Microsoft

Meet your presenter

Sumit Sengupta

Cloud Solution Architect

Focus:

NoSQL, AI, Security

Fun-Fact:

Gardener, Volunteer Teacher, Amateur Theater

ocp | One Commercial Partner

Azure Cosmos DB

Microsoft's globally-distributed database service

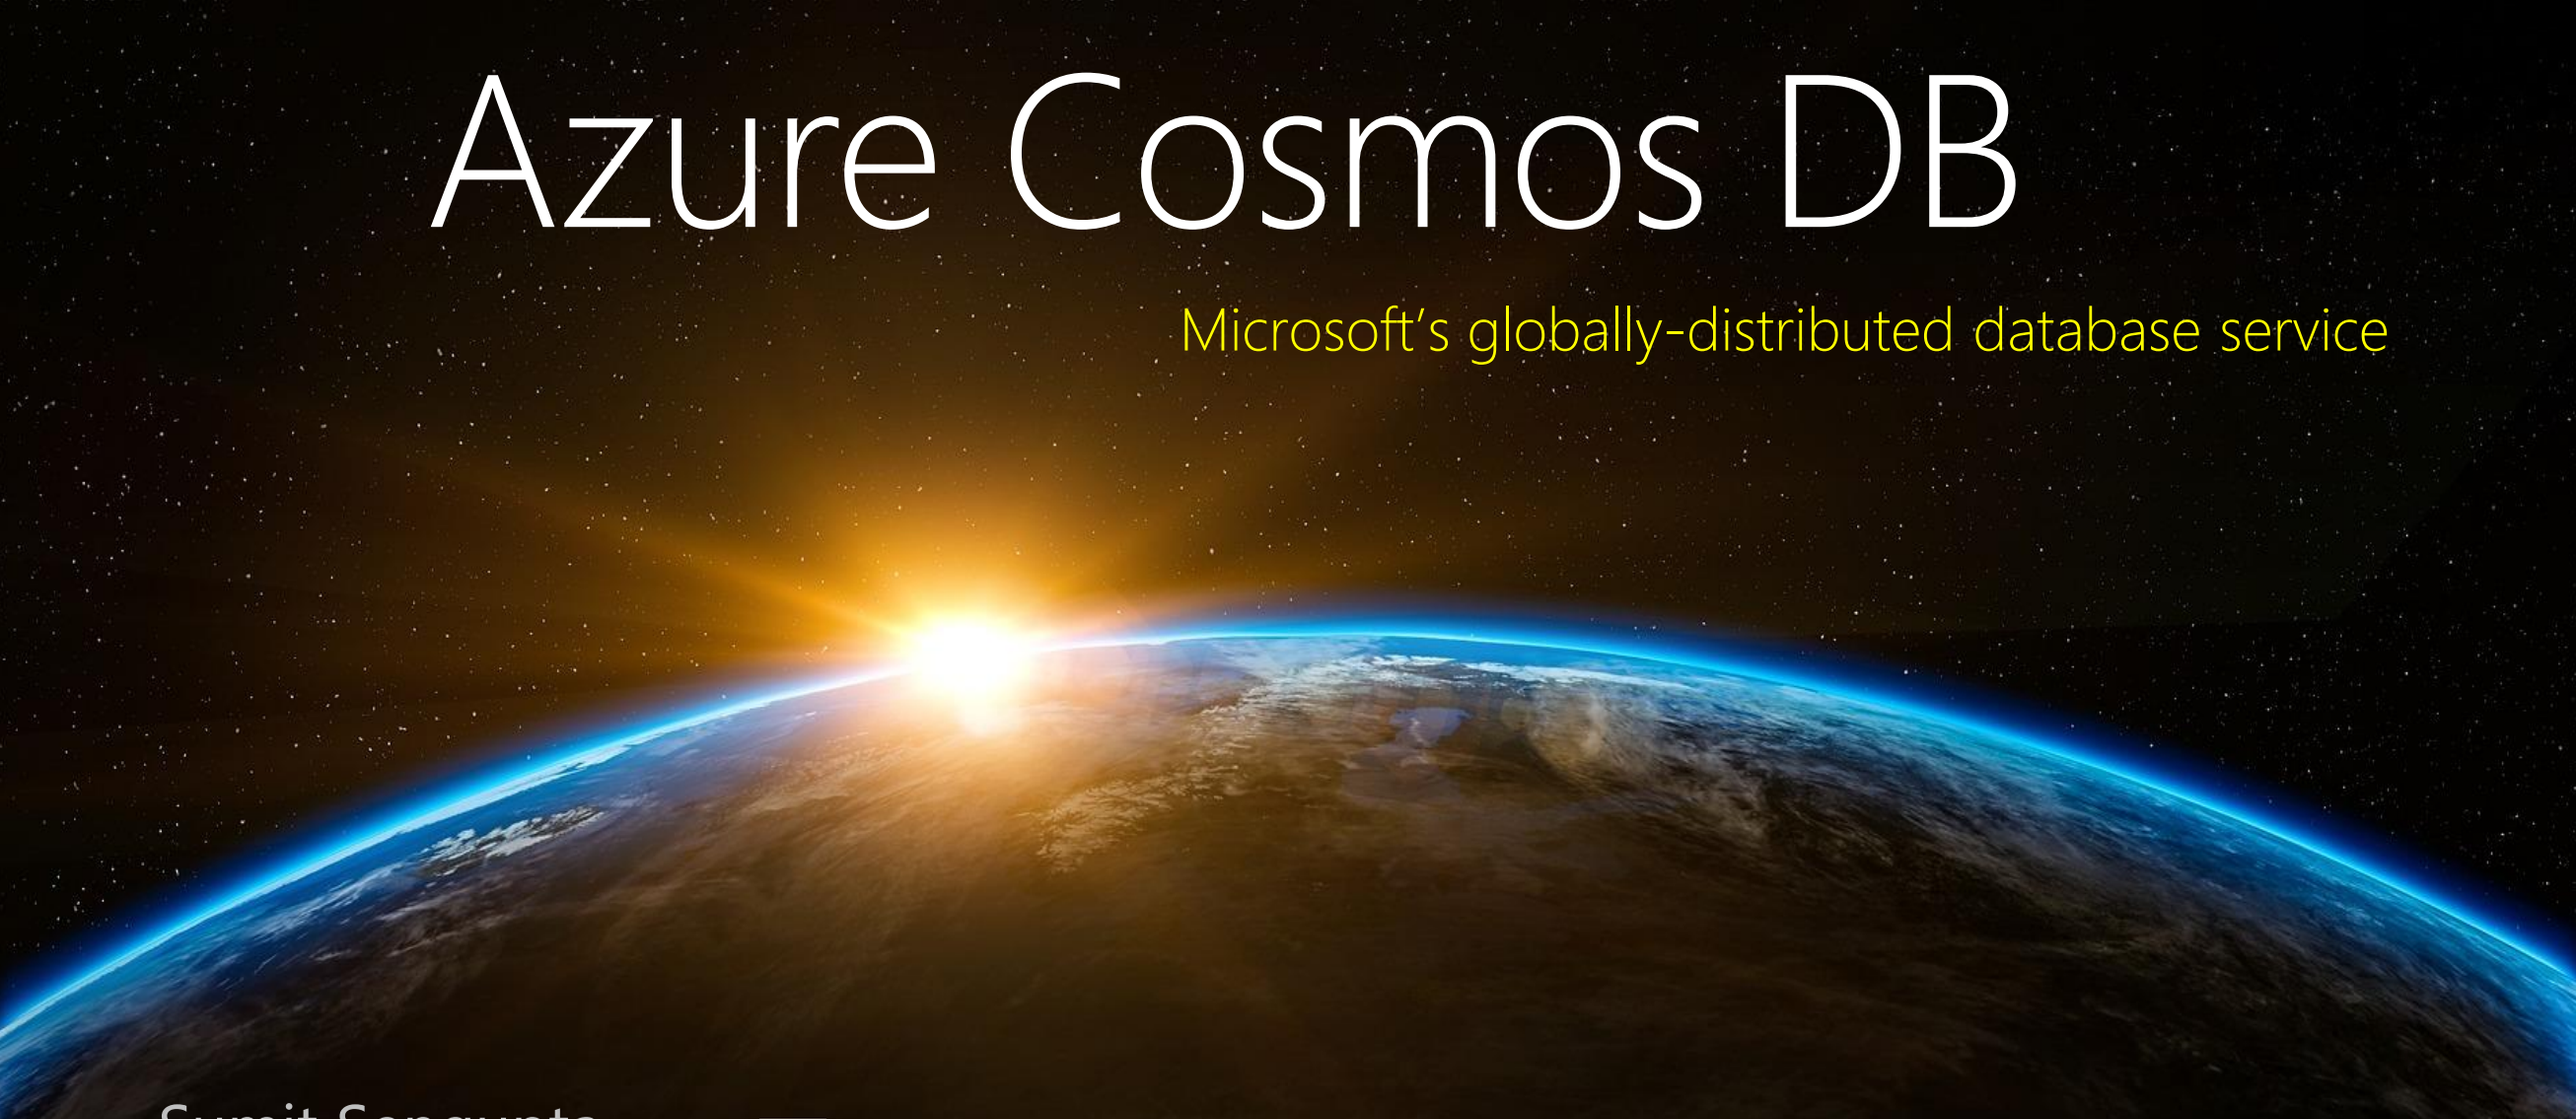

Sumit Sengupta

Cloud Solution Architect

US One Commercial Partner

https://ahttps://twitter.com/ssengupta/aka.ms/ss_twitter

sumit.sengupta@microsoft.com

https://aka.ms/ss_linkedin

https://twitter.com/ssengupta

Azure Cosmos DB

Blog: https://blogs.technet.microsoft.com/msuspartner/2018/04/05/azure-cosmos-db/

This is What a Typical Activity of an App Looks Like

Azure Cosmos DB

A globally distributed, massively scalable, multi-model database service

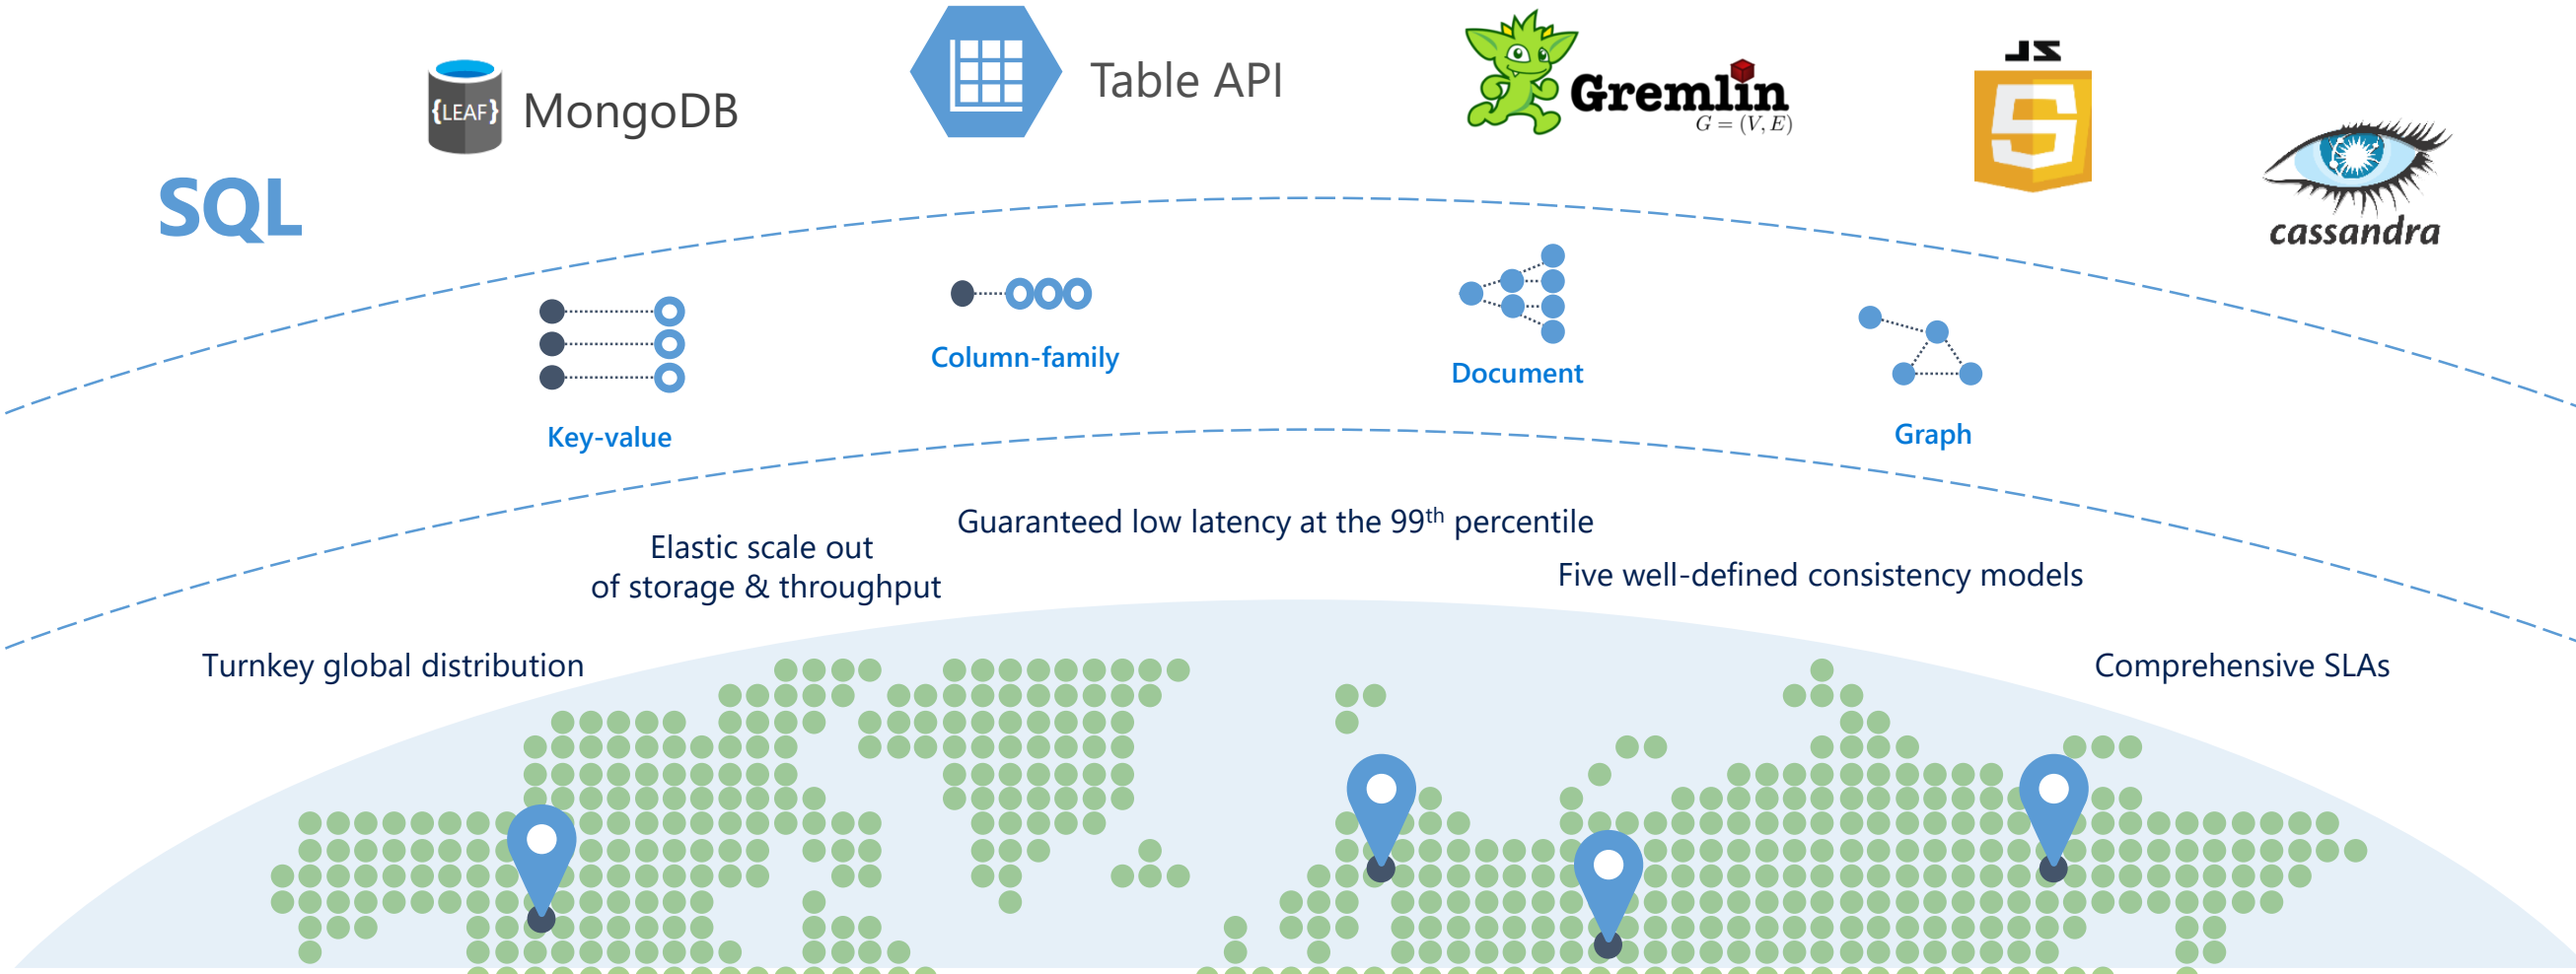

From validation to momentum

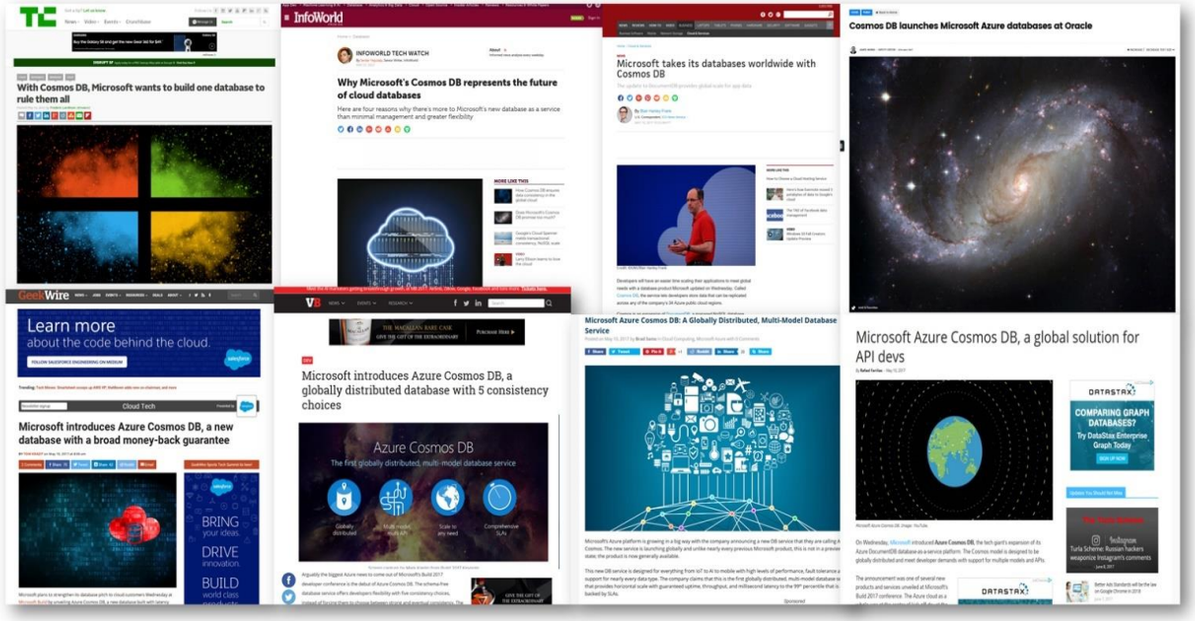

DB-ENGINES

Overall

31 +23

2018 March

1 Year

Key Value	Wide Column	Document	Graph
4 +5	3 +1	5 +4	2 +2

Cameron Vetter @poshporcupine · May 10
I would have never guessed that Microsoft was going to disrupt the database market today, **amazing**, **cosmosdb** #Build2017

Microsoft Business follows

Kincho Guerrero @kincho_guerrero · May 10
Azure **#CosmosDB** multi-model, multi-api, multi-region DB in #Azure. Wow a real **gamechanger** for the economics of cloud storage!👌👌👌#MSBuild

marco m blasio @marco_m_blasio · May 11
#Azure **#CosmosDB** eventually a true **game changer**. Brilliant choice of 5 Consistency Models #Database #Databasedesign

Azure Cosmos DB: The industry's first globally-dist...
Today, we're excited to announce the general availability of Azure Cosmos DB.
azure.microsoft.com

Dave Walden @DBADaveKC · Jun 5
The more I read about **#CosmosDB**, the more I **love** it!

> **The New Stack** @thenewstack
> Cosmos DB: Microsoft Azure's All-in-One Distributed Database Service
> bit.ly/2qPiapt

50/50 Pledge and 9 others follow

SAGE FRANCH @theTrendyTechie · May 17
When they replicated the database with just one click at **#MSBuild**, all our jaws dropped! **#CosmosDB** is super cool, here's why.

Cortana Intelligence and 3 others follow

Rami Sarieddine @RamieSays · May 10
#Azure **#CosmosDB** the answer to all your database needs! A **game changer**!! #MSBuild

Sean Werick @seanwerick · May 11
WOW. **#CosmosDB** is an **amazing** technology.

Microsoft Business follows

Adam Jamal @ONI_Assassin · May 10
OMG, **#CosmosDB** is **#amazing** and quite an achievement (Achievement Unlocked), well done @Microsoft Can't wait to get my hands on it. #MSBuild

Microsoft Business follows

Bogdan Grozoiu @bogdangr · May 20
Microsoft Build 2017 highlights. The **Cosmos** DB is **outstanding**!

Yasser Asmi @yasserasmi · May 11
Microsoft's Azure **Cosmos DB** looks **incredible**! #Build2017

Comparison

Capabilities	Azure Cosmos DB	AWS DynamoDB	Google Cloud Spanner
Global distribution	A single entity can span any number of Azure regions (30+).	No	No, currently single region. Multi-region support is coming later in 2017.
Horizontal scalability of both storage and throughput worldwide	Yes, across any number of Azure regions	No	No
Multi-homing APIs	Yes	No	No
Guaranteed low read/write latency worldwide	Yes, at 99th percentile worldwide	No	No
Consistency models	5 (*Strong, Eventual, Bounded Staleness, Session, Consistent-Prefix*)	2 (*Strong, Eventual*)	1 (*Strong*)
Multi-model + multi-API	Yes	Yes	No
Schema management	Schema agnostic; no management required	Index management required	Schema and index management required
Automatic indexing	Yes	No	No
Comprehensive SLAs	Yes (latency, consistency, throughput, high availability)	No	No (availability only)

Azure Cosmos DB

Deeply Exploits Cloud Core Properties and Economies of Scale

Lowest Cost

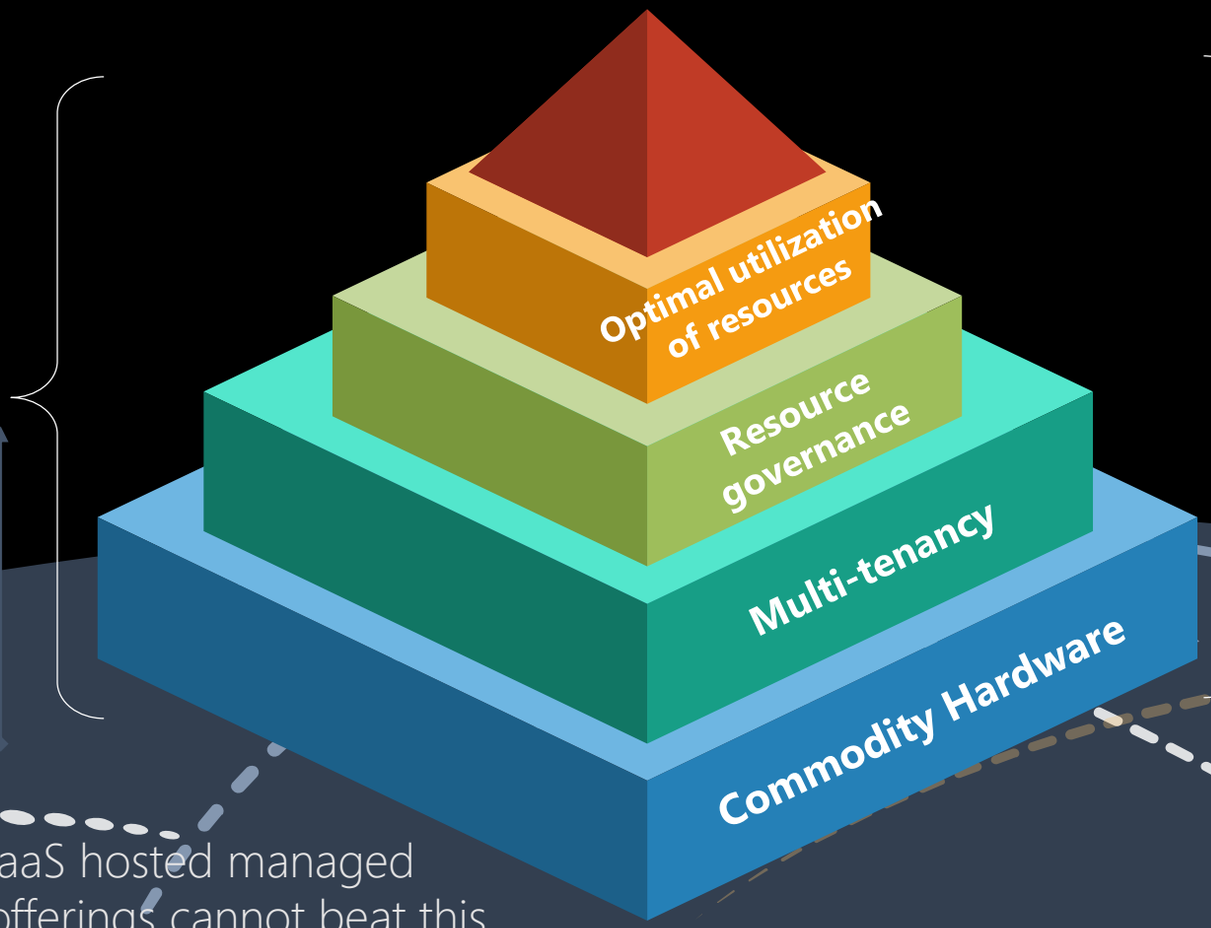

IaaS hosted managed offerings cannot beat this

Gigabytes of data	Petabytes of data
Hundreds transactions / sec	Millions of trans/sec

Scale-out Architecture

Fully-managed and Secure

Global Distribution from the Ground Up

Azure Cosmos DB: Value to Customer

- Save money
- Become more productive
- Become more flexible
- Become more responsive
- Become more innovative

Turnkey global distribution

Azure Cosmos DB

Globally-distribute data around the world

Automatically replicate all your data around the world – across more regions than Amazon and Google combined

Global Distribution From The Ground-Up

- Azure Cosmos DB is a Foundational (Ring 0) Azure service
  - Available in *all* Azure regions by default, including sovereign/government clouds
- Transparent and automatic multi-region replication
  - Associate *any* number of regions with your database account, *at any time*
  - Policy based geo-fencing
- Multi-homing APIs
  - All endpoints are logical, by default
  - Apps don't need to be redeployed during regional failover
  - Apps can also access physical endpoints if needed
- Support for both manual and automatic failover
- Designed for high availability
  - Allows for dynamically setting *priorities* to regions
  - *Simulate* regional disasters via API
  - Test the *end-to-end availability* for the entire app (beyond just the database)

- Comprehensive SLAs
  - First and only to offer comprehensive SLA for latency, throughput, availability and consistency

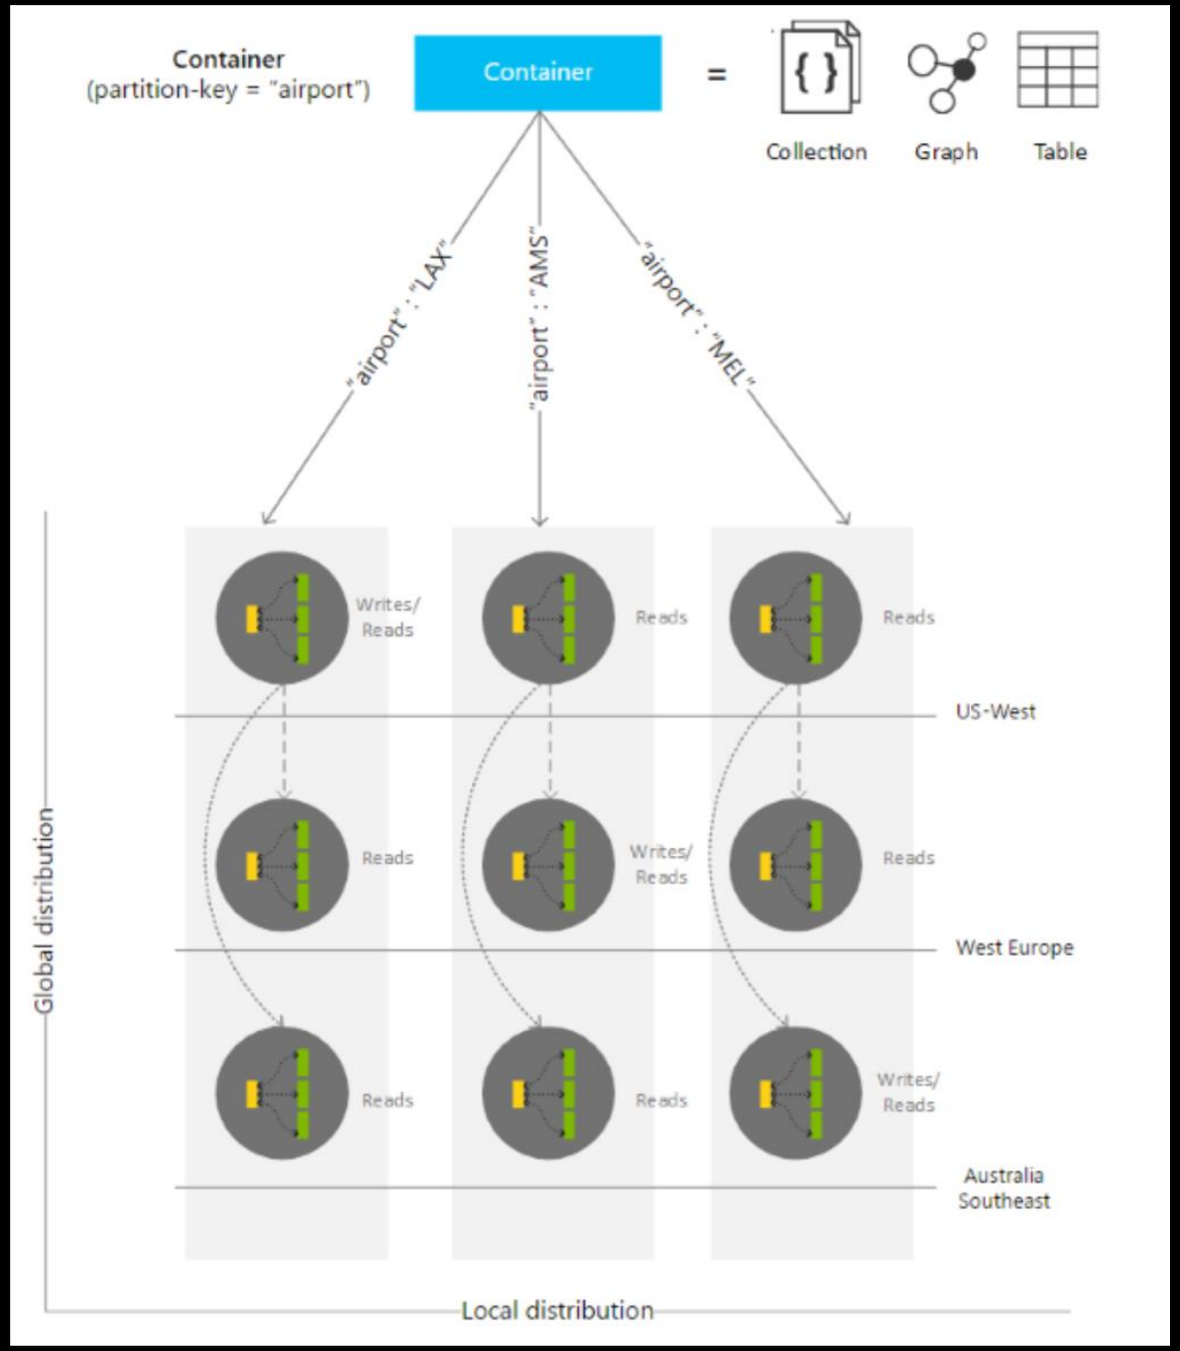

Replication

- Choice of partition key is important
- Each partition is replicated separately
- Each partition can have separate write region
- This makes local read and write possible

