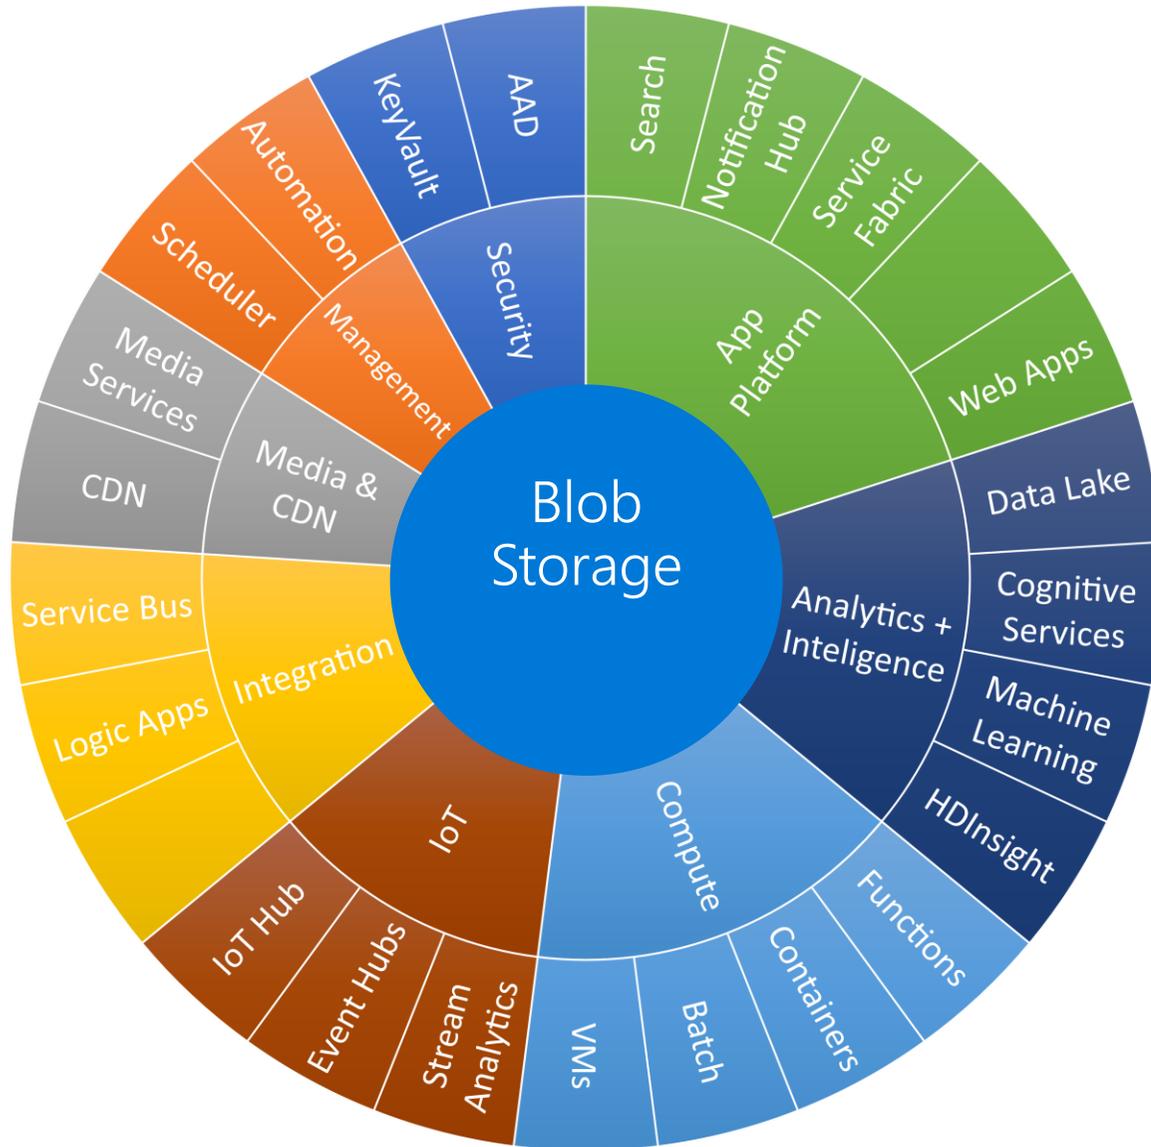


Azure Storage

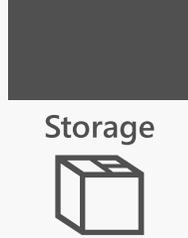
Azure Ecosystem and Blob Storage



Broad integration for Blobs across Azure services

Enables many scenarios

Azure Storage



IaaS



Storage



Virtual machines



Networking

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

Sync with on-premises (New!)

Blobs

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

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

>30 million transactions per second, trillions of objects

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

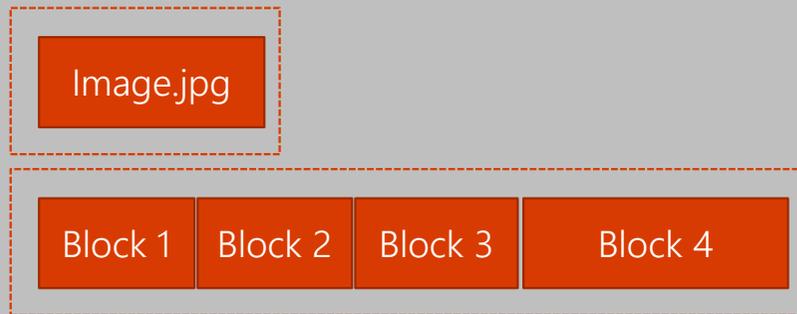
Extensive partner ecosystem. Azure Stack for private/hosted clouds.

Blob Storage Service



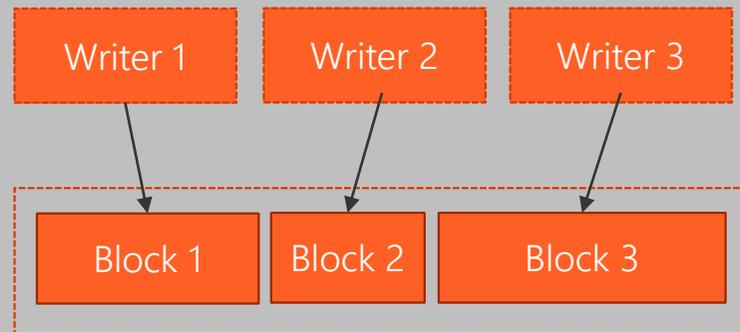
Block Blobs

Most object storage scenarios



Append Blobs

Multi-writer append only scenarios



Page Blobs

Page aligned random reads & writes



What is the Blob Storage Service?

