



# Azure Storage Overview and New Capabilities

Dennis Mulder  
Senior Program Manager  
Azure Storage



# What we'll cover today

Azure Storage Introduction

Azure Storage Architecture

Azure Object (Blob) Storage Overview

Storage Solutions

Working with Object Storage

# Azure Storage

## Foundational Building Block of Azure

Azure Services: SQL Data Warehouse, HDInsight, Data Lake Store, Event Hubs, IoT Hubs...

Microsoft Services: Office 365, OneDrive, Xbox, Skype...

### Hyper Scale

> 120 Trillion Objects, > 19 million transactions per second

### Durable

Never lose your data. Multiple redundancy options. Automatic data checks

### Secure

Encryption at Rest. Client side Encryption. Integration with KeyVault

### Highly Available

Fault tolerance to hardware/software issues. Automatic load balancing

### Open

REST API, Open sourced Client Libraries – .NET, Java, C++, Python, Node.js, iOS, Android, Xamarin...

### Hybrid

Extensive partner ecosystem. Azure Stack for private/hosted clouds.

# Azure Scale



32 GA, 6 coming soon – Storage is available in every region







# Azure Storage Services

IaaS



Storage



Virtual machines



Networking

PaaS



Existing frameworks



Web and mobile



Microservices



Serverless Compute

## Disks

Persistent disks for Azure IaaS VMs

Premium Storage Disks option: SSD based, high IOPS, low latency

## Files

Fully Managed File Shares in the Cloud

SMB and REST access

"Lift and shift" legacy apps

## Objects

Highly scalable, REST based cloud object store

Block Blobs: Sequential file I/O

Cool Tier Available

Page Blobs: Random-write pattern data

Append Blobs

## Tables

Massive auto-scaling NoSQL store

Dynamic scaling based on load

Scale to PBs of table data

Fast key/value lookups

## Queues

Reliable queues at scale for cloud services

Decouple and scale components

Message visibility timeout and update message to protect against unreliable dequeuers

**Built on a unified Distributed Storage System**

Durability, Encryption at Rest, Strongly Consistent Replication, Fault Tolerance, Auto Load-Balancing



# Azure Storage Platform Capabilities

## Scalability

Per account: 500TB of data , 20000 tps, 10Gbps ingress/20Gbps egress

## Durability Options

Locally Redundant (LRS), Zone Redundant (ZRS), Geo-Redundant (GRS)  
Read access from geo-secondary (RA-GRS)

## Availability

99.9% for reads/writes, 99.99% for reads with RA-GRS. Backed by SLA.

## Access Control

Symmetric shared key (AuthN)  
Shared Access Signatures (SAS) - Delegated AuthZ  
Public (blob service only)

## Analytics

Rich perf-counter like metrics and detailed logs

# Azure Object (Blob) Storage

# What is the Blob Storage Service?

Object Storage: Store and serve unstructured data

App and Web scale data

Big Data from IoT, Genomics, etc.

Backups and Archive

## Key Features

Store PBs of data

Cost Effective

Dynamically scale up on bandwidth and TPS

Strongly consistent

Highly Available - 99.9% SLA, 99.99% for reads (RA-GRS)

Durable – Choose level of durability: LRS, ZRS, GRS

# The Dilemma

I understand the benefits of the cloud

I want to begin a journey

I don't want to make new on-premises investments



I can't give up existing features

I can't re-platform my workload


How do I preserve my existing investments?



# Azure Storage + ISV Offerings = Solutions



COMMVAULT®











COHESITY



# Featured Partner Solutions

## Providing Backup and Disaster Recovery with Azure

Partner	Product Solution	Key Workloads
 <b>COMMAVAULT</b>	<a href="#">CommVault</a>	Backup and DR, Workload and Data Migration, Endpoint Data Protection
 <b>VERITAS</b>	<a href="#">Veritas NetBackup (in Beta)</a>	Backup and DR, Workload and Data Migration, Endpoint Data Protection
 <b>Hewlett Packard Enterprise</b>	<a href="#">HPE Data Protector</a> (October 9.08 release) <a href="#">HPE VM Explorer</a> (September 6.2 release)	Backup and DR
 <b>NetApp</b>	<a href="#">NetApp AltaVault</a> <a href="#">NetApp AltaVault Cloud-Based Appliance</a>	Backup and DR
 <b>EMC<sup>2</sup></b>	<a href="#">EMC Isilon CloudPools</a> <a href="#">EMC Avamar Virtual Edition</a> <a href="#">EMC Data Protection Suite CloudBoost</a>	Hybrid Storage Backup and DR Long-Term Retention
 <b>VEEAM</b>	<a href="#">Veeam® Cloud Connect™ for the Enterprise</a> <a href="#">Veeam® Cloud Connect™ for Service Providers</a> <a href="#">Veeam® Direct Restore to Azure</a>	Backup and DR, Workload and Data Migration, Endpoint Data Protection
 <b>DELL</b>	<a href="#">Dell Rapid Recovery Replication Target for Azure</a>	Backup and DR, Archiving
 <b>Datacastle</b>		