Demo – Global Distribution

Guaranteed Single Digit ms Latency at the 99th Percentile

Guaranteed low latency – Value to Customer

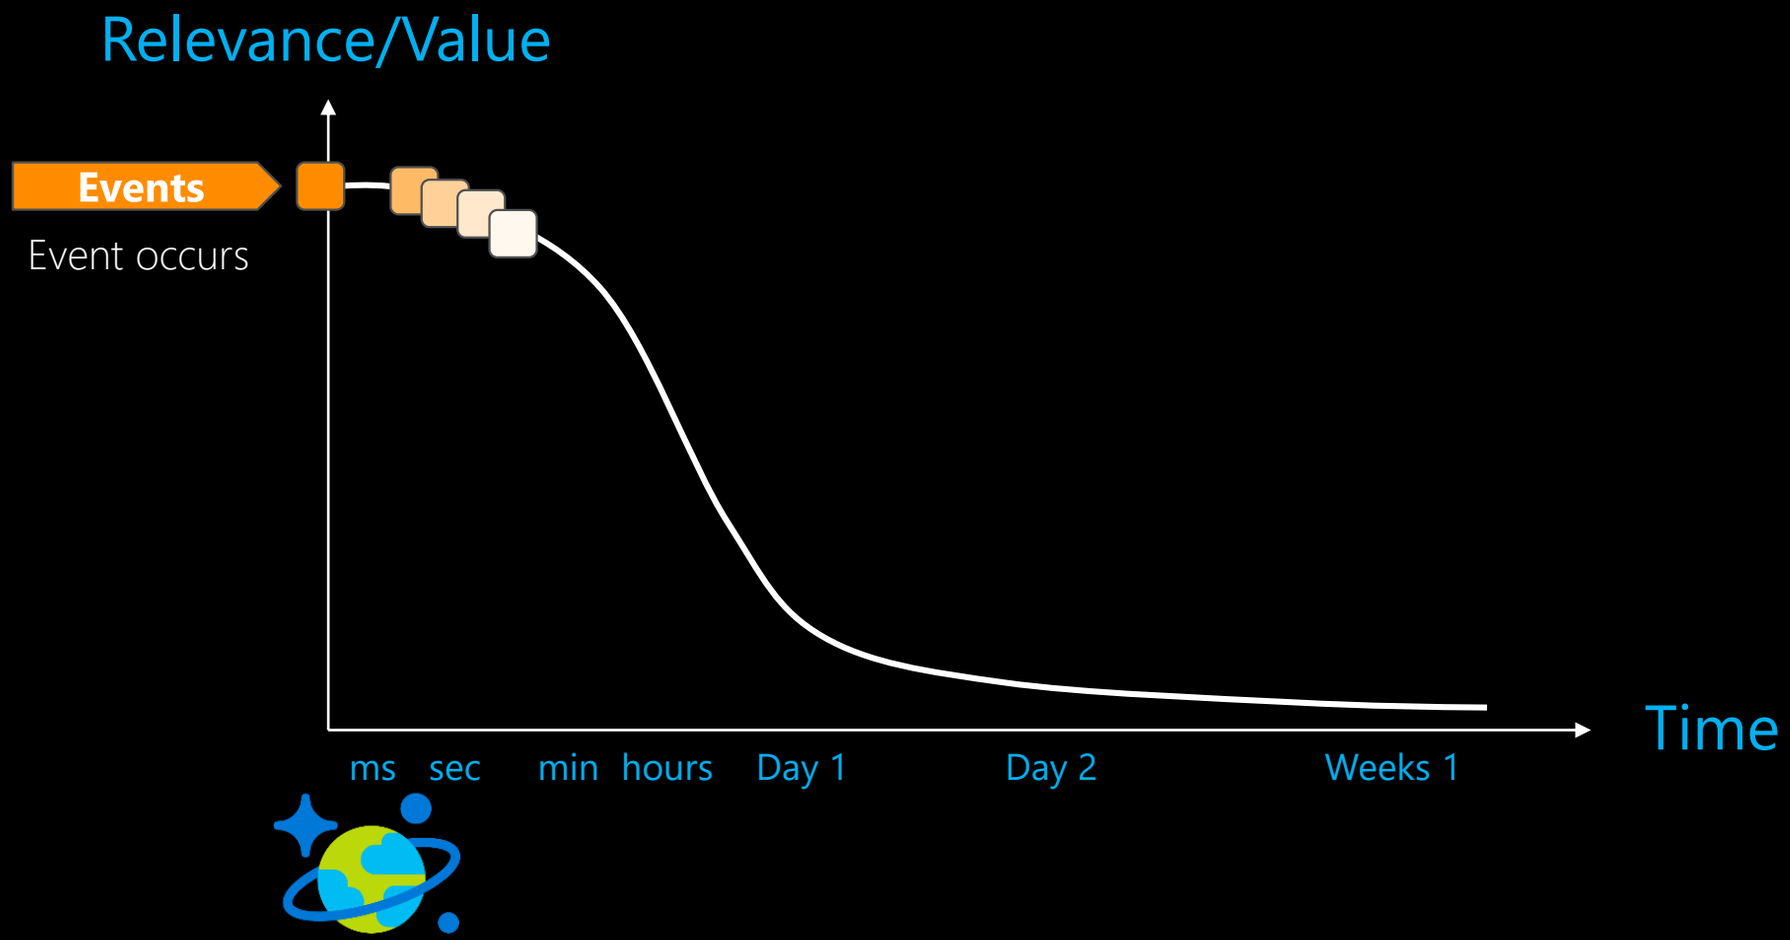

where Cosmos DB fits

Guaranteed Low Latency

- Reads and writes served from local region
- Guaranteed millisecond latency worldwide
- Write optimized, latch-free database engine
- Automatically indexed SSD storage
- Synchronous and automatic indexing at sustained ingestion rates
  - No schema or index management needed
  - No schema versioning needed
  - No schema migration needed
  - All of this is highly relevant for rapidly evolving apps in a globally distributed setup

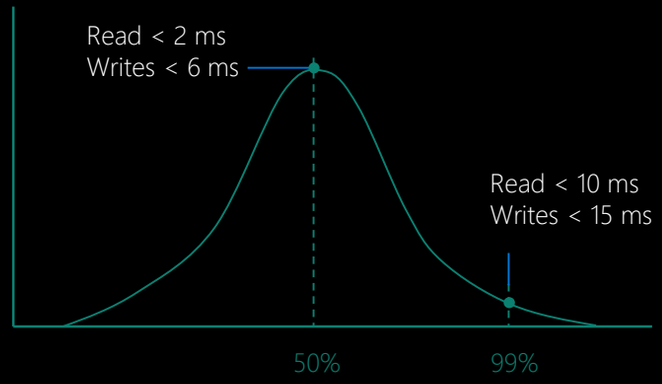

	Reads (1KB)	Indexed writes (1KB)
50th	<2ms	<6ms
99th	<10ms	<15ms

3

Elastic Scaleout

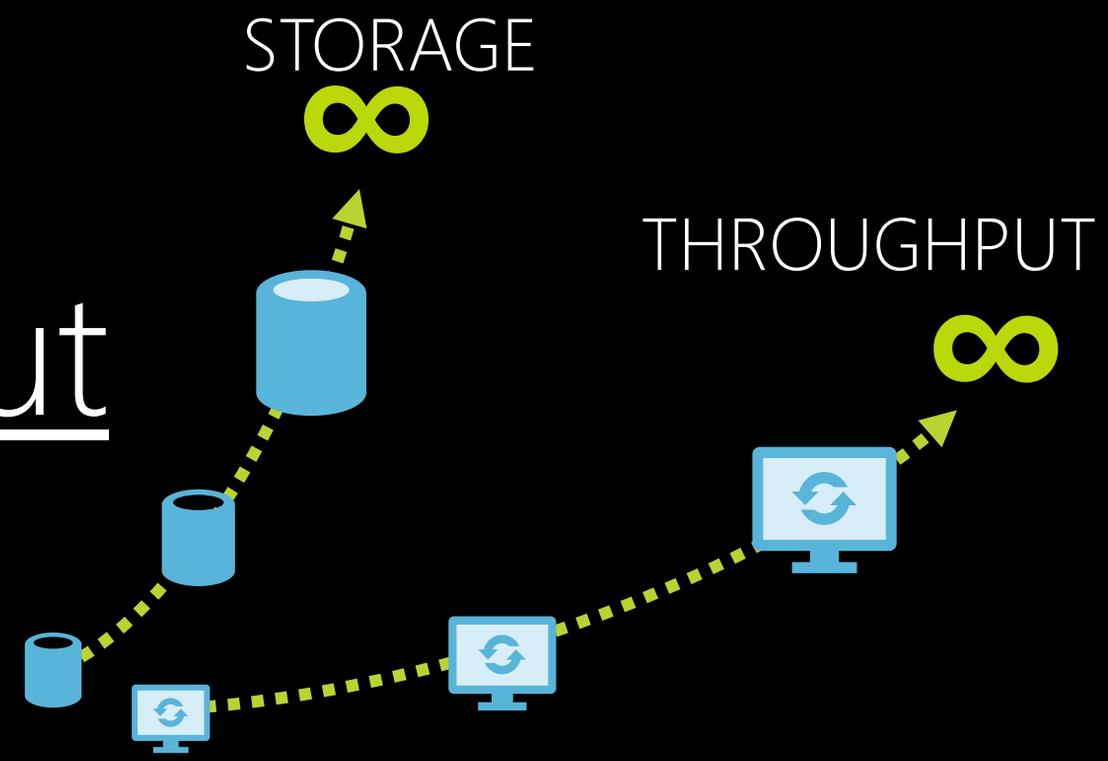

Azure Cosmos DB: Elastically Scalable Storage

- Single machine is never a bottleneck
- A single table can scale from GB-PBs, across many machines, and regions
- Transparent server side partition management and routing
- Optionally delete old data using built-in support for TTL
  - Policy based, automatic tiering to any HDFS compatible data lake (e.g. ADLS or Azure Storage)
  -
- Customers pay only for the throughput and storage they need

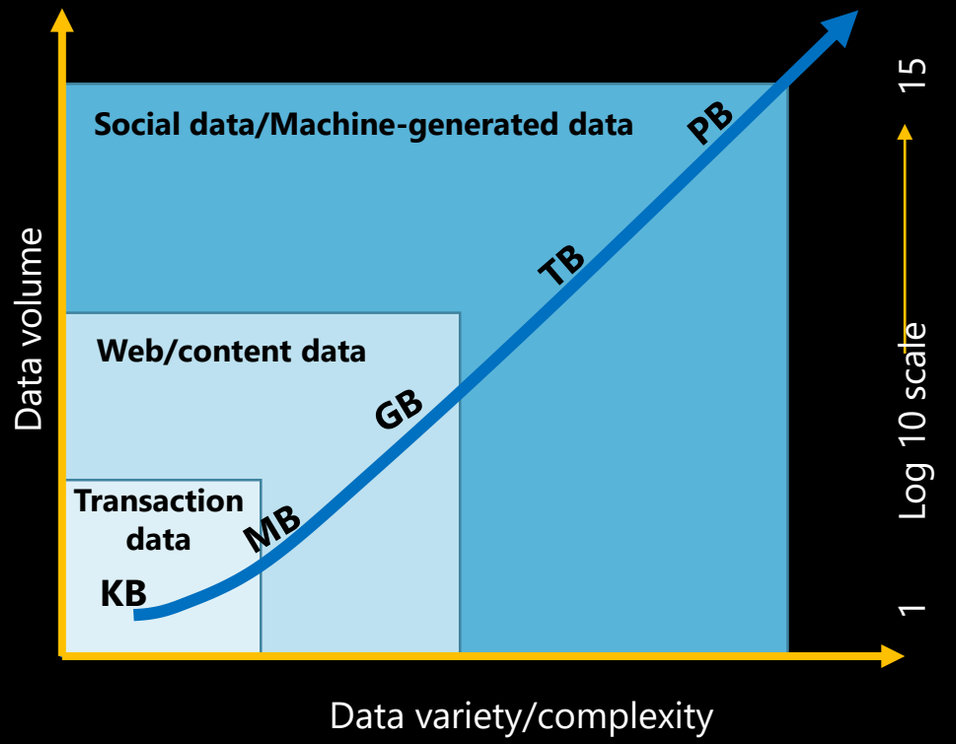

Azure Cosmos DB: Elastically Scalable Throughput

- Elastically scale throughput from 10 to 100s of millions of requests/sec across multiple regions
- Support for requests/sec for different workloads
  - This ensures that never have to provision for the peak
- Customers pay only for the throughput and storage they need
- Customers pay by the hour for the provisioned throughput

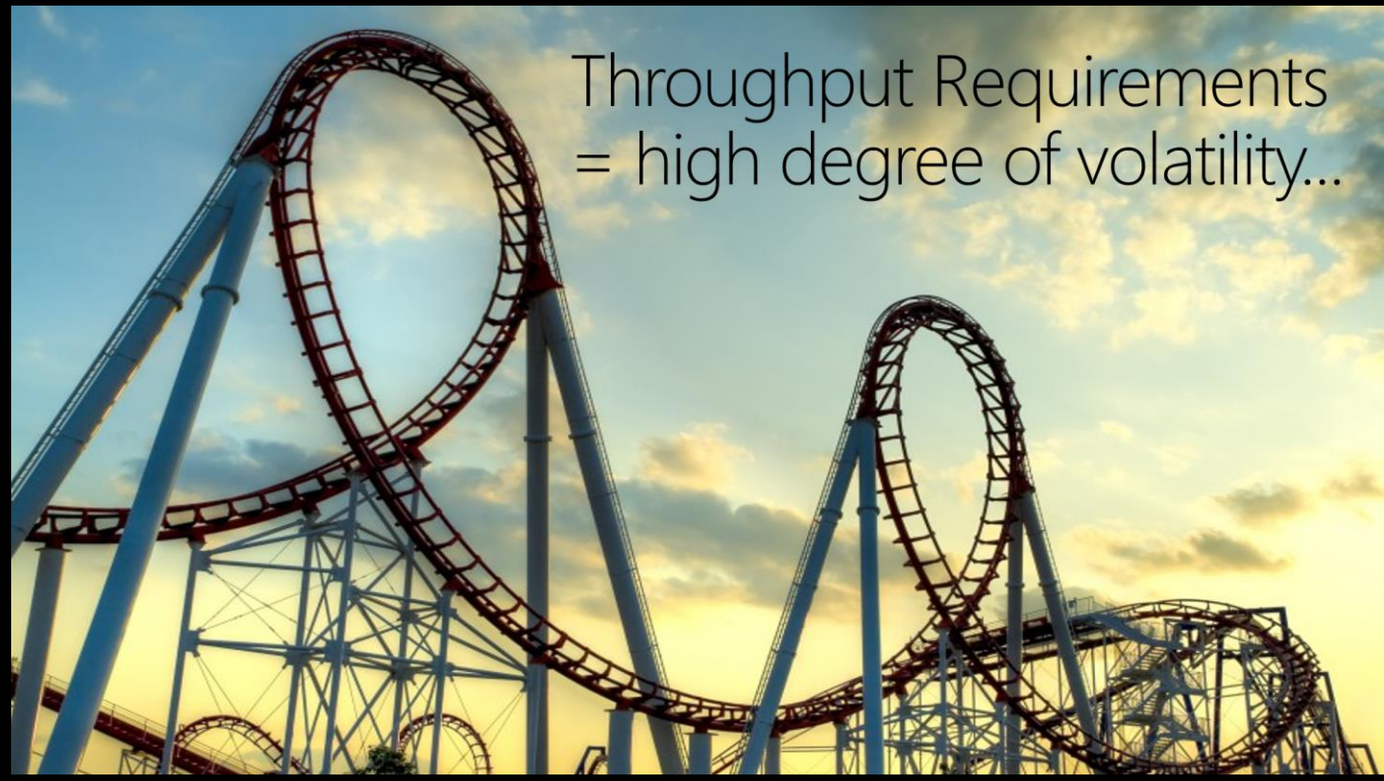

Azure Cosmos DB Total Cost of Ownership (TCO)

Cosmos DB – Lowest TCO

Deeply exploits cloud core properties and economies of scale

- Commodity hardware
- Fine-grained multi-tenancy
- End to end resource governance
- Optimal utilization of resources

Cosmos DB: 5-10X more cost-effective

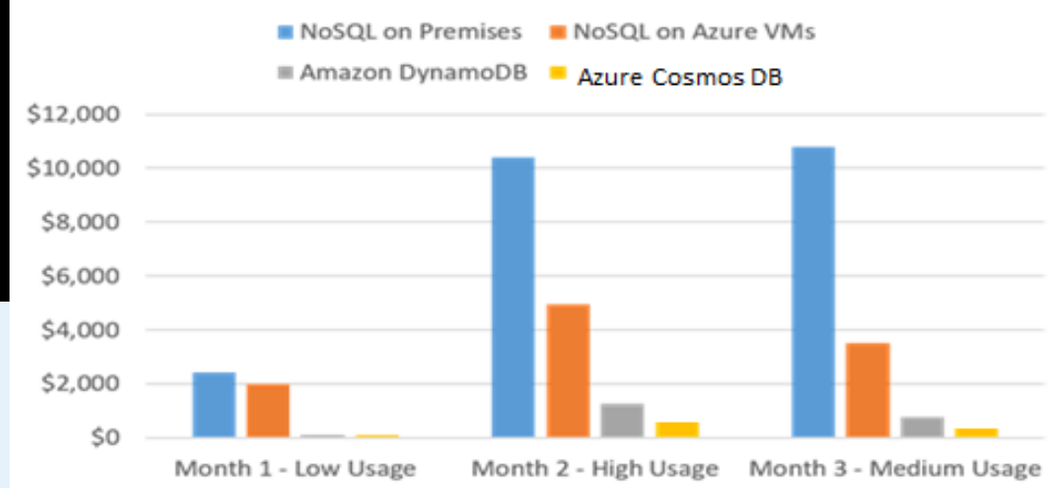

	Month 1 - Low Usage	Month 2 - High Usage	Month 3 - Medium Usage
Reads (per sec)	50	3500 (peak) 1200 (off-peak)	1200 (peak) 700 (off peak)
Writes (per sec)	25	1500 (peak) 800 (off-peak)	800 (peak) 300 (off-peak)
Data stored	200 GB	900 GB	1200 GB
Data Egress	200 GB	1500 GB	600 GB

RU/m - Predictable Performance For Unpredictable Needs

Customers save 60-73% in provisioning cost!

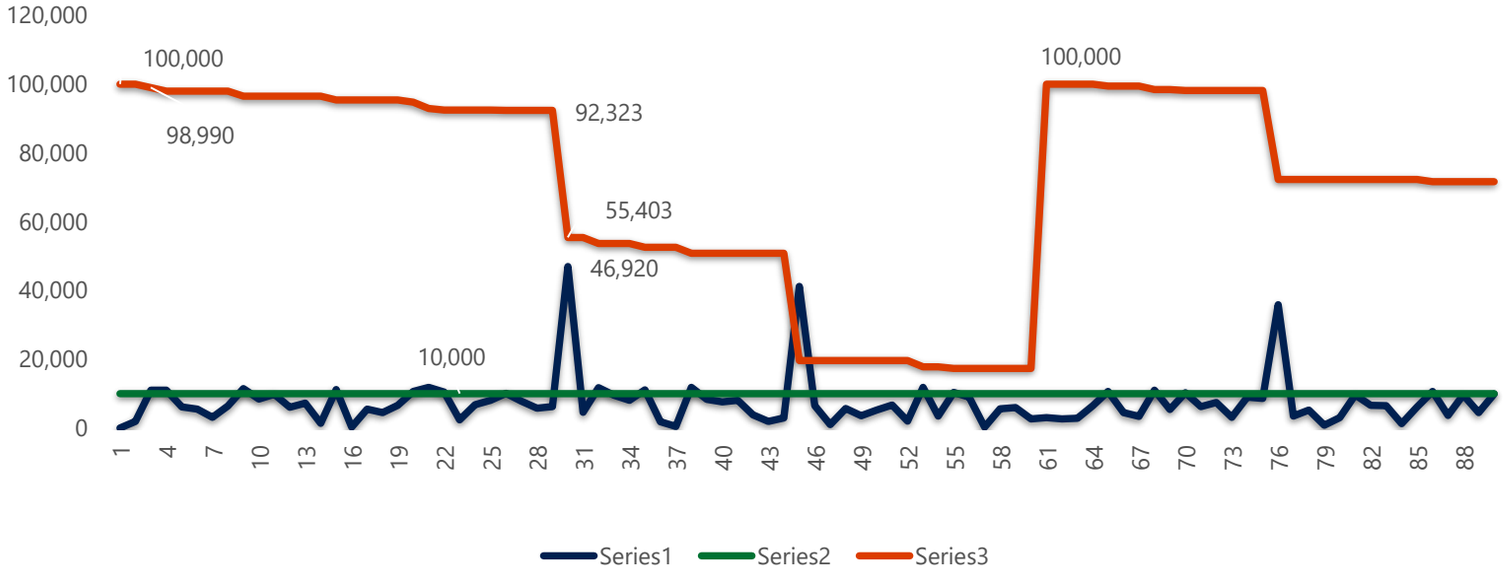

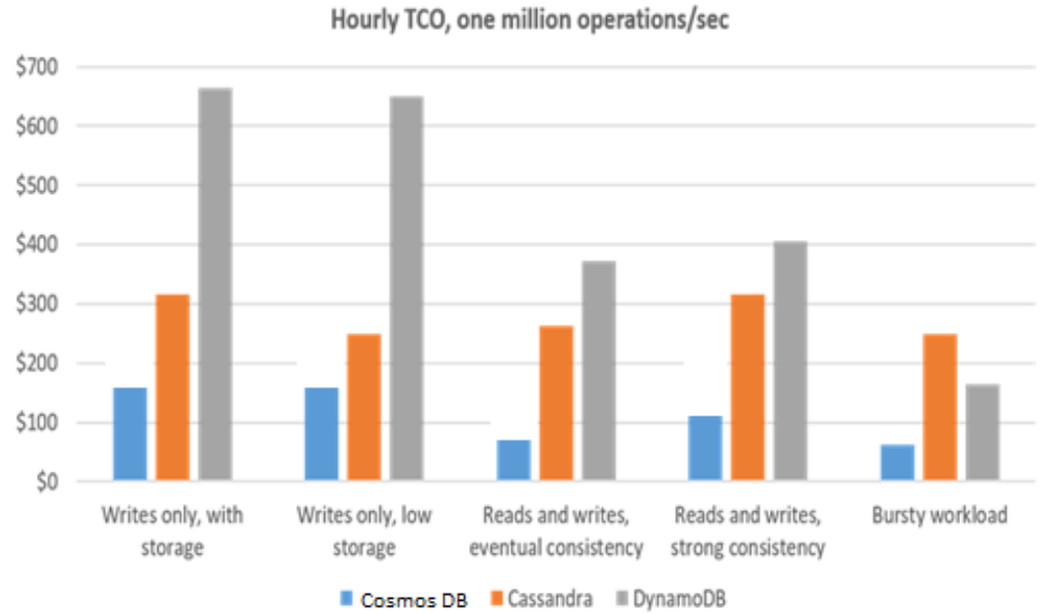

TCO Factors

- Fine-grained multi-tenancy
  - Optimal utilization of resources coupled with highly responsive cluster and datacenter wide load balancing across heterogeneous workloads of customers, enables tremendous amounts of cost savings.
- No administration dev/ops required
  - Cosmos DB = a managed cloud service, you do not need to employ a dev/ops team to handle deployments, maintenance, scale, patching and other day-to-day tasks required with an OSS NoSQL cluster hosted on-premises or on cloud infrastructure.
- Superior elasticity
  - Cosmos DB throughput can be scaled up and down within seconds, allowing you to reduce the cost of ownership during non-peak times.
- Economy of Scale
  - Managed services like Cosmos DB are operating really large number of nodes, and are able to pass on savings to the customer.

Resource Model And Getting Started With CosmosDB

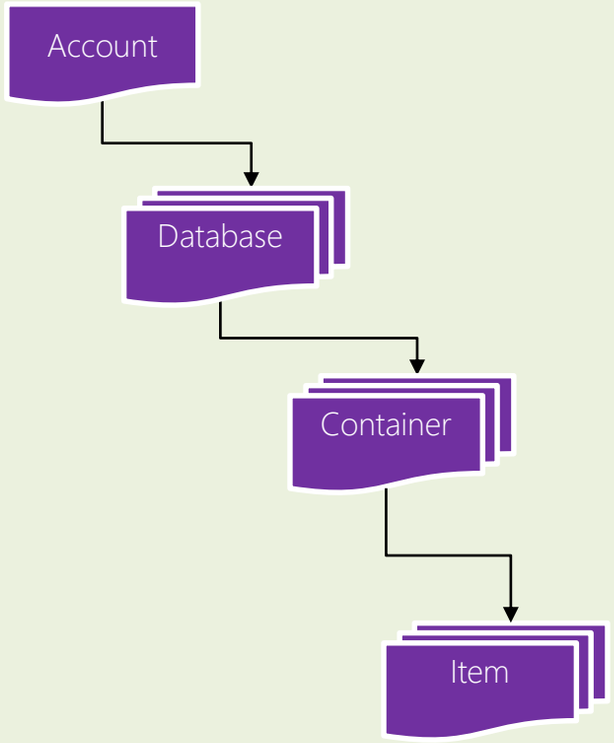
Account
Database
Container
Item

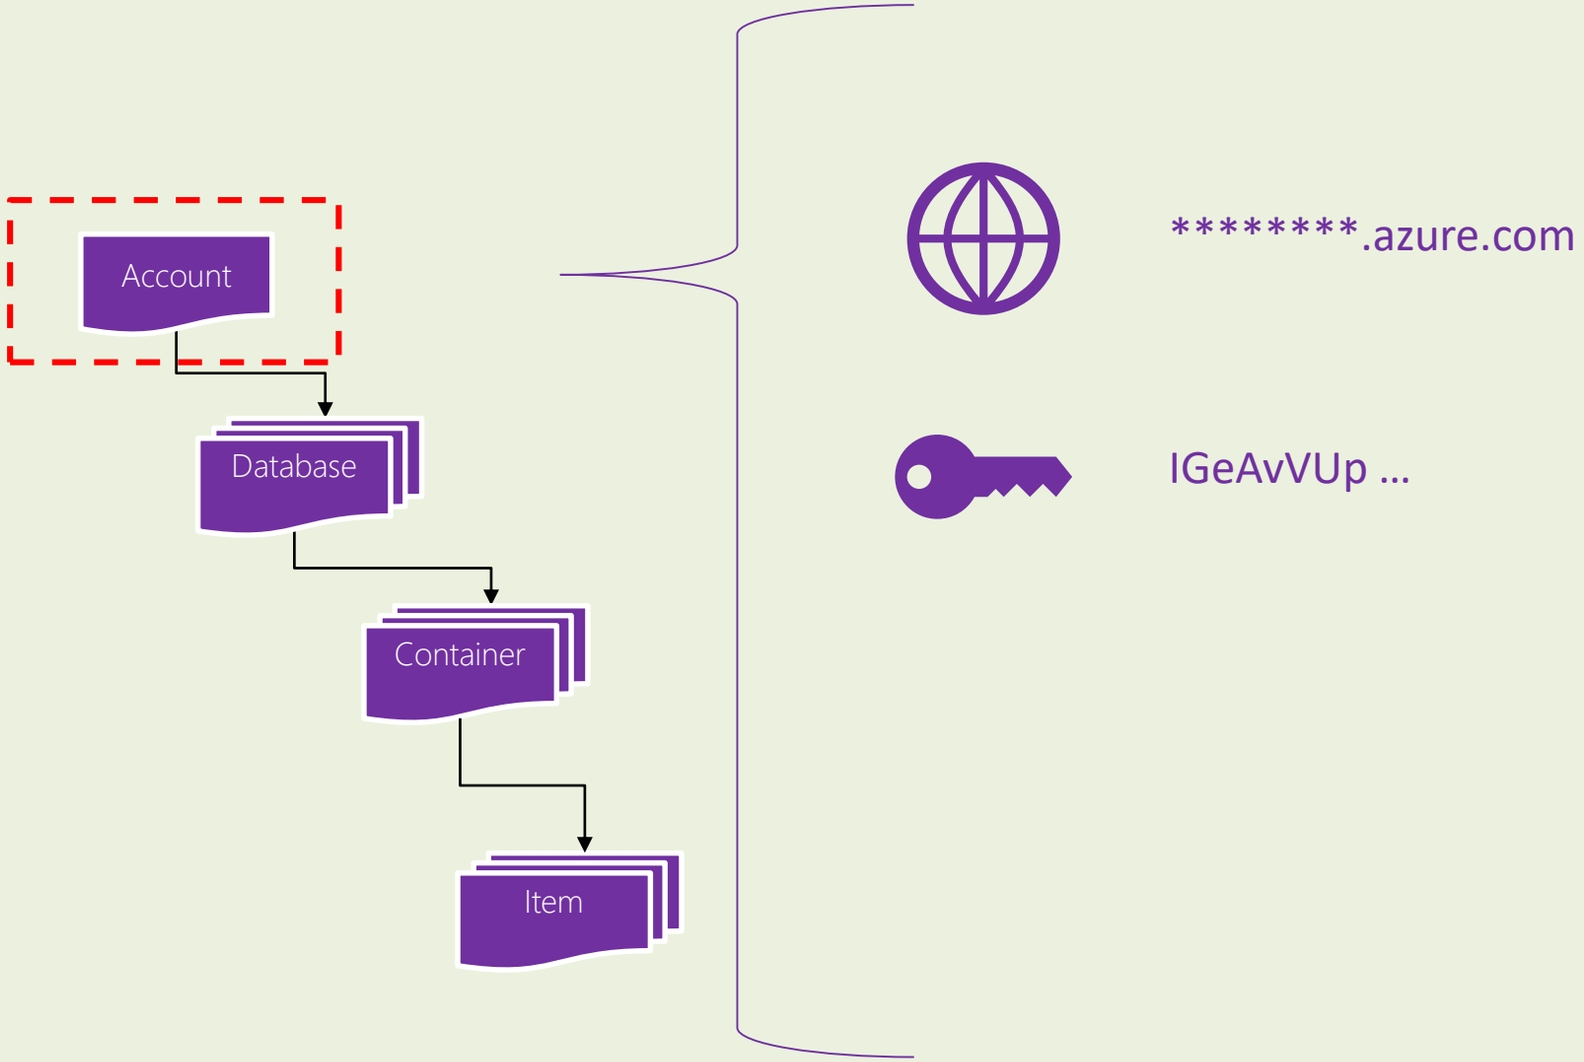
Account
Database
Container
Item
********.azure.com
IGeAvVUp ...

