

Microsoft Cloud Storage for Enterprise Architects

What IT architects need to know about storage for Microsoft cloud services and platforms

This topic is 1 of 5 in a series



Designing storage for the Microsoft cloud

Article version of this poster

Integrating your storage by using Microsoft cloud services gives you access to a broad range of services and cloud platform options.

Why cloud storage?

Speed to market

- Faster configuration for high availability and disaster recovery
- No storage hardware to purchase
- Built-in plumbing provided by Microsoft's cloud offerings
- Available from anywhere in the world

Lower costs to maintain

- Elasticity to scale up and down your storage demands
- No storage hardware to maintain or migrate
- Built-in plumber to maintain and improve infrastructure (Microsoft)
- Best storage security in the marketplace with ongoing improvements

Microsoft cloud storage options

Move-in ready

See topic 2 for more information and resources

Use these prepackaged solutions that are bundled with existing services. Use immediately and with minimal configuration.

Office 365	Microsoft Intune
OneDrive for Business	Dynamics 365
Visual Studio Team Services	Azure Site Recovery
Yammer file sharing	Azure Backup

Some assembly required

See topic 3 for more information and resources

Use these existing services as a starting point for your storage solution with additional configuration or coding for a custom fit.

Azure Content Delivery Network	Azure Media Services
HDInsight	Azure Redis Cache
Azure SQL Database	SQL Server on an Azure VM
Azure Cosmos DB	StorSimple
Azure SQL Data Warehouse	Azure Data Lake Store

Build from the ground up

See topic 4 for more information and resources

Use these storage building blocks, along with coding, to create your own storage solution or apps from scratch.

Azure Storage (files)
Azure Storage (blobs)
Azure Storage (queues)
Azure Storage (tables)

Key storage scenarios

Cache data

Accelerate access to commonly used data by storing it in a high-speed cache.

Collaborate with team members

Grant permission to multiple users to allow access to data in cloud storage.

Manage data

Store, move, or delete internal or external bulk data.

Manage source code

Upload, collaborate, and run application code files in the cloud.

Backup files

Store copies of internal or external data offsite in multiple cloud locations.

Publish company communications

Create a single point of publication for internal or external messages.

Distribute millions of events

Create storage for telemetry ingestion from websites, apps, and devices.

Manage/serve videos

Store and serve video content to customers or organization users.

Microsoft Cloud Storage for Enterprise Architects

What IT architects need to know about storage for Microsoft cloud services and platforms

This topic is 2 of 5 in a series











Move-in ready

✔ Use prepackaged solutions immediately

✔ Ready to go, bundled with existing services

✔ Limited customization

Product	Features	Common uses	Key storage scenarios
 SharePoint Online in Office 365	<ul style="list-style-type: none"> Files are secure, reliable, and accessible from multiple devices and from anywhere Content is easily discoverable through search Workflows and retention policies make it easy to use and maintain content Content can be accessed through web APIs or a browser, or offline 	<ul style="list-style-type: none"> Corporate or organization portal Team-based collaboration Enterprise search <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Collaborate Company communications Manage videos
 Exchange Online in Office 365	<ul style="list-style-type: none"> Mailboxes are secure, reliable, and accessible from multiple devices and from anywhere Mailbox content can be found through search 	<ul style="list-style-type: none"> Enterprise email storage <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Collaborate
 OneDrive for Business	<ul style="list-style-type: none"> Storage is secure, reliable, and accessible from multiple devices and from anywhere Content is easily discoverable through search by the user Content can be accessed through web APIs or a browser, or offline Each user gets 1 TB of storage 	<ul style="list-style-type: none"> Individual users' document store Ad hoc collaboration on documents <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Backup files
 Visual Studio Team Services	<ul style="list-style-type: none"> Has built-in redundancy Is based upon Team Foundation Server Create backlogs, work in sprints, and track changes and bugs The best option for storing source code (it's best to store documentation in SharePoint Online) 	<ul style="list-style-type: none"> Source code repository <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage source code
 Yammer file sharing	<ul style="list-style-type: none"> 5 GB file size limit for each uploaded file for premium networks No limit on the number of files that can be uploaded Best for quick sharing and discussion Accessible through a browser 	<ul style="list-style-type: none"> Enterprise social Share documents or links relevant to a conversation or topic <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Collaborate
 Dynamics 365	<ul style="list-style-type: none"> Subscription includes 10 GB of storage For each increment of 20 Professional USLs, 5 GB of storage capacity is added Storage is capped at 5 TB per customer 	<ul style="list-style-type: none"> Store and manage customer relationships <p>More Information</p>	<ul style="list-style-type: none"> Collaborate
 Microsoft Intune	<ul style="list-style-type: none"> All software that you deploy is packaged and uploaded to Microsoft Intune cloud storage Subscription includes 20 GB of storage with the option to purchase additional space, various software and hardware inventory reports, policy management, mobile app publishing, alerts, monitoring, and 99.99% uptime SLA 	<ul style="list-style-type: none"> Software deployment <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Backup files
 Azure Site Recovery	<ul style="list-style-type: none"> Simple, automated protection, replication, and recovery to Azure Continuous health monitoring and orchestrated recovery Supports managed disks and Ubuntu virtual machines Supports Windows Server 2016 and Storage Spaces 	<ul style="list-style-type: none"> Virtual machines running in an environment using Hyper-V can be replicated between two datacenters or to Azure InMage Scout enables the protection of VMware and physical environments <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Backup files
 Azure Backup	<ul style="list-style-type: none"> Seamlessly protects data received from Windows Server and Windows Server Essentials Includes System Center Data Protection Manager The data is stored in geo-replicated storage, which maintains six copies of your data across two Azure datacenters Supports delta changes, defined frequency, compression, encryption, longer retention, and bandwidth throttling 	<ul style="list-style-type: none"> Protect data from a variety of on-premises (SharePoint, SQL, Exchange) and Azure (VMs) sources <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Backup files

