Microsoft

Enabling Data Residency and Data Protection in Microsoft Azure Regions

April 2021

Authors

Christoph Siegert
Debra Shinder
David Burt

Reviewers and Contributors

Derek Harris
Andres Juarez
Karina Juarez
Arnaud Jumelet
Ole Kjeldsen
Peter Koen
Shont Miller
Tia Sargent
Stevan Vidich
Mikko Viitaila
Girish Viriyur
Martin Vliem
Ralf Wigand

Contents

Continued next page.

Introduction

As cloud computing has become increasingly essential for business success, it has brought into sharper focus heightened concerns about protecting data and its privacy and the attendant risks. A key concern is managing data residency—the requirement that data be stored in a specific geographic location. Organizations may want their data to stay within a specific location for a variety of reasons such as taking advantage of a superior tax regime, reducing latency, or avoiding running afoul of a country's privacy laws.

An added layer of complexity is the matter of data sovereignty, the concept that data—particularly personal data—is subject to the laws and regulations of the country in which it is physically collected, held, or processed. Data protection laws have been in effect for many decades, but with the advent of the EU's GDPR and the significant fines that it can impose for noncompliance, organizations have begun to take a much more serious look at the implications of their data sovereignty requirements and capabilities. Customers want to know how services, such as Microsoft Azure, to which they have entrusted their data, will manage where it resides, how it is processed and protected, who has access to it, and where it goes.

Grounded in the firm Microsoft belief in and commitment to the principle that our customers own their data and have a right to control it, Azure is built with the tools and protections customers need to help address their data residency and data sovereignty concerns.

Azure is available in over 140 countries, and offers customers more than 60 datacenter regions worldwide to facilitate data management and transfer in compliance with regional data privacy laws. Most Azure services are deployed regionally and enable the customer to specify the region into which the service will be deployed and control where the customer data will be stored. (Certain services and regions have exceptions and limitations to these rules, which are described in this paper.) In addition, Azure regions unlock cloud adoption, particularly for restricted and regulated industries, by reducing latency, facilitating high availability, and enhancing performance.

When customers move workloads to Azure, they have a number of choices, such as datacenter regions, high availability and disaster recovery architecture, and encryption models. To make the right decisions, customers need to consider technical as well as regulatory requirements. To reduce latency, customers should determine the appropriate region based on the location of their users or customer base. For customers who are targeting a global user base, Azure offers services that ease global deployment.

When it comes to compliance, data residency regulations may govern the physical locations where data can be stored and how and when it can be transferred, processed, or accessed internationally. These regulations can differ significantly depending on jurisdiction. Azure regions and service features give customers choices so they can select and limit data residency as well as control their access to their own data. This enables customers in regulated industries to successfully run mission-critical workloads in the cloud and leverage all the advantages of the Microsoft hyperscale cloud.

This paper covers the information Azure customers need to help them understand how to better control data residency, and meet their data protection obligations within Azure datacenter regions. Specifically, it is structured to address the following:

- Understanding the Azure regional infrastructure, including high availability, disaster recovery, and latency and service availability considerations.
- Data residency assurances, and how customers can control data residency.
- Details about how customers can access diagnostic, service-generated, and support data, and how customers can manage their access to their own data.
- How Microsoft protects customer data from unauthorized access, and how Microsoft handles and challenges government requests and other third-party orders.
- Tools customers can use to restrict, protect, and encrypt data at rest, in transit, and in some cases, in use.
- The strict policies and practices that Microsoft follows for the retention and deletion of customer data.
- How Microsoft compliance with privacy regulations and standards helps protect the privacy of customer data.

I. Infrastructure of Azure regions

The global infrastructure of Azure enables you to deliver services and reach customers and partners wherever they are, while supporting data residency requirements. Azure enables you to choose the location of your data. Azure infrastructure comprises Availability Zones, regions, and geographies.

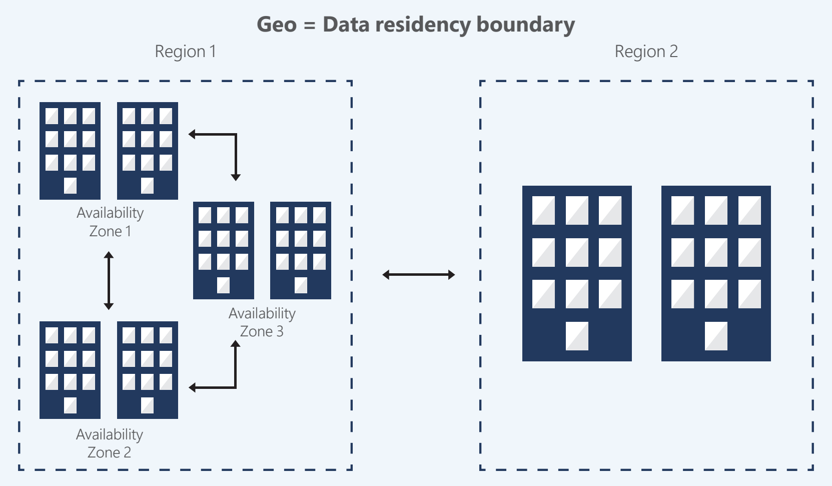

The relationship between Availability Zones, regions, and geographies

Availability Zones

Availability Zones are unique physical locations within an Azure region. Each zone consists of one or more datacenters equipped with independent power, cooling, and networking. Physical separation of Availability Zones within a region protects applications and data from datacenter failures.

An Availability Zone in an Azure region is a combination of a fault domain and an update domain. Zone-redundant services replicate your applications and data across Availability Zones to protect from single points of failure. This architecture also protects against unplanned downtime as well as potential downtime from planned maintenance events. If one datacenter or one Availability Zone fails, zone-redundant Azure services automatically replicate and continue in the other Availability Zones without impacting the customer's zonal applications. Moreover, if the Azure platform is updating for faults or maintenance, the Azure platform recognizes this distribution across update domains to make sure that VMs in different zones are not updated at the same time.

With Availability Zones, Azure offers an industry-best 99.99% VM uptime service-level agreements.

Regions

A region is what the customer typically sees in the Azure portal or command line interface (CLI) as a selectable scope for a deployment location. For example, customers can choose to deploy their VMs into the region US West 2, which will create VMs in the physical location of the Azure US West 2 datacenters. As illustrated in the graphic below, a region can consist of several Availability Zones; for example, US West 2 consists of three Availability Zones. A region can also consist of several datacenters even if the region does not have multiple Availability Zones.

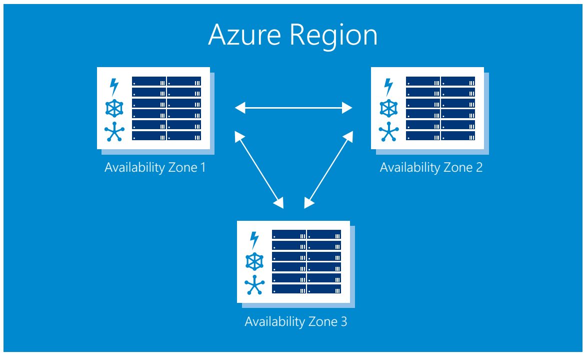

A region can have several Availability Zones.

Geographies

Azure regions are organized into “geographies” or for short, “geos.” An Azure geography ensures that data residency, sovereignty, compliance, and resiliency requirements are honored within geographic boundaries.

A geo can be a country or a set of countries. For example, Canada Central and Canada East regions are in the “Canada” geography and Korea Central and Korea South regions are in the “Korea” geography, while North Europe and West Europe regions are in the “Europe” geography.

- **Geographies that consist of more than one region**. Many Azure regions are paired with another region within the same geography, and together they make a regional pair in a geography. Paired regions are typically hundreds of kilometers apart, providing long-distance disaster recovery within the same geography.
  > Additional details are available in Business continuity and disaster recovery: Azure Paired Regions.[1]
- **Geographies that consist of a single region**. Azure also has geographies that consist of a single region—for example, UAE North is a single region in the United Arab Emirates geography. The UAE geography still provides data residency, although disaster recovery is managed either within the single region, or as out of the country by connecting to other Azure regions.
  > For more information, see Disaster recovery in a geography with a single region (page 11).

Because the geography determines the data residency boundary, it is important to understand the location of each region in a geo. The full list of Azure geographies, including which regions map to which geography, is shown below and on Find the Azure geography that meets your needs.[2]

Azure geography and data residency boundary	Azure regions
Africa	South Africa North South Africa West
Asia Pacific	East Asia Southeast Asia
Australia	Australia Central Australia Central 2 Australia East Australia Southeast
Austria	Austria East (announced)
Brazil (Customer data in Brazil South may be replicated to South Central US for disaster recovery purposes.)	Brazil South Brazil Southeast
Canada	Canada Central Canada East
Chile	Chile Central (announced)
China (Dedicated sovereign cloud with special with special data residency)	China East China East 2 China North China North 2 China North 3 (announced)
Denmark	Denmark East (announced)
Europe	North Europe West Europe
France	France Central France South
Germany	Germany West Central Germany North
Greece	Greece Central (announced)
India	Central India South India West India
Indonesia	Indonesia Central (announced)
Israel	Israel Central (announced)
Italy	Italy North (announced)

Continued next page.

[1] https://aka.ms/paired-regions

[2] https://aka.ms/Azure-geos

Azure geography and data residency boundary	Azure regions
Japan	Japan East Japan West
Korea	Korea Central Korea South
Mexico	Mexico Central (announced)
New Zealand	New Zealand North (announced)
Norway	Norway East Norway West
Poland	Poland Central (announced)
Qatar	Qatar Central (announced)
Spain	Spain Central (announced)
Sweden	Sweden Central (announced) Sweden South (announced)
Switzerland	Switzerland North Switzerland West
Taiwan	Taiwan North (announced)
United Arab Emirates	UAE North UAE Central
United Kingdom	UK South UK North
United States	Central US East US East US 2 East US 3 (announced) North Central US South Central US West Central US West US West US 2 West US 3 (announced)

The Azure regional concept allows customers to achieve two aspects of business continuity—high availability and disaster recovery—while keeping the customer in control of data residency:

- High availability, for example, via Availability Zone SLAs with 99.99% VM uptime
- Disaster recovery via multiple zones in a region, a second region in the Azure geo, or pairing to a region outside of an Azure geo

High availability

High availability refers to solutions that provide service availability, data availability, and automatic recovery from failures that affect the service or data. service-level agreements (SLAs) describe Microsoft commitments for uptime and connectivity. Availability Zones, as described above, provide the highest uptime availability SLAs. In regions without Availability Zones, Availability Sets (a logical grouping of VMs that provides for redundancy and availability) provide 99.95% uptime SLAs.

> See the full list of Azure SLAs in Microsoft Service-level agreements.[3]

High availability in regions with Availability Zones

VMs in an Availability Zone are synchronously replicated across the Availability Zone. If one zone should fail, the VMs in the other zones will continue to run and Azure will load balance without impacting the customer's applications.

Managed Disks, which are like physical disks in an on-premises server but virtualized, deliver consistent performance and high availability within Availability Zones as documented in the Disk FAQs[4] and in Azure premium storage: design for high performance.[5] Note that Managed Disks also ensure that the placement of disks for VMs within an availability set (detailed below) honors fault domain semantics, as documented in Availability options for Azure Virtual Machines.[6] Managed Disks provide redundancy within an Availability Zone, with three replica instances spread across storage stamps in the same datacenter. They also support designing for a recovery point objective of zero hours within an Availability Zone for resiliency and high availability.

[3] https://aka.ms/Azure-SLA

[4] https://aka.ms/Azure-Disc-Res

[5] https://aka.ms/Azure-Prem-Stor

[6] https://aka.ms/Azure-VM-Avail

Zone-redundant storage (ZRS) is a service that replicates Azure Storage data synchronously across three Availability Zones within the same region. ZRS offers durability for Azure Storage data objects of at least 99.9999999999% (twelve 9s) over a given year. With ZRS, data would still be accessible for both read and write operations in the event a zone becomes unavailable. More specifically, in the hypothetical situation where one Availability Zone fails, the Azure platform would undertake networking updates, such as DNS repointing, to enable the other Availability Zones to take on the storage workloads that are kept synchronous, allowing customers to design for zero data loss.

> Refer to Zone-redundant storage[7] for further details.

High availability in regions without Availability Zones

In regions without Availability Zones, availability sets[8] provide 99.95% uptime SLAs. An availability set is a logical grouping capability for isolating VM resources from each other when they are deployed. Azure makes sure that the VMs you place within an availability set run across multiple physical servers, compute racks, storage units, and network switches. If a hardware or software failure occurs, only a subset of your VMs are impacted and your overall solution stays operational.