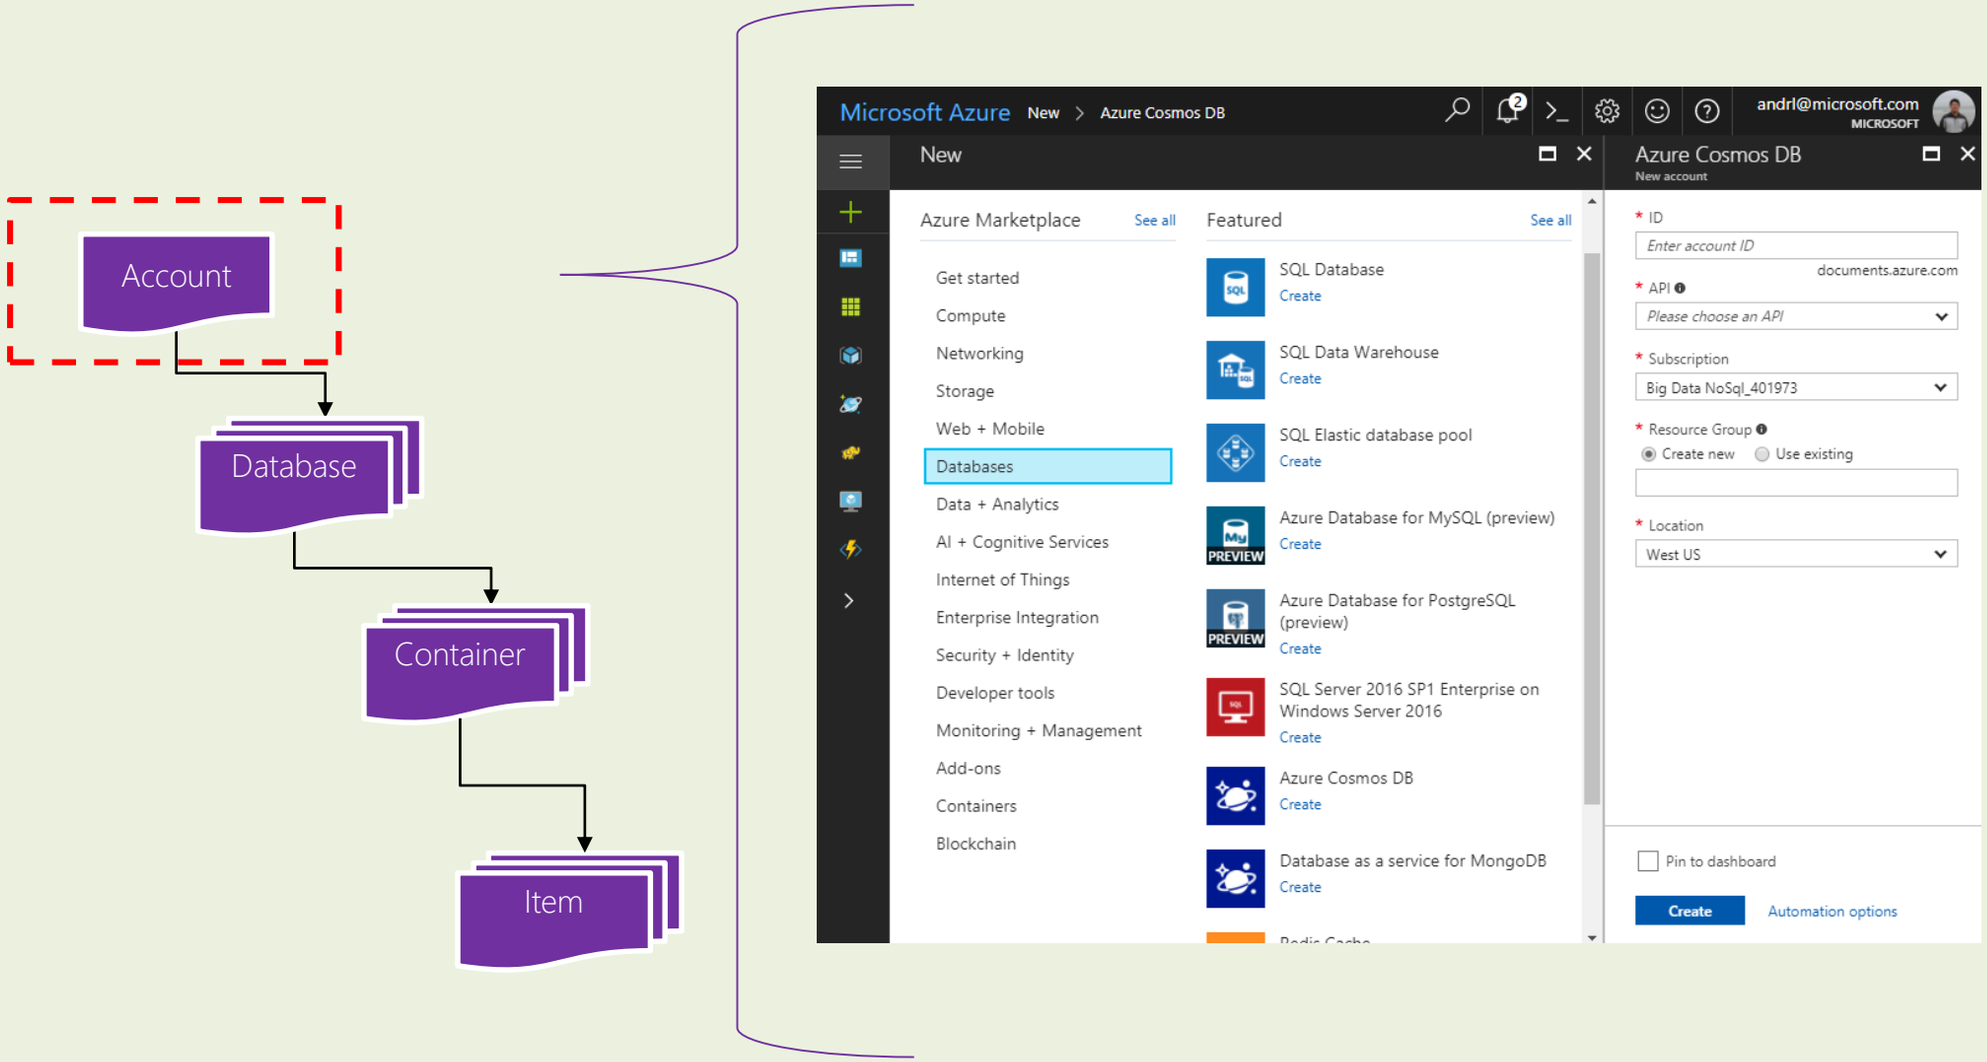
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

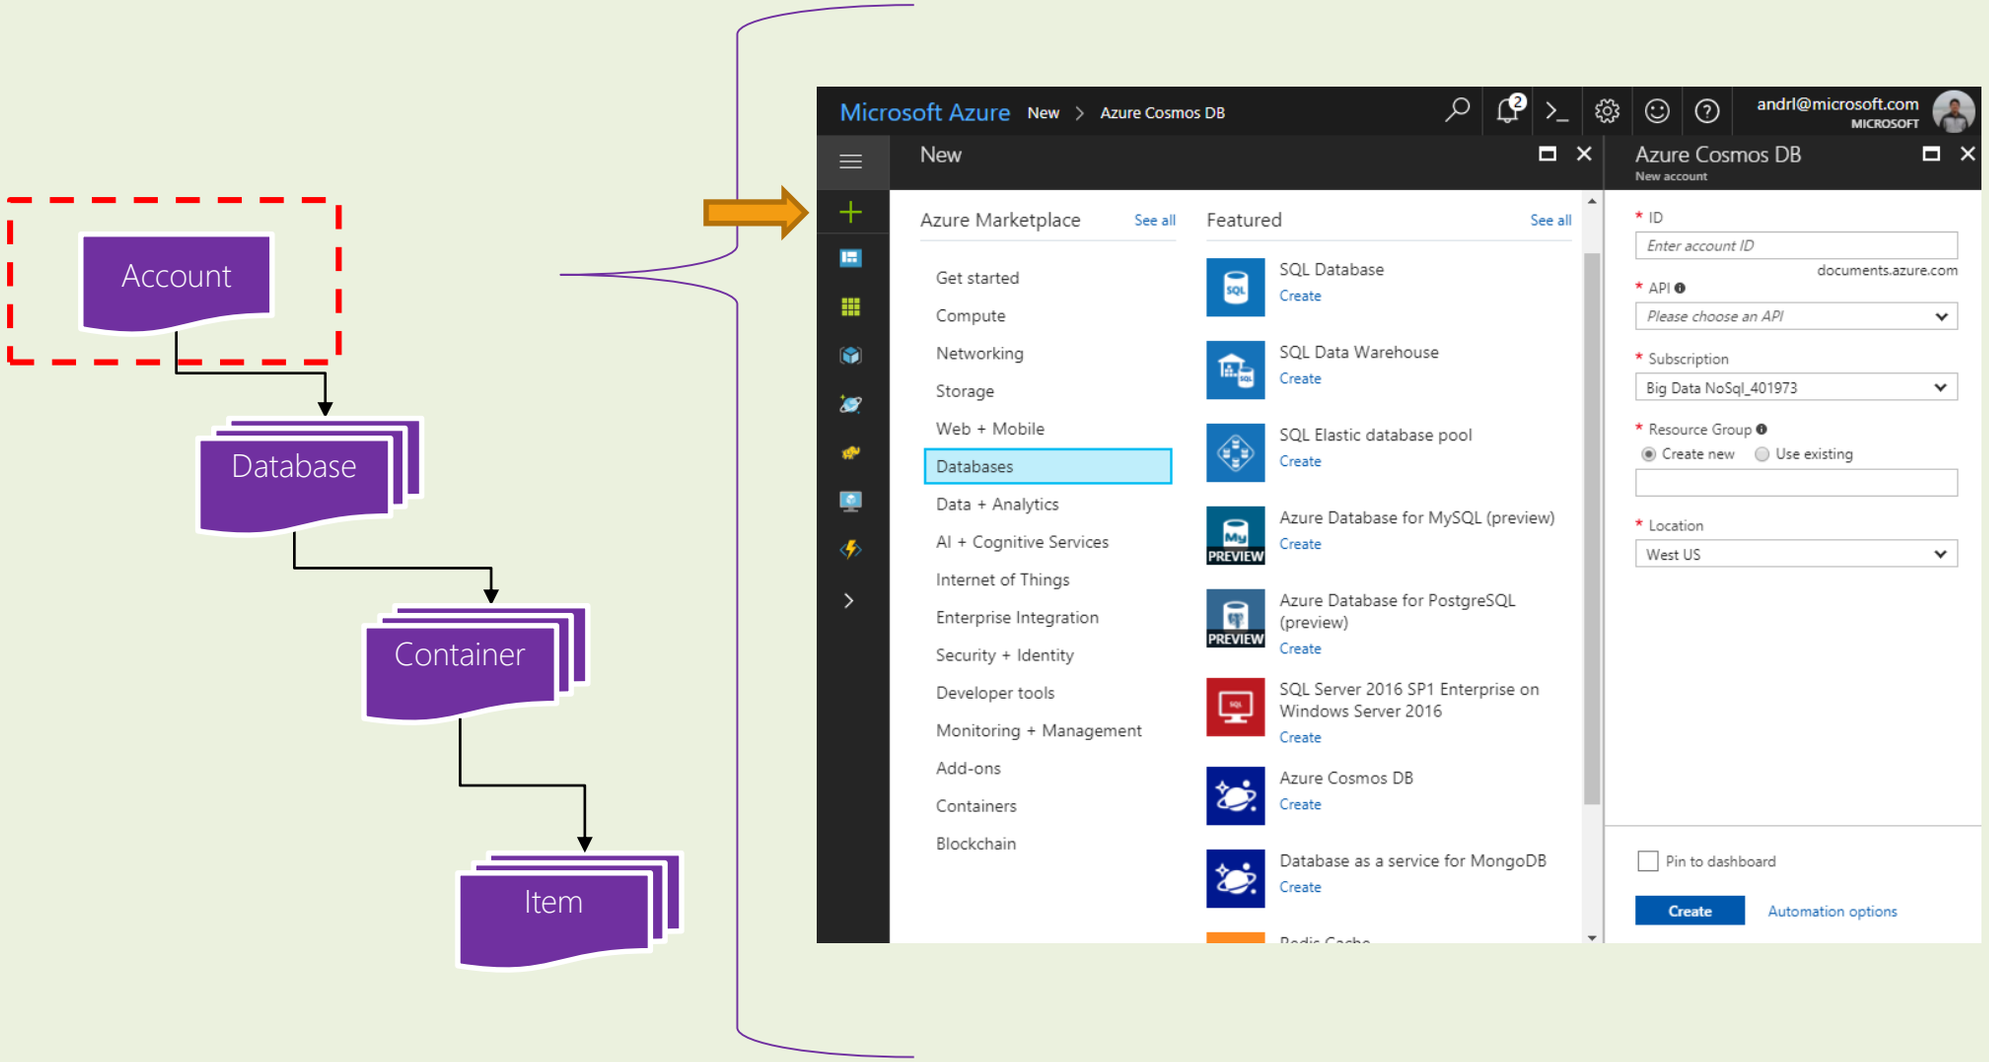
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

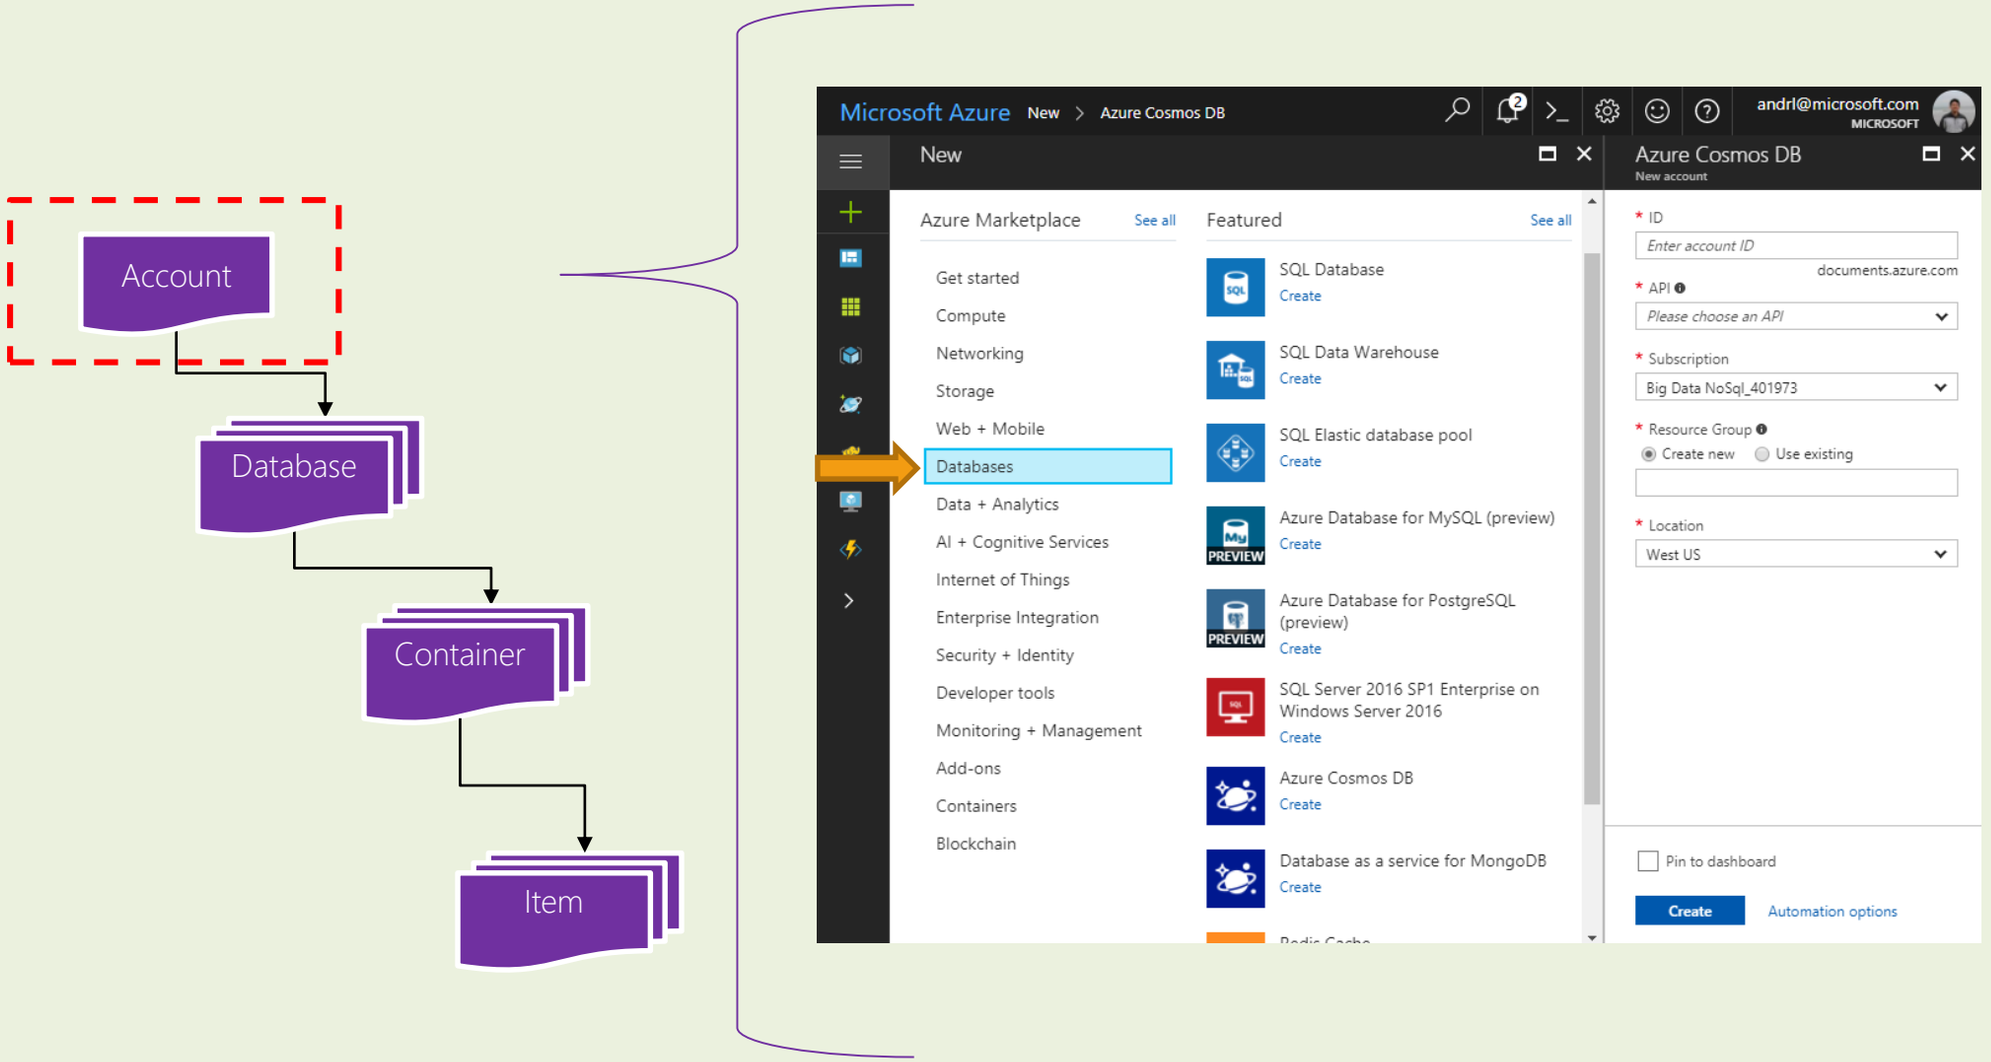
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

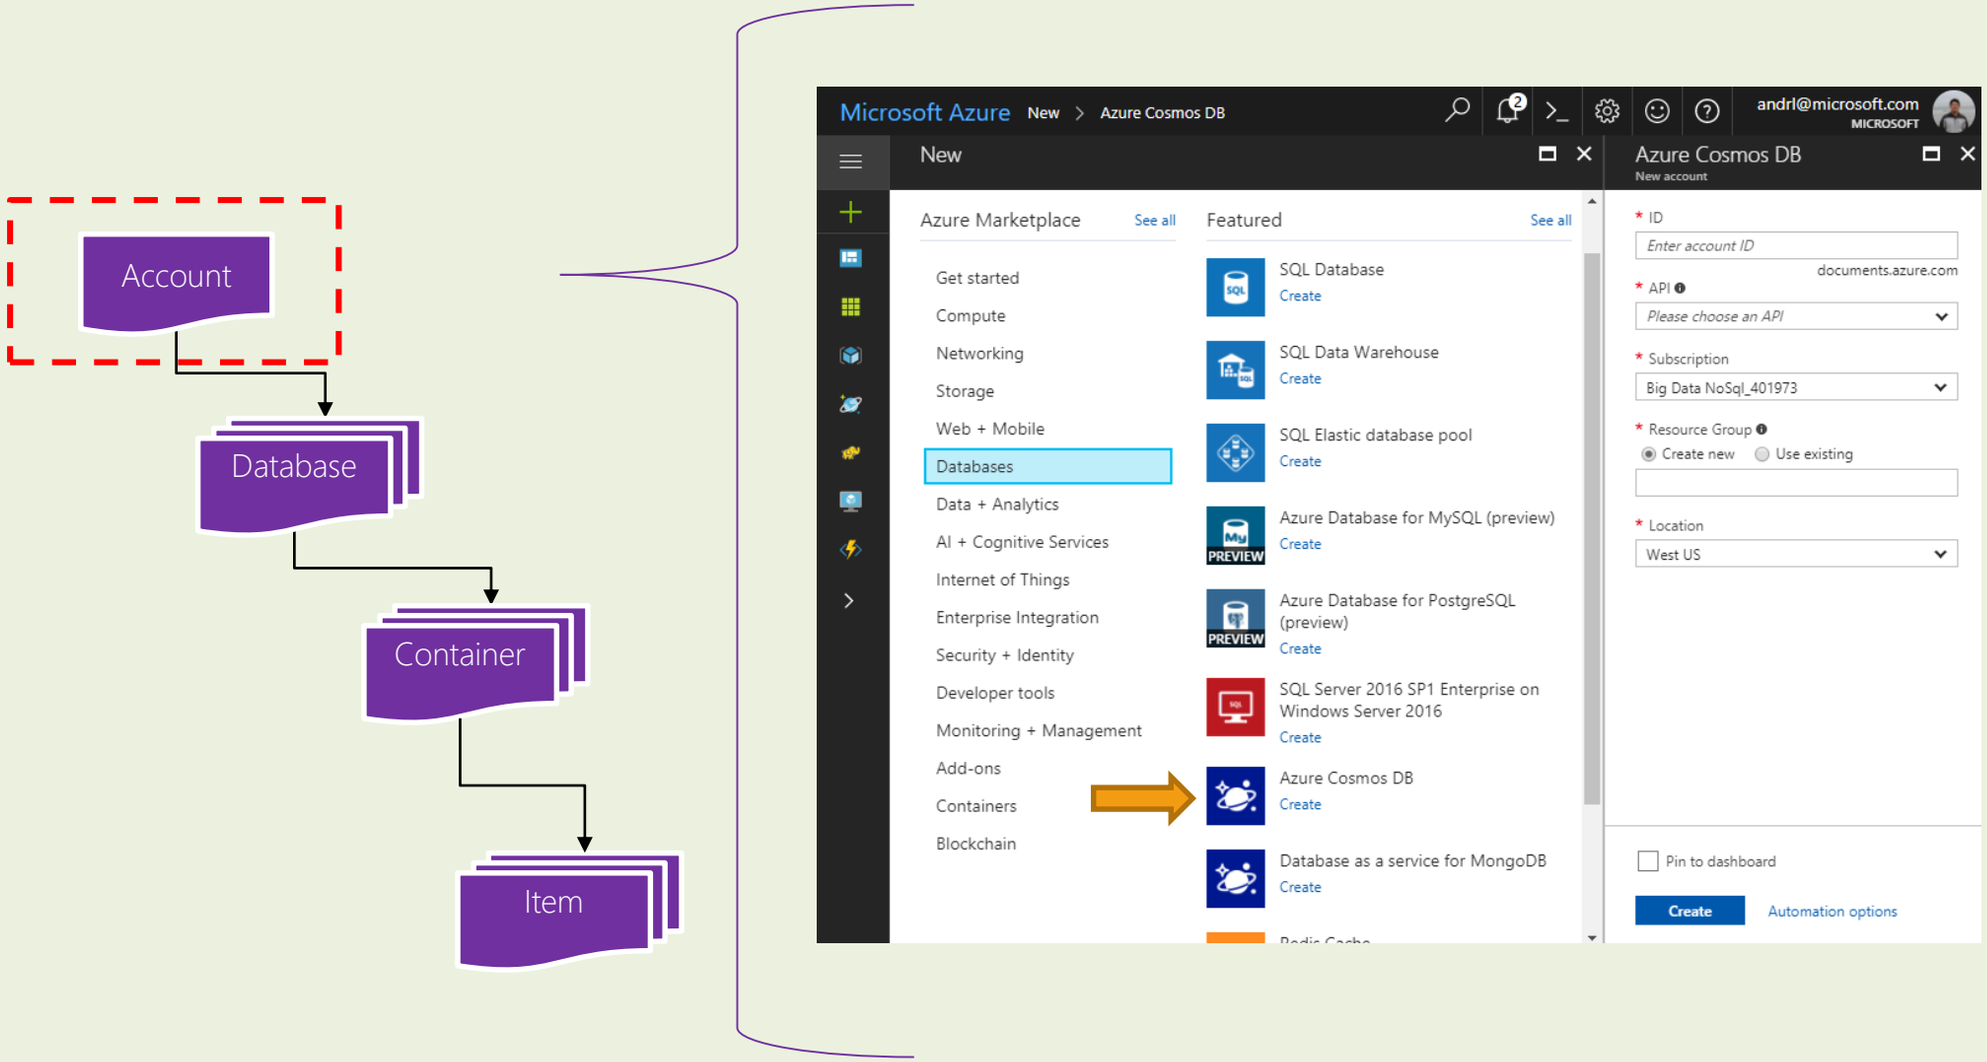
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

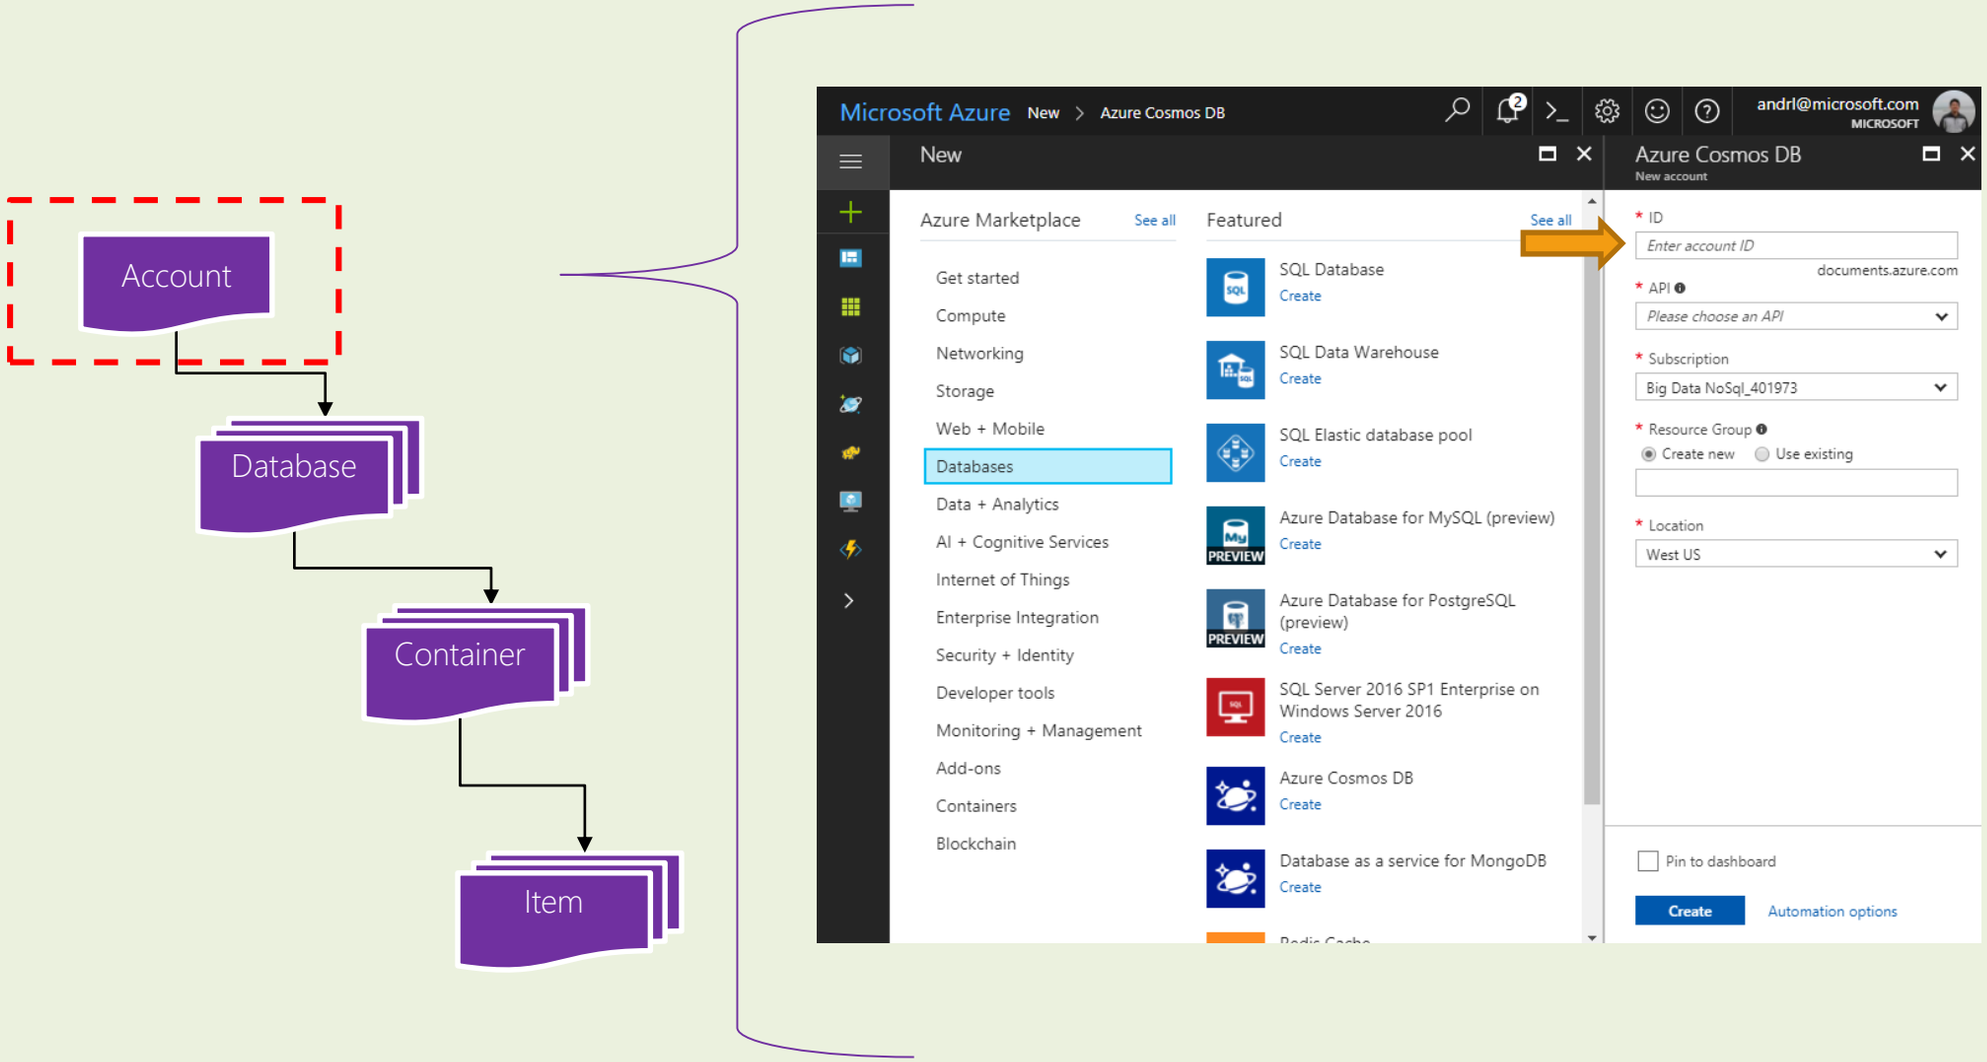
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

