BUILDING A .NET APPLICATION USING AZURE COSMOS DB

Who is Chad Green

Director of IT Architecture
Atria Senior Living / Glennis Solutions

Microsoft® Most Valuable Professional

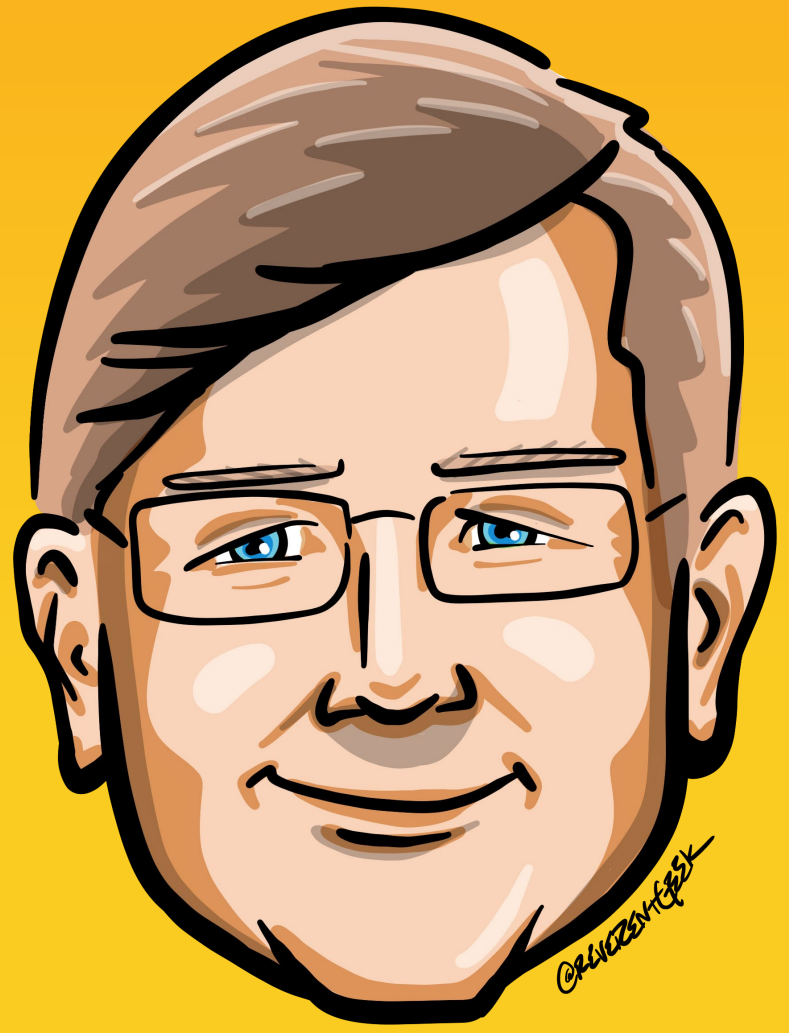

chadgreen@chadgreen.com

TaleLearnCode

ChadGreen.com

ChadGreen & TaleLearnCode

ChadwickEGreen

WHAT IS COSMOS DB

Azure Cosmos DB

What is Cosmos DB

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

Azure Cosmos DB

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

Turnkey global distribution

Azure Cosmos DB

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

Turnkey global distribution

Azure Cosmos DB

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

Elastic scale out of storage & throughput

Turnkey global distribution

Azure Cosmos DB

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

Elastic scale out of storage & throughput

Turnkey global distribution

Azure Cosmos DB

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

Guaranteed low latency at the 99th percentile

Elastic scale out
of storage & throughput

Turnkey global
distribution

Azure Cosmos DB

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

Guaranteed low latency at the 99th percentile

Elastic scale out
of storage & throughput

Turnkey global
distribution

Azure Cosmos DB

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

Five well-defined consistency models

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Turnkey global distribution

Azure Cosmos DB

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

Five well-defined consistency models

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Turnkey global distribution

Azure Cosmos DB

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

Comprehensive SLAs

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Five well-defined consistency models

Turnkey global distribution

Azure Cosmos DB

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

Comprehensive SLAs

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Five well-defined consistency models

Turnkey global distribution

Azure Cosmos DB

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

Battle Tested

XBOX LIVE

Office 365

skype™

Guaranteed low latency
at the 99th percentile

Elastic scale out
of storage & throughput

Five well-defined
consistency models

Turnkey global
distribution

Comprehensive
SLAS

Azure Cosmos DB

Battle Tested

Coca-Cola · AVEVA Schneider Electric · Albertsons SAFEWAY · Johnson Controls · SIEMENS Healthineers · BMT

Liberty Mutual INSURANCE · jet · Rolls-Royce · ExxonMobil · Cincinnati Children's

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Five well-defined consistency models

Turnkey global distribution

Comprehensive SLAS

Azure Cosmos DB

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

Ubiquitous Regional Presence

Turnkey global distribution

rehensive
LAS

Azure Cosmos DB

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

Secure by default and enterprise ready

Guaranteed low latency at the 99th percentile

Elastic scale out of storage & throughput

Five well-defined consistency models

Turnkey global distribution

Comprehensive SLAS

Azure Cosmos DB

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

Key-value

Column-family

Document

Graph

Guaranteed low latency
at the 99th percentile

Elastic scale out
of storage & throughput

Five well-defined
consistency models

Turnkey global
distribution

Comprehensive
SLAs

COSMOS DB
ARCHITECTURE

Elements in an Azure Cosmos DB Account

- Provision Azure Cosmos DB Account
- Create database in that account
- Add containers on those databases
- Container can be realized based upon the data API

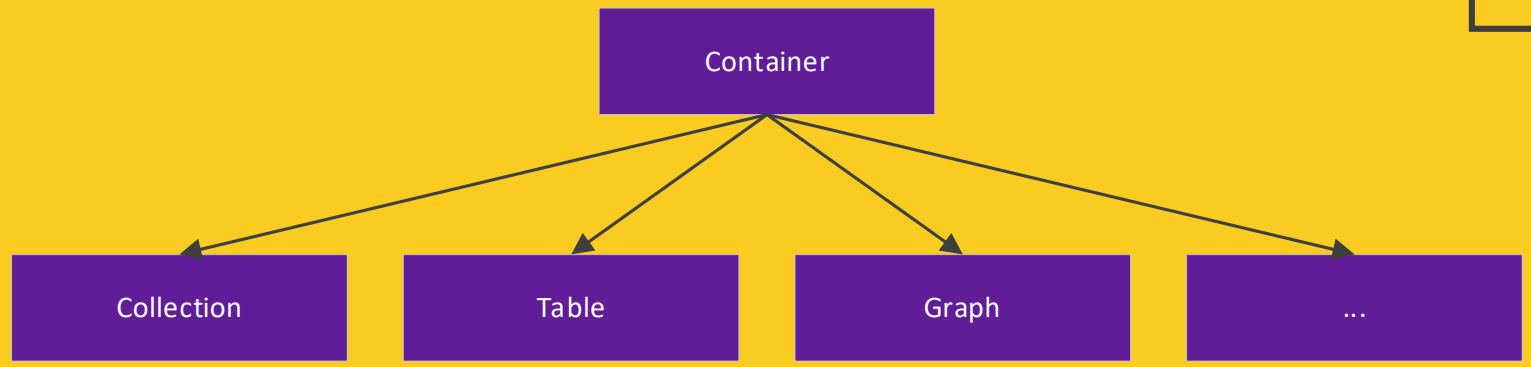

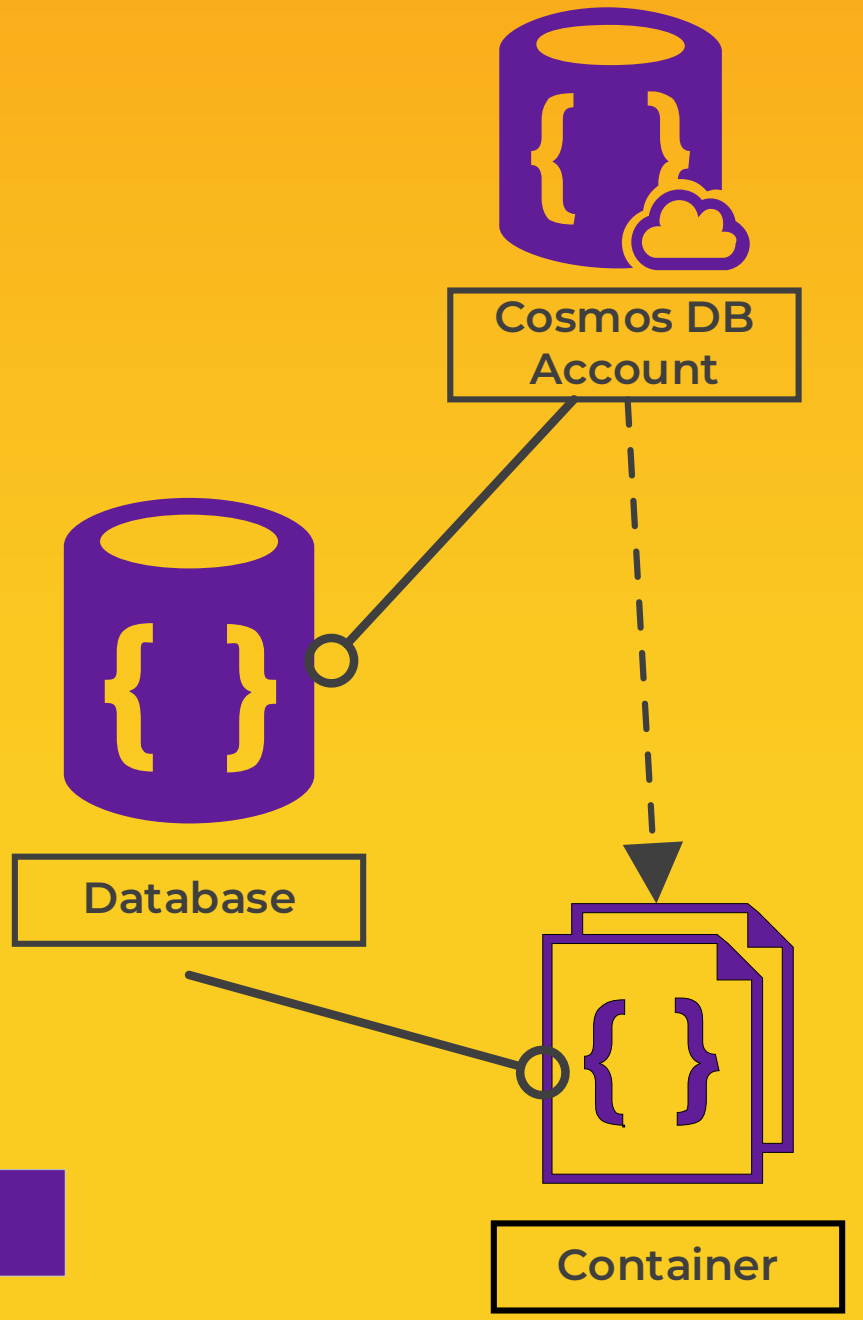

Elements in an Azure Cosmos DB Account

- Provision Azure Cosmos DB Account
- Create database in that account
- Add containers on those databases
- Container can be realized based upon the data API

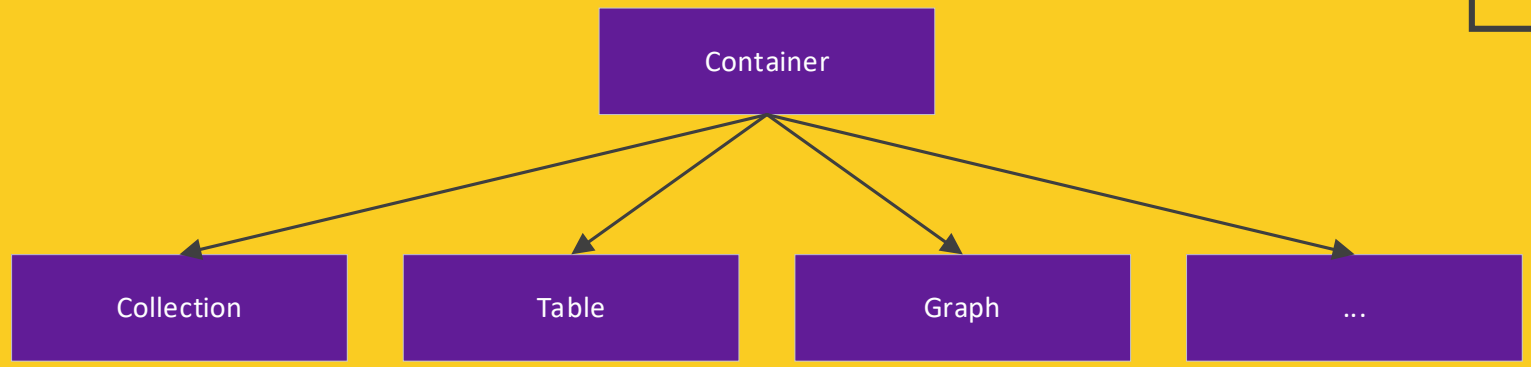

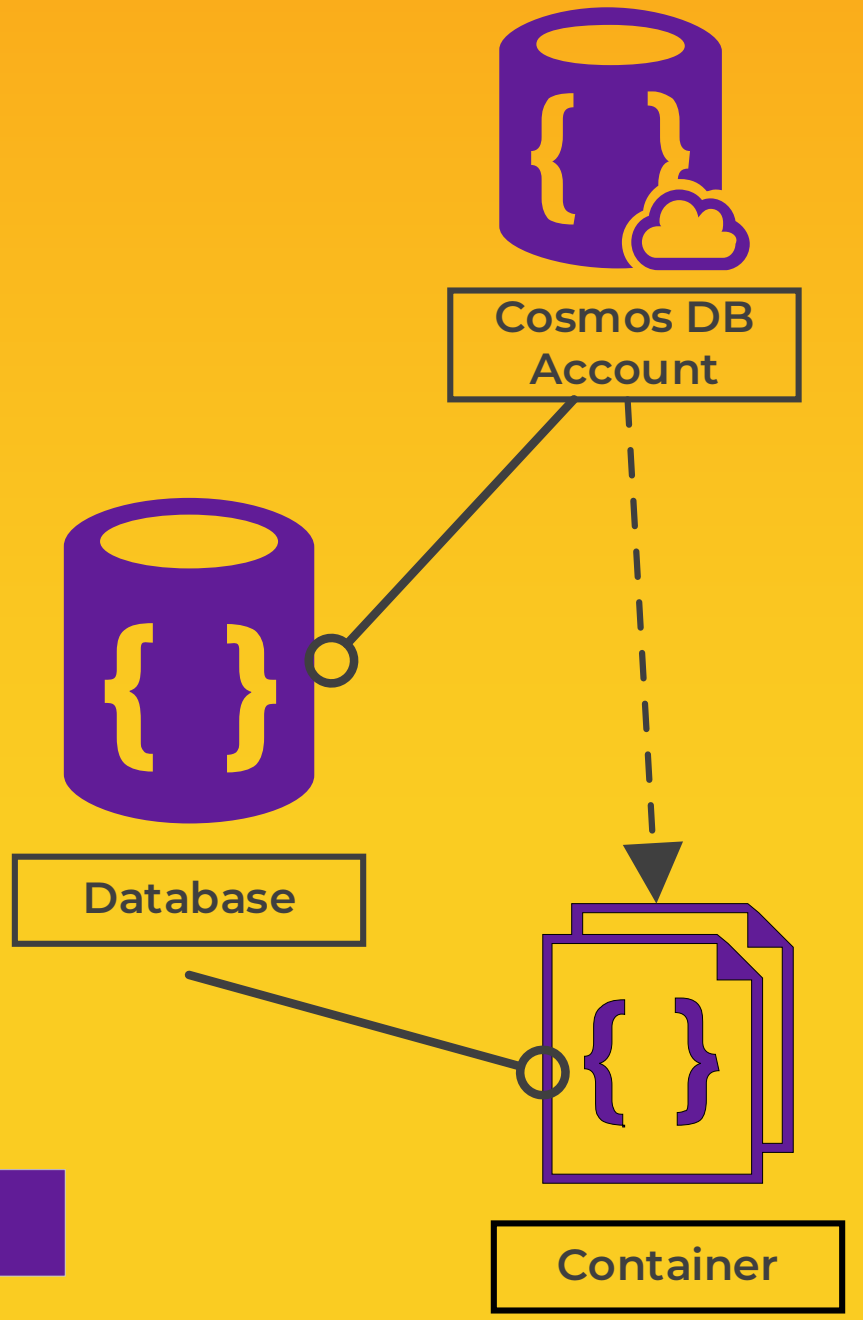

Elements in an Azure Cosmos DB Account

- Provision Azure Cosmos DB Account
- Create database in that account
- Add containers on those databases
- Container can be realized based upon the data API
- Items are realized based upon the data API

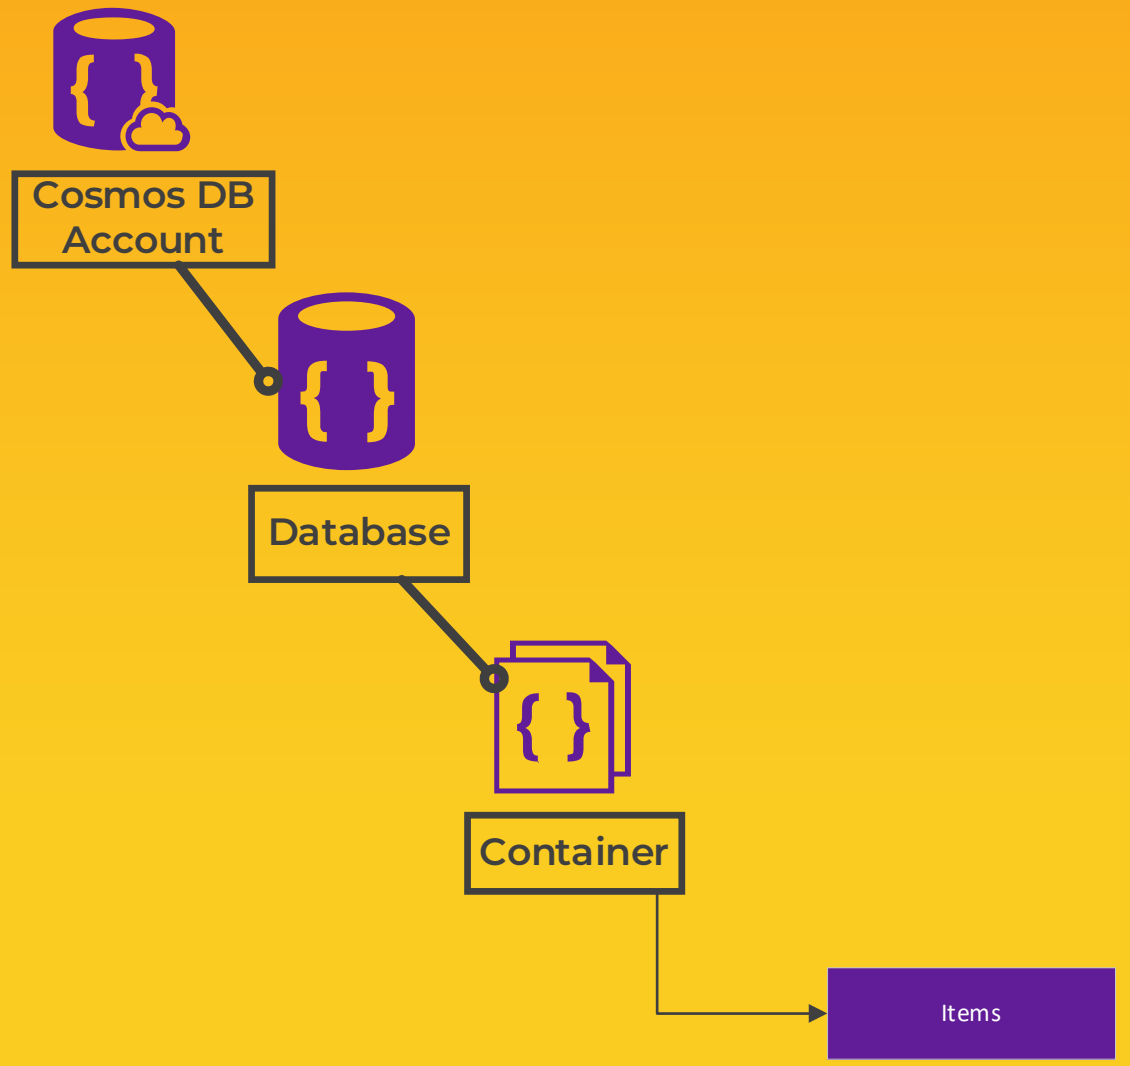

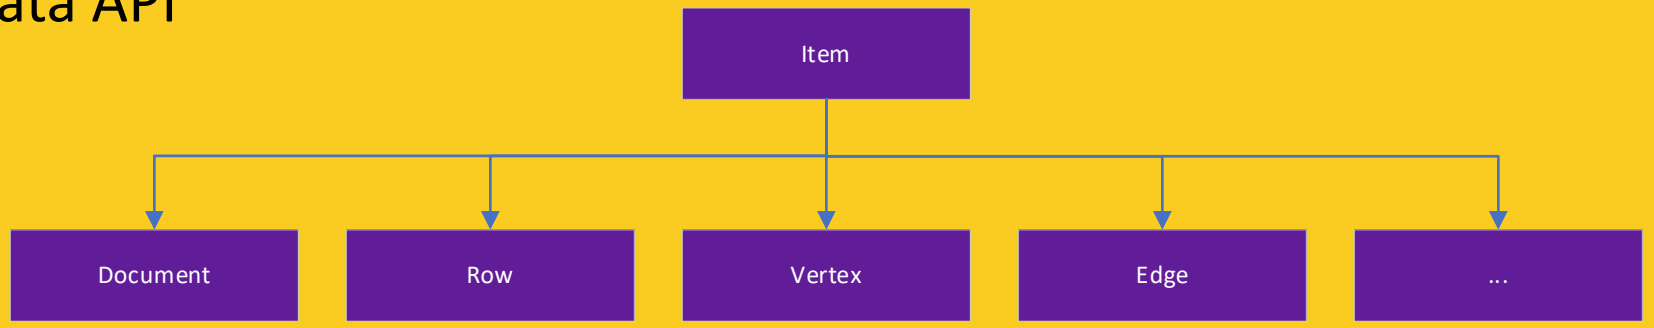

Elements in an Azure Cosmos DB Account

- Provision Azure Cosmos DB Account
- Create database in that account
- Add containers on those databases
- Container can be realized based upon the data API
- Items are realized based upon the data API

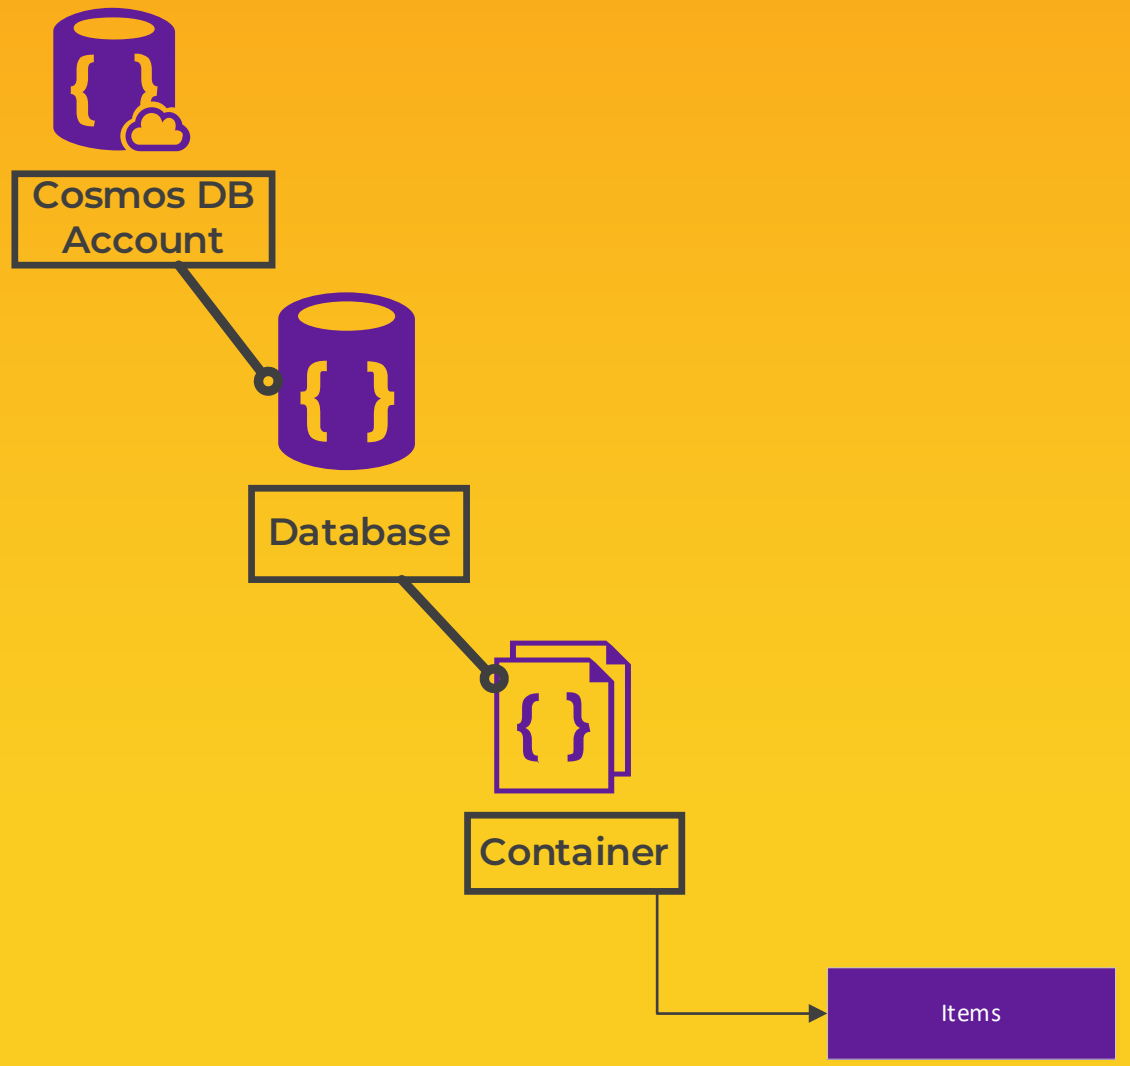

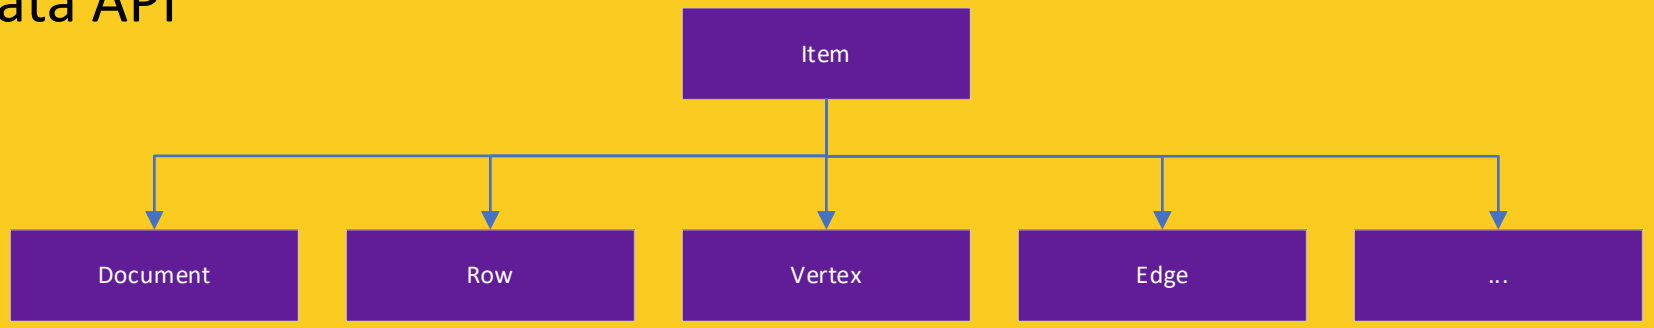

Elements in an Azure Cosmos DB Account

- Provision Azure Cosmos DB Account
- Create database in that account
- Add containers on those databases
- Container can be realized based upon the data API
- Items are realized based upon the data API

