

# Move mainframe compute and storage to Azure

By Larry Mead  
AzureCAT

March 2019

# Contents

Introduction .....	3
Mainframe compute and storage at a glance .....	4
Mainframe compute .....	4
Mainframe storage .....	4
Mainframe versus IO latency and IOPS .....	5
Azure compute and storage at a glance .....	5
Azure compute .....	5
Azure storage .....	8
Optimizing Azure compute and storage .....	9
Azure Government .....	10
Summary .....	10
Learn more .....	11
Azure compute and storage .....	11
IBM resources .....	11
Azure Government .....	11
Mainframe migration .....	11

Authored by Larry Mead. Edited by Nanette Ray. Reviewed by AzureCAT.

© 2019 Microsoft Corporation. This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

# Introduction

Mainframes have a reputation for high reliability and availability and continue to be the trusted backbone of many enterprises. They are often thought to have nearly limitless scalability and computing power as well. However, some enterprises have outgrown the capability of the largest available mainframes and need to know how to make the most of the cloud’s agility, reach, and infrastructure savings.

To run mainframe workloads on Microsoft Azure, you need to know how your mainframe’s capabilities compare to Azure. This guide focuses on compute capacity and storage. Based on an IBM z14 mainframe (the most current model as of this writing), this guide tells you how to get comparable results on Azure.

To get started, it helps to see both environments side by side. Figure 1 compares a mainframe environment for running applications to an Azure hosting environment. This diagram serves as a useful point of reference throughout this guide.

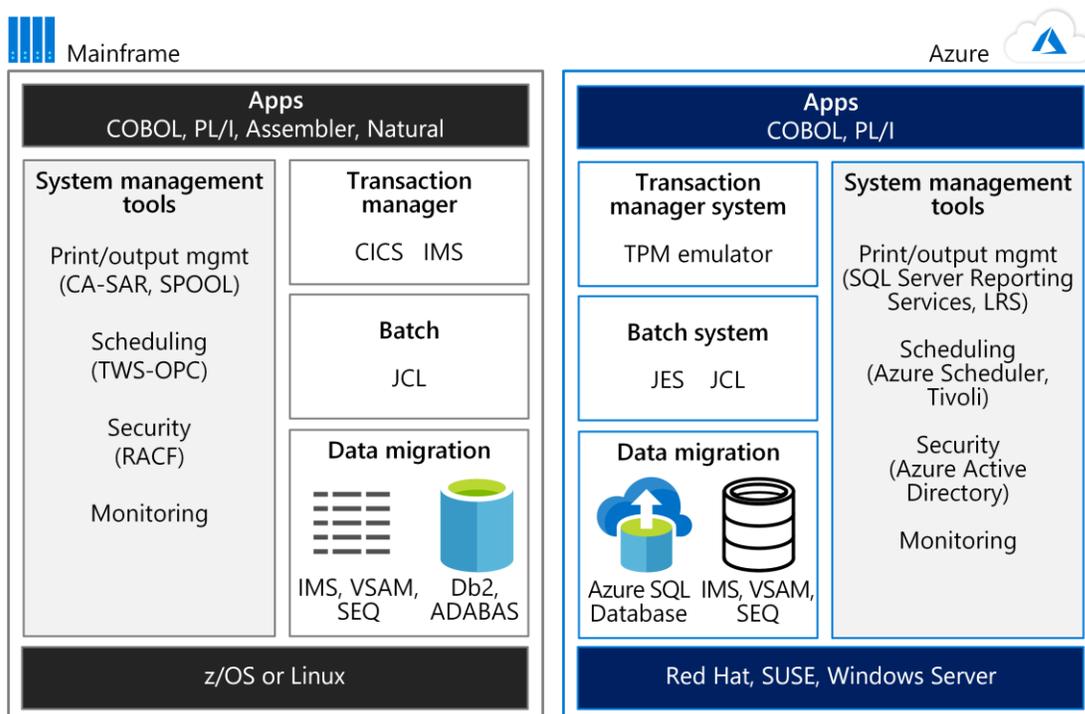


Figure 1. Azure services and emulation environments offer comparable support and streamlines migration