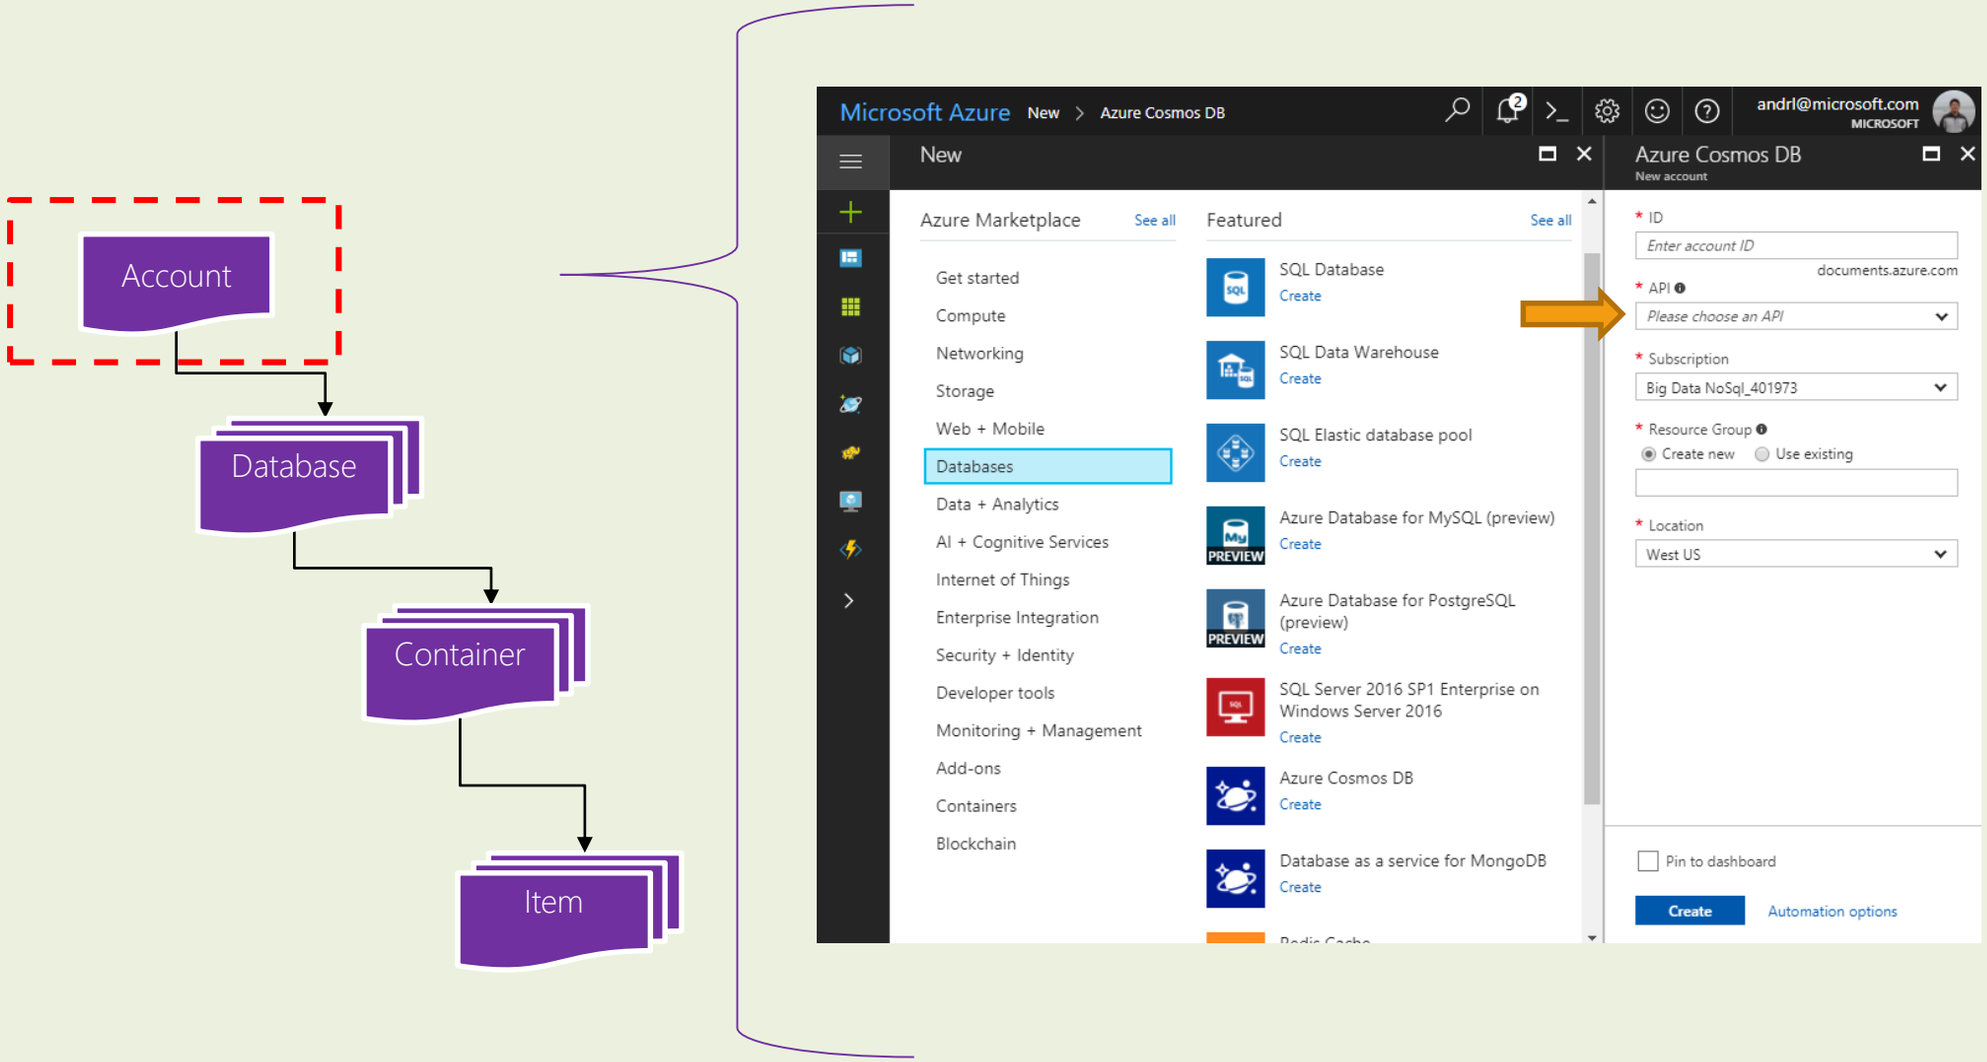
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

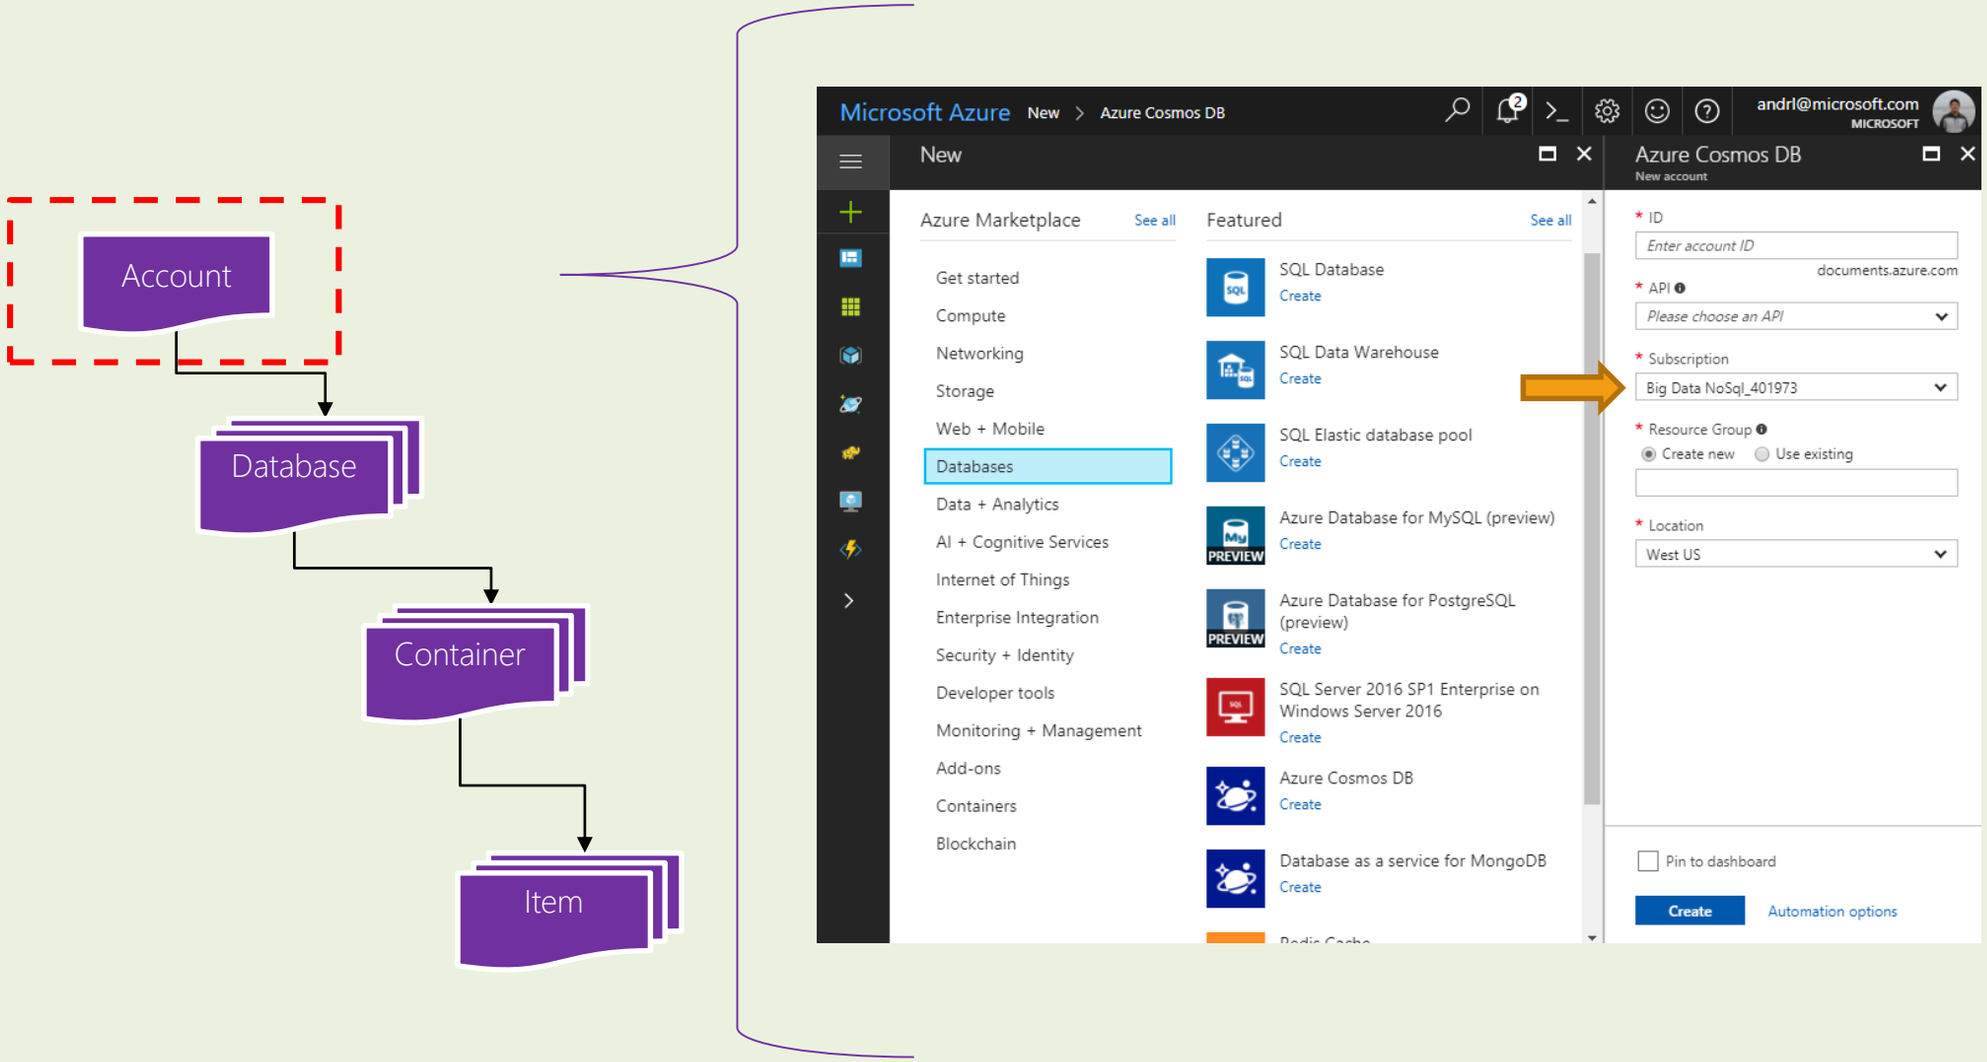
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

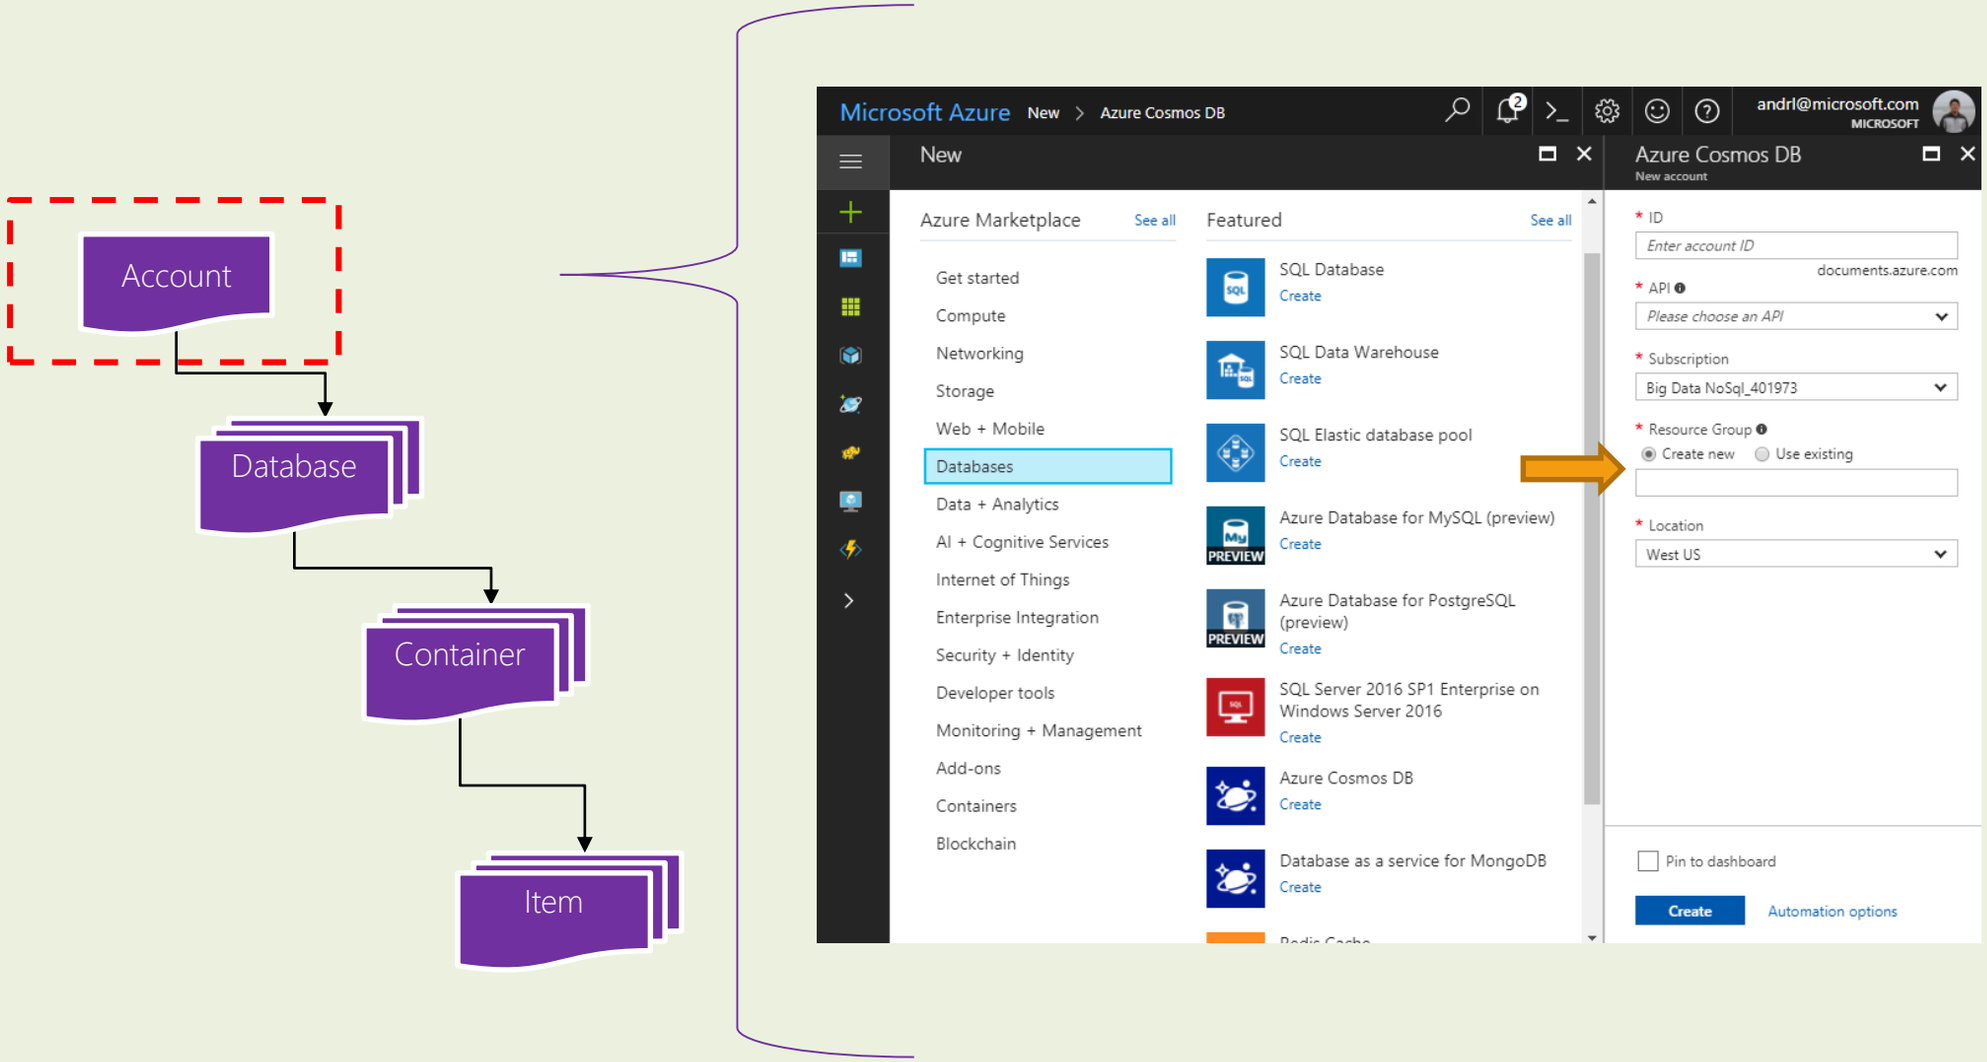
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

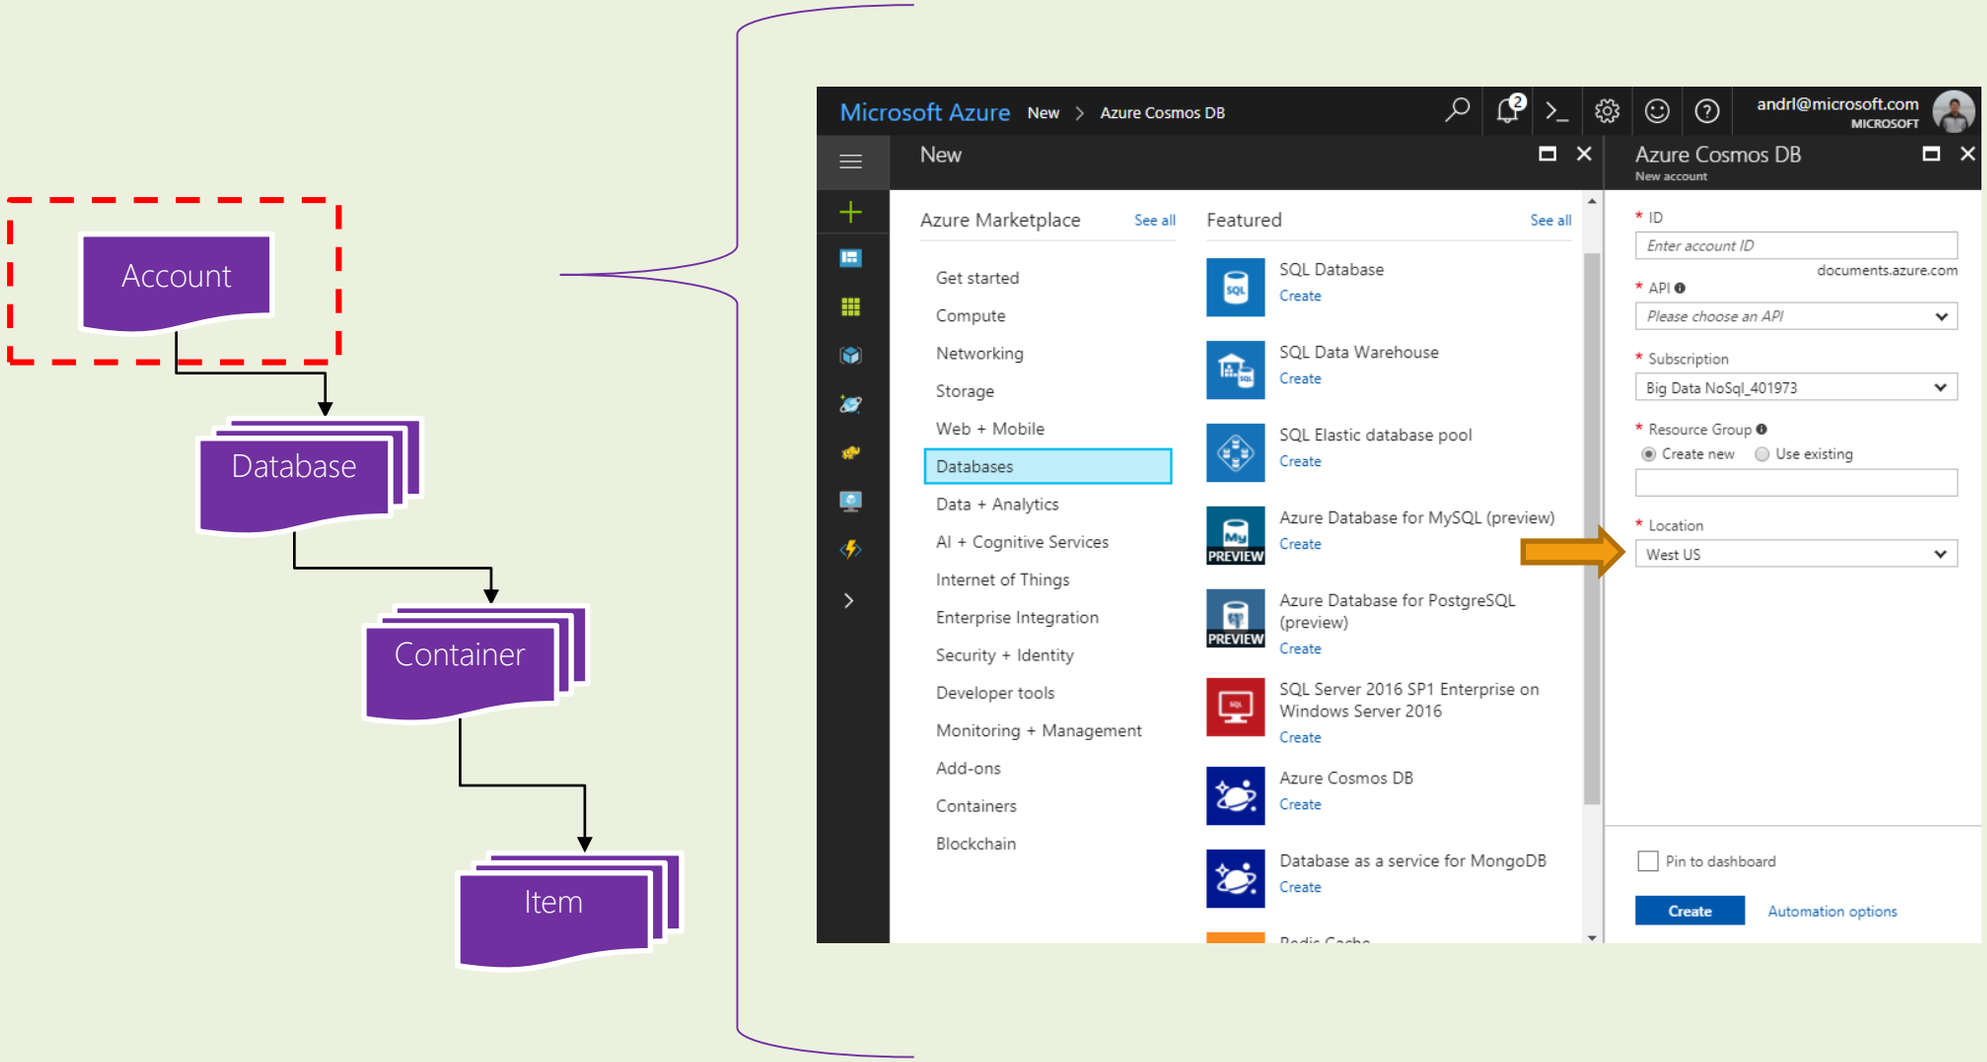
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