Microsoft Cloud Storage for Enterprise Architects

What IT architects need to know about storage for Microsoft cloud services and platforms

This topic is 3 of 5 in a series










Some assembly required




✔ Use existing services as a starting point for your storage solution

✔ Requires some configuration or coding

✔ Can be customized to fit your needs

Product	Features	Common uses	Key storage scenarios
 Azure Content Delivery Network	<ul style="list-style-type: none"> Advanced and real time analytics Robust security against DDoS Gets content automatically from an Azure Website or Azure Cloud Service once you set up the integration New partnership with Akamai Can handle sudden traffic spikes and heavy loads 	<ul style="list-style-type: none"> Distribute audio, video, applications, images, and other files faster and more reliably to customers by using the servers that are closest to them <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data Manage videos
 HDInsight	<ul style="list-style-type: none"> Apache Hadoop distribution powered by the cloud A Data Lake service Scale to petabytes on demand Process unstructured and semi-structured data Develop in Java, .NET, and more Skip buying and maintaining hardware Connect on-premises Hadoop clusters with the cloud Flexibility to deploy arbitrary Hadoop projects through custom scripts (e.g. R, Giraph, Solr) 	<ul style="list-style-type: none"> Data analytics workloads In-memory data processing framework for big data (Spark) Real-time stream processing (Storm) Large transactional processing (OLTP) of non-relational data (HBase) <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data
 Azure SQL Database	<ul style="list-style-type: none"> Optimized to reduce management and costs Automatic high availability, disaster recovery, and upgrade Recommended for organizations managing hundreds or thousands of databases of up to 1 TB in size Sharding techniques can split data across databases for increased storage Stretch database with SQL Server 2016 	<ul style="list-style-type: none"> New cloud-designed applications with relational data Data processing over schematic, highly structured data sets with relationships Spatial data or rich data types <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data <p><i>See topic 5 for information about Elastic Database and Stretch Database.</i></p>
 Azure Cosmos DB	<ul style="list-style-type: none"> Guaranteed low latency, 99.99% availability SLA with limitless, elastic scale of storage and throughput All data is globally replicated across any number of regions with transparent failover and four well-defined consistency levels Automatically indexes all your data without requiring schemas or secondary indices Rich SQL and JavaScript queries and multi-item transactions 	<ul style="list-style-type: none"> IoT, Mobile and Social Gaming Retail Content management <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data <p><i>See topic 5 for a Cosmos DB vs. Azure Tables vs. Azure SQL Database comparison.</i></p>
 Azure Media Services	<ul style="list-style-type: none"> Live and video on demand (VOD) delivery with scale Highly available encoding and streaming Supports Flash, iOS, Android, HTML5, and Xbox Studio-certified DRM support Rich content monetization Broad ecosystem of pre-integrated partners 	<ul style="list-style-type: none"> Encode, store, and stream audio and video at scale Real time streaming and VOD Streamlined video content management <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage videos
 Azure Redis Cache	<ul style="list-style-type: none"> Secure, dedicated Redis server with high-availability with data replication and failover managed by MS Recommended for any app needing high-throughput Available in sizes up to 530 GB and beyond (with Premium and automatic sharding) Redis Persistence persists in-memory cached data to Azure Storage Redis Clustering allows you to achieve maximum scale and throughput Enhanced security and network isolation with Azure Virtual Network support 	<ul style="list-style-type: none"> Reverse lookup for data in any storage service in Azure, such as Cosmos DB and Azure SQL Database Synchronized content from other data stores <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Cache data Message broker for high-throughput applications
 SQL Server on an Azure VM	<ul style="list-style-type: none"> SQL Server running as an installed application on an Azure virtual machine Use a gallery image with SQL Server installed or bring your own SQL Server license 	<ul style="list-style-type: none"> Manage data for applications <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data