# Solution: Backup and Disaster Recovery

## Backup

On-premises data  
Cloud data and VMs

## Target

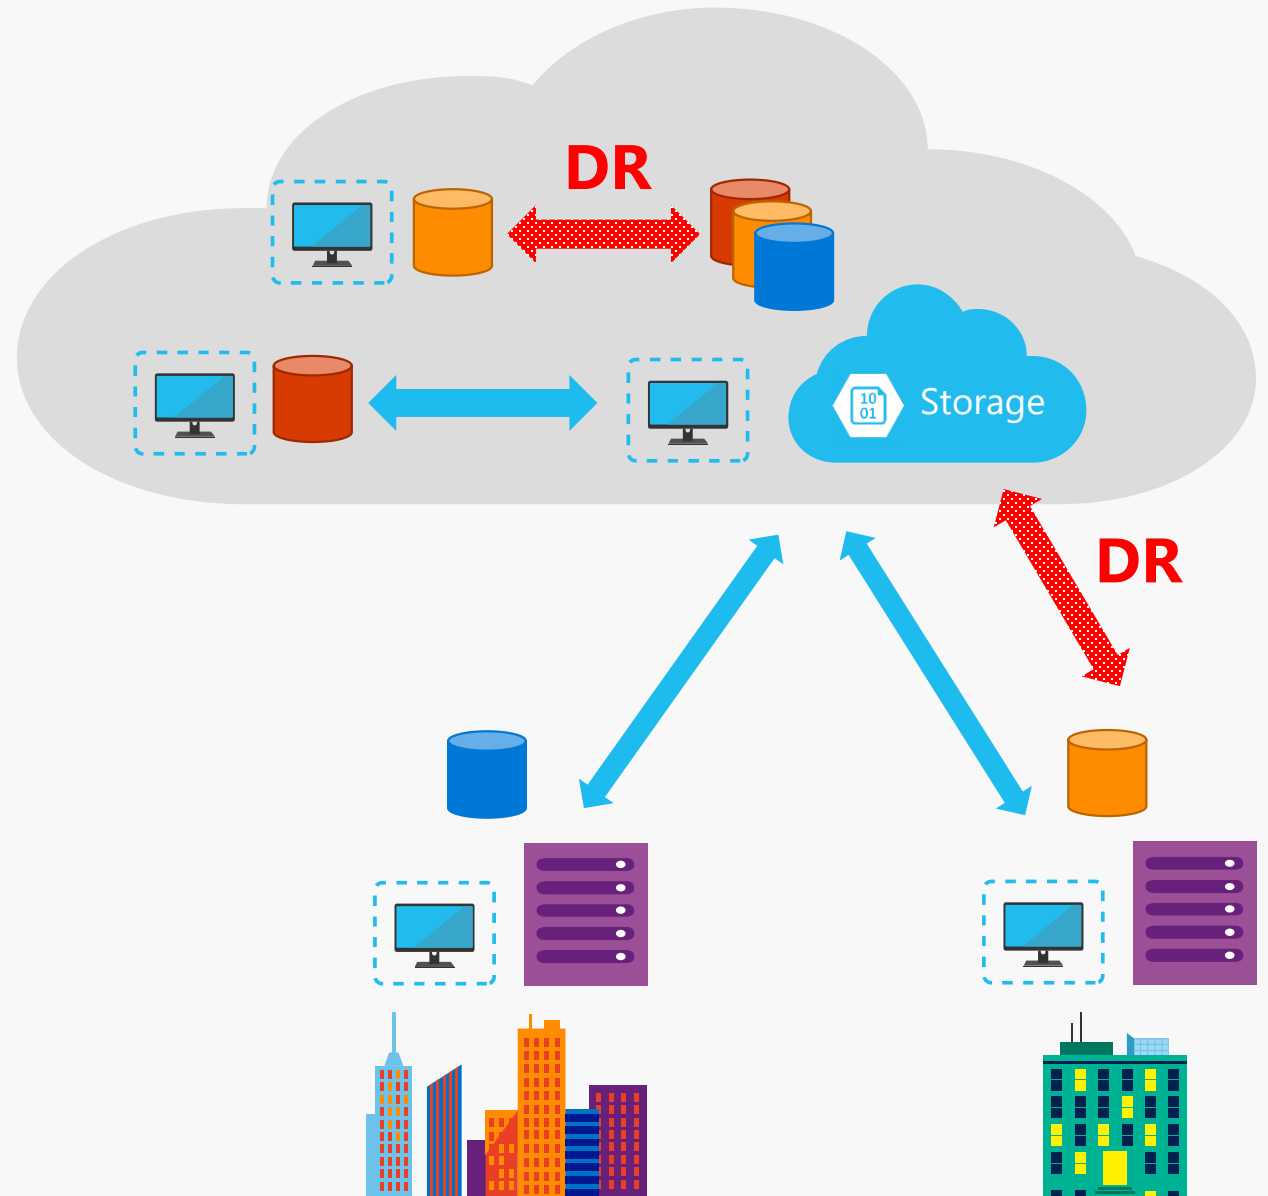
Object storage directly  
Ideal workload for cool storage  
Virtual Backup Appliance

