


Enabling Data Residency and Data Protection in Microsoft **Azure** Regions



Authors

Christoph Siegert

Debra Shinder

David Burt

Reviewers and Contributors

Karina Juarez

Andres Juarez

Fotini Kaklamanou

Stevan Vidich

Deepali Bhardwaj

Ralf Wigand

Peter Koen

Adrian Maziak

Contents

| | |
|--|----|
| Introduction..... | 4 |
| I. Azure Region infrastructure | 5 |
| High availability | 7 |
| High availability in regions with Availability Zones | 7 |
| High availability in regions without Availability Zones | 7 |
| Maintenance downtime | 7 |
| Disaster recovery | 8 |
| Disaster recovery in a geography with paired regions | 8 |
| Disaster recovery in a geography with a single region | 9 |
| Latency considerations | 11 |
| Latency from customer VMs to Azure VMs | 12 |
| Latency from Azure Region to Azure Region | 13 |
| Regional service availability | 13 |
| II. Data residency for customer data | 14 |
| Data residency for regional services..... | 14 |
| Data residency for regional services in region without second region in country | 15 |
| Data storage for non-regional services. | 15 |
| Data storage for Azure Security Center | 16 |
| Data storage for Cognitive Services | 16 |
| Data storage for Azure Active Directory identity data | 16 |
| Use Azure Policy to control data residence | 17 |
| Use Azure Blueprints to enforce compliance and data residency | 17 |
| III. Access to telemetry data | 18 |
| Billing telemetry..... | 18 |
| Service health telemetry..... | 18 |
| Transparency and customer control of service health telemetry | 18 |
| Customer access to telemetry data | 19 |
| Support telemetry..... | 19 |
| Diagnostics data | 20 |
| Memory dump | 20 |
| Enabling boot diagnostics | 20 |
| Support data retention | 20 |
| Elevated DevOps data access for support cases | 20 |
| IV. How Microsoft protects data from unauthorized access..... | 21 |
| Who can access your data and on what terms? | 21 |
| How Microsoft manages your data | 21 |
| Tenant separation | 22 |
| When Microsoft deletes your data | 22 |
| Implementation of policies | 22 |
| Subcontractors and subprocessors | 22 |
| How Microsoft handles government requests | 23 |
| The CLOUD Act | 23 |
| Law Enforcement Requests Reports | 24 |
| V. How customers can protect data from unauthorized access | 25 |
| Encryption of data at rest | 25 |
| Key management | 26 |
| Key management – Server-side encryption | 27 |
| Key management – Client-side encryption (Customer-managed on-premises encryption keys) | 27 |
| Encryption of data in transit | 28 |
| ExpressRoute encryption | 29 |
| Encryption during processing of data (confidential computing) | 29 |
| Customer Lockbox for Azure | 29 |
| VI. Data retention and deletion..... | 30 |
| Data retention..... | 30 |
| Data deletion..... | 30 |
| Data disk destruction..... | 31 |
| VII. Compliance..... | 31 |
| Conclusion | 33 |
| Appendix A: Selected resources | 33 |

Introduction

Azure is available in over 140 countries, and offers customers more than 60 datacenter regions worldwide from which to choose. These Azure datacenter regions provide customers with data residency and latency optimization, and may enable regional compliance. Azure's regions unlock cloud adoption, particularly for restricted and regulated industries and for latency-sensitive applications.

Customers in countries that don't have a local Azure region can still use Azure. If there is a regulatory requirement to keep data in country, there are multiple paths that a customer can take to achieve benefits from Azure:

- Review whether data can be categorized and whether selected sets of data categories can be processed outside of the country (e.g. unclassified customer content data).
- Review whether the encryption options outlined in this paper can lead to regulatory approval to process data outside of the country.
- Review whether the regulator allows storing secondary data copy outside of country as long as primary data copy stays inside of country.

Selecting Azure datacenter regions for workload deployment should be based on technical and regulatory requirements. To optimize 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 the global deployment.

When it comes to compliance considerations, data residency regulations govern in which physical locations data can be stored and how and when it can be transferred internationally. These regulations differ depending on jurisdiction. Azure regions and service features provide customers with choices so they can select and limit data residency and data access. 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 necessary information and available tools to optimize data residency and data access using Azure's datacenter regions. Specifically, it addresses the following:

- The Azure regional infrastructure, including high availability, disaster recovery, latency, and service availability considerations
- Data residency assurances and how customers can control data residency
- Data access to telemetry data, including elevated access for support data, and how customers can manage data access
- How Microsoft protects customer data from unauthorized access, and how Microsoft handles government requests
- Tools customers can use to protect against unauthorized and support authorized data access
- Data retention and deletion

The paper is structured into these sections, with each describing how Microsoft helps customers to meet data residency and data protection requirements. Additional resources are available in the appendix.

I. Azure Region infrastructure

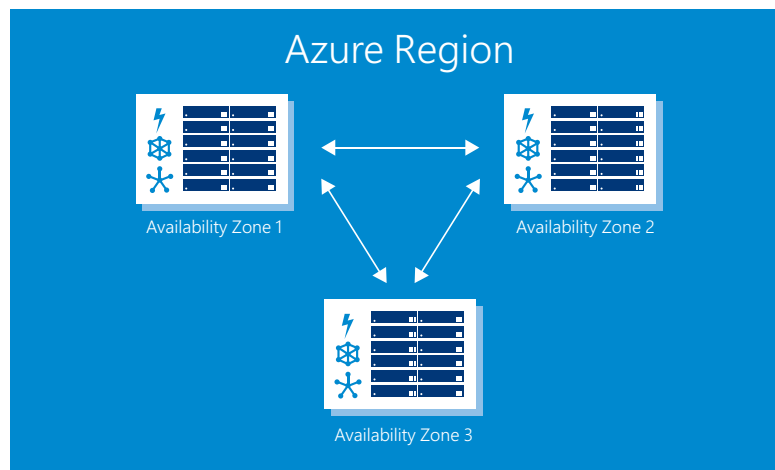
The global infrastructure of Azure enables you to deliver services and reach customers and partners wherever they are, while meeting data residency requirements. Azure allows customers to choose the location of their data. Azure's location taxonomy is focused on Availability Zones, Regions, and Geographies. These are defined and described as follows:

- **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 protects against unplanned downtime as well as 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 Agreement (SLA).

- **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 consist of several datacenters even if the region does not have multiple 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 geographical 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.
 - Germany North and Germany West Central regions are in the "Germany" geography.
 - North Europe and West Europe regions are in the "Europe" geography.

Many Azure regions are paired with another region within the same geography; together they make a regional pair in a geography. These regions are typically hundreds of kilometers apart within the same geography. This pairing provides geographic separation within a country, and provides long-distance disaster recovery (DR) within a country. Additional details are available in [Business continuity and disaster recovery \(BCDR\): Azure Paired Regions](#)

Azure also has geographies with single regions; for example, Qatar Central is a single region in the Qatar geography. The Qatar geography still provides data residency, although disaster recovery has to be managed either within the single region, or as out-of-country DR (through connecting to other Azure regions).

Because the geography determines the data residency boundary, it is important to understand the geo of each region. The full list of Azure geographies, including which region maps to which geography, is shown below and outlined on the [Azure geographies](#) page.

| Azure Regions | | Azure geography and data residency boundary |
|------------------------|---|---|
| Americas | East US - West US East US 2 - Central US North Central US - South Central US West US 2 - West Central US | USA |
| Americas | Canada Central - Canada East | Canada |
| Americas | Brazil South | Brazil |
| Americas | Mexico Central | Mexico |
| Europe | North Europe - West Europe | Europe |
| Europe | France Central - France South | France |
| Europe | UK South - UK West | United Kingdom |
| Europe | Germany West Central - Germany North | Germany |
| Europe | Switzerland North - Switzerland West | Switzerland |
| Europe | Norway East - Norway West | Norway |
| Europe | Spain Central | Spain |
| Europe | Italy North | Italy |
| Europe | Poland Central | Poland |
| Asia-Pacific | East Asia - Southeast Asia | Asia-Pacific |
| Asia-Pacific | Australia East - Australia Southeast Australia Central - Australia Central 2 | Australia |
| Asia-Pacific | China North - China East China North 2 - China East 2 | China (dedicated sovereign cloud with special data residency) |
| Asia-Pacific | Central India - South India West India - South India | India |
| Asia-Pacific | Japan East - Japan West | Japan |
| Asia-Pacific | Korea Central - Korea South | Korea |
| Asia-Pacific | New Zealand North | New Zealand |
| Middle East and Africa | South Africa North - South Africa West | South Africa |
| Middle East and Africa | Israel Central | Israel |
| Middle East and Africa | UAE Central - UAE North | United Arab Emirates |
| Middle East and Africa | Qatar Central | Qatar |

The Azure regional concept allows customers to achieve two aspects of business continuity, while keeping the customer in control of data residency:

- High availability (e.g. via AZ's 99.99% VM uptime SLAs)
- 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)

These are explained in more detail in the following sections.

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 describe Microsoft commitments for uptime and connectivity. Availability Zones, as described in the above section, provide the highest uptime availability SLAs. The full list of Azure SLAs is available on the [Service Level Agreements](#) page.

High availability in regions with Availability Zones

VMs in Availability Zones are synchronously replicated across the Availability Zones. 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 [FAQ](#) and in [Azure premium storage: design for high performance](#). This also ensures that the placement of disks for virtual machines within an Availability Set honors fault domain semantics, as documented in [Availability options for virtual machines in Azure](#). Managed Disks provide redundancy within an Availability Zone, with three replica instances spread across storage stamps in the same datacenter and support designing for a recovery point objective of zero hours within an Availability Zone for resiliency and high availability purposes.

Zone-Redundant Storage (ZRS) is a service that replicates Azure Storage data synchronously across three Azure availability zones within the same region. ZRS offers durability for Azure Storage data objects of at least 99.999999999% (12 9's) 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 synchronously kept in synch, allowing customers to design for an RPO of 0. Please refer to [Zone-Redundant Azure Blob Storage](#) for further details.

High availability in regions without Availability Zones

In regions without availability zones, Availability Sets 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 downtime

There are three high-level scenarios that can impact Virtual Machines in Azure: planned maintenance, unplanned hardware maintenance, unexpected downtime.

Planned maintenance events are periodic updates made by Microsoft to the underlying Azure platform to improve the overall reliability, performance, and security of the platform infrastructure on which virtual machines run. This includes applying security patches or bug fixes to the hosting environment or upgrading and decommissioning hardware. More detailed information about maintenance updates for VMs is available in [Maintenance for virtual machines in Azure](#).

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 impact to virtual machines hosted on affected hardware. Azure uses [Live Migration](#) technology to migrate the Virtual Machines from the failing hardware to a healthy physical machine.

Unexpected downtime occurs when the hardware or the physical infrastructure for the virtual machine fails unexpectedly. When this is detected, the Azure platform automatically migrates (heals) customer virtual machines to a healthy physical machine in the same datacenter. During the healing procedure, virtual machines experience downtime (reboot) and in some cases loss of the temporary drive; however, the attached OS operating system and data disks are always preserved.

Disaster recovery

All Azure Regions are built for hyperscale production workloads. Azure services are designed with redundancy to sustain 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 comply with industry standards for availability and are designed to run 24x7x365, 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 24x7 uninterruptible power supplies (UPS) and emergency power support, which includes on-site 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 the 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 the business continuity plans are updated accordingly.

Microsoft Azure provides customers with tools that can be used to design highly available services, employing features such as load balancing, Azure Paired Regions, Azure Backup, AzCopy, Azure Storage replication, and more. Systems are proactively monitored to ensure service performance and availability is achieved in accordance with financially backed service level agreements.

Audit Reports of the Azure platform operations detailing the processes and procedures to ensure proper operation of the Azure platform are available on the [Audit Reports page](#) of the Azure Trust Center Portal.

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.

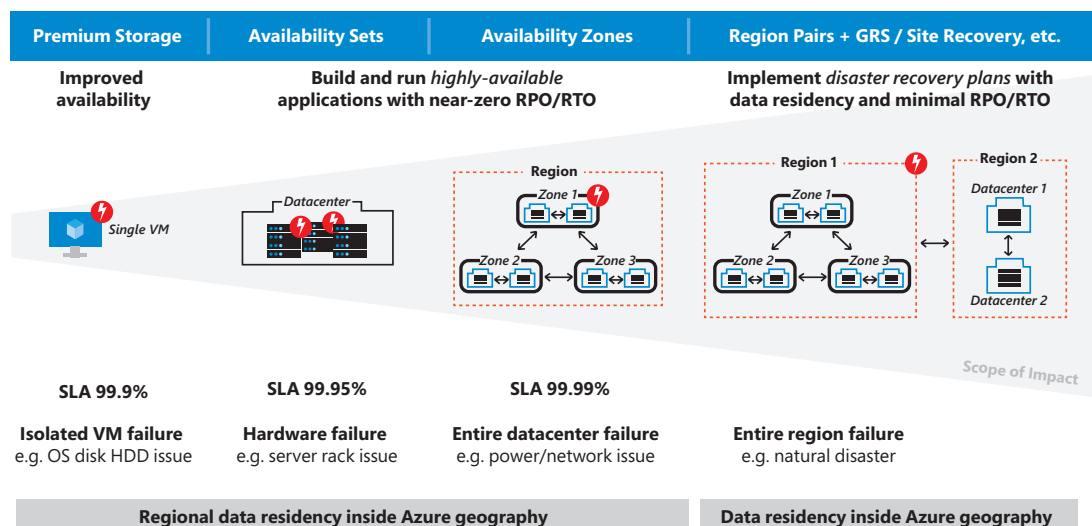
For virtual machine workloads, [Azure Site Recovery](#) provides an Azure native replication approach from a primary region to a secondary region. When an outage occurs at the primary region, the customer can trigger a failover to the secondary region, with the ability to fail back when the primary region 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 offers to 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) (in preview)** 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 are available, including AzCopy and Backup service, which are described in the following section. The graphic below illustrates the characteristics of the solutions discussed here.



Not all Azure services automatically replicate data, nor do all Azure services automatically fall back from a failed region to its pair. In such cases, recovery and replication must be configured by the customer.

You can find a listing of Azure regions pairing within the same geography in [Business continuity and disaster recovery \(BCDR\): Azure Paired Regions](#).