The power of mainframes is often used for online transaction processing (OLTP) systems that handle millions of updates for thousands of users. These applications often use software for transaction processing, screen handling, and form entry, such as a Customer Information Control System (CICS), Information Management System (IMS), and Transaction Interface Package (TIP). As Figure 1 shows, a TPM emulator on Azure can handle CICS and IMS workloads. A batch system emulator on Azure performs the role of Job Control Language (JCL). Mainframe data is migrated to Azure databases, such as Azure SQL Database. Azure services or other software hosted in Azure Virtual Machines can be used for system management.

# Mainframe compute and storage at a glance

For those unfamiliar with mainframes, this section describes the basic units of compute power and storage.

## Mainframe compute

In the z14 mainframe, processors are arranged in up to four *drawers*. A *drawer* is simply a cluster of processors and chipsets. Each drawer can have six active central processor (CP) chips, and each CP has 10 system controller (SC) chips. In Intel x86 terminology, there are six sockets per drawer, 10 cores per socket, and four drawers. This architecture provides the rough equivalent of 24 sockets and 240 cores, maximum, for a z14.

The z14 CP is extremely fast with a 5.2 GHz clock speed.

A z14 is usually delivered with all the CPs in the box. They are activated as needed. A customer is typically charged for at least four hours of compute time per month despite actual usage.

A mainframe processor can be configured as a General Purpose (GP) processor, System z Integrated Information Processor (zIIP), Integrated Facility for Linux (IFL) processor, System Assist Processor (SAP), and Integrated Coupling Facility (ICF) processor.

## Scaling mainframe compute up and out

IBM mainframes offer the ability to scale up to 240 cores (the current z14 size for a single system). Additionally, IBM mainframes can scale out through a feature called the Coupling Facility (CF). The CF allows multiple mainframe systems to access the same data simultaneously. Using the CF, the mainframe Parallel Sysplex technology groups mainframe processors in clusters. When this guide was written, the Parallel Sysplex feature supported 32 groupings of 64 processors each. Up to 2,048 processors can be grouped in this manner to scale out compute capacity.

A CF allows the compute clusters to share data with direct access, and it is used for locking information, cache information, and the list of shared data resources. A Parallel Sysplex using one or more CFs can be thought of as a “shared everything” scale-out compute cluster. For more information about these features, see [Parallel Sysplex on IBM Z](#) on the IBM website.

Applications can use these features to provide both scale-out performance and high availability. For information about how CICS can use Parallel Sysplex with CF, download the [IBM CICS and the Coupling Facility: Beyond the Basics](#) redbook.

## Mainframe storage

The IBM mainframe characterizes storage in two ways. The first is a direct access storage device (DASD). The second is sequential storage. For managing storage, the mainframe provides the Data Facility Storage Management Subsystem (DFSMS). It manages data access to the various storage devices.

According to Wikipedia<sup>1</sup>, DASD refers to a separate device for secondary (not in-memory) storage that allows a unique address to be used for direct access of data. Originally, the term DASD applied to spinning disks, magnetic drums, or data cells. However, now the term can also apply to solid-state storage devices (SSDs), storage area networks (SANs), network attached storage (NAS), and optical drives. For the purposes of this document, DASD refers to spinning disks, SANs, and SSDs.

In contrast to DASD storage, sequential storage on a mainframe refers to devices like tape drives where data is accessed from a starting point, then read or written in a line.

Storage devices are typically attached using a fiber connection (FICON) or are accessed directly on the mainframe's IO bus using *HiperSockets*, an IBM technology for high-speed communications between partitions on a server with a hypervisor.

Most mainframe systems separate storage into two types:

- *Online storage* (also known as hot storage) is needed for daily operations. DASD storage is usually used for this purpose. However, sequential storage, such as daily tape backups (logical or physical), can also be used for this purpose.
- *Archive storage* (also known as cold storage) is not guaranteed to be mounted at a given time. Instead, it is mounted and accessed as needed. Archive storage is often implemented using sequential tape backups (logical or physical) for storage.

## Mainframe versus IO latency and IOPS

Mainframes are often used for applications that require high performance IO and low IO latency. They can do this using the FICON connections to IO devices and HiperSockets. When HiperSockets are used to connect applications and devices directly to a mainframe's IO channel, latency in the microseconds can be achieved.

# Azure compute and storage at a glance

Although some people think that Intel-based servers are not as powerful as mainframes, the new core-dense, Intel-based systems have as much compute capacity as mainframes. This section describes the Azure infrastructure-as-a-service (IaaS) options for computing and storage. Azure provides platform-as-a-service (PaaS) options as well, but the focus here are the IaaS choices that provide comparable mainframe capacity.

## Azure compute

Azure Virtual Machines provide compute power in a range of sizes and types. In Azure, a virtual CPU (vCPU) roughly equates to a core on a mainframe.

Currently, the range of Azure Virtual Machine sizes provides from 1 to 128 vCPUs. Virtual machine (VM) types are optimized for particular workloads. For example, the following list shows the VM

---

<sup>1</sup> Direct-access storage device topic, [https://en.wikipedia.org/wiki/Direct-access\\_storage\\_device](https://en.wikipedia.org/wiki/Direct-access_storage_device).

types (current as of this writing) and their recommended uses:

Size	Type and description
D-Series	General purpose with 64 vCPU and up to 3.5 GHz clock speed
E-Series	Memory optimized with up to 64 vCPUs
F-Series	Compute optimized with up to 64 vCPUs and 3.7 GHz clock speed
H-Series	Optimized for high performance computing (HPC) applications
L-Series	Storage optimized for high-throughput applications backed by databases such as NoSQL
M Series	Largest compute and memory optimized VMs with up to 128 vCPUs

For details about available VMs, see [Virtual Machine series](#).

As mentioned earlier, a z14 mainframe can have up to 240 cores. However, z14 mainframes almost never use all the cores for a single application or workload. Instead, a mainframe segregates workloads into logical partitions (LPARs), and the LPARs have ratings—MIPS (Millions of Instructions Per Second) or MSU (Million Service Unit). When determining the comparable VM size needed to run a mainframe workload on Azure, factor in the MIPS (or MSU) rating.

The following are general estimates:

- 150 MIPS per vCPU
- 1,000 MIPS per processor

To determine the correct VM size for a given workload in a LPAR, first optimize the VM for the workload. Then determine the number of vCPUs needed. A conservative estimate is 150 MIPS per vCPU. Based on this estimate, for example, an F-series VM with 16 vCPUs could easily support an IBM Db2 workload coming from an LPAR with 2,400 MIPS.

### Azure compute scale-up

As this document was written, the M-series VMs could scale up to 128 vCPUs. Using the conservative estimate of 150 MIPS per vCPU, the M-series VM equates to about 19,000 MIPS. The general rule of thumb for estimating MIPS for a mainframe is 1,000 MIPS per processor. A z14 mainframe can have up to 24 processors. This provides about 24,000 MIPS for a single mainframe system.

While the largest single z14 mainframe has approximately 5,000 MIPS more than the largest VM available in Azure, it is important to compare how workloads are deployed on a mainframe and on an Azure Virtual Machine. If a system has both an application and a relational database, both are typically deployed on the same physical mainframe, using separate LPARs for the application and database. On Azure, the same solution is typically deployed using one VM for the application and a separate, suitably sized VM for the database. For example, if a M64 vCPU system supports the application, and a M96 vCPU is used for the database, approximately 150 vCPUs are needed—or about 24,000 MIPS as the following figure shows.

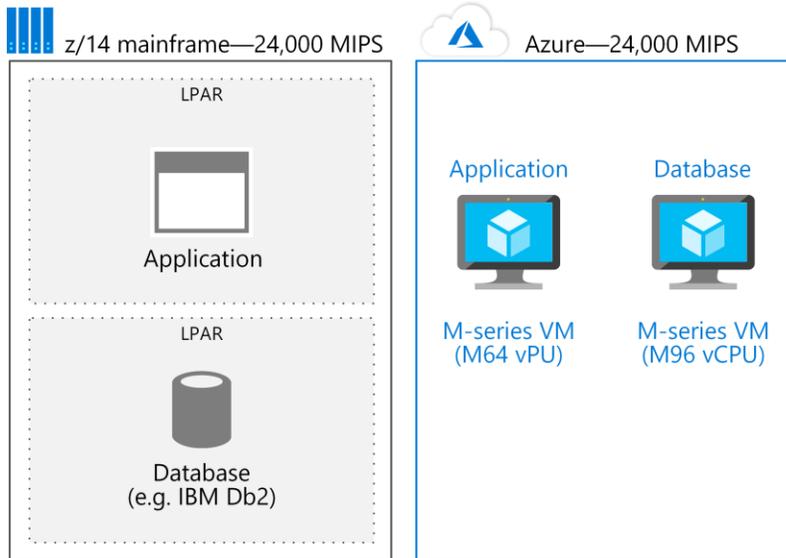


Figure 2. Comparing workload deployments

When you take the approach of migrating LPARs to individual VMs, Azure can easily scale up to the size needed for most applications that are deployed on a single mainframe system.

### Azure compute scale-out

One of the big advantages of an Azure-based solution is the ability to scale out. Scaling makes nearly limitless compute capacity available to an application. Azure supports multiple methods to scale-out compute power:

- **Scale-out with multiple VMs in a cluster.** In this scenario, an application can use a load balancer or resource manager to spread the workload out among multiple VMs in a cluster. If more compute capacity is needed, additional VMs are added to the cluster.
- **VM scale sets.** In this burst scenario, an application can scale to additional compute resources based on VM utilization. When demand decreases, the number of VMs in a scale set can also decrease, ensuring efficient use of compute power.
- **PaaS scaling.** In addition to VMs, Azure PaaS offerings scale compute resources as needed. For example, Azure Service Fabric, a platform for building and running microservices and containers, allocates additional compute resources to support increases in the volume of requests.
- **Kubernetes clusters.** Applications on Azure can use Kubernetes clusters for compute services for specified resources. Azure Kubernetes Service (AKS) is a managed service that orchestrates Kubernetes nodes, pools, and clusters on Azure.

In choosing the right method for scaling out compute resources, it's important to understand the difference in how Azure and mainframes scale out. The key is how—or if—data is shared by compute resources. In Azure, data (by default) is not typically shared by multiple VMs. If data sharing is required by multiple VMs in a scale-out compute cluster, the shared data must reside in a resource that supports this functionality. On Azure, data sharing involves storage as the following section discusses.

For more information, see the following Azure documentation articles:

- [What are virtual machine scale sets?](#)
- [What is a load balancer?](#)
- [Kubernetes core concepts for AKS](#)
- [Overview of Azure Service Fabric](#)

## Azure storage

Microsoft offers petabytes worth of storage for applications hosted in Azure, and you have several storage options. These range from SSD storage for high performance to low-cost blob storage for mass storage and archives. Additionally, Azure provides a data redundancy option for storage—something that takes more effort to set up in a mainframe environment.

Azure storage is available as Azure Disks, Azure Files, and Azure Blobs as the following table summarizes. For more information about Azure Storage, see [Deciding when to use Azure Blobs, Azure Files, or Azure Disks](#) in the Azure documentation.

Type	Description	Use when you want to:
<b>Azure Files</b>	Provides an SMB interface, client libraries, and a <a href="#">REST interface</a> that allows access from anywhere to stored files.	<ul style="list-style-type: none"> <li>• Lift and shift an application to the cloud when the application uses the native file system APIs to share data between it and other applications running in Azure.</li> <li>• Store development and debugging tools that need to be accessed from many virtual machines.</li> </ul>
<b>Azure Blobs</b>	Provides client libraries and a <a href="#">REST interface</a> that allows unstructured data to be stored and accessed at a massive scale in block blobs.  Also supports <a href="#">Azure Data Lake Storage Gen2</a> for enterprise big data analytics solutions.	<ul style="list-style-type: none"> <li>• Support streaming and random-access scenarios in an application.</li> <li>• Have access to application data from anywhere.</li> <li>• Build an enterprise data lake on Azure and perform big data analytics.</li> </ul>
<b>Azure Disks</b>	Provides client libraries and a <a href="#">REST interface</a> that allows data to be persistently stored and accessed from an attached virtual hard disk.	<ul style="list-style-type: none"> <li>• Lift and shift applications that use native file system APIs to read and write data to persistent disks.</li> <li>• Store data that is not required to be accessed from outside the virtual machine to which the disk is attached.</li> </ul>

## Azure hot (online) and cold (archive) storage

The type of storage for a given system depends on the requirements of the system, including

storage size, throughput, and IOPS. For DASD-type storage on a mainframe, applications on Azure typically use Azure Disks drive storage instead. For mainframe archive storage, blob storage is used on Azure.

SSDs provide the highest storage performance on Azure. The following options are available (as of the writing of this document):

Type	Size	IOPS
Ultra SSD	4 GB to 64 TB	1,200 to 160,000 IOPS
Premium SSD	32 GB to 32 TB	12 to 15,000 IOPS
Standard SSD	32 GB to 32 TB	12 to 2,000 IOPS

Blob storage provides the largest volume of storage on Azure. In addition to storage size, Azure offers both managed and unmanaged storage. With managed storage, Azure takes care of managing the underlying storage accounts. With unmanaged storage, the user takes responsibility for setting up Azure storage accounts of the appropriate size to meet the storage requirements.

## Optimizing Azure compute and storage

You can optimize each tier of processing in an Azure architecture using the most suitable type of VMs and features for that environment. Figure 3 shows one potential pattern for deploying VMs in Azure to support a CICS application that uses Db2. In the primary site, the production, preproduction, and testing VMs are deployed with high availability. The secondary site is for backup and disaster recovery.

Each tier can also provide appropriate disaster recovery services. For example, production and database VMs might require a hot or warm recovery, while the development and testing VMs support a cold recovery.

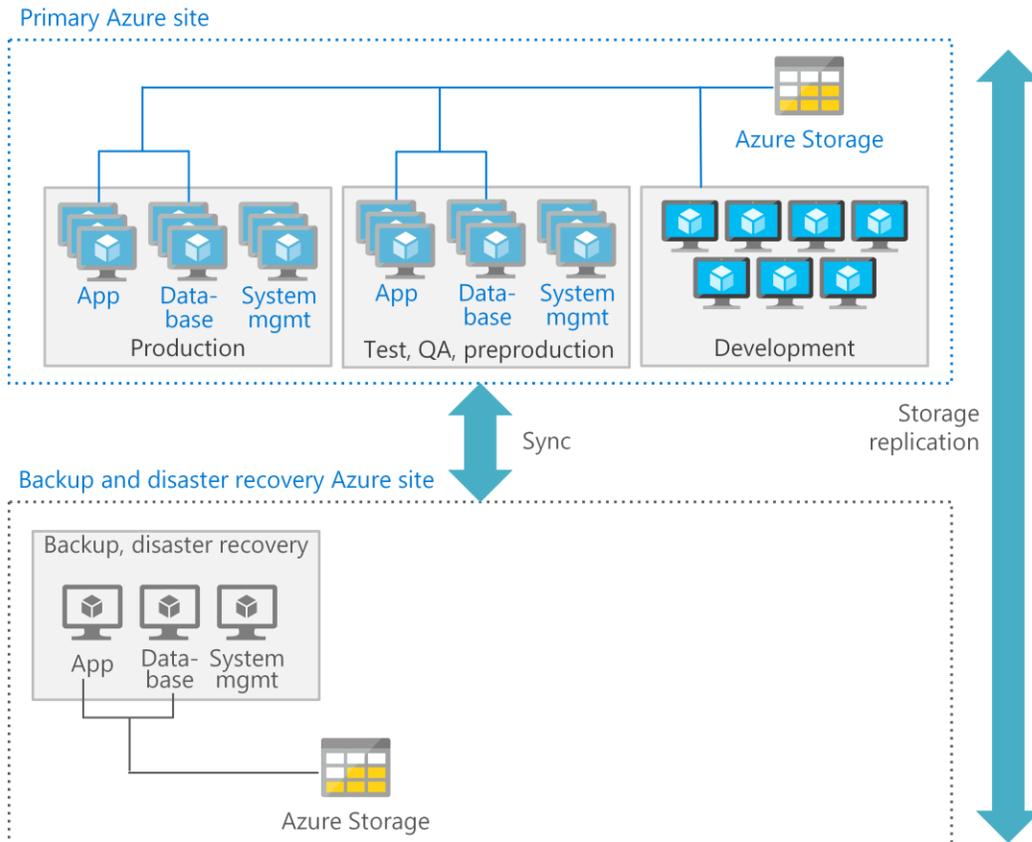


Figure 3. One pattern for a highly available deployment that supports disaster recovery across a primary and a secondary site

## Azure Government

Mainframe workloads that support the government and public sector can be rehosted on Azure Government, the cloud platform specifically for the United States government agencies and their partners. Azure Government services handle data that is subject to certain government regulations and requirements, such as the Federal Risk and Authorization Management Program (FedRAMP), National Institute of Standards and Technology (NIST) 800.171 (DIB), International Traffic in Arms Regulations (ITAR), Internal Revenue Service (IRS) 1075, Department of Defense (DoD) L4, and Criminal Justice Information Services (CJIS). For details, download the [Microsoft Azure Government cloud for mainframe applications](#) white paper.

In addition, Azure Government earned a Provisional Authority to Operate (P-ATO) for FedRAMP High Impact for systems that need this type of environment. Microsoft continues to add features to Azure Government when they meet the requirements for FedRAMP High Impact. For more information, see the [Microsoft Trust Center: FedRAMP](#) website.

## Summary

The massively scalable compute and storage resources on Azure provide immediate benefits to organizations running mainframe workloads. Organizations can begin to modernize without

abandoning the applications they rely on. As this guide shows, Azure provides mainframe-like features and compute and storage capacity that is comparable to IBM z14-based systems.

## Learn more

### Azure compute and storage

- [VM series](#)
- [Deciding when to use Azure Blobs, Azure Files, or Azure Disks](#)
- [Standard SSD Managed Disks for Azure Virtual machine workloads](#)
- [Virtual Machine Scale Sets](#)
- [Load Balancer](#)
- [Kubernetes core concepts for AKS](#)
- [Azure Service Fabric](#)

### IBM resources

- [Parallel Sysplex on IBM Z](#)
- [IBM CICS and the Coupling Facility: Beyond the Basics](#)
- [Creating required users for a Db2 pureScale Feature installation](#)
- [Db2icrt - Create instance command](#)
- [Db2 pureScale Clustered Database Solution](#)
- [IBM Data Studio](#)

### Azure Government

- [Microsoft Azure Government cloud for mainframe applications](#)
- [Microsoft and FedRAMP](#)

### Mainframe migration

- [Platform Modernization Alliance: IBM Db2 on Azure](#)
- [Azure Virtual Data Center Lift and Shift Guide](#)
- [GlusterFS iSCSI](#)