## Disaster Recovery

Recover workloads in the cloud  
Recover data to alternate sites








## Migration

Move data into the cloud



# Featured Partner Solutions

## Global Data Access with Azure

Partner	Product Solution	Key Differentiators
	<a href="#">Enterprise File Service</a>	SMB/NFS, global file sharing/locking, global deduplication and compression, Extensive Snapshot Capabilities, Replication
	<a href="#">Panzura Global NAS Appliance</a>	SMB/NFS, global file sharing/locking, Compression and Deduplication, Extensive Snapshot Capabilities
	<a href="#">CloudFAST</a>	SMB, global file sharing/locking, intelligent file caching
	<a href="#">SoftNAS Cloud NAS</a>	Unified storage (iSCSI, SMB/NFS), Deduplication and Compression, Extensive Snapshot/Cloning Capabilities, Replication
	<a href="#">CloudArray</a> <a href="#">CloudPools</a>	iSCSI, SMB, NFS, Deduplication and Compression
	<a href="#">DataPlatform</a>	SMB, NFS, global data deduplication, snapshots, replication
	<a href="#">Cloud Storage Gateway</a>	SMB, NFS, AFP, FTP, WebDAV, rsync, iSCSI, CTERA Mobile, CTERA Sync desktop application, web browser





# Working with Blob Storage

# Types of Blobs


## Block Blobs

Most object storage scenarios  
Documents, images, video, etc.




## Append Blobs

Multi-writer append only scenarios  
Logging, Big Data Analytics output



## Page Blobs

Page aligned random reads and writes  
IaaS Disks, Event Hub, Block level backup



# Cloud Tiered Storage

## Two Tiers

Hot – for commonly used data

Cool – for rarely used data

API is 100% identical; similar throughput and latency

Durability options: LRS, GRS, RA-GRS

Availability: Cool - 99%, Hot - 99.9%. Higher for RA-GRS reads

## Pricing to match your workload

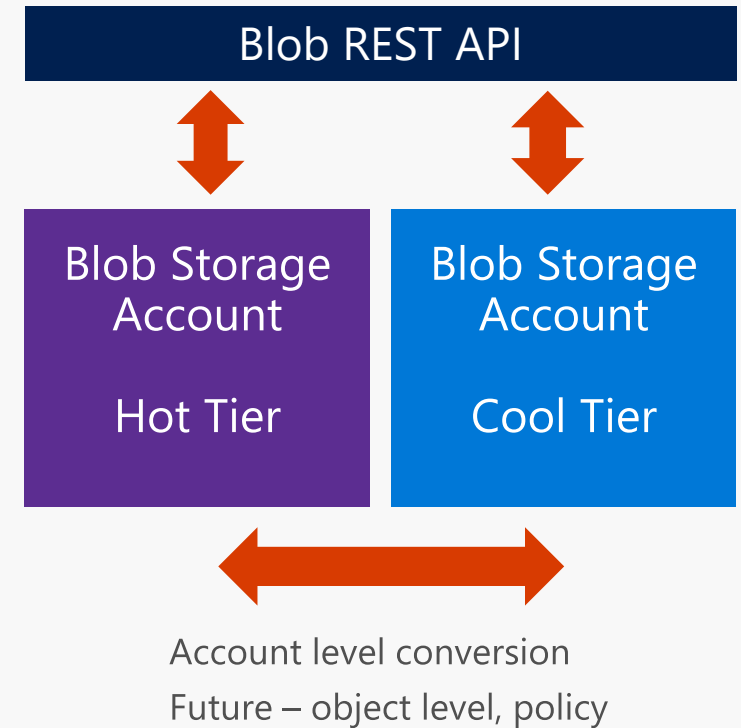
Hot: Lower access prices for frequent use

