane access is needed for the application.
- Store authentication certificates in the Key Vault: Azure Resource Manager(ARM) can securely deploy the certificates stored inside the Key Vault to Azure VMs when the VMs are deployed. You can set up required access policies for the Key Vault to grant access to the certificate.

You can use the Azure disk encryption in a virtual machine to encrypt the attached disks on the Windows or Linux VMs. Azure Storage uses Azure storage encryption to encrypt data at rest in the Azure Storage. Encryption, decryption, and key management are transparent to end users. Azure SQL databases and Azure Synapse Analytics use Transparent Data Encryption (TDE) to execute the real-time encryption and decryption of the database, related backup, and transaction log files without requiring any changes to the application. SQL databases also have a feature called Always Encrypted to protect sensitive data at rest and on the server. This prevents Database Administrators (DBAs), cloud database operators, and other high privileged non-authorized users from accessing encrypted data directly from the server as well.

Network Security for Data Access

In order to protect data in transit, you can use Secure Socket Layer (SSL)/Transport Layer Security (TLS) certificates while exchanging data across locations. You can isolate the communication channel between the on-premises and cloud infrastructure using the Virtual Private Network (VPN) or Express Route (ER).

Using the network security groups (NSGs), you can reduce the number of potential attacks, as the network security groups contain a list of security rules that allow or deny inbound or outbound network traffic based on destination address, source ports, destination ports, or protocol. VMs within the two Azure virtual networks can easily talk to each other using VNET peering. Network traffic between the peered virtual networks is private.

Monitoring

Microsoft Defender for Cloud automatically collects, analyzes, and integrates logs from the Azure resources, networks, and solutions such as firewalls. See Figure 5-16.

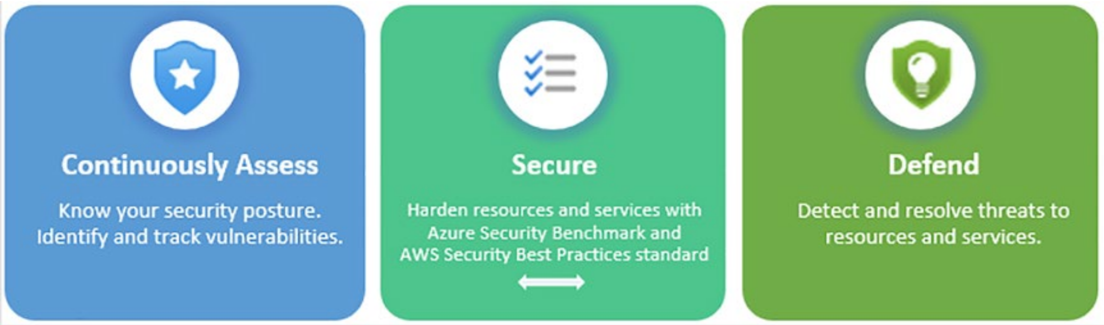
Continuously Assess
Know your security posture.
Identify and track vulnerabilities.
Secure
Harden resources and services with
Azure Security Benchmark and
AWS Security Best Practices standard
Defend
Detect and resolve threats to
resources and services.

Figure 5-16. *Microsoft Defender for Cloud*

Log analytics provide centralized access to the logs and help analyze the log data to create custom alerts for proactive monitoring. It is the primary tool for editing the log queries and interactively analyzing the results. From the portal, you can use and interactively execute queries and sort, filter, and analyze the logs. See Figure 5-17.

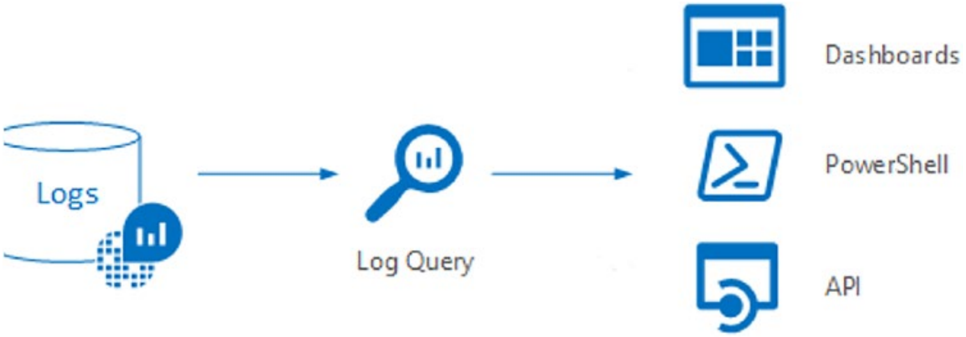

Figure 5-17. *Log analytics workspace*

Azure SQL Database and threat protection detects abnormal activities and identifies attempts to get unauthorized access to the data, database, or system. Security officers and administrators get immediate notifications about abnormal activities so they can take required actions in a timely manner. See Figure 5-18.

