Cosmos DB

Prashant Kumar Mishra
Engineering Architect

Instructions:

Please keep your microphones on mute during the entire session

Type in your questions in chat window, it will be answered towards the end of our session

Cosmos DB

Welcome

Introduction
Section 1

- Welcome to Azure Cosmos DB
- Global Data Distribution
- Consistency Level
- Choose the right consistency level
- High Availability
- Regional Presence
- Partitioning
- Provisioned Throughput

Security
Section 2

- Security Overview
- Data Encryption
- Secure Access to Data
- IP Firewall
- Access from Virtual Network
- Advanced Threat Protection

How-To Guide
Section 3

- Create Cosmos DB resources
- Create and Manage Data
- Migrate Data
- Create a Notebook
- Distribute Data Globally

Introduction

Welcome to Azure Cosmos DB

Introduction
Section 1

Welcome to Azure Cosmos DB

Global Data Distribution

Consistency Level

Choose the right consistency level

High Availability

Regional Presence

Partitioning

Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

Azure Cosmos DB is Microsoft's globally distributed, multi-model database service for operational and analytics workloads. It offers multi-mastering feature by automatically scaling throughput, compute, and storage.

Key benefits

- Turnkey global distribution
- Always On
- Elastic scalability of throughput and storage, worldwide
- Guaranteed low latency at 99th percentile, worldwide
- Precisely defined, multiple consistency choices
- No schema or index management

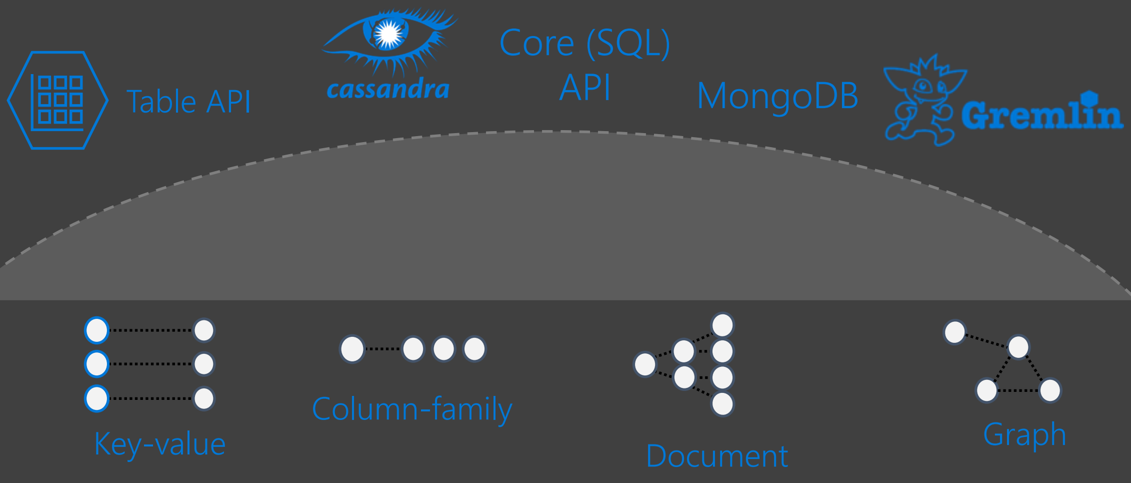

Introduction

Global Data Distribution

Introduction
Section 1

Welcome to Azure Cosmos DB

Global Data Distribution

Consistency Level

Choose the right consistency level

High Availability

Regional Presence

Partitioning

Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

Cosmos DB transparently replicates the data to all the regions associated with your Cosmos account. It provides a single system image of your globally distributed Azure Cosmos database and containers that your application can read and write to locally.
With Azure Cosmos DB, you can add or remove the regions associated with your account at any time. Your application doesn't need to be paused or redeployed to add or remove a region.

Benefits

- Unlimited elastic write and read scalability.
- 99.999% read and write availability all around the world.
- Guaranteed reads and writes served in less than 10 milliseconds at the 99th percentile.