Cool: Lower storage prices for high volume

## Switch account tiers as needed

No charge for Hot to Cool switch

Future – Object level switch with automatic policy based management



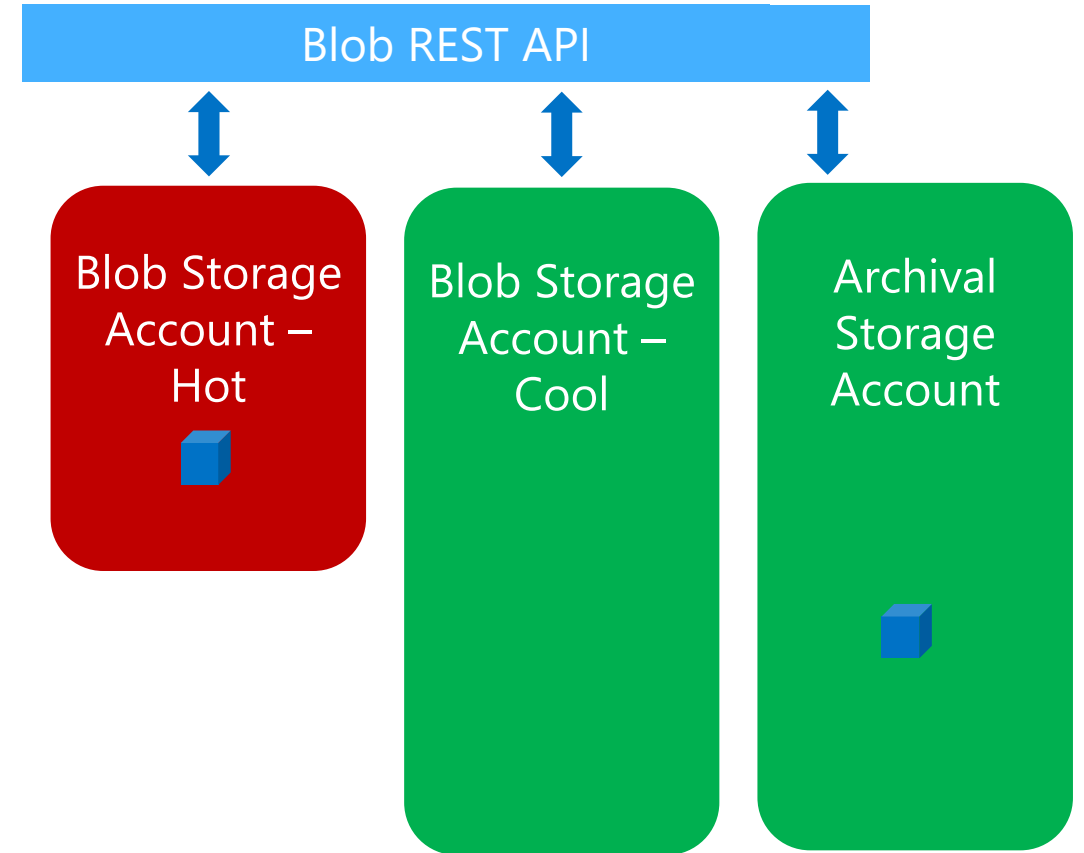
# Archival Storage

## New archival account for object storage\*

- Offline data – Hours to retrieve
- Consistent object level APIs and separate archival API
- Highly Durable – multiple copies
- Lowest storage cost, higher access costs