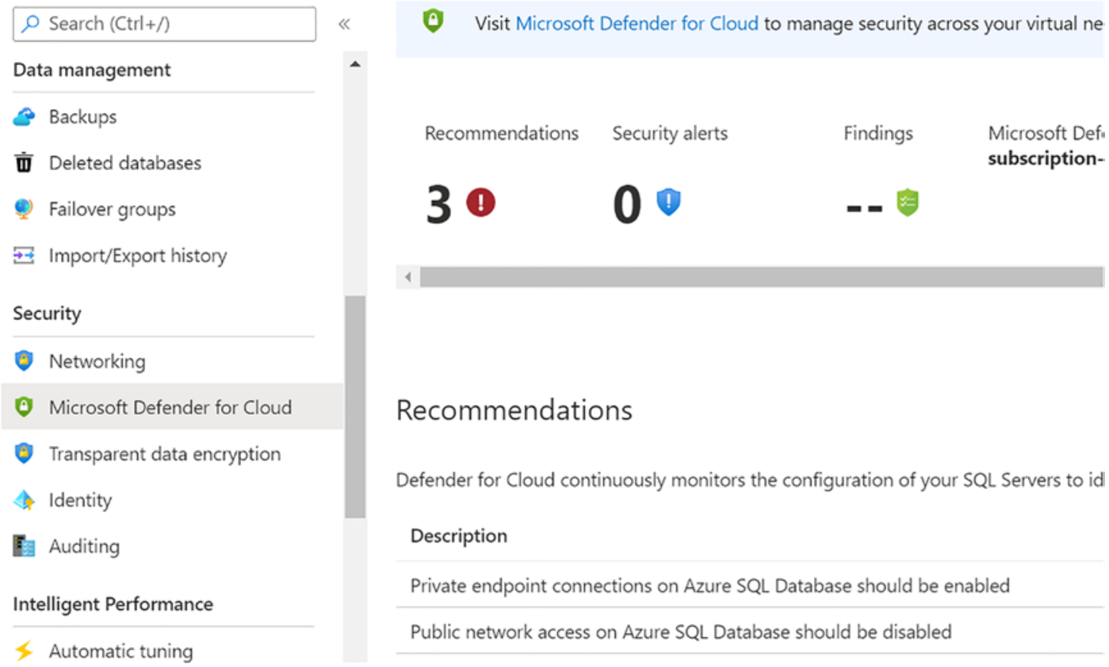

Visit Microsoft Defender for Cloud to manage security across your virtual ne

Recommendations

Defender for Cloud continuously monitors the configuration of your SQL Servers to id

Description
Private endpoint connections on Azure SQL Database should be enabled
Public network access on Azure SQL Database should be disabled

Figure 5-18. *Microsoft Defender for Cloud*

Advanced Threat Protection (ATP) for Azure SQL Database can identify potential SQL injections, access attempts from unusual location or data centers, and brute force attacks on the SQL database. Once Advanced Threat Protection is configured, it will automatically detect these threats via email notifications or via the Azure Portal. Advanced Threat Protection is part of Microsoft Defender for SQL. It can be accessed from the SQL Server Portal or from the central Microsoft Defender portal. See Figure 5-19.