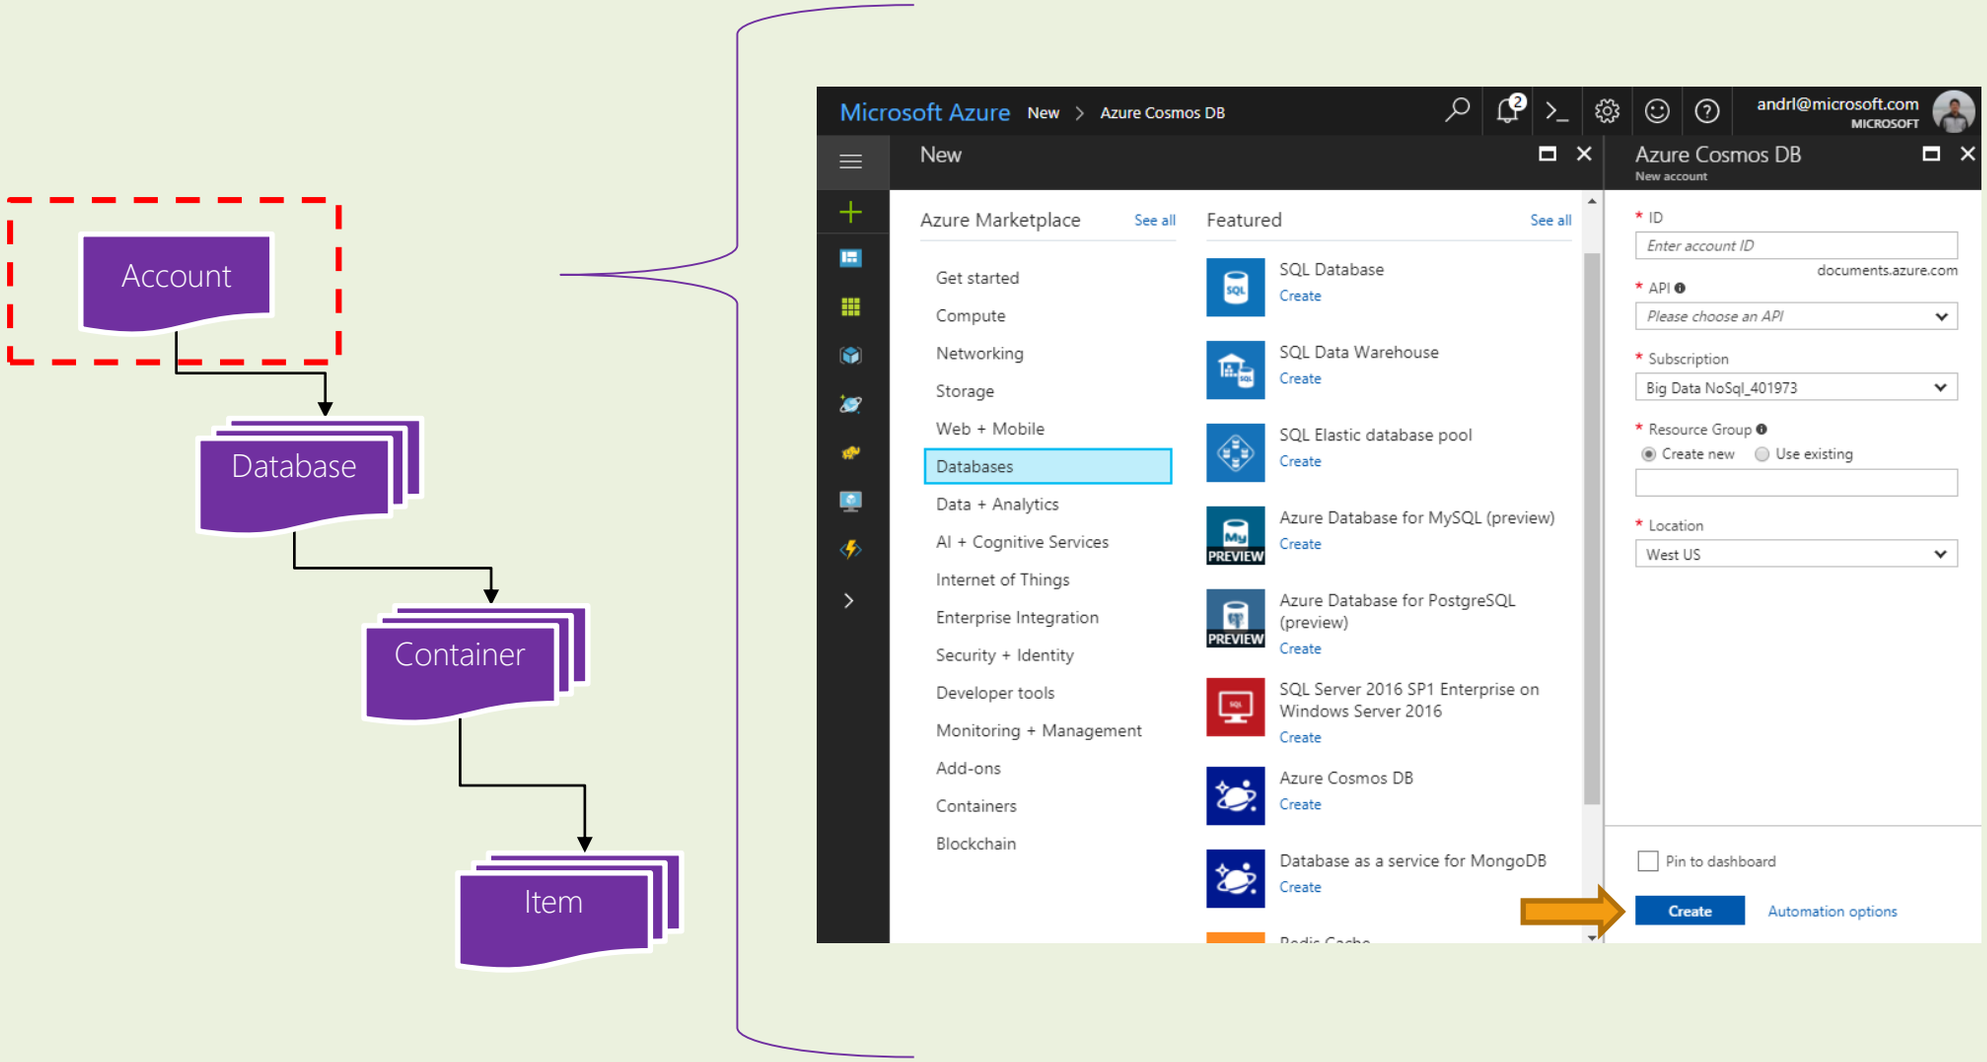
Account
Database
Container
Item
Microsoft Azure
New
Azure Cosmos DB
andrl@microsoft.com
MICROSOFT
New
Azure Marketplace
See all
Get started
Compute
Networking
Storage
Web + Mobile
Databases
Data + Analytics
AI + Cognitive Services
Internet of Things
Enterprise Integration
Security + Identity
Developer tools
Monitoring + Management
Add-ons
Containers
Blockchain
Featured
See all
SQL Database
Create
SQL Data Warehouse
Create
SQL Elastic database pool
Create
Azure Database for MySQL (preview)
Create
Azure Database for PostgreSQL (preview)
Create
SQL Server 2016 SP1 Enterprise on Windows Server 2016
Create
Azure Cosmos DB
Create
Database as a service for MongoDB
Create
Azure Cosmos DB
New account
* ID
Enter account ID
documents.azure.com
* API
Please choose an API
* Subscription
Big Data NoSql_401973
* Resource Group
Create new
Use existing
* Location
West US
Pin to dashboard
Create
Automation options

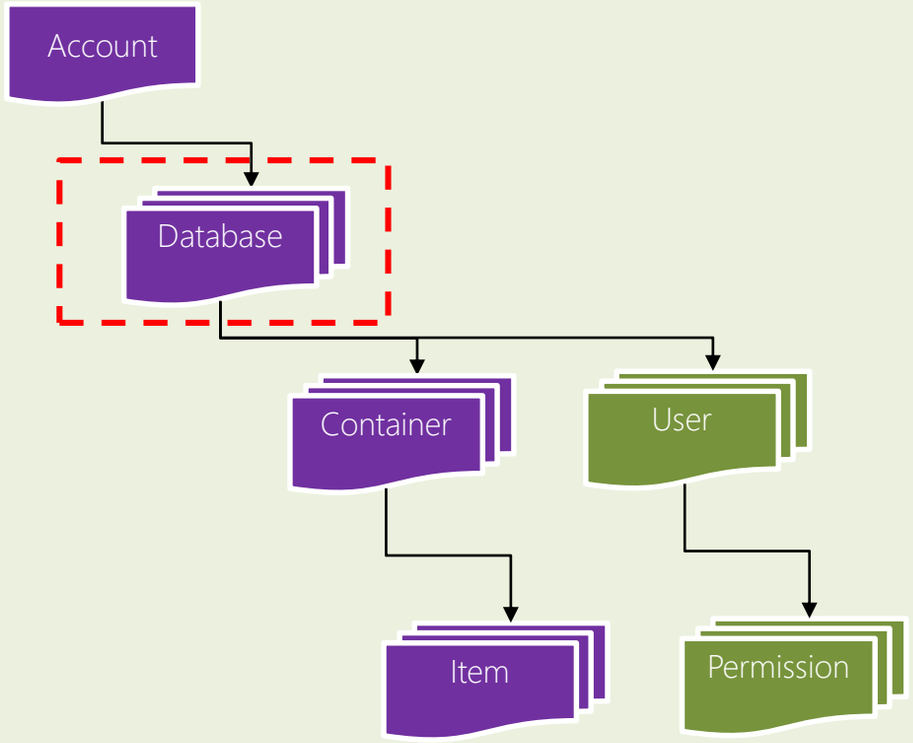
Account
Database
Container
User
Item
Permission

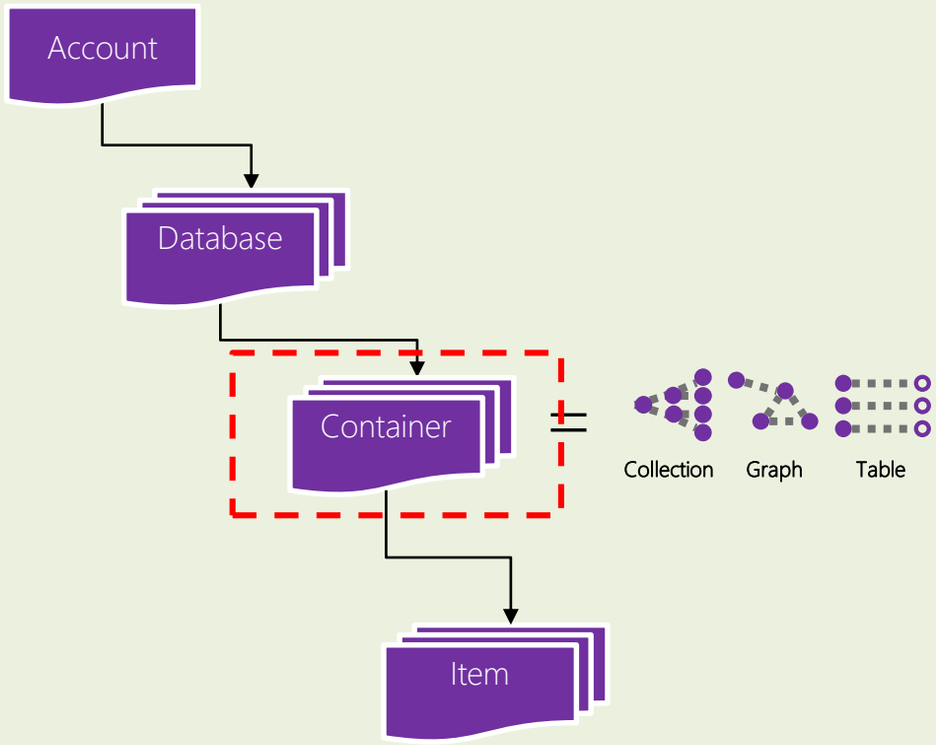
Account
Database
Container
Item
Collection
Graph
Table

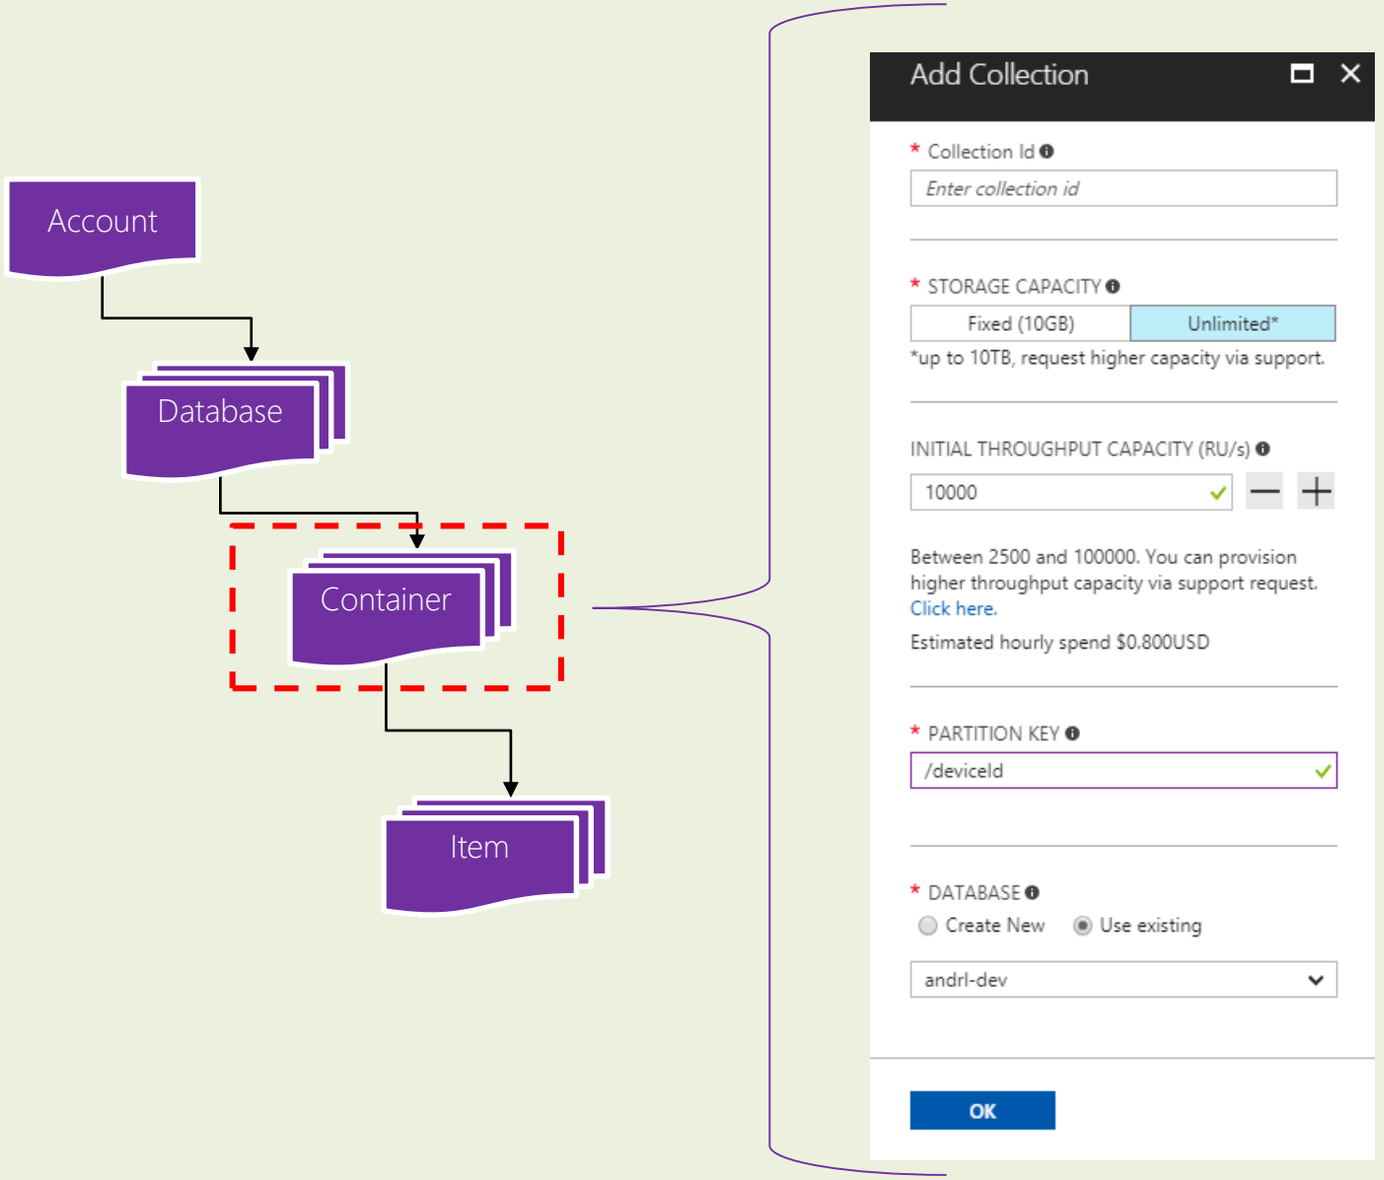
Account
Database
Container
Item
Add Collection
* Collection Id
Enter collection id
* STORAGE CAPACITY
Fixed (10GB)
Unlimited*
*up to 10TB, request higher capacity via support.
INITIAL THROUGHPUT CAPACITY (RU/s)
10000
Between 2500 and 100000. You can provision higher throughput capacity via support request.
Click here.
Estimated hourly spend $0.800USD
* PARTITION KEY
/deviceId
* DATABASE
Create New
Use existing
andrl-dev
OK

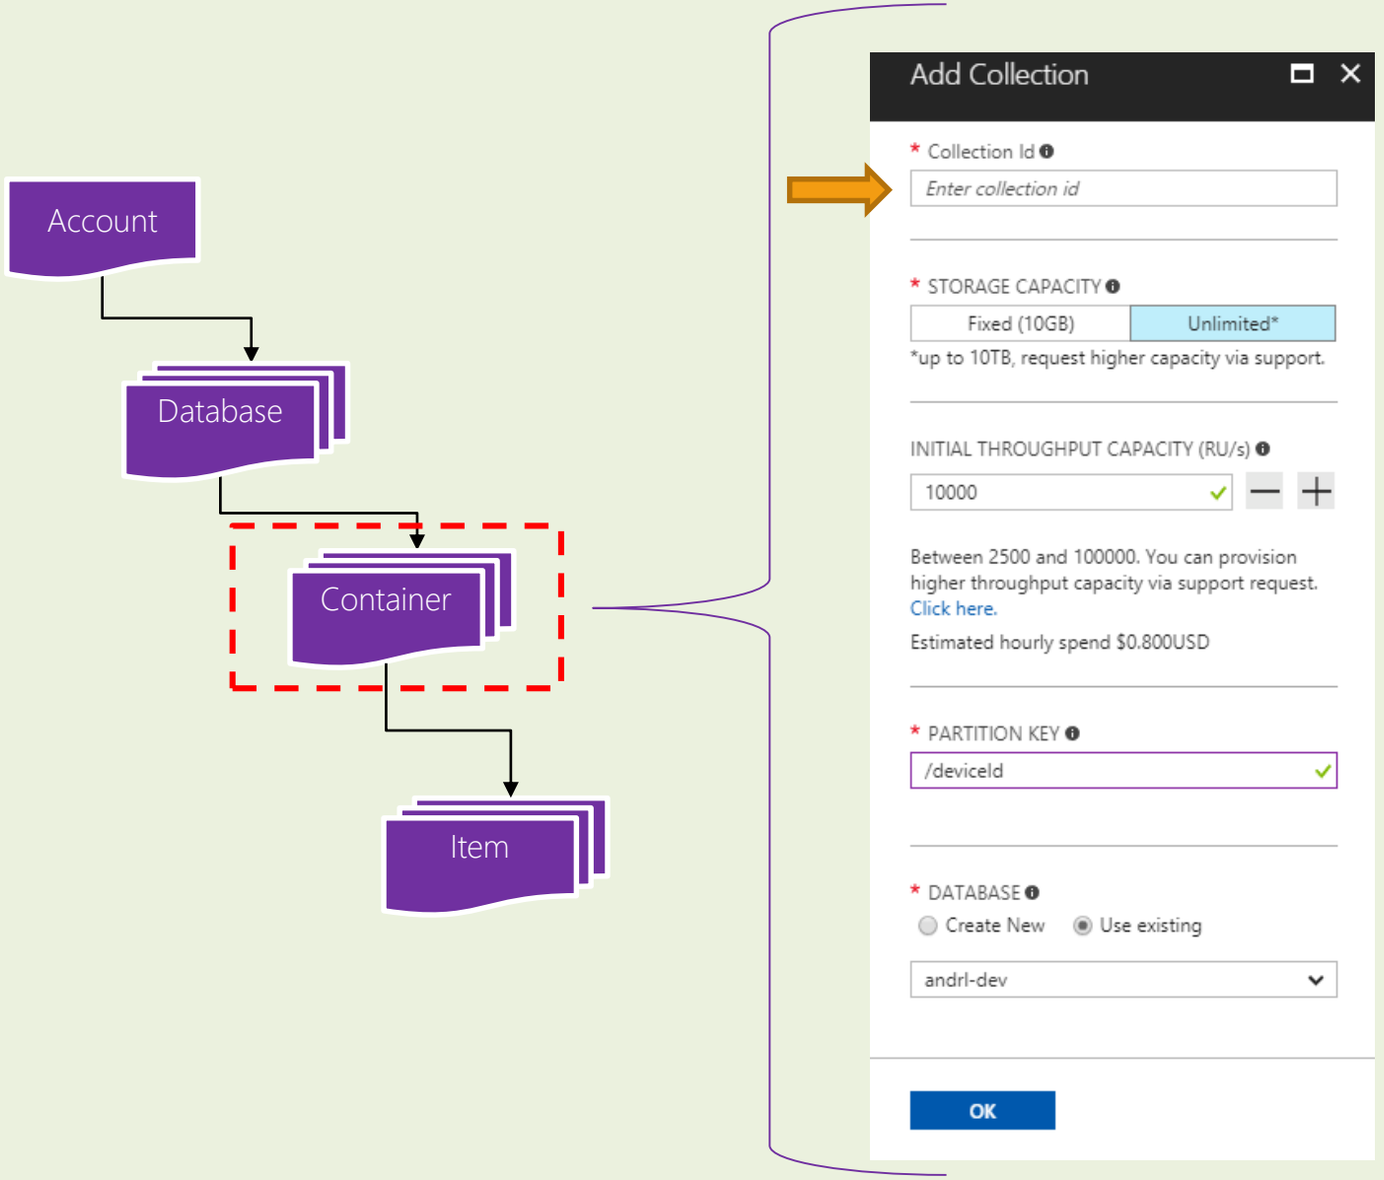
Account
Database
Container
Item
Add Collection
* Collection Id
Enter collection id
* STORAGE CAPACITY
Fixed (10GB)
Unlimited*
*up to 10TB, request higher capacity via support.
INITIAL THROUGHPUT CAPACITY (RU/s)
10000
Between 2500 and 100000. You can provision higher throughput capacity via support request.
Click here.
Estimated hourly spend $0.800USD
* PARTITION KEY
/deviceId
* DATABASE
Create New
Use existing
andrl-dev
OK

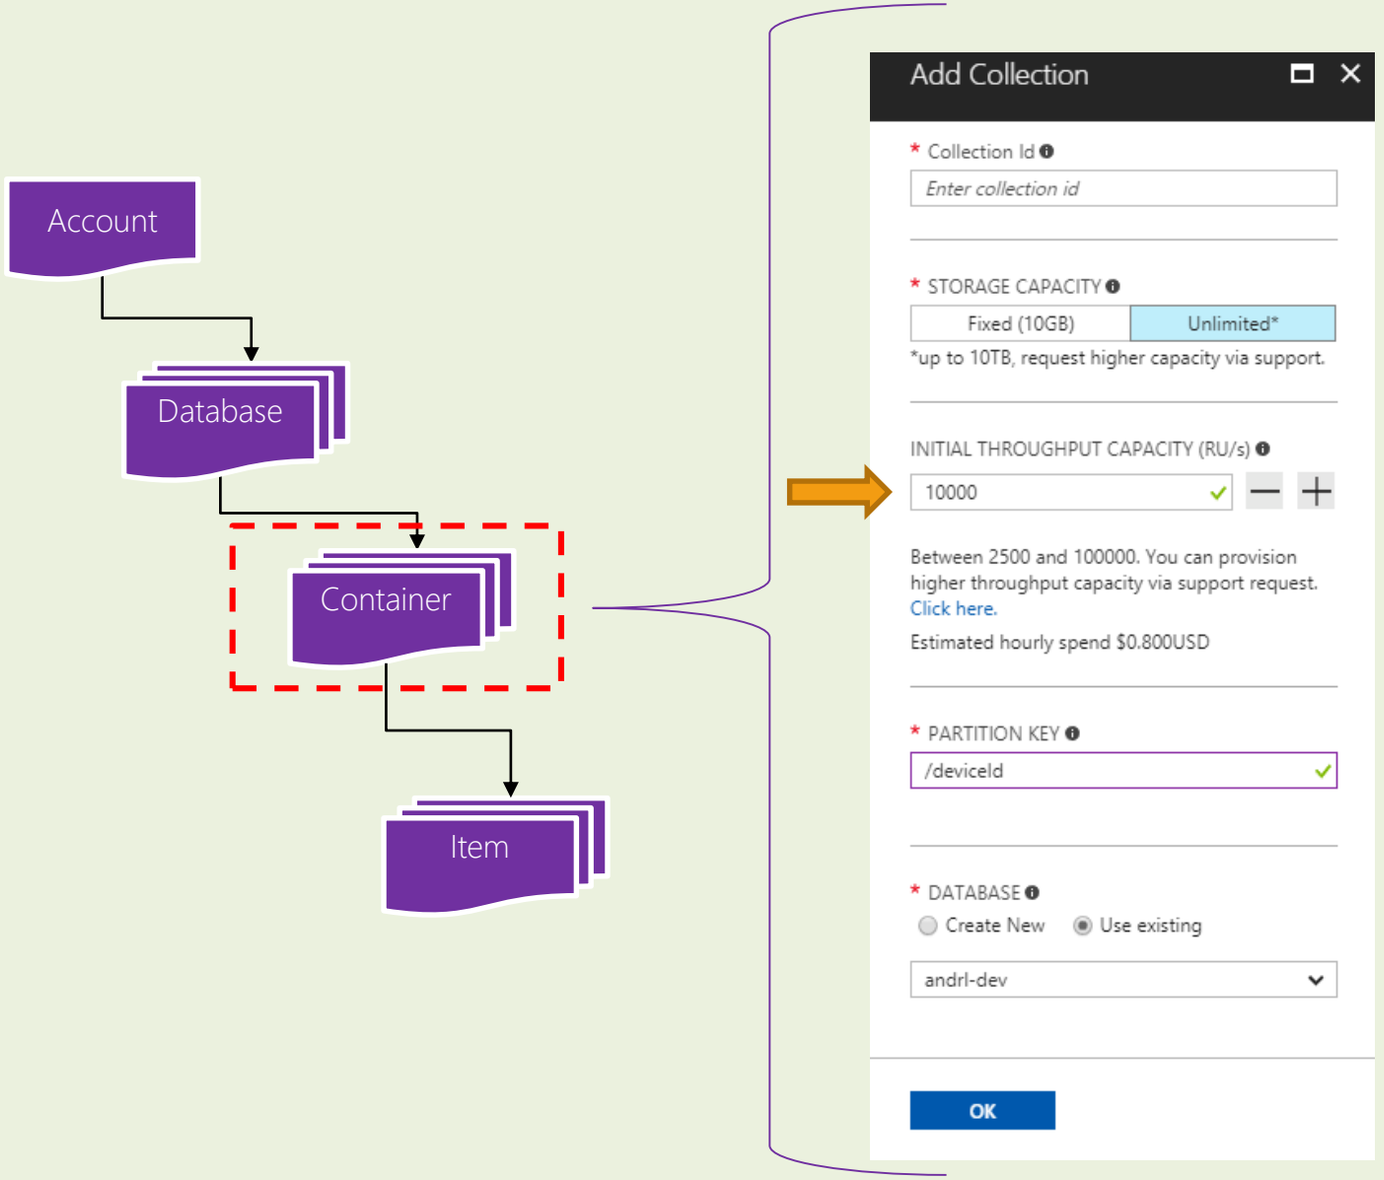
Account
Database
Container
Item
Add Collection
* Collection Id
Enter collection id
* STORAGE CAPACITY
Fixed (10GB)
Unlimited*
*up to 10TB, request higher capacity via support.
INITIAL THROUGHPUT CAPACITY (RU/s)
10000
Between 2500 and 100000. You can provision higher throughput capacity via support request.
Click here.
Estimated hourly spend $0.800USD
* PARTITION KEY
/deviceId
* DATABASE
Create New
Use existing
andrl-dev
OK

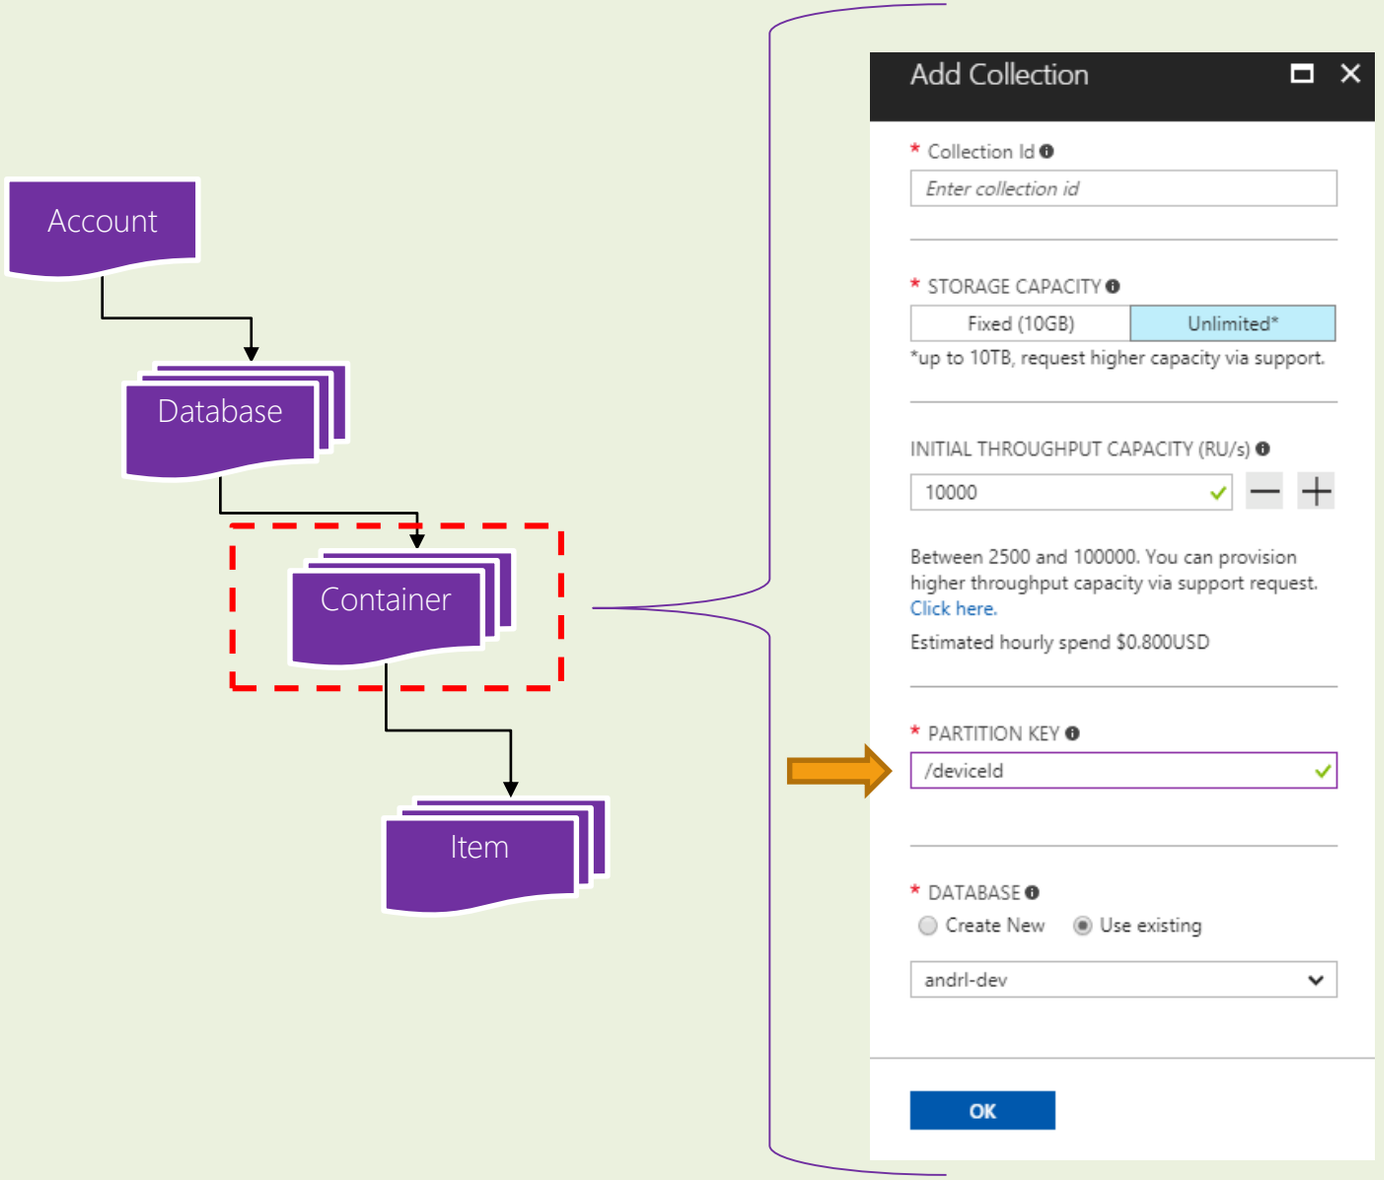
Account
Database
Container
Item
Add Collection
* Collection Id
Enter collection id
* STORAGE CAPACITY
Fixed (10GB)
Unlimited*
*up to 10TB, request higher capacity via support.
INITIAL THROUGHPUT CAPACITY (RU/s)
10000
Between 2500 and 100000. You can provision higher throughput capacity via support request.
Click here.
Estimated hourly spend $0.800USD
* PARTITION KEY
/deviceId
* DATABASE
Create New
Use existing
andrl-dev
OK

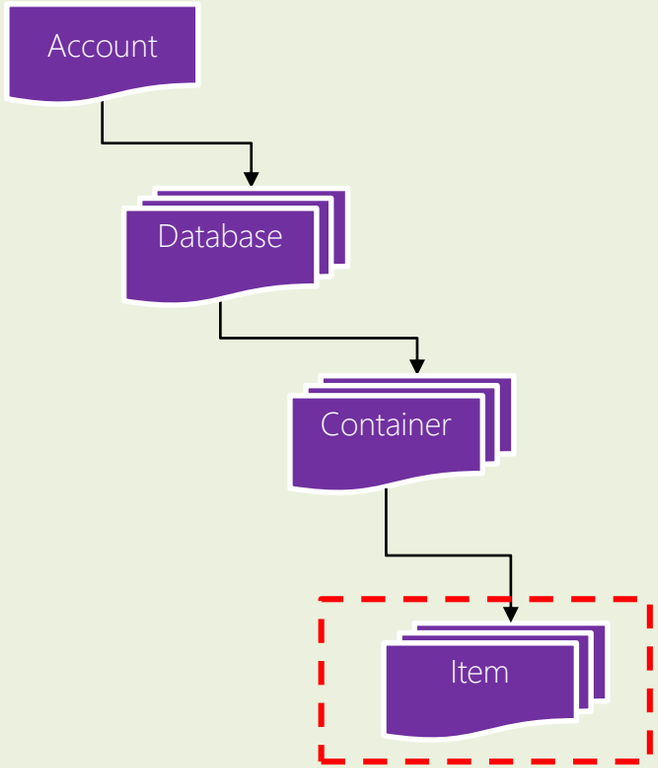
Account
Database
Container
Item

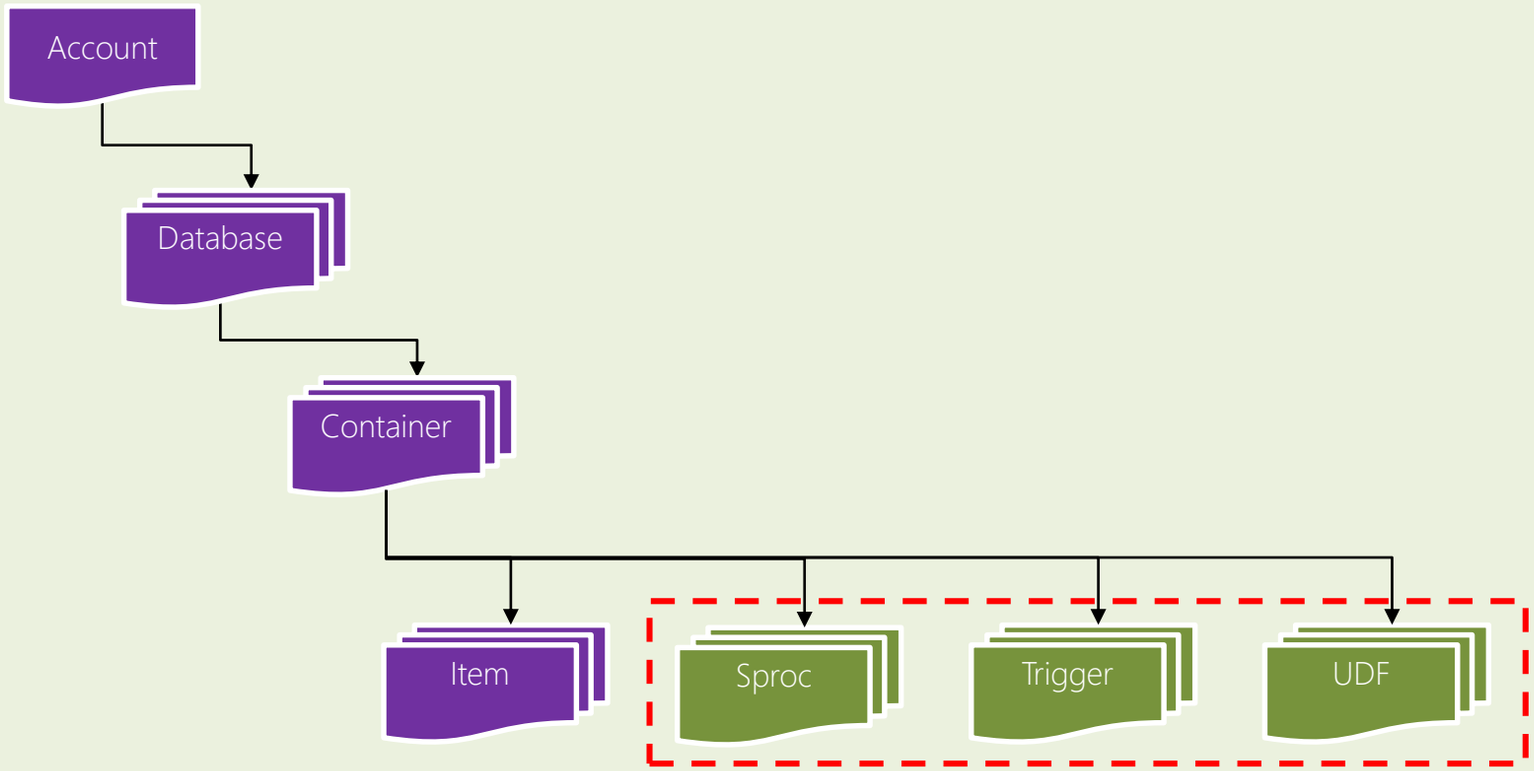
Account
Database
Container
Item
Sproc
Trigger
UDF

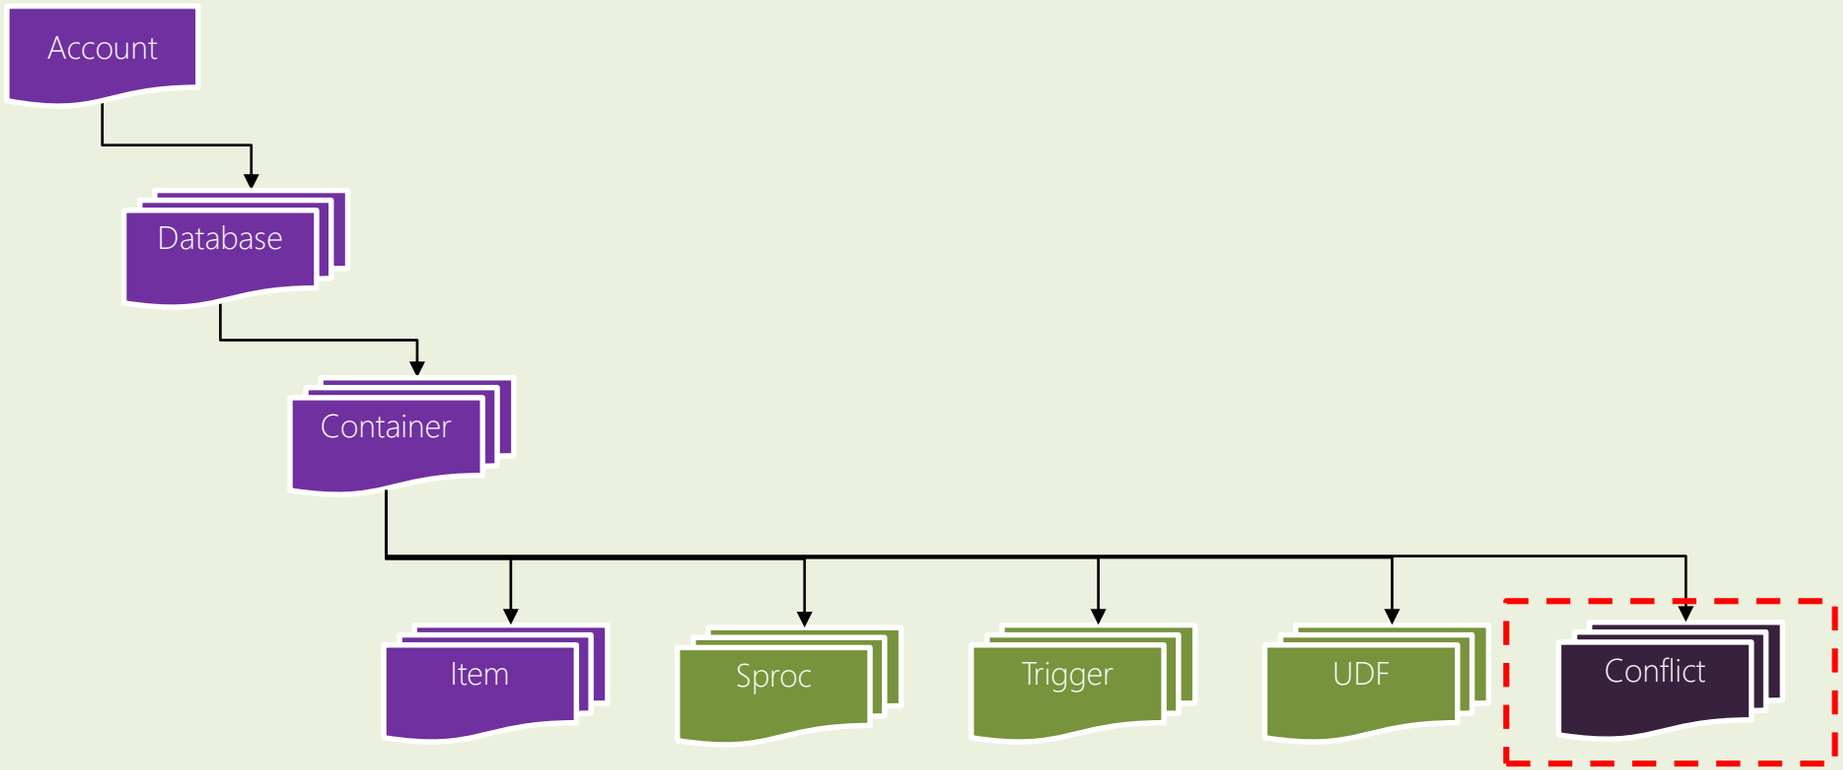
Account
Database
Container
Item
Sproc
Trigger
UDF
Conflict

5 Programmable Data Consistencies

Navigating CAP theorem
Consistent data worldwide

Brewer's CAP Theorem

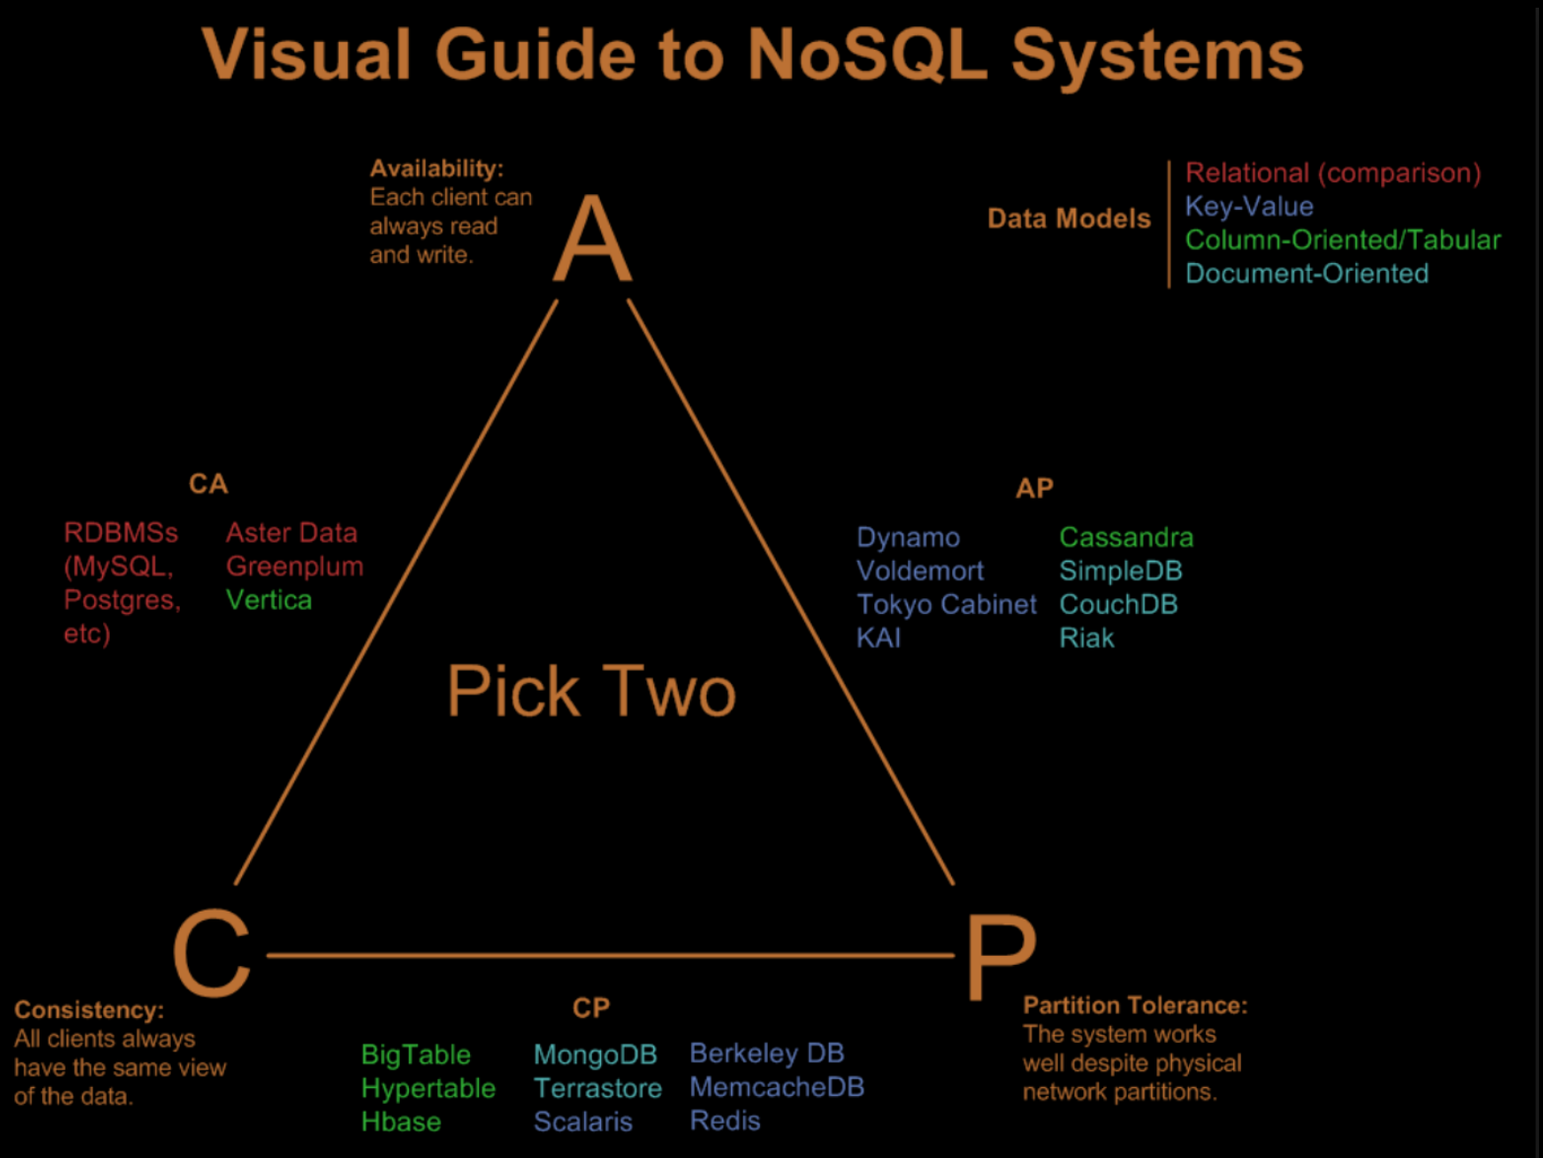

Programmable Data Consistency

- Databases are divided into two categories
  - Provide extreme choices – strong vs. eventual consistency (e.g., DynamoDB)
  - Leave everything for developers to configure (e.g., Cassandra)
    - Read repair, Hinted handoff, quorum sizes, replication topologies etc
    - Developers have to make precise tradeoffs between
      - Consistency and availability (during failures)
      - Consistency and latency (during steady state)
      - Consistency and throughput (this is important for TCO reasons)