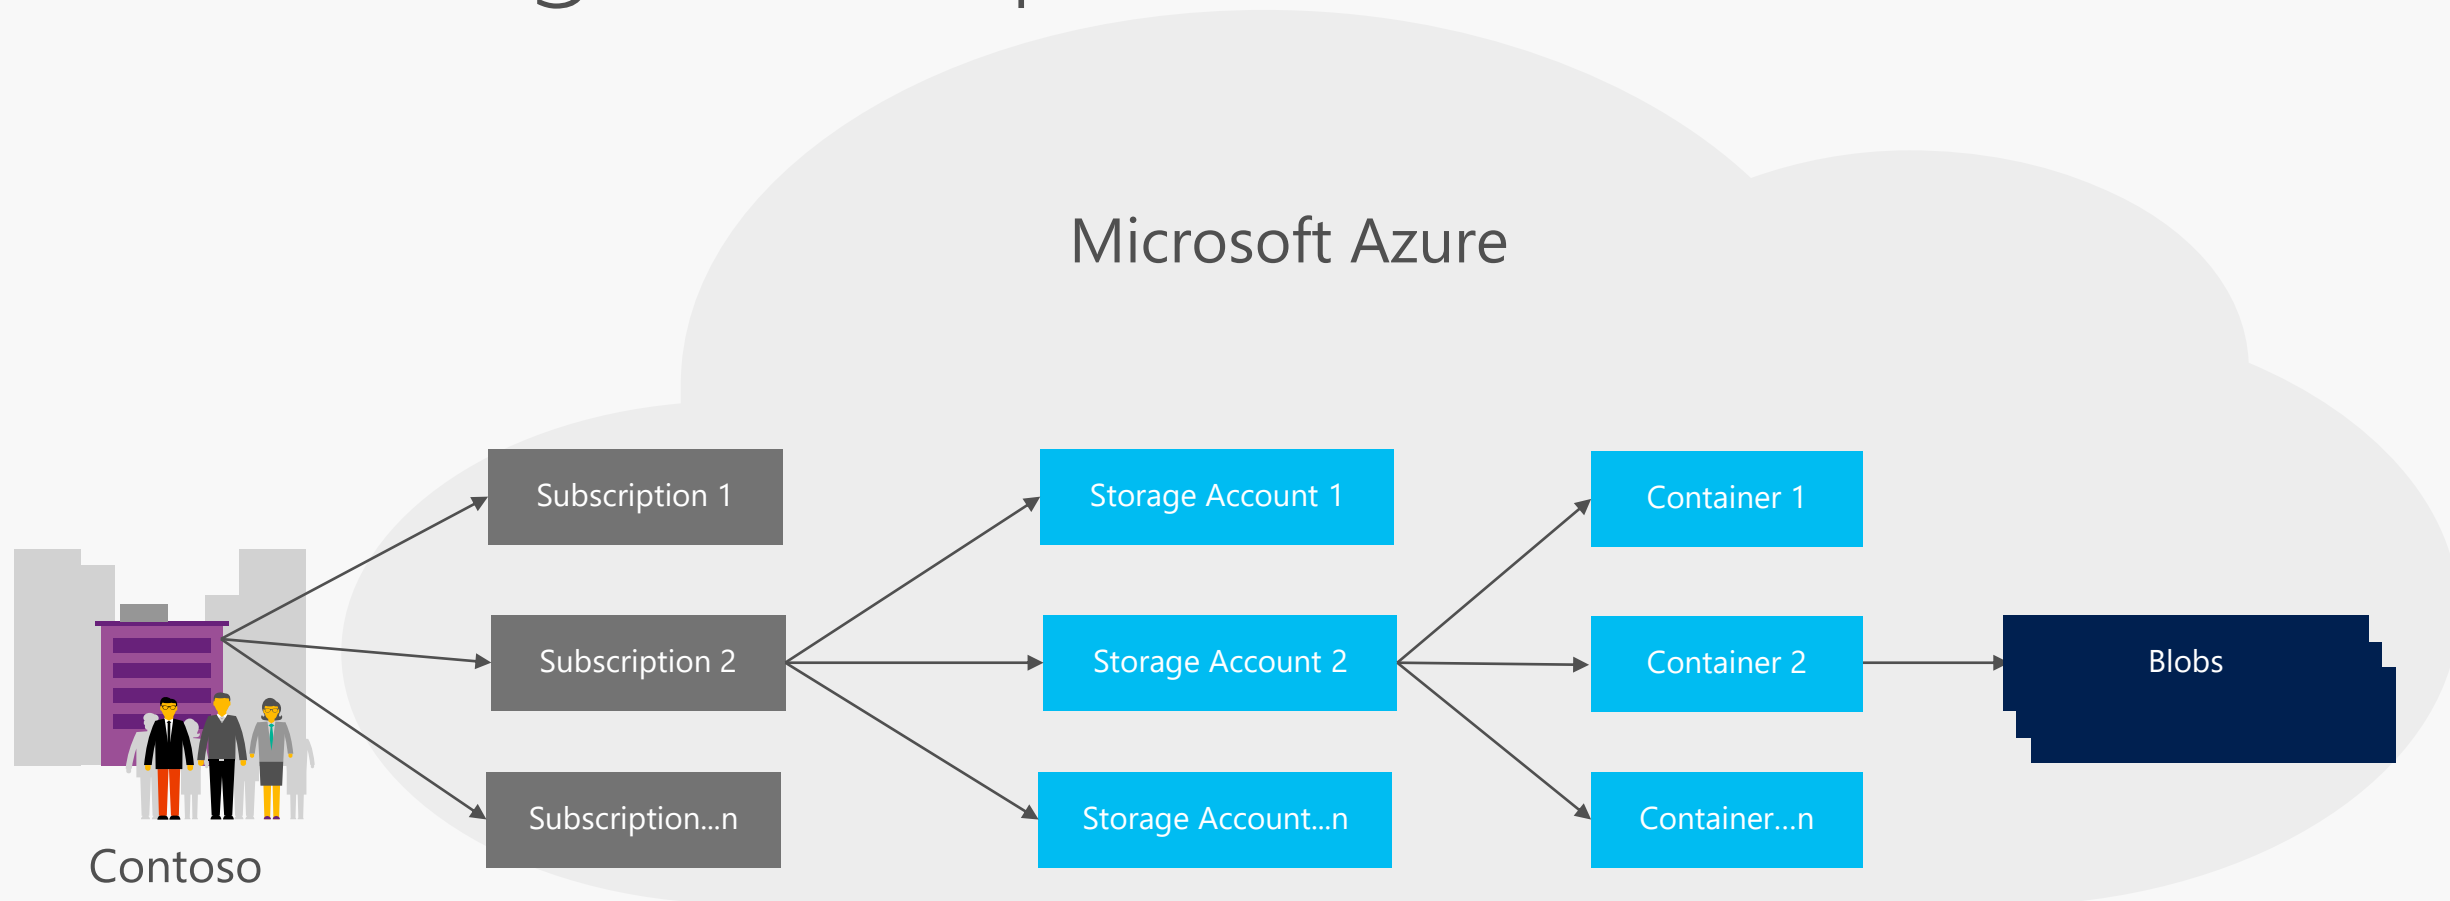
Lifecycle management between hot, cool and archival