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 permitted from the regulatory side, customers can use any other Azure regions for DR and store secondary data set outside of the primary country with the services outlined in this section.

If the regulator doesn't allow even secondary data copies to leave the country, customers should explore whether AZs offer sufficient disaster recovery and business continuity. Azure offers data redundancy through Availability Zones by using zonal deployments (ZRS for storage, zonal VMs, zone-to-zone DR, etc.).

Note that customers can't define their own regional pairs to take advantage of fully managed geo-replicated services for all services. Customers can also 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: Site Recovery** helps ensure business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary region to a secondary region. When an outage occurs at your primary region, you fail over to the secondary location and access apps from there. The service is application-agnostic, allowing 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 running again, you can fail back to it.

You can find a description of how to enable zone to zone disaster recovery in [Enable Zone to Zone Disaster Recovery for Azure virtual machines](#).

Additionally, customers can validate their disaster recovery strategies as described in [Create and customize recovery plans](#).

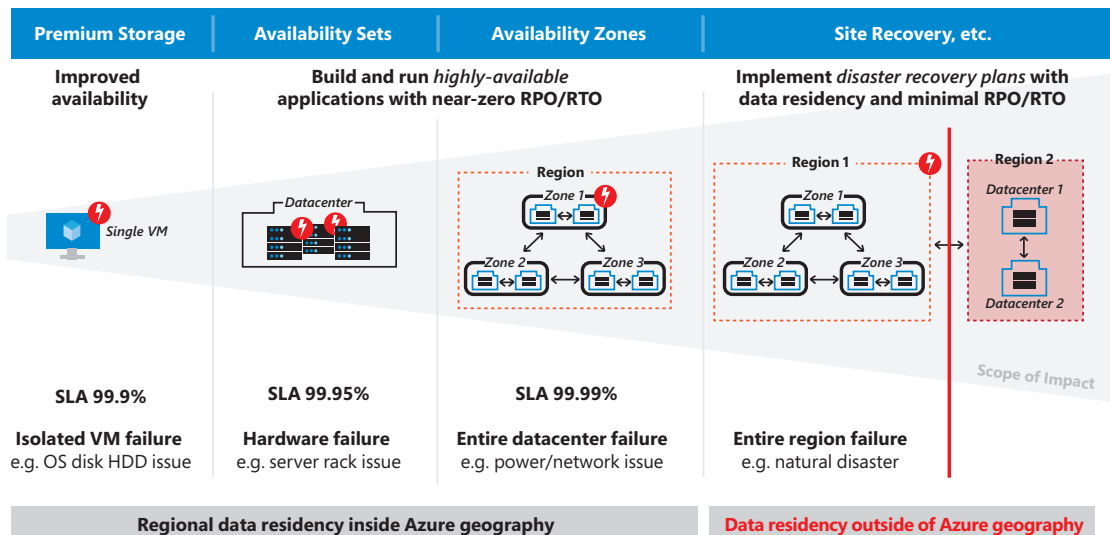
- **Multi-VM consistency** is a capability provided by Azure Site Recovery, which creates a replication group of all the machines. All of the machines in a replication group have shared crash-consistent and app-consistent recovery points when failed over. Enabling multi-VM consistency can impact workload performance as it is CPU intensive. The maximum number of virtual machines 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 cord was pulled from the server 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, and this setting cannot be modified.
 - 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 in progress transactions. 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 (Site Recovery) to read the snapshot data and proceed.

These are more complex and take longer to complete than crash-consistent snapshots and affect the performance of applications running on a virtual machine 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. Site Recovery keeps recovery points for twenty-four hours by default, but this can be configured to be a value between one and seventy-two hours.

- **Azure Backup 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 [Azure Recovery Services \(MARS\) agent](#). Backup is also available for Azure Files shares, SQL Server in Azure VMs, and SAP HANA databases running on Azure VMs.

Additional disaster recovery services include [Azure DNS](#) and [Azure Traffic Manager](#), which allow customers to design a resilient architecture that will survive the loss of the primary region, or [AzCopy](#) to schedule data backups to a Storage account in a different region.

If the DR plan uses a secondary region outside of the primary region's geography, data will be moved outside of the geography or country. The graphic below illustrates how this works.



As Microsoft provides service level SLAs, and recovery time objective (RTO) and recovery point objective (RPO) are dependent on customer architecture, Azure does not provide RTO and RPO SLAs.

At the time of this writing, Azure regions without a DR region in the same geography are the following:

- Israel Central
- Italy North
- Mexico Central
- New Zealand North
- Poland Central
- Qatar Central
- Spain Central

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 customers 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. Azure does offer 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.

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 more reliability, faster speeds, consistent latencies, 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 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 is designed to improve the data path performance between your on-premises network and your virtual network. It is available on all ExpressRoute circuits. When enabled, FastPath sends network traffic directly to virtual machines in the virtual network, bypassing the gateway. Read more in [About ExpressRoute FastPath](#).

Accelerated networking 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.

Read more about accelerated networking in [Create a Windows VM with accelerated networking using Azure PowerShell](#).

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.

While Azure does not publish a latency SLA, Ventsislav Ivanov, Development Architect for SAP, participated in the early preview program of proximity placement groups and had this to say:

"It is really great to see this feature now publicly available. We are going to make use of it in our standard deployments. My team is automating large scale deployments of SAP landscapes. To ensure best performance of the systems it is essential to ensure low-latency between the different components of the system. Especially critical is the communication between Application server and the database, as well as the latency between HANA VMs when synchronous replication has to be enabled. In the late 2018 we did some measurements in various Azure regions and found out that sometimes the latency was not as expected and not in the optimal range. While discussing this with Microsoft, we were offered to join the early preview and evaluate the Proximity Placement Groups (PPG) feature. During our evaluation we were able to bring down the latency to less than 0.3 ms between all system components, which is more than sufficient to ensure great system performance. Best deterministic results we achieved when PPGs were combined with Network acceleration of VM NICs, which additionally improved the measured latencies."

Read more in [Introducing proximity placement groups](#).

To test Azure VM network latency, see [Test VM network latency](#).

Latency from Azure region to Azure region

Azure continuously monitors the latency (speed) of core areas of its network, using internal monitoring tools as well as measurements collected by a third-party synthetic monitoring service (ThousandEyes). You can read more in [Azure network round trip latency statistics](#).

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 is a region that provides the broadest range of service capabilities and is designed to support Availability Zones now or in the future.
- An alternate (other) region is a region that extends Azure’s 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. Azure’s general 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 region or service general availability; 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 of services into these categories is available in [Regions and Availability Zones in Azure](#).

The list of services by region is available in [Products available by region](#). Note that 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.