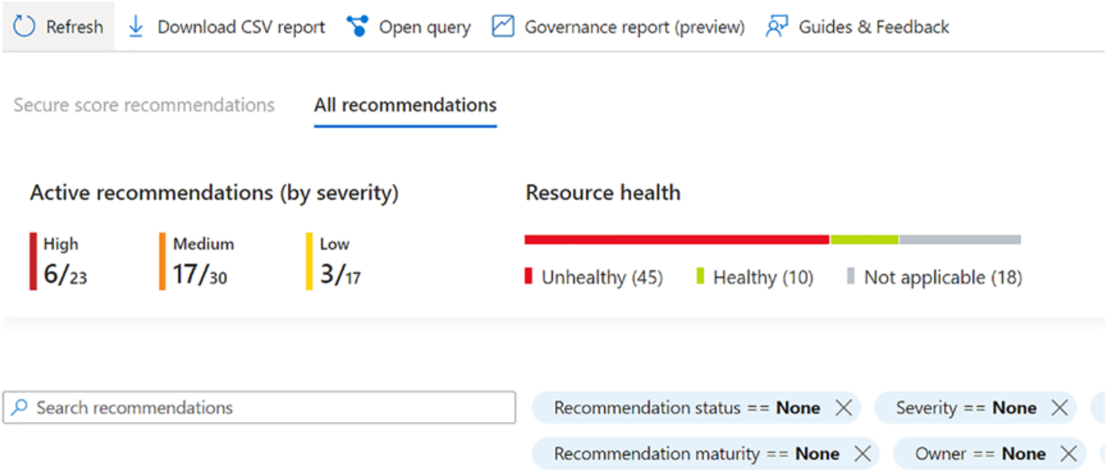

Recommendation status == **None**

Severity == **None**

Recommendation maturity == **None**

Owner == **None**

Figure 5-19. *Advanced Threat Protection*

Data Encryption Patterns

When data is in the cloud, it resides inside the cloud storage, in memory, or on the network during the transit. With a cloud solution, a single transaction can lead to multiple data operations whereby the data moves from one storage medium to another. In order to provide full data protection and security, the data must be encrypted on the storage volumes. See Figure 5-20.

Figure 5-20. *Data encryption*

A few major points to consider with data encryption:

- Use identity-based storage access controls
- Encrypt the virtual disks for the virtual machines
- Use secure hash algorithms for data encryption
- Protect data in transit by using encrypted network channels like HTTPS or TLS for client-server communication
- Use additional key encryption key (KEK) to protect the data encryption key (DEK)

Microsoft Azure has built-in data encryption features in many layers and also participates in data processing. Microsoft recommends you enable the encryption capability for all Microsoft services to protect and secure your data.

Cloud storage is very well architected and implemented differently than traditional on-premises solutions. It enables massive scaling and modern access via the REST API and isolation between the tenants. All Azure storage services have various built-in data encryption features to protect and secure data (see Table 5-1).

- Identity based access control: You can enable access to the storage service using the Azure Active Directory and key-based authentications, such as storage account access key and shared access signature (SAS).
- Built-in storage encryption: All cloud storage data is by default encrypted. Data can't be read by the tenant if it hasn't been written by the tenant. With this feature, you can make sure that data isn't leaked.
- Region-based controls: Data only remains in the selected region and three copies of the data is maintained in other regions. Azure storage has detailed activity logging available based on the configuration.
- Firewall features: Azure Firewall provides an additional layer of access control and storage threat protection to detect abnormal activities related to access.