Preview: Q1 CY17



\*Design shown is subject to change

# Blob Storage Concepts



<http://<StorageAccount>.blob.core.windows.net/<Container>/<Blob>>

<http://foo.blob.core.windows.net/images/2016/September/Annika.jpg>

Account Name

Container Name

Virtual directories + Blob

# Demo

Block Blobs: Getting Started with Code and Tools

# Blob Structure

## Standard Properties

Blobtype: BlockBlob|PageBlob|AppendBlob

Etag & Last-Modified – Used for conditional operations

Content-Length & Content-MD5 – Size of blob in bytes and checksum

## User Defined Metadata

Up to 8KB of unique names and associated values (valid HTTP headers)

No partial updates – Retrieve, then set

Can set on containers as well

## Blob Data

Up to 200GB per blob, increasing to 5TB soon

# Uploading Block Blobs

## PutBlob

Single REST call for data up to 64MB in size  
New PutBlob on existing object name will replace it

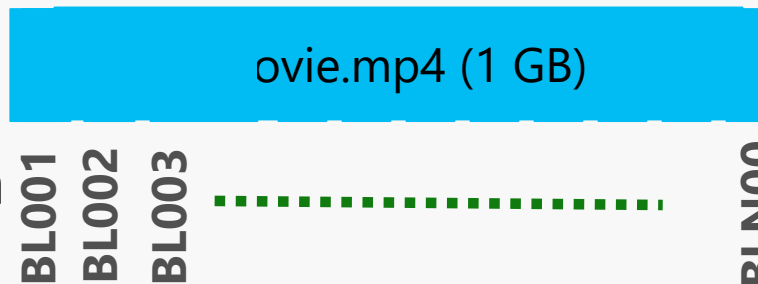
Image.jpg  
(<64 MB)

REST Calls

```
PutBlob(Image.jpg)
```

## PutBlock/PutBlockList

Split data to be uploaded into blocks (up to 4MB each)  
Efficient Continuation and Retry  
Parallel, out of order uploads  
Allows for intra-object de-duplication



REST Calls

```
name = "Movie.mp4";  
PutBlock(BL001);  
PutBlock(BL002);  
.....  
PutBlock(BLN00);  
PutBlockList(  
    BL001,...,BL00N);
```

## Block Blobs are mutable

New PutBlock calls to add to object or replace selective blocks  
Combine existing and new blocks to create a modified object



# Versioning

## Snapshots

Used to create a read only copy of a blob at a point in time

Same name as blob with a datetime stamp

Can be read, copied/restored or deleted

CopyBlob to make a write-able new blob from a snapshot

CopyBlob to restore a snapshot over the current version

Snapshots are not copied with a normal CopyBlob operation

Blob Deletion requires deletion of all snapshots

# Concurrency

## Optimistic concurrency – Timestamps/ETags

Timestamp based – If-Modified-Since and If-Unmodified-Since

ETag based – If-Match and If-None-Match

Conditional update with supplied Timestamp or ETag will fail if conditions not met

## Pessimistic Concurrency - Leases

Lease Blob for exclusive write and delete access

15-60s lease duration (can be renewed) or infinite lease (locks)

Can change lease id to acquire ownership in a chain/workflow

Can also acquire on containers to prevent container deletion

## Last Writer wins

# Data Movement

## AzCopy tool simplifies data migration at scale

Efficient means of copying millions of files – large or small

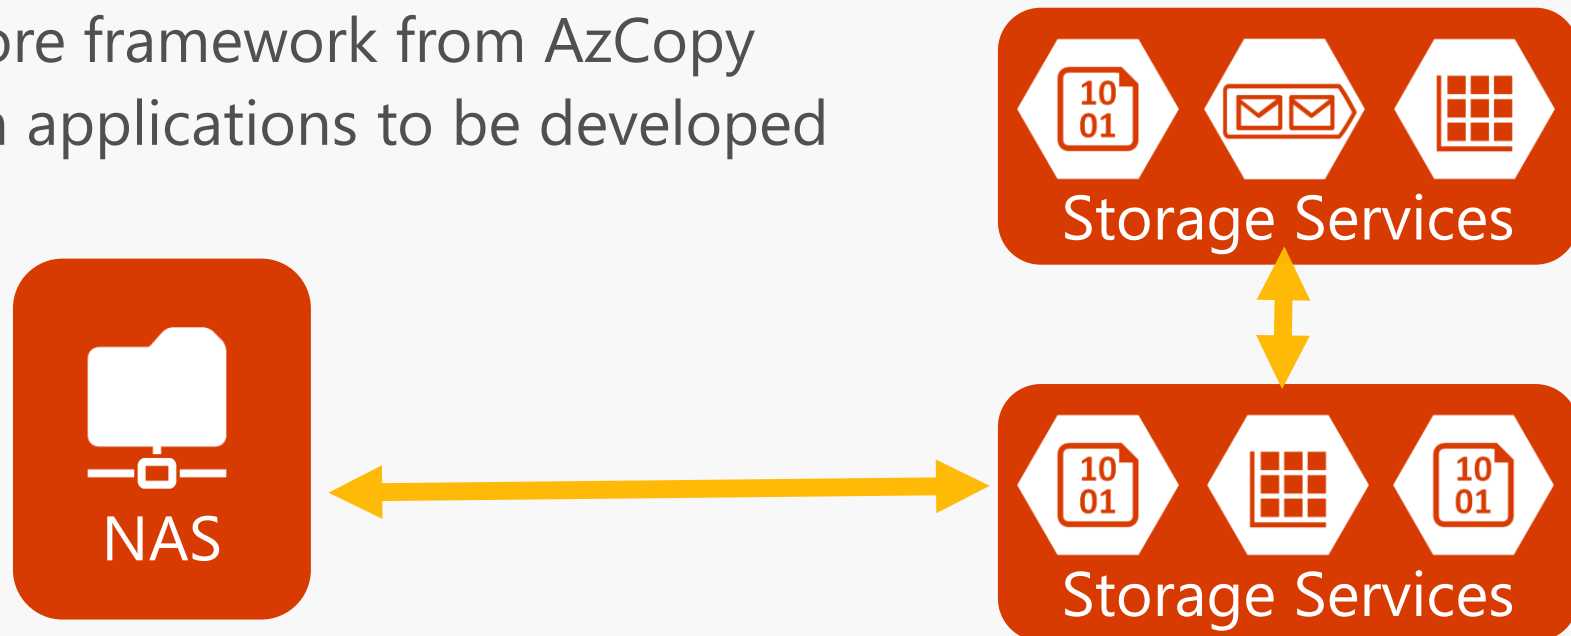
Includes journaling for reliability

Supports Blob Storage, Table Storage and File Storage

## Data Movement library supports developers

Open source core framework from AzCopy