Maintenance and other downtime

There are three high-level scenarios that can impact the performance of VMs in Azure: planned maintenance, unplanned hardware maintenance, and unexpected downtime.

- **Planned maintenance events** are periodic updates Microsoft makes to the underlying Azure platform to improve the overall reliability, performance, and security of the infrastructure on which VMs run. This includes applying security patches or bug fixes to the hosting environment, or upgrading and decommissioning hardware.
- **Unplanned hardware maintenance events** occur when the Azure platform predicts that the hardware or any platform component is about to fail. When the platform predicts a failure, it will issue an unplanned hardware maintenance event to reduce the impact to VMs hosted on affected hardware. Azure uses live migration technology[9] to migrate the VMs from the failing hardware to a healthy physical machine.
- **Unexpected downtime** occurs when the hardware or the physical infrastructure for the VM fails unexpectedly. When this is detected, the Azure platform automatically migrates customer VMs to a healthy physical machine in the same datacenter. During the healing procedure, VMs experience downtime due to rebooting and in some cases loss of the temporary drive; however, the attached operating system and data disks are always preserved.

> Get the details about maintenance updates for VMs in Maintenance for virtual machines in Azure.[10]

Disaster recovery

All Azure regions are built for hyperscale production workloads. Azure services are designed with redundancy to tolerate faults and minimize disruptions. Azure follows a rigorous testing and production rollout process, which ensures that all technical components are in alignment, and that customer solutions are not negatively impacted by deployment of new versions or processes.

Azure datacenters are designed to run 365 days a year, employing measures to protect operations from physical intrusion, network failures, and power outages. Azure provides hardware, network, local data redundancy, and Distributed Denial of Service (DDoS) protection. Datacenters have dedicated uninterruptible power supplies (UPS) and emergency power support, which includes onsite generators that provide backup power. Regular maintenance and testing are conducted for both the UPS and generators, and operations teams have contractual agreements with local vendors for emergency fuel delivery. Datacenters also have a dedicated Facility Operations Center to monitor power systems, including critical electrical components.

Datacenters are required to test continued operation and resumption of critical datacenter processes in the event of a disruption. Each critical service maintains and tests a disaster recovery plan against each loss scenario to ensure restoration of service within recovery time and recovery point objectives. Any issues identified during testing are resolved, goals are set for continued improvement, and business continuity plans are updated accordingly.

[7] https://aka.ms/AZ-Stor-Redund
[8] https://aka.ms/avail-sets
[9] https://aka.ms/AZ-VM-Migr
[10] https://aka.ms/AZ-VM-main

Azure offers tools that customers can use to design highly available services, employing features such as load balancing, Azure paired regions, Azure Backup, AzCopy, and Azure Storage replication. Systems are proactively monitored to ensure service performance, and achieve availability in accordance with financially backed service-level agreements.

> Customers can validate their disaster recovery strategies following the guidelines in Create and customize recovery plans.[11]

Disaster recovery in a geography with paired regions

For regions that have a regional pair within the same geography, Azure offers convenient disaster recovery options.

For VM workloads, Azure Site Recovery[12] provides an Azure native replication approach from a primary region to a secondary region. When an outage occurs at the primary region, the customer can trigger a failover to the secondary region, with the ability to fail back when the primary region returns to a healthy state. The Azure Site Recovery service also provides a multi-VM consistency group option, which creates a replication group with shared crash-consistent and app-consistent recovery points when failed over.

Azure Platform as a Service (PaaS) services, such as Azure SQL Database, Cosmos DB, or Key Vault also offer native capabilities to replicate data or state to a secondary Azure region.

Several storage solutions take advantage of paired regions to ensure data availability. For example, on top of locally redundant and zonally redundant storage, Azure can copy the data in your storage account to a secondary region that is hundreds of kilometers away from the primary region. If your storage account is copied to a secondary region, then your data is durable even in the case of a complete regional outage or a disaster in which the primary region isn't recoverable. These options work as follows:

- **Geo-redundant storage (GRS)** copies your data synchronously three times in the primary region using locally redundant storage (LRS). It then copies your data asynchronously to a secondary paired region that is hundreds of kilometers away from the primary region. GRS offers durability for Azure Storage data objects of at least 99.99999999999999% (sixteen 9's) over a given year.
- **Geo-zone-redundant storage (GZRS)** combines the high availability provided by redundancy across Availability Zones with protection from regional outages provided by geo-replication. Data in a GZRS storage account is copied across three Azure Availability Zones in the primary region and is also replicated to a secondary region for protection from regional disasters.

Additional cross-region services, such as Azure Backup, are described in the following section.

> For more information on these options, see Redundancy in a secondary region.[13]

[11] https://aka.ms/AZ-Site-recovery

[12] https://aka.ms/AZ-Site-recover-over

[13] https://aka.ms/redund-2nd-region

Not all Azure services automatically replicate data, nor do all Azure services automatically fail back from a failed region to its pair. In such cases, customers must configure recovery and replication. Explore which Azure high availability, disaster recovery, and backup capabilities to use with your apps through the infographic below, Reliability with Azure.[14]

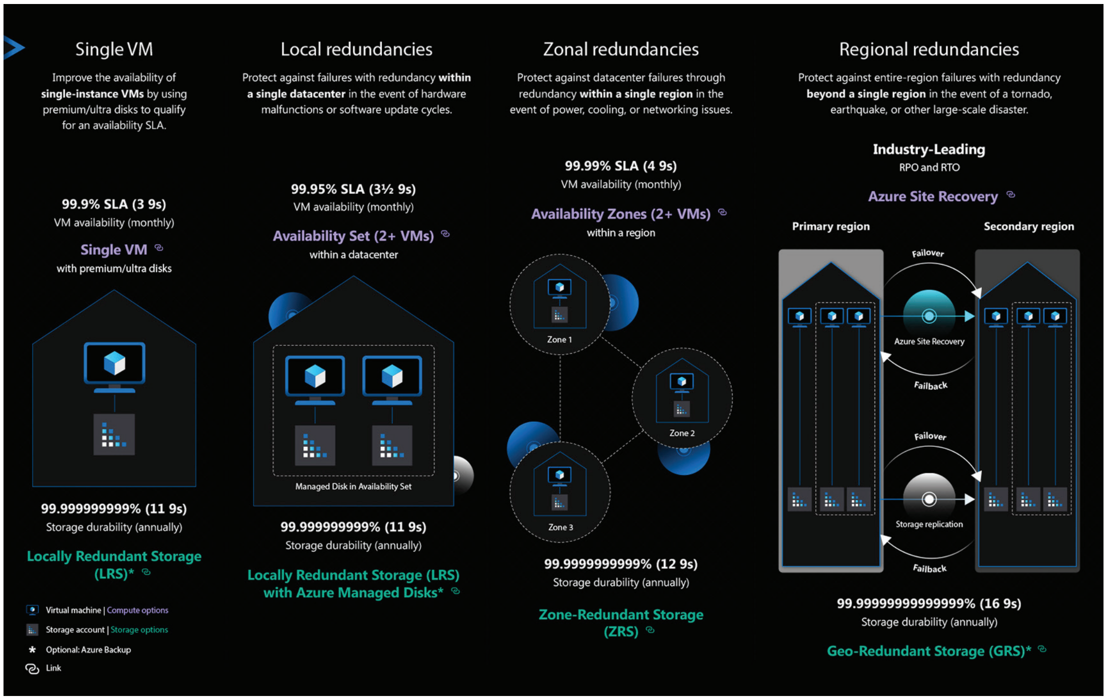

Single VM

Improve the availability of **single-instance VMs** by using premium/ultra disks to qualify for an availability SLA.

99.9% SLA (3 9s)
VM availability (monthly)

Single VM
with premium/ultra disks

99.999999999% (11 9s)
Storage durability (annually)

Locally Redundant Storage (LRS)*

Local redundancies

Protect against failures with redundancy **within a single datacenter** in the event of hardware malfunctions or software update cycles.

99.95% SLA (3½ 9s)
VM availability (monthly)

Availability Set (2+ VMs)
within a datacenter

99.999999999% (11 9s)
Storage durability (annually)

Locally Redundant Storage (LRS) with Azure Managed Disks*

Zonal redundancies

Protect against datacenter failures through redundancy **within a single region** in the event of power, cooling, or networking issues.

99.99% SLA (4 9s)
VM availability (monthly)

Availability Zones (2+ VMs)
within a region

99.9999999999% (12 9s)
Storage durability (annually)

Zone-Redundant Storage (ZRS)

Regional redundancies

Protect against entire-region failures with redundancy **beyond a single region** in the event of a tornado, earthquake, or other large-scale disaster.

Industry-Leading
RPO and RTO

Azure Site Recovery

99.99999999999999% (16 9s)
Storage durability (annually)

Geo-Redundant Storage (GRS)*

Virtual machine | Compute options
Storage account | Storage options
★ Optional: Azure Backup
Link

> You can find a listing of Azure regions pairing within the same geography in Business continuity and disaster recovery: Azure Paired Regions.[15]

Disaster recovery in a geography with a single region

Whether a region without a secondary region in country provides disaster recovery functionality depends on customer architecture and regulatory requirements. If there are no regulatory concerns, customers can use any Azure region for disaster recovery (DR) and store a secondary data set outside the primary country using the services outlined in this section. If the disaster recovery plan uses a secondary region outside of the primary region’s geography, data will be moved outside of the geography or country.

If regulations limit even secondary data copies from leaving the country, customers should explore whether Availability Zones offer sufficient disaster recovery and business continuity. Azure offers data redundancy through Availability Zones by using zonal deployments such as ZRS for storage, zonal VMs, and zone-to-zone DR.

Note that customers cannot define their own regional pairs to take advantage of fully managed geo-replicated services for all services. Customers always have the option to create their own disaster recovery solutions by building services in any number of regions and leveraging Azure services to pair them. Disaster recovery services include Azure Site Recovery (which includes multi-VM consistency) and multi-VM consistency, and Azure Backup.

Azure regions without a DR region in the same geography can rely on the use of Availability Zones to achieve DR with Azure Site Recovery if they want to implement a DR plan without a secondary region. While this won’t protect against full region outages, it can support localized DR scenarios.

[14] https://aka.ms/Reliability-graphic

[15] https://aka.ms/DR-paired-regions

At the time of this writing, Microsoft has announced the following Azure regions without a DR region in the same geo.

- Austria East
- Chile Central
- Denmark East
- Greece Central
- Indonesia Central
- Israel Central
- Italy North
- Mexico Central
- New Zealand North
- Poland Central
- Qatar Central
- Spain Central
- Sweden Central
- Taiwan North

> See the availability by region of any Azure service at Products available by region.[16]

Azure disaster recovery services

Azure Site Recovery

Azure Site Recovery[17] helps ensure business continuity by keeping business apps and workloads running during outages. It replicates workloads running on physical and virtual machines from a primary region to a secondary region. When an outage occurs at the primary region, the customer can trigger a failover to the secondary location and access applications from there. The service is application-agnostic, enabling customers to build disaster recovery for any application hosted on VMs to another zone, within the region or to another region. After the primary region is healthy again, the customer can fail back to it.

Multi-VM consistency, a capability provided by Azure Site Recovery, creates a replication group of all the machines. All of the machines in a replication group have shared crash-consistent and app-consistent recovery points when failed over. Enabling multi-VM consistency can impact workload performance as it is CPU intensive. The maximum number of VMs in a replication group is sixteen.

- Crash-consistent recovery points capture data that was on the disk when the snapshot was taken. This doesn’t include anything in memory, but it contains the equivalent of the on-disk data that would be present if the VM crashed or the power was lost at the instant that the snapshot was taken. A crash-consistent recovery point does not guarantee data consistency for the operating system, nor for apps on the VM. Today, most apps can recover well from crash-consistent recovery points. Azure Site Recovery creates crash-consistent recovery points every five minutes by default.
- Application-consistent recovery points are created from app-consistent snapshots. An app-consistent snapshot contains all the information in a crash-consistent snapshot, plus all the data in memory and transactions in progress. App-consistent snapshots use the Volume Shadow Copy Service (VSS): 1) when a snapshot is initiated, VSS performs a copy-on-write (COW) operation on the volume; 2) before it performs the COW, VSS informs every app on the machine that it needs to flush its memory-resident data to disk; and 3) VSS then allows the backup/disaster recovery application (Azure Site Recovery) to read the snapshot data and proceed.

  App-consistent snapshots are more complex and take longer to complete than crash-consistent snapshots and affect the performance of applications running on a VM enabled for replication. By default, Azure Site Recovery takes an app-consistent snapshot every 4 hours, but it is possible to configure any value between 1 and 12 hours. Azure Site Recovery keeps recovery points for 24 hours by default, but this can be configured to be a value between 1 and 72 hours.

>Learn how to move Azure VMs to another region[18] and enable zone-to-zone disaster recovery[19] for Azure VMs.

Azure Backup

The Azure Backup [20] service keeps your data safe and recoverable in case of a disaster. It allows backups of entire Windows/Linux VMs (using backup extensions), or it can back up files, folders, and system state using the Microsoft Azure Recovery Services (MARS) agent.[21] Backup is also available for Azure Files shares, SQL Server in Azure VMs, and SAP HANA databases running on Azure VMs.

[16] https://aka.ms/Products-by-Region

[17] https://aka.ms/AZ-Site-recover-over

[18] https://aka.ms/Azure-to-Azure

[19] https://aka.ms/site-recovery-zone

[20] https://aka.ms/azure-backup

[21] https://aka.ms/MARS-agent

Additional disaster recovery services

These include Azure DNS and Azure Traffic Manager,[22] which enable customers to design a resilient architecture that will survive the loss of the primary region, as well as AzCopy[23] to schedule data backups to a storage account in a different region.

Because recovery time objectives (RTO) and recovery point objectives (RPO) are dependent on customer architecture, Azure does not provide service-level agreements for RTO and RPO.

Latency considerations

Besides data residency, minimization of latency is also a key value proposition of local Azure regions. Latency is a significant factor in the time it takes to transfer data between a primary and secondary region. In disaster recovery, this is important because it determines whether and how much data will be lost if the primary region fails. Latency between VMs affects the performance of many applications. Bringing cloud services closer to customer workloads enables adoption of latency-sensitive applications.

Latency from customer VMs to Azure VMs

Latency from customer VMs to Azure VMs is highly dependent on customer architecture, and thus Azure does not publish or guarantee a latency SLA. However, Azure offers services to minimize latency from customers' facilities to Azure, and also between resources within Azure. ExpressRoute, accelerated networking, and proximity placement groups all help improve performance by reducing latency.

> Find out how to test Azure VM network latency.[24]

ExpressRoute

ExpressRoute lets you extend your on-premises networks into the Microsoft cloud over a private connection facilitated by a connectivity provider. Connectivity can be from an any-to-any (IP VPN) network, a point-to-point Ethernet network, or a virtual cross-connection through a connectivity provider at a co-location facility. ExpressRoute connections do not go over the public internet. This allows ExpressRoute connections to offer consistent latencies, greater reliability, faster speeds, and higher security than typical connections over the internet.

ExpressRoute connections can be made to an Azure region if available, or alternatively to an ExpressRoute peering location, which then connects to an Azure region via the Microsoft network backbone.

- ExpressRoute Direct[25] gives you the ability to connect directly into the Microsoft global network at peering locations strategically distributed across the world. Each ExpressRoute circuit consists of two redundant connections to two Microsoft enterprise edge routers from the connectivity provider. Microsoft requires dual connections from the connectivity provider, with a redundant Layer 3 connectivity. Microsoft guarantees a minimum of 99.95% ExpressRoute (ER) Dedicated Circuit availability.
- ExpressRoute FastPath,[26] available on all ExpressRoute circuits, is designed to improve the data path performance between your on-premises network and your virtual network. When enabled, FastPath sends network traffic directly to VMs in the virtual network, bypassing the gateway.

Accelerated networking

This is a feature whereby network traffic arrives at the VM's network interface card (NIC), and the NIC forwards network traffic directly to the VM, bypassing the host and the virtual switch.

> Get information about accelerated networking in Create a Windows VM with accelerated networking using Azure PowerShell.[27]

Proximity placement groups

Proximity placement groups offer co-location of Azure VMs in the same datacenter. A proximity placement group is a logical grouping used to make sure that Azure compute resources are located physically close to each other, reducing latency.

> Read more in Introducing proximity placement groups.[28]

[22] https://aka.ms/Disaster-DNS
[23] https://aka.ms/Storage-AZcopy
[24] https://aka.ms/vnet-latency
[25] https://aka.ms/expressroute-er
[26] https://aka.ms/about-fastpath
[27] https://aka.ms/vm-powershell
[28] https://aka.ms/proximity-groups

Latency from Azure region to Azure region

Azure continuously monitors the latency of core areas of its network, using internal monitoring tools as well as measurements collected by a third-party synthetic monitoring service (ThousandEyes).

> Read more in Azure network round-trip latency statistics.[29]

In addition, a helpful latency test website is available to run a generic measurement of latency from your location to any Azure region around the world.

Regional service availability

Customers should consider regional service availability before deploying workloads into an Azure region. Note that in the Azure portal, regions are marked as "recommended region" and "alternate (other) region."

- A recommended region provides the broadest range of service capabilities and is designed to support Availability Zones now or in the future.
- An alternate (other) region extends the Azure footprint within a data residency boundary where a recommended region also exists. Alternate regions help to minimize latency and provide a second region for disaster recovery needs. They are not designed to support Availability Zones, although Azure conducts regular assessment of these regions to determine whether they should become recommended regions.

Azure services are grouped into three categories: foundational, mainstream, and specialized services. The general Azure policy on deploying services into any given region is primarily driven by region type, service categories, and customer demand.

- **Foundational.** Available in all recommended and alternate regions when the region is generally available, or within twelve months of a new foundational service becoming generally available.
- **Mainstream.** Available in all recommended regions within twelve months of the general availability of the region or service; demand-driven in alternate regions. (Many are already deployed into a large subset of alternate regions.)
- **Specialized.** Targeted services offerings that are usually aligned with specific industries or specialized hardware, or in response to specific regional demand.

> Mapping services into these categories is available in Regions and Availability Zones in Azure.[30]

The list of services by region is available in Products available by region.[31] Not all services are available in every region. For services that are not yet available, but are on the deployment roadmap, the website provides an estimated availability date.

[29] https://aka.ms/AZ-latency

[30] https://aka.ms/AZ-overview

[31] https://aka.ms/Products-by-Region

II. Data residency for customer data

Data residency for regional services

As a customer, you retain all right, title, and interest in and to customer data—personal data and other content—that you provide for storing and hosting in Azure services. Microsoft will not store or process customer data outside the geography you specify, except for certain services and scenarios that are discussed below. You are also in control of any additional geographies where you decide to deploy your solutions or replicate your data. In addition, you and your users may move, copy, or access your customer data from any location globally.

Most Azure services are deployed regionally and enable you to specify where your customer data will be stored and processed. Examples of such regional services include VMs, storage, and SQL Database. For a complete list, see Products available by region.[32]

For regional services, customers preselect the region in which the service will be deployed. (Note that when you select the region in Azure, you automatically choose the geography that contains that region.) Deployment location (and thereby data residency) can be defined by the region variable in the Azure portal or via the command line interface.

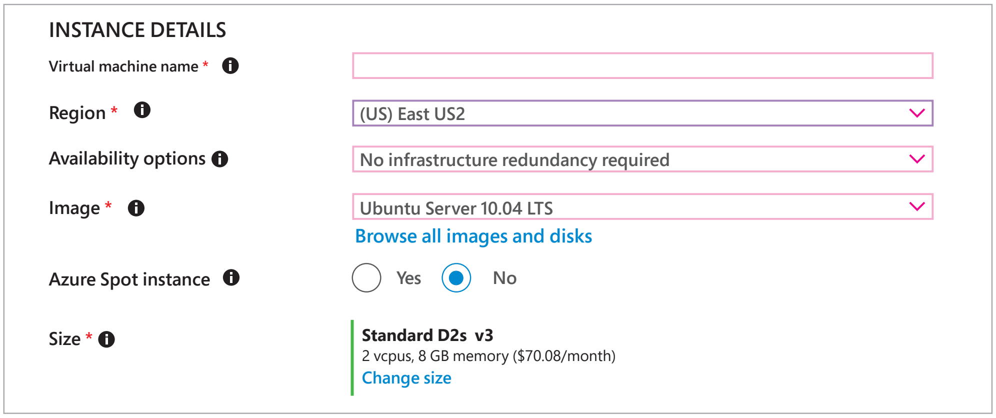

INSTANCE DETAILS

Virtual machine name *	
Region *	(US) East US2
Availability options	No infrastructure redundancy required
Image *	Ubuntu Server 10.04 LTS Browse all images and disks
Azure Spot instance	○ Yes ● No
Size *	**Standard D2s v3** 2 vcpus, 8 GB memory ($70.08/month) Change size

How to define a region in the Azure portal

Azure regions are often paired with another region within the same geography, which together make a regional pair. The regional pair is always inside the specified Azure geography; for example, Canada Central and Canada East are the regional pair in the Canada geo. Because the geography determines the data residency boundary, it is important to understand the location of each region in a geo. The full list of Azure geographies, including the alignment of regions to geographies, is available in Find the Azure geography that meets your needs.[33]

Microsoft may replicate customer data to other regions within the same geo for data resiliency. For example, if a customer deploys Azure Blob Storage in Japan East, the customer data may be replicated to Japan West for disaster recovery purposes, but will remain stored inside Japan.

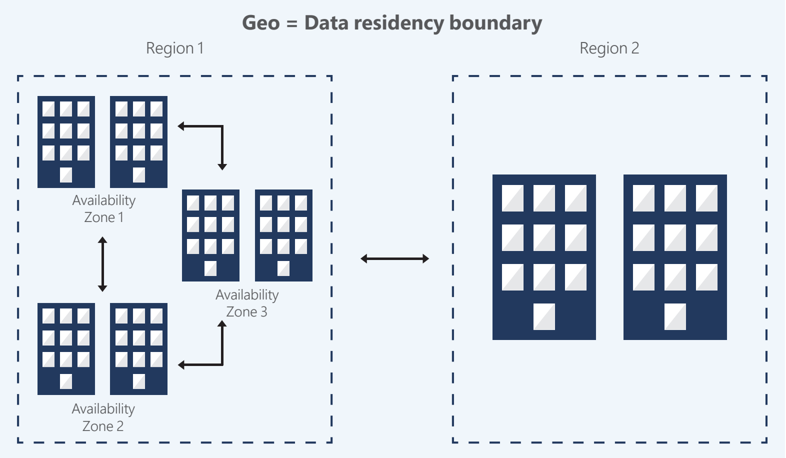

Relationship between Azure Availability Zones and regions

[32] https://aka.ms/Products-by-Region

[33] https://aka.ms/Azure-geos

To maintain resiliency, Microsoft uses variable network paths that sometimes cross geo boundaries; however replication of customer data between regions is always transmitted over encrypted network connections.

Microsoft personnel (including subprocessors) located outside the geography may remotely access data processing systems in the geo, but will not access customer data without authorization by the customer.

> For a list of all regional services, go to Data residency in Azure.[34]

Regional services that enable data residency in a single region

Customers can configure certain Azure services, tiers, or plans to store customer data only in a single region, with certain exceptions. These include Azure Backup, Azure Data Factory, Azure Site Recovery, Azure Stream Analytics, and locally redundant storage (LRS).

> To see the full list of these, go to Data residency in Azure.[35]

Data residency exceptions for regional services

Certain regional services store or process customer data outside the customer-specified geography.

- **Azure Cloud Services**, which back up web and worker-role software deployment packages to the United States regardless of the deployment region.
- Language Understanding,[36] which may store active learning data in the United States, Europe, or Australia depending on the authoring regions that the customer uses.
- **Azure Machine Learning**, which may store freeform text that the customer provides (such as names for workspaces, resource groups, experiments, files, and images) and experiment parameters in the United States.
- **Azure Databricks**, which stores identity data, and certain table names and object path information in the United States.
- **Azure Sentinel**
- Azure Serial Console,[37] which stores all customer data at rest in the geography selected by the customer, but when used through the Azure Portal may process console commands and responses outside of the geo for the sole purpose of providing a console experience inside the portal.
- **Preview, beta, or other pre-release services**, which typically store customer data in the United States but may store it globally.

Data residency for non-regional services

Certain Azure services do not allow the customer to specify the region where the service will be deployed. These services may store or process customer data in any Microsoft datacenter unless specified otherwise below:

- **Content Delivery Network (CDN)** provides a global caching service and stores customer data at edge locations around the world.
- **Azure Active Directory (AAD)** may store Azure AD data globally. This does not apply to Azure AD deployments in the United States (where Azure AD data is stored solely in the United States) and in Europe (where Azure AD data is stored in Europe or the United States). > For more information, see Data storage for AAD identity data (page 17).
- **Azure multifactor authentication (MFA)** stores authentication data in the United States. > For the details, see Data residency and customer data for Azure multifactor authentication.[38]

[34] https://aka.ms/Azure-Data-Res

[35] https://aka.ms/Azure-Data-Res

[36] https://aka.ms/LUIS-location

[37] https://aka.ms/AZ-serial-con

[38] https://aka.ms/Data-residency-AADMFA

- **Azure Security Center** stores a copy of security-related customer data, collected from or associated with a customer resource (such as a VM or an Azure AD tenant), in the same geo as that resource, except in those geos where Microsoft has yet to deploy Azure Security Center, in which case a copy of such data will be stored in the United States. And where Azure Security Center uses another Microsoft online service to process such data, it may store it in accordance with the geolocation rules of that other online service.
  > For more information, see Azure Security Center.[39]
- **Services that provide global routing functions and do not themselves process or store customer data**. These services include Traffic Manager,[40] which provides load balancing between different regions, and Azure DNS,[41] which provides domain name services that route to different regions.

> For a complete list of non-regional services, see Products available by region[42] and select **Non-regional** for **Region**.

Data residency for AAD identity data

Azure Active Directory (AAD) stores most identity data in the geographic location based on the address that an organization uses when subscribing to a Microsoft online service.

> Get detailed information on where AAD data is located in the AAD section of Where is your data located?[43]

For example, if the address provided is in Europe, AAD keeps most of the identity data within European datacenters, although the following data may be stored outside of Europe:

- Multifactor authentication phone calls or SMS.
- Push notifications using the Microsoft Authenticator app.
- AAD B2C policy configuration data and Key Containers.
- B2B invitations with redeem link and redirect URL information and email addresses of users that unsubscribe from receiving B2B invitations.
- Microsoft Exchange Server 2013 Application Identifier (AppID), the approved federated domains list for application, and the application's token signing Public Key.

> For more information, see Identity data storage for European customers in AAD.[44]

Using Azure Policy to control data residency

Microsoft provides Azure Policy to implement governance over cloud infrastructure and data, including but not limited to regions in which resources can be deployed, which services can be deployed, and resource monitoring requirements. To restrict the data and resources to certain Azure regions, such as for data residency, customers can use the Allowed Locations[45] policy.

Once policies are established, not only will new resources that are deployed be checked against the policies, but all resources will be periodically scanned to help ensure ongoing compliance.

> Get more information in What is Azure Policy?[46]

[39] https://aka.ms/AZ-sec-center

[40] https://aka.ms/AZ-traffic

[41] https://aka.ms/DNS-overview

[42] https://aka.ms/Products-by-Region

[43] https://aka.ms/AAD-Data-Residency

[44] https://aka.ms/AAD-Storage-EU

[45] https://aka.ms/AZPolicySamples

[46] https://aka.ms/WhatIsAzurePolicy

Using Azure Blueprints to enforce compliance and data residency

Azure Blueprints[47] is a free service that supplies templates to create, deploy, and update fully governed cloud environments to consistent standards, which helps customers comply with regulatory requirements. It differs from Azure Resource Manager (ARM) and Azure Policy in that Blueprints is a package that contains different types of artifacts—including ARM templates, resource groups, policy assignments, and role assignments—all in one container, so you can quickly and easily deploy all these components in a repeatable configuration. Blueprints help to simplify large-scale Azure deployments by packaging policies in a single blueprint definition.

The Azure Blueprints service provides built-in blueprints mapped to key portions of common standards such as HIPAA, FedRAMP, and PCI DSS. For example, the ISO 27001 Blueprint[48] and the PCI DSS Blueprint[49] map a core set of policies for those respective standards to any Azure environment. Blueprints can be deployed to multiple Azure subscriptions and managed from a central location, and are scalable to support production implementations for large-scale migrations.

You can use the built-in blueprints or create your own custom blueprints. Blueprints can be created in the Azure portal or using the REST API with tools such as PowerShell. If the latter method is used, you can create blueprint parameters to prevent conflicts when reusing certain blueprints.

Blueprints can be used to help manage data residency for specific compliance needs by specifying both allowed locations and allowed locations for resource groups—for example, control mapping of the UK OFFICIAL and UK NHS blueprint samples.[50] You can also use the regulatory compliance dashboard for insight into your compliance posture based on how you're meeting specific compliance requirements.

[47] https://azure.microsoft.com/services/blueprints/

[48] https://aka.ms/ISO-27001-Blueprint

[49] https://aka.ms/PCI-DSS-Blueprint

https://aka.ms/bp-parameters

[50] https://aka.ms/UKOFFICIAL-Blueprint

https://aka.ms/comp-dashboard

III. Access to diagnostic, service-generated, and support data

Microsoft relies on a set of definitions for types of data outlined in the Microsoft Online Services Data Protection Addendum[51] and Microsoft Online Services Terms.[52]

- Diagnostic data refers to data collected or obtained by Microsoft from software that is locally installed by the customer for use with an Azure service. Diagnostic data may also be referred to as telemetry. Diagnostic data does not include customer data, service-generated data, or professional services data.
- Service-generated data refers to data generated or derived by Microsoft through the operation of an Azure service. Service-generated data does not include customer data, diagnostic data, or professional services data.
- Support data refers to all data, including all text, sound, video, image files, or software, that are provided to Microsoft by or on behalf of an Azure customer to obtain technical support for Azure services. Support data is a subset of professional services data.

For services where customer data is stored and processed in the cloud, service-generated, diagnostic, and support data include application and server logs that are required to maintain modern applications and platforms. These logs provide customers with the information they need to operate and troubleshoot their workloads, and provide Microsoft with the information it needs to operate, troubleshoot, and improve the platform.

Microsoft policies regarding data at rest in some cases do not apply to the three types of data defined above. Microsoft uses such data for clearly defined scenarios, in line with GDPR requirements and aligned to the best practices described in ISO/IEC 19944[53] (the standard that describes data flow, data categories, and data use in the cloud).

> See Who can access customer data and on what terms (page 23) for further information regarding limitations on access to customer data, as well as diagnostic, service-generated, and support data.

Diagnostic data

This includes data collected or obtained by Microsoft from software that is locally installed by the customer for use with an Azure Service, such as log files, system-generated event logs, registry keys, debug logs, server and database information, console screenshots, and basic network and storage disk information.

For App Service-related issues, HTTP logs, detailed errors, KUDU trace, transform logs, FREB logs, winsock logs, event logs, DAAS logs, and Webjob logs are collected to help with troubleshooting.

For Azure AD Connect-related issues, information about Active Directory objects (such as user and device properties), your synchronization configuration, and related log files (such as Sign-In, Audit, or synchronization logs) are collected to help with troubleshooting.

> Get a detailed list of diagnostic data that Microsoft collects in the following:

- Windows Server logs [54]
- Azure PaaS VM logs [55]
- Azure IaaS logs [56]
- Azure Service Fabric logs [57]
- StorSimple support package and device logs [58]
- SQL Server on Azure VM logs [59]
- Azure Active Directory diagnostic logs [60]

[51] https://aka.ms/MS-DPA

[52] https://aka.ms/Online-Services-Terms

[53] https://www.iso.org/standard/79573.html

[54] https://aka.ms/WinServer-Logs-Diag

[55] https://aka.ms/AZ-PaaS-logs

[56] https://aka.ms/AZ-IaaS-logs

[57] https://aka.ms/AZ-Serv-Fab-logs

[58] https://aka.ms/StorSimple-logs

[59] https://aka.ms/SQL-VM-logs

[60] https://aka.ms/AAD-dia-logs

Transparency and customer control of diagnostic data

Azure provides transparency and control functions for some of the most common types of diagnostic data scenarios for Windows VMs, Linux VMs, and virtual networks.

> Get an overview of Azure Monitor agents.[61]

Windows VMs on Azure

Virtual machines and other compute resources require an agent to collect monitoring data to measure the performance and availability of their guest operating system and workloads.

- The Azure Diagnostics extension[62] collects monitoring data from the guest operating system and workloads of Azure VMs and other compute resources.
- The Log Analytics agent[63] collects monitoring data from the guest operating system and workloads of VMs in Azure.
- The Dependency Agent[64] collects discovered data about processes running on the VM and external process dependencies.

	Diagnostics extension (WAD)	Log Analytics agent	Dependency agent
Data collected	Event Logs ETW events Performance File-based logs IIS logs .NET app logs Crash dumps Agent diagnostics logs	Event Logs Performance File-based logs Insights and solutions Other services	Process details and dependencies Network connection metrics
Data sent to	Azure Storage Azure Monitor Metrics Event Hub	Azure Monitor Logs	Azure Monitor Logs

A quick comparison of the Azure Monitor agents for Windows

For Windows VM diagnostic data, Windows Server images on Azure are set up similarly to off-the-shelf products. Customers can control the diagnostic data they share with Microsoft—for example, Windows Security Baselines can be used to efficiently configure Windows 10 and Windows Server settings for best security practices.

> Get more information on:

- How to use security baselines in Windows security baselines.[65]
- The Windows Restricted Traffic Limited Functionality Baseline in Manage connections from Windows operating system components to Microsoft services.[66]

Linux VMs on Azure

For customers who run Linux VMs on Azure, Microsoft provides the Azure Linux Agent[67] (WALinuxagent) as open source software for Linux. The Azure Linux Agent manages Linux and FreeBSD provisioning and VM interaction with the Azure Fabric Controller. Microsoft provides full transparency so administrators will know which data is sent from Linux to the Azure platform. This information can be correlated and used for further analysis, to monitor important system metrics and perform data-based decisions.

Additionally, a diagnostics extension, log analytics, and additional agents can be implemented to analyze application level logs, as outlined in Overview of Azure Monitor Agents.[68]

Service-generated data

Service health is an important aspect of operating Azure services, and so Microsoft collects schematized service-generated data to diagnose and perform root-cause analysis on incidents for the platform. Every service uses this data to trigger self-healing processes, which reduces the human intervention required. As an example, if the load on a specific component increases, the platform assigns more resources to manage the load. Microsoft has integrated anonymized service-generated data in the Azure DevOps tools without accessing customer data.

[61] https://aka.ms/AZ-Mon-Agents
[62] https://aka.ms/AZ-Mon-Diag
[63] https://aka.ms/AZ-Mon-log
[64] https://aka.ms/AZ-Dep-Agent
[65] https://aka.ms/Win-Sec-Baselines
[66] https://aka.ms/Win-Sec-Connect
[67] https://aka.ms/AZ-Linux-Agent
[68] https://aka.ms/AZ-Mon-Agents

Customer access to service-generated data

- **Azure Monitor.**[69]
  Customers can use Azure Monitor to manage the health of their workloads. This service provides a 360-degree view of applications, infrastructure, and networks with advanced analytics, dashboards, and visualization maps. Azure Monitor gives the customer a centralized hub that helps to identify network glitches, CPU spikes, memory leaks in code, and other issues before they impact the customer's workload. Azure Monitor also offers various ways to notify administrators in case of alerts, like Action rules, e-mail, SMS, or Logic Apps.
- **Application Insights.**[70]
  Azure Monitor also includes Application Insights which provides an extensible Application Performance Management (APM) service for web developers on multiple platforms. It can be used to monitor web applications during runtime. It will automatically detect performance anomalies, and it includes powerful analytics tools to help diagnose issues and understand user interactions with applications.

  Application Insights is designed to help continuously improve performance and usability by sending service-generated data from the customer's web applications to the Azure portal. It works for apps on a wide variety of platforms, including .NET, Node.js, and J2EE, hosted on premises or in the cloud. Collectors are designed to provide a schematized output of data, limit the transmission of personal data as much as possible, and transmit data securely. A data retention policy must be defined by the user. The customer can also use this data to build high-availability workloads, which could detect an incident based on the data and perform automated predefined actions to mitigate the incident.
- **Azure Network Watcher.**[71]
  Customers can leverage Azure Network Watcher to monitor network traffic associated with their IaaS workloads. This service provides tools to monitor, diagnose, view metrics, and enable or disable logs for resources in an Azure virtual network. It is designed to monitor and repair the network health of IaaS products, which include VMs, Virtual Networks, Application Gateways, and Load Balancers. Note that Network Watcher is not intended for and will not work for Platform as a Service (PaaS) monitoring or Web analytics.

Support data

Customers initiating a support request can give Microsoft engineers access to support data—data that a customer provides to Microsoft to obtain technical support for Azure services. Through the "Share diagnostic information" feature in the Azure portal, you can authorize a Microsoft support engineer to remotely collect data related to the current support incident from your Azure Virtual Machines or Azure Cloud to troubleshoot your issue. Support data can be retained for up to 90 days.

> Get more information in Azure Support diagnostic information and memory dump collection.[72]

Memory dump

When a customer VM crashes, customer data may be contained inside a memory dump file on the VM. Customer data remains in the region where the VM is deployed, and does not leave the Azure data residency boundary.