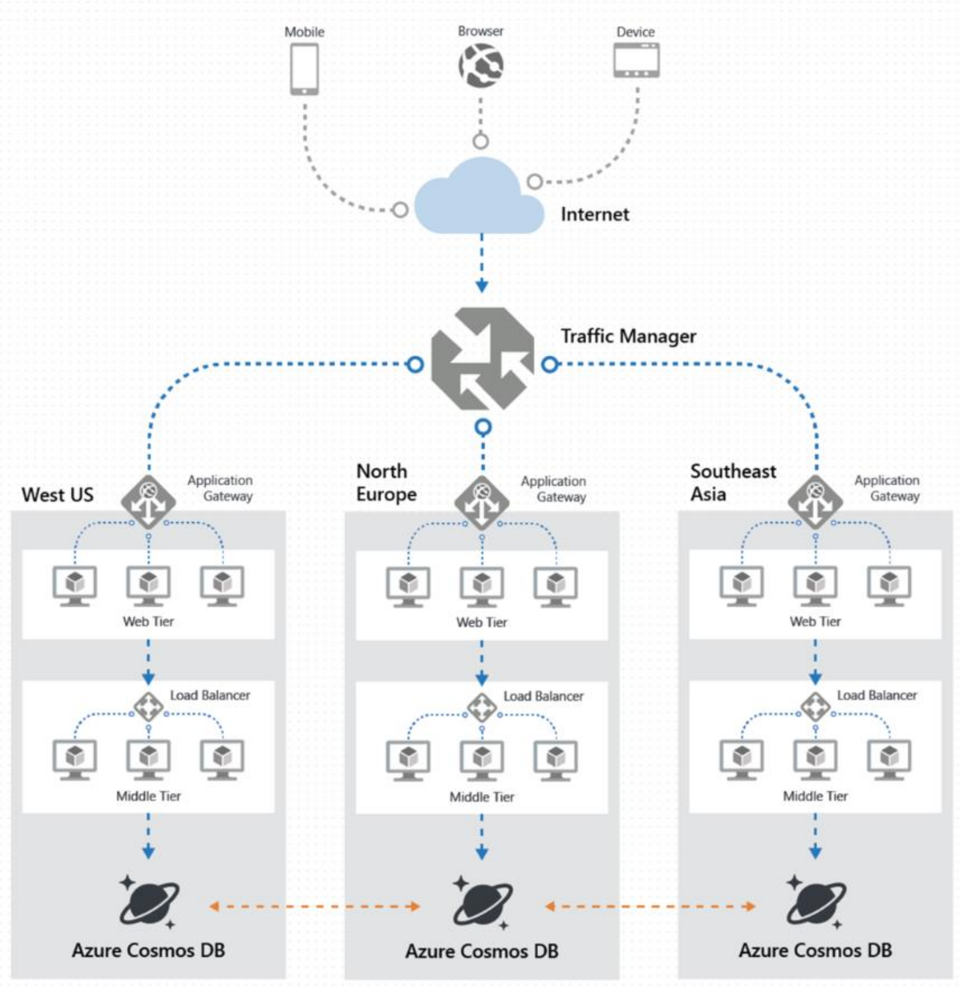

Introduction
Section 1

Welcome to Azure Cosmos DB
Global Data Distribution
Consistency Level
Choose the right consistency level
High Availability
Regional Presence
Partitioning
Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

Azure Cosmos DB approaches data consistency as a spectrum of choices instead of two extremes. Developers can use these options to make precise choices and granular tradeoffs with respect to high availability and performance.

You can configure the default consistency level on your Azure Cosmos account at any time. The default consistency level configured on your account applies to all Azure Cosmos databases and containers under that account. All reads and queries issued against a container or a database use the specified consistency level by default.

- Strong
- Bounded Staleness
- Session
- Consistent Prefix
- Eventual

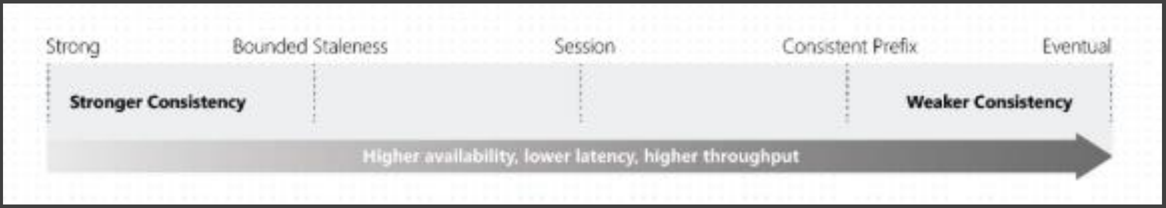

Introduction

Choose the right Consistency Level

Introduction
Section 1

Welcome to Azure Cosmos DB

Global Data Distribution

Consistency Level

Choose the right consistency level

High Availability

Regional Presence

Partitioning

Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

- When the consistency level is set to bounded staleness, Cosmos DB guarantees that the clients always read the value of a previous write, with a lag bounded by the staleness window.
- When the consistency level is set to strong, the staleness window is equivalent to zero, and the clients are guaranteed to read the latest committed value of the write operation.
- For the remaining three consistency levels, the staleness window is largely dependent on your workload. For example, if there are no write operations on the database, a read operation with eventual, session, or consistent prefix consistency levels is likely to yield the same results as a read operation with strong consistency level.

