



A Quickstart Guide to Azure Data Storage: From Data Lakes to Blobs, Files and Tables

Table of Contents

Effective disaster recovery is a cornerstone of modern business resilience. This Azure SQL Server Disaster Recovery Plan outlines essential strategies and protocols to safeguard your operations against disruptions, ensuring continuity, data integrity, and adaptability in the face of challenges.

01 Introduction.....	4
02 Azure Data Storage Primer.....	5
03 The Benefits of Azure Storage.....	7
04 Roadmap to Organizational Maturity.....	10
05 How to Prepare for Azure Storage Adoption.....	12



Chapter 1: Introduction

Data makes the modern world go round, and companies understand they need to be in the business of collecting data if they want to stay competitive. It's why seemingly every company these days has an app and/or loyalty program, all designed to collect information so they can make data-informed decisions on how to best run their business.

A survey commissioned by Dell Technologies and conducted by Forrester found that 70% of data decision-makers are gathering data faster than they can analyze and use it, yet 67% say they constantly need more data than their current capabilities provide.

So where do companies store all this data? Increasingly, they're relying less on on-premises storage and turning to cloud solutions. In 2015, just 30% of corporate data was stored in the cloud, but by 2022 this figure had surged to over 60%.

We expect more companies will rely on cloud storage as they embark on or continue their data modernization journey, and increasingly they'll turn to one of the leaders in the space—Azure Data Storage.



In 2015, just 30% of corporate data was stored in the cloud, but by 2022 this figure had surged to over 60%.

Chapter 2: Azure Data Storage Primer

Azure Data Storage is a suite of cloud storage solutions offered by Microsoft within the Azure Cloud platform. It provides highly available, scalable and secure storage for various data needs. Here's a breakdown of its features, benefits and actual use cases.

Azure Data Storage Primer Features

Durability and Availability

Azure Storage ensures your data is reliably stored and accessible, reducing the risk of data loss and costly downtime.

Global Reach

With data centers worldwide, Azure Storage enables you to deliver content and services with low latency to users globally, enhancing user experience and supporting international operations.

Scalability

Azure Storage allows you to adjust resources according to your current needs, ensuring that your applications can handle growth and variable loads efficiently without any degradation in performance.

Security

Azure Storage protects your data against threats and unauthorized access, helping you meet compliance requirements and safeguarding your business reputation.

Competitive Pricing Models

Azure Storage allows your business to optimize storage costs based on your specific usage patterns and needs with pay-as-you-go and reserved capacity options.

Manageability

Azure handles the underlying infrastructure, freeing you to focus on your core business.

Unparalleled reliability Azure Storage is the sole cloud provider among the major CSPs that maintains a 0% annual failure rate since launch.

Chapter 2: Azure Data Storage Primer

Services Within Azure Storage

With Azure Storage, you get access to an array of services designed to meet diverse storage needs.

Ideal for

Azure Data Lake Storage (ADLS), Azure Blob Storage, Azure Files, Azure Tables and Azure Queues.

Azure Data Lake Storage

ADLS is an enterprise-wide hyper-scale repository for big data analytics workloads. ADLS integrates with Azure Blob Storage but is optimized for processing large data lakes that require massive parallel processing. It supports advanced capabilities like fine-tuned access control and file system semantics, making it suitable for Big Data analytics and operational data scenarios.

Ideal for

Big data analytics workloads involving massive volumes of unstructured data in its native format (e.g., CSV, JSON, parquet).

Azure Blob Storage

Blobs are best suited for large datasets or media files accessed infrequently. Data is organized in containers (virtual folders) for better manageability. Blob storage supports three types of blobs: block blobs, append blobs and page blobs, catering to different use cases like storing documents or supporting large-scale workloads like virtual machines.

Ideal for

Large unstructured data like media files (videos, images, audio), backups, archives, log files and scientific datasets.

Azure Files

Azure Files offers fully managed file shares in the cloud, accessible via the industry-standard Server Message Block (SMB) and Network File System (NFS) protocols. It is a simple way to migrate legacy applications that rely on file shares to the cloud without major changes. Azure Files supports both hot and cool storage tiers to optimize costs based on access patterns.

Ideal for

Sharing files across platforms (Windows, Linux, macOS) in the cloud or on-premises environments.

Azure Tables

Azure Table Storage is a NoSQL data store for semi-structured data. It's a great option for storing datasets that don't require complex joins, foreign keys or stored procedures, and where schemas need to evolve over time. It's designed for fast access to large volumes of data, making it great for web-scale applications.