Azure's Object Storage platform

Store and serve unstructured data

App and Web scale data

Backups and Archive

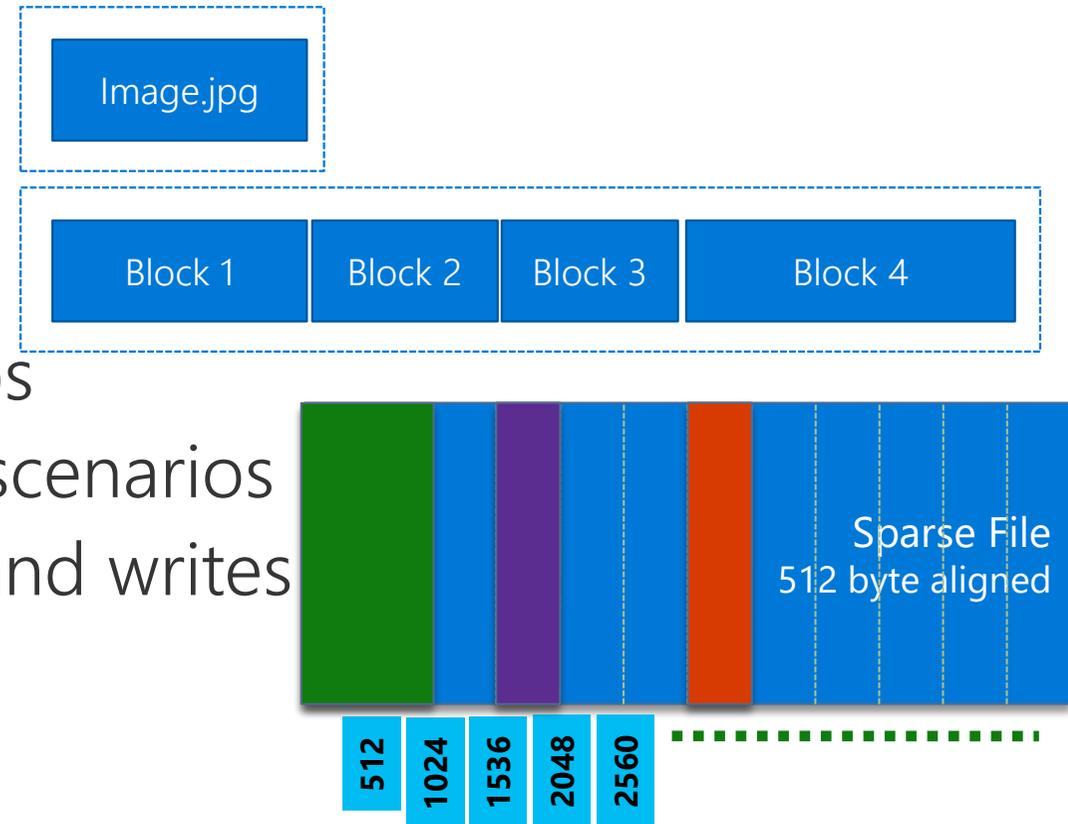
Big Data from IoT, Genomics, etc.

Types of Blobs

Block Blobs - Most object storage scenarios

Append Blobs - Multi-writer append only scenarios

Page Blobs - Page aligned random reads and writes



Tiered Storage

Introducing Blob-Level Tiering

Individual blob can move between tiers

All tiers of blobs co-exist in the a storage account

New Storage Tier – Archive Storage

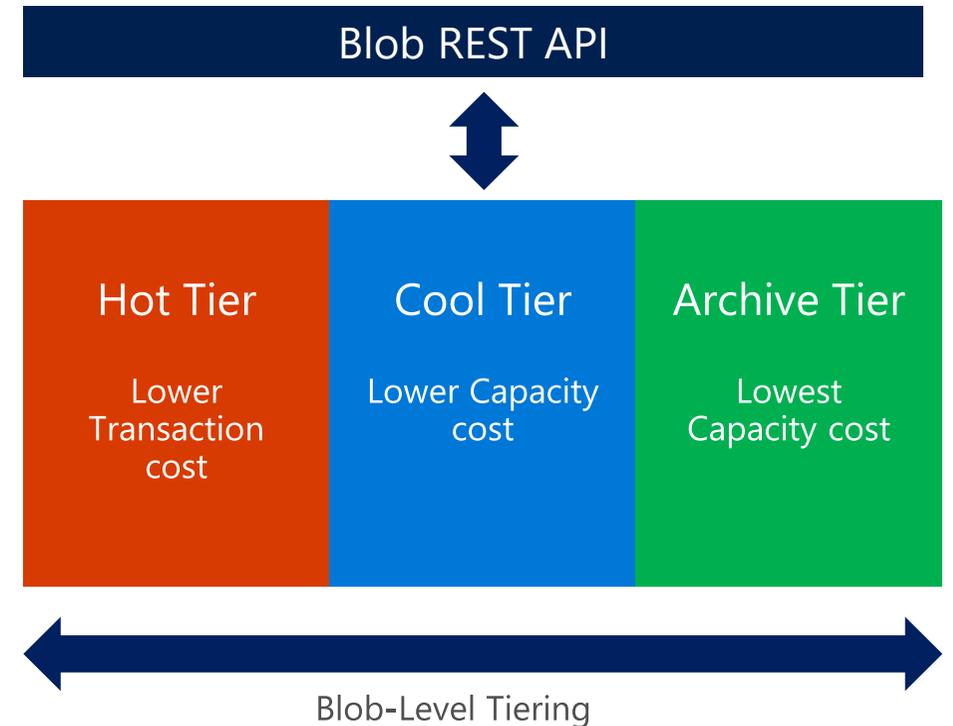
Cold storage for long term data

Retrieval latency is hours

Consistent API Among Storage Tiers

Access through Blob REST API

Support direct writes to Archive



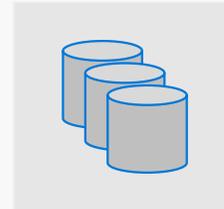
Azure Storage Durability



Typically >300mi



Async



Primary

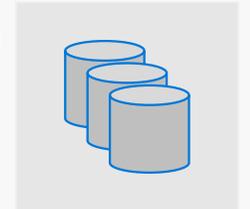
Secondary



Typically >300mi



Async



Primary

Secondary

LRS

3 replicas, 1 region
Protect against disk, node, rack failures
Write is ack'd when all replicas are committed
Superior to dual-parity RAID