Table 5-1. *Azure's Key Management Parameters*

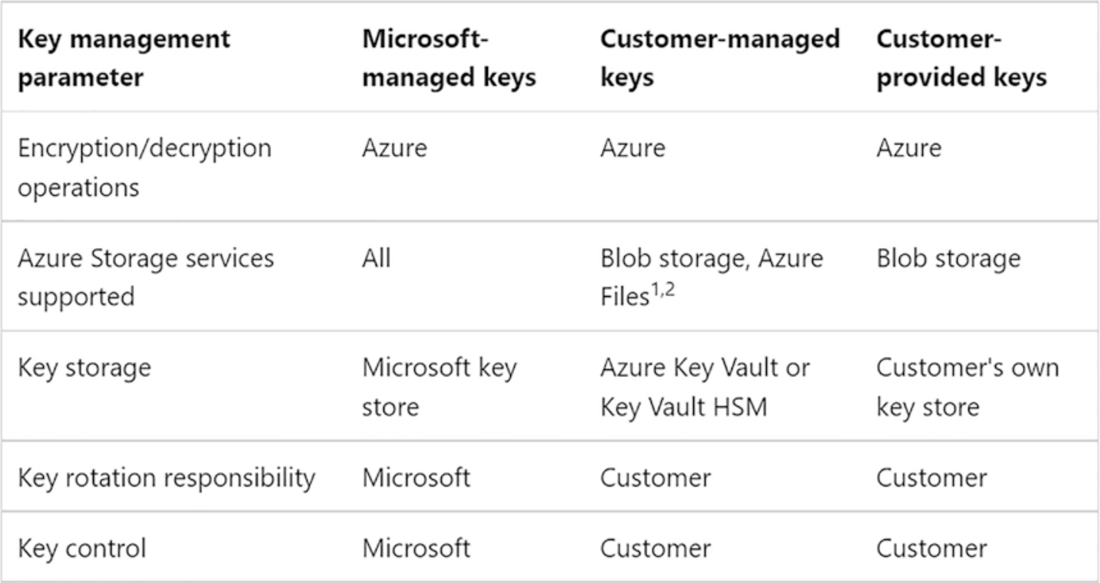

Key management parameter	**Microsoft-managed keys**	**Customer-managed keys**	**Customer-provided keys**
Encryption/decryption operations	Azure	Azure	Azure
Azure Storage services supported	All	Blob storage, Azure Files[1,2]	Blob storage
Key storage	Microsoft key store	Azure Key Vault or Key Vault HSM	Customer's own key store
Key rotation responsibility	Microsoft	Customer	Customer
Key control	Microsoft	Customer	Customer

In order to better understand how Microsoft Azure implements encryption at rest, you need to understand the various encryption models. These definitions are shared across all resource providers to ensure a common language:

- Server-side encryption using service-managed keys: In this method of encryption, encryption is mainly performed by the Azure service, which is basically done by the Microsoft Azure cloud resource provider. Consider an example where Azure Storage receives the data in plain text format and encryption and decryption are performed automatically by the cloud service providers when you write or read data to the Azure storage account. Resource providers can use their own encryption key or they can use a custom encryption key, depending on the storage encryption configuration. See Figure 5-21.

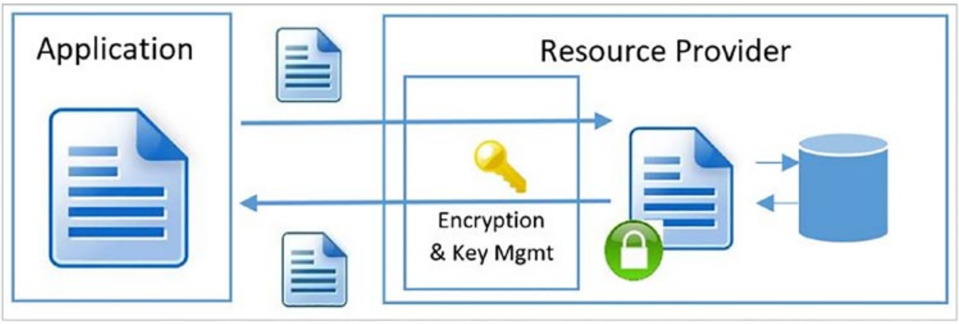

Figure 5-21. *Server-side encryption*

- Client encryption model: This encryption model is performed outside Azure by the service or client application. With this setup, Azure resource providers can encrypt the blob of data, but they can't decrypt the data or access the encryption keys. See Figure 5-22.

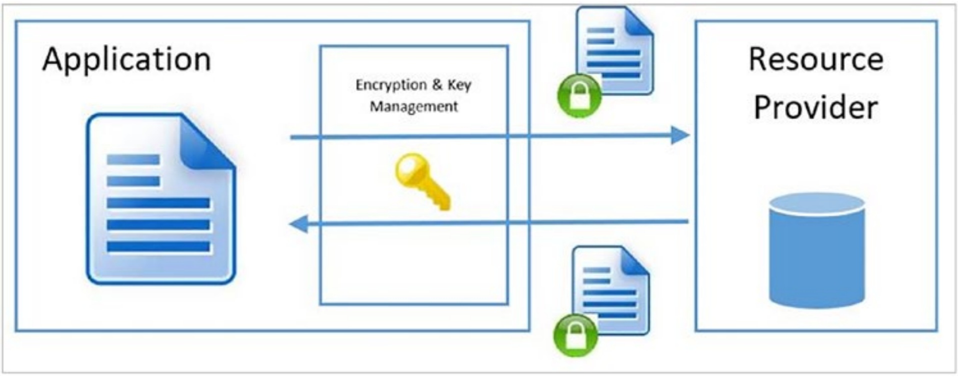

Figure 5-22. Client-side encryption

Some Azure services store the root key encryption key (KEK) in the Azure Key Vault and store the encrypted data encryption key (DEK) in the internal location where the data resides.

Conclusion

This chapter discussed securing the data stored in the cloud. You also learned about the various ways to classify the data and make it available for downstream users and applications in a secure manner. You also learned about the various data encryption patterns and related models when working with public cloud providers like Azure, Google, AWS, and so on.