II. Data residency for customer data

Data residency for regional services

Most Azure services are deployed regionally and enable the customer to specify the region into which the service will be deployed, and therefore control where the customer data will be stored. For these services, customers preselect the region in which the service will be deployed. For a complete list of regional services, see [Services by Region](#). For regional services, deployment location (and thereby data residency) can be defined by the region variable in the Azure portal or via the CLI. The graphic below shows how to define the region in the portal:

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 ⓘ

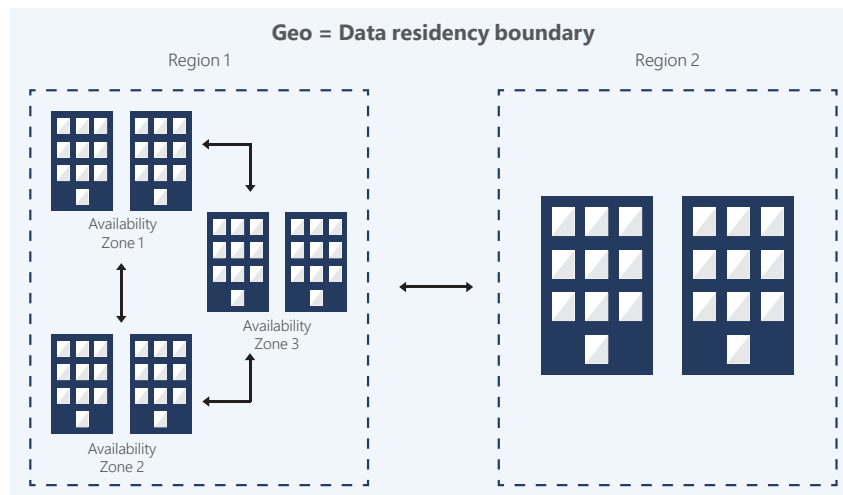
☐ Yes ☒ No

Size * ⓘ

Standard D2s v3
2 vcpus, 8 GB memory (\$70.08/month)
[Change size](#)

As mentioned in the previous section, 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, France Central and France South are the regional pair in the France geo.

The service may replicate customer data stored in that service to other regions in that geo for data resiliency, but for regional services, Microsoft will not replicate or move customer data outside the geo.



For example, if a customer deploys Azure Blob Storage in Germany West Central, the customer data may be replicated to Germany North for disaster recovery purposes, but will remain stored inside Germany. Customers and their end users may still move, copy, or access customer data from any location globally. The Microsoft Trust Center outlines where customer

data is stored and notes some limited exceptions to this rule on the [Azure datacenter map](#). Examples of exceptions to the data residency for regional services include:

- **Language Understanding** may store active learning data in the United States, Europe, or Australia based on the authoring regions that the customer uses.
- **Azure Machine Learning service** may store free-form text that the customer provides (e.g. names for workspaces, resource groups, experiments, files, and images) and experiment parameters in the United States.
- **Azure Sentinel** generates new security data such as incidents, alert rules, and bookmarks, that may themselves contain customer data from the customer's instances of Azure Log Analytics. Such security data will be stored at rest in Europe (for security data generated from the customer's Log Analytics workspaces located in Europe) or in the United States (for security data generated from the customer's Log Analytics workspaces located elsewhere).
- **Preview, beta, or other prerelease services** typically store customer data in the United States but may store it globally.

Data residency for regional services – in region without second region in country

There are selected countries where Azure offers only one region; an example is Qatar Central.

As with all Azure regions, control over residency of the customer data resides with the customer, even in Azure regions without a second region in country. For applications and services deployed into such a region, Azure's data residency promise for regional services holds.

In regions without a second region in country, that primary region will be the Azure geography and thus the data residency boundary. Customers can configure certain Azure services, tiers, or plans to store customer data only in a single region. These include Locally-Redundant Storage (LRS), Zone Redundant Storage (ZRS), Azure App Service Environment (AZ enabled ILB ASEs), Azure Backup, Azure Bastion, Azure Cache for Redis, Azure Data Explorer (ADX), Azure Data Factory, Azure Data Lake Storage Gen2, Azure Event Hubs, Azure Functions (Durable Functions), Azure HDInsight, Azure Kubernetes Service (AKS), Azure Monitor (Application Insights and Log Analytics), Azure Red Hat Openshift, Azure Service Bus (Premium), Azure Service Fabric, Azure Site Recovery, Azure SQL Database for MySQL, Azure SQL Database for PostgreSQL, and Azure Stream Analytics.

For these services, Microsoft will not replicate or move customer data outside the primary region.

A paired region inside the geography for in-country DR (beyond the single region with Availability Zones), will not be available. Customers deploying to a single-region location should consider that Azure Site Recovery (ASR) provides an application-agnostic replication technology that enables customers to build disaster recovery for any application hosted on VMs to another zone, within the region or to another region.

Additional detail on service data residency is available in [Where your data is located](#). Up-to-date information on services with single-region data residency is available on the [Azure datacenter map](#).

Data storage for non-regional services

A limited set of Azure non-regional services do not enable the customer to specify the region where the service will be deployed because they rely on a global architecture. Information about these non-regional services is provided on the [Azure datacenter map](#) and a complete list of non-regional services can be found at [Services by Region](#).

Azure services that do not enable the customer to specify the region where the service will be deployed may store customer data in any Microsoft datacenter unless specified otherwise below:

- **Content Delivery Network (CDN)**, which provides a global caching service and stores customer data at edge locations around the world.
- **Azure Active Directory**, which may store Active Directory data globally. This does not apply to Active Directory deployments in the United States (where Active Directory data is stored solely in the United States) and in Europe (where Active Directory data is stored in Europe or the United States). See also “Data storage for AAD identity data” in this paper.
- **Azure Multi-Factor Authentication (MFA)**, which stores data as outlined in [Data residency and customer data for Azure Multi-Factor Authentication](#).
- **Services that provide global routing functions and do not themselves process or store customer data**. This includes [Traffic Manager](#), which provides load balancing between different regions, and [Azure DNS](#), which provides domain name services that route to different regions.

Data storage for Azure Security Center

Azure Security Center is a global (non-regional) service. It may store a copy of security-related customer data, collected from or associated with a customer resource (e.g., a virtual machine or Azure Active Directory tenant): (a) 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 (b) where Azure Security Center uses another Microsoft online service to process such data, it may store such data in accordance with the geolocation rules of that other online service.

Data storage for Cognitive Services

Cognitive Services cover a large set of Azure services. Many are non-regional so customers cannot specify the region where the service will be deployed.

Non-regional Cognitive Services include: Bing Autosuggest, Bing Custom Search, Bing Entity Search, Bing Image Search, Bing Local Business Search, Bing News Search, Bing Speech, Bing Spell Check, Bing Video Search, Bing Visual Search, Bing Web Search, Ink Recognizer, and Translator Text. These services may store customer data in any Microsoft global datacenter.

Regional cognitive services include: Anomaly Detector, Computer Vision, Content Moderator, Custom Vision, Face, Form Recognizer, Immersive Reader, Language Understanding, Personalizer, QnA Maker, Speaker Recognition, Speech Services, Text Analytics, and Video Indexer. For these services, Microsoft will not store customer data outside the customer-specified geo.

A current list of which Cognitive Services are regional and which are non-regional is available at [Products available by region](#).

Additional information on compliance with cognitive services is available at [Azure Cognitive Services](#).

Data storage for AAD identity data

Azure Active Directory (AAD) stores identity data in the geographic location based on the address that an organization uses in its Microsoft online service subscription.

You can get detailed information on where your AAD data is stored in the Azure Active Directory section of [Where is your data located?](#)

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

- Multifactor authentication phone calls or SMS
- Push notifications using the Microsoft Authenticator app
- Validation of OAuth codes
- Azure AD 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
- Azure AD DS stores user data (same location as the customer-selected Azure Virtual Network)
- Microsoft Exchange Server 2013 Application Identifier (AppID), approved federated domains list for Application, and Application's token signing Public Key

Use Azure Policy to control data residence

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, or resource monitoring requirements), Microsoft provides [Azure Policy](#). To restrict the policy to certain Azure regions (e.g. for data residency), the [Allowed Locations](#) policy can be used.

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. Additional information can be found in the overview of [What is Azure Policy?](#)

Use Azure Blueprints to enforce compliance and data residency

[Azure Blueprints](#) is a free service that provides you with templates to create, deploy, and update fully governed cloud environments to consistent standards, and to 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 Resource Manager 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.

Azure Blueprints helps customers build Azure applications that are secure and can become compliant with common industry standards, such as HIPAA, FedRAMP, and PCI DSS. Blueprints help to simplify large-scale Azure deployments by packaging policies in a single blueprint definition that includes artifacts such as Azure Resource Manager templates, resource groups, role-based access controls, and policies.

Microsoft provides built-in blueprints to deploy common compliance certification scenarios. For example, the [ISO 27001 Blueprint](#) and the [PCI DSS Blueprint](#) 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 define [blueprint parameters](#) to prevent conflicts when reusing certain blueprints.

Blueprints can be used to enforce data residency for your specific compliance needs by specifying allowed locations and allowed locations for resource groups. You can find an example of [control mapping of the UK OFFICIAL and UK NHS blueprint samples](#).

III. Access to telemetry data

Telemetry refers to the automated collection of data and can take on many forms. For services where customer data is stored and processed in the cloud, telemetry consists of 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 telemetry data. Microsoft is utilizing telemetry data for clearly defined usage scenarios, in line with GDPR requirements and aligned to the best practices described in ISO/IEC 19944. Telemetry data consists of billing, service health, and support data. These are referred to in the Online Services Terms and the Data Protection Addendum as “legitimate business operations.”

Billing telemetry

Billing data for pay-by-use is sent from regional locations to central billing systems. This data is used exclusively for billing purposes. Billing data flow cannot be prevented in public cloud regions.

Billing telemetry includes Azure subscription ID, date and time for a usage event, resource/licensing identifier, quantity to be billed, and the region where the resource is located. Telemetry may include some additional data fields such as an order number.

Service health telemetry

Service health is an important aspect of telemetry analysis. Microsoft collects schematized telemetry data to diagnose and perform root cause analysis on incidents for the platform. Every service utilizes this data to trigger self-healing processes, and this reduces the amount of manual 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 telemetry in the Azure DevOps tools, without customer data access.

Transparency and customer control of service health telemetry

Azure provides transparency and control functions for some of the most common types of telemetry scenarios for Windows Server, Linux virtual machines, and virtual networks.

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](#) collects monitoring data from the guest operating system and workloads of Azure virtual machines and other compute resources. The [Log Analytics agent](#) collects monitoring data from the guest operating system and workloads of virtual machines in Azure. The Dependency agent collects discovered data about processes running on the virtual machine and external process dependencies.

The following table provides a quick comparison of the Azure Monitor agents for Windows:

| | 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 |

Additional detail is available in [Overview of Azure Monitor agents](#).

For Windows VM diagnostic data, note that 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. Further information on how to use security baselines can be found in [Windows security baselines](#). Information about the Windows Restricted Traffic Limited Functionality Baseline can be found in [Manage connections from Windows operating system components to Microsoft services](#).

Linux on Azure: For customers who run Linux virtual machines on Azure, Microsoft provides the Linux agent (WALinuxagent) as open source software for Linux. The Azure Linux Agent manages Linux & FreeBSD provisioning, and VM interaction with the Azure Fabric Controller. It is available on [GitHub](#). Microsoft provides full transparency so administrators will know which data are sent from Linux to the Azure platform. This information can be correlated and used for further analysis, to monitor important system metrics and to 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](#).

Infrastructure as a Service (IaaS) network telemetry: Customers can leverage [Network Watcher](#) to monitor network traffic associated with their workloads. Azure Network Watcher provides tools to monitor, diagnose, view metrics, and enable or disable logs for resources in an Azure virtual network. Network Watcher is designed to monitor and repair the network health of IaaS products, which include Virtual Machines, 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.

Customer access to service health telemetry

Customers can use Azure tools to manage the health of their workloads. For example, telemetry data is collected by the following services:

- **Azure Monitor:** This service provides a 360-degree view of applications, infrastructure, and network 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, mail, SMS or Logic Apps.
- **Application Insights** is a feature of Azure Monitor. It 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 what users actually do with the app. Application Insights is designed to help continuously improve performance and usability by sending telemetry 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 the data securely. A data retention policy must be defined by the user. The customer can also utilize this data to build high-availability workloads, which could detect an incident based on the telemetry information and perform automated predefined actions to mitigate the incident.

Support telemetry

DevOps and support do not have direct access to customer data. Customers initiating a support request can give Microsoft support engineers access to telemetry data. Through the "Share diagnostic information" feature in the Azure portal, you give your consent to allow a Microsoft Support engineer to remotely collect data from your Azure Virtual Machines or Azure Cloud Services that are subjects of the current support incident in order to troubleshoot your issue.

Diagnostics data

This includes common 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) and your synchronization configuration and related log files (such as Sign-In, Audit, or synchronization logs) are collected to help with troubleshooting.

A detailed list of diagnostics data that is collected can be found in the following articles:

- [Windows Server logs](#)
- [Azure PaaS logs](#)
- [IaaS logs](#)
- [Service Fabric logs](#)
- [StorSimple logs](#)
- [SQL Server Windows VM logs](#)
- [Azure Active Directory logs](#)

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 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 actions are logged and audited. Additional information on memory dump is available in [Azure for Secure Worldwide Public Sector Cloud Adoption](#) (page 19).

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 virtual machine 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, custom scripts, and more.

Support data retention

Support data can be retained for up to 90 days. Additional information on support data telemetry is available in [Azure Support diagnostic information and memory dump collection](#).

Elevated DevOps data access for support cases

When access to customer data is granted, leadership approval is required and then access is carefully managed and logged. The access-control requirements are established by the following Azure Security Policy:

- No access to customer data by default.
- No user or administrator accounts on customer virtual machines (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 accounts by Microsoft. Azure relies on Microsoft corporate Active Directory, 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 telemetry setup. In Azure, the most important use case for telemetry is its use as a sensor for the automated operation of the cloud. Based on the telemetry information 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” in this paper.

See “Who can access your data and on what terms” in this paper for further information regarding limitations on access to customer data, including telemetry data.

IV. How Microsoft protects data from unauthorized access

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 thorough requirements for responding to government requests for customer data.

Who can access your data and on what terms?

Only in rare cases does a Microsoft engineer need access to customer data to resolve a customer issue. 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. (See the section “Support telemetry” for more information on access.).

Access to customer data 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.

Customers can use customer-managed keys to further 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. In either case, Microsoft would not have access to encryption keys and cannot decrypt the 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 SOC 2 Type 2 attestation report published by the auditor explains under what circumstances access to customer data can occur and how. For more information, see the [Azure and Azure Government SOC 2 Type 2 Report](#). By far, the most common scenario involves a customer opening a troubleshooting ticket with Azure Support, and Support subsequently obtaining consent to access customer resources that could potentially include customer data. For the majority of support scenarios, access to customer data is not needed. When access to customer data is needed, customers can manage access, as outlined in the above section “Support telemetry” in this paper.

Microsoft employs rigorous operational controls and processes to prevent unauthorized physical access to datacenters, including 24×7 video monitoring, trained security personnel and processes, and smart card / biometric multifactor 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 the section, “[Data Disk Destruction](#)”.

How Microsoft manages your data

With Azure, you are the owner of your customer data and you retain control over how it is used. You can access your own customer data at any time and for any reason without assistance from Microsoft.

Microsoft does not share customer data for advertising. Your data is your business. Microsoft does not share business customer data with Microsoft advertiser-supported services, or mine it for marketing or advertising. This policy is backed by agreements and adoption of the international code of practice for cloud privacy, ISO/IEC 27018, and by contractual commitments in the Online Services Terms.

Microsoft processes your data only as authorized by you and as required to provide the service and for purposes compatible with providing the service, including day-to-day operations and troubleshooting.

Tenant separation

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. The Azure platform uses a virtualized environment, whereby workloads from different tenants run in isolation on shared physical servers, to keep customers' data secure in the multitenant environment. User instances operate as standalone virtual machines 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 virtual machine.

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 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.

More information can be found in [Isolation in the Azure Public Cloud](#).

When Microsoft deletes your data

If you leave the Azure service or your subscription terminates, Microsoft abides by its commitment in the Online Service 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 the section titled "Data Deletion" in this paper.

Implementation of policies

Microsoft has created a set of internal policies and technical controls that govern data handling. The controls and policies are designed to conform with ISO/IEC 27018, the code of practice for protection of personally identifiable information (PII) in public clouds processing 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.

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. 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 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 Online Services Terms. 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 customers notice (by updating the subprocessor list on its website and providing the customer with a way to obtain notice of that update) at least 14 days in advance of providing that subprocessor with access to customer data or personal data. With respect to Core Online Services, Microsoft will give the customer notice of any new subprocessor at least six months in advance of providing that subprocessor with access to customer data.

If a customer does not approve of a new subprocessor, the customer may terminate its 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-approval. 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.

Further reading:

- [Microsoft Online Services Terms](#)
- [Subcontractors list](#)
- [Microsoft Services Supplier List](#)
- [Microsoft Online Services Subprocessors list](#)

How Microsoft handles government requests

Microsoft has taken a firm public stand on protecting customer data from inappropriate government access, and where necessary, it has advanced its position through the courts.

Microsoft believes customers have a right to know when law enforcement requests their email or documents, and we have a right to tell them. The reason is simple – we believe our customers own their data and have the right to control it. Absent extraordinary circumstances, government agents should seek data directly from our enterprise customers, and if they seek our customers' data from us, they should allow us to tell our customers when demands are made. Unless legally forbidden from doing so, Microsoft will seek to notify customers of law enforcement requests for data access.

Microsoft has dedicated resources to evaluate the sufficiency and legality of law enforcement requests for access to data.

- **Non-content data requires subpoena or court order.** Non-content data includes basic subscriber information, such as email address, name, state, country, ZIP code, and IP address at time of registration. Other non-content data may include IP connection history, an Xbox gamertag, and credit card or other billing information. Microsoft requires a valid legal demand, such as a subpoena or court order, before Microsoft will consider disclosing non-content data to law enforcement.
- **Content data requires a warrant.** Content is what customers create, communicate, and store on or through Microsoft services, such as the words in an email exchanged between friends or business colleagues or the photographs and documents stored on OneDrive or cloud offerings such as Microsoft 365 and Azure. Microsoft requires a warrant (or its local equivalent) before Microsoft will consider disclosing content to law enforcement.

This framework applies to the United States and abroad as applicable (Microsoft will adhere to local laws and regulations). Microsoft conducts a local legal review of each request it receives against local laws and standards. Microsoft also periodically reviews its screening processes around the world to ensure local judicial procedures are being followed and its global human rights statement is being applied.

The CLOUD Act

The adoption of the Clarifying Lawful Overseas Use of Data Act (CLOUD Act) is a result of United States Supreme Court proceedings which were based on the question of the legality of a search warrant issued by a New York court.

The CLOUD Act does not reduce the foregoing protections. The CLOUD Act amends US law to make clear that law enforcement may compel US-based service providers to disclose data that is in their “possession, custody, or control” regardless of where the data is located. This law, however, 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. Microsoft adheres to the same principles and customer commitments related to government demands for user data.

The CLOUD Act serves the investigation of crimes. The CLOUD Act does not oblige cloud service providers to disclose customer information to US law enforcement agencies in any case. It merely provides a legal framework for resolving conflicts of law by enabling the United States and encouraging foreign governments to conclude bilateral agreements on dealing with requests in cross-border investigations.

The UK already concluded such a bilateral agreement with the United States in October 2019. The European Union has not yet concluded such an agreement with the US.

To protect the privacy of its business customers in the future, Microsoft complies with the following five principles:

1. Microsoft will continue to refer US authorities to the respective business customers instead of providing data to the authorities by choice.
2. Microsoft will continue to go to court to defend the local rights of its customers if their rights are violated by the US government.
3. Microsoft will continue to push for new international agreements that strengthen the rights of its customers.
4. Microsoft will be transparent about the number of international search warrants Microsoft receives.
5. Microsoft will continue to offer its customers several alternatives for storing their data.

Further information about the CLOUD Act can be found on the [US Department of Justice web site](#).

Microsoft does not give any government (including law enforcement or other government entities) direct or unfettered access to customer data. Microsoft does not provide any government with customer encryption keys or the ability to break encryption. In many circumstances, Microsoft also offers the option for consumers or enterprises to keep their own keys, and in those instances, Microsoft does not maintain copies.

Law Enforcement Requests Reports

Microsoft is committed to transparency. Its Law Enforcement Requests Report site brings together in one place the reports that Microsoft issues regularly on requests for customer data made by law enforcement, as well as government 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 related to criminal law or national security. For enterprise customers, that number drops further to a mere handful. For example, in the second half of 2019:

- **For consumer services:** In the United States, Microsoft received 4,315 legal demands for consumer data from U.S. law enforcement. Of those, only 145 warrants sought content data which was stored outside of the United States. This included all Microsoft products. The number of legal demands for customer data in Azure is extremely small.
- **For enterprise services:** In the same time frame, Microsoft received only 39 legal demands from US law enforcement for enterprise customers (defined as those who purchased more than 50 seats). Of those demands, one warrant resulted in disclosure of content data related to a non-US enterprise customer whose data was stored outside of the United States.

Twice yearly, Microsoft publishes the number of worldwide official investigation requests on its website. This report only covers law enforcement requests, but Microsoft follows the same principles for responding to government requests for all customer data. You can find these trust reports on the [Law Enforcement Requests Report](#) page under the category “Digital trust reports.” You can also refer to the FAQs on that page, which deal in more detail with the number of investigation requests relating to enterprise cloud customers.

V. How customers can protect data from unauthorized access

Azure offers multiple ways for customers to customize the degree of additional protection for cloud workloads. Encryption is a fundamental component that helps ensure the confidentiality of cloud workloads. Azure provides customers with several offerings to manage and control the security of customer data, including the means to encrypt all of the following:

- Data at rest
- Data in transit
- Data during processing

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. Further information can be found in the [Azure Encryption Overview](#) document.

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:

- Managed by Azure in Key Vault.
- Managed by the customer in Key Vault.
- Managed and stored on-premises.

Storing the customer keys on-premises eliminates the possibility for Azure to decrypt workloads, though it also limits some Azure functionality, such as the search functionality. Details are described in the following sections.

Encryption of data at rest

Customer data at rest is automatically encrypted when it’s written to Azure Storage, including Azure Managed Disks, Azure Blob, Queue, Table Storage, or 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. Azure Storage encryption is similar to BitLocker encryption on Windows.

Key management is discussed below 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.

Further information regarding Storage Service Encryption can be found in [Azure Storage encryption for data at rest](#).

Disk encryption for VMs and Virtual Machine Scale Sets (VMSSs): Customers can configure operating system (OS) and data disks used by Azure virtual machines (VMs) to be encrypted using [Azure disk encryption](#). Azure offers multiple options to encrypt OS and data discs for Windows Server and Linux instances. Azure Disk Encryption encrypts Windows and Linux infrastructure as a service (IaaS) virtual machine disks by using the BitLocker feature of Windows and the DM-Crypt feature of Linux to provide volume encryption for the operating system disk and the data disk. This protects both operating system disks and data disks with full volume encryption. 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 virtual machine. The Shielded VMs solution includes the new Host Guardian Service feature, which is used for virtualization host attestation and encryption key release. Further information regarding encrypting Windows and Linux VM disks can be found in [Azure Disk Encryption for virtual machines and virtual machine scale sets](#).

If you want to use a key encryption key (KEK) for an additional layer of security for encryption keys, you can add a 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](#), customers can back up and restore encrypted virtual machines (VMs) that use the key encryption key (KEK) configuration.

Encryption technologies for other data at rest

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

- **Transparent Data Encryption (TDE)** encrypts data at rest when it's stored in Azure Synapse Analytics.
- **Always Encrypted** supports the ability to encrypt data within client applications prior to storing it in Azure SQL Database.
- **Azure Data Lake Storage (ADLS)** is protected by transparent encryption of data at rest similar to what is provided with Azure SQL Database. Azure Data Lake Storage 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** is encrypted by default, using secure key storage systems, encrypted networks, and cryptographic APIs. The encryption keys are managed by Microsoft and rotated per its internal guidelines.

Key management

Key management can be performed by Azure or by the customer, and encryption can be performed client-side or server-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** provide a combination of control and convenience with low overhead. Azure Resource Providers perform the encryption and decryption operations. Microsoft manages the keys.
 - **Customer-managed keys** give you control over the keys, including Bring Your Own Keys (BYOK) support, or allows you to generate new ones. Azure Resource Providers perform the encryption and decryption operations. Customer controls keys via Azure Key Vault
 - **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, cloud service providers don't 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](#).

Key management – Server-side encryption

Azure Key Vault 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. This supports Microsoft customers in performing additional analysis, such as threat detection.

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 it 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.

Customer-managed keys: Azure Key Vault also provides a bring-your-own-key capability. Customers can generate the keys on premises using an offline workstation equipped with a nCipher HSM and then transmit the keys securely to the Azure HSMs in the cloud. The nCipher software used for the 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.

Customer-provided keys 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 enable you to pass an encryption key as part of read or write operation to a storage service using [blob APIs](#). 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, 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](#) 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 400 Bad Request. As the encryption key itself is provided in the request, a secure connection must be established to transfer the key.

Further information about how to use server-side encryption to protect secrets, certificates, and keys can be found in the following documents:

- [Azure Storage encryption for data at rest](#)
- [Configure customer-managed keys with Azure Key Vault](#)
- [Bring your own key \(BYOK\) Azure Key Vault](#)
- [Customer-provided keys](#)

Key management – Client-side encryption (Customer-managed on-premises encryption keys)

The client-side encryption model refers to encryption that is performed outside of the Resource Provider or Azure. It includes: 1) Data encrypted by an application that's running in the customer's datacenter or by a service application. 2) Data that is already encrypted when it is received by Azure.

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/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).

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 Hardware Security Module. 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.

Read more:

- [Azure Data Encryption-at-Rest.](#)
- [Client-Side Encryption and Azure Key Vault for Azure Storage](#)

Encryption of data in transit

Protecting data in transit should be an essential part of your 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 you should always use Secure Sockets Layer/Transport Layer Security (SSL/TLS) protocols to exchange data across different locations.

Microsoft enables and encourages Azure customers to encrypt customer data in transit to Microsoft 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 provide 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 encryption of data in transit. Every Azure server contains Azure SmartNICs, which are based on the 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](#).

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

- For data moving between your on-premises infrastructure and Azure, you should 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](#).
- For organizations that need to secure access from one workstation located on premises to Azure, use [Point-to-Site VPN](#).

- Larger data sets can be moved over a dedicated high-speed WAN link such as [ExpressRoute](#). If you choose to use ExpressRoute, you can also encrypt the data at the application-level using [SSL/TLS](#) or other protocols for added protection.
- If you are interacting with Azure Storage through the Azure Portal, all transactions occur via HTTPS. [Storage REST API](#) over HTTPS can also be used to interact with [Azure Storage](#) and [Azure SQL Database](#).

ExpressRoute encryption

ExpressRoute supports the following encryption technologies to ensure 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.

You can read more about how these technologies work in [ExpressRoute Encryption](#).

Encryption during processing of data (confidential computing)

[Azure confidential computing](#) protects customer data at runtime, bringing Intel SGX and Virtualization Based Security (VBS) 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. The TEE enforces these protections against viewing and modification, including access by Microsoft personnel, during the time the data is being processed.

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 DCsv2 series VMs. The DCsv2-series machines are backed by the latest generation of Intel XEON E-2288G processors with SGX technology. (See [DCsv2-series](#) for sizing information.)

Further information about Azure encryption technologies and options can be found in the following documents:

- [Azure Encryption Overview](#)
- [Azure Data Encryption-at-Rest](#)
- [Azure Data Security and Encryption Best Practices](#)
- [Azure Cosmos DB Encryption](#)
- [Storage Service Encryption using customer-managed keys in Azure Key Vault](#)
- [Azure Storage Security Guide](#)
- [Azure Confidential Computing](#)

Customer Lockbox for Azure

To further safeguard customer data, Microsoft has introduced the Customer Lockbox for Azure. Customer Lockbox is a service that provides customers with the capability 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 the support workflow, a Microsoft engineer

may require elevated access to customer content. Customer Lockbox puts the customer in charge of that decision by enabling the customer to review and approve or deny such elevated requests.

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 Customer Lockbox via the Azure portal and integrate them into their SIEM systems.

Note that for the majority of support scenarios, access to customer data is not needed and the workflow should not require Customer Lockbox.

Read more about how this works in [Customer Lockbox for Azure](#).

VI. Data retention and deletion

Data retention and deletion policies and practices are important to protecting data. Privacy compliance requires that fundamental principles for minimizing data processing and retention be followed. Microsoft follows strict guidelines for retaining and deleting data.

Data retention

In the [Online Services Data Protection Addendum](#) (DPA), 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.

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) to enable customers to extract their data or renew the subscription. During this period, Microsoft provides multiple notices so the customer will be amply forewarned of the upcoming deletion of its data from Microsoft systems.

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 Online Services Terms).

Data retention for Cognitive Services: For the purposes of data retention and deletion, a Cognitive Services configuration or custom model that has been inactive may, 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.

Data deletion

Microsoft uses different types of data deletion techniques depending on the type of data object that is being deleted; for example, whole subscriptions, storage, virtual machines, or databases.

- **Subscriptions:** As earlier referenced, 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 subscription.
- **Azure Storage:** In Azure Storage, all disk writes are sequential. 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 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 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 Virtual Machines** are stored in Azure Storage as blobs, and the deletion rules apply as explained 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 occur 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 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.

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. Records of the destruction are retained and reviewed as part of the Microsoft audit and compliance process. All Microsoft Azure services utilize approved media storage and disposal management services.

VII. Compliance

Microsoft Trust Center provides a centralized location where you can find resources for information about security, privacy, compliance, and transparency in regard to Microsoft products and services, including data residency information. This is where you can go for guidance on government and industry-specific compliance, audit reports, security assessments, and more.

The Microsoft Trust Center provides a list of Microsoft compliance offerings. Azure possesses an industry-leading compliance portfolio. This helps customers meet their compliance commitments. Azure compliance offerings are grouped into four categories: globally applicable, US government, industry specific, and region specific. Compliance offerings are based on various types of assurances, including formal certifications, attestations, validations, authorizations, and assessments produced by independent third-party auditing firms, as well as contractual amendments, self-assessments, and customer guidance documents produced by Microsoft.

Learn more about the extensive [Azure offerings and certification coverage](#).

Data privacy compliance and the GDPR: Microsoft is committed to complying with local privacy laws and protecting the privacy of its users. The Azure compliance portfolio includes conformity to cloud privacy practices such as ISO/IEC 27018. That commitment includes ensuring that Microsoft products and services comply with laws that are applicable to cloud providers, such as the GDPR.

Microsoft will comply with all laws and regulations applicable to its provision of its online services, including security breach notification laws. However, Microsoft is not responsible for compliance with any laws or regulations applicable 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.

Microsoft believes data privacy is a fundamental right, and that the GDPR is an important step forward for clarifying and enabling individual privacy rights. Also, Microsoft recognizes that the GDPR required significant changes by organizations all over the world with regards to the discovery, management, protection, and reporting of personal data that is collected, processed, and stored within an organization.

Microsoft contractually commits to meet GDPR requirements not only in the European Union but in all public cloud regions. Azure is also certified for the ISO/IEC 27701 standard. ISO/IEC 27701 is an extension of the widely-used ISO/IEC 27001 standard for information security management, making the implementation of a privacy information management system a helpful compliance extension for the many organizations that rely on ISO/IEC 27001, as well as creating a strong integration point for aligning security and privacy controls.

Most countries allow the transfers outside of their boundaries, though often with certain restrictions. EU countries are governed by the GDPR, which regulates transfers of personal data of European residents to destinations outside the European Economic Area. The transfer of data outside of EU boundaries requires a specific legal mechanism, such as a contract, or adherence to a certification mechanism to enable these transfers. Microsoft details the mechanisms they use in the Online Services Terms.

Service Trust Platform: Microsoft provides guidelines to help customers match requirements for regulated data types. These documents are available for download in the [Service Trust Platform \(STP\) portal](#). The STP is a companion feature to the Microsoft Trust Center, and it enables Microsoft customers to:

- Access audit reports across Microsoft cloud services on a single page.
- Access compliance guides that help customers understand how they can use Microsoft cloud service features to manage compliance with various regulations.
- Access trust documents to help in understanding how Microsoft cloud services help protect data.
- Conduct assessments and pre-audits to prepare for external audits.

Customers who have active paid or trial subscriptions with Azure accounts can access the STP directly. New customers and those who are evaluating Microsoft online services can access the STP with any Live-ID or O365 account.

Microsoft commitments, including Online Services Term and Online Services Data Protection Addendum: When customers subscribe to an online service through a Microsoft Volume Licensing program, the terms that control how they can use the service are defined in the Volume Licensing Online Services Terms (OST) and the [Online Services Data Protection Addendum](#) (DPA). Due to the frequency with which Microsoft adds new services, the OST is updated monthly. Additional amendments exist to cover restricted industries, including financial services. The OST is available in 35 languages. An accessible archive contains older versions for reference.

Microsoft makes commitments regarding important aspects of cloud usage by Microsoft customers, including the following:

- **Data privacy commitments:** Microsoft has implemented operational processes to meet the exacting requirements of the GDPR, as defined in the Data Protection Addendum. Microsoft also offers customers EU Model Clauses, referred to as Standard Contractual Clauses, that make specific guarantees around transfers of personal data for in-scope Azure services.
- **Technical and organizational measures:** Microsoft describes and documents technical and organizational measures in the Online Services Terms. Those measures are set forth in the Microsoft Security Policy. These comply with the requirements for Information Security Management Systems (ISMS) set forth in ISO27001, ISO 27002, and ISO 27018. Additional descriptions of the technical and organizational measures can be found in the Azure SOC 2 Type 2 assessment report.
- **Service level agreements (SLA):** The detailed service level terms for Azure can be found on the [Service Level Agreements](#) page of the Azure website.

Learn more about the OST on the [Licensing Terms](#) webpage.

Conclusion

This whitepaper provides insights into Microsoft data residency promises, Microsoft access to data, how customers and Microsoft protect data access, and Azure's data retention policies.

Microsoft continuously adds features and documentation to assist customers in addressing their compliance needs and to provide transparency regarding Microsoft practices and processes.

Appendix A: Selected Resources

Microsoft currently hosts datacenters in over 60 regions across multiple countries. Microsoft provides compliance guidance for customers through a series of documents. These documents address data residency requirements generally, along with special emphasis on the financial services and healthcare sectors in over 40 countries.

- **Navigating Your Way to the Cloud in Europe: A Compliance Guide:** Belgium, Bulgaria, Croatia, Czech Republic, Estonia, Finland, France, Germany, Italy, Ireland, Latvia, Lithuania, Luxembourg, Netherlands, Norway, Poland, Slovenia, Spain, Sweden, Switzerland, U.K.
- **Navigating your way to the cloud in Asia: A Guide for the Legal & Compliance Professionals:** 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:** Angola, Jordan, Kenya, Mauritius, Morocco, Nigeria, Rwanda, South Africa, UAE
- **Service Trust Portal Cloud Compliance Guides for Financial Services:** Australia, Belgium, Brazil, Canada, China, Czech Republic, Denmark, European Union, France, Germany, Greece, Hong Kong, Hungary, India, Israel, Italy, Japan, Luxembourg, Mexico, Netherlands, New Zealand, Nigeria, Norway, Poland, Saudi Arabia, Singapore, South Africa, Spain, Sweden, Switzerland, United Kingdom, United States

Recommended whitepaper on data residency and security:

- [Azure for Secure Worldwide Public Sector Cloud Adoption](#)