Ideal for

Storing structured or semi-structured data.

Azure Queues

Azure Queues provides a cloud-based messaging service within the Azure Storage platform, offering a reliable way to asynchronously communicate messages between application components, even if they're not online at the same time.

Ideal for

Creating a backlog of work to process asynchronously.

Chapter 3: The Benefits of Azure Storage

Data makes the modern world go round, and companies understand they need to be in the business of collecting data if they want to stay competitive. It's why seemingly every company these days has an app and/or loyalty program, all designed to collect information so they can make data-informed decisions on how to best run their business.

Business Benefits

Scalability and Flexibility

Azure Storage offers almost unlimited scalability, allowing businesses to expand their storage capabilities without the need for significant capital expenditure on additional hardware. Businesses in expansion mode can rely on this adaptability and feel safe in the knowledge that their storage solution will scale with them.

Cost-Effectiveness

To the surprise of many customers who assume on-premises storage is always cheaper, Azure Storage is very affordable. In fact, moving a terabyte of data into Azure Storage can cost as little as \$200 a month. This affordability makes the offering an attractive option for companies looking to manage their data storage costs effectively.

Reduced Capital Expenditure

With Azure Storage, there is no need for businesses to make significant capital outlays for additional storage hardware. This not only saves on the direct costs of purchasing hardware but also on the indirect costs associated with hardware maintenance and upgrades.

Governance and Security

Azure provides robust governance tools, allowing businesses to set up guardrails and role-based access controls. This ensures data is accessed only by authorized personnel and in accordance with business policies, thereby enhancing security and compliance.



Chapter 3:

The Benefits of Azure Storage

Technical Benefits

Performance

Azure Storage boasts superior performance, with lightning-fast data access and processing. This translates to smooth operations and a delightful user experience for your customers and employees alike. Businesses that thrive on speed, like those providing customer-facing services or relying on real-time analytics, will find Azure Storage a great fit.

Configurability

The ability to tailor Azure Storage solutions to meet specific business requirements means you can optimize your storage setup for performance, cost or security according to your priorities. This flexibility supports diverse use cases, from archiving to active data processing, ensuring that your storage environment aligns with your business strategy.

Reliability

The proven reliability of Azure Storage ensures that your critical data is always available when needed. Its built-in redundancy features automatically replicate your data across geographically distributed servers. This means that even in the event of a hardware failure, your data remains available, preventing downtime and minimizing disruption.

Data Recovery

The data recovery features of Azure Storage act as a safety net for your business, protecting against data loss from accidental deletions or other errors. This capability makes your business resilient, safeguarding your valuable data assets and maintaining regulatory compliance.

Access Flexibility

Azure's multiple data access methods provide your business the flexibility to interact with data in the most efficient way possible, whether through applications, automated workflows or direct access. This versatility supports a wide range of use cases and operational needs, enhancing productivity and innovation.

Lifecycle Management

Lifecycle management automates data tiering based on access frequency, seamlessly moving data to cost-effective tiers and even deleting expired data. This intelligent approach promotes optimal data availability while keeping costs in check, making Azure Storage a powerful and reliable solution for businesses of all sizes.

Granular Security Controls

Detailed security controls allow your business to protect sensitive information effectively, ensuring that only authorized users can access specific data. This granular access control maintains data confidentiality, integrity and compliance with regulatory standards.

Integration with Azure Ecosystem

The seamless integration of Azure Storage with other Azure services enables your business to build and deploy sophisticated solutions that combine storage with advanced analytics, machine learning and data visualization. This ecosystem synergy drives innovation, enhances operational insights and creates competitive advantages.

Support for Popular Programming Languages

Compatibility with widely used programming languages like Python and Scala simplifies the integration of Azure Storage into your existing development projects and workflows. This ease of integration accelerates development cycles, reduces learning curves for your technical teams and enhances productivity.

WHY MOVE TO THE CLOUD?

In a survey conducted by Foundry regarding the transition to cloud computing, IT leaders provided these reasons for making the shift.

REPLACING ON-PREMISES LEGACY TECHNOLOGY

40%

ENABLING DISASTER RECOVERY AND BUSINESS CONTINUITY

40%

LOWER TOTAL COST OF OWNERSHIP

34%

IMPROVING EMPLOYEE PRODUCTIVITY

32%

GREATER FLEXIBILITY TO REACT TO CHANGING MARKET CONDITIONS

32%

Azure Storage Use Cases

Azure Storage caters to a broad spectrum of industry needs. Here are some real-world applications, showcasing how businesses leverage Azure Storage to overcome challenges and achieve operational excellence.

Mining Industry: Drone Footage Storage

A mining company faced significant challenges with its on-premise storage infrastructure, which crashed under the weight of 12 terabytes of drone footage. The

incident highlighted issues related to capital expenditure and the inherent limitations of physical servers. Transitioning to Azure Files offered a robust solution, eliminating server costs and ensuring unparalleled reliability. The move not only provided cost savings but also enhanced data accessibility, allowing the company to analyze drone footage from any location with internet access, a critical factor for operational agility in the mining sector.

Manufacturing Sector: File Share Modernization

A manufacturing company sought to divest from the data center business, prioritizing the migration of its file shares to a more reliable and cost-effective platform. Azure Storage, with its SMB file shares, emerged as an ideal solution. This transition enabled employees to map their drives to Azure Storage, moving away from traditional, hardware-dependent setups. The shift not only reduced the company's reliance on physical infrastructure but also eliminated the need for managing backups, as these were seamlessly integrated into the Azure service, resulting in significant cost savings and operational efficiency.

Software Development: Event Log Management

A manufacturing firm utilized a high-cost Azure SQL instance for storing event logs, leading to a staggering monthly expense of \$10,000. Adopting Azure Table Storage for NoSQL storage of event logs marked a transformative shift, drastically reducing the cost to a mere \$100 a month. This strategic move not only exemplified cost efficiency but also underscored Azure Storage's capability to provide scalable, high-performance solutions for data management challenges commonly faced in software development.

Analytics Optimization: Data Lake and Synapse Integration

A business grappling with performance issues due to continuous analytics queries on its production database found relief in ADLS and Synapse Analytics. By batch processing data transfers to a data lake and conducting analytics in Synapse, the company alleviated stress on its primary database, enhancing overall system performance. This integration not only optimized analytics workflows but also ensured that core business operations remained uninterrupted, demonstrating the power of Azure Storage in facilitating efficient data management and analytics strategies.

Chapter 4:

Roadmap to Organizational Maturity

Initiating a data modernization project requires a thorough understanding of the entire journey so you can take the actions required to transition to a more sophisticated platform. Below is a table outlining a framework for adopting Azure's data modernization services.

Early Adoption

Azure Storage: Organizations typically start their Azure journey by migrating data to Azure Blob Storage or Azure Data Lake Storage for scalable, secure cloud storage.

Azure SQL Database: Often one of the first steps in cloud adoption. A typical entry point for cloud migration. It transitions on-prem relational databases to a managed cloud service with familiar SQL capabilities, simplifying cloud migration.

Azure Data Factory: Used for initial data integration and orchestration efforts. It helps move and refine data from multiple sources for Azure-based processing and analytics.

Azure Data Pipelines: Refer to the data processing workflows in Azure Data Factory (ADF). These pipelines orchestrate and automate the movement and transformation of data from various sources to different targets.

Mid Adoption

Azure Synapse Analytics: As orgs look to scale their analytics capabilities, Azure Synapse Analytics enables complex queries and extensive data processing for scalable analytics.

Azure Databricks: Ideal for businesses expanding their use of analytics and AI. Azure Databricks provides a unified workspace optimal for enhancing analytics and AI initiatives, using Apache Spark for rapid data science and machine learning development.

Advanced Adoption

Azure Cosmos DB: Fits into the more advanced stages of adoption. Its ability to scale massively, offer low latency and support various data models makes it suitable for complex, high-performance applications.

Power BI: Power BI can be introduced at various stages. It is used to visualize and share insights derived from complex analytics processes supported by other Azure services. It tends to align more with advanced adoption, enabling deep insights and data-driven decision-making across the organization.

Recommended for the Early Adoption Stage

Azure Storage is recommended for businesses at the early stages of their data maturity journey. At this foundational level, your organization is just beginning to recognize the limitations of traditional, on-premises data storage solutions, such as scalability issues, high maintenance costs and inadequate disaster recovery capabilities.

Adopting Azure Storage at this stage sets a solid foundation for your data storage needs, and enables easy expansion into more advanced Azure services, such as Azure Synapse Analytics and Power BI.