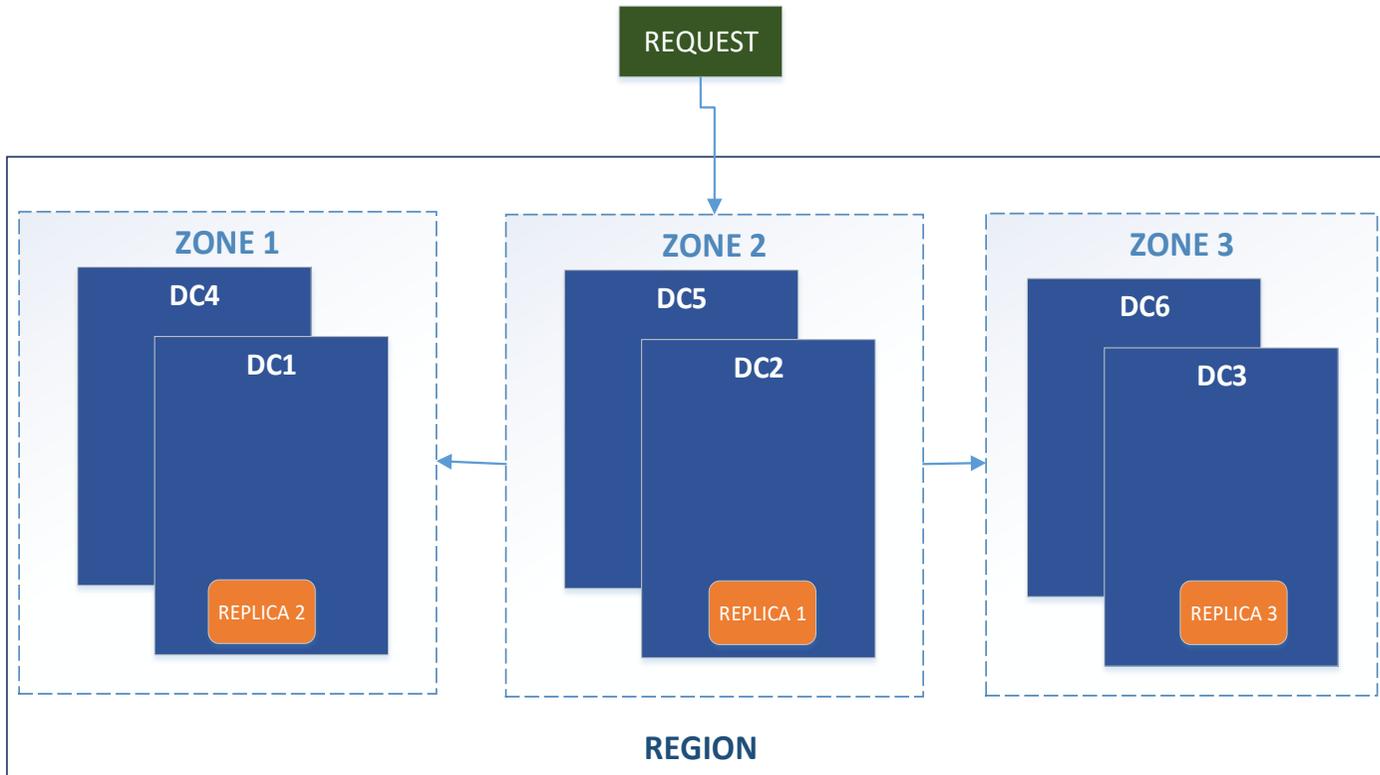
GRS

6 replicas, 2 regions (3/region)
Protects against major regional disasters
Asynchronous to secondary

RA-GRS

GRS + Read access to secondary
Separate secondary endpoint
RPO delay to secondary can be queried

NEW - Azure Zone Redundant Storage (ZRS)



Read / write resilience against single cluster / DC unavailability

Support for Blob, Table, File, Queue Storage

Public Preview in Q4 CY 2017 in multiple regions.

GA in H1 CY 2018

Synchronous data replication across [Azure Availability Zones](#) within region

LRS

Resilient to disk/node/rack failures

ZRS

Resilient to single cluster / datacenter outage

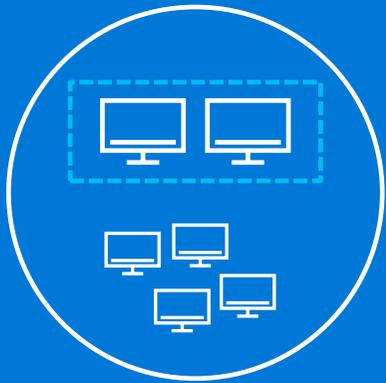
GRS

Resilient to regional outage

RA-GRS

Resilient to regional outage
Read access to second region

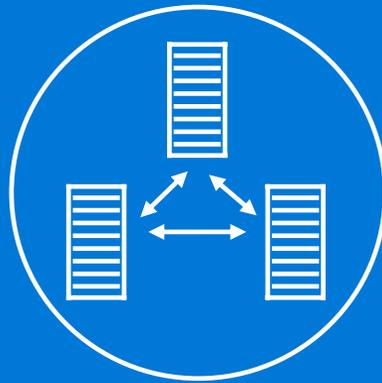
Azure Availability Zones



Availability Sets

High Availability protection from hardware failures in a datacenter.

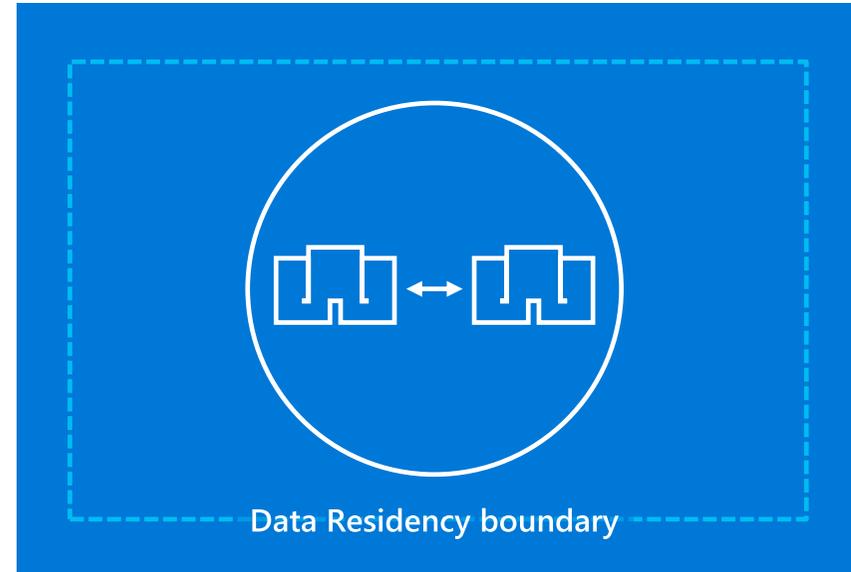
Hardware



Availability Zones

High Availability protection against hardware, network, cooling, and power failures within a region.

Datacenter



Region Pairs

Protection for your data and applications from the loss of an entire region with GRS backup and Azure Site Recovery.

Region

Virtual Machine Storage Basics

Compute



Storage



Networking



Management



Premium disks

SSD based, high IOPS,
low latency

Up to 64 TB, 80000 IOPs

Appropriate for data base, file servers, and interactive applications that have a dependency on high IOPs and low latency storage.



Standard disks

Magnetic disks, HDD based.
Up to 500 IOPs/disk

Appropriate for web and application servers that do not have a dependency on high IOPs or low latency storage.



Files

"SMB File Share"

SMB 3.0

Shared storage (read/write multiple virtual machines), replace file servers, support legacy applications and support lift and shift style workloads

Azure Storage

Azure IaaS: Understanding Storage



Azure hosts multiple types of storage

Blobs

"Highly scalable, REST based cloud object store"

Tables

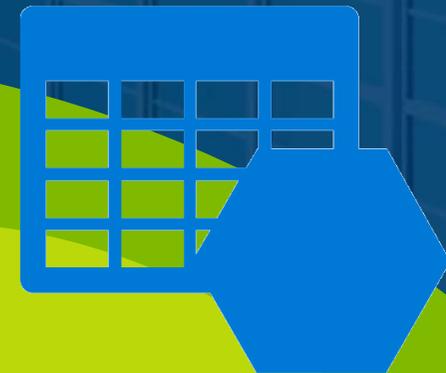
"Massive auto-scaling NoSQL store"

Queues

"Reliable messaging at scale for cloud services"

Files

"SMB Access to Azure Storage"



Azure Storage Offerings



Blobs

"Highly scalable, REST based cloud object store"

Tables

"Massive auto-scaling NoSQL store"

Queues

"Reliable messaging at scale for cloud services"

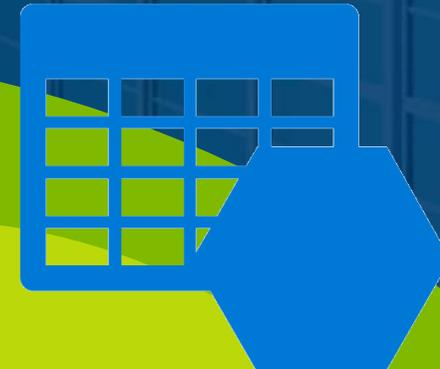
Files

"SMB Access to Azure Storage"

Block Blob – Optimized for streaming and storing cloud objects such as: Documents, Media and Backups

Page Blob – Optimized for IaaS disks and supporting random writes

Append Blob – Optimized for append operations such as logs



Azure Storage Offerings



Blobs

"Highly scalable, REST based cloud object store"

Tables

"Massive auto-scaling NoSQL store"

Queues

"Reliable messaging at scale for cloud services"

Files

"SMB Access to Azure Storage"

The Table service is Microsoft's NoSQL key/attribute store utilising a tabular format to store flexible datasets. This includes user data, address books, device information and many other metadata types.