Enables custom applications to be developed



# Delegated Authorization

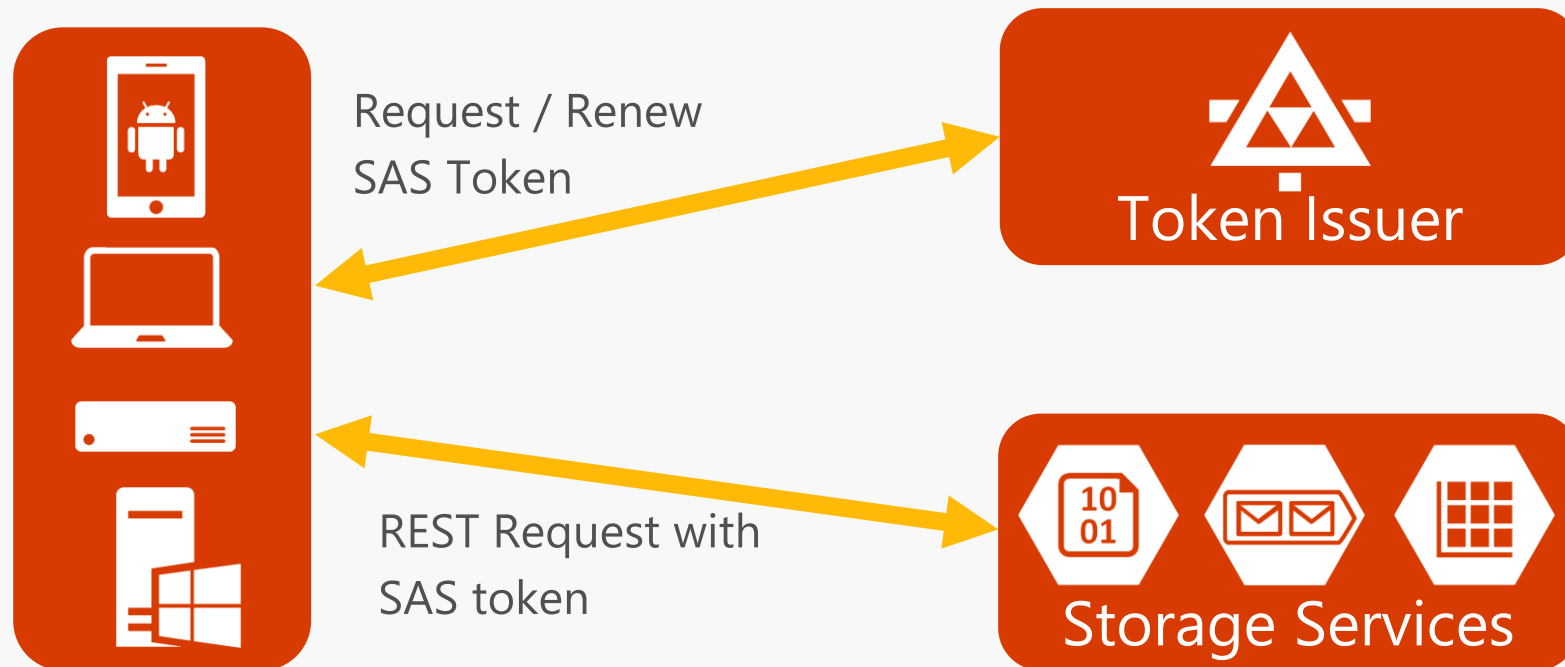
## SAS Tokens enable direct access by client apps

Supports least privilege principle

Removes the need to manage a data access layer

Mobile / IOT – Enable millions of untrusted clients

Web / Javascript – CORS + SAS enables website support



# Full Text Search

## Azure Search - Full-text search for Blobs (& Tables)

Index Data: Formats: PDF, Office, HTML, XML, ZIP, plain text, CSV, JSON

Index Metadata: Storage Metadata only or Storage & Custom Metadata

<https://azure.microsoft.com/en-us/documentation/articles/search-howto-indexing-azure-blob-storage/>

Wrap Up

# Summary

## Azure Blob Storage:

Highly Scalable, Durable and Available

Secure, Open and Multi-platform

Hybrid

Partner Ecosystem for Turnkey Usage

Cost Effective

Documentation: <http://aka.ms/AzureStorageDocs>

Azure Storage Feedback email:

AzureStorageFeedback@microsoft.com

# Microsoft Ignite Conference Sessions

[Dive into scenarios and customer success stories with Azure Storage](#)

[Build tiered cloud storage in Microsoft Azure](#)

[Optimize IaaS VM Storage using Azure Disks and Files](#)

[Secure your data in the cloud with Microsoft Azure Storage](#)

\*Search on Channel9 for Azure Storage Ignite 2016