The Role of a Data Lake in Data Modernization

Integrating a data lake into your tech stack represents a shift towards a more flexible, scalable and analytics-driven approach to data management. As organizations amass vast quantities of structured and unstructured data from diverse sources, the need for a centralized, resilient and accessible repository becomes imperative. ADLS meets this need, serving as a key element within a data modernization framework.

Centralized Data Repository

ADLS provides a centralized platform to store vast amounts of data in its native format. This consolidation facilitates efficient data management, eliminating data silos and enabling comprehensive data analysis and intelligence. By serving as a single source of truth, ADLS enhances data integrity and accessibility across the organization.

Scalability and Performance

ADLS is designed for enterprise-level scalability, accommodating exponential data growth without compromising performance.

Its ability to handle massive workloads and perform high-speed analytics makes it essential to modern data architectures as it supports a range of analytics, from batch processing to real-time analytics.

Advanced Analytics and AI Integration

A data lake's structure allows for the seamless integration of advanced analytics and AI capabilities. ADLS, particularly when combined with Azure Databricks or Azure Synapse Analytics, empowers organizations to unlock deep insights, predictive analytics and machine learning opportunities. This integration enables your business to leverage its data for strategic decision-making and innovation.

Data Democratization

By facilitating easy access to data and integrating with popular analytics and BI tools, ADLS plays a vital role in democratizing data within your organization. This accessibility encourages a culture of data-driven decision-making, where insights derived from the data lake can inform strategies across all levels of the business.

Chapter 5:

How to Prepare for Azure Storage Adoption

1) Access Requirements

Defining access requirements is key to a successful Azure Storage implementation. Consider taking the following actions:

Identify Stakeholders

Determine who needs access to the Azure Storage and for what purposes. This includes IT personnel, data analysts and business users.

Access Controls

Plan how to implement role-based access control so users have the right level of access to data stored in Azure Storage. Azure provides fine-grained access control mechanisms to secure data.

Network Configuration

Assess your current network infrastructure to ensure it can support connectivity to Azure services. Consider options such as Azure ExpressRoute for a more reliable and secure connection to Azure Storage.

Performance and Scalability

Evaluate your performance requirements and how they align with Azure Storage options. Azure offers different storage tiers to accommodate various performance and access frequency needs.

Initiating a data modernization project requires a thorough understanding of the entire journey so you can take the actions required to transition to a more sophisticated platform. Below is a table outlining a framework for adopting Azure's data modernization services.

2) Data and Schema

The effectiveness of Azure Storage is shaped by how your data is organized and structured.

Data Assessment

Conduct a thorough audit of your current data assets. Understand the types, volumes and sources of data that will be moved to Azure Storage.

Schema Design

Plan how to organize your data in Azure Storage. This includes deciding between different storage options like Blob, File and Table storage based on your data types and usage patterns.

Data Migration Strategy

Develop a comprehensive migration plan to minimize downtime and maintain data integrity during the migration. Consider using Azure Data Factory or Azure Data Box for large datasets.

3) Security and Privacy

Keeping data secure and private is non-negotiable.

Encryption

Familiarize yourself with Azure's encryption options for data at rest and in transit. Azure Storage automatically encrypts your data before storing it and decrypts it for retrieval.

Compliance

Review Azure's compliance offerings to ensure they meet your industry's regulatory requirements. Azure has a comprehensive set of compliance certifications that can help meet your legal and regulatory obligations.

Data Privacy

Understand the tools and policies Azure provides to manage data privacy, such as Azure Policy and Azure Blueprints, which can help enforce your organization's privacy standards.

4) Sovereignty and Residency

Data sovereignty and residency are important considerations, especially for businesses operating in multiple countries or regions.

Data Residency

Choose the Azure region(s) where your data will be stored, considering factors like latency, availability and specific data residency requirements your organization might have due to local regulations.

Sovereignty Considerations

Be aware of the legal implications of storing data in different regions. Azure's global infrastructure offers options to ensure data is stored and managed according to regional laws and regulations.

Geo-Replication

For high availability and disaster recovery, consider Azure's geo-replication options. Understand how geo-replication can impact data sovereignty and ensure compliance with cross-border data transfer policies.





Digital Infrastructure Solutions Built for Your Business



US Signal, established in 2001, is a premier national digital infrastructure company that operates a fully owned fiber network to deliver a wide range of advanced digital solutions. Our offerings include robust cloud services, secure colocation facilities, high-performance connectivity, comprehensive hardware resale, and managed IT services, empowering businesses to enhance their operational efficiency through tailored network, data center, data protection, and cybersecurity solutions.