Note: If you need even higher data durability without sacrificing performance, you can create a custom consistency level at the application layer. For more information see, How-to implement custom synchronization in your applications.

Introduction

High Availability

Introduction
Section 1

Welcome to Azure Cosmos DB

Global Data Distribution

Consistency Level

Choose the right consistency level

High Availability

Regional Presence

Partitioning

Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

Azure Cosmos DB transparently replicates your data across all the Azure regions associated with your Cosmos account

- The data within Cosmos containers is horizontally partitioned.
- Within each region, every partition is protected by a replica-set with all writes replicated and durably committed by most replicas. Replicas are distributed across as many as 10-20 fault domains.
- Each partition across all the regions is replicated. Each region contains all the data partitions of a Cosmos container and can accept writes and serve reads.

If your Cosmos account is distributed across N Azure regions, there will be at least N x 4 copies of all your data.

In addition to cross region resiliency, you can now enable zone redundancy when selecting a region to associate with your Azure Cosmos database.

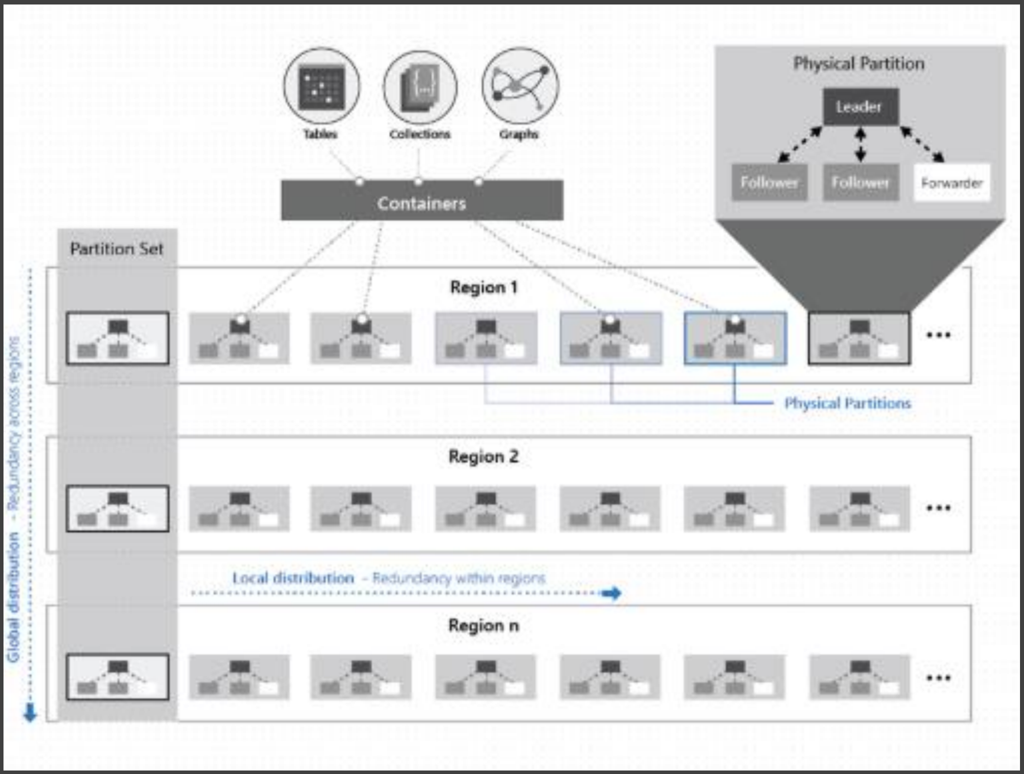

Introduction
Section 1

Welcome to Azure Cosmos DB
Global Data Distribution
Consistency Level
Choose the right consistency level
High Availability
Regional Presence
Partitioning
Provisioned Throughput

Security
Section 2

How-To Guide
Section 3

Azure Cosmos DB is a foundational service in Azure, and, by default, is always available in all regions, where Azure is available. Currently, Azure is available in 54 regions worldwide.

Cosmos DB is available in all five distinct Azure cloud environments available to customers:

- **Azure public cloud**, which is available globally.
- **Azure China** 21Vianet is available through a unique partnership between Microsoft and 21Vianet, one of the country's largest internet providers in China.
- **Azure Germany** provides services under a data trustee model, which ensures that customer data remains in Germany under the control of T-Systems International GmbH, a subsidiary of Deutsche Telecom, acting as the German data trustee.
- **Azure Government** is available in four regions in the United States to US government agencies and their partners.
- **Azure Government for Department of Defense (DoD)** is available in two regions in the United States to the US Department of Defense.

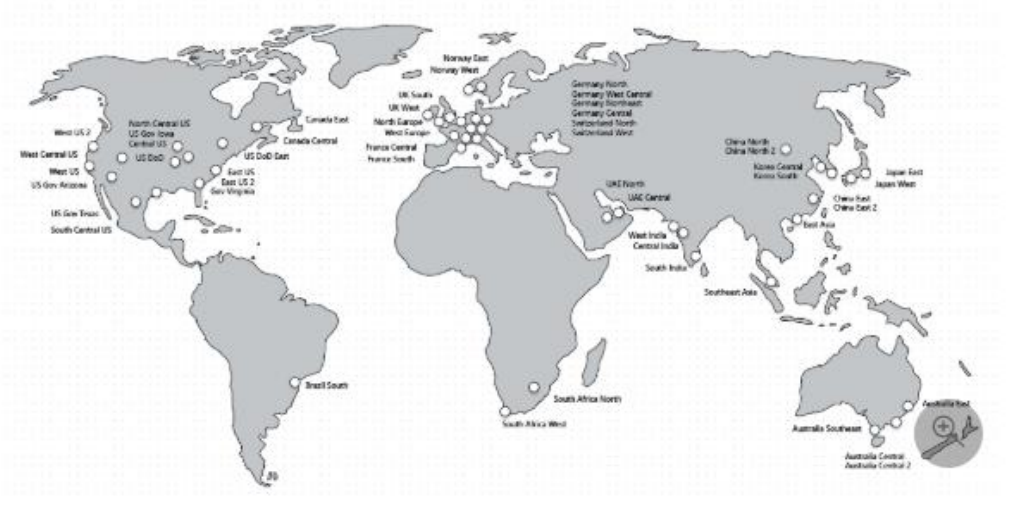

Azure Cosmos DB uses partitioning to scale individual containers in a database to meet the performance needs of your application. In partitioning, the items in a container are divided into distinct subsets called logical partitions. Logical partitions are formed based on the value of a partition key that is associated with each item in a container. All items in a logical partition have the same partition key value.

For all containers, your partition key should:

- Be a property that has a value which does not change. If a property is your partition key, you can't update that property's value.
- Have a high cardinality. In other words, the property should have a wide range of possible values.
- Spread request unit (RU) consumption and data storage evenly across all logical partitions. This ensures even RU consumption and storage distribution across your physical partitions.

There is no limit to the number of logical partitions in your container. Each logical partition can store up to 20GB of data.

The number of physical partitions in your Cosmos container depends on the following:

- Amount of provisioned throughput (each individual physical partition can provide a throughput of up to 10,000 request units per second)
- Total data storage (each individual physical partition can store up to 50GB)

Request Units per second (RU/s) : Cost of a request in terms of CPU, memory and I/O

With Azure Cosmos DB, you can provision throughput at two granularities:

- Azure Cosmos containers
- Azure Cosmos databases

Note: In general, container level throughput is a good choice. This leads to predictable performance since each container is guaranteed its provisioned RU's

Latest update: Auto scale to scale based on usage. You need to define the maximum throughput, but you will be charged for 10% of the maximum throughput defined

Identify number of Rus

1 RU = 1 Read of 1 KB document
10 RU= 1 read of 100 KB document
5 RU = 1 write of 1 KB document
50 RU=1 write of a 100 KB document

Note: Indexing also affects RU cost

Introduction
Section 1

Security
Section 2

Security Overview

Data Encryption

Secure Access to Data

IP Firewall

Access from Virtual Network

Advanced Threat Protection

How-To Guide
Section 3

Encryption at rest is now available for documents and backups stored in Azure Cosmos DB in all Azure regions. Encryption at rest is applied automatically for both new and existing customers in these regions.

Network Security: Using an IP firewall is the first layer of protection to secure your database

Authorization: Azure Cosmos DB uses hash-based message authentication code (HMAC) for authorization

Users and Permissions: Using the master key for the account, you can create user resources and permission resources per database.

Active Directory Integration (RBAC): You can also provide or restrict access to the Cosmos account, database, container, and offers (throughput) using Access control (IAM) in the Azure portal.

Global Replication: global replication ensures data protection against regional failures.

HTTPS/SSL/TLS encryption: All connections to Azure Cosmos DB support HTTPS. Azure Cosmos DB also supports TLS 1.2.

Security and data protection certifications: SOCS 1/2 Type 2, HITRUST, PCI DSS Level 1, ISO 27001, HIPAA, FedRAMP High, and many others.

Security

Data Encryption

Introduction
Section 1

Security
Section 2

Security Overview
Data Encryption
Secure Access to Data
IP Firewall
Access from Virtual Network
Advanced Threat Protection

How-To Guide
Section 3

Encryption at rest is implemented by using several security technologies, including secure key storage systems, encrypted networks, and cryptographic APIs. Systems that decrypt and process data must communicate with systems that manage keys.

There is no impact or changes to the performance SLAs now that encryption at rest is enabled for all existing and new accounts.

There is no additional cost if Storage Service Encryption is enabled

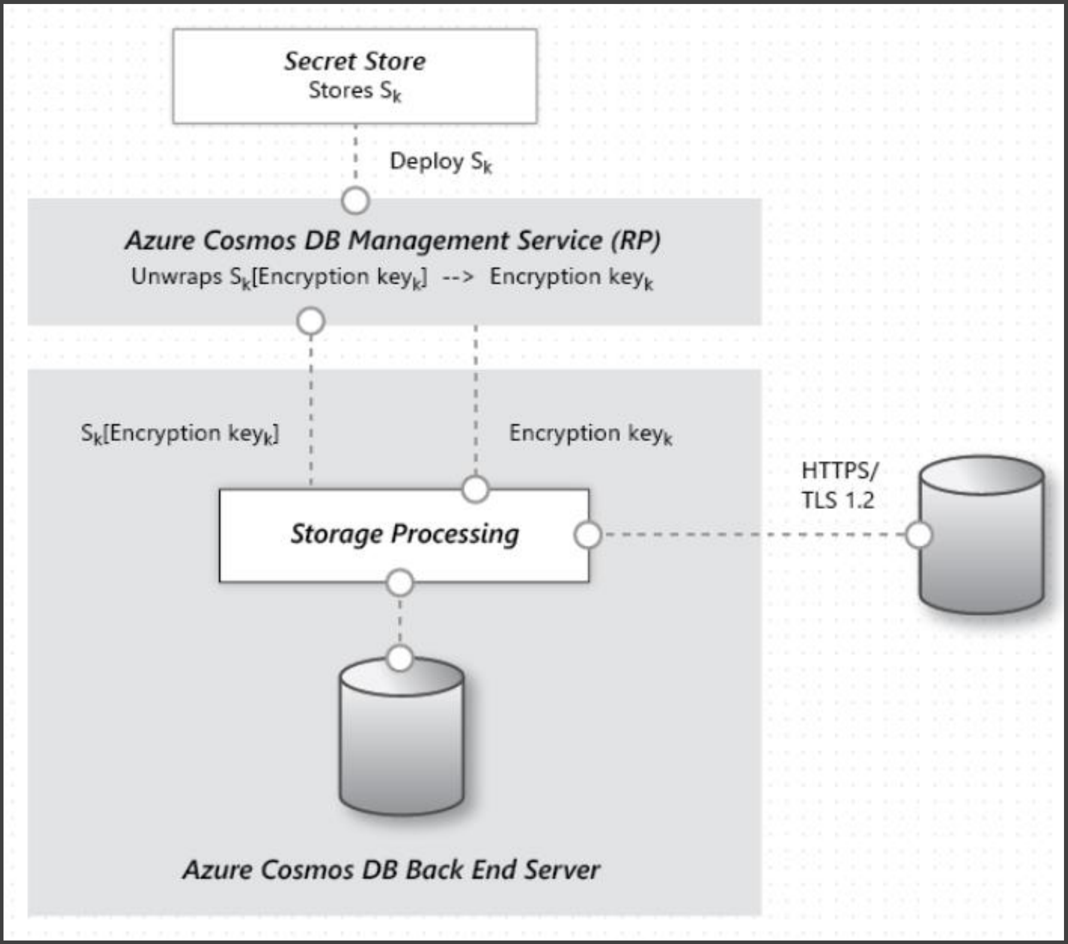
Secret Store
Stores S_k
Deploy S_k
Azure Cosmos DB Management Service (RP)
Unwraps S_k[Encryption key_k] --> Encryption key_k
S_k[Encryption key_k]
Encryption key_k
HTTPS/
TLS 1.2
Storage Processing
Azure Cosmos DB Back End Server