Tables store data as collections of entities which are similar to rows. Each entity has a primary key and a set of properties that uniquely identify it along with a timestamp that is used to track when the entity was last updated.



Azure Storage Offerings



Blobs

"Highly scalable, REST based cloud object store"

Tables

"Massive auto-scaling NoSQL store"

Queues

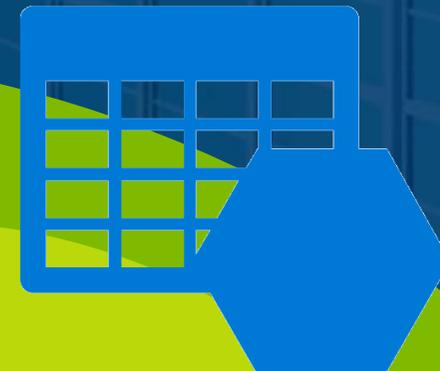
"Reliable messaging at scale for cloud services"

Files

"SMB Access to Azure Storage"

Azure Queues provide a reliable messaging solution for asynchronous communication between application components, these could be running in the cloud, on the desktop, on an on-premises server, or on a mobile device.

Messages typically have a small body with attributes such as a time to live which are used for configuring the service and representing tasks created by the front end "producer" which are in turn processed by the back end "consumer".



Azure Storage Offerings



Blobs

"Highly scalable, REST based cloud object store"

Tables

"Massive auto-scaling NoSQL store"

Queues

"Reliable messaging at scale for cloud services"

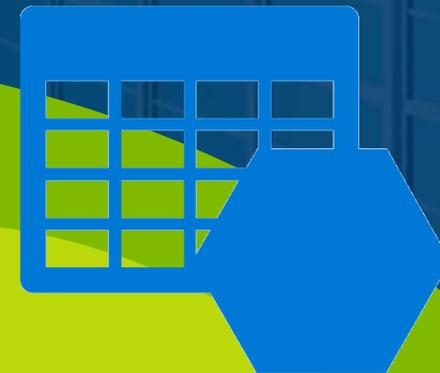
Files

"SMB Access to Azure Storage"

Azure File storage offers cloud-based encrypted SMB file shares allowing customers to easily migrate legacy applications to Azure quickly and without costly rewrites.

In addition to this, Virtual Machines and cloud services can share file data across application components and with on-premises applications via a simple File service REST API.

Azure File Storage offers availability, scalability and Geo Redundancy allowing you to confidently store your server data, log files, event data and backups in the cloud.



Xbox Live and Azure Storage Example



One of the earliest and largest users of the Azure platform

Blobs:

- Activity Feed
- Game DVR
- Leaderboards
- Player Stats
- Saved Games
- Social Graph
- Thunderhead

Tables:

- Achievements
- Comments
- Friends
- Matchmaking
- Multiplayer
- Push Notifications
- Security



Azure IaaS: Replication and High Availability



When you create a storage account, you must select one of the following replication options



LRS

Maintains three copies of your data within a single facility. LRS protects your data from normal hardware failures, but not from the failure of a single facility.



ZRS

Maintains three copies of your data replicated across two to three facilities, either within a single region or across two regions, providing higher durability than LRS.



GRS

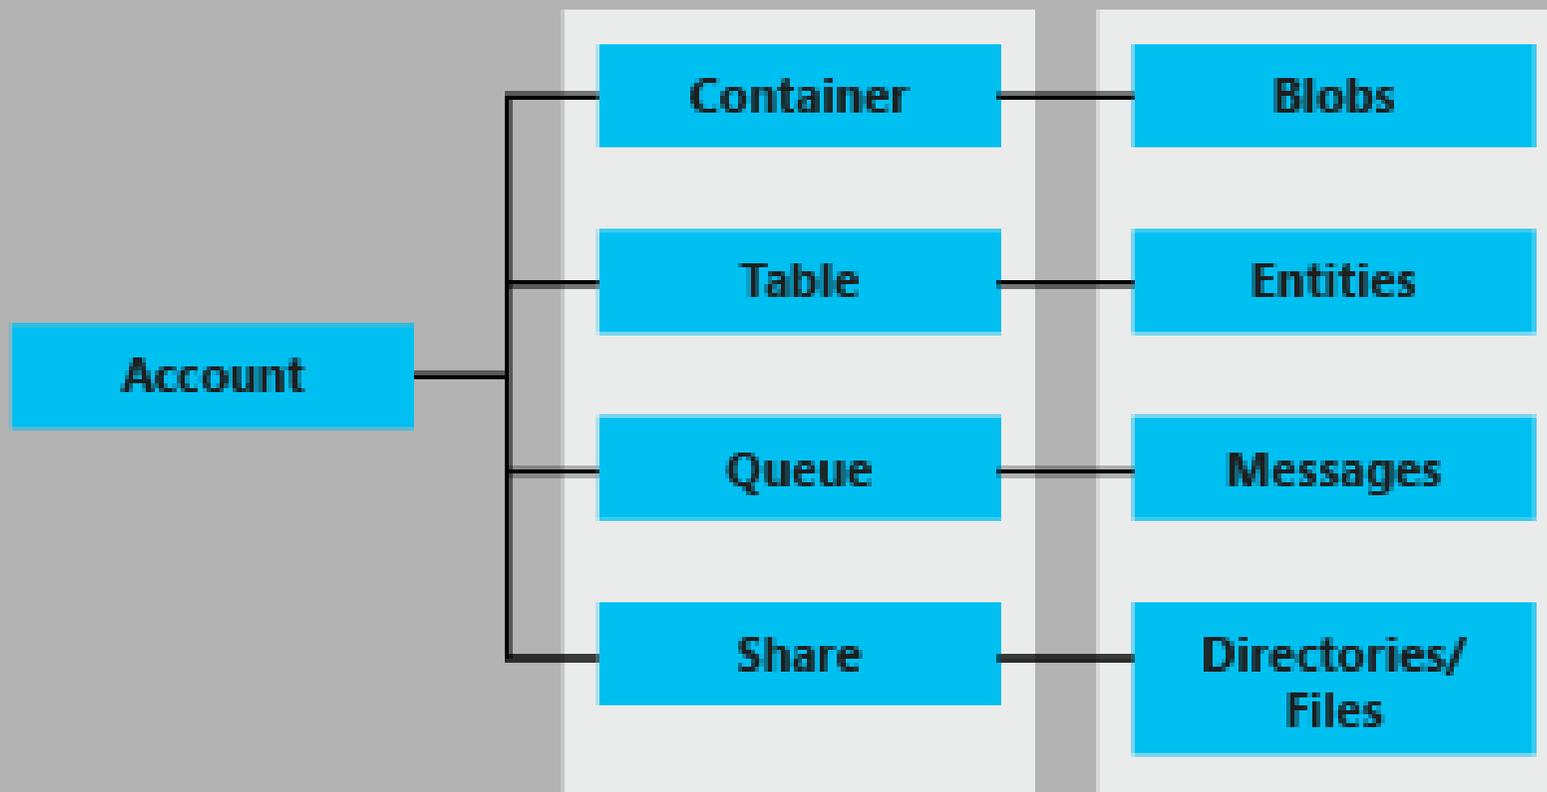
Maintains six copies of your data, three within the primary region, and three replicated into a secondary region hundreds of miles away. In the event of a failure at the primary region, Azure Storage will failover to the secondary region automatically.