Choices of Consistency

Most real-life applications do not fall into these two extremes

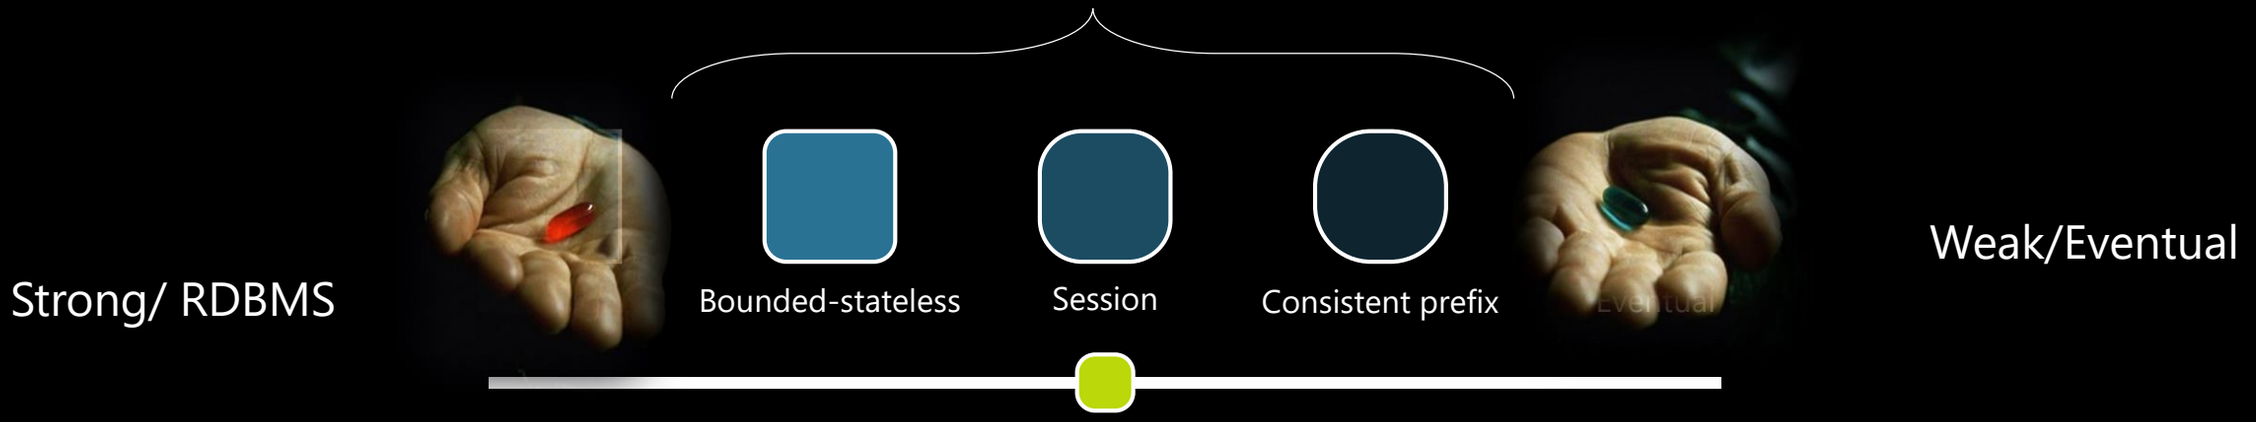

5 well-defined consistency levels for low latency and high availability

Azure Cosmos DB

5 well-defined consistency models

Clear Tradeoffs

- Latency
- Availability
- Throughput

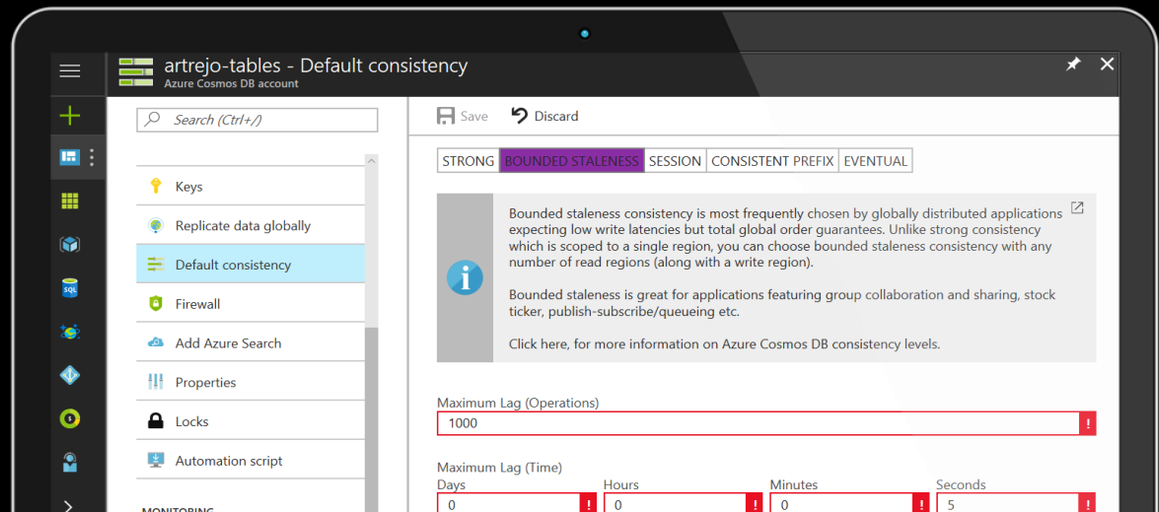

6 Industry-Leading, Comprehensive SLAs

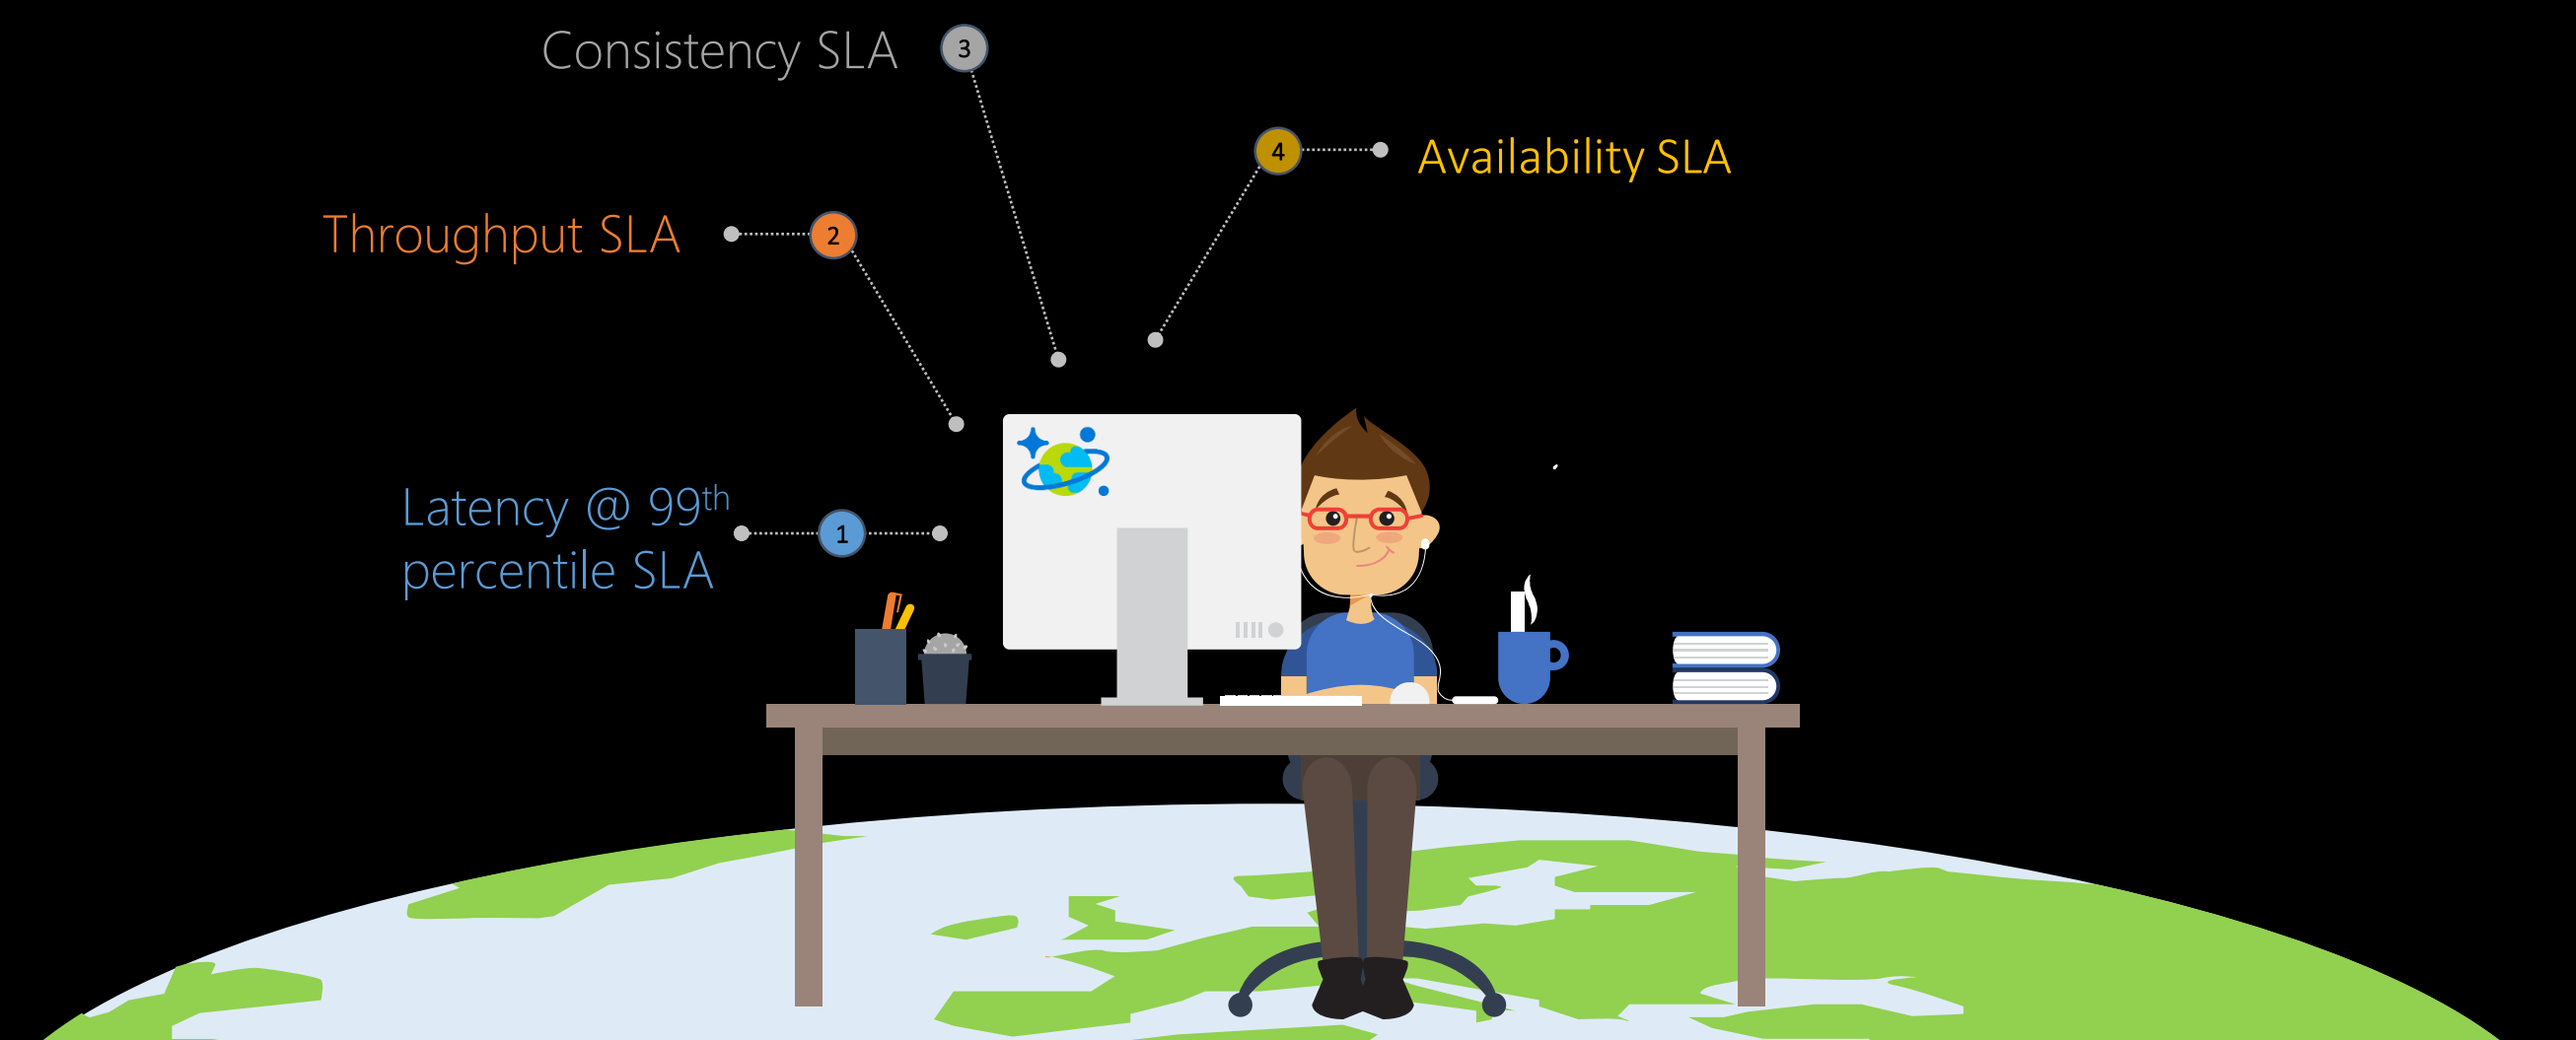

Azure Cosmos DB

Only database with comprehensive SLAs across 4 dimensions

High Availability

Performance Latency

Performance Throughput

Data Consistency

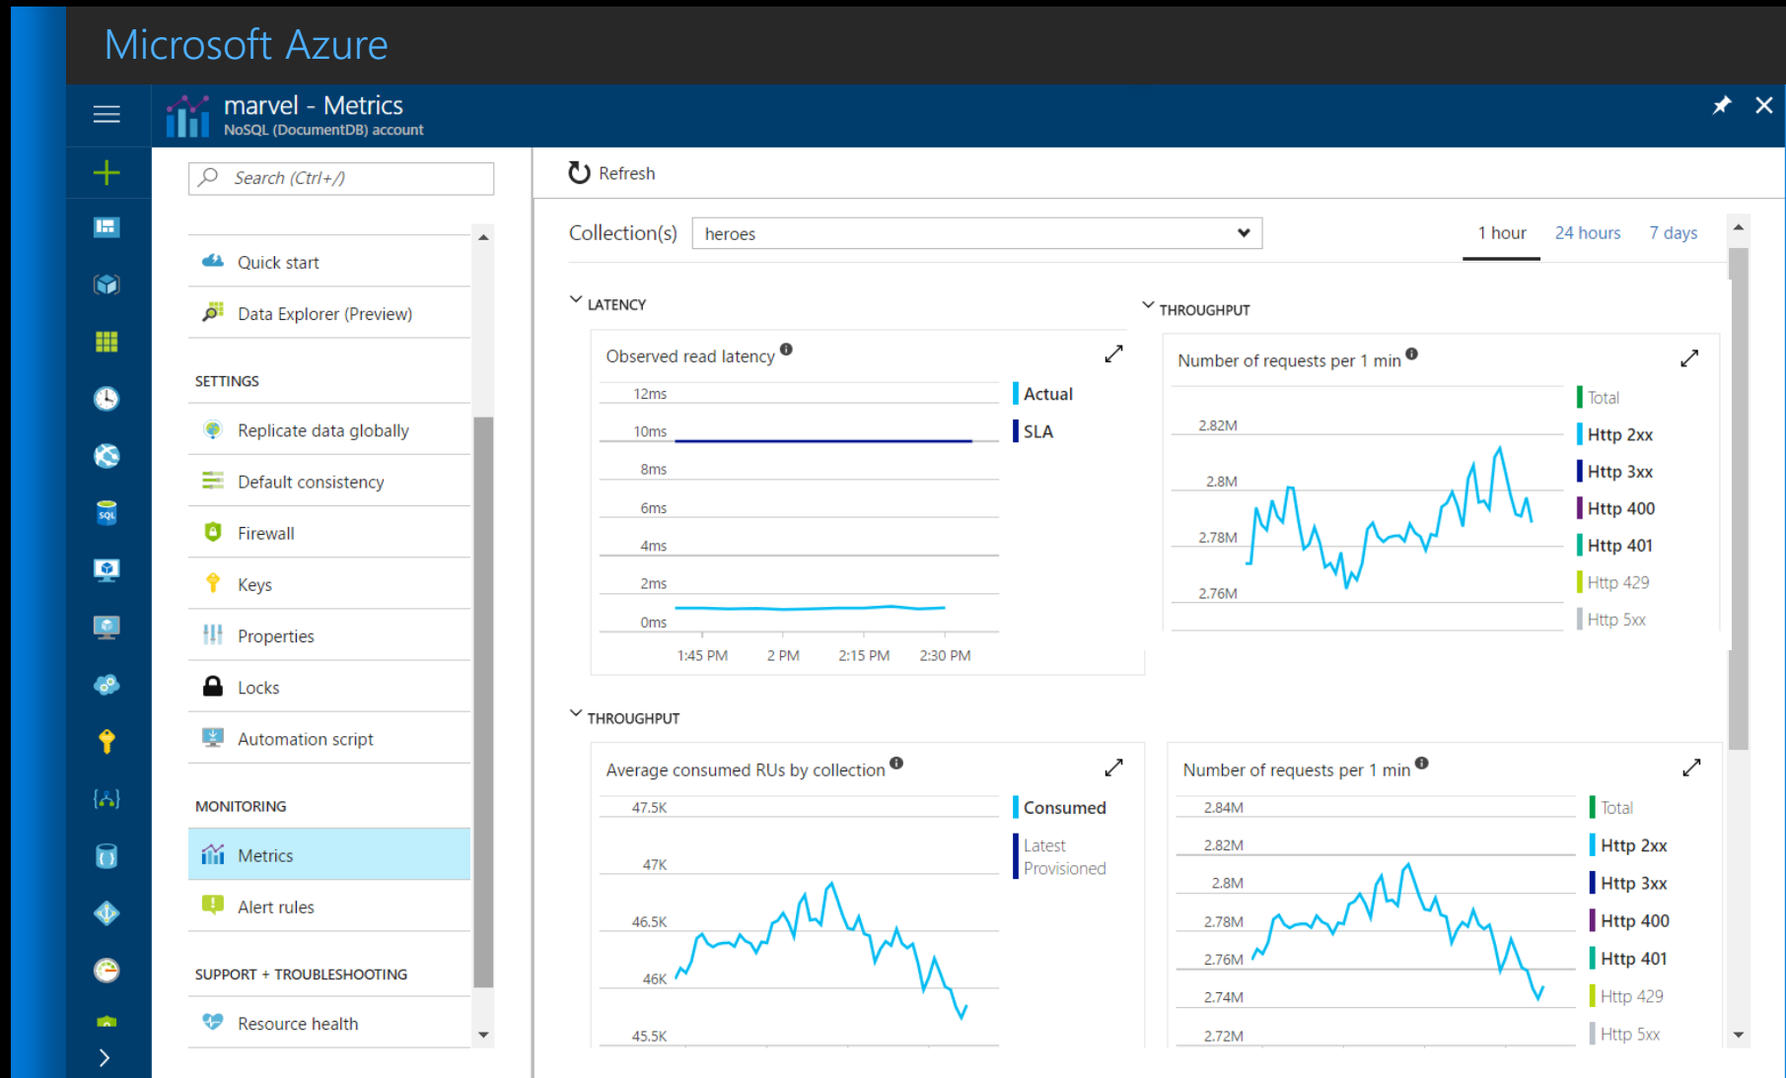

Demo – SLA

Schema-agnostic, automatic indexing

Schema-agnostic, automatic indexing

- At global scale, schema/index management is hard
- Automatic and synchronous indexing of all ingested content - hash, range, geo-spatial (Point, Polygon or Linestring)
  - No schemas or secondary indices ever needed
  - Choice of indexing mode – consistent, lazy, none (_id only)
  - Defined per collection – this is in addition to consistency for collection themselves
- Online and in-situ index transformations
- While the database is fully schema-agnostic, schema-extraction is built in
  - Customers can get Avro schemas from the database

Native Multi-Model

Why Multi-Model?

Who Wants to Store Their Data Like This?

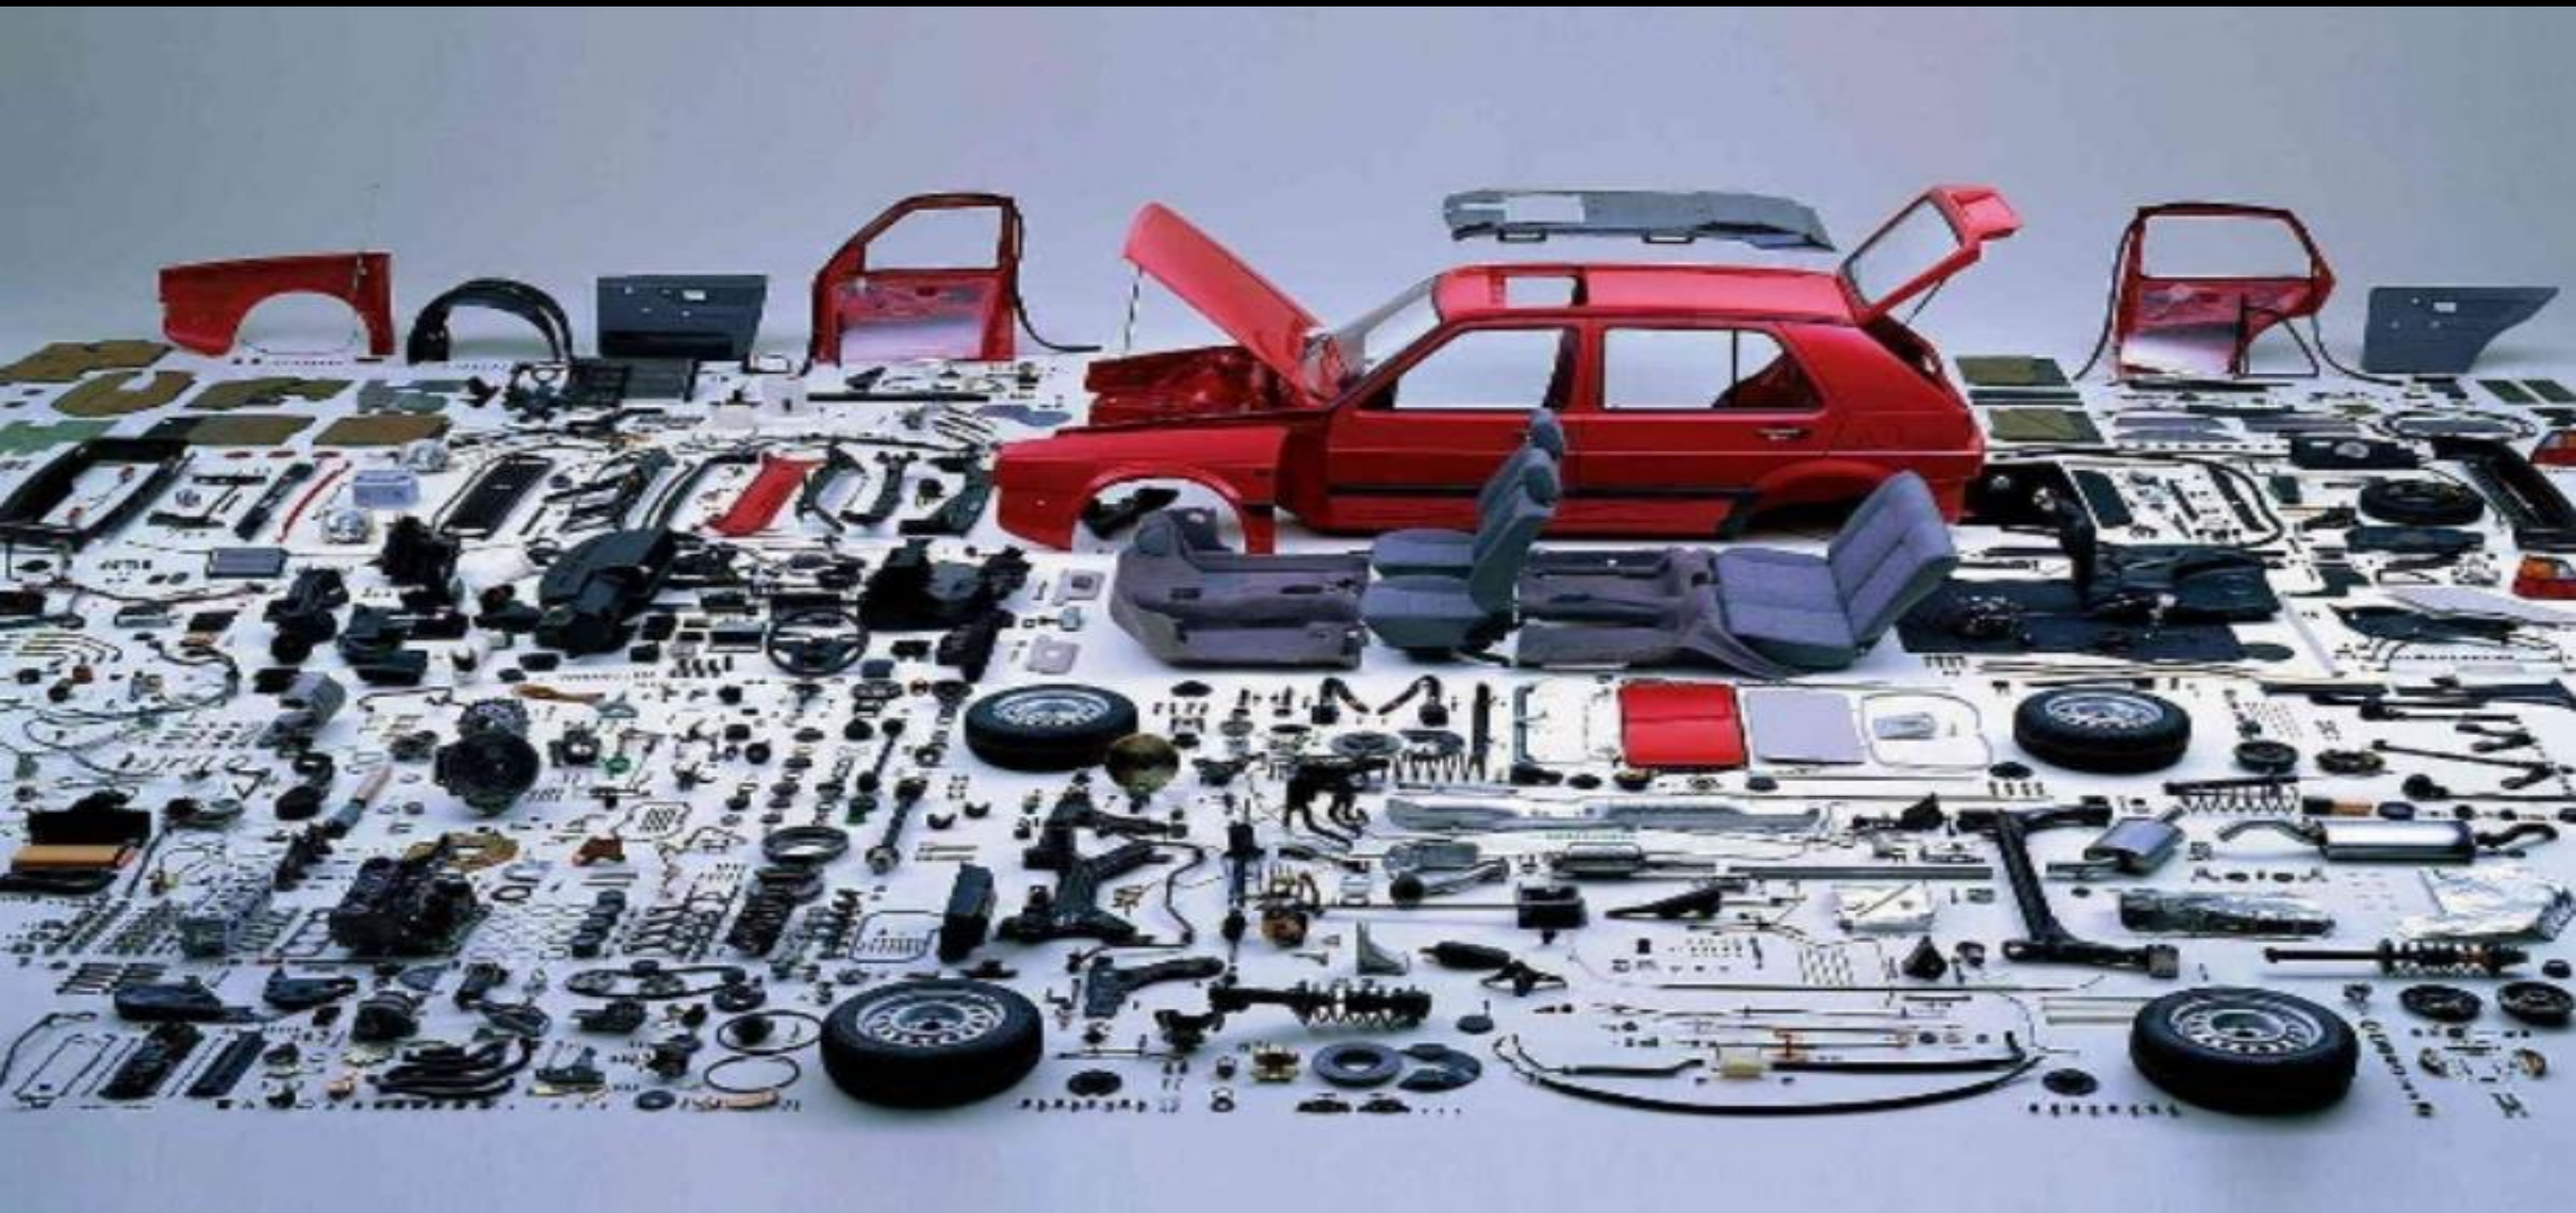

Why Multi-Model?

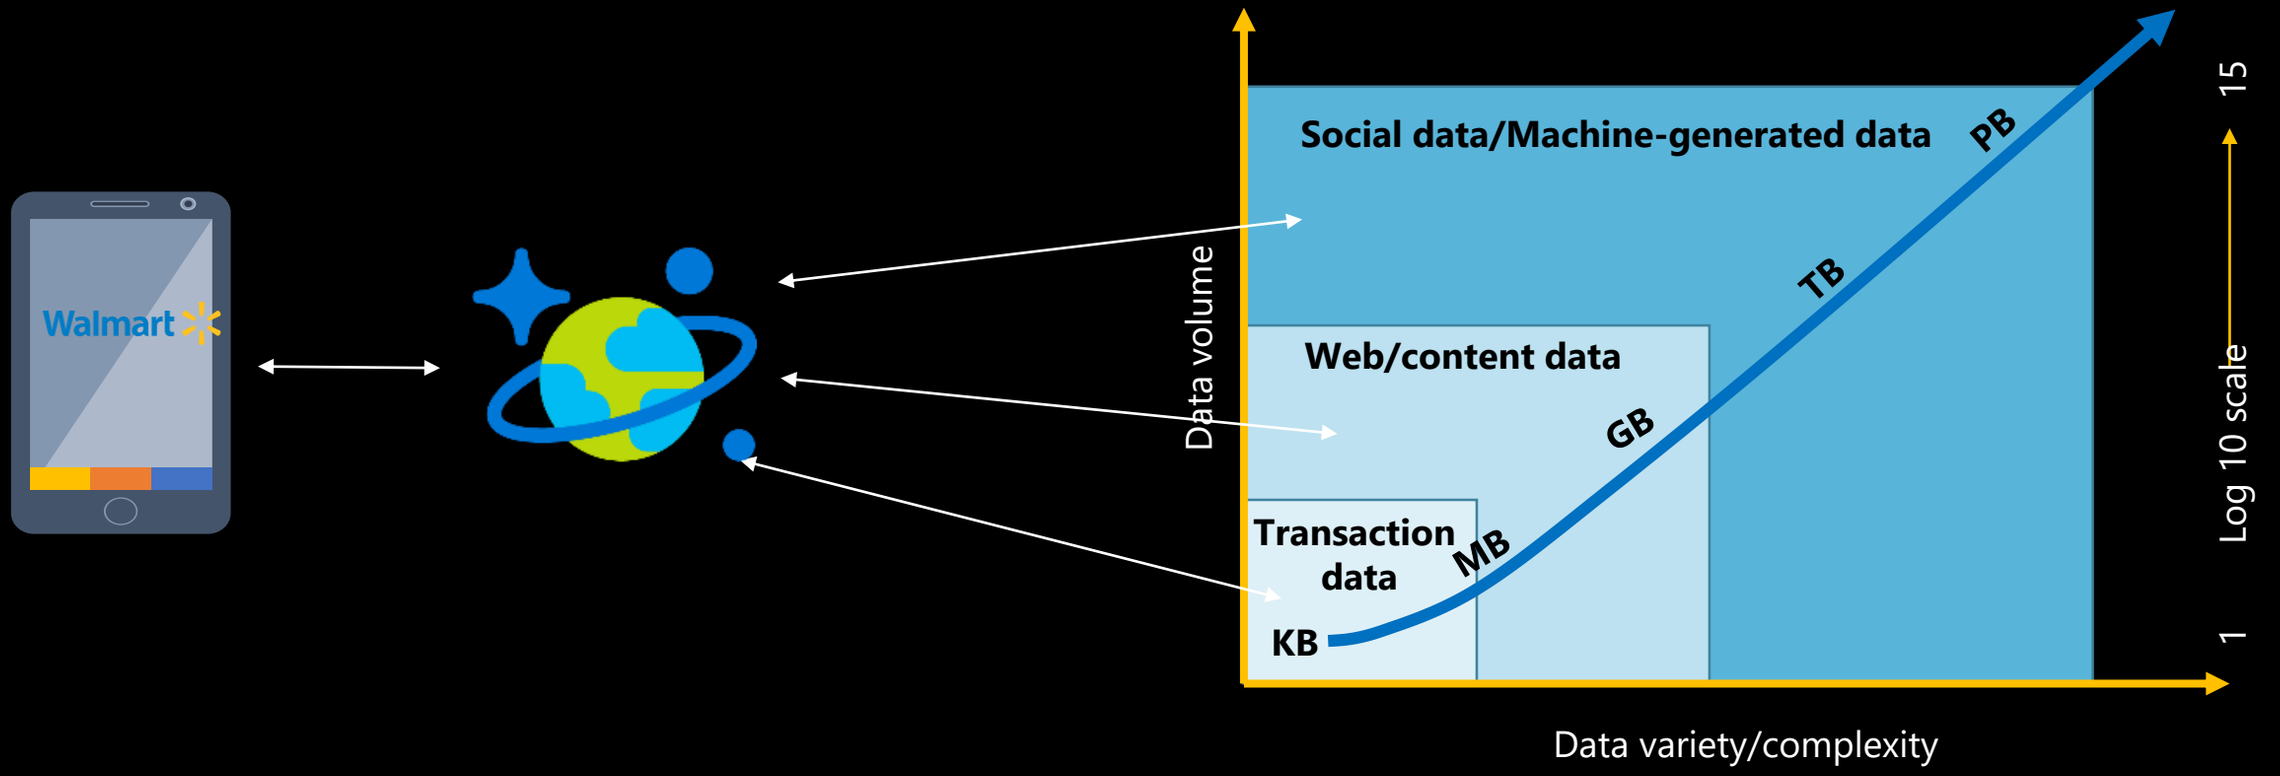

Who Wants to Have 3-5 Different Backend Databases?

Native Support for Multiple Data Models

- Database engine operates on atom-record-sequence (ARS) based type system
  - All data models are **efficiently** translated to ARS
- API and wire protocols are supported via extensible modules
- Instance of a given data model can be materialized as trees
- Graph, documents, key-value, column-family, … *more to come*

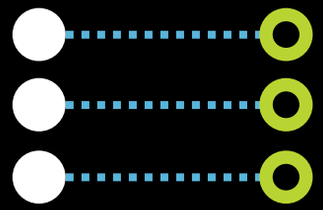

KEY-VALUE

COLUMN-FAMILY

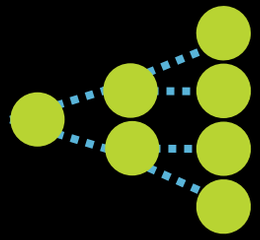

DOCUMENT

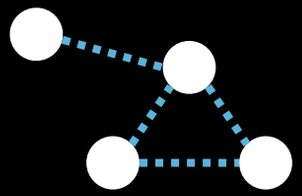

GRAPH

Native Multi-API

Native Support for Multiple APIs, formats & Wire Protocols

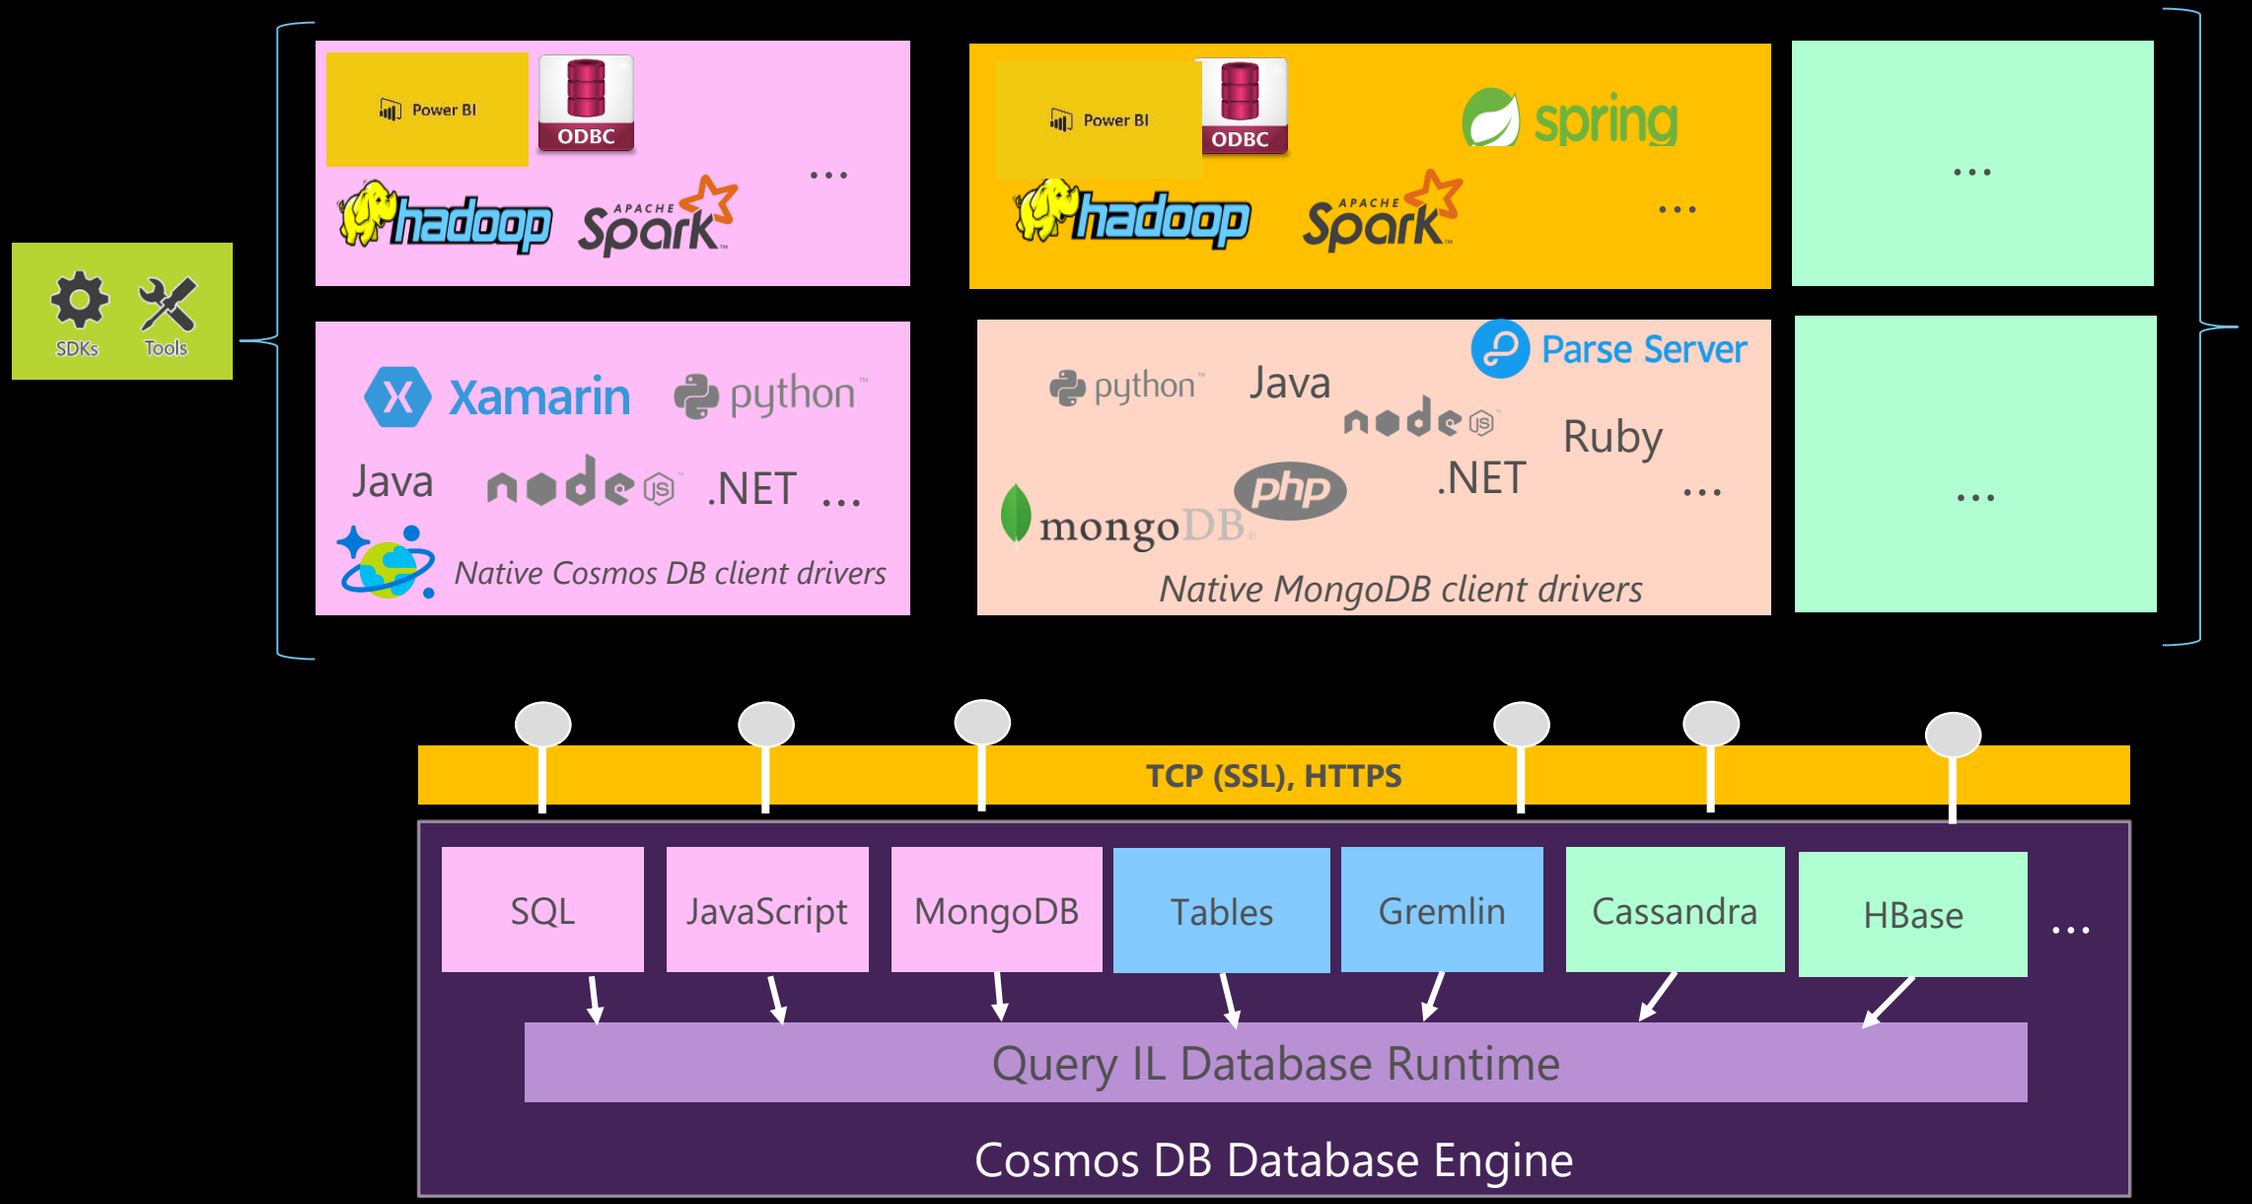

10 Security, Encryption, Compliance

Security & Compliance

Enterprise grade security

Encryption at Rest

- Always encrypted at rest and in motion
- Data, index, backups, and attachments encrypted

Encryption is enabled automatically by default

- No impact on performance, throughput or availability
- Transparent to your application

Comprehensive Azure compliance certification

- ISO 27001, ISO 27018, EUMC, HIPAA, PCI
- SOC1 and SOC2 (Audit complete, Certification in Q2 2017)
- FedRAMP, IRS 1075, UK Official (IL2) (Q2 2017)
- HITRUST (H2 2017)

11

Resource Governance -RU Estimator

Cosmos DB Compliance

Office 365 Compliance	Status
Level A	No mining of customer data ✔ No voluntary disclosure to law enforcement ✔
Level B	+ISO 27001, ISO 27018, EUMC ✔
Level C	+HIPAA ✔ +SOC1 and SOC2 (Audit complete, Certification in March 2017) +PCI (Audit complete; Certification in June 2017)
Level D	+Data residency requirements ✔ +GovCloud (April 2017) ✔ +FedRAMP, IRS 1075, UK Official (IL2) (June 2017) +HITRUST (H2 2017)

- Encryption at rest and PCI-DSS compliant since 2017

Demo – RU and Cost Estimate


Thank You

Developing planet-scale apps comes with planet-scale challenges

- ⚠ **Write accurate, globally distributed apps**
- ⚠ **Managing and versioning complex schemas**
- ⚠ **Scaling both throughput and storage based on global demand**
- ⚠ **Balancing the needs for strong and eventual consistency**
- ⚠ **Delivering highly-responsive experiences**
- ⚠ **Ensuring an always-on system**

Operational Workloads

Versatile Workloads

Analytical Workloads

mongoDB

Gremlin

$G = (V, E)$

Multi-API

Azure Cosmos DB

SQL

ANSI SQL

APACHE HBASE

DynamoDB

Multi-Model

Key-Value

Tabular

Documents

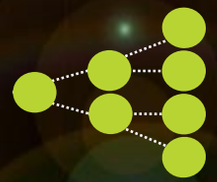

Graph

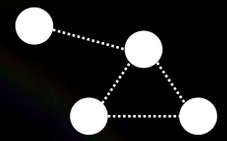

Relational

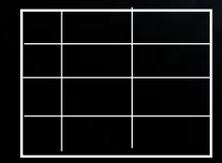

Planet-Scale

Global Distribution from the ground-up

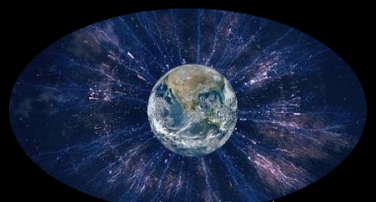

Limitless Scale

Extremely Low Latency

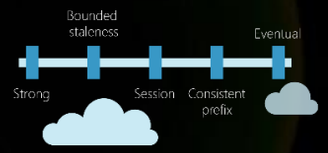

Multiple Consistency Levels

ARS model

Comprehensive SLAs

Customers using Cosmos DB already today

Common Architecture Patterns

Content Management Systems

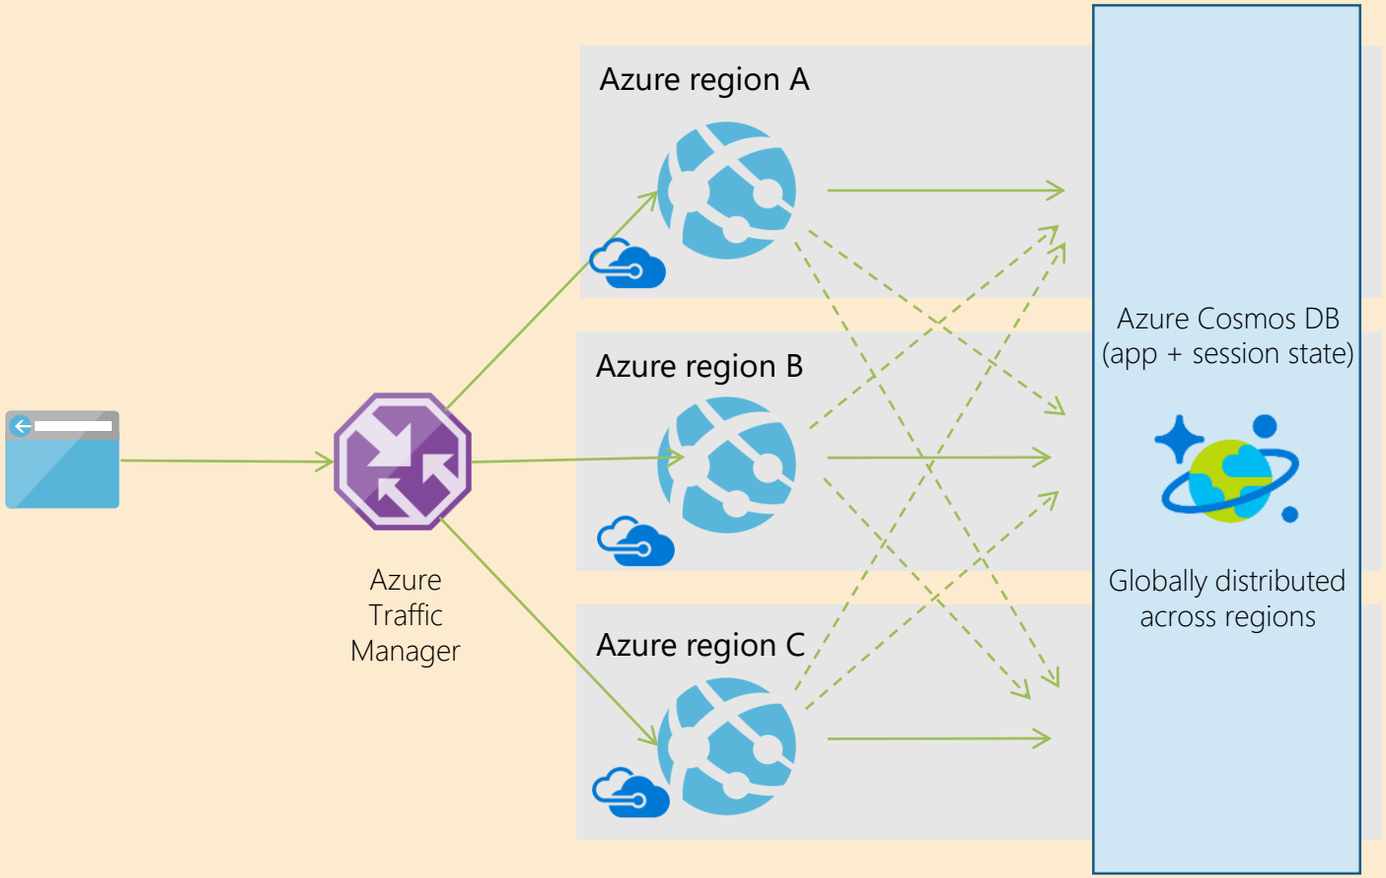

Internet of Things – Telemetry & Sensor Data

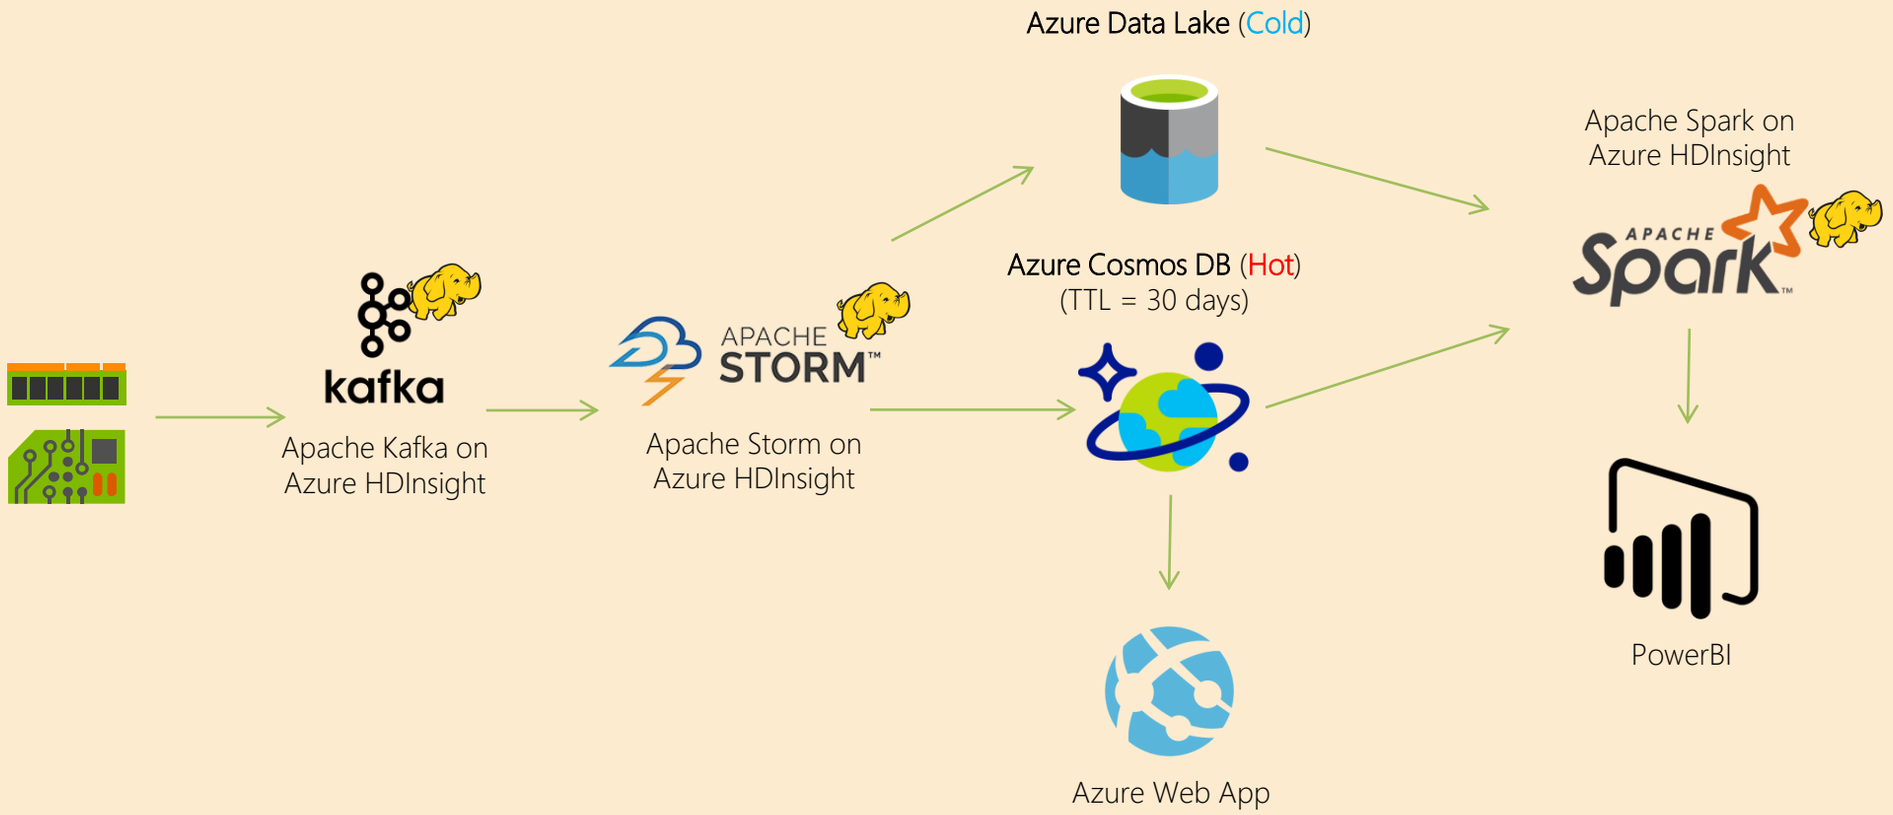

Retail Product Catalogs

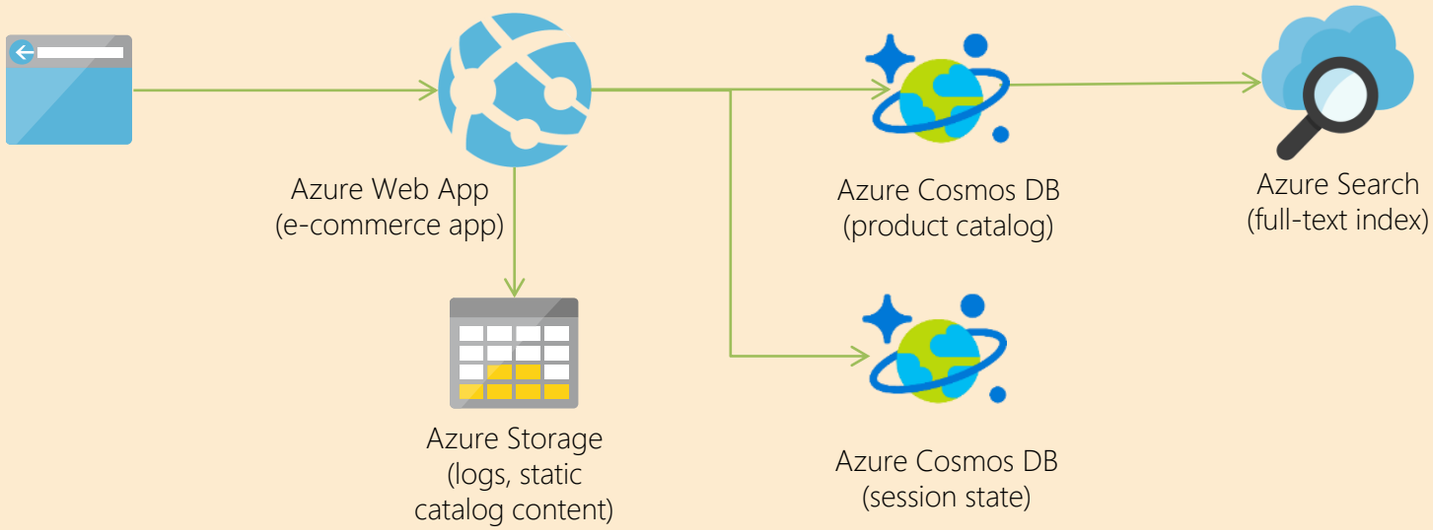

Retail Order Processing Pipelines

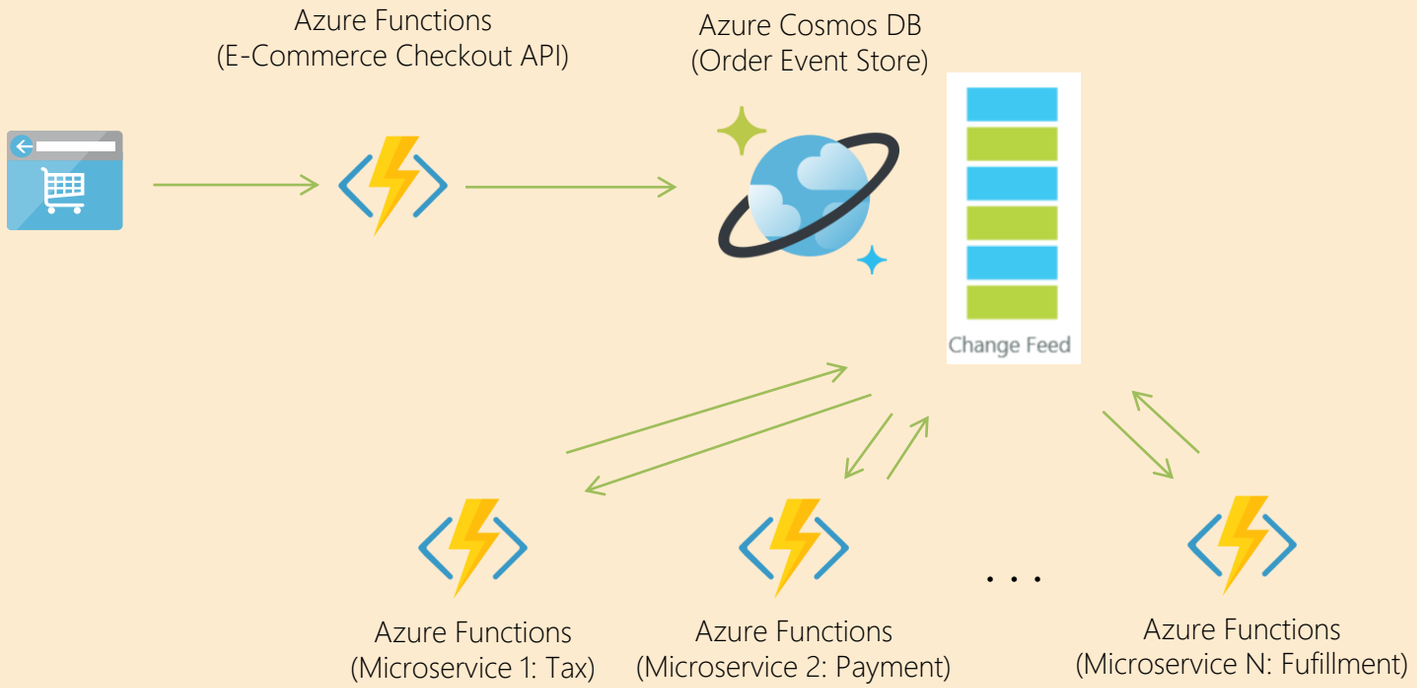

Real-time Recommendations

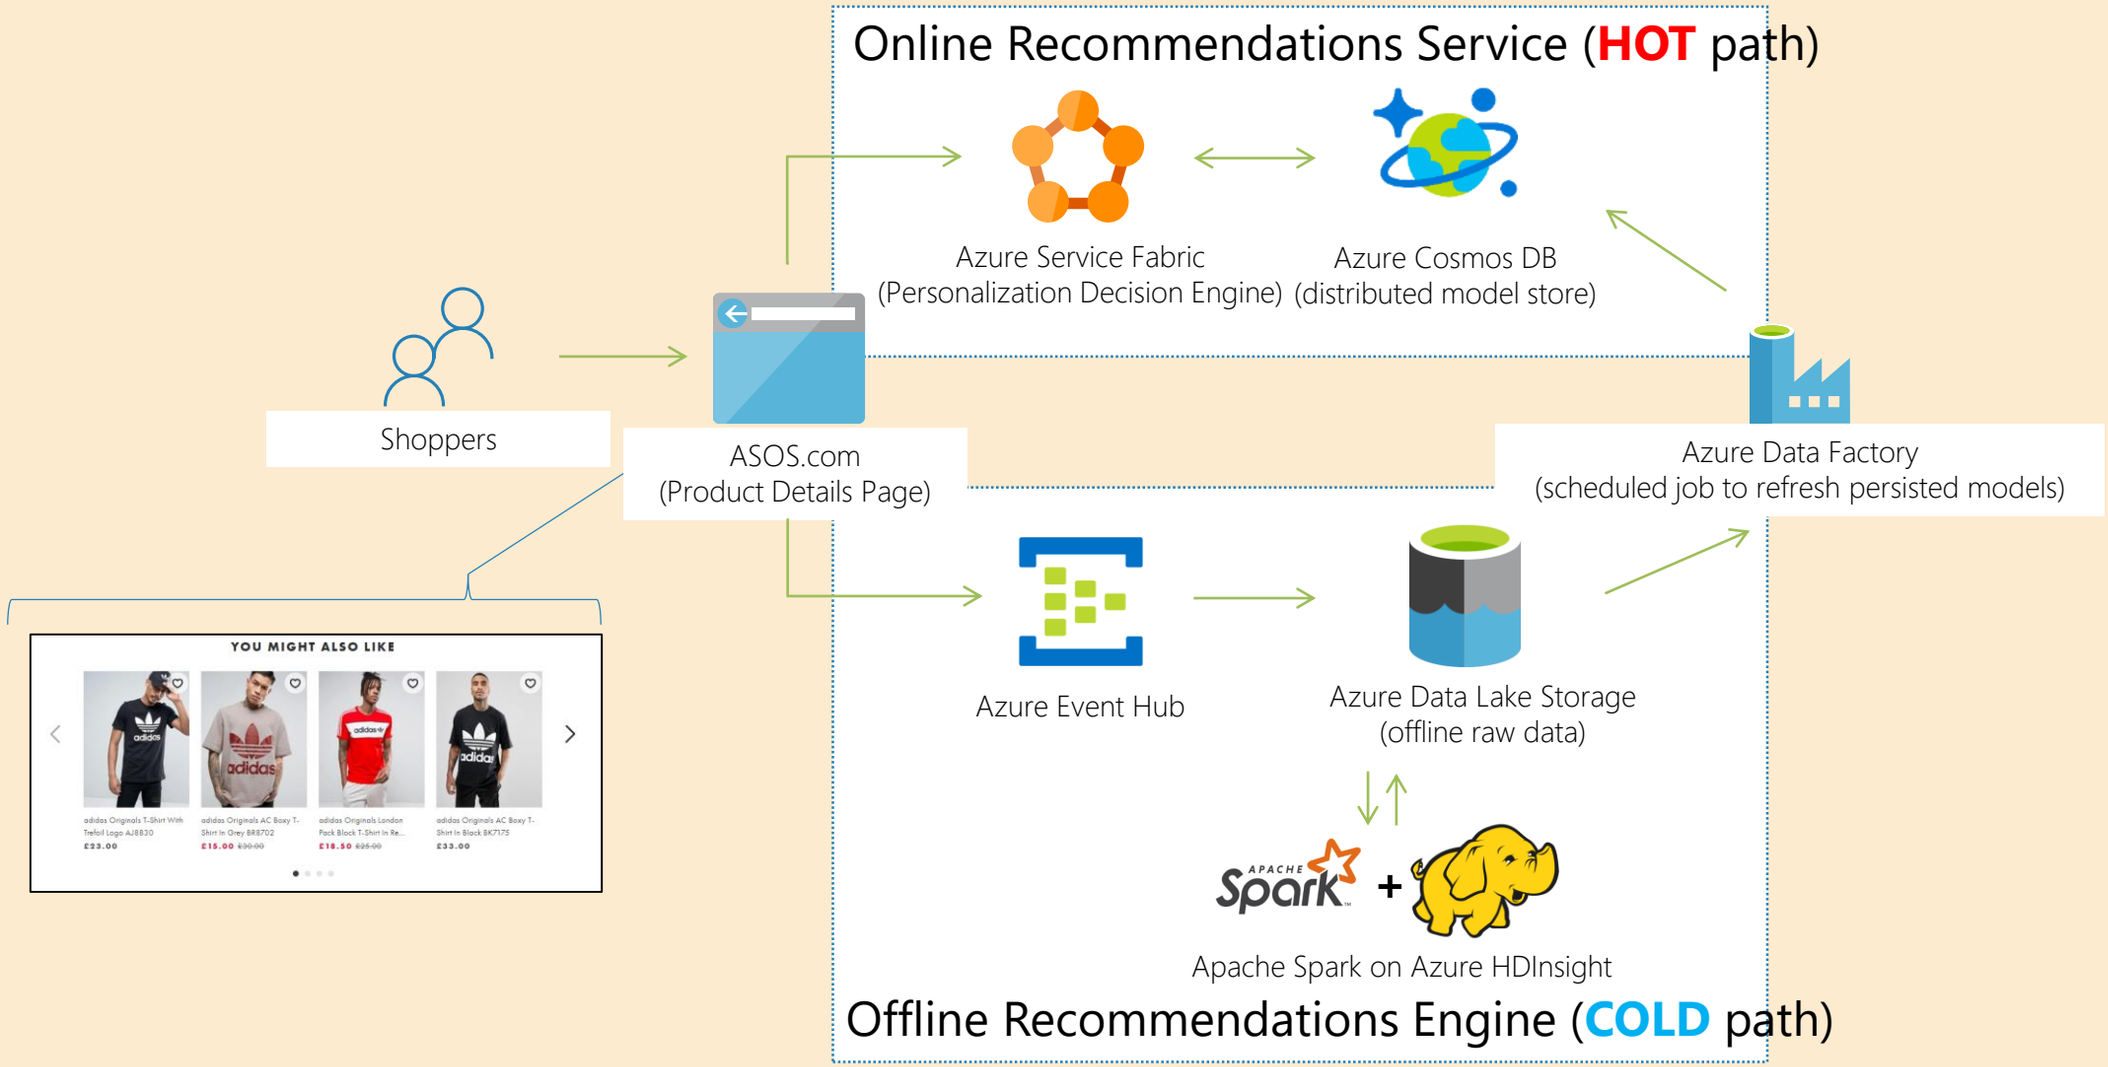

Multiplayer Gaming

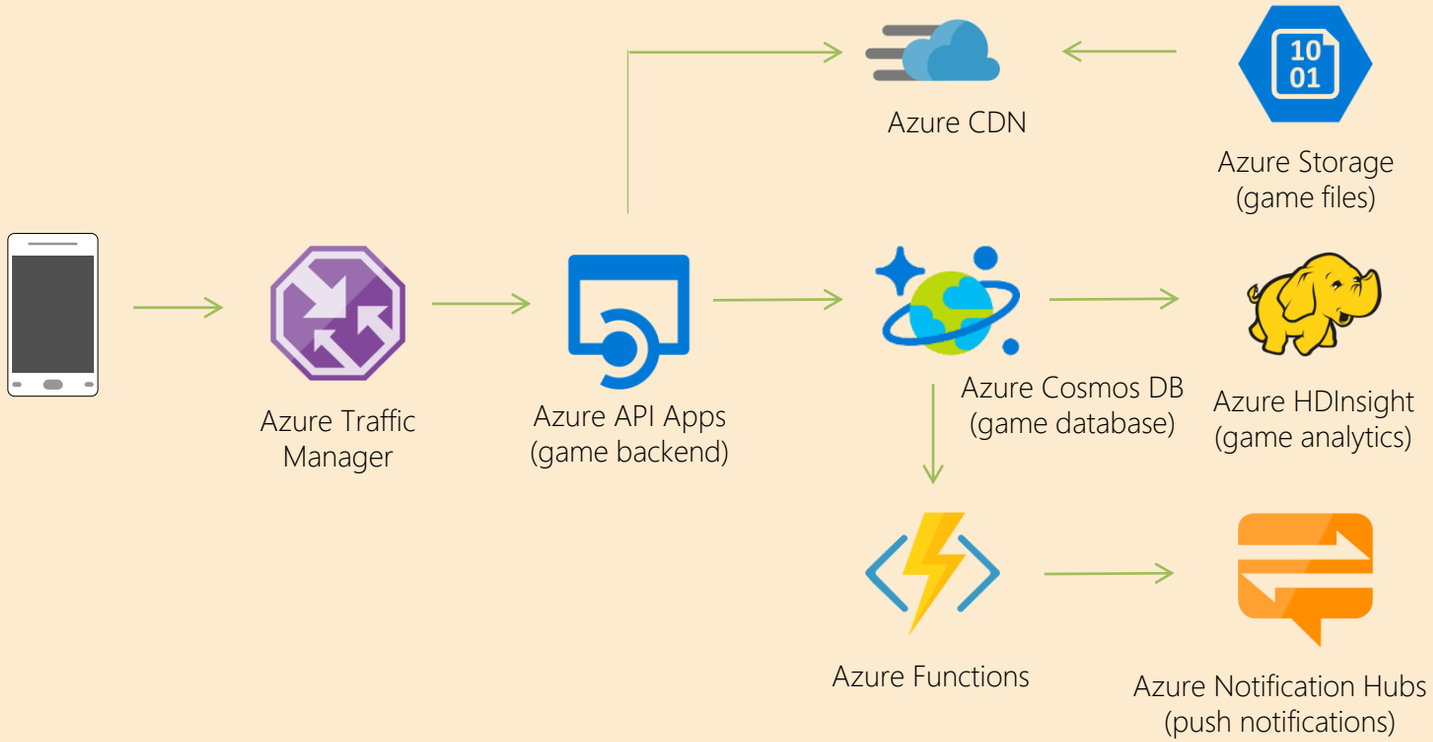

Customer 360 + Social Analytics

Messaging

Azure Web App
(messaging app)

Week 0

Azure Cosmos DB
(message data – partitioned by userid)

Azure Search
(full-text index)

Messaging

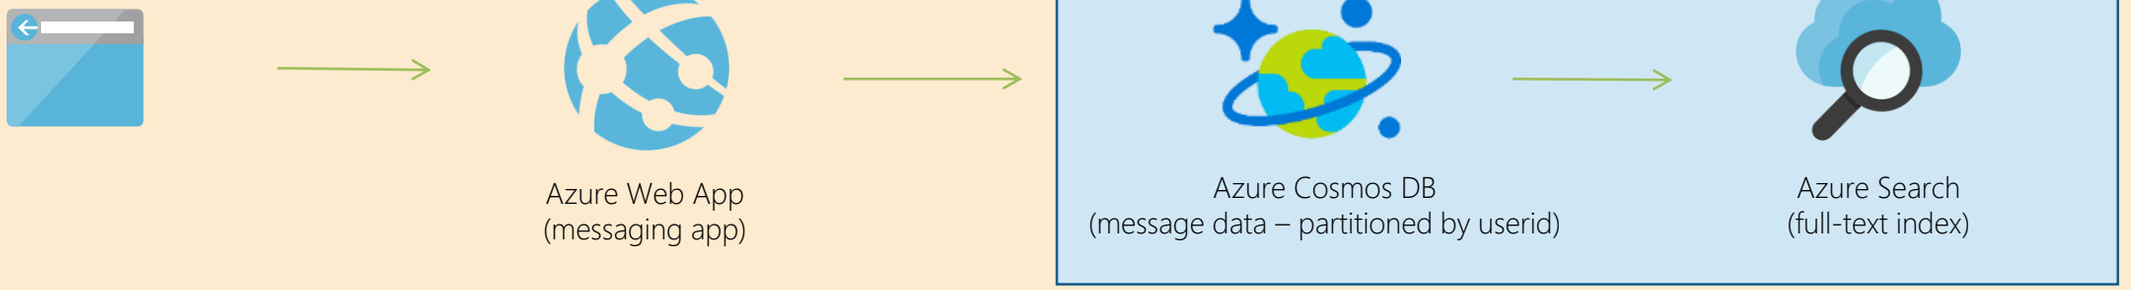
Week 0
Azure Cosmos DB
(message data – partitioned by userid)
Azure Search
(full-text index)
Azure Web App
(messaging app)
Week 1
Azure Cosmos DB
(message data – partitioned by userid)
Azure Search
(full-text index)

Messaging

Azure Web App
(messaging app)

Week 0

Azure Cosmos DB
(message data – partitioned by userid)

Azure Search
(full-text index)

Week 1

Azure Cosmos DB
(message data – partitioned by userid)

Azure Search
(full-text index)

Week 2

Azure Cosmos DB
(message data – partitioned by userid)

Azure Search
(full-text index)

Azure Cosmos DB

Other stellar capabilities

Getting Started

- Web
  - cosmosdb.com
  - portal.azure.com
  - aka.ms/cosmosdb
  - aka.ms/cosmosdb-Tables
  - aka.ms/cosmosdb-Graph
  - aka.ms/cosmosdb-MongoDB
  - aka.ms/cosmosdb-DocumentDB
  - cosmosdb.com/capacityplanner
- Download
  - aka.ms/CosmosDB-emulator
- Re-visit Build session recordings on Channel 9.
- Read, comment, like, share the blog

#MSBuild
#CosmosDB