Continued on next page

Product	Features	Common uses	Key storage scenarios
 StorSimple	<ul style="list-style-type: none"> Scalable, enterprise hybrid SAN storage with SSD and HDD in the on-premises hybrid storage array, with cloud storage as an integrated extension of the solution Inline deduplication, compression, automatic tiering, and encryption unstructured and semi structured data Automated offsite data protection using cloud snapshots Highly-efficient, location-independent disaster recovery Data mobility for enterprise data with StorSimple Virtual Appliance in Azure 	<ul style="list-style-type: none"> Manage data growth related to file shares, archives, and other data repositories Offsite data protection and disaster recovery for file shares, virtual machines, SQL, and SharePoint (using Remote Blob Storage) Utilize cloud snapshots to clone data in Azure and increase business agility <p style="text-align: right;"> More Information Cost Information </p>	<ul style="list-style-type: none"> Manage data Collaborate
 Azure SQL Data Warehouse	<ul style="list-style-type: none"> Elastic data warehouse that scales to petabytes Up to 32 concurrent queries Manage large volumes of structured data with fast analytics Dynamically grow and shrink compute in seconds Supports Transparent Data Encryption Backed up every 8 hours for 7 days 	<ul style="list-style-type: none"> Sales reports Usage reports Lots of data <p style="text-align: right;"> More Information Cost Information </p>	<ul style="list-style-type: none"> Manage data
 Azure Data Lake Store	<ul style="list-style-type: none"> A hyper-scale repository for big data analytics workloads A Hadoop Distributed File System for the cloud No fixed limits on file size No fixed limits on account size Unstructured and structured data in their native format Massive throughput to increase analytic performance High durability, availability, and reliability (99.9% enterprise-grade SLA and 24/7 support) Azure Active Directory access control 	<ul style="list-style-type: none"> Enterprise-wide repository to store every type of data collected in a single place <p style="text-align: right;"> More Information Cost Information </p>	<ul style="list-style-type: none"> Manage data

Note

Thanks to Oleg Kofman and Jon Epstein, both Senior Consultants for Microsoft, for their original material on which this model is based.

Microsoft Cloud Storage for Enterprise Architects

What IT architects need to know about storage for Microsoft cloud services and platforms

This topic is 4 of 5 in a series







Build from the ground up

✔ Create your own storage solution from scratch

✔ Requires programming using the REST APIs

✔ Ultimate in customization and flexibility

Product	Features	Common uses	Key storage scenarios
 <p>Azure Storage (files)</p>	<ul style="list-style-type: none"> Makes it easier to move legacy applications to the cloud Blob storage preferred for new applications Can mount from an Azure virtual machine Can mount on-premises with SMB 3.0 Works with Linux and Windows Doesn't support Azure AD-based authentication or ACLs (Azure Storage account keys provide authentication and authorized access to the file share) 	<ul style="list-style-type: none"> Migrating legacy applications to the cloud that rely on file shares Share development and testing tools Distributed apps can store logs, diagnostic data, and crash dumps <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Backup files
 <p>Azure Storage (blobs)</p>	<ul style="list-style-type: none"> Each storage account can hold up to 500 TB (one subscription can have multiple storage accounts) Storage accounts are organized into containers, which can have security applied to them and can contain blobs Block blobs are optimized for streaming and storing cloud objects, up to 200 GB in size Page blobs are optimized for representing IaaS disks and supporting random writes, up to 1 TB in size Append blobs are optimized for append operations, up to 195 GB Premium Storage provides faster IOPS through SSD storage 	<ul style="list-style-type: none"> Backups of files, computers, databases, and devices Images and text for web applications Configuration data for cloud applications Big data, such as logs and other large datasets Azure uses blob storage for its own services, such as HDInsight and virtual machine disks. <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data
 <p>Azure Storage (queues)</p>	<ul style="list-style-type: none"> Storage account can contain any number of queues Queue can contain any number of messages (until the storage account is full) Queue messages are automatically deleted after seven days if not retrieved and deleted by an application Messages may be up to 64 KB in size Secured at storage account level Queues are intended to pass control messages, not raw data 	<ul style="list-style-type: none"> Create a backlog of work to process asynchronously Processing log messages Decouple applications <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Distribute events
 <p>Azure Storage (tables)</p>	<ul style="list-style-type: none"> Best for semi-structured datasets Typically lower cost than traditional SQL Very fast if querying for key, slow if querying for value Massively scalable; any amount of tables up to the limits of the storage account Accessible through REST API, limited oData protocol, .NET Values must be serialized 	<ul style="list-style-type: none"> User data for web applications Address books Device information <p>More Information Cost Information</p>	<ul style="list-style-type: none"> Manage data

Microsoft Azure Storage recommendations

- Leverage multiple storage accounts for greater scalability, either for increased size (> 100 TB) or for more throughput (> 5,000 operations per second).
- Design the ability for adding additional storage accounts as a configuration change, not as a code change.
- Carefully select partitioning functions for table storage to enable the desired scale in terms of insert and query performance.
- Choose short column names for table properties as the metadata (property names) are stored in-band (the column names also count towards the maximum row size of 1 MB).
- When possible, batch operations into storage.
- Aggressively cache information in the configuration database into a distributed cache.
- If application performance or reliability is dependent on having a certain segment of data available in the cache, your application should refuse incoming requests until the cache has been prepopulated.
- Partition the data in either vertically (by table) or horizontally (segment table across multiple shards) to spread the load across multiple databases.

More information

Microsoft Azure Storage

<http://azure.microsoft.com/services/storage/>

Azure Storage Documentation

<https://docs.microsoft.com/azure/storage/>

Microsoft Cloud Storage for Enterprise Architects

What IT architects need to know about storage for Microsoft cloud services and platforms

This topic is 5 of 5 in a series



Appendices

Elastic database

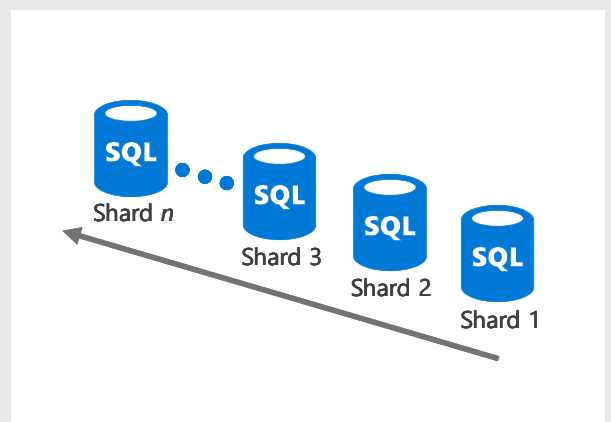
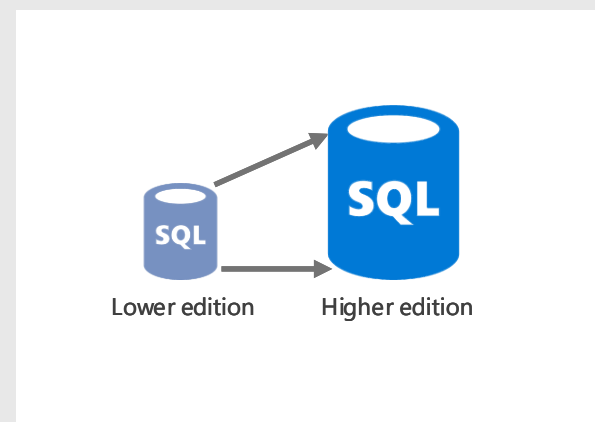
Use the virtually unlimited resources of Azure SQL Database when:

- The total amount of data is too large to fit within the constraints of a single database.
- The transaction throughput of the overall workload exceeds the capabilities of a single database.
- Tenants require physical isolation from each other, so separate databases are needed for each tenant.
- Different sections of a database need to reside in different geographies for compliance, performance, or geopolitical reasons.

[More Information](#)

Vertical scaling: You can change Azure database performance level/edition or by using elastic database pools.

Horizontal scaling: You can add new databases as needed.



Cosmos DB vs. Azure Tables vs. Azure SQL Database

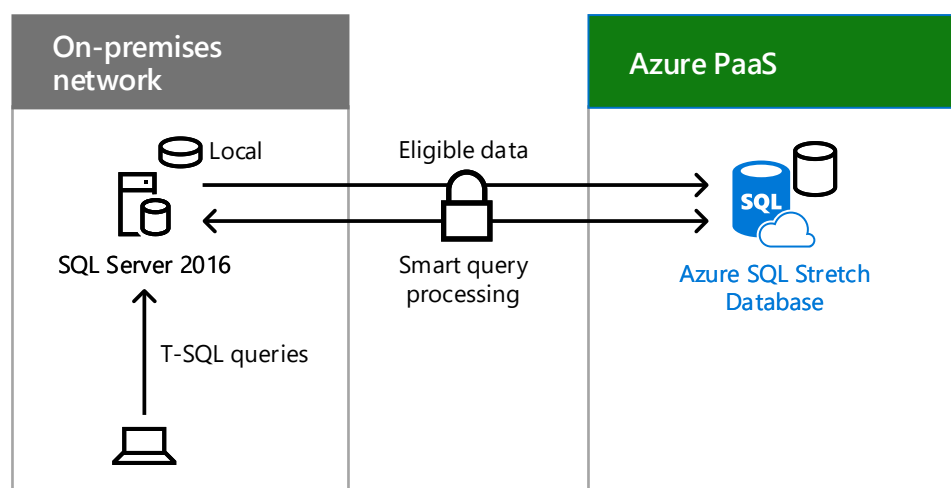
Common attributes of Cosmos DB, Azure Table Storage, and Azure SQL Database

- 99.99 availability SLA
- Fully managed database services
- ISO 27001, HIPAA and EU Model Clauses Compliant

Uncommon attributes

	Azure Cosmos DB	Azure Table Storage	Azure SQL Database
Multi-model (documents, key-value, ...)	Multi-model (documents, key-value, ...)	Key-value	Relational
SQL and JavaScript queries	SQL and JavaScript queries	No support for SQL or JavaScript queries	SQL queries
All data globally distributed across any number of regions	All data globally distributed across any number of regions	Single Region with DR	Single Region with DR, User Initiated failover support
Limitless storage and throughput	Limitless storage and throughput	Limitless storage	Max 1TB/database *Premium

Stretch Database with SQL Server 2016



Stretch database is a feature of SQL Server 2016 that allows you to transparently and securely move cold data, such as closed business data in a large table that contains customer order information, to a SQL Stretch database in Azure.

When stretched, the contents of a SQL Server instance, a database, or even a single table is the combination of local data in SQL Server 2016 server and remote data in Azure. Data that becomes eligible for stretch is automatically moved to Azure by SQL Server 2016.

User queries that include the historical data are transparently forwarded to Azure SQL Stretch database. The queries do not need to be re-written, even though the table is stretched.

Stretch database provides a cost-effective option for long-term storage and transparent access to historical data. It also solves performance and availability problems that arise when tables become very large.

[More Information](#)

More Microsoft cloud IT resources

<p>Services and Platform Options aka.ms/cloudarchoptions</p>	<p>Security aka.ms/cloudarchsecurity</p>	<p>Identity aka.ms/cloudarchidentity</p>	<p>Networking aka.ms/cloudarchnetworking</p>
<p>Mobility aka.ms/cloudarchmobility</p>	<p>Hybrid aka.ms/cloudarchhybrid</p>	<p>Contoso in the Microsoft Cloud aka.ms/cloudarchcontoso</p>	