Security

Secure Access to Data

Introduction
Section 1

Security
Section 2

Security Overview

Data Encryption

Secure Access to Data

IP Firewall

Access from Virtual Network

Advanced Threat Protection

How-To Guide
Section 3

Azure Cosmos DB uses two types of keys to authenticate users and provide access to its data and resources.

- **Master keys:**
  - Provide access to accounts, databases, users, and permissions.
  - Cannot be used to provide granular access to containers and documents.
  - Are created during the creation of an account.
  - Can be regenerated at any time.

- **Resource tokens:**
  - Provide access to specific containers, partition keys, documents, attachments, stored procedures, triggers, and UDFs.
  - Are created when a user is granted permissions to a specific resource.
  - Are recreated when a permission resource is acted upon on by POST, GET, or PUT call.
  - Use a hash resource token specifically constructed for the user, resource, and permission.
  - Are time bound with a customizable validity period. The default valid time span is one hour. Token lifetime, however, may be explicitly specified, up to a maximum of five hours.
  - Provide a safe alternative to giving out the master key.
  - Enable clients to read, write, and delete resources in the Cosmos DB account according to the permissions they've been granted.

Security

IP Firewall

Azure Cosmos DB supports IP-based access controls for inbound firewall support. This model is like the firewall rules of a traditional database system and provides an additional level of security to your account.

With firewalls, you can configure your Azure Cosmos account to be accessible only from an approved set of machines and/or cloud services.

To configure IP policy-based access control, the user must provide the set of IP addresses or IP address ranges in CIDR (Classless Inter-Domain Routing) form to be included as the allowed list of client IPs to access a given Azure Cosmos account. Once this configuration is applied, any requests originating from machines outside this allowed list receive 403 (Forbidden) response.

Note that firewall changes may take up to 15min to propagate.

You can combine IP-based firewall with subnet and VNET access control.

Introduction
Section 1

Security
Section 2

Security Overview
Data Encryption
Secure Access to Data
IP Firewall
Access from Virtual Network
Advanced Threat Protection

How-To Guide
Section 3

You can configure the Azure Cosmos account to allow access only from a specific subnet of virtual network (VNet). By enabling Service endpoint to access Azure Cosmos DB on the subnet within a virtual network, the traffic from that subnet is sent to Azure Cosmos DB with the identity of the subnet and Virtual Network. Once the Azure Cosmos DB service endpoint is enabled, you can limit access to the subnet by adding it to your Azure Cosmos account.

By default, an Azure Cosmos account is accessible from any source if the request is accompanied by a valid authorization token. When you add one or more subnets within VNets, only requests originating from those subnets will get a valid response. Requests originating from any other source will receive a 403 (Forbidden) response.

Security

Advanced Threat Protection

Introduction
Section 1

Security
Section 2

Security Overview
Data Encryption
Secure Access to Data
IP Firewall
Access from Virtual Network
Advanced Threat Protection

How-To Guide
Section 3

Advanced Threat Protection for Azure Cosmos DB detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. It can currently trigger the following alerts:

- **Access from unusual locations:** This alert is triggered when there is a change in the access pattern to an Azure Cosmos account, where someone has connected to the Azure Cosmos DB endpoint from an unusual geographical location. In some cases, the alert detects a legitimate action, meaning a new application or developer's maintenance operation. In other cases, the alert detects a malicious action from a former employee, external attacker, etc.

- **Unusual data extraction:** This alert is triggered when a client is extracting an unusual amount of data from an Azure Cosmos DB account. This can be the symptom of some data exfiltration performed to transfer all the data stored in the account to an external data store.

Security

Advanced Threat Protection

- Create Cosmos DB resources
- Migrate Data
- Create a Notebook
- Distribute Data Globally