By default, Microsoft engineers do not have access to customer VMs and cannot review crash dumps without the customer's approval. Before investigating a VM crash dump, engineers must gain explicit customer authorization to access customer crash dump data. Access is gated by the Just-in-time (JIT) privileged access management system and Customer Lockbox so that all but extraordinary actions, such as major outages and confidential law enforcement requests, are logged and audited.

> Get more information on which services are covered by Customer Lockbox[73] and on processes for memory dump in Azure for Secure Worldwide Public Sector Cloud Adoption.[74]

[69] https://aka.ms/AZ-Monitor

[70] https://aka.ms/APP-Insights-Over

[71] https://aka.ms/AZ-Netwatch

[72] https://aka.ms/AZ-Legal-Diags

[73] https://aka.ms/msazurelockbox

[74] https://aka.ms/Azure-WWPS

Enabling boot diagnostics

Customers can also choose to enable boot diagnostics, which captures logs, the serial console output, and screenshots from the host running the VM. Enabling boot diagnostics also allows the Azure platform to inspect the Operating System Virtual Hard Disk (OS VHD) for VM provisioning errors, helping to provide deeper information on the root causes of failures. Access to the OS VHD includes guest operating system information, system files on the OS VHD, and custom scripts.

Elevated DevOps data access for support cases

When access to customer data is granted, Microsoft internal leadership approval is required and then access is carefully managed and logged.[75] The access-control requirements are established by the following Azure security policies:

- There is no access to customer data by default.
- There are no user or administrator accounts on customer VMs.
- Grant the least privilege that's required to complete the task, and audit and log access requests.
- Azure support personnel are assigned unique corporate Active Directory (AD) accounts by Microsoft. Azure relies on Microsoft corporate AD, managed by Microsoft Information Technology (MSIT), to control access to key information systems. Multifactor authentication is required, and access is granted only from secure consoles.
- All access attempts are monitored.

The design principles defined for the development of Azure services require a schematized setup for support data. In Azure, the most important use case for this data is its use as a sensor for the automated operation of the cloud. Based on that data and desired state configuration, remediation activities are triggered via automation, thus reducing the additional risk caused by manual human intervention.

To manage DevOps data access, solutions include Customer Lockbox and customer-managed encryption keys, which are described in How customers can protect data from unauthorized access (page 28).

[75] https://aka.ms/audit-logs

IV. How Microsoft protects access to customer data

Microsoft takes strong measures to help protect your customer data from unauthorized access. In addition to the physical and technological protections, there are access restrictions for Microsoft personnel and subcontractors, as well as stringent requirements for responding to government requests for customer data.

Who can access customer data and on what terms

Controlling Microsoft access to customer data

Only in rare cases will a Microsoft engineer access or be exposed to customer data. Such access may be required to resolve a support request initiated by a customer, or as part of a Microsoft-initiated maintenance or troubleshooting operation on underlying software. Nearly all service operations performed by Microsoft are fully automated and human involvement is highly controlled and abstracted away from customer data.

Access to customer data by Microsoft operations and support personnel is denied by default. Should access to customer data be required, it is restricted based on business need by role-based access controls, multifactor authentication, minimization of standing access to production data, and other controls. Access to the platform of DevOps personnel is requested via the Just-in-time (JIT) access tool. All access to customer data is strictly logged, and both Microsoft and third parties perform regular audits (as well as sample audits) to attest that any access is appropriate.

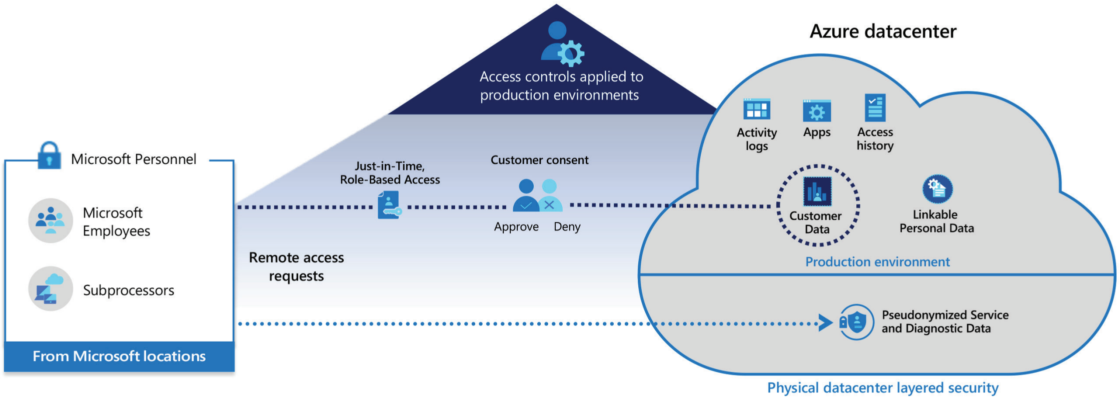

How Microsoft controls access to customer data

Azure undergoes a SOC audit by an AICPA-accredited auditor twice a year to verify the effectiveness of its security controls in audit scope. The Azure and Azure Government SOC 2 Type 2[76] attestation report published by the auditor explains the circumstances when access to customer data can occur and how.

For the majority of customer-initiated support requests, access to customer data is not needed. However, the most common scenario by far involves a customer opening a troubleshooting ticket with Azure Support, and Support subsequently obtaining an authorization to access customer resources that could potentially include customer data. When access to customer data is needed, customers can manage that access, as outlined in Support data (page 21).

Preventing unauthorized access to customer data

Microsoft employs rigorous operational controls and processes to prevent unauthorized physical access to datacenters, including video monitoring, trained security personnel, as well as smart card and biometric access controls. Since data in Azure is 1) encrypted, and 2) stored across multiple physical disks, even in the highly unlikely scenario that someone could remove selected physical disks (and knew which disks to remove), the data would be unreadable. Upon end of life, data disks are shredded and destroyed as outlined in Data disk destruction (page 34).

[76] https://aka.ms/AZ-SOC

Customers can use customer-managed keys to further help prevent their data from being readable in case of unauthorized access. Both server-side and client-side encryption can rely on customer-managed keys or customer-provided keys.

> For more information about encryption and key management, see [How customers can protect data from unauthorized access](#) (page 28).

How Microsoft manages customer data

With Azure, you are the owner of your customer data and retain all right, title, and interest in and to customer data. Microsoft provides contractual commitments that make this clear.

Your data is your business, and you can access, modify, or delete it at any time. Microsoft only processes your data based on your authorization, and in accordance with the strict policies and procedures that we have contractually agreed to. Microsoft does not share customer data with Microsoft advertiser-supported services or for similar commercial purposes, process it for user profiling, or mine it for any purposes such as marketing research or advertising.

Your control over your data is reinforced by the Microsoft commitment to comply with broadly applicable data protection and privacy laws. To that end, Microsoft has created a set of internal policies and technical controls that govern data handling and are designed to conform with both ISO/IEC 27018 Code of Practice for Protecting Personal Data in the Cloud[77] and ISO/IEC 27701 Privacy Information Management System (PIMS).[78]

- ISO/IEC 27018 protects personally identifiable information (PII) in public clouds that process PII. For example, for application code, any output written to the logfiles goes through data scrubbers that remove customer data before the data is sent to central systems. These measures minimize the risk of customer data being replicated to analysis or operations repositories.
- Microsoft has also achieved certification as a data processor for the new international standard ISO/IEC 27701 (PIMS). The PIMS certification demonstrates that Azure provides a comprehensive set of management and operational controls that can help organizations demonstrate compliance with privacy laws and regulations. The Azure implementation creates a strong integration point for aligning security and privacy controls through a framework for managing personal data that can be used by both data controllers and data processors, a key distinction for GDPR compliance.

You have access to independent audit reports of our compliance with these privacy standards, which in turn offer support for meeting your own privacy obligations.

> Learn more about how to access audit reports in the Sharing responsibility for compliance (page 36).

How Microsoft keeps customer data separate

Azure is a multitenant service, which means that multiple customer deployments and VMs are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's data from the data of others. To keep customers' data secure in the multitenant environment, the Azure platform uses a virtualized environment, whereby workloads from different tenants run in isolation on shared physical servers. User instances operate as standalone VMs that do not have access to a physical host server, and this isolation is enforced by using physical processor privilege levels.

Hypervisors are designed to be as small as possible and undergo rigorous security reviews to prevent a workload from being able to detect other workloads. Each workload sees a virtual storage device containing only the files associated with its own data. Moreover, the hypervisor has complete control to start, stop, and pause workloads. It also controls the physical network cards, so it can filter all the network packets based on the workload identity and tenant. The physical storage media contents are tagged with the tenant owner and associated VM.

In addition, tenants can control their network connectivity between servers and the internet, and they can create separate virtual networks for different purposes such as production, development, and testing. The hosting provider's fabric controller coordinates with the hypervisors hosting workloads for each tenant to make sure only workloads on the same virtual networks of a tenant can see each other's traffic or have connectivity to the internet.

[77] https://aka.ms/AZ-ISO27018
[78] https://aka.ms/AZ-ISO27701

Microsoft also offers the Azure Dedicated Host[79] service that supplies physical servers—able to host one or more VMs—dedicated to one Azure subscription and providing hardware isolation at the physical server level. No other VMs will be placed on your hosts. Dedicated hosts are deployed in the same datacenters and share the same network and underlying storage infrastructure as other, non-isolated hosts and you can provision them within a region, Availability Zone, and fault domain.

> Get more information on Isolation in the Azure Public Cloud.[80]

When Microsoft will delete customer data

If you leave the Azure service or your subscription terminates, Microsoft abides by its commitment in the Online Services Terms and follows specific processes for:

- Removing customer data from cloud systems under its control within specified time frames.
- Overwriting storage resources before reuse.
- Physical destruction of decommissioned hardware.

> Learn more about how Microsoft handles data upon service termination in Data deletion (page 33).

How Microsoft manages subcontractors and subprocessors

Where Microsoft hires a subcontractor to perform work that may require access to customer data, the subcontractor is considered a subprocessor. Microsoft publicly discloses these subprocessors in the Microsoft Online Services Subprocessor List.[81]

Subprocessors may access customer data only to deliver the functions in support of online services that Microsoft has hired them to provide and are prohibited from using customer data for any other purpose. They are required to maintain the confidentiality of this data and are contractually obligated to meet strict privacy requirements that are equivalent to or stronger than the contractual commitments Microsoft makes to its customers in the Microsoft Online Services Data Protection Addendum[82] and the Standard Contractual Clauses.[83] Subprocessors are also required to meet EU General Data Protection Regulation (GDPR) requirements, including those related to implementing appropriate technical and organizational measures to protect personal data.

When engaging new subprocessors, Microsoft will provide notice to customers of any new subprocessor at least six months in advance of providing them with access to customer data, including personal data in customer data. It does this by updating the Microsoft Online Services Subprocessor List, which identifies authorized subprocessors who have been audited against a set of stringent security and privacy requirements in advance, and by providing the customer with a mechanism to be notified of that update. In addition, Microsoft will notify customers (using the same methods described above) of any new subprocessor at least 30 days in advance of providing them with access to personal data other than that contained in customer data. If Microsoft engages a new subprocessor for a new online service, Microsoft will give the customer notice before making that service available to the customer.

> To receive notifications of updates to the Subprocessor List, follow the instructions that described in the Service Trust Portal: My library.

If a customer does not accept a new subprocessor, the customer may terminate their subscription for the affected online service without penalty by providing, before the end of the relevant notice period, written notice of termination that includes an explanation of the grounds for non-acceptance. If the affected online service is part of a suite (or similar single purchase of services), then the termination will apply to the entire suite.

[79] https://aka.ms/vm-dedicated

[80] https://aka.ms/AZ-isolation-choices

[81] https://aka.ms/Azure-subprocessors

[82] https://aka.ms/MS-DPA

[83] https://aka.ms/AZ-SCC

https://aka.ms/receive-notific

How Microsoft handles government requests

Microsoft has taken a firm public stand on protecting customer data from inappropriate government access. Through clearly defined and well-established response policies and processes, strong contractual commitments, and if need be, the courts, Microsoft defends your data. We believe that all government requests for your data should be directed to you. We do not give any government direct or unfettered access to customer data.

Microsoft is principled and transparent about how we respond to requests for data. Because we believe that you have control over your own data, we will not disclose data to a government except as you direct or where required by law. Microsoft scrutinizes all government demands to ensure they are legally valid and appropriate.

If Microsoft receives a demand for a customer's data, we will direct the requesting party to seek the data directly from the customer. If compelled to disclose or give access to any customer's data, Microsoft will promptly notify the customer and provide a copy of the demand unless legally prohibited from doing so.

> For data about enterprise customers who were impacted by law enforcement requests, see Law enforcement requests (page 27).

Microsoft contractual commitments to our commercial and public sector customers include Defending Your Data,[84] which builds on our existing protections and demonstrates our confidence in our ability to protect customer data against exposure to inappropriate disclosure:

- Microsoft will challenge every government request for commercial and public sector customer data—from any government—where there is a lawful basis for doing so. We have a proven track record of successfully using the courts to challenge government demands that are inconsistent with the rule of law. We have more experience than any other company taking the US government to court to challenge orders seeking access to an individual's data and to protect our ability to tell customers about those orders, even taking one case[85] to the US Supreme Court. Our challenges have led to greater protections and transparency for our customers worldwide, which include enabling us to disclose reports about the number of US national security orders we receive.
- We stand behind the strength of our GDPR compliance and other data protection safeguards. However, to provide added reassurance against liability for our commercial and public sector customers, we will provide monetary compensation to these customer's users if we disclose their data in response to a government request in violation of the GDPR.

  > For in-depth information about how Microsoft handles government requests, see About our practices and your data.[86]

The CLOUD Act

The Clarifying Lawful Overseas Use of Data Act—better known as the CLOUD Act—was passed into law in the United States in 2018. This was the culmination of many years of legislative work by the tech industry, Congress, and the administration to balance the privacy rights of our customers around the world with the law enforcement need to access data stored outside the United States.

The CLOUD Act is not a mechanism for greater government surveillance; it is a mechanism for assisting in specific criminal investigations. It aims to ensure that customer data is ultimately protected by the laws of each customer's home country while continuing to facilitate lawful access to evidence for legitimate criminal investigations. This does not change any of the legal and privacy protections that previously applied to law enforcement requests for data—and those protections continue to apply. US law enforcement must still obtain a warrant demonstrating probable cause of a crime from an independent court before seeking the contents of communications.

The CLOUD Act preserves existing protections:

- It preserves the common law right of Microsoft and other cloud service providers to go to court to challenge search warrants when there is a conflict of laws.
- Microsoft retains the legal right to object to a law enforcement order in the United States when the order clearly conflicts with the laws of the country where our customer's data is hosted.

[84] https://aka.ms/defendingyourdataterms

[85] https://aka.ms/MS-SupCourt-data

[86] https://blogs.microsoft.com/datalaw/our-practices/

- Due in large part to important litigation brought by Microsoft, including our case challenging overbroad and indefinite secrecy orders, the law continues to support the freedom of Microsoft and other cloud providers to inform our customers about law enforcement requests for their data.
- It reinforces our existing policies that, for legitimate enterprise customers, US law enforcement will in most instances now go directly to the customer rather than to Microsoft for information requests.

The CLOUD Act creates new protections:

- The CLOUD Act creates the authority and framework for the United States to establish international agreements that, on a reciprocal basis, will enable law enforcement agencies to access data in each other's countries to investigate and prosecute crimes. The UK concluded a bilateral agreement[87] with the United States in October 2019; the European Union has yet to finalize such an agreement.
- The CLOUD Act preserves and expands the direct legal rights of cloud service providers such as Microsoft, to protect privacy under these agreements in two forms. It gives providers the right to inform foreign governments that have these agreements when their citizens are impacted by US warrants, and to go directly to court to raise comity concerns when the United States seeks a warrant that goes beyond the scope of an agreement and that conflicts with a foreign law. (Comity is the legal convention that political entities, such as courts from different jurisdictions, will reciprocally recognize each other's judicial acts.)

To protect the privacy of its business customers into the future, Microsoft complies with the following five principles and will continue to:

- Carefully evaluate every law enforcement request and exercise our rights to protect our customers where appropriate.
- Refer US authorities to the respective business customers instead of providing data to the authorities by choice.
- Go to court to defend the local rights of our customers if their rights are violated by the US government.
- Push for new international agreements that strengthen the rights of our customers.
- Be transparent about the number of international search warrants Microsoft receives.

> Get more information in the Customer Guide to Microsoft's Position on the CLOUD Act.[88]

Law enforcement requests

Microsoft is committed to transparency, and twice yearly Microsoft publishes the Law Enforcement Requests Report.[89] The report brings together in one place the reports that Microsoft issues regularly on requests for customer data made by law enforcement, as well as civil legal requests related to US national security.

The aggregate data Microsoft has published shows clearly that only a small fraction of a percent of Microsoft customers have ever been subjected to a government request. For enterprise customers, that number drops further to a mere handful. For example, in the second half of 2020:

- **For consumer services:** Microsoft received 5,682 legal demands for consumer data from law enforcement in the United States. Of those, 165 warrants sought content data which was stored outside of the United States. In the same time frame, Microsoft received 71 legal demands from law enforcement in the United States for commercial enterprise customers who purchased more than 50 seats.
- **For enterprise services:** Microsoft received 109 requests from law enforcement around the world for accounts associated with enterprise cloud customers. In 69 cases, these requests were rejected, withdrawn, had no data, or law enforcement was successfully redirected to the customer. In 40 cases, Microsoft was compelled to provide responsive information: 19 of these cases required the disclosure of some customer content and in 21 of the cases we were compelled to disclose non-content information only. Of the 19 instances that required disclosure of content data, 10 of those requests were associated with U.S. law enforcement.

> Refer to the FAQs at the bottom of the Law Enforcement Requests Report for information concerning the requests we receive from law enforcement agencies around the world.

[87] https://aka.ms/US-UK-DOJ

[88] https://aka.ms/mscloudactguidance

[89] https://aka.ms/LERR

V. How customers can protect data from unauthorized access

Encryption is fundamental to helping ensure the confidentiality of cloud workloads. Microsoft uses multiple encryption methods, protocols, and algorithms across its products and services to help provide a secure path for data to travel through the Azure infrastructure, and to help protect the confidentiality of data that is stored within the infrastructure. Microsoft uses some of the strongest, most secure encryption protocols in the industry to provide protection against unauthorized access to its customers' data.

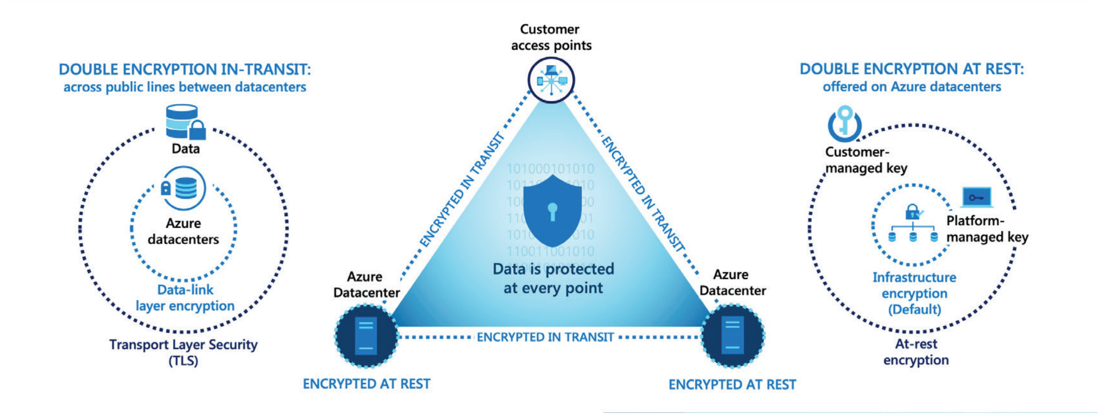
DOUBLE ENCRYPTION IN-TRANSIT:
across public lines between datacenters
Data
Azure datacenters
Data-link layer encryption
Transport Layer Security (TLS)
Customer access points
ENCRYPTED IN TRANSIT
ENCRYPTED IN TRANSIT
Data is protected at every point
Azure Datacenter
Azure Datacenter
ENCRYPTED IN TRANSIT
ENCRYPTED AT REST
ENCRYPTED AT REST
DOUBLE ENCRYPTION AT REST:
offered on Azure datacenters
Customer-managed key
Platform-managed key
Infrastructure encryption (Default)
At-rest encryption

Encryption for Azure services

Azure also offers many ways for customers to manage and control the security of customer data, including the means to encrypt data at rest, data in transit, and data during processing.

Proper key management is an essential element in encryption best practices, and Azure Key Vault helps ensure that encryption keys are properly secured. Encryption keys can be managed three ways: by Azure in Key Vault, by the customer in Key Vault, or managed and stored on premises. Key management is discussed in Key management (page 29) and includes options ranging from server-side service-managed keys to client-side encryption in which Azure services do not have access to encryption keys and cannot decrypt customer data.

> Get more information in Azure Encryption Overview[90] and in Azure data security and encryption best practices.[91]

Encryption of data at rest

Customer data at rest is automatically encrypted when it is written to Azure Storage, including Azure Managed Disks; Azure Blob, Queue, and Table Storage; and Azure Files. All data written to the Azure storage platform is encrypted through 256-bit AES encryption, which is one of the strongest block ciphers available, and is FIPS 140-2 compliant.

> Get more information on Azure Storage encryption for data at rest.[92]

To manage disk encryption for VMs and Virtual Machine Scale Sets (VMSSs), customers can configure operating system (OS) and data disks used by Azure VMs to be encrypted. Azure offers multiple options to encrypt OS and data disks for Windows Server and Linux instances. Azure Disk Encryption encrypts Windows and Linux infrastructure as a service (IaaS) VM disks by using the BitLocker feature of Windows and the DM-Crypt feature of Linux to provide full volume encryption for the operating system disk and the data disk. Azure Disk Encryption requires an Azure Key Vault to control and manage disk encryption keys and secrets. The key vault and VMs must reside in the same Azure region and subscription.

BitLocker also encrypts Shielded VMs in Windows Server 2016 to ensure that fabric administrators can't access the information inside the VM. The Shielded VMs solution includes the Host Guardian Service, which is used for virtualization host attestation and encryption key release.

> Get more information regarding encrypting Windows and Linux VM disks in Azure Disk Encryption for VMs and VM scale sets.[93]

[90] https://aka.ms/AZ-encryption-overview

[91] https://aka.ms/AZ-encryption-best

[92] https://aka.ms/AZ-Storage-Encrypt

[93] https://aka.ms/AZ-encryption-vms

Encryption technologies for specific storage types

Additional encryption technologies for specific storage types are available, including the following:

- Transparent Data Encryption (TDE)[94] encrypts data at rest when it's stored in Azure Synapse Analytics.
- The Always Encrypted[95] feature supports the ability to encrypt data within client applications before storing it in Azure SQL Database. Always Encrypted with secure enclaves expands the confidential computing capabilities of Always Encrypted by enabling in-place encryption and richer confidential queries.
- Azure Data Lake Storage (ADLS)[96] is protected by transparent encryption of data at rest similar to what is provided with Azure SQL Database. ADLS is on by default and performs key management by default, but there is an option to self-manage the keys if desired. In ADLS Gen 2, similar to Blob storage, Storage Service Encryption (SSE) automatically encrypts data at rest using Microsoft-managed keys or the customer's own encryption keys.
- Azure Cosmos DB[97] is encrypted by default, using secure key storage systems, encrypted networks, and cryptographic APIs. Microsoft manages the encryption keys, rotating them per our internal guidelines.
  > Learn more about data encryption in Azure Cosmos DB.[98]

Key management

Key management can be performed by Azure or by the customer, and encryption can be performed server-side or client-side.

- **Server-side encryption:** There are three server-side encryption models that offer different key management characteristics from which you can choose according to your organization's requirements:
  - **Service-managed keys** use Azure Key Vault to provide a combination of control and convenience with low overhead. Azure resource providers perform the encryption and decryption operations, and Microsoft manages the keys.
  - **Customer-managed keys** give you control over the keys, including bring-your-own-key (BYOK) support, or allow you to generate new ones. Azure resource providers perform the encryption and decryption operations. The customer controls keys using Azure Key Vault.
    > Learn more about configuring customer-managed keys with Azure Key Vault.[99]
  - **Customer-provided keys** (CPK) enable you to store and manage keys in on-premises or key stores other than Azure Key Vault.
- **Client-side encryption:** With client-side encryption, Microsoft does not have access to the encryption keys and cannot decrypt the data. Customers encrypt data and upload the data as an encrypted blob. The customer maintains complete control of the keys and keeps keys on premises (or in other secure stores); keys are not available to Azure services. This model is supported by Azure, but not by all Azure services.

> For more information, see Azure Encryption Overview.[100]

Key management: Server-side encryption

Service-managed keys

Azure Key Vault[101] is a cloud-hosted service that provides centralized storage and management of cryptographic keys and other secrets that are used in customers' cloud applications. This Azure service enables customers to safeguard cryptographic keys, certificates, and application passwords, and helps protect secrets from accidental leakage.

Azure Key Vault uses specialized hardware security modules (HSMs) for maximum protection and is designed in a way that allows customers to maintain control of keys and data. Usage of customers' stored keys can be monitored and audited in different ways, including Azure logging and the import of these logs into Azure HDInsight. Customers can also incorporate this information into their existing security information and event management (SIEM) systems, which supports Microsoft customers in performing additional analysis, such as threat detection.

[94] https://aka.ms/AZ-SQL-Encrypt

[95] https://aka.ms/AZ-SQL-Always

[96] https://aka.ms/AZ-DataLake

[97] https://aka.ms/AZ-Cosmos-intro

[98] https://aka.ms/AZ-Cosmo-Encrypt

[99] https://aka.ms/AZ-CMK

[100] https://aka.ms/AZ-encryption-overview

[101] https://aka.ms/AZ-keyvault-over

Azure Key Vault allows segregation of secrets in multiple vaults. This helps reduce the chances of accidental loss of security information by centralizing the storage of application secrets. Azure Key Vault can handle requests and renewals of TLS certificates. It also provides features that enable robust certificate lifecycle management. Note that Azure Key Vault is designed to support application keys and secrets and is not intended to be a store for user passwords.

Access to a key vault is controlled through two separate interfaces: the management plane and the data plane. Access controls for the management plane and data plane work independently. Customers should use dedicated role definitions in Azure Active Directory to manage role-based access. This approach implements an effective segregation of duties.

If you want an additional layer of security for encryption keys, you can add a key encryption key (KEK) to your key vault. When a key encryption key is specified, Azure Disk Encryption uses that key to wrap the data encryption key (DEK). The entity that has access to the KEK may be different than the entity that requires the DEK. The DEK is cached and accessed by the resource provider for efficient encryption as close to the data as possible. The KEK is under customer control in Key Vault. By using the Azure Backup service,[102] customers can back up and restore encrypted VMs that use the KEK configuration.

Customer-managed keys

Azure Key Vault also provides a bring-your-own-key (BYOK) capability. Customers can generate the keys on premises using an offline workstation equipped with an nCipher HSM, and then transmit the keys securely to the Azure HSMs in the cloud. The nCipher software used for key submission ensures that the keys are bound to this environment and can never be extracted out of the HSMs. Customers who require additional functions such as enterprise key management processes or hybrid cloud setups can use the CipherTrust Cloud Key Manager.

Storing the customer keys on premises eliminates the possibility for Azure to decrypt workloads, though it also limits some Azure functionality, such as search.

> Get information about how Azure supports BYOK functionality in Import HSM-protected keys to Key Vault.[103]

Customer-provided keys

These enable you to store and manage keys in on-premises or key stores other than Azure Key Vault to meet corporate, contractual, and regulatory compliance requirements for data security. Customer-provided keys (CPK) enable you to pass an encryption key as part of a read or write operation to a storage service using blob APIs.[104] Since the encryption key is defined at the object level, you can have multiple encryption keys within a storage account. When you create a blob with a customer-provided key, the storage service persists the SHA-256 hash of the encryption key with the blob to validate future requests. When you retrieve an object, you must provide the same encryption key as part of the request. For example, if a blob is created with Put Blob[105] using CPK, all subsequent write operations must provide the same encryption key. If a different key is provided or if no key is provided in the request, the operation will fail with a 400 Bad Request. As the encryption key itself is provided in the request, a secure connection must be established to transfer the key.

> Get more information about Customer Provided Keys with Azure Storage Service Encryption.[106]

Key management: Client-side encryption

The client-side encryption model refers to customer-managed on-premise encryption keys—that is, encryption that is performed outside of the resource provider or Azure. It includes data encrypted by an application that's running in the customer's datacenter or by a service application; and data that is already encrypted when Azure receives it.

In either case, when leveraging this encryption model, the Azure resource provider receives an encrypted blob of data without the ability to decrypt the data in any way or have access to the encryption keys. In this model, the key management is done by the calling service or application and is opaque to the Azure service. Azure services cannot see decrypted data, keys are not available to Azure services, and customers manage and store keys on premises (or in other secure stores).

[102] https://aka.ms/azure-backup
[103] https://aka.ms/AZ-HMS-Keys
[104] https://aka.ms/AZ-Blob-APIs
[105] https://aka.ms/AZ-Put-Blob
[106] https://aka.ms/AZ-Stor-Encrpt

Client-side encryption works only for some (not all) Azure services. Azure Blobs, Tables, and Queues, as well as Azure DevOps Services and Azure Repos, Service Bus, IoT Hub, Media Services, StorSimple, Azure Backup, and Data Box support client-side encryption.

Client-side encryption of Azure SQL Database data is supported through the Always Encrypted feature. Customers can store the master key in a Windows certificate store, Azure Key Vault, or a local HSM. Using SQL Server Management Studio, SQL users choose which key they'd like to use to encrypt which column.

At the time of this writing, client-side encryption does not work for artificial intelligence (AI) and machine learning services, analytics services, containers, compute services, identity services, management and governance services, security services, and many storage services.

> Get more information about Client-Side Encryption and Azure Key Vault for Azure Storage.[107]

Encryption of data in transit

Protecting data in transit should be an essential part of any data protection strategy. Organizations that fail to protect data in transit are more susceptible to man-in-the-middle attacks, eavesdropping, and session hijacking. Since data is moving back and forth from many locations, the general recommendation is that customers always use Secure Sockets Layer/ Transport Layer Security (SSL/TLS) protocols to exchange data across different locations.

Microsoft enables and encourages customers to encrypt customer data in transit to Azure datacenters through TLS, which uses a combination of asymmetric (TLS handshake) and symmetric (shared secret) cryptography to encrypt communications as they travel over the network. Microsoft also uses Internet Protocol Security (IPsec), an industry-standard set of protocols, to protect the authentication, integrity, and confidentiality of data at the IP packet level as the data is transferred across the network.

Investment by Microsoft in research and development has brought about a breakthrough in the encryption of data in transit. Every Azure server contains Azure SmartNICs, which are based on Field Programmable Gate Array (FPGA) technology. These FPGAs are programmable hardware modules, which significantly speed up the processing of data including encryption of data in transit. This enables high performance for all workloads, along with low latency. Microsoft publishes the hardware design under an open source license, enabling the community and its customers to benefit from this innovation.

> Read more in Azure Accelerated Networking: SmartNICs in the Public Cloud.[108]

In some circumstances, customers may want to isolate the entire communication channel between on-premises and cloud infrastructures by using a virtual private network (VPN). The following are some ways to protect data in transit in that situation:

- For data moving between on-premises infrastructure and Azure, consider appropriate safeguards such as HTTPS or VPN.
- For organizations that need to secure access from multiple workstations located on premises to Azure, use Azure Site-to-Site VPN.[109]
- For organizations that need to secure access from one workstation located on premises to Azure, use Point-to-Site VPN.[110]
- Larger data sets can be moved over a dedicated high-speed WAN link such as ExpressRoute.[111] Customers choosing to use ExpressRoute can also encrypt the data at the application-level using SSL/TLS or other protocols for added protection.
- For customer interactions with Azure Storage through the Azure Portal, all transactions occur via HTTPS. The Storage REST API[112] over HTTPS can also be used to interact with Azure Storage[113] and Azure SQL Database.[114]

[107] https://aka.ms/AZ-client-encrypt

[108] https://aka.ms/AZ-SmartNICs

[109] https://aka.ms/ite-to-site-gateway

[110] https://aka.ms/point-to-site-gateway

[111] https://aka.ms/expressroute

[112] https://aka.ms/AZ-REST-API

[113] https://azure.microsoft.com/services/storage/

[114] https://azure.microsoft.com/services/sql-database/

ExpressRoute encryption

ExpressRoute supports the following encryption technologies to ensure the confidentiality and integrity of the data traversing between your network and the Microsoft network:

- **Point-to-point encryption using MACsec**. MACsec is an IEEE standard that encrypts data at the Media Access control (MAC) level, which is Layer 2 of the OSI networking model. You can use MACsec to encrypt the physical links between your network devices and Microsoft network devices when you connect to Microsoft via ExpressRoute Direct.
- **End-to-end encryption using IPsec**. IPsec is an IETF standard that encrypts data at the Internet Protocol (IP) level, which is Layer 3 of the OSI networking model. You can use IPsec to encrypt an end-to-end connection between your on-premises network and your virtual network (VNET) on Azure.

> Read more about how these technologies work in ExpressRoute encryption.[115]

Encryption during processing of data

To encrypt customer data during processing, Azure confidential computing[116] protects it at runtime, bringing Intel SGX and Virtualization Based Security to the cloud. Confidential computing helps ensure that when data needs to be "in the clear" (unencrypted) for efficient processing, the data is protected inside a Trusted Execution Environment (TEE).

TEEs help to ensure that no one on the outside can view the data or the operations inside the TEE, even with a debugger. While data is being processed, the TEE enforces these protections against viewing and modification, including access by Microsoft personnel. This also helps to ensure that only authorized code is permitted to access data. If the code is altered or tampered with, the operations are denied, and the environment is disabled.

To use confidential computing, customers choose the DCsv-2series, which are backed by the latest generation of Intel XEON E-2288G processors with SGX technology.

> For sizing information, see DCsv2-series.[117]

Customer Lockbox for Azure

To further safeguard customer data, Microsoft has introduced the Customer Lockbox for Azure customers, which is used in cases where a Microsoft engineer needs to access customer data whether in response to a customer-initiated support ticket or a problem identified by Microsoft, except in cases of emergency or external legal demands for data. Customer Lockbox is a service that enables such customers to control how a Microsoft engineer can access the customer's content stored in an Azure service in those rare instances when it's necessary. As part of this support workflow, a Microsoft engineer may require elevated access to customer content. Customer Lockbox puts the customer in charge by enabling them to review and approve or deny such elevated Microsoft requests to access customer data.

Customer Lockbox is an extension of the Just-in-time (JIT) workflow and also comes with full audit logging enabled. Customers can access the logs related to this service via the Azure portal and integrate them into their SIEM systems.

> Read more about how the Customer Lockbox for Azure[118] works. For details on external legal demands for data, see the Microsoft Law Enforcement Requests Report.

[115] https://aka.ms/AZ-expressroute-Enc

[116] https://aka.ms/confidential-compute

[117] https://aka.ms/AZ-dcv2-series

[118] https://aka.ms/msazurelockbox

VI. Data retention and deletion

Data retention and deletion policies and practices are important for protecting data. Privacy compliance requires following fundamental principles for data retention. Microsoft follows strict guidelines for retaining and deleting data.

In the Online Services Data Protection Addendum,[119] Microsoft contractually commits to specific processes when a customer terminates an online service or its subscription expires. This includes deleting customer data from systems under Microsoft control.

Data retention

If a customer terminates a cloud subscription or it expires (with the exception of free trials), Microsoft will store the customer's data in a limited-function account for 90 days (the retention period—a safety net, in effect) to enable customers to extract their data or renew their subscription. During this period, Microsoft provides multiple notices so the customer will be amply forewarned of the upcoming deletion of their data from Microsoft systems. Note, too, that at all times during a subscription, customers can access, extract, and delete customer data stored in our online services.

After this 90-day retention period, Microsoft will disable the account and delete the customer data, including any cached or backup copies. For in-scope services, that deletion will occur within 90 days after the end of the retention period. (In-scope services are defined in the Data Processing Terms section of the Microsoft Online Services Terms.[120]

For Cognitive Services, a configuration or custom model that has been inactive may, for the purposes of data retention and deletion, at Microsoft discretion, be treated as an online service for which the customer's subscription has expired. A configuration or custom model is inactive if for 90 days: (1) no calls are made to it; (2) it has not been modified and does not have a current key assigned to it and; (3) the customer has not signed in to it.

> For information in support of creating an exit plan for Microsoft cloud projects, see Exit Planning for Microsoft Cloud Services.

Data deletion

Microsoft uses different types of data deletion techniques depending on the type of data object that is being deleted—whole subscriptions, Azure Storage, Azure VMs, Azure SQL Database, or Azure Active Directory.

- **Subscriptions**. As noted above, when a subscription is canceled or terminated, Microsoft retains customer data for 90 days to permit the customer to extract its data. Microsoft will then delete all customer data within another 90 days after the retention period (i.e., by day 180 after cancellation or termination). If a storage account is deleted within an existing subscription (or when a subscription deletion has reached its timeout), the storage account is not actually deleted for two weeks; this is to allow recovery from mistakes. When a storage account is finally deleted, or when blob or table data is deleted outside the context of a storage account deletion, the data is no longer available. To make storage data unrecoverable faster, customers should delete tables and blobs individually before deleting the storage account or canceling a subscription.
- **Azure Storage**. All disk writes are sequential in Azure Storage. This minimizes the number of disk "seeks," but requires updating the pointers to objects every time they are written. (New versions of pointers are also written sequentially.) A side effect of this design is that if there is a secret on disk, you can't ensure it is gone by overwriting with other data. The original data will remain on the disk and the new value will be written sequentially. Pointers will be updated such that there is no longer any way to find the deleted value.

  When the disk is full, the system has to write new logs onto disk space that has been freed up by the deletion of old data. Instead of allocating log files directly from disk sectors, log files are created in a file system running the New Technology File System (NTFS). A background thread running on Azure Storage nodes frees up space by going through the oldest log file, copying blocks that are still referenced from that oldest log file to the current log file (and updating all pointers as it goes). It then deletes the oldest log file. Thus, there

[119] https://aka.ms/MS-DPA

[120] https://aka.ms/Online-Services-Terms

are two categories of free space on the disk: 1) space that NTFS knows is free, where it allocates new log files from this pool and 2) space within those log files that Azure Storage knows is free because there are no current pointers to it.

Customers can access only virtual disks and are never provided with access to the underlying physical storage, so other customers and Microsoft personnel cannot read a customer's deleted data.

- **Azure VMs** are stored in Azure Storage as blobs, and the deletion rules apply as explained in "Subscriptions" above. The virtualization mechanism is designed to ensure that those spots on the disk cannot be read by another customer (or even by the same customer) until data is written again. This mitigates the threat of data leakage. When a new virtual disk is created for a VM, it will appear to the VM to be zeroed; however, the explicit zeroing of the data buffers occurs when a portion of the virtual disk is read before it is written. If a VM instance is reinitialized in place, it's the same as if it had been moved to new hardware.
- **Azure SQL Database**. With Azure SQL Database, deleted data is marked for deletion. If an entire database is deleted, it is the equivalent of deleting the database's entire contents. The SQL Database implementation is designed to ensure that user data is never leaked by disallowing all access to the underlying storage except via the SQL Database API. That API allows users to read, write, and delete data, but does not have a way to express the reading of data that the user has not previously written.
- **Azure Active Directory (AAD)**. When an administrator or the AAD service deletes a user object, it is first moved into the recycle bin where the data remains intact. This gives a customer the ability to easily recover user objects if they are accidentally deleted. After 30 days, the user object becomes a deleted object where the attribute data is removed from AAD, except for the subset of unique identifier data that is required for replicating deletions among Microsoft datacenters. At this point, no personal data remains. After another 30 days, the user object is removed from the AAD scale unit.

Data disk destruction

If a disk drive used for storage suffers a hardware failure or reaches its end of life, it is securely erased or destroyed. The data on the drive is completely overwritten to ensure the data cannot be recovered by any means. When such devices are decommissioned, they are shredded and destroyed in line with NIST SP 800-88 R1 Guidelines for Media Sanitization.[121] Records of the destruction are retained and reviewed as part of the Microsoft audit and compliance process. All Azure services use approved media storage and disposal management services.

[121] https://aka.ms/NIST-800-88

VII. Data residency and privacy compliance

Compliance with the requirements of formal technical standards, regulations, and laws plays a critical role in providing assurance for customers that the privacy of customer data is respected and protected while giving customers support to help them meet their own compliance commitments.

Azure possesses a broad, industry-leading portfolio of compliance offerings with stringent and widely recognized formal standards. This provides the evidence customers need to assert their own compliance with government regulations and organizational policies. This evidence also helps customers maintain a record of the measures that Microsoft (the data processor) has implemented to mitigate potential risks. Rigorous audits (many of which require annual review of Azure facilities and capabilities) are conducted by independent accredited third parties such as BSI and Deloitte, which validate Azure adherence to these standards.

Azure compliance offerings include those that apply globally and to the US government. In addition, customers can take advantage of more than 20 region-specific and country-specific offerings—from Argentina and China to the UAE and United Kingdom—and over 35 offerings specific to the needs of such key industries as finance, healthcare, and manufacturing. Compliance offerings are based on various types of assurances, including formal certifications, attestations, validations, authorizations, and assessments produced by independent auditing firms, as well as contractual amendments, self-assessments, and customer guidance documents produced by Microsoft.

> Learn more about Azure Compliance Offerings[122] and get a complete listing in Azure compliance documentation[123].

Data privacy compliance

The Azure compliance portfolio includes conformity to such cloud privacy laws and regulations as the General Data Protection Regulation, Standard Contractual Clauses, ISO/IEC 27018, and ISO/IEC 27701.

> Get an overview of Privacy in Azure.[124]

General Data Protection Regulation

The GDPR is a European privacy law that applies to the European Economic Area (EEA), which includes all EU countries plus Iceland, Liechtenstein, and Norway. The GDPR imposes new rules on organizations that offer goods and services to people in the EEA or that collect and analyze data belonging to EEA individuals. The GDPR requires that data controllers (such as organizations using Azure) use only data processors (such as Microsoft) that provide sufficient guarantees to meet key GDPR requirements. Microsoft contractually commits to meet GDPR requirements not only in the EU but in all public cloud regions, and all Azure services can be used in compliance with the GDPR. Microsoft offers extensive documentation and tools[125] to help customers comply with GDPR requirements.

Standard Contractual Clauses

Most countries allow data transfers outside of their boundaries, though often with certain restrictions. The GDPR also regulates transfers of the personal data of European residents to destinations outside the EEA. These transfers require a specific legal mechanism, such as a contract, or adherence to a certification mechanism to enable them. If customers using Azure services choose to transfer content containing personal data across borders, they will need to consider the legal requirements that apply to such transfers.

In July 2020, the Court of Justice for the European Union invalidated the Privacy Shield Frameworks for transfers of personal data from the EU to the United States. However, the Standard Contractual Clauses (also known as EU Model Clauses) continue to provide a lawful mechanism for the transfer of all customer and personal data out of the EU, the European Economic Area (EEA), Switzerland, and the United Kingdom.

[122] https://aka.ms/AzureCompliance

[123] https://docs.microsoft.com/azure/compliance/

[124] https://aka.ms/AZ-Privacy

[125] https://aka.ms/AZ-GDPR

While Microsoft minimizes transfers of data from the geography in which customers choose to store their data, when such transfers are necessary, Microsoft relies on its longstanding use of the Standard Contractual Clauses that make specific guarantees around transfers of personal data for in-scope Azure services. We have updated our contractual commitment, the Data Protection Addendum for Online Services[126] (in Attachment 2), to reflect that transfers of personal data are now governed by the Standard Contractual Clauses (controller to processor).

Further, Attachment 3 provides additional safeguards to our customers and additional redress to customers' data subjects: Microsoft commits that we will challenge every government request for customer data where there is a lawful basis for doing so, and we will provide monetary compensation to customers' users if we disclose their data in response to a government request in violation of the GDPR. The scope of the Data Protection Addendum has also expanded to fully cover data transfers for all of Microsoft online services (rather than simply Core Online Services).

ISO/IEC 27018 Code of Practice for Protecting Personal Data in the Cloud

Azure complies with ISO/IEC 27018,[127] which specifically targets controls that apply to processing data in compliance with GDPR Article 9 requirements. At least once a year, Azure is audited for its compliance with ISO/IEC 27018 by an accredited third-party certification body, providing independent validation that applicable security controls are in place and operating effectively.

ISO/IEC 27701 Privacy Information Management System (PIMS)

Azure is also certified for the ISO/IEC 27701 standard,[128] an extension of the widely-used ISO/IEC 27001 standard[129] for information security management. This makes the implementation of a PIMS a helpful compliance continuation of the many organizations that rely on ISO/IEC 27001, as well as creating a strong integration point for aligning security and privacy controls. ISO/IEC 27701 accomplishes this integration through a framework for managing personal data that can be used by both data controllers and data processors, a key distinction for GDPR compliance.

Sharing responsibility for compliance

Microsoft complies with all laws and regulations applicable to providing its online services, including laws governing notification of security breaches. However, Microsoft is not responsible for compliance with any laws or regulations that apply to the customer or the customer's industry that are not generally applicable to information technology service providers. Microsoft does not determine whether customer data includes information subject to any specific law or regulation.

And while it is up to you to determine whether Azure services comply with the specific laws and regulations that are applicable to your business, Microsoft can help you make these assessments by providing the specifics of our compliance programs, including audit reports and certificates. Your auditors can compare Azure results with your own legal and regulatory requirements, and you can verify the Azure implementation of controls. Azure customers (and trial customers) can access these in the Azure Portal Audit Reports.[130] Available documents include:

- Cyber Essentials Plus (UK)
- ENS (Spain)
- FedRAMP
- GSMA (France)
- HDS (France)
- HITRUST
- IRAP (Australia)
- ISO 27001, 27018, 27017, and 27701 (PIMS) frameworks
- MTCS (Singapore)
- NIST 800-53
- OSPAR (Singapore)
- PCI DSS and PCI 3DS
- SOC 1, 2, and 3 (Including bridge letters)

> Shared Responsibility for Cloud Computing[131] explains the roles and responsibilities that cloud service providers and customers share in cloud computing.

[126] https://aka.ms/MS-DPA
[127] https://aka.ms/AZ-ISO27018
[128] https://aka.ms/AZ-ISO27701
[129] https://aka.ms/AZ-ISO27001
[130] https://aka.ms/AZ-audit-reports
[131] https://aka.ms/AZ-Shared

Microsoft contractual commitments

When customers subscribe to an online service through a Microsoft Volume Licensing program, we contractually guarantee our commitments in our standard contracts for commercial and public sector customers. The terms that control how customers can use the service are defined in the Microsoft Online Services Terms[132] (OST) and the Online Services Data Protection Addendum[133] (DPA), both of which are available in 34 languages. Due to the frequency with which Microsoft adds new services, the OST is updated monthly and the DPA is updated as needed. Additional amendments exist to cover restricted industries, including financial services, and will be available to customers where applicable. An archive contains older versions for reference.

> Learn more about all the contractual commitments Microsoft makes in Licensing Terms,[134] and get the detailed terms that describe Microsoft commitments for Azure uptime and connectivity in our service-level agreements.[135]

[132] https://aka.ms/Online-Services-Terms

[133] https://aka.ms/MS-DPA

[134] https://aka.ms/MS-product-licensing

[135] https://azure.microsoft.com/support/legal/sla/

Appendix: Selected resources

Compliance guidance

Microsoft currently hosts datacenters in over 60 regions across multiple countries. Microsoft provides compliance guidance for customers in a series of documents that address data residency requirements generally, along with special emphasis on the financial services and healthcare sectors.

- Navigating your way to the cloud in Europe:[136] Belgium, Bulgaria, Croatia, Czech Republic, Estonia, Finland, France, Germany, Italy, Ireland, Latvia, Lithuania, Luxembourg, Netherlands, Norway, Poland, Slovenia, Spain, Sweden, Switzerland, United Kingdom
- Navigating your way to the cloud in Asia: A Guide for the Legal & Compliance Professional:[137] Australia, Bangladesh, Brunei, Hong Kong, India, Indonesia, Japan, Korea, Malaysia, Nepal, New Zealand, Philippines, Singapore, Sri Lanka, Thailand, Vietnam
- Navigating your way to the cloud in the Middle East and Africa: A Guide for Legal and Compliance Professionals:[138] Angola, Jordan, Kenya, Mauritius, Morocco, Nigeria, Rwanda, South Africa, UAE

Microsoft Trust Center

The Microsoft Trust Center[139] is a central location for information about security, privacy, compliance, and transparency for Microsoft products and services, including data residency information. There's also guidance on government and industry-specific compliance, audit reports, security assessments, and more.

Recommended data residency and security white papers

- Azure for Secure Worldwide Public Sector Cloud Adoption[140]
  This paper addresses common data residency and security concerns pertinent to public sector customers around the world. It also explores technologies available in Azure to safeguard both unclassified and classified workloads in the public multitenant cloud in combination with Azure Stack and Data Box Edge deployed on premises and at the edge.
- Protecting Data Privacy using Microsoft Azure[141]
  This paper discusses the Azure tools and services that organizations can use and the steps they can take to protect both customer data and personal data.

[136] https://aka.ms/TrustedCloud-Europe

[137] https://aka.ms/TrustedCloud-APAC

[138] https://aka.ms/TrustedCloud-MEA

[139] https://www.microsoft.com/trust-center

[140] https://aka.ms/AzureWWPS

[141] https://aka.ms/AZ-Data-Privacy

Microsoft