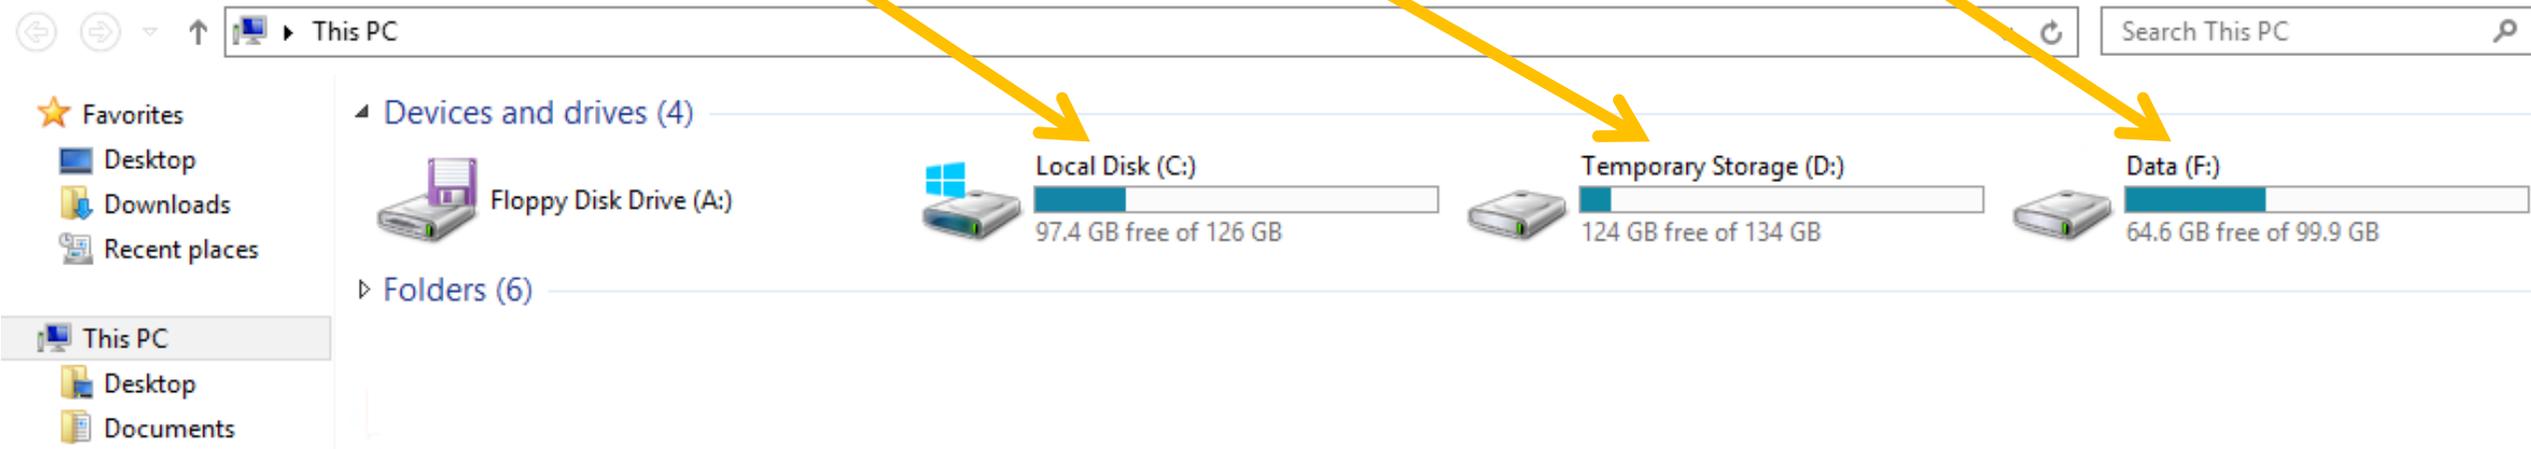
RAGRS

Replicates data to a secondary geographic location, and also provides read access to the data hosted in the secondary location.

Azure Storage Concepts

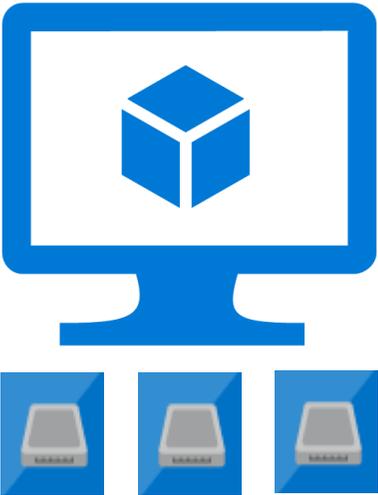


VM storage

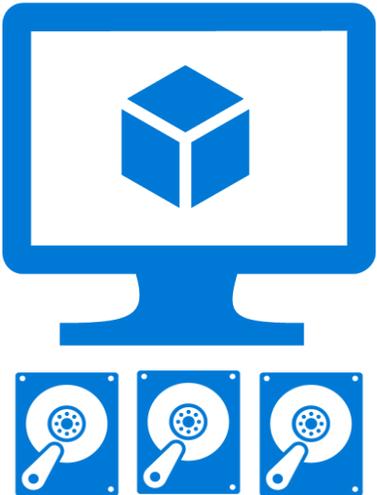


Azure Disks

Performance Tiers



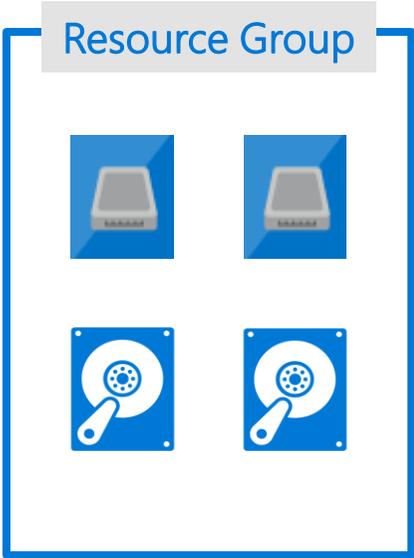
Premium Disks (SSD)



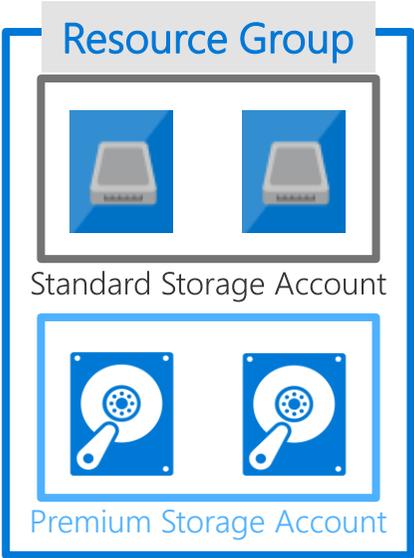
Standard Disks (HDD)

- ✓ Premium Disks: SSD based, provisioned performance
- ✓ Standard Disks: HDD based, cost effective

Management Options



Managed Disks



Unmanaged Disks
/Page Blob

- ✓ Managed Disks: highly available & manageable
- ✓ Unmanaged Disks: legacy with Storage Account

Differentiators

Industry leading
**ZERO % Annual
Failure Rate**

Enterprise grade
durability with 3 replicas

Best in class
High I/O
performance VMs

80,000 IOPS & 2,000 MB/s
Disk throughput per VM

< 1ms latency for
cached operations

Blob Cache technology
Up to 160,000 IOPS

Easy migration
from other
platforms

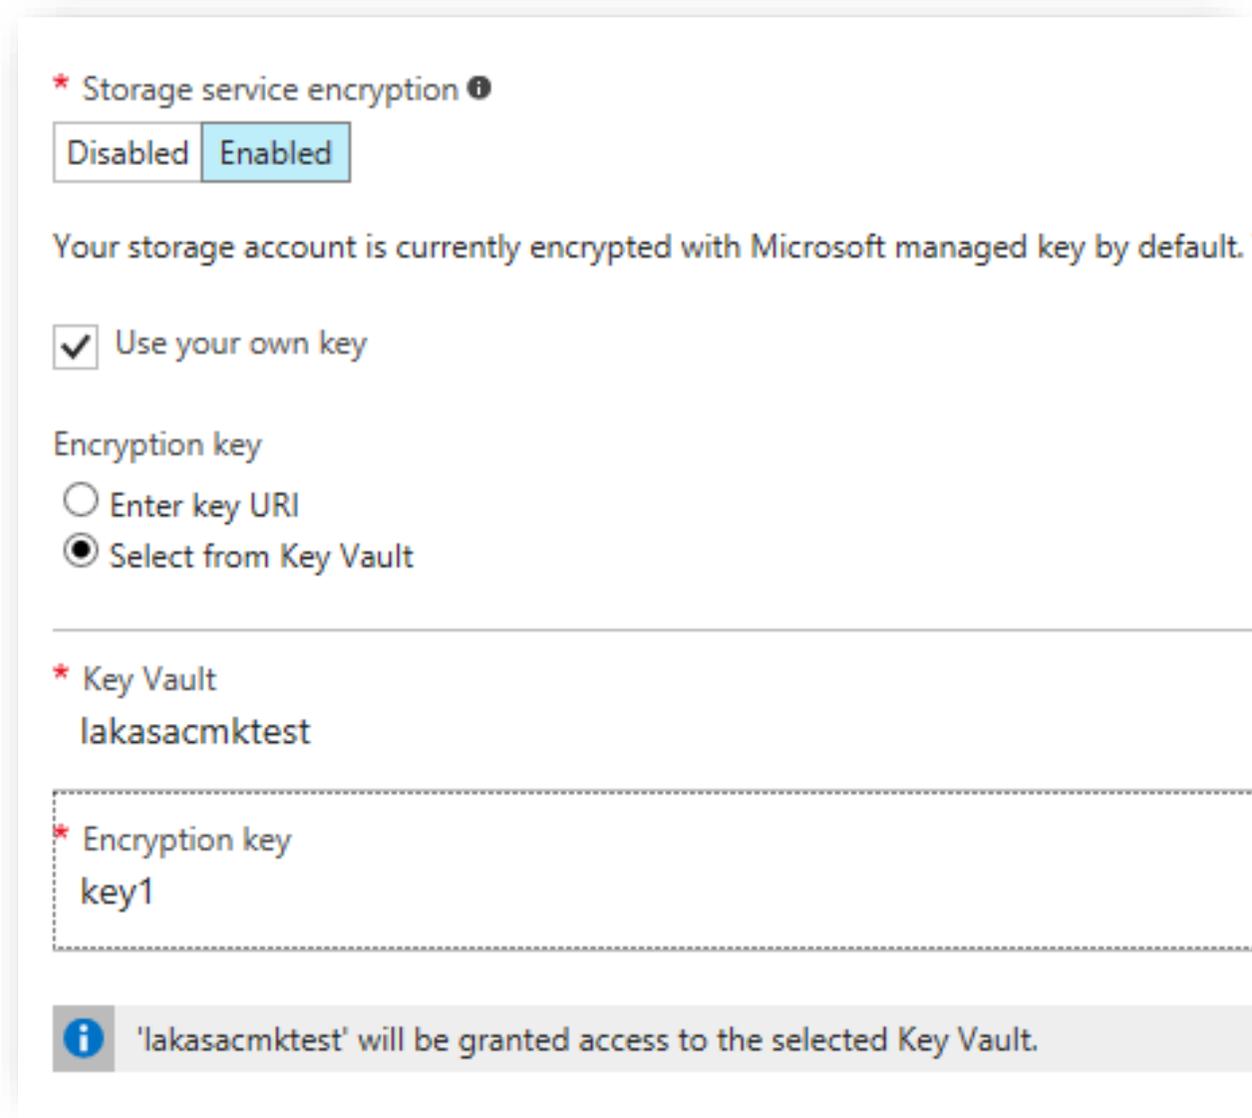
REST API support
Rich partner ecosystem

Encryption At Rest – Storage Service Encryption

Prevents data disclosure from physical disk compromise

Available for Blobs with MS managed keys

Customer managed encryption keys available



* Storage service encryption ⓘ

Disabled Enabled

Your storage account is currently encrypted with Microsoft managed key by default.

Use your own key

Encryption key

Enter key URI

Select from Key Vault

* Key Vault

lakasacmktest

* Encryption key

key1

ⓘ 'lakasacmktest' will be granted access to the selected Key Vault.

AAD Authentication and RBAC

Currently support AAD, OAuth and RBAC on Storage Resource Provider via ARM

Adding AAD and OAuth support to Blob Storage REST API

Integrating with Azure RBAC for Authorization

Role assignments down to container scope

RESOURCE TYPE	READ	WRITE	DELETE
 Microsoft Storage			
 Location			
 Name Availability	✓		
 Operations	✓		
 Storage Accounts	✓	✓	✓
 Blob Services	✓	✓	
 Containers	✓	✓	✓
 Blobs	✓	✓	✓

Archive Storage

New Storage Tier – Archive Storage

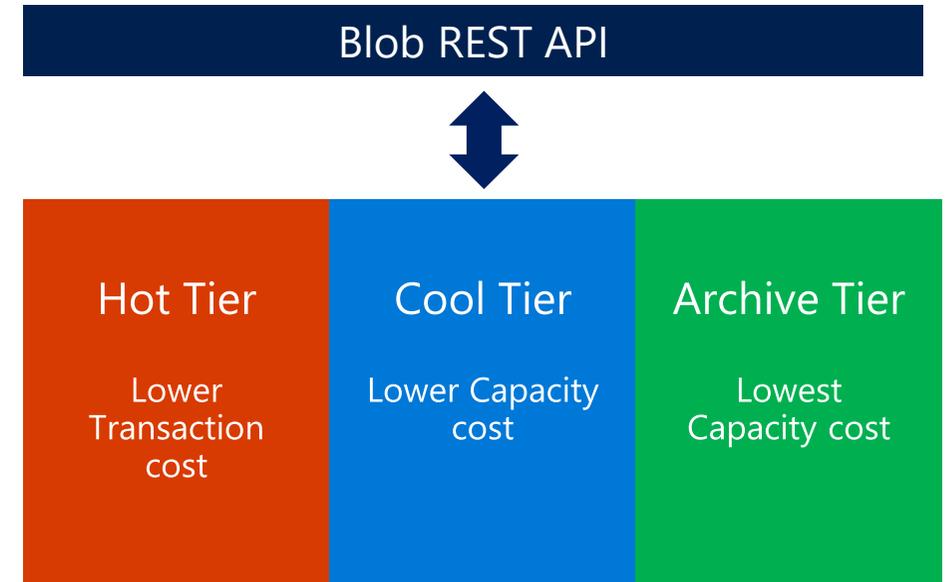
Cold storage for long term data

Offline data - Retrieval latency is hours

Consistent API Among Storage Tiers

Access through Blob REST API

Metadata for archive tier is always online



Blob Level Tiering

Introducing Blob-Level Tiering

Individual blobs can move between tiers

All tiers co-exist in the same storage account

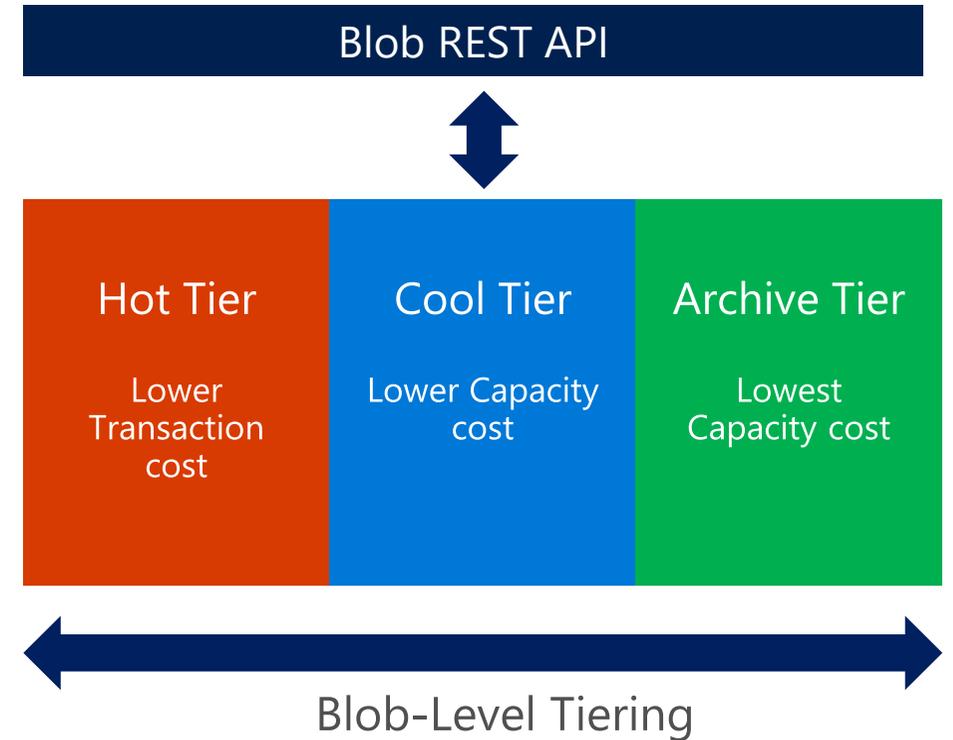
New API to set blob tier: *SetBlobTier*

Acknowledged immediately from service

Get APIs (*GetBlobProperties* and *ListBlobs*) return current tier and archive status

New headers "*x-ms-access-tier*" and "*x-ms-archive-status*"

Future: Automated Lifecycle Management



Tiered Storage

acsk8
Container

Upload Refresh Delete container Properties Access policy

Location: [acsk8](#)

Search blobs by prefix (case-sensitive)

NAME	MODIFIED	ACCESS TIER	BLOB TYPE	SIZE
CreateUiDefinition.json	2/23/2017 11:08:42 AM	Cool	Block blob	11.42 KiB
mainTemplate.json	2/21/2017 9:09:40 AM	Archive	Block blob	9.07 KiB
myStuff.txt	2/21/2017 9:09:40 AM	Hot	Block blob	9.07 KiB
myOtherStuff.txt	2/21/2017 9:09:40 AM	Default	Block blob	9.07 KiB

Tiered Storage

The image displays three sequential screenshots of the 'Blob properties' page in the Azure portal for a file named '02.04.15 SA Hussain OB 10.jpg'. The screenshots illustrate the process of changing the storage tier from 'Hot' to 'Archive'.

Common Properties (Visible in all screenshots):

- NAME:** 02.04.15 SA Hussain OB 10.jpg
- URL:** <https://kumailtestbn.blob.core.windows...>
- LAST MODIFIED:** 7/4/2017 11:40:20 AM
- TYPE:** Block blob
- SIZE:** 68.48 KiB
- ETAG:** 0x8D4C30C1FCCD7F3
- CONTENT-MD5:** -
- LEASE STATUS:** Unlocked
- LEASE STATE:** Available
- LEASE DURATION:** -

ACCESS TIER (Visible in all screenshots):

- Hot (Selected in the first two screenshots)
- Archive (Selected in the third screenshot)

Rehydration Warning (Visible in the third screenshot):

This blob is currently rehydrating. This may take several hours to complete and during this time your blob is inaccessible. Archive status 'rehydrate-pending-to-hot'.

Implement Storage

Blobs & Azure Files

Manage Access

Diagnostics, Monitoring,
Analytics

Implement SQL
Databases

Implement Recovery
Services



Implement Blobs

Highly scalable, REST interface based object store in the cloud

Data sharing – share documents, pictures, video, music, etc.

Big Data – store raw data/logs and compute/map reduce over data

Backups – data and device backups

Block blobs - (read/write/update blocks of data, great for sequential IO like files). Up to 200GB each.
Most cost effective storage.

Page Blobs - (read and write in 512 byte pages, sparse files and random access, e.g. for disks). Up to 1TB each

[AZCopy](#) cli high-performance uploading, downloading, and copying data to and from Microsoft Azure Blob, File, and Table storage.

FORMAT = `https://{storage account}.blob.core.windows.net{blob container} KNOW /Pattern,Source,Dest`

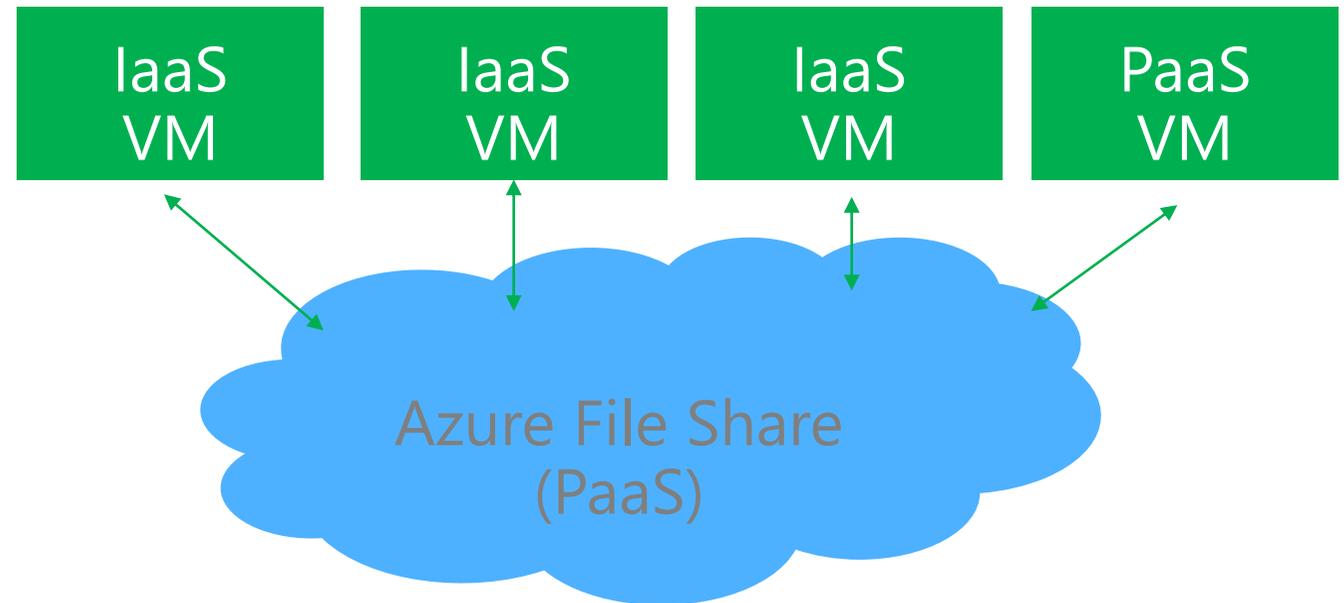
Azure Files

Shared Network File Storage for Azure

Availability, durability, scalability are managed automatically

Supports two interfaces: SMB and REST

aka.ms/Azure/Files



Azure Files - SMB 3.0/2.1 Protocol

Enables moving on-premises applications that rely on shared file storage to Azure

Azure VMs can "net use" to a share

[SMB Features NOT supported](#)

Natively supported by OS APIs, libraries, and tools

Windows (CreateFile, ReadFile, WriteFile, ...)

CRTs (fopen, fread, fwrite, ...)

.Net (FileStream.Read, FileStream.Write, ...)

Many more

Supports standard file system semantics

Storage Decision Matrix

How To Store	Why use This?
Azure Files	<ul style="list-style-type: none">• SMB Interface, Client Libraries and REST• Lift & Shift an Application <u>which already uses the native file system APIs to share data between it and other applications running in Azure.</u>• Store Development and Debugging Tools for use by many VMs
Azure Blobs	<ul style="list-style-type: none">• Client Libraries and REST• Want APP to support Streaming and Random Access Scenarios• Access Application Data from Anywhere
Azure Data Disks	<ul style="list-style-type: none">• Client Libraries and REST• Lift and Shift applications <u>that use native file system APIs to read and write data to persistent disks.</u>• Store data that is <u>not required</u> to be accessed from outside the virtual machine to which the disk is attached.

Manage Access

SAS – Shared Access Signatures

2 Types : [Ad Hoc SAS](#) & [SAS controlled by Stored Access Policy](#)

[Delegated access](#) to Storage Account Resources > [Blobs, Queues, Tables](#)

[URI format](#) with [permissions](#) and [specified time](#) | [signedidentifier](#) specifies [Stored Access Policy](#)

Client then passes the SAS to constructor or method

Stored Access Policies

Groups SASs + provide additional restrictions | up to 64 char

Greater control | [Best Practice to use with SAS](#)

[5 policies per Container](#), queue or table. Each policy-unlimited SASs

Regenerate Keys

WHY? Increase security

Affects virtual machines, media services, and any applications dependent on the storage account.

Must update all clients to use the new key.

[Share Access Signatures, Pt 1](#) [Stored Access Policies](#)

Configure diagnostics, monitoring and analytics

Configure Retention Policies

Days (1-365) | zero = set no policy

Logging Levels

Minimal e.g. ingress/egress, availability, latency, & success %'s
Aggregated for the **Blob**, **Table**, and **Queue** services.

Verbose – Same as above + collects same metrics

per each storage operation in Azure Storage Service API. Enables closer analysis of issues occurring during application operations.

Off - Turns off monitoring. Existing monitoring data persisted till end of retention period.

Analyze Logs

Logs saved in **blob container \$logs** in storage account.

Use **Blob svc API** to access [Read more on logs!](#)

[Monitor Storage Account](#)

Implement SQL databases

Im/Export Data

[Geo-Restore & Point in Time](#) preferred

Can use for Archiving

Can combine with Database Copy

Temp increase perf level to decrease export times

Export is in bulk | no guarantee on transactional consistency

Export = [BACPAC](#) files | requires [Storage Account](#) | Use [Export Data-tier Application Wizard](#)

Can Schedule [Automated Exports](#) & Also Can [Import/Export using REST API](#)

Im/Export Schema

A [DAC package](#) vs [BACPAC](#) target different scenarios.

A [BACPAC](#) contains [both schema and data](#), but does not support being imported to a database project for schema modification. [DAC packages](#) contain [only schema](#) information import into an [SSDT](#) database project for further development work. The primary use for a DAC package is in deploying a database schema to development, testing, and then production environments.

[Read More](#)

Implement Azure Recovery Services

Create Backup Vault

Backs up files/data from [Win Server](#) to [Azure](#)

Create a backup vault in [geographic region](#)

[Vault Credentials](#) Replace Certificates

Backup & Restore Data

"Protected Items" = been backed up

Recover 2012 or 2008 R2 SP1

Alternate Server Recovery

Know the [Passphrase](#) used to encrypt original

New server to recover [does not](#) need to be same [DNS/NetBIOS](#) name

[Start-OBRecovery](#) -RecoverableItem \$FinalItem -RecoveryOption \$secureString -Credential \$cstrial

[Configure Azure Back Up to back up Windows Server](#)
[Azure Backup Overview](#)

Implement recovery services

Deploy Backup Agent

REQUIRES: WIF and PS

WABInstaller.exe

Can install on:

- Servers: [2012 R2](#), [2012](#), [2008 R2 SP1](#)
- [64 bit Win 7, 8, 8.1](#)
- Ext available for Server 2012 Essentials

If using DPM, requires [Update Roll up 2](#) for SCDPM SP1

Recovery Services → Quick Start → to generate and download credential

Select Agent Type:

- [Azure Backup Agent](#)
- [Windows Server and System Center Data Protection Manager](#)
- [Windows Server Essentials](#)

[Install Backup Agent and upload vault credential](#)
[Administer Azure Backup with Windows PowerShell](#)

Implement Storage – *know these 5 things now*

Implement Blobs

Block Blobs (Sequential IO) up to 200GB each | **Page Blobs** (Random Access) up to 1 TB

Shared Access Storage

Delegated Access | Limit Permissions to Blobs, Queues, Tables | URI format w\perms & spec. time

Logging Levels

For Blobs, Tables and Queue Services | Off , Minimal, Verbose - > per Storage operations

SQL Import/Export | 2 File Types & Scenarios

BACPAC contains both schema and data | **DAC** packages contain only schema

Deploy Backup Agent | can install on....

Servers 2012 R2, 2012, 2008 R2 SP1, 64 bit Win 7, 8, 8.1, Ext available - Server 2012 Essentials

Design Data Storage

Data Storage Options

Table Storage	SQL Database	My SQL	DocumentDB	Blob Storage	MongoDB
NoSQL	Relational	Relational	NoSQL	Unstructured Data	NoSQL
Key/attribute	Tables		JSON/JavaScript	Big Data - logs	JSON-like
Schemaless			Web/Mobile!	Config Data- apps	Dynamic schemas
Fast Data Access		Fast DB			High Performance
Lower \$\$ versus>		Open Source			Written in C#

Data Security Options

Azure Storage | [Share Access Signature](#) - [SAS](#)

SQL Database – [same model](#) as on-premises

See aka.ms/azure/storage and for the real world... <http://storageexplorer.com/>

Security Options for Storage

Management Plane Security

RBAC – least privileges and Scope

Data Plane Security

Storage Account Keys & SAS

Encryption in transit / at rest

HTTPS & SAS / [Disk Encryption](#) &

Storage Analytics

Shows logging , Metrics, Authentic

Cross-Origin Resource Shari

Cross Domain, but not Authentica

Disabled by default

Enable w/ REST API or Client Library

The screenshot displays the Azure portal interface for a storage account named 'tr24'. The top section shows the 'Access keys' settings, including a search bar, a list of settings (Access keys, Configuration, Shared access signature, Properties, Locks, Automation s), and a list of services (BLOB SERVICE: Containers, Encryption, Metrics, Usage; FILE SERVICE: ...). The 'Allowed services' section includes checkboxes for Blob, File, Queue, and Table. The 'Allowed resource types' section includes checkboxes for Service, Container, and Object. The 'Allowed permissions' section includes checkboxes for Read, Write, Delete, List, Add, Create, Update, and Process. The 'Start and expiry date/time' section shows a start date of 2017-02-02 and a time of 9:14:23 PM. The bottom section shows the 'Encryption' settings, including a search bar, a list of settings (Encryption, Metrics, Usage), and a description of storage service encryption. A red circle highlights the text: 'Currently, this feature is available only for the Azure Blob service'. Below this, a note states: 'Note that after enabling Storage Service Encryption, only new data will be encrypted, and any existing blobs in this storage account will remain unencrypted.' At the bottom, there is a 'Storage service encryption' section with a 'Disabled' button and an 'Enabled' button.