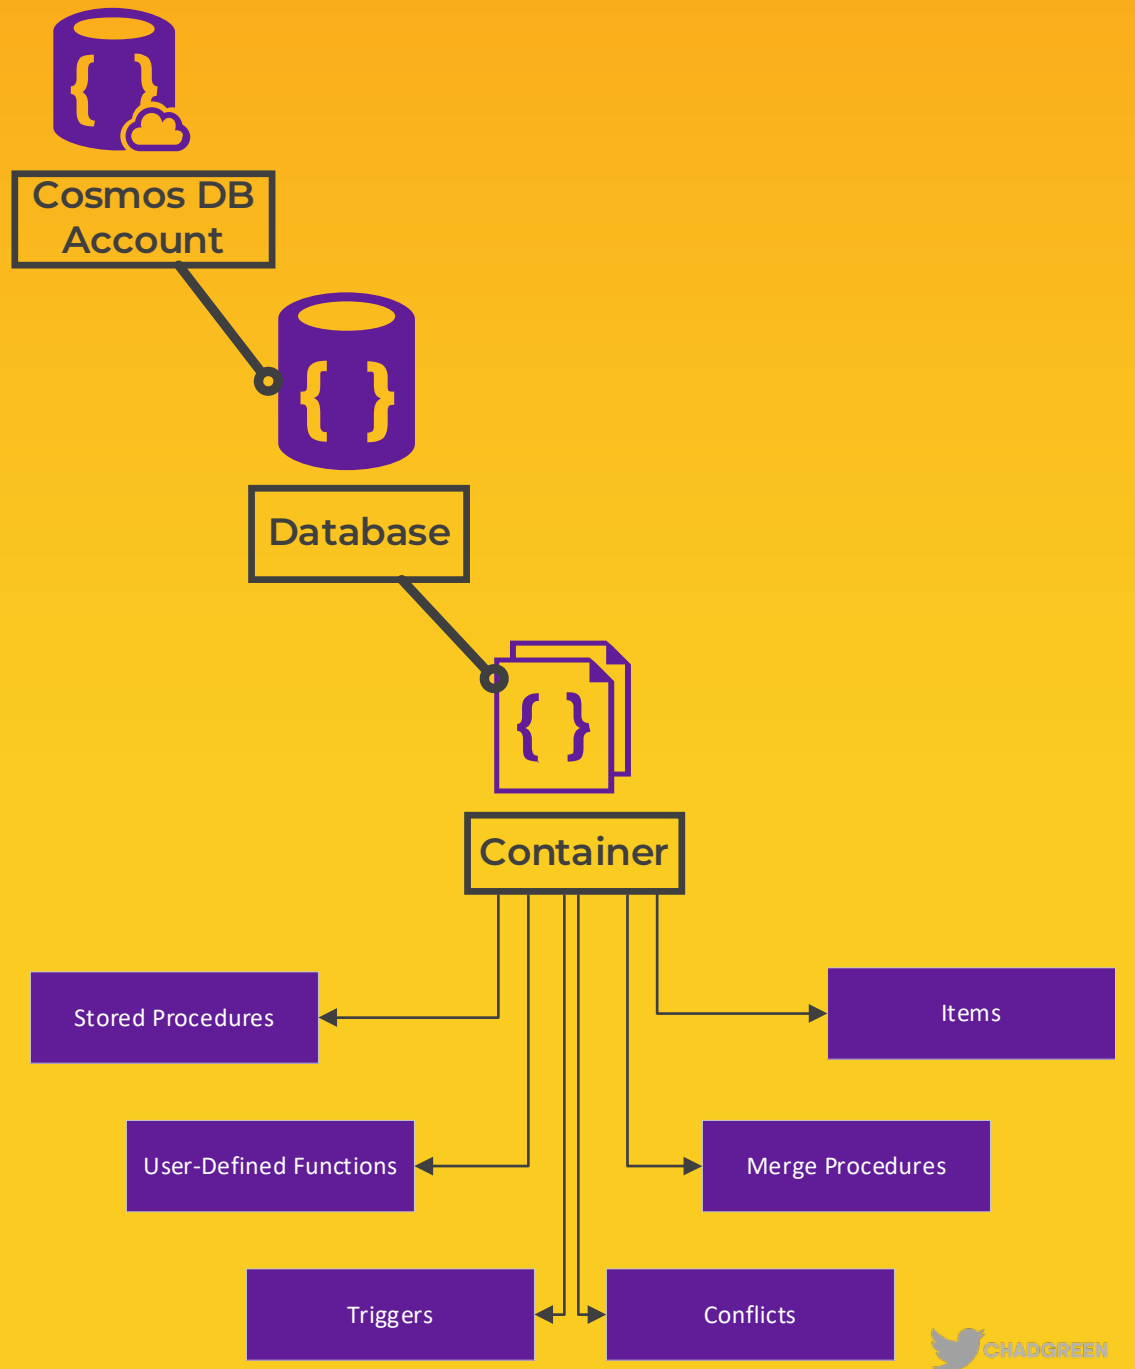

Azure Cosmos DB Architecture

Atom

Record

Sequence

DEMO

CREATE A COSMOS DB ACCOUNT

CAPACITY

Request Units (RUs)

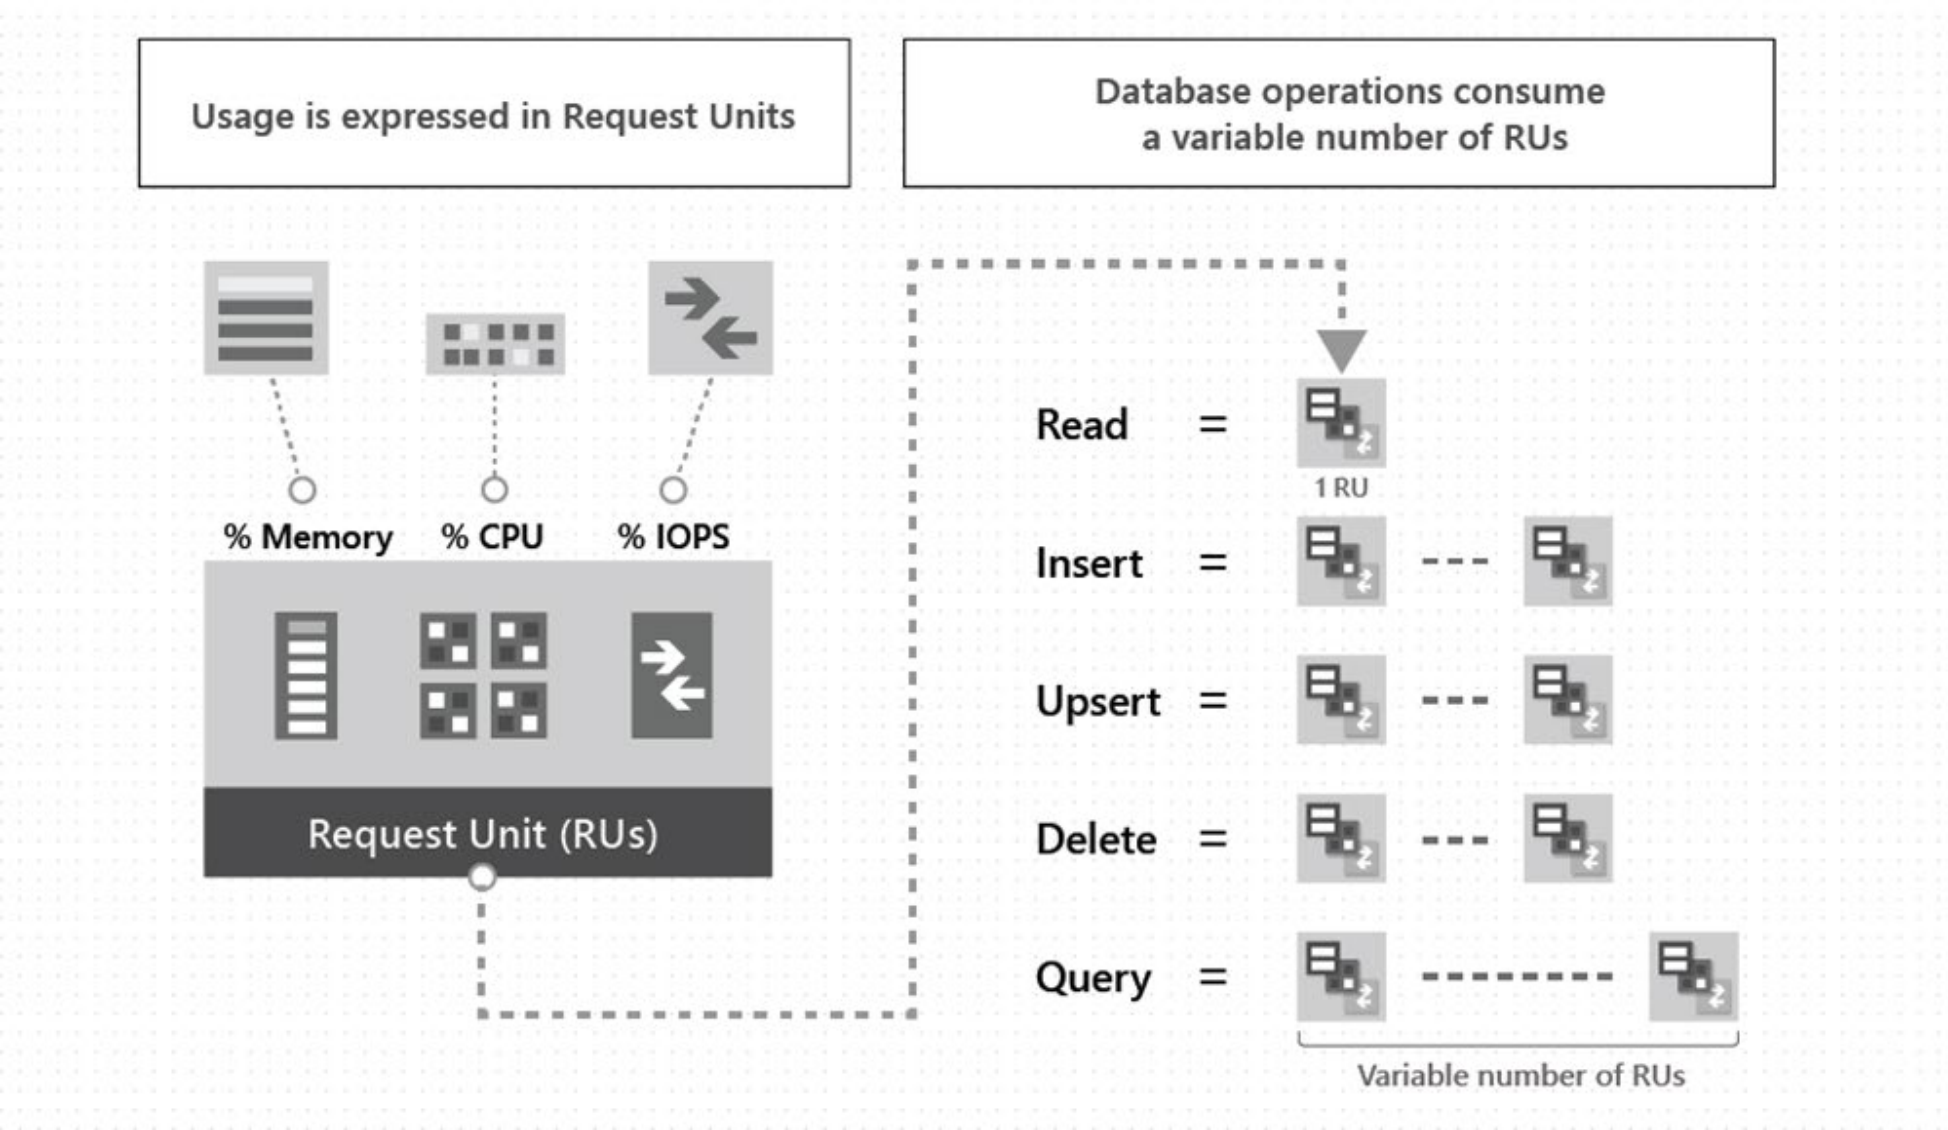

Request Unit Considerations

Item Size

Request Unit Considerations

Item Size

Item Indexing

Request Unit Considerations

- Item Size
- Item Indexing
- Item Property Count

Request Unit Considerations

Item Size

Item Indexing

Item Property Count

Indexed Properties

Request Unit Considerations

- Item Size
- Item Indexing
- Item Property Count
- Indexed Properties
- Data Consistency

Request Unit Considerations

- Item Size
- Item Indexing
- Item Property Count
- Indexed Properties
- Data Consistency
- Type of Reads

Request Unit Considerations

- Item Size
- Item Indexing
- Item Property Count
- Indexed Properties
- Data Consistency
- Type of Reads
- Query Patterns

Request Unit Considerations

- Item Size
- Item Indexing
- Item Property Count
- Indexed Properties
- Data Consistency
- Type of Reads
- Query Patterns
- Script Usage

Estimate RU Costs

Operation	Estimated Costs
Create an item	5 RUs
Update an item	10 RUs
Read an item (point read)	1 RU
Delete an Item	5 RUs
Execute a query	10 RUs

Operational RU Costs

Operation	Estimated Cost	Notes
Create an item	5 RUs	Average cost for a 1-Kb item with less than 5 properties to index
Update an item	10 RUs	Average cost for a 1-Kb item with less than 5 properties to index
Read an item (point-read)	1 RU	Average cost for a 1-Kb item
Delete an item	5 RUs	
Execute a query	10 RUs	Average cost for query that takes full advantage of indexing and returns 100 results or less

PARTITIONING

Partitioning

Logical Partitions

Partitioning

Logical Partitions

Physical Partitions

Choosing a partition key

Unchanging
Property Value

Choosing a partition key

Unchanging Property Value

High Cardinality

Choosing a partition key

Unchanging Property Value

High Cardinality

Spreads RU Consumption

Choosing a partition key

Unchanging Property Value

High Cardinality

Spreads RU Consumption

Common Filter

Using Item Id as the Partition Key

Wide Range of Possible Values

Balances RU Consumption

Point Reads Become Easier

Using Item Id as the Partition Key

Wide Range of Possible Values

Balances RU Consumption

Point Reads Become Easier

Partition key will become unique identifier

Should have an equality filter with the *item Id*

Stored Procedures/Triggers Cannot Run Across Multiple Partitions

PROVISIONING

Provisioning – The Basics

Database Provisioning

Provisioning – The Basics

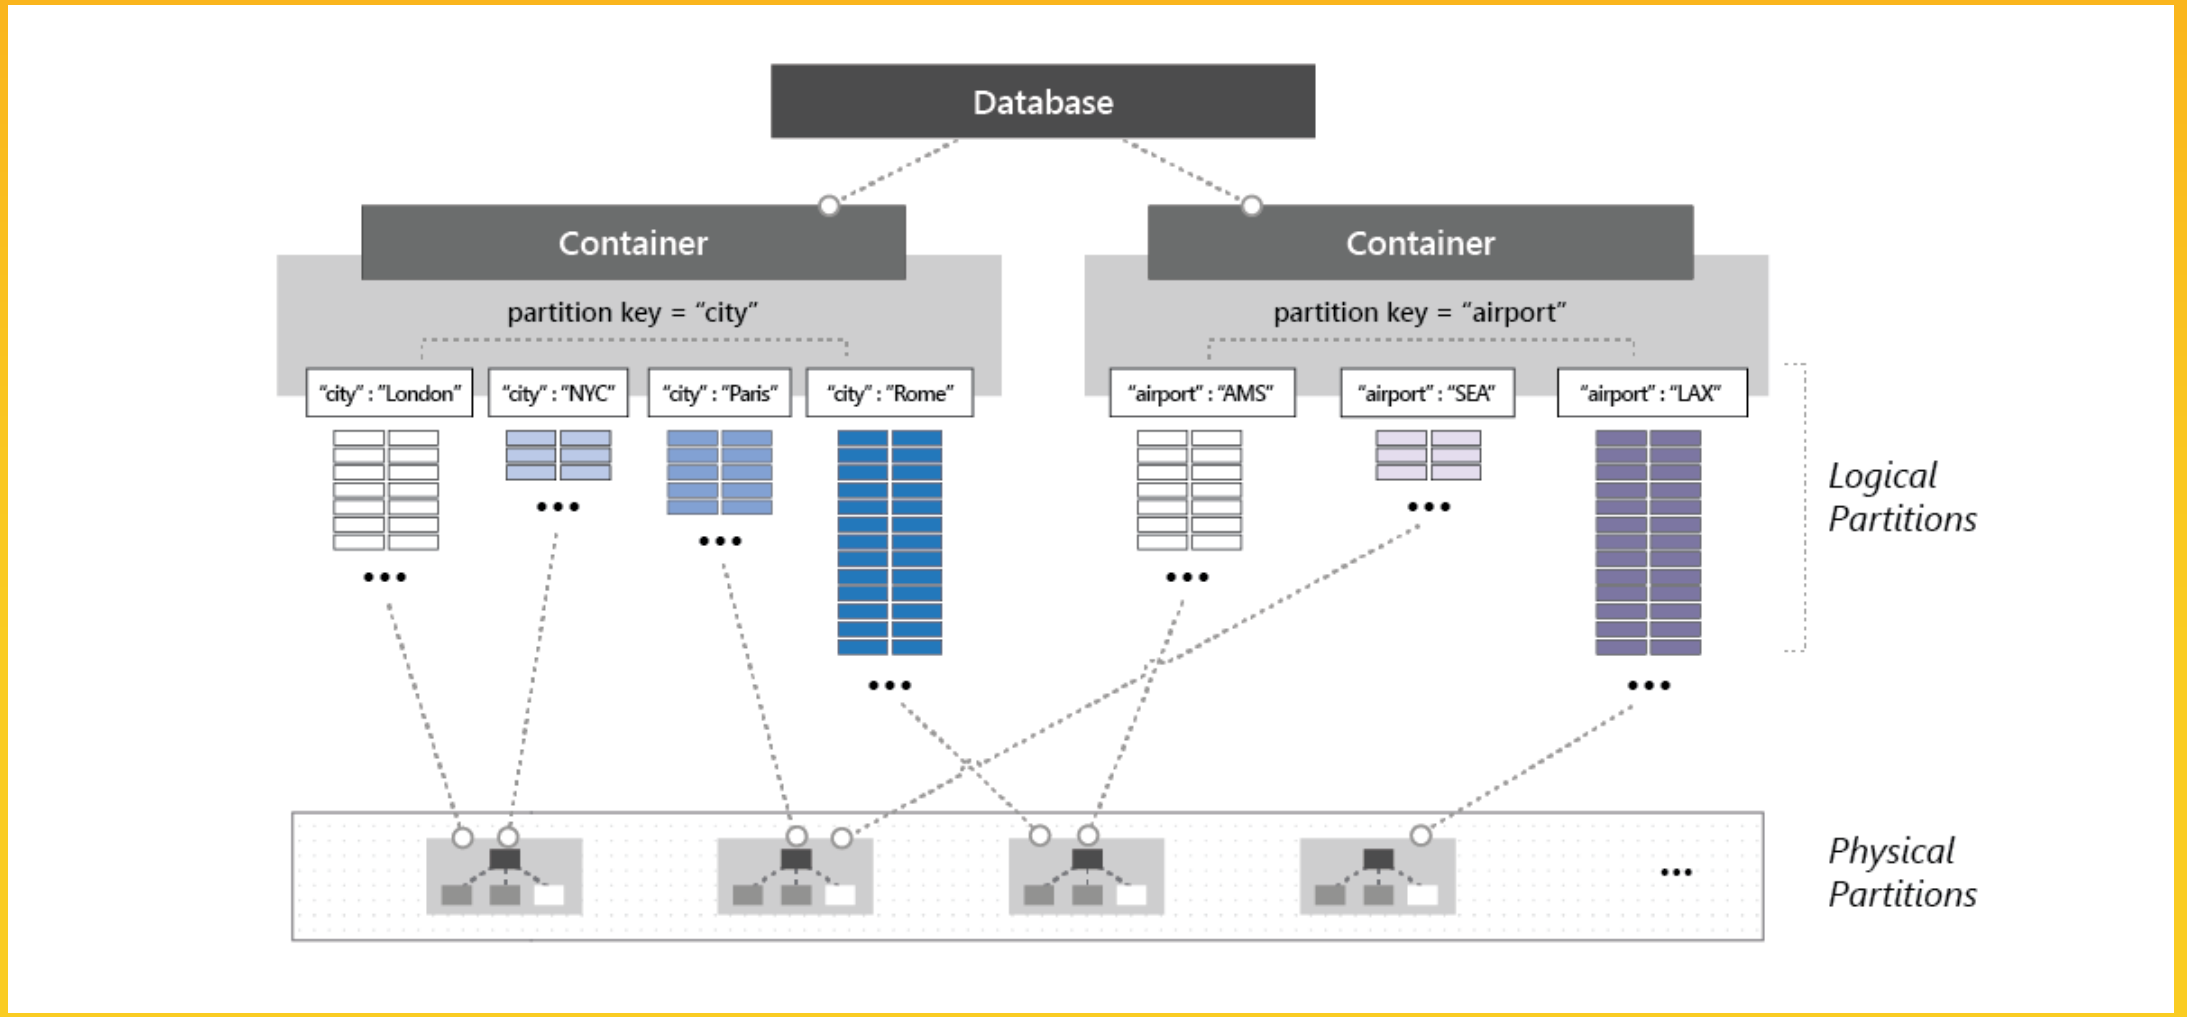

Database Provisioning

Provisioning – The Basics

Database Provisioning

Container Provisioning

Provisioning – The Basics

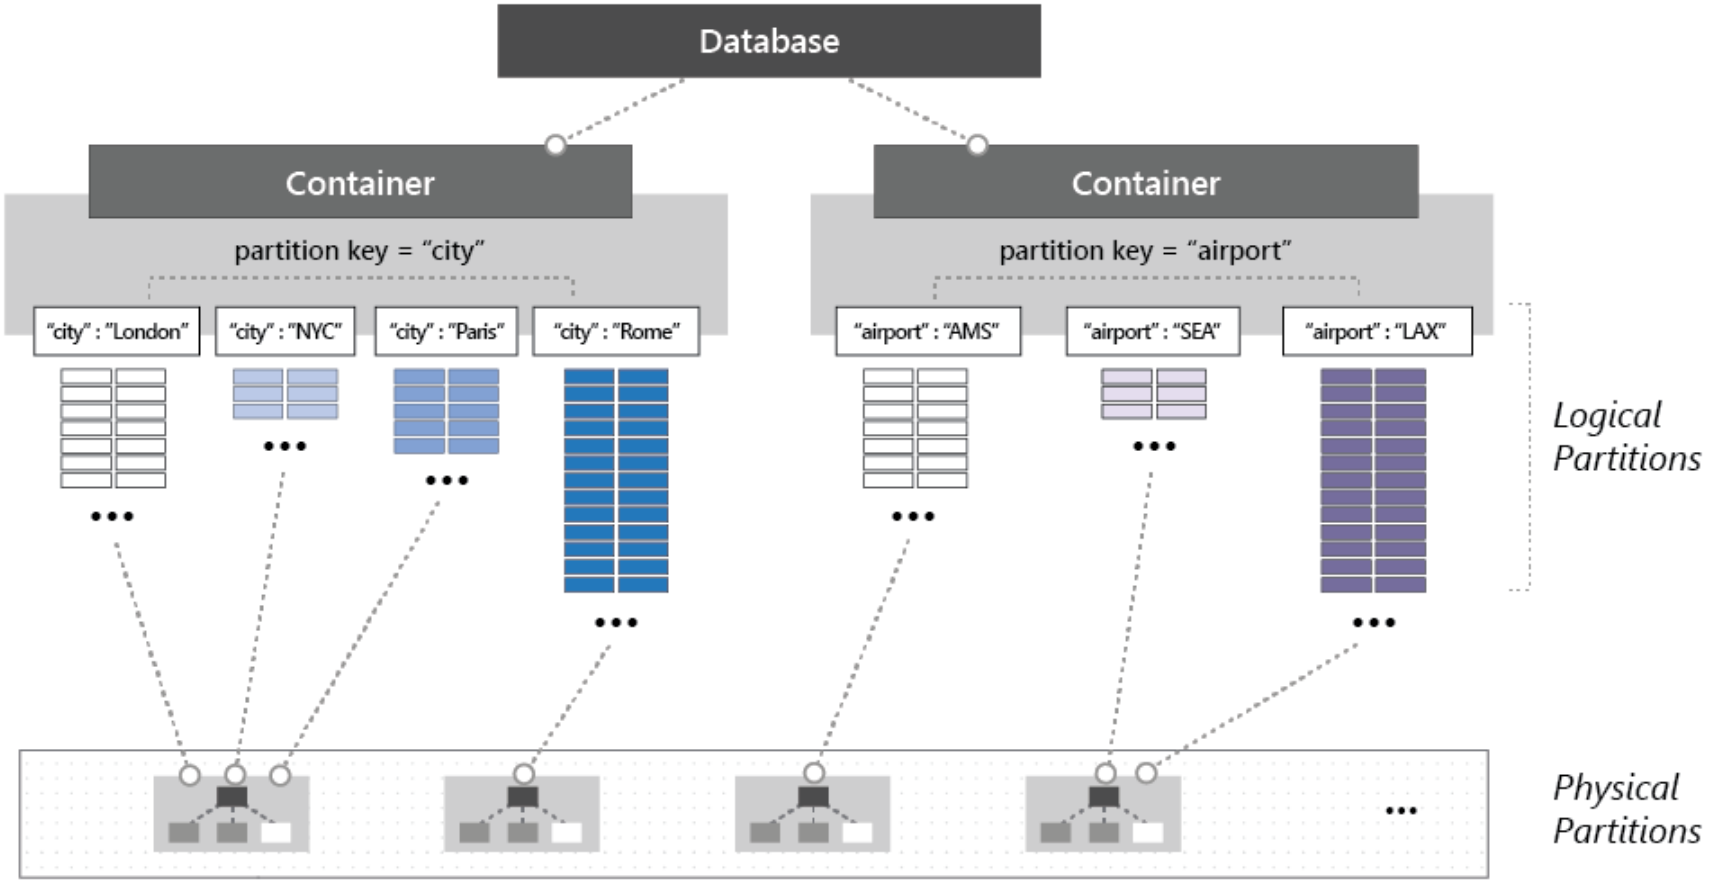

Container Provisioning

Provisioning – The Basics

Database Provisioning

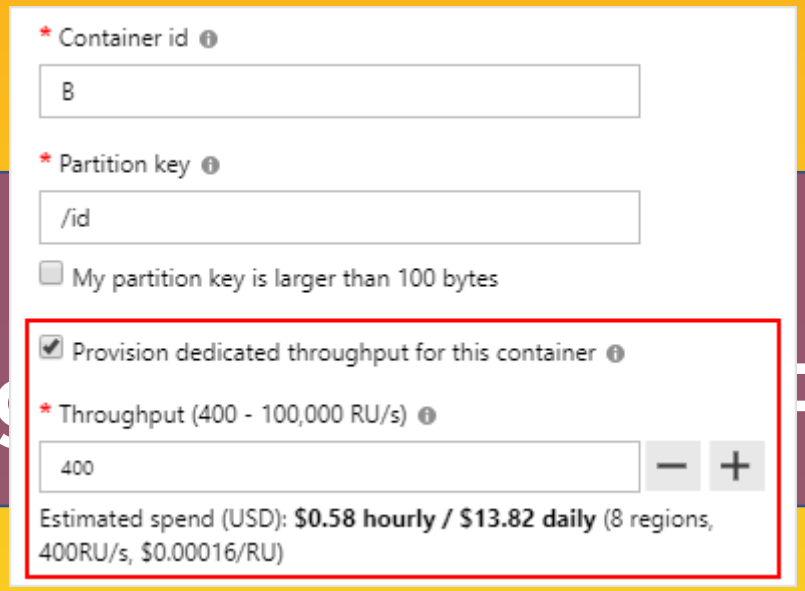

Container Provisioning

Combination

Provisioning – Autoscale

Provisioning – Autoscale Benefits

Simple

Provisioning – Autoscale Benefits

Simple

Scalable

Provisioning – Autoscale Benefits

Simple

Scalable

Cost-Effective

Provisioning – Autoscale Benefits

Simple

Scalable

Cost-Effective

Highly Available

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

New Applications

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

New Applications

Infrequently Used Applications

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

New Applications

Infrequently Used Applications

Development and Test Workloads

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

New Applications

Infrequently Used Applications

Development and Test Workloads

Scheduled Production Workloads/Queries

Provisioning – Autoscale Use Cases

Variable/Unpredictable Workloads

New Applications

Infrequently Used Applications

Development and Test Workloads

Scheduled Production Workloads/Queries

Provisioning – Autoscale Use Cases

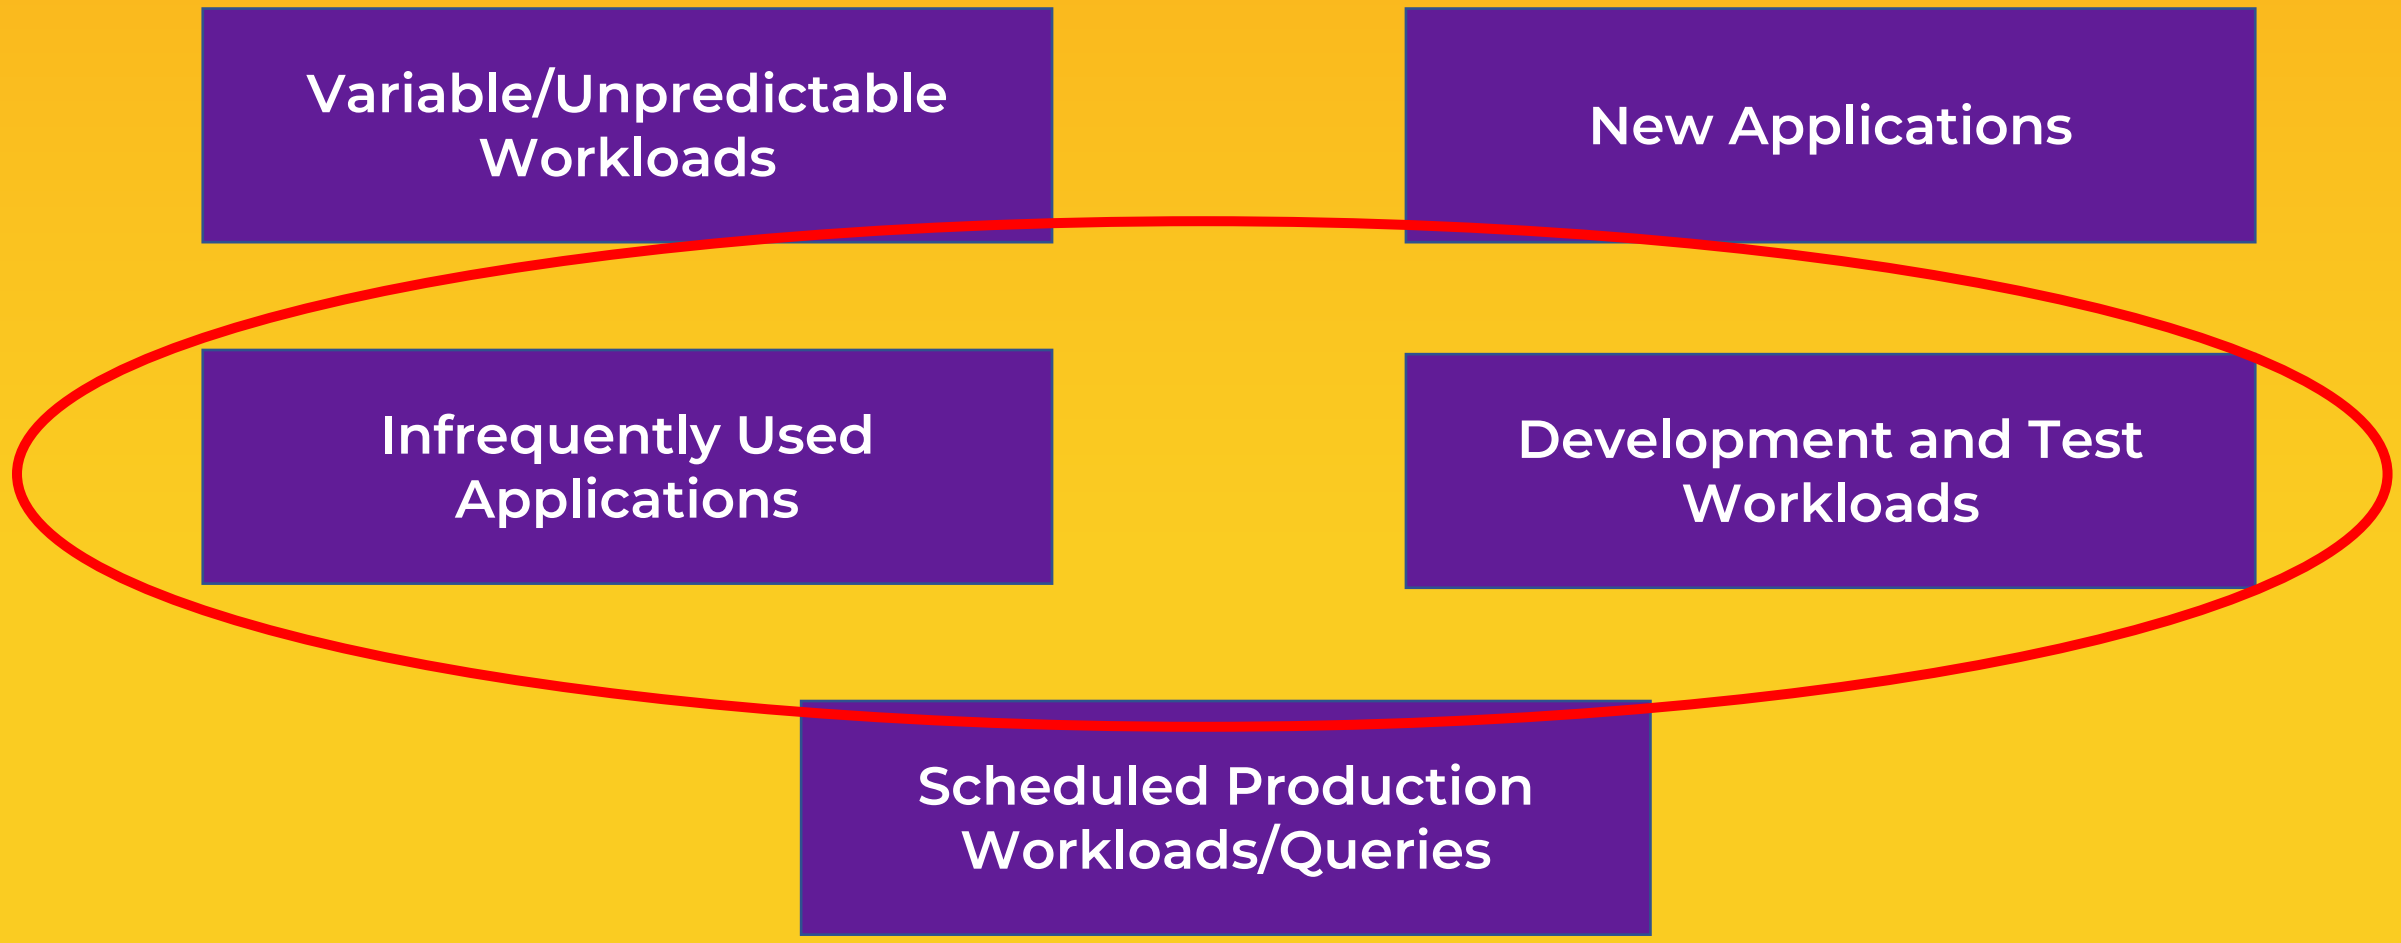

Provisioning – Serverless

Provisioned Throughput
Guarantee Based Billing

Provisioning – Serverless

Preview

Provisioned Throughput
Guarantee Based Billing

Serverless Throughput
Consumption Based Billing

Provisioning – Serverless Performance

Availability

Provisioning – Serverless Performance

Availability

Latency

Provisioning – Serverless Performance

Availability

Latency

Burstability

Provisioning – Serverless Use Cases

Light Traffic

Provisioning – Serverless Use Cases

Light Traffic

Moderate Burstability

Provisioning – Serverless Use Cases

Light Traffic

Moderate Burstability

Moderate Performance

Provisioning – Serverless Use Cases

Light Traffic

Moderate Burstability

Moderate Performance

Provisioning – Serverless Use Cases

Light Traffic	Moderate Burstability	Moderate Performance

- **Development**
- **Testing**
- **Prototyping**
- **Proof of concept**
- **Non-critical application with light traffic**

Provisioning – Serverless Limitations

Single Region

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Max of 5,000 RU/s

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Max of 5,000 RU/s

Maximum of 50-Gb Storage

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Max of 5,000 RU/s

Maximum of 50-Gb Storage

Core (SQL) API Only

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Max of 5,000 RU/s

Maximum of 50-Gb Storage

Core (SQL) API Only

Unable to migrate to/from

Provisioning – Serverless Limitations

Single Region

Synapse Link Unavailable

Unable to Specify RU Provisioning

Max of 5,000 RU/s

Maximum 50 Gb Storage

Core (SQL) API Only

Unable to migrate to/from

Preview

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region
Limitations per Container	Unlimited throughput Unlimited storage	Maximum of 5,000 RU/s Maximum of 50-Gb Storage

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region
Limitations per Container	Unlimited throughput Unlimited storage	Maximum of 5,000 RU/s Maximum of 50-Gb Storage
Availability Guarantee	99.99% to 99.999%	99.9 to 99.99%

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region
Limitations per Container	Unlimited throughput Unlimited storage	Maximum of 5,000 RU/s Maximum of 50-Gb Storage
Availability Guarantee	99.99% to 99.999%	99.9 to 99.99%
Latency Guarantee	< 10-ms for point-reads and writes (SLA)	< 10-ms for point-reads and < 30-ms for writes (SLO)

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region
Limitations per Container	Unlimited throughput Unlimited storage	Maximum of 5,000 RU/s Maximum of 50-Gb Storage
Availability Guarantee	99.99% to 99.999%	99.9 to 99.99%
Latency Guarantee	< 10-ms for point-reads and writes (SLA)	< 10-ms for point-reads and < 30-ms for writes (SLO)
Throughput Guarantee	99.99% (SLA)	95% Burstability (SLO)

Provisioning – Choosing

Criteria	Provisioned	Serverless
Status	Generally Available	In Preview
Best Suited For	Mission-critical workloads requiring predictable performance (and cost)	Small-to-medium non-critical workloads with light traffic
Limitations per Account	Unlimited Azure regions	Limited to one Azure region
Limitations per Container	Unlimited throughput Unlimited storage	Maximum of 5,000 RU/s Maximum of 50-Gb Storage
Availability Guarantee	99.99% to 99.999%	99.9 to 99.99%
Latency Guarantee	< 10-ms for point-reads and writes (SLA)	< 10-ms for point-reads and < 30-ms for writes (SLO)
Throughput Guarantee	99.99% (SLA)	95% Burstability (SLO)
Billing Model	Per-hour basis for RU/s provisioned, regardless of how many RUs consumed	Per-hour bases for the amount of RUs consumed by your database operations

Provisioning – Choosing

Burstability

Expected Consumption

Provisioning – Choosing

Burstability

Expected Consumption

Provisioning – Choosing

Burstability

Expected Consumption

Workload expected to burst to a maximum of 500 RU/s and consume a total of 20,000,000 RUs over a month

Provisioned	Serverless
$29.20	$5.00

Provisioning – Choosing

Burstability

Expected Consumption

Workload is expected to burst to a maximum 500 RU/s and consume a total of 250,000,000 RUs over a month.

Provisioned	Serverless
$29.20	$62.50

CONSISTENCY LEVELS

Consistency Levels – Strong

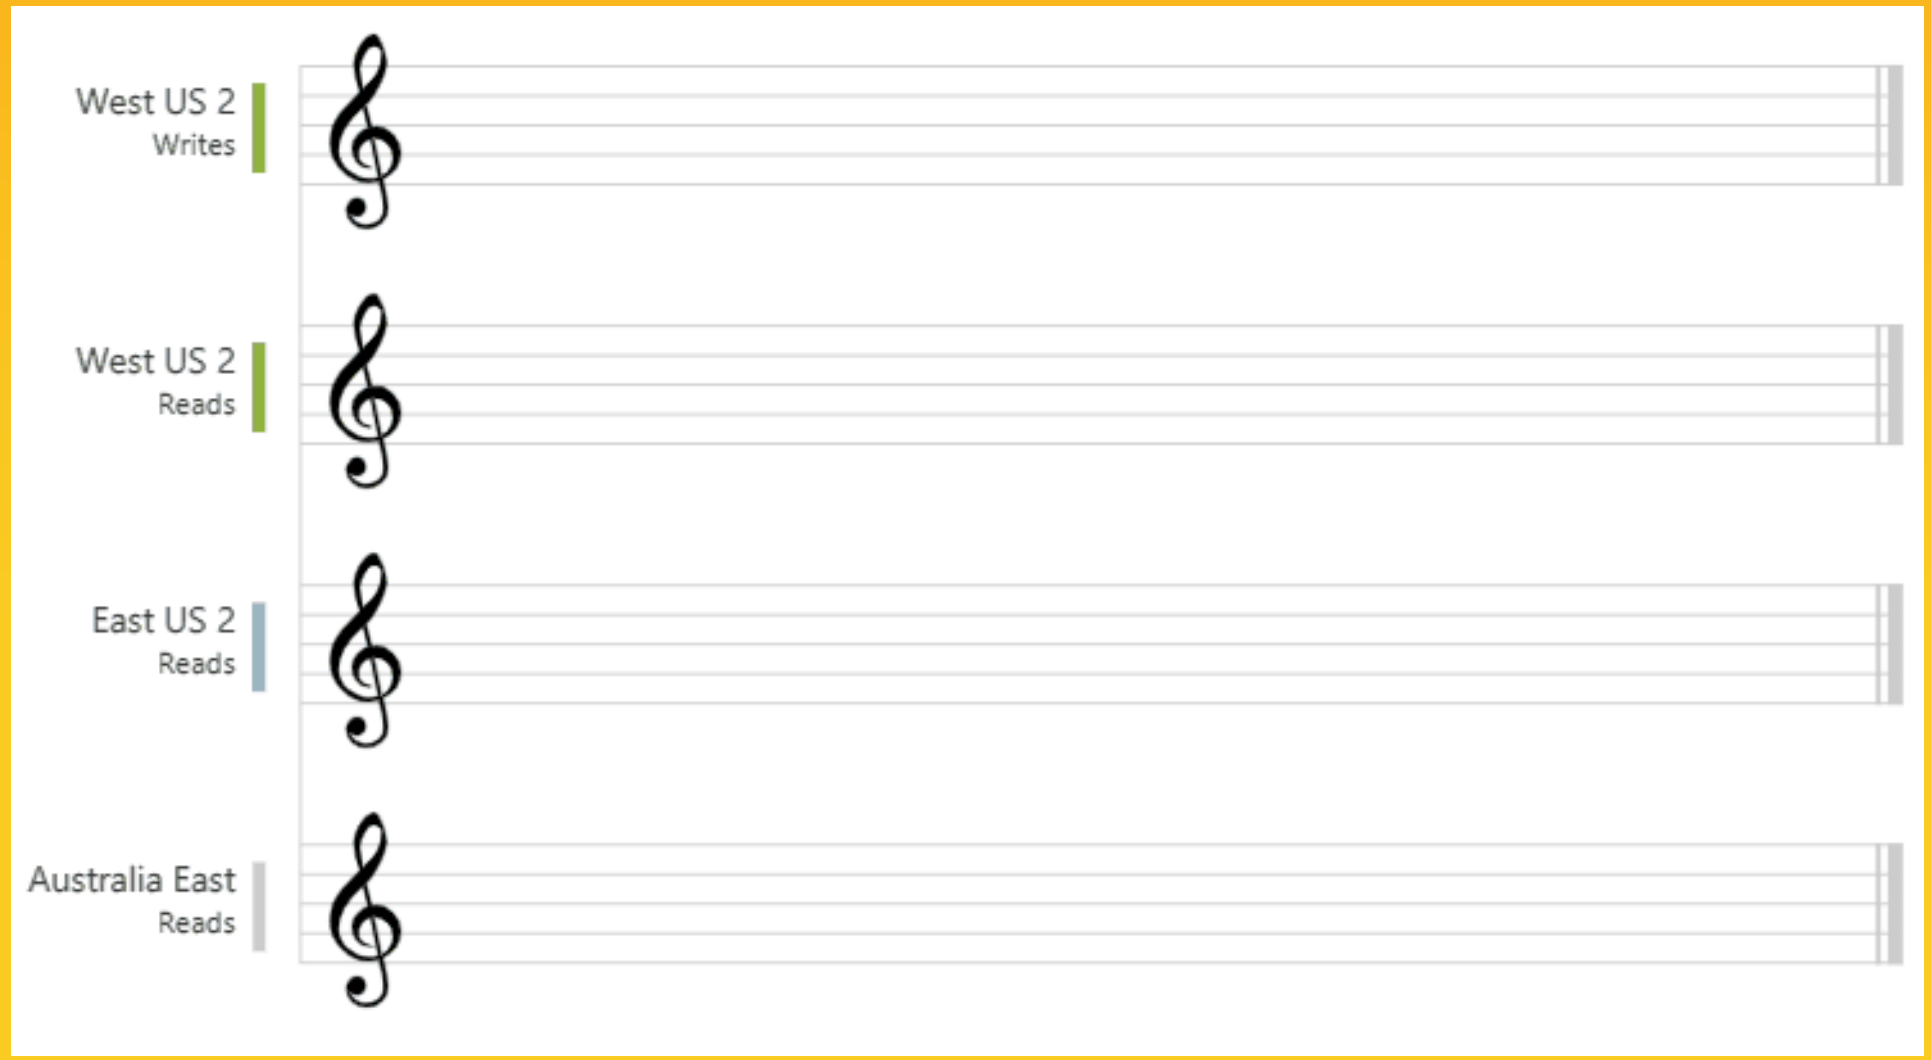

Consistency Levels – Bounded Staleness

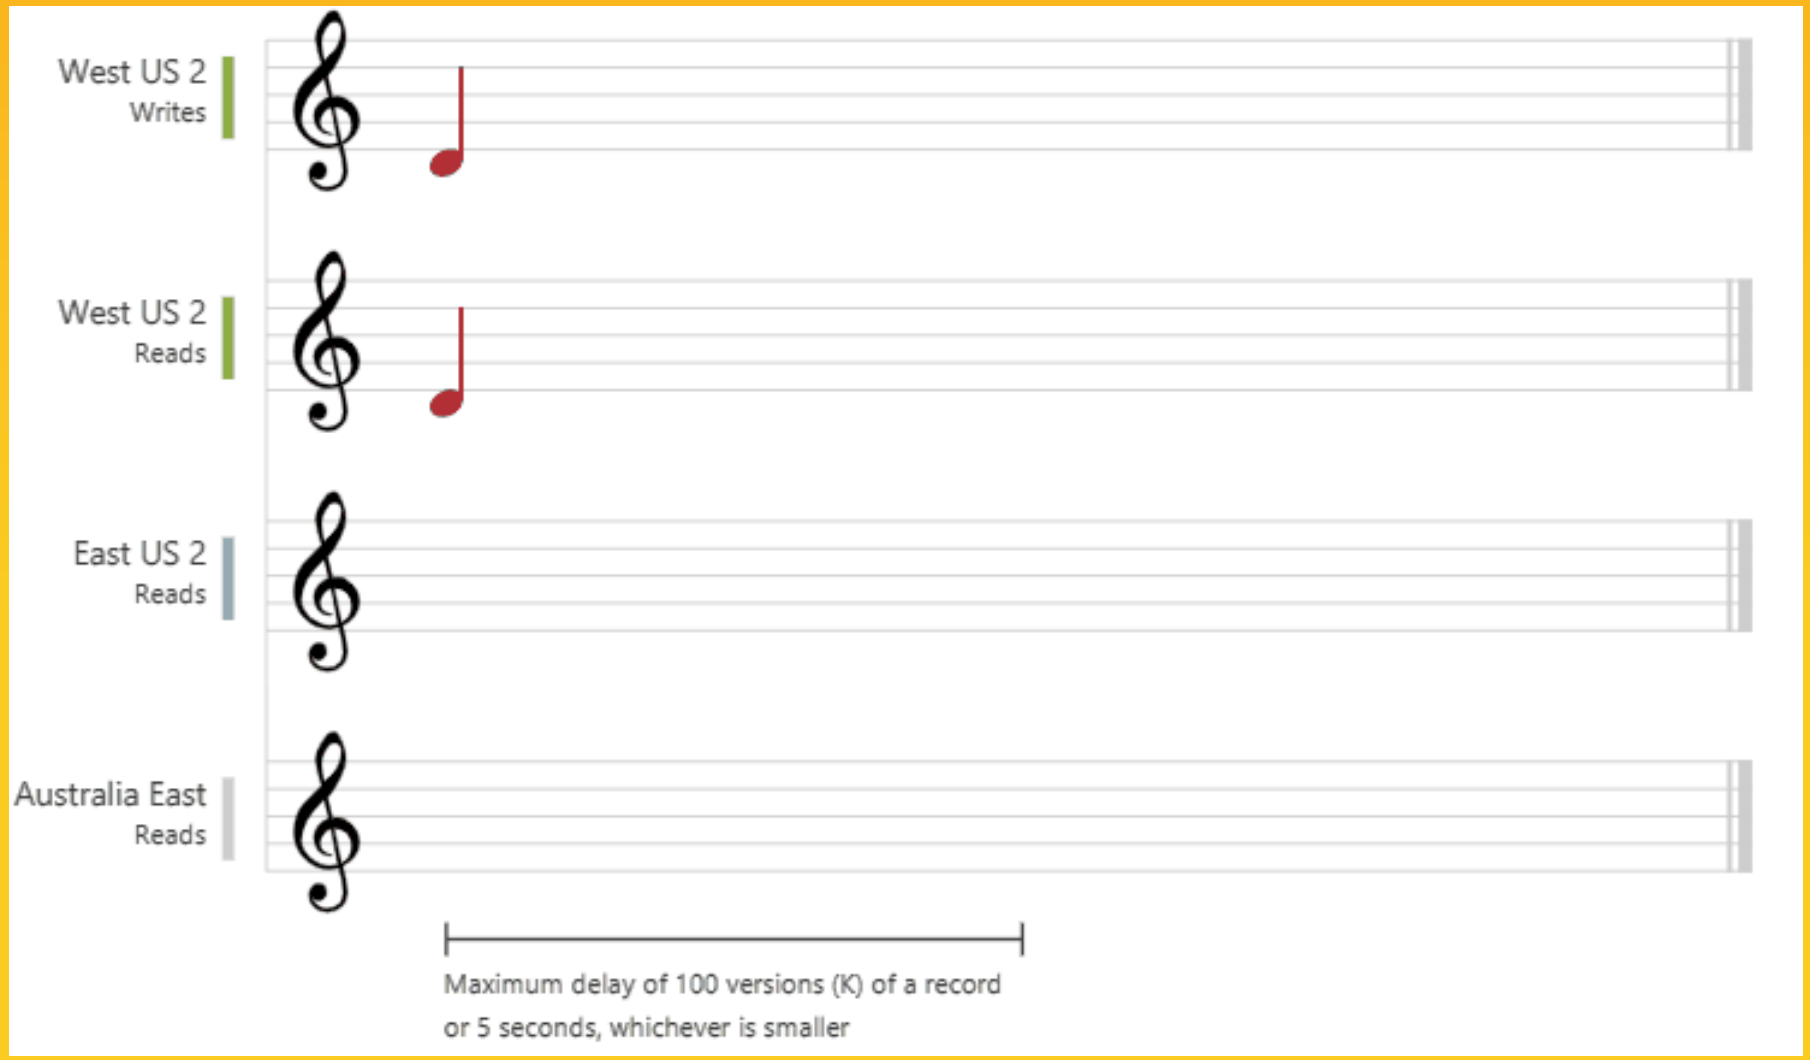

Consistency Levels – Session

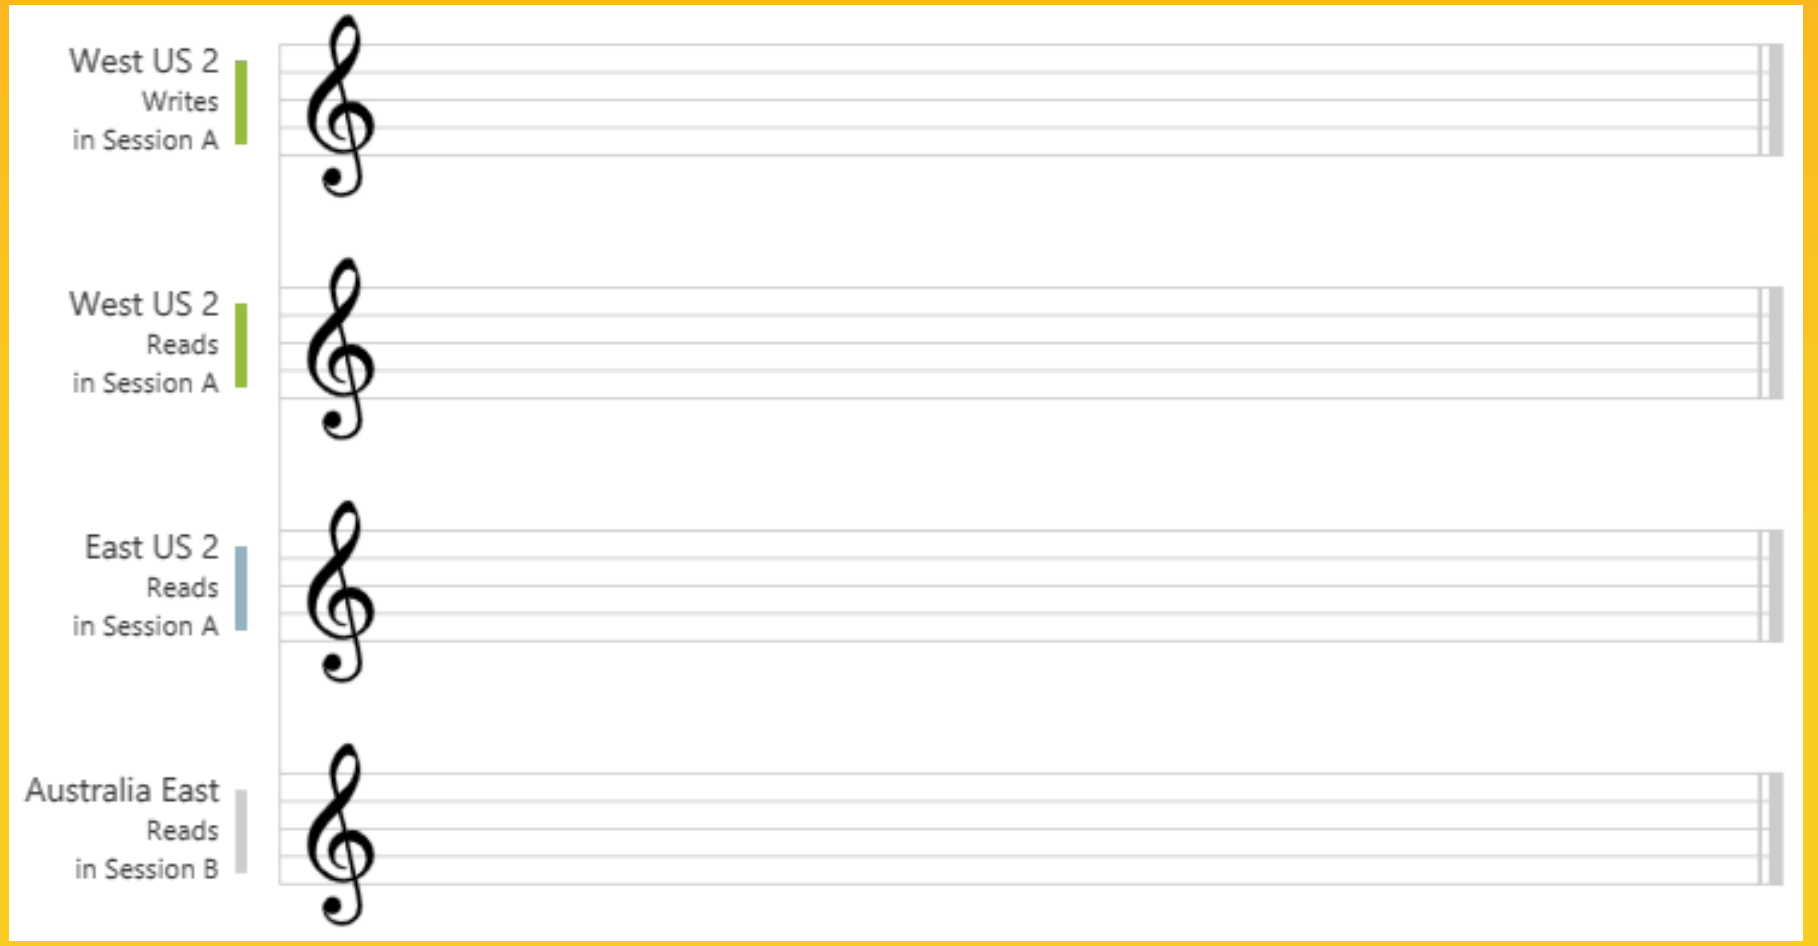

Consistency Levels – Consistent Prefix

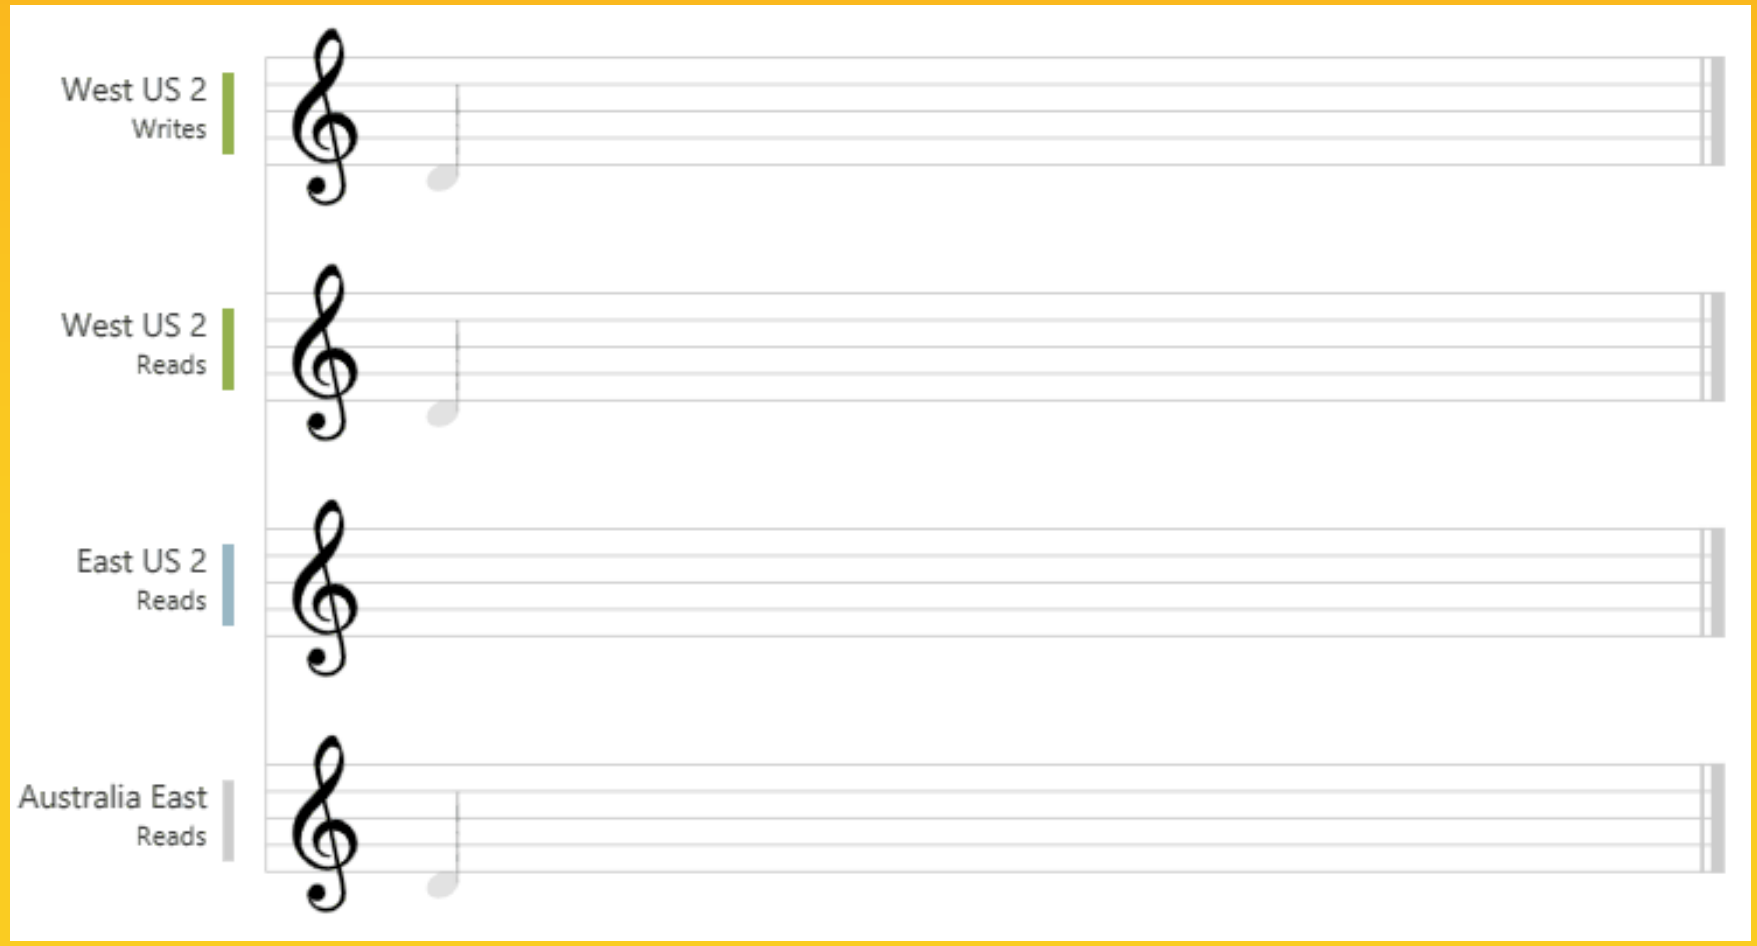

Consistency Levels – Eventual

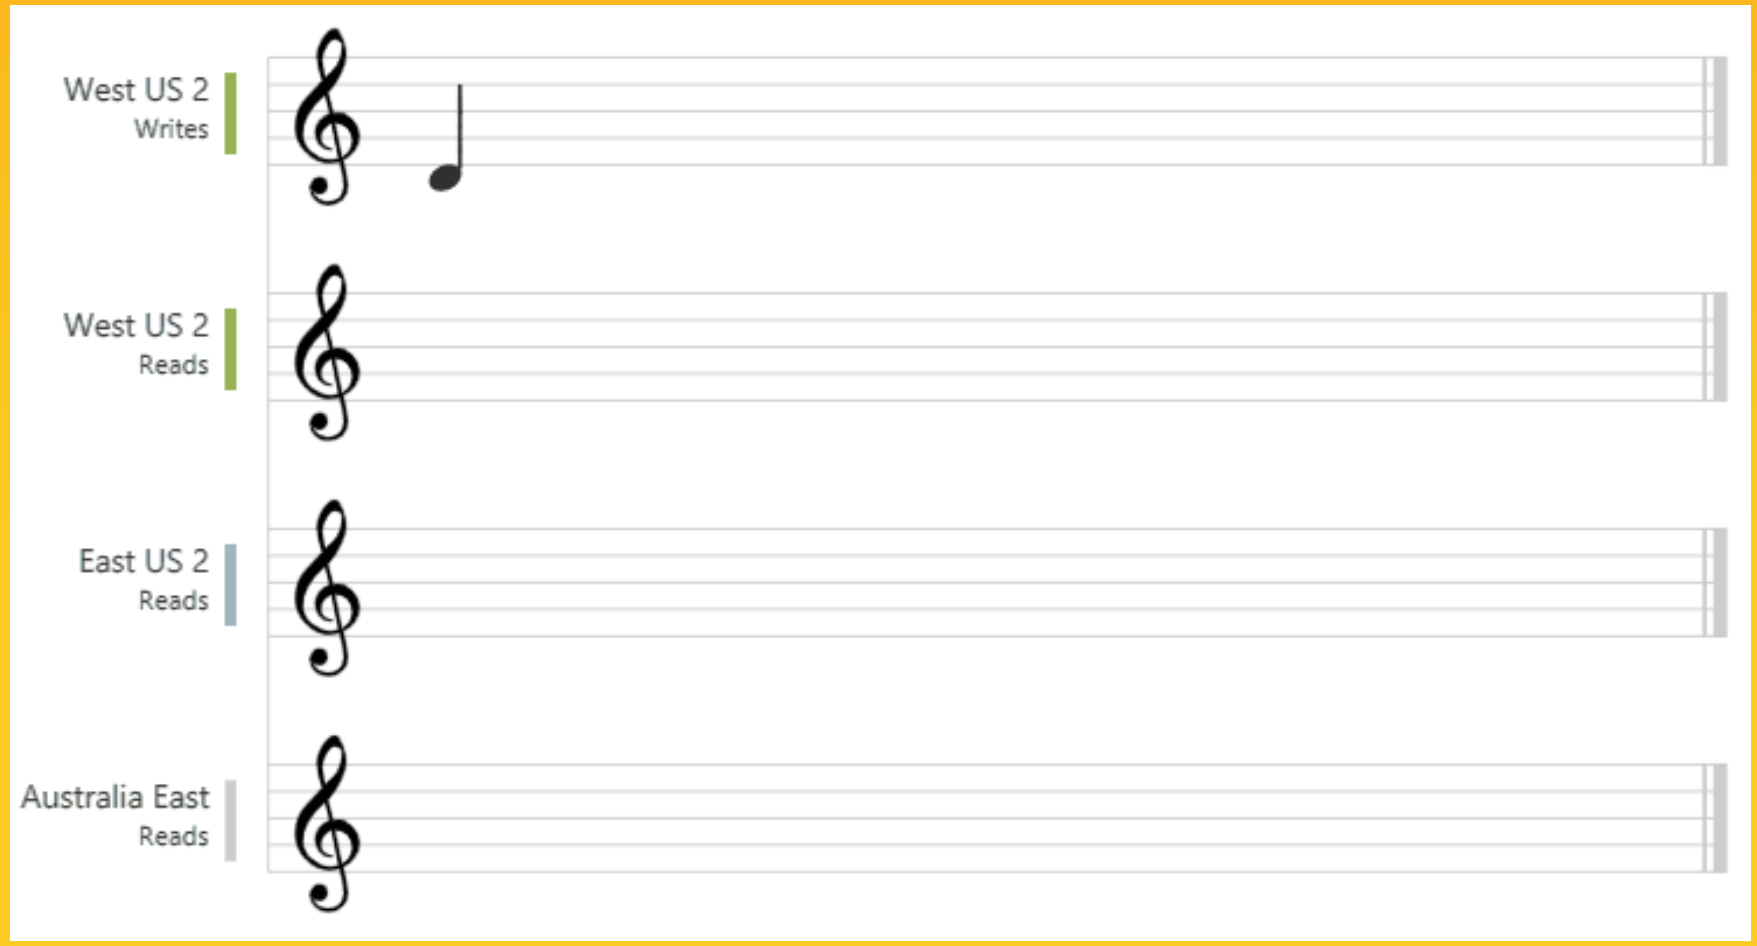

Consistency Guarantees in Practice

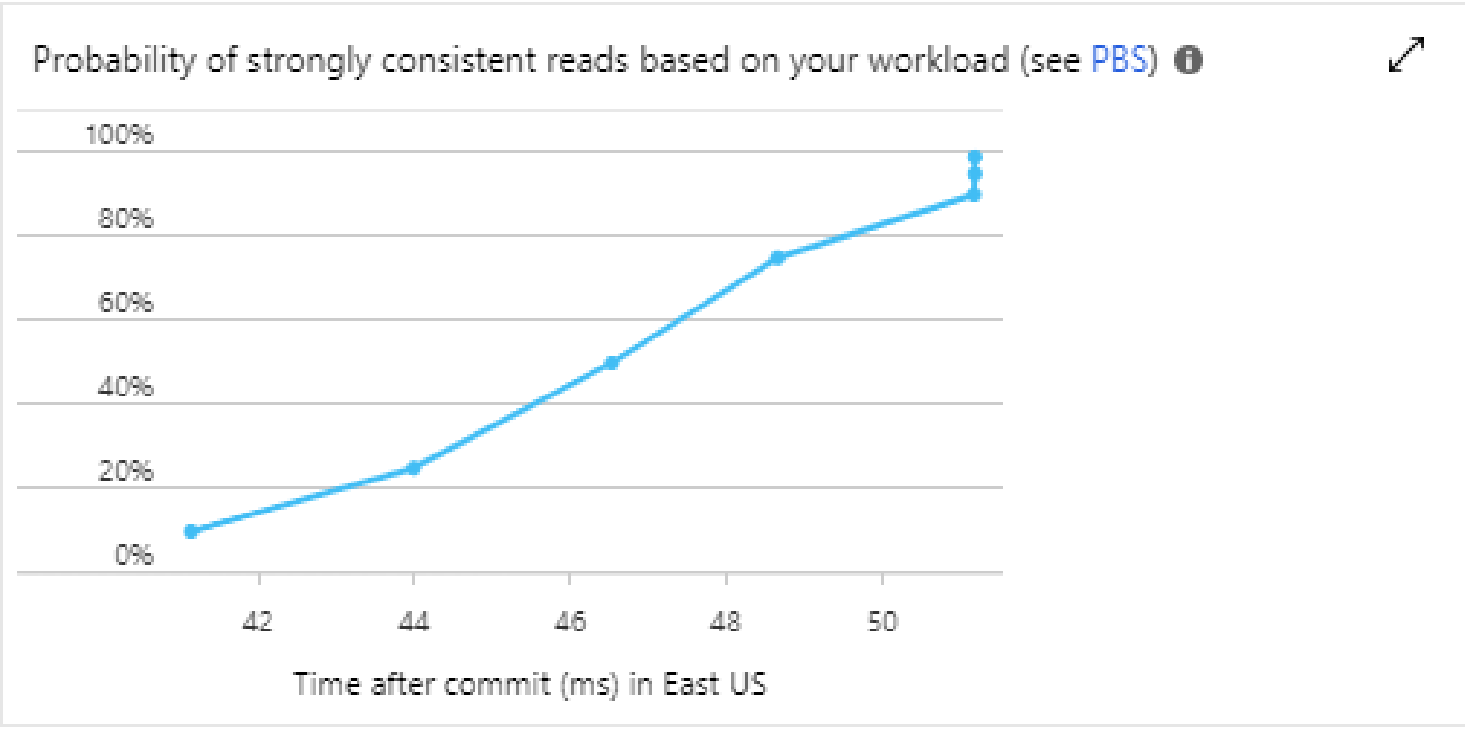

SERVICE QUOTAS

Service Quotas – Provisioned Throughput

Resource	Default Limit
Maximum RUs per container	1,000,000 (by default)
Maximum RUs per database	1,000,000 (by default)
Maximum RUs per (logical) partition	10,000
Maximum storage across all items per (logical) partition	20-Gb
Maximum number of distinct (logical) partition keys	Unlimited
Maximum storage per container	Unlimited
Maximum storage per database	Unlimited
Minimum RU/s required pre 1-Gb	10 RU/s

Service Quotas – Serverless

Resource	Default Limit
Maximum RU/s per container	5,000
Maximum RU/s per (logical) partition	5,000
Maximum storage across all items per (logical) partition	20-Gb
Maximum number of distinct (logical) partition keys	Unlimited
Maximum storage per container	50-Gb

Service Quotas – Per-Account Limits

Resource	Provisioned Throughput	Serverless
Maximum number of databases	Unlimited	Unlimited
Maximum number of containers	Unlimited per account 25 per database	100 per account
Maximum number of regions	No limit (All Azure regions)	1 (Any Azure region)

Service Quotas – Per-Item Limits

Resource	Default Limit
Maximum size of an item	2-MB
Maximum length of partition key value	2048 bytes
Maximum length of ID value	1023 bytes
Maximum number of properties per item	No practical limit
Maximum length of property name	No practical limit
Maximum length of property value	No practical limit
Maximum length of string property value	No practical limit

Service Quotas – Per-Request Limits

Resource	Default Limit
Maximum execution time for single operation	5 Seconds
Maximum request size	2-Mb
Maximum response size	4-Mb
Maximum number of operations in a transaction batch	100

Service Quotas – Try Cosmos DB Free Limits

Resource	Default Limit
Duration of the trial	30 days
Maximum containers per subscription	1 (SQL, Gremlin, Table) 3 (MongoDB)
Maximum throughput per container	5,000
Maximum throughput shared-throughput database	20,000
Maximum total storage per account	10-Gb

Service Quotas – Free Tier Account Limits

Resource	Default Limit
Number of free tier accounts per Azure subscription	1
Duration of free-tier discount	Lifetime of the account
Maximum RU/s for free	400 RU/s
Maximum storage for free	5-Gb
Maximum number of shared throughput databases	5
Maximum number of containers in a shared throughput database	25

CHANGE FEED

Change Feed

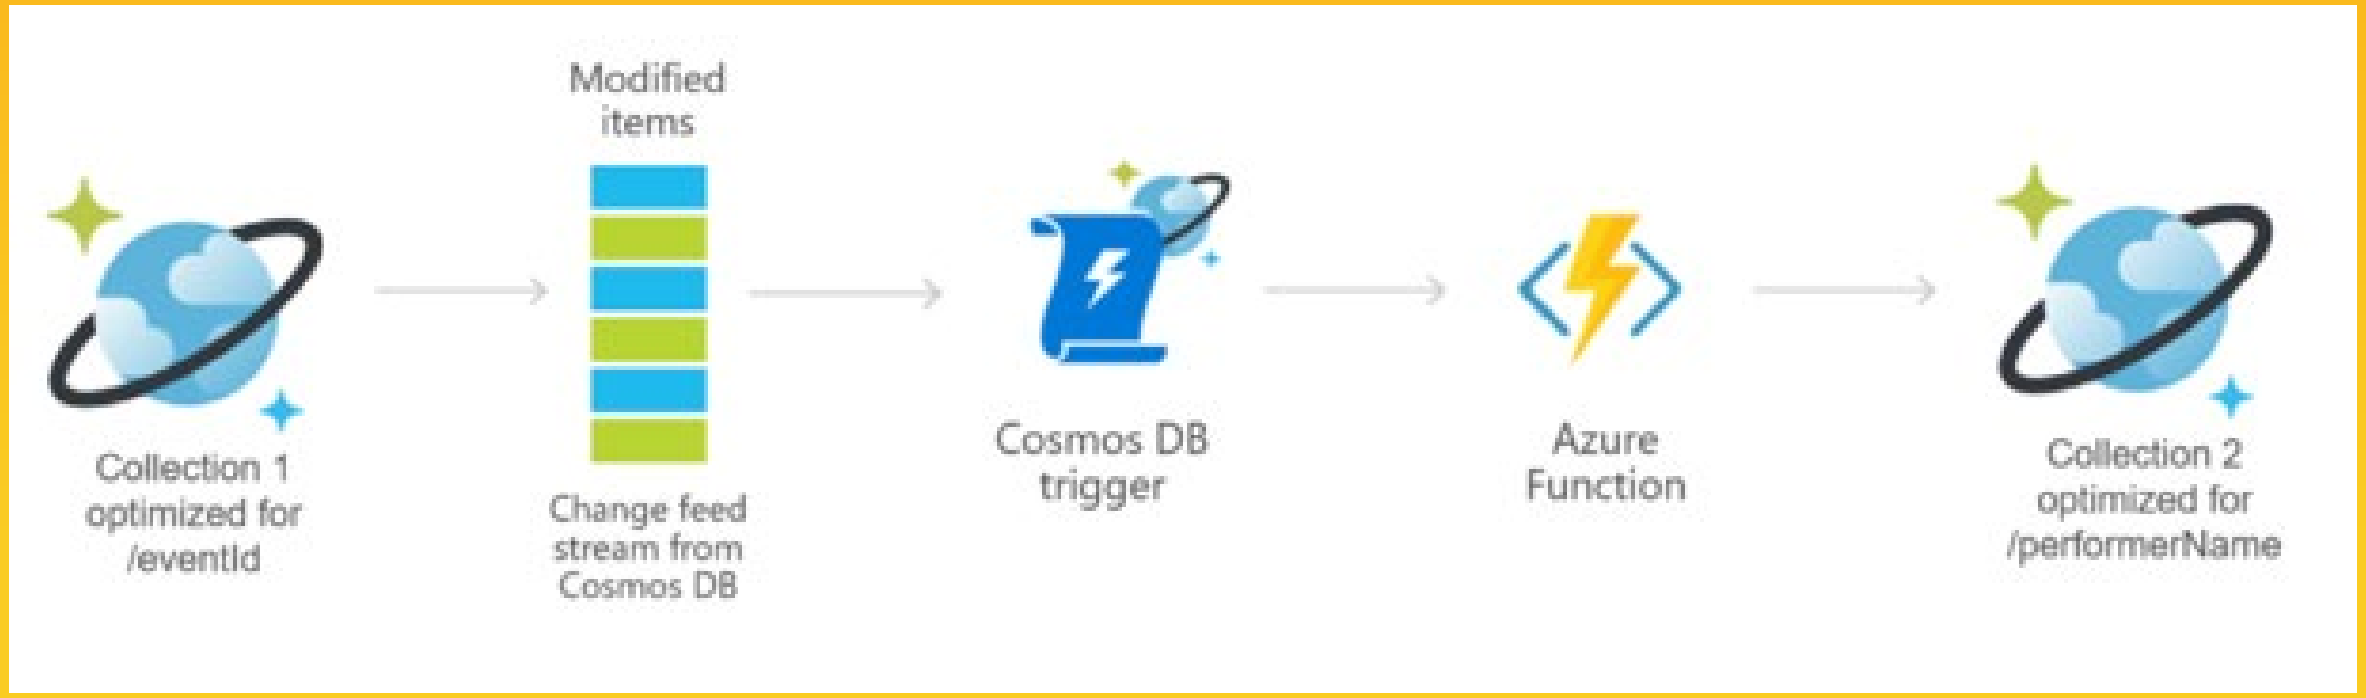
Modified items
Collection 1 optimized for /eventId
Change feed stream from Cosmos DB
Cosmos DB trigger
Azure Function
Collection 2 optimized for /performerName

Change Feed

Enabled by default

Change Feed

Enabled by default

Change Feed

Enabled by default

Includes insert and update operations

Change Feed

Enabled by default

Includes insert and update operations

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Sorted by order of modification

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Sorted by order of modification

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

No guaranteed order of logical partitions

Sorted by order of modification

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

No guaranteed order of logical partitions

Sorted by order of modification

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Synchronized from any point-in-time

Sorted by order of modification

No guaranteed order of logical partitions

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Synchronized from any point-in-time

Sorted by order of modification

No guaranteed order of logical partitions

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Changes available in parallel for logical partitions

Sorted by order of modification

No guaranteed order of logical partitions

Synchronized from any point-in-time

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Changes available in parallel for logical partitions

Clients manage checkpointing logic

Sorted by order of modification

No guaranteed order of logical partitions

Synchronized from any point-in-time

Change Feed

Enabled by default

Includes insert and update operations

Each change appears exactly once

Clients manage checkpointing logic

Applications can request multiple change feeds

Sorted by order of modification

No guaranteed order of logical partitions

Synchronized from any point-in-time

Changes available in parallel for logical partitions

Change Feed Options

Change Feed Processor

Azure Function

DEMO
BUILD AN APP

Thank You!

chadgreen@chadgreen.com

TaleLearnCode

www.ChadGreen.com

ChadGreen & TaleLearnCode

ChadwickEGreen