

Microsoft Azure

Azure for Secure Worldwide Public Sector Cloud Adoption



Abstract

Microsoft Azure is a multi-tenant cloud services platform that government agencies can use to deploy a variety of solutions. A multi-tenant cloud platform implies that multiple customer applications and data are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's applications and data from those of others. This approach provides the scale and economic benefits of multi-tenant cloud services while rigorously helping prevent customers from accessing one another's data or applications. Azure is available globally in more than 50 regions and can be used by government entities worldwide to meet rigorous data protection requirements across a broad spectrum of data classifications, including unclassified and classified data.

This paper addresses common security and isolation concerns pertinent to worldwide public sector customers. It also explores technologies available in Azure to safeguard both unclassified and classified workloads in the public multi-tenant cloud in combination with Azure Stack and Data Box Edge deployed on-premises and at the edge.

February 2019

<https://aka.ms/AzureWWPS>

Acknowledgments

Author: Stevan Vidich

Contributors: LeeAnn Gunther, Nate Johnson, James Kavanaugh, Diego Lapiduz, Dan Quintiliano, Kamran Zargahi

Reviewers: AT Ball, Lindsay Berg, Dave De Bie, Andreas Ebert, Lily Kim, Erin Nord, Mahesh Punyamurthula, Jack Richins, Rahul Savdekar, Christoph Siegert, Shigenori Tanaka, Karthik Thirumalai, Devendra Tiwari, Seth Varty

(c) 2019 Microsoft Corporation. All rights reserved. This document is provided "as-is". Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

Contents

- Introduction 4
- Key considerations for public sector customers 4
 - Data residency and government requests for customer data 5
 - Data encryption in transit, at rest, and in use 7
 - Access to Customer Data by Microsoft personnel..... 10
 - Threat detection and prevention..... 11
 - Private and hybrid cloud with Azure Stack and Data Box Edge 14
 - Compliance and certifications 17
 - Logical isolation considerations 18
 - Physical versus logical security considerations..... 22
- Conceptual architecture 25
 - Azure for Confidential data..... 26
 - Azure for Sensitive data 27
 - Azure for Highly Sensitive data 27
- Select workloads and use cases 28
 - Processing highly sensitive or regulated data on Azure Stack..... 28
 - Machine Learning model training 29
 - IoT analytics 29
 - Precision Agriculture with Farm Beats..... 29
 - Unleashing the power of analytics with Synthetic Data 30
 - Knowledge mining 30
 - Scenarios for Confidential Computing 30
- Frequently asked questions 31
 - Data residency and data sovereignty..... 31
 - Safeguarding of customer data..... 33
 - Operations 34
 - Transparency and audit 35

Introduction

Governments around the world are in the process of a digital transformation, actively investigating solutions and selecting architectures that will help them transition many of their workloads to the cloud. There are many drivers behind the digital transformation, including the need to engage citizens, empower employees, transform government services, and optimize government operations. Governments across the world are also looking to improve their cybersecurity posture to secure their assets and counter the evolving threat landscape.

For governments and the public sector industry worldwide, Microsoft provides Azure – a public multi-tenant cloud services platform that government agencies can use to deploy a variety of solutions. A multi-tenant cloud platform implies that multiple customer applications and data are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's applications and data from those of others. This approach provides the scale and economic benefits of multi-tenant cloud services while rigorously helping prevent customers from accessing one another's data or applications.

A hyperscale public cloud provides resiliency in time of natural disaster and warfare. The cloud provides capacity for failover redundancy and empowers sovereign nations with flexibility regarding global resiliency planning. Hyperscale public cloud also offers a feature rich environment incorporating the latest cloud innovations such as artificial intelligence, machine learning, IoT services, intelligent edge, and many more to help government customers increase efficiency and unlock insights into their operations and performance.

Using Azure's public cloud capabilities, customers benefit from rapid feature growth, resiliency, and the cost-effective operation of the hyperscale cloud while still obtaining the levels of isolation, security, and confidence required to handle workloads across a broad spectrum of data classifications, including unclassified and classified data. Leveraging Azure isolation technologies, as well as intelligent edge capabilities (such as Azure Stack and Data Box Edge, as described later in the document), customers can process confidential and sensitive data in secure isolated infrastructure within Azure public multi-tenant regions or highly sensitive data at the edge under the customer's full operational control.

This paper addresses common security and isolation concerns pertinent to worldwide public sector customers. It also explores technologies available in Azure to help safeguard unclassified, confidential, and sensitive workloads in the public multi-tenant cloud in combination with Azure Stack and Data Box Edge deployed on-premises and at the edge for fully disconnected scenarios involving highly sensitive data. Given that unclassified workloads comprise the majority of scenarios involved in worldwide public sector digital transformation, Microsoft recommends that customers start their cloud journey with unclassified workloads and then progress to classified workloads of increasing data sensitivity.

Key considerations for public sector customers

This section provides a high-level overview of key Azure public cloud features and capabilities that enable the processing of unclassified and classified workloads in combination with Azure Stack and Data Box Edge for private and hybrid cloud deployment models. Also provided is an overview of logical isolation in the public cloud, as well as the physical versus logical isolation considerations that are likely to be of interest to government customers considering cloud adoption.

Data residency and government requests for customer data

This section addresses key cloud implications for data residency and data sovereignty, as well as the fundamental principles guiding Microsoft's handling of worldwide law enforcement requests for Customer Data, including CLOUD Act provisions.

Microsoft defines [Customer Data](#) as all data, including text, sound, video, or image files and software that customers provide to Microsoft to manage on customer's behalf through customer's use of Microsoft online services. Microsoft provides [strong customer commitments](#) regarding cloud services data residency and transfer policies:

- **Data storage for regional services:** Most Azure services are deployed regionally and enable the customer to specify the region into which the service will be deployed, e.g., Europe. Microsoft will not store Customer Data outside the customer-specified Geo except for [Cloud Services](#), [Cognitive Services](#), [Azure Databricks](#), and Preview services as described on the [data location page](#). This commitment helps ensure that Customer Data stored in a given region will remain in the corresponding Geo and will not be moved to another Geo for the majority of regional services, including Storage, SQL Database, Virtual Machines, etc.
- **Data storage for non-regional services:** Certain Azure services do not enable the customer to specify the region where the services will be deployed as described on the [data location page](#). For a complete list of non-regional services, customers should see [Services by Region](#).

Customer Data in an Azure Storage account is [always replicated](#) to help ensure durability and high availability. Azure Storage copies Customer Data to protect it from transient hardware failures, network or power outages, and even massive natural disasters. Customers can choose to replicate their data within the same data center, across availability zones within the same region, or across geographically separated regions. Specifically, when creating a storage account, customers can select one of the following redundancy options:

- Locally redundant storage (LRS)
- Zone redundant storage (ZRS)
- Geo redundant storage (GRS), including option for read-access GRS

Azure Storage redundancy options can have implications on data residency as Azure relies on [paired regions](#) to deliver GRS. For example, customers concerned about geo-replication across regions that span country boundaries, may want to choose LRS or ZRS to keep Azure Storage data at rest within the geographic boundaries of the country in which the primary region is located. Similarly, [geo replication for Azure SQL Database](#) can be obtained by configuring asynchronous replication of transactions to any region in the world, although it is recommended that paired regions be used for this purpose as well. If customers need to keep relational data inside the geographic boundaries of their country, they should not configure Azure SQL Database asynchronous replication to a region outside that country.

Data sovereignty implies data residency; however, it also introduces rules and requirements that define who has control over and access to Customer Data stored in the cloud. In many cases, data sovereignty mandates that Customer Data be subject to the laws and legal jurisdiction of the country in which data resides. These laws can have direct implications on data access even for service troubleshooting or customer-initiated support requests.

Customers can use Azure public multi-tenant cloud in combination with Azure Stack and Data Box Edge for on-premises and edge solutions to meet their data sovereignty requirements, as described later in this document. These additional products can be deployed to put customers solely in control of their data, including storage, processing, transmission, and remote access.

Government requests for Customer Data follow a strict procedure. Microsoft takes strong measures to help protect Customer Data from inappropriate access or use by unauthorized persons. This includes restricting access by Microsoft personnel and subcontractors and carefully defining requirements for responding to government requests for Customer Data. Microsoft ensures that there are no back-door channels and no direct or unfettered government access to Customer Data. Microsoft imposes special requirements for government and law enforcement requests for Customer Data.

As stated in the [Online Services Terms](#), Microsoft will not disclose Customer Data to law enforcement unless required by law. If law enforcement contacts Microsoft with a demand for Customer Data, Microsoft will attempt to redirect the law enforcement agency to request that data directly from the customer. If compelled to disclose Customer Data to law enforcement, Microsoft will promptly notify the customer and provide a copy of the demand unless legally prohibited from doing so.

Government requests for Customer Data must comply with applicable laws. A subpoena or its local equivalent is required to request non-content data and a warrant, court order, or its local equivalent is required for content data. Every year, Microsoft rejects a number of law enforcement requests for Customer Data. Challenges to government requests can take many forms. In many of these cases, Microsoft simply informs the requesting government that it is unable to disclose the requested information and explains the reason for rejecting the request. Where appropriate, Microsoft challenges requests in court. Customers should review the [Law Enforcement Requests Report](#) that Microsoft publishes twice a year.

The **CLOUD Act** is a United States law that was enacted in March 2018. Customers should refer to the following [blog](#) for more information, as well as the follow-up [posting](#) that describes Microsoft's call for principle-based international agreements governing law enforcement access to data. Key points of interest to government customers procuring Azure services are captured below.

- The CLOUD Act enables governments to negotiate new government-to-government agreements that will result in greater transparency and certainty for how information is disclosed to law enforcement agencies across international borders.
- The CLOUD Act is not a mechanism for greater government surveillance; it is a mechanism toward ensuring that customer data is ultimately protected by the laws of each customer's home country while continuing to facilitate lawful access to evidence for legitimate criminal investigations. Law enforcement in the U.S. still needs to obtain a warrant demonstrating probable cause of a crime from an independent court before seeking the contents of communications. The CLOUD Act requires similar protections for other countries seeking bilateral agreements.
- While the CLOUD Act creates new rights under new international agreements, it also preserves the common law right of cloud service providers to go to court to challenge search warrants when there is a conflict of laws – even without these new treaties in place.

- Microsoft retains the legal right to object to a law enforcement order in the United States where the order clearly conflicts with the laws of the country where Customer Data is hosted. Microsoft will continue to carefully evaluate every law enforcement request and exercise its rights to protect customers where appropriate.
- For legitimate enterprise customers, U.S. law enforcement will, in most instances, now go directly to the customer rather than Microsoft for information requests.

Microsoft does not disclose additional data as a result of the CLOUD Act. This law does not practically 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.

As described in this white paper, Azure offers an unmatched variety of public, private, and hybrid cloud deployment models to address each customer’s concerns regarding the control of their data. Government customers worldwide expect to be fully in control of protecting their data in the cloud. As described in the next section, Azure enables customers to protect their data through its entire lifecycle whether at rest, in transit, or in use.

Data encryption in transit, at rest, and in use

Azure has extensive support to safeguard Customer Data using [data encryption in transit and at rest](#), as well as [data encryption while in use](#).

Data encryption in transit: Azure uses the Transport Layer Security (TLS) protocol to help protect data when it is traveling between Customers and Azure services. TLS provides strong authentication, message privacy, and integrity. Perfect Forward Secrecy (PFS) protects connections between Customer’s client systems and Microsoft cloud services by generating a unique session key for every session a Customer initiates. PFS protects past sessions against potential future key compromises. Connections also use RSA-based 2,048-bit encryption key lengths. This combination makes it more difficult to intercept and access data that is in transit. Customers should review Azure [best practices](#) for the protection of data in transit and properly configure HTTPS endpoints for their resources provisioned in Azure to help ensure that all traffic going to and from their Virtual Machines is encrypted. For key Azure services (e.g., Azure SQL Database), data encryption in transit is [enforced by default](#).

Data encryption at rest: Azure provides extensive options for [data encryption at rest](#) to help customers safeguard their data and meet their compliance needs using both Microsoft managed encryption keys, as well as customer managed encryption keys.

Azure [Storage Service Encryption for Data at Rest](#) ensures that data is automatically encrypted before persisting it to Azure Storage and decrypted before retrieval. All data written to Azure Storage is encrypted through 256-bit AES encryption, and the handling of encryption, decryption, and key management in Storage Service Encryption is transparent to customers. However, customers can also [use their own encryption keys for Azure Storage](#) encryption at rest and manage their keys in [Azure Key Vault](#). Storage Service Encryption is enabled by default for all new and existing storage accounts and cannot be disabled.

Azure SQL Database provides [Transparent Data Encryption](#) (TDE) at rest by [default](#). TDE performs real-time encryption and decryption operations on the data and log files. Database Encryption Key (DEK) is a symmetric key stored in the database boot record for availability during recovery. It is secured via a certificate stored in the master database of the server or an asymmetric key called TDE Protector stored under customer control in [Azure Key Vault](#), which is Azure's cloud-based external key management system. Azure Key Vault supports [Bring Your Own Key](#) (BYOK), which enables customers to store TDE Protector in Key Vault and control key management tasks including key rotation, permissions, deleting keys, enabling auditing/reporting on all TDE Protectors, etc. See [Always Encrypted](#) for more information. Note that Always Encrypted is a feature of Azure SQL Database designed specifically to help protect sensitive data by allowing clients to encrypt data inside client applications and never reveal the encryption keys to the [Database Engine](#). In this manner, Always Encrypted provides separation between those who own the data (and can view it) and those who manage the data (but should have no access).

Encryption support is also in place for customer [IaaS Virtual Machines](#), enabling customers to encrypt their Windows and Linux IaaS VM disks. Disk encryption leverages the industry standard [BitLocker](#) feature of Windows and the [DM-Crypt](#) feature of Linux to provide volume encryption for the OS and data disks. The solution is integrated with Azure Key Vault to help customers control and manage the disk encryption keys.

[Azure Key Vault](#) is a multi-tenant secrets management service that uses Hardware Security Modules (HSMs) to store and control access to secrets, encryption keys, and certificates. Key Vault HSMs are FIPS 140-2 Level 2 validated, which includes requirements for physical tamper evidence and role-based authentication. With Key Vault, customers can import or [generate encryption keys](#) in HSMs that never leave the HSM boundary to support Bring Your Own Key (BYOK) scenarios. Key Vault is designed, deployed, and operated such that Microsoft and its agents do not see or extract customer keys.

For customers who require single-tenant HSMs, Azure provides [Dedicated HSMs](#) that have FIPS 140-2 Level 3 validation, as well as Common Criteria EAL4+ certification and conformance with eIDAS requirements. Azure Dedicated HSM is most suitable for scenarios where customers require full administrative control and sole access to their HSM device for administrative purposes. Dedicated HSMs are provisioned directly on customer's virtual network and can also connect to on-premises infrastructure via Virtual Private Network (VPN).

Data encryption in use: [Azure Confidential Computing](#) is a set of new data security capabilities that offers encryption of data while in use. This means that data can be [processed in the cloud](#) with the assurance that it is always under customer control. Confidential computing ensures that when data is in the clear, which is needed for efficient data processing in memory, the data is protected inside a Trusted Execution Environment (TEE, also known as an enclave), as depicted in Figure 1. TEE helps ensure that there is no way to view data or the operations from outside the enclave and that only the application designer has access to TEE data; access is denied to everyone else including Azure administrators. Moreover, TEE helps 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.

Azure supports two TEEs: 1) Virtualization Based Security – VBS, which is a software based TEE that is implemented by the Hyper-V hypervisor, and 2) Intel SGX, which is a hardware based TEE available on a new family of Azure [DC-series Virtual Machines](#) that have the latest generation of Intel Xeon processors

with [Intel Software Guard Extensions](#) (SGX) technology. Intel SGX isolates a portion of physical memory to create an enclave where select code and data are protected from viewing or modification. The protection offered by Intel SGX, when used appropriately by application developers, can [prevent compromise](#) due to attacks from privileged software and many hardware-based attacks. An application leveraging Intel SGX needs to be re-factored into trusted and untrusted components. The untrusted part of the application sets up the enclave, which then allows the trusted part to run inside the enclave. No other code, irrespective of the privilege level, has access to the code executing within the enclave or the data associated with enclave code. Design best practices call for the trusted partition to contain just the minimum amount of content required to protect customer’s secrets.

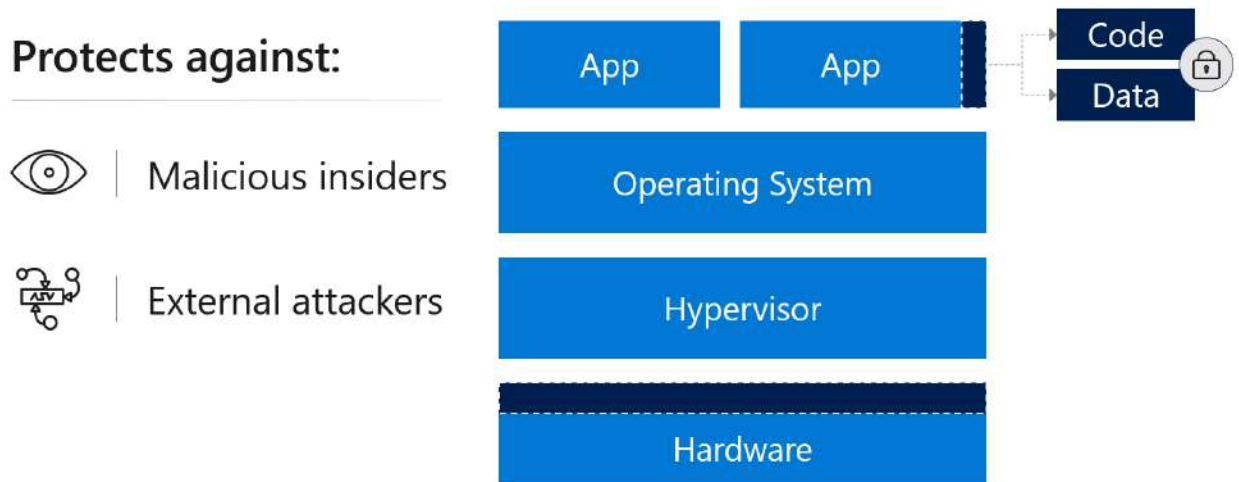


Figure 1. Trusted Execution Environment protection

Microsoft has published the [Open Enclave](#) SDK as an open-source code on GitHub, which developers can use to build C/C++ enclave applications targeting Intel SGX technology. The Open Enclave SDK is intended to be [portable across enclave technologies](#), cross platform – cloud, hybrid, edge, or on-premises, and designed with architectural flexibility in mind. In addition, Microsoft is working on tooling and debugging support for developing and testing confidential applications.

Customers can now build, deploy, and run applications that protect data confidentiality and integrity through the entire data lifecycle whether at rest, in transit, or in use. To get started, customers can deploy a DC-series VM through the custom deployment flow in [Azure Marketplace](#). The custom deployment flow deploys and configures the VM and installs the Open Enclave SDK for Linux VMs if selected. Many of the basic VM deployment configurations are [supported through the Confidential Computing](#) VM Deployment workflow, including: 1) Windows/Linux VM, 2) New or existing resource group, 3) New or existing VNet, 4) Storage/disk type, 5) Enabled diagnostics, and other properties.

Based on the feedback from private preview customers, Microsoft has started to invest in higher level scenarios of confidential computing, including:

- **Confidential querying in databases:** Existing [Always Encrypted](#) technology in SQL Server and Azure SQL Database will be extended using [Always Encrypted with Secure Enclaves](#) to allow computations on plaintext data inside a secure enclave on the server side. While existing Always Encrypted technology offers protection from malware and high-privileged unauthorized

users (e.g., database administrators, cloud admins, etc.), it restricts SQL Server operations on encrypted columns inside the database as data encryption originally takes place on the client side without allowing the data or the corresponding cryptographic keys to appear in plaintext inside the SQL Server engine.

- **Securing the intelligent edge:** [Azure IoT Edge security with enclaves](#) helps customers solve a very challenging security problem in the Internet of Things (IoT) – protecting code and data while in use at the edge. The solution includes securing compute workloads within the TEE confines by leveraging a platform for developing edge applications that execute in enclaves.
- **Creating confidential consortium networks that scale:** The [Confidential Consortium Blockchain Framework](#) is an open-source system that enables high-scale, confidential blockchain networks that meet key enterprise requirements such as performance, enhanced confidentiality, and distributed governance. Leveraging the power of existing blockchain protocols, TEEs, distributed systems, and cryptography, the Confidential Consortium Blockchain Framework offers a trusted foundation for blockchain protocols to digitally transform business.
- **Secure multi-party Machine Learning (ML):** Privacy-preserving [multi-party ML](#) allows multiple organizations to perform collaborative data analytics while guaranteeing the privacy of their individual datasets. Using Intel SGX technology, multiple parties can agree on a joint ML task to be executed on their aggregate data. Although they do not trust one another, they can each review the corresponding ML code, deploy the code into a TEE, upload their encrypted data just for this task, perform remote attestation, securely upload their encryption key in the enclave, run ML code, and finally download the encrypted ML model.

Data encryption in the cloud is an important risk mitigation requirement expected by government customers worldwide. As described in this section, Azure helps customers protect their data through its entire lifecycle whether at rest, in transit, or even in use. Moreover, Azure offers comprehensive encryption key management to help customers control their keys in the cloud, including key rotation, key deletion, permissions, etc. End-to-end data encryption using advanced ciphers is fundamental to ensuring confidentiality and integrity of customer data in the cloud.

Nonetheless, government customers worldwide also expect to receive assurances regarding any potential customer data access by Microsoft engineers for troubleshooting, customer support, or other scenarios. These controls are described in the next section.

Access to Customer Data by Microsoft personnel

Microsoft takes strong measures to protect [Customer Data](#) from inappropriate access or use by unauthorized persons. Microsoft engineers [do not have default access](#) to Customer Data in the cloud. Instead, they are granted access, under management oversight, only when necessary. Using the [restricted access workflow](#), access to Customer Data is carefully controlled, logged, and revoked when it is no longer needed. For example, access to Customer Data may be required to resolve customer-initiated troubleshooting requests. The access control requirements are [established by the following policy](#):

- No access to Customer Data, by default.
- No user or administrator accounts on customer Virtual Machines (VMs).
- Grant the least privilege that is required to complete task; audit and log access requests.

Microsoft engineers can be granted access to Customer Data using temporary credentials via **Just-in-Time Access (JIT)**. There must be an incident logged in the Azure Incident Management system that describes the reason for access, approval record, what data was accessed, etc. This approach ensures that there is appropriate oversight for all access to Customer Data and that all JIT actions (consent and access) are logged for audit. Evidence that procedures have been established for granting temporary access for Azure personnel to Customer Data and applications upon appropriate approval for customer support or incident handling purposes is available from the Azure [SOC 2 Type 2 attestation report](#) produced by an independent third-party auditing firm.

[Azure Customer Lockbox](#) is a service that provides customers with the capability to control how a Microsoft Engineer accesses their data. As part of the Support workflow, a Microsoft Engineer may require elevated access to Customer Data. Azure Customer Lockbox puts the customer in charge of that decision by enabling the customer to Approve/Deny such elevated requests. Azure Customer Lockbox is an extension of the JIT workflow and comes with full audit logging enabled. It is important to note that Azure Customer Lockbox capability is not required for support cases that do not involve access to Customer Data. For the majority of support scenarios, access to Customer Data is not needed and the workflow should not require Azure Lockbox. Azure Lockbox is available to customers from all Azure public regions.

Aside from controls implemented by Microsoft to safeguard Customer Data, government customers deployed in Azure derive considerable benefits from security research that Microsoft conducts to protect the cloud platform. Microsoft global threat intelligence is one of the largest in the industry and it is derived from one of the most diverse sets of threat telemetry sources. It is both the volume and diversity of threat telemetry that makes Microsoft machine learning algorithms applied to that telemetry so powerful. All Azure customers benefit directly from these investments as described in the next section.

Threat detection and prevention

The Microsoft [Intelligent Security Graph](#) uses advanced analytics to synthesize massive amounts of threat intelligence and security signals obtained across Microsoft products, services, and partners to combat cyberthreats. Millions of unique threat indicators across the most diverse set of sources are generated every day by Microsoft and its partners and shared across Microsoft products and services (Figure 2). Across its portfolio of global services, each month Microsoft scans more than 400 billion email messages for phishing and malware, processes 450 billion authentications, executes more than 18 billion page scans, and scans more than 1.2 billion devices for threats. Importantly, this data always goes through strict privacy and compliance boundaries before being used for security analysis.

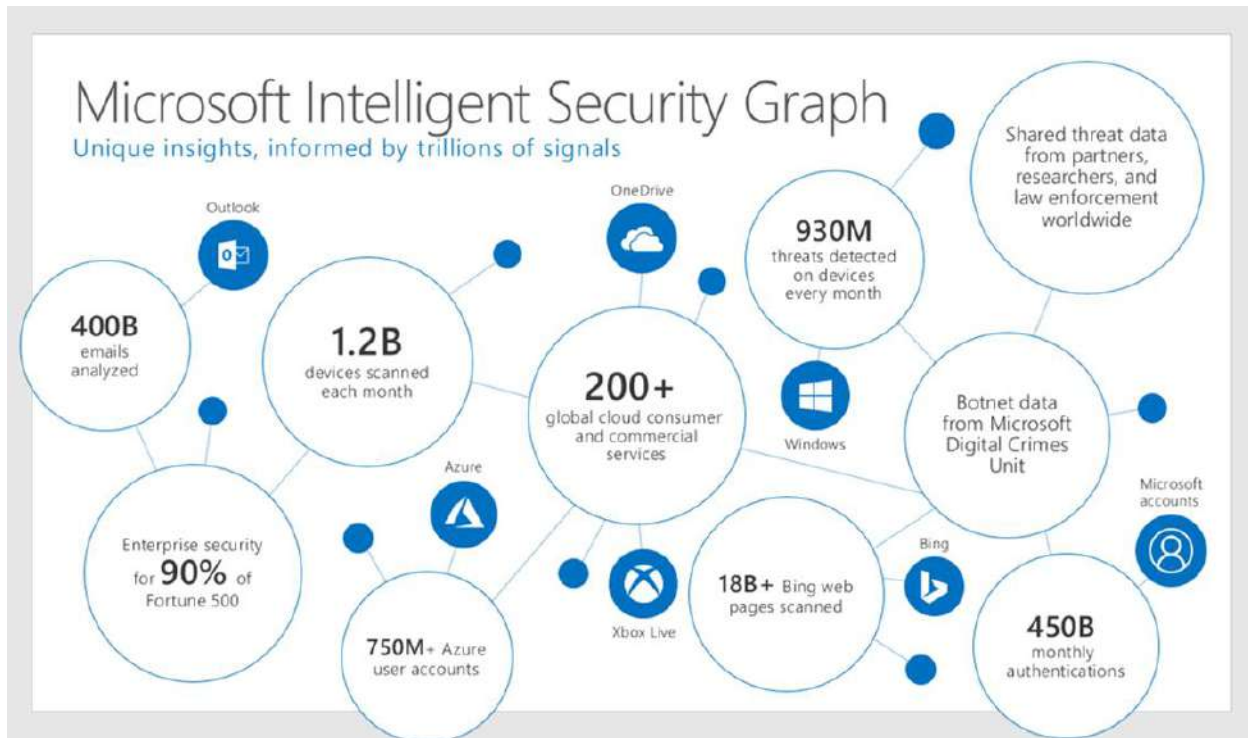


Figure 2. Microsoft global threat intelligence is one of the largest in industry

The Intelligent Security Graph provides an unparalleled view into the evolving threat landscape and enables rapid innovation to detect and respond to threats. Machine Learning models and Artificial Intelligence reason over vast security signals to identify vulnerabilities and threats. The [security API](#) provides a common gateway to share and act on security insights across the Microsoft platform and partner solutions. Insights from the Intelligent Security Graph power real-time threat protection in Microsoft products and services.

Azure customers benefit directly from the Intelligent Security Graph as Microsoft makes the vast threat telemetry and advanced analytics available in Azure services, including Azure Security Center.

[Azure Security Center](#) provides unified security management and advanced threat protection across hybrid cloud workloads. This is an essential service for government customers to limit their exposure to threats, protect cloud resources, respond to incidents, and improve their regulatory compliance posture.

With Azure Security Center, customers can:

- Monitor security across on-premises and cloud workloads
- Leverage advanced analytics and threat intelligence to detect attacks
- Use access and application controls to block malicious activity
- Find and fix vulnerabilities before they can be exploited
- Simplify investigation when responding to threats
- Apply policy to ensure compliance with security standards

To assist customers with Azure Security Center usage, Microsoft has published extensive [online documentation](#), as well as numerous blog posts covering specific security topics:

- [How Azure Security Center detects a Bitcoin mining attack](#)
- [How Azure Security Center detects DDoS attack using Cyber Threat Intelligence](#)
- [How Azure Security Center aids in detecting good applications being used maliciously](#)
- [How Azure Security Center unveils suspicious PowerShell attack](#)
- [How Azure Security Center helps reveal a Cyberattack](#)
- [How Azure Security Center helps analyze attacks using Investigation and Log Search](#)
- [Azure Security Center adds Context Alerts to aid threat investigation](#)
- [How Azure Security Center automates the detection of cyber-attack](#)
- [Heuristic DNS detections in Azure Security Center](#)
- [Detect the latest ransomware threat \(aka Bad Rabbit\) with Azure Security Center](#)
- [Petya ransomware prevention & detection in Azure Security Center](#)
- [Detecting in-memory attacks with Sysmon and Azure Security Center](#)
- [How Security Center and Log Analytics can be used for Threat Hunting](#)
- [How Azure Security Center helps detect attacks against your Linux machines](#)
- [Leverage Azure Security Center to detect when compromised Linux machines attack](#)

[Azure Monitor](#) helps customers maximize the availability and performance of applications by delivering a comprehensive solution for collecting, analyzing, and acting on telemetry from both cloud and on-premises environments. It helps customers understand how their applications are performing and proactively identifies issues affecting deployed applications and resources they depend on. Azure Monitor integrates the capabilities of Log Analytics and [Application Insights](#) that were previously branded as standalone services.

Azure Monitor collects data from each of the following tiers:

- **Application monitoring data:** Data about the performance and functionality of the code customers have written, regardless of its platform.
- **Guest OS monitoring data:** Data about the operating system on which customer application is running. The application could be running in Azure, another cloud, or on-premises.
- **Azure resource monitoring data:** Data about the operation of an Azure resource.
- **Azure subscription monitoring data:** Data about the operation and management of an Azure subscription, as well as data about the health and operation of Azure itself.
- **Azure tenant monitoring data:** Data about the operation of tenant-level Azure services, such as Azure Active Directory.

With Azure Monitor, customers can get a 360-degree view of their applications, infrastructure, and network with advanced analytics, dashboards, and visualization maps. Azure Monitor provides intelligent insights and enables better decisions with AI. Customers can analyze, correlate, and monitor data from various sources using a powerful query language and built-in machine learning constructs. Moreover, Azure Monitor provides out-of-the-box integration with popular DevOps, IT Service Management (ITSM), and Security Information and Event Management (SIEM) tools.

Microsoft provides additional services intended to help customers protect their Azure resources and implement effective governance strategies to enforce compliance with internal policies, including:

- [Web Application Firewall](#) – Protects customer web applications from common exploits and vulnerabilities as part of [Application Gateway](#). It comes preconfigured to handle the [OWASP](#) (Open Web Application Security Project) top 10 common vulnerabilities.
- [Azure DDoS Protection](#) – Provides extensive Distributed Denial of Service (DDoS) mitigation capability to help customers protect their Azure resources from attacks. Always-on traffic monitoring provides near real-time detection of a DDoS attack, with automatic mitigation of the attack as soon as it is detected. In combination with Web Application Firewall, DDoS Protection defends against a comprehensive set of network layer attacks, including SQL injection, cross-site scripting attacks, and session hijacks. Integrates with Azure Monitor for analytics and insight.
- [Azure Advisor](#) – Helps customers follow best practices to optimize their Azure deployments. It analyzes resource configurations and usage telemetry and then recommends solutions that can help customers improve the cost effectiveness, performance, high availability, and security of Azure resources.
- [Azure Governance](#) – Enables customers to implement and audit [policy-based management](#) for their Azure services, create compliant environments using [Azure Blueprints](#), and manage costs by gaining insights into cloud spend.

Microsoft has implemented extensive protection for the Azure cloud platform and made available a wide range of Azure services to help customers monitor and protect their provisioned cloud resources from attacks. Nonetheless, for certain types of workloads and data classifications, government customers expect to have full operational control over their environment and even operate in a fully disconnected mode. Azure Stack and Data Box Edge are Microsoft products that enable customers to provision private and hybrid cloud deployment models that can accommodate highly sensitive data, as described in the next section.

Private and hybrid cloud with Azure Stack and Data Box Edge

Azure Stack and Data Box Edge represent key enabling technologies that allow customers to process highly sensitive data using a private or hybrid cloud and pursue digital transformation leveraging Microsoft [intelligent cloud and intelligent edge](#) approach. For many government customers, enforcing data sovereignty, addressing custom compliance requirements, and applying maximum available protection to highly sensitive data are the primary driving factors behind these efforts.

[Azure Stack](#) is an integrated system of software and validated hardware that customers can purchase from Microsoft hardware partners, deploy in their own data center, and then operate entirely on their own or with the help from a managed service provider. With Azure Stack, the customer is always fully in control of access to their data. Azure Stack can accommodate up to 16 physical servers per Azure Stack scale unit. It represents an extension of Azure, enabling customers to provision a variety of IaaS and PaaS services and effectively bring multi-tenant cloud technology to on-premises and edge environments. Customers can run many types of VM instances, App Services, Containers (including Cognitive Services containers), Functions, Azure Monitor, Key Vault, Event Hubs, and other services while using the same development tools, APIs, and management processes they use in Azure. Azure Stack is not dependent on connectivity to Azure to run deployed applications and enable operations via local connectivity.

In addition to Azure Stack which is intended for on-premises deployment (e.g., in a data center), a ruggedized and field-deployable version called [Tactical Azure Stack](#) is also available to address tactical edge deployments for limited or no connectivity, fully mobile requirements, harsh conditions requiring military specification solutions, etc. Azure Stack can be operated completely disconnected from Azure or the Internet. Customers can run the next generation of AI-enabled hybrid applications where their data lives. For example, government agencies can rely on Azure Stack to bring a trained AI model to the edge and integrate it with their applications for low-latency intelligence, with no tool or process changes for local applications.

Azure and Azure Stack can help government customers unlock new hybrid use cases for customer-facing and internal line of business application, including edge and disconnected scenarios, cloud applications intended to meet data sovereignty and custom compliance requirements, and cloud applications deployed on-premises in customer data center. These may be mobile scenarios or fixed deployments within highly secure data center facilities. Figure 3 shows Azure Stack capabilities and key usage scenarios.

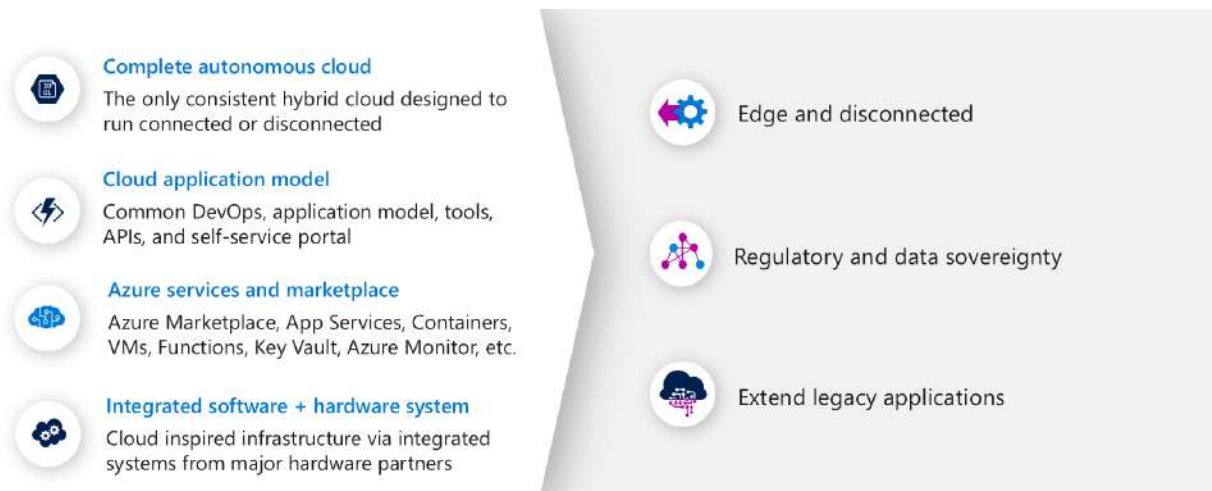


Figure 3. Azure Stack capabilities

Azure Stack brings the following value proposition for key scenarios shown in Figure 3:

- **Edge and disconnected solutions:** Address latency and connectivity requirements by processing data locally in Azure Stack and then aggregating in Azure for further analytics, with common application logic across both, connected or disconnected. Aircraft, ship, or truck-delivered, Azure Stack meets the tough demands of exploration, construction, agriculture, oil and gas, manufacturing, disaster response, government and military efforts in the most extreme conditions and remote locations. Government customers can leverage Azure Stack architecture for [edge and disconnected solutions](#), e.g., bring the next generation of AI-enabled hybrid applications to the edge where the data lives and integrate it with existing applications for low-latency intelligence.
- **Cloud applications to meet data sovereignty:** Deploy a single application differently depending on the country or region. Customers can develop and deploy applications in Azure, with full flexibility to deploy on-premises with Azure Stack based on the need to meet data sovereignty or custom compliance requirements. Customers can leverage Azure Stack architecture for [data](#)

[sovereignty](#), e.g., transmit data from Azure VNet to Azure Stack VNet over private connection and ultimately store data in SQL Server database running in a VM on Azure Stack. Government customers can use Azure Stack to accommodate even more restrictive requirements such as the need to deploy solutions in a **completely disconnected environment** managed by security-cleared, in-country personnel. These disconnected environments may not be permitted to connect to the Internet for any purpose because of the security classification they operate at.

- **Cloud application model on-premises:** Use Azure Stack to update and extend legacy applications and make them cloud ready. With App Service in Azure Stack, customers can create a web front-end to consume modern APIs with modern clients while taking advantage of consistent programming models and skills. Customers can leverage Azure Stack architecture for [legacy system modernization](#), e.g., apply a consistent DevOps process, Azure Web Apps, containers, serverless computing, and microservices architectures to modernize legacy applications while integrating and preserving legacy data in mainframe and core line of business systems.

Azure Stack uses either Azure Active Directory (Azure AD) or Active Directory Federation Services (AD FS) as an identity provider. Customers can use [Role Based Access Control](#) (RBAC) to grant system access to authorized users, groups, and services by assigning them roles at a subscription, resource group, or individual resource level. Each role defines the access level a user, group, or service has over Microsoft Azure Stack resources. Customers can store and manage their secrets including cryptographic keys on an external Hardware Secure Module (HSM) by using Thales [CipherTrust Cloud Key Manager](#) (available via the [Azure marketplace](#)), which allows customers to integrate an HSM with [Key Vault service running on Azure Stack](#).

[Azure Data Box Edge](#) is an AI-enabled edge computing device with network data transfer capabilities. It enables customers to pre-process data at the edge and also move data to Azure efficiently. Azure Data Box Edge uses advanced Field-Programmable Gate Array (FPGA) hardware natively integrated into the appliance to run Machine Learning algorithms at the edge efficiently. The size and portability allow customers to run Azure Data Box Edge as close to users, apps, and data as needed.

Key uses cases for Data Box Edge include:

- **Preprocess data:** Analyze data from on-premises or IoT devices to quickly obtain results while staying close to where data is generated. Data Box Edge transfers the full data set (or just the necessary subset of data when bandwidth is an issue) to the cloud to perform more advanced processing or deeper analytics. Preprocessing can be used to aggregate data, modify data (e.g., remove Personally Identifiable Information or other sensitive data), transfer data needed for deeper analytics in the cloud, and analyze and react to IoT events.
- **Inference Azure Machine Learning:** Inference is part of deep learning that takes place after model training, e.g., the prediction stage resulting from applying learned capability to new data. For example, it's the part that recognizes a vehicle in a target image after the model has been trained by processing many tagged vehicle images, often augmented by computer synthesized images (aka synthetics). With Data Box Edge, customers can run Machine Learning (ML) models to get results quickly and act on them before the data is sent to the cloud. The necessary subset of data (in case of bandwidth constraints) or the full data set is transferred to the cloud to continue to retrain and improve customer's ML models.

- **Transfer data over network to Azure:** Use Data Box Edge to transfer data to Azure to enable further compute and analytics or for archival purposes.

Being able to gather, discern, and distribute mission data is essential for making critical decisions. Tools that help process and transfer data directly at the edge make this possible. For example, Data Box Edge, with its light footprint and built-in hardware acceleration for ML inferencing, is useful to further the intelligence of forward-operating units or similar mission needs with AI solutions designed for the tactical edge. Data transfer from the field, which is traditionally complex and slow, is made seamless with the [Data Box](#) family of products.

These products unite the best of edge and cloud computing to unlock never-before-possible capabilities like synthetic mapping and ML model inferencing. From submarines to aircraft to remote bases, Azure Stack and Data Box Edge allow customers to harness the power of cloud at the edge.

Leveraging Azure in combination with Azure Stack and Data Box Edge, government customers can process confidential and sensitive data in secure isolated infrastructure within Azure public multi-tenant regions or highly sensitive data at the edge under the customer's full operational control. Customers deploying these types of workloads typically seek assurances from Microsoft that the underlying cloud platform security controls for which Microsoft is responsible are operating effectively. To address the needs of customers across regulated markets worldwide, Azure maintains a comprehensive compliance portfolio based on formal third-party certifications and other types of assurance documents to help customers meet their own compliance obligations.

Compliance and certifications

Azure has the broadest [compliance coverage](#) in the industry, including key independent certifications and attestations such as ISO 27001, ISO 27017, ISO 27018, ISO 22301, ISO 9001, ISO 20000-1, SOC 1/2/3, PCI DSS Level 1, HITRUST, CSA STAR Certification, CSA STAR Attestation, US FedRAMP Moderate, Australia IRAP Protected, Germany C5, Japan CS Mark Gold, Singapore MTCS Level 3, Spain ENS High, UK G-Cloud and Cyber Essentials Plus, and many more. Azure compliance portfolio includes more than 90 compliance offerings spanning globally applicable certifications, US Government specific programs, industry assurances, and regional / country specific offerings. Government customers can leverage these offerings when addressing their own compliance obligations across regulated industries and markets worldwide.

When deploying applications to Azure that are subject to regulatory compliance obligations, customers seek assurances that all cloud services comprising the solution be included in the cloud service provider's audit scope. Azure offers industry leading depth of compliance coverage judged by the number of cloud services in audit scope for each Azure certification. Customers can build and deploy realistic applications and benefit from extensive compliance coverage provided by Azure independent third-party audits.

Azure Stack also provides [compliance documentation](#) to help customers integrate Azure Stack into solutions that address regulated workloads. Customers can download the following Azure Stack compliance documents:

- PCI DSS assessment report produced by a third-party Qualified Security Assessor

Azure for Secure Worldwide Public Sector Cloud Adoption

- Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) assessment report, including Azure Stack control mapping to CCM domains and controls
- FedRAMP High System Security Plan (SSP) precompiled template to demonstrate how Azure Stack addresses applicable controls, Customer Responsibility Matrix for FedRAMP High baseline, and FedRAMP assessment reports produced by an independent Third-Party Assessor Organization (3PAO)

[Compliance Manager](#) is an online tool available from the Microsoft Service Trust Portal that customers can use to get insight into how Microsoft implements controls specified in leading standards such as ISO 27001, ISO 27018, NIST SP 800-53 Rev 4, and others. For example, Compliance Manager provides insight into Microsoft control implementation and test details for controls that are part of Microsoft responsibility. Moreover, customers can use this interactive tool to track progress for control implementations that they own. There is a 2-minute [video](#) to introduce key Compliance Manager features.

[Azure Blueprints](#) is a set of reference architectures with supporting deployment automation and guidance, security control mappings, and customer responsibility matrices to assist customers with deploying applications to Azure that meet established compliance standards. Customer responsibility matrices outline which controls are part of customer's responsibility, and they can be downloaded as part of the Azure Blueprint [documentation](#).

Azure compliance and certification resources are intended to help customers address their own compliance obligations with various regulations. Some governments across the world have already established cloud adoption mandates for their agencies and the corresponding regulation to facilitate cloud onboarding. However, there are many government customers that still operate traditional on-premises datacenters and are in the process of formulating their cloud adoption strategy. The following section covers key enabling cloud technologies such as virtualization, logical isolation, and multi-tenancy to help customers prepare for the cloud. It addresses common security and isolation concerns pertinent to worldwide public sector customers pursuing cloud adoption.

Logical isolation considerations

A multi-tenant cloud platform implies that multiple customer applications and data are stored on the same physical hardware. Azure uses [logical isolation](#) to segregate each customer's applications and data from those of others. This approach provides the scale and economic benefits of multi-tenant cloud services while rigorously helping enforce controls designed to keep customers from accessing one another's data or applications.

Identity and access

[Azure Active Directory](#) (AD) is an identity repository and cloud service that provides authentication, authorization, and access control for an organization's users, groups, and objects. Azure AD can be used as a standalone cloud directory or as an integrated solution with existing on-premises Active Directory to enable key enterprise features such as directory synchronization and single sign-on.

Each Azure subscription has an Azure AD. Using [Role-Based Access Control](#) (RBAC), users, groups, and applications from that directory can be granted access to resources in the Azure subscription. For example, a storage account can be placed in a resource group to control access to that specific storage

account using Azure AD. In this manner, only specific users can be given the ability to access the Storage Account Key, which controls access to storage.

All data in Azure irrespective of the type or storage location is associated with a subscription. Customers may have multiple subscriptions and multiple deployments/tenants within each subscription; however, the account used to create and manage the subscription has full rights over any data stored in it. Authentication to the Management Portal is performed through Azure AD using an identity created either in Azure AD or federated with an on-premises Active Directory. The identity and access stack helps enforce isolation among subscriptions, including limiting access to resources within a subscription only to authorized users. The concept of logical isolation is also deeply embedded by design across Azure services such as compute, storage, and networking.

Compute isolation

- Customer VMs do not have access to a physical host server
- By design, customer VMs cannot generate spoofed traffic or receive traffic not addressed to them, direct traffic to protected infrastructure endpoints, or send/receive inappropriate broadcast traffic
- Azure provides VM instances that are isolated to hardware dedicated to a single customer

Microsoft Azure compute platform is based on [machine virtualization](#). This means that customer code – whether it's deployed in a PaaS Worker Role or an IaaS Virtual Machine – executes in a Windows Server Hyper-V virtual machine. On each Azure node, there is a Hypervisor that runs directly over the hardware and divides the node into a variable number of Guest Virtual Machines (VMs). Each node also has one special Root VM, which runs the Host OS. Isolation of the Root VM from the Guest VMs and the Guest VMs from one another is a key concept in Azure security architecture that forms the basis of Azure [compute isolation](#).

Dedicated hosts: Azure E64is v3, E64i v3, GS5, G5, DS15 v2, and D15 v2 Virtual Machine instances allow customer workloads to be deployed on [hardware dedicated to a single customer](#). Utilizing isolated VMs essentially guarantees that customer VM will be the only one running on that specific server node. Customers can also choose to further subdivide the resources on these isolated VMs by using [Azure support for nested Virtual Machines](#).

Storage isolation

- By design, a customer cannot read data belonging to another customer
- Data deleted by a customer is not accessible to other customers
- Data destruction and disposal follows NIST SP 800-88 R1
- Azure administrators do not have default access to Customer Data

Microsoft Azure separates customer VM-based computation resources from storage as part of its [fundamental design](#). The separation allows computation and storage to scale independently, making it easier to provide multi-tenancy and isolation.

Consequently, Azure Storage runs on separate hardware with no network connectivity to Azure Compute except logically. All requests run over HTTP or HTTPS based on customer's choice. Storage is allocated sparsely. This means that when a virtual disk is created, disk space is not allocated for its entire capacity. Instead, a table is created that maps addresses on the virtual disk to areas on the physical disk and that table is initially empty. The first time a customer writes data on the virtual disk, space on the physical disk is allocated and a pointer to it is placed in the table.

When the customer deletes a blob or table entity, it will immediately get deleted from the index used to locate and access the data on the primary location, and then the deletion is done asynchronously at the geo-replicated copy of the data (for customers who provisioned [geo-redundant storage](#)). At the primary location, the customer can immediately try to access the blob or entity, and they won't find it in their index, since Azure provides strong consistency for the delete. So, the customer can verify directly that the data has been deleted.

By design, a customer cannot read deleted data of another customer. If anyone tries to read a region on virtual disk that they have not yet written to, physical space will not have been allocated for that region and therefore only zeroes would be returned.

Conceptually, this applies regardless of the software that keeps track of reads and writes. In the case of [Azure SQL Database](#), it is the SQL Database software that does this enforcement. In the case of Azure Storage, it is the Azure Storage software. In the case of non-durable drives of a VM, it is the VHD handling code of the host OS. Since customer software only addresses virtual disks (the mapping from virtual to physical address takes place outside of the customer VM), there is no way to express a request to read from or write to a physical address that is allocated to a different customer or a physical address that is free.

If a disk drive used for storage suffers a hardware failure, it is securely [erased or destroyed](#) before Microsoft returns it to the manufacturer for replacement or repair. The data on the drive is overwritten to ensure that the data cannot be recovered by any means. When such devices are decommissioned, Microsoft follows the [NIST SP 800-88 R1](#) disposal process with data classification aligned to FIPS 199 Moderate. Magnetic, electronic, or optical media are purged or destroyed in accordance with the requirements established in NIST SP 800-88 R1 where the terms are defined as follows:

- Purge: "a media sanitization process that protects the confidentiality of information against a laboratory attack" which involves "resources and knowledge to use nonstandard systems to conduct data recovery attempts on media outside their normal operating environment" using "signal processing equipment and specially trained personnel." Note: For hard disk drives (including ATA, SCSI, SATA, SAS, etc.) a firmware-level secure-erase command (single-pass) is acceptable, or a software-level three-pass overwrite and verification (ones, zeros, random) of the entire physical media including recovery areas, if any. For solid state disks (SSD), a firmware-level secure-erase command is necessary.
- Destroy: "a variety of methods, including disintegration, incineration, pulverizing, shredding, and melting" after which the media "cannot be reused as originally intended."

Microsoft engineers [do not have default access](#) to Customer Data. As described previously, they are granted access, under management oversight, only when necessary. Controls for the protection of

customer secrets are audited on a regular basis as part of existing Azure audits, including SOC 2 Type 2 and FedRAMP. Customers also have several options for [encrypting their data at rest](#), including using [Azure Key Vault](#) to manage their own encryption keys for [Azure Storage](#) and [Azure SQL Database](#). Encryption keys are stored in tamper-resistant Hardware Security Modules that are FIPS 140-2 Level 2 validated.

Networking isolation

- Private IP addresses are isolated from other customers
- Firewalls limiting traffic to VMs
- No local accounts on PaaS VMs for remote logins
- Encrypted communications

The logical isolation of customer infrastructure in a public cloud is fundamental to maintaining security ([Palekar, 2015](#)). The overarching principle for a virtualized solution is to allow only connections and communications that are necessary for that virtualized solution to operate, blocking all other ports and connections by default. Azure [Virtual Network](#) (VNet) helps ensure that each customer's private network traffic is logically isolated from traffic belonging to other customers.

Network access to VMs is limited by packet filtering at the network edge, at load balancers, and at the Host OS level. Customers can, in addition, configure their host firewalls to further limit connectivity. Microsoft allows customers to specify for each listening port whether connections are accepted from the Internet or only from role instances within the same cloud service or VNet. For each VM, the Fabric Controller composes (and keeps up to date) a list of IP addresses of VMs in the same cloud service. This list of IP addresses is used by the Fabric Agent to program the packet filters to only allow intra-service or virtual network communication to those IP addresses. Some PaaS roles (e.g., Web role) are normally allowed to initiate communication to Internet addresses. This enables them to communicate with the Internet and send traffic to any other role that can be reached from the Internet.

Azure provides network isolation for each deployment. Using input endpoints, customers decide which ports can be accessed from the Internet.

- Traffic between VMs always traverses through trusted packet filters.
 - Protocols such as Address Resolution Protocol (ARP), Dynamic Host Configuration Protocol (DHCP), and other OSI Layer-2 traffic from a VM are controlled using rate-limiting and anti-spoofing protection.
 - VMs cannot capture any traffic on the network that is not destined to them.
- Customer VMs cannot send traffic to Azure private interfaces and infrastructure services, or to other customers' VMs. Customer VMs can only communicate with other VMs owned or controlled by the same customer and with Azure infrastructure service endpoints meant for public communications.
- When customers put VMs on a virtual private network, those VMs get their own address spaces that are completely invisible, and hence, not reachable from VMs outside of a deployment or virtual network (unless configured to be visible via public IP addresses). Customer environments

are open only through the ports that customers specify for public access; if the VM is defined to have a public IP address, then all ports are open for public access.

For PaaS Web and Worker roles, remote access is not permitted by default. It is possible for customers to enable Remote Desktop Protocol (RDP) access as an explicit option. For IaaS VMs created using the Azure Management Portal, RDP and remote PowerShell ports are opened by default; however, port numbers are assigned randomly. For IaaS VMs created via PowerShell, RDP and remote PowerShell ports must be opened explicitly. If the administrator chooses to keep the RDP and remote PowerShell ports open to the Internet, the account allowed to create RDP and PowerShell connections should be secured with a strong password.

The cumulative effect of these restrictions is that each cloud service acts as though it were on an isolated network where VMs within the cloud service can communicate with one another, identifying one another by their source IP addresses with confidence that no other parties can impersonate their peer VMs. They can also be configured to accept incoming connections from the Internet over specific ports and protocols.

Virtual Network also provides a means for Azure VMs to act as part of a customer’s internal (on-premises) network. It expands the nature of intranet connectivity beyond a single cloud service to include any set of internal addresses of other cloud services on Azure or other machines on a customer’s own network (presumably behind the customer’s datacenter firewall). With VNet, customers choose the address ranges of non-globally-routable IP addresses to be assigned to the VMs so that they will not collide with addresses the customer is using elsewhere. A cryptographically protected “tunnel” is established between Azure and the customer’s internal network, allowing the VM to connect to the customer’s back-end resources as though it was directly on that network. Customers have options to securely connect to a Virtual Network – choose an [IPsec protected VPN](#) (e.g., point-to-site VPN or site-to-site VPN) or a private connection by using Azure [ExpressRoute](#).

When migrating to the cloud, government customers accustomed to traditional on-premises data center deployment will usually conduct a risk assessment to gauge their threat exposure and formulate mitigating measures. In many of these instances, security considerations for traditional on-premises data center deployment tend to be well established and understood whereas the corresponding cloud options tend to be new. The next section is intended to help government customers with this comparison.

Physical versus logical security considerations

Table 1 provides a summary of key security considerations for physically isolated on-premises deployments (e.g., bare metal) versus logically isolated cloud-based deployments (e.g., Azure). It’s useful to review these considerations prior to examining risks identified to be specific to shared cloud environments.

Table 1: Key security considerations for physical versus logical isolation

Security Consideration	On-Premises	Azure
Firewalls, networking	• Physical network enforcement (switches, etc.)	• Physical network enforcement (switches, etc.)

Security Consideration	On-Premises	Azure
	<ul style="list-style-type: none"> Physical host-based firewall can be manipulated by compromised application 2 layers of enforcement 	<ul style="list-style-type: none"> Hyper-V host virtual network switch enforcement cannot be changed from inside VM VM host-based firewall can be manipulated by compromised application 3 layers of enforcement
Attack surface area	<ul style="list-style-type: none"> Large hardware attack surface exposed to complex workloads, enables firmware based advanced persistent threat (APT) 	<ul style="list-style-type: none"> Hardware not directly exposed to VM, no potential for APT to persist in firmware from VM Small software-based Hyper-V attack surface area with low historical bug counts exposed to VM
Side channel attacks	<ul style="list-style-type: none"> Side channel attacks may be a factor, although reduced vs. shared hardware 	<ul style="list-style-type: none"> Side channel attacks assume control over VM placement across applications; may not be practical in large cloud service
Patching	<ul style="list-style-type: none"> Varied effective patching policy applied across host systems Highly varied/fragile updating for hardware and firmware 	<ul style="list-style-type: none"> Uniform patching policy applied across host and VMs
Security analytics	<ul style="list-style-type: none"> Security analytics dependent on host-based security solutions, which assume host/security software has not been compromised 	<ul style="list-style-type: none"> Outside VM (hypervisor based) forensics/snapshot capability allows assessment of potentially compromised workloads
Security policy	<ul style="list-style-type: none"> Security policy verification (patch scanning, vulnerability scanning, etc.) subject to tampering by compromised host Inconsistent security policy applied across customer entities 	<ul style="list-style-type: none"> Outside VM verification of security policies Possible to enforce uniform security policies across customer entities
Logging and monitoring	<ul style="list-style-type: none"> Varied logging and security analytics solutions 	<ul style="list-style-type: none"> Common Azure platform logging and security analytics solutions Most existing on-premises / varied logging and security analytics solutions also work
Malicious insider	<ul style="list-style-type: none"> Persistent threat caused by system admins having elevated access rights typically for the duration of employment 	<ul style="list-style-type: none"> Greatly reduced threat because admins have no default access rights

Listed below are key risks that are unique to shared cloud environments that may need to be addressed when accommodating data and workloads common to worldwide government customers.

Exploitation of vulnerabilities in virtualization technologies

Compared to traditional on-premises hosted systems, Azure provides a greatly **reduced attack surface** by using a locked-down Windows Server core for the Host OS layered over the Hypervisor. Moreover, by default, guest PaaS VMs do not have any user accounts to accept incoming remote connections and the default Windows administrator account is disabled. Customer software in PaaS VMs is restricted by default to running under a low-privilege account, which helps protect customer's service from attacks by its own end users. These permissions can be modified by customers, and they can also choose to configure their VMs to allow remote administrative access.

PaaS VMs offer more advanced **protection against persistent malware** infections than traditional physical server solutions, which if compromised by an attacker can be difficult to clean, even after the vulnerability is corrected. The attacker may have left behind modifications to the system that allow re-entry, and it is a challenge to find all such changes. In the extreme case, the system must be reimaged from scratch with all software reinstalled, sometimes resulting in the loss of application data. With PaaS VMs, reimaging is a routine part of operations, and it can help clean out intrusions that have not even been detected. This makes it much more difficult for a compromise to persist.

When VMs belonging to different customers are running on the same physical server, it is the Hypervisor's job to ensure that they cannot learn anything important about what the other customer's VMs are doing. As described previously, blocking unauthorized direct communication is straightforward; however, there are subtle effects where one customer might be able to characterize the work being done by another customer. The most important of these are timing effects when different VMs are competing for the same resources. By carefully comparing operations counts on CPUs with elapsed time, a VM can learn something about what other VMs on the same server are doing. Known as **side-channel attacks**, these exploits have received plenty of attention in the academic press where researchers have been seeking to learn much more specific information about what is going on in a peer VM. Of particular interest are efforts to learn the cryptographic keys of a peer VM by measuring the timing of certain memory accesses and inferring which cache lines the victim's VM is reading and updating. Under controlled conditions with VMs using hyper-threading, successful attacks have been demonstrated against commercially available implementations of cryptographic algorithms. There are several mitigations in Azure that reduce the risk of such an attack:

- The standard Azure cryptographic libraries have been designed to resist such attacks by not having cache access patterns depend on the cryptographic keys being used.
- Azure uses an advanced VM host placement algorithm that is highly sophisticated and nearly impossible to predict, which helps reduce the chances of adversary controlled VM being placed on the same host as the target VM.
- All Azure servers have at least 8 physical cores and some have substantially more. Increasing the number of cores that share the load placed by various VMs adds noise to an already weak signal.

Potential for providing back door connections and CSP privileged user access to customer's systems and data (insider threat)

No default access rights and Just-in-Time (JIT) access provisions reduce greatly the risks associated with traditional on-premises administrator elevated access rights that typically persist throughout the

duration of employment. Microsoft makes it considerably more difficult for malicious insiders to tamper with customer applications and data.

Conceptual architecture

Figure 4 shows a conceptual architecture using products and services that support various data classifications. Azure public multi-tenant cloud is the underlying cloud platform that makes this solution possible. Customers can augment Azure with on-premises and edge products such as Azure Stack and Data Box Edge to accommodate critical workloads over which customers seek increased or exclusive operational control. For example, Azure Stack is intended for on-premises deployment in customer-owned data center where the customer has full control over service connectivity. Moreover, Azure Stack can be deployed to address tactical edge deployments for limited or no connectivity, including fully mobile scenarios.

For classified workloads, customers can provision key enabling Azure services to secure target workloads while mitigating identified risks.

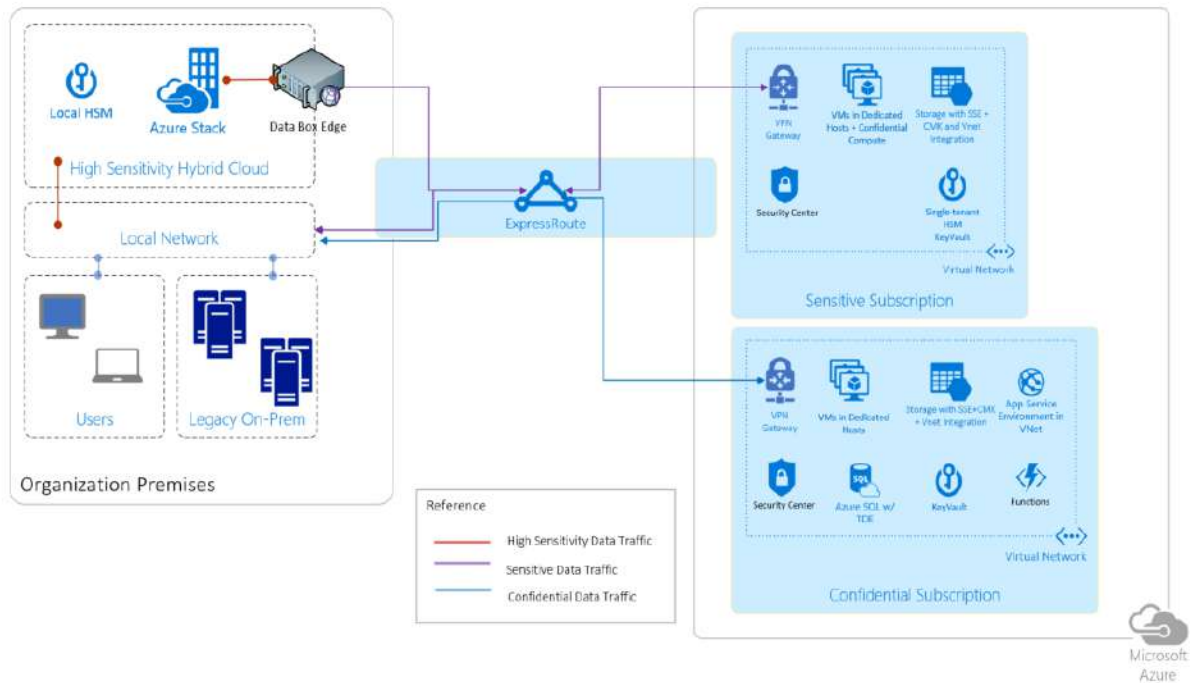


Figure 4. Conceptual architecture for classified workloads

Azure, in combination with [Azure Stack](#) and [Data Box Edge](#), can accommodate private and hybrid cloud scenarios and it is suitable for a wide range of government workloads involving both unclassified and classified data. The following data classification taxonomy is used in this document:

- Confidential
- Sensitive
- Highly Sensitive

Similar data classification schemes exist in many countries. For example, Confidential data can also be referred to as Protected or Restricted data whereas Sensitive data is frequently referred to as Secret data. For Highly Sensitive data, customers can deploy Azure Stack, which can operate fully disconnected from Azure and the Internet. [Tactical Azure Stack](#) is also available to address tactical edge deployments for limited or no connectivity, fully mobile requirements, harsh conditions requiring military specification solutions, etc. Figure 5 depicts key enabling services that customers can provision to accommodate a variety of workloads in Azure.

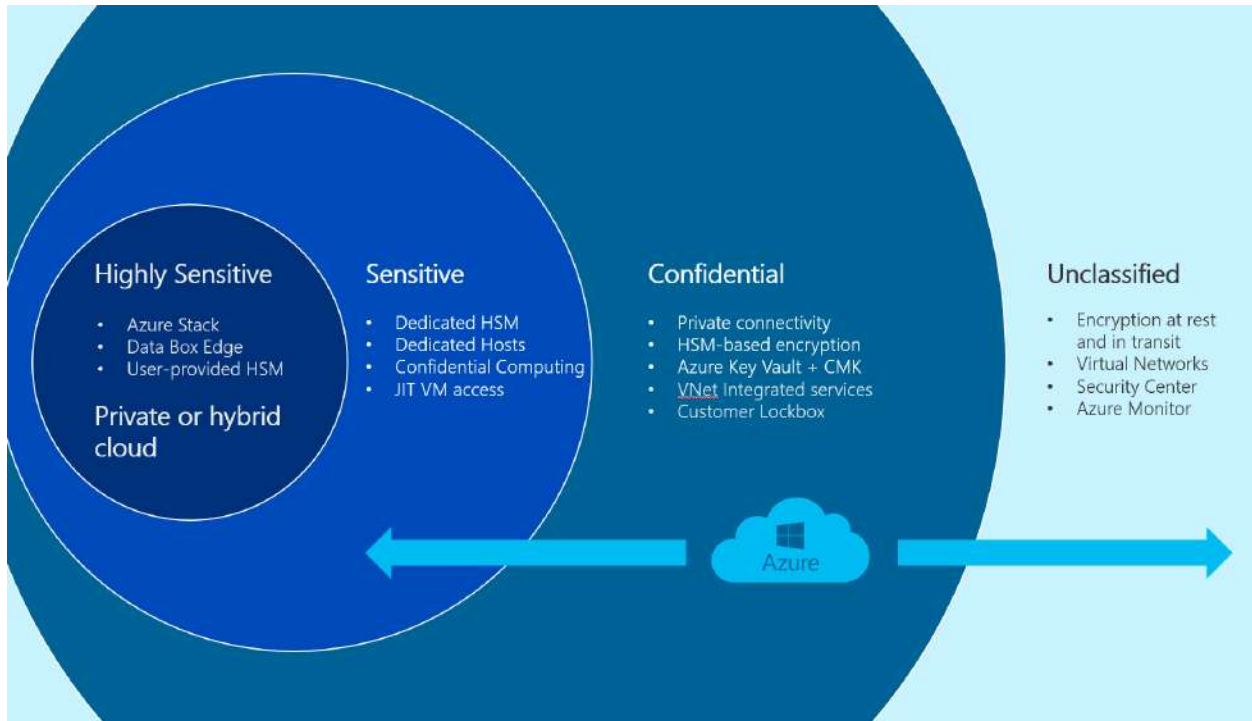


Figure 5. Azure support for various data classifications

Azure for Confidential data

Listed below are key enabling technologies and services that customers may find helpful when deploying Confidential data and workloads in Azure:

- All recommended technologies used for Unclassified data, especially services such as [Virtual Network](#), [Azure Security Center](#), and [Azure Monitor](#).
- Public IP addresses are disabled allowing only traffic through private connections, including [ExpressRoute](#) and [Virtual Private Networks](#) (VPNs).
- Data encryption at rest and in transit is recommended with Customer Managed Keys (CMK) stored in tamper-resistant [Azure Key Vault](#) Hardware Security Modules (HSMs), currently with FIPS 140-2 Level 2 overall validation.
- Only services that support [VNet Integration](#) options are enabled. Azure VNet enables customers to place Azure resources in a non-internet routable network, which can then be connected to customer on-premises network using VPN technologies. VNet Integration gives web apps access to resources in the virtual network.

Azure for Secure Worldwide Public Sector Cloud Adoption

- [Azure Customer Lockbox](#) enables customers to Approve/Deny elevated access requests for Customer Data in support scenarios. It's an extension of the Just-in-Time (JIT) workflow that comes with full audit logging enabled.

Using Azure public multi-tenant cloud capabilities, customers can achieve the level of isolation, security, and confidence required to store Confidential data. Customers should leverage Azure Security Center and Azure Monitor to gain visibility into their Azure environments including security posture.

Azure for Sensitive data

Listed below are key enabling technologies and services that customers may find helpful when deploying Sensitive data and workloads in Azure:

- All recommended technologies used for Confidential data.
- Use [Dedicated HSMs](#) (aka single-tenant HSMs), which have FIPS 140-2 Level 3 validation. Customers have full administrative and cryptographic control over Dedicated HSMs that are provisioned directly on customer's virtual network (can also connect to on-premises infrastructure via VPN).
- Dedicated Hosts via Azure E64is v3, E64i v3, GS5, G5, DS15 v2, and D15 v2 Virtual Machine instances allow customers to be deployed on [hardware dedicated to a single customer](#). Utilizing isolated VMs guarantees that customer VM will be the only one running on that specific server node.
- Accelerated FPGA networking based on [Azure SmartNICs](#) enables customers to offload host networking to dedicated hardware, enabling tunneling for VNets, security, and load balancing. Offloading network traffic to a dedicated chip guards against side-channel attacks on the main CPU.
- [Confidential Computing](#) offers encryption of data while in use, ensuring that data is always under customer control. Data is protected inside a Trusted Execution Environment (TEE) and there is no way to view data or operations from outside the enclave.
- [Just-in-time \(JIT\) virtual machine \(VM\) access](#) can be used to lock down inbound traffic to Azure VMs by creating Network Security Group (NSG) rules. Customer selects ports on the VM to which inbound traffic will be locked down and when a user requests access to a VM, Azure Security Center checks that the user has proper Role Based Access Control (RBAC) permissions.

To accommodate Sensitive data in the Azure public multi-tenant cloud, customers can deploy additional technologies and services on top of those used for Confidential data and limit provisioned services to those that provide sufficient isolation. These services offer isolation options at run time and support data encryption at rest using customer managed keys in dedicated single tenant HSMs that are solely under customer control.

Azure for Highly Sensitive data

Listed below are key enabling products that customers may find helpful when deploying Highly Sensitive data and workloads in Azure:

- All recommended technologies used for Sensitive data.
- [Azure Stack](#) enables customers to run workloads using the same architecture and APIs as in Azure while having a physically isolated network for their highest classification data.

Azure for Secure Worldwide Public Sector Cloud Adoption

- Azure [Data Box Edge](#) allows the storage and processing of highest classification data but also enables customers to upload resulting information or models directly to Azure. This approach creates a path for information sharing between domains that makes it easier and more secure.
- User-provided Hardware Security Modules (HSMs) allow customers to store their encryption keys and other secrets in HSMs deployed on-premises and controlled solely by customers.

Accommodating Highly Sensitive data will likely require a disconnected environment, which is what Azure Stack provides. Even though “air-gapped” networks do not necessarily increase security, many governments may be reluctant to store data with this classification in an Internet connected environment.

Select workloads and use cases

This section provides an overview of select use cases that showcase Azure capabilities for workloads that might be of interest to worldwide governments. In terms of capabilities, Azure is presented via a combination of public multi-tenant cloud and on-premises + edge capabilities provided by [Azure Stack](#) and [Data Box Edge](#).

Processing highly sensitive or regulated data on Azure Stack

Microsoft provides Azure Stack as an on-premises, cloud-consistent experience for customers who do not have the ability to directly connect to the Internet, or where certain workload types are required to be hosted in-country due to law, compliance, or sentiment. Azure Stack offers IaaS and PaaS services and shares the same APIs as the public Azure cloud. Azure Stack is available in scale units of 4, 8, and 16 servers in a single server rack, as well as 4 servers in a military-specification, ruggedized set of transit cases, or multiple racks in a modular data center configuration.

Azure Stack is a solution for customers who operate in scenarios where:

- Microsoft does not have an in-country cloud presence and therefore cannot meet data sovereignty requirements
- For compliance reasons, the customer cannot connect their network to the public Internet
- For geo-political or security reasons Microsoft cannot offer connectivity to other Microsoft clouds
- For geo-political or security reasons, the host organization may require cloud management by non-Microsoft entities, or in-country by security-cleared personnel
- Cloud management would pose significant risk to the physical well-being of Microsoft resources operating the environment

For the majority of these customers, Microsoft and its partners offer a customer-managed, Azure Stack-based private cloud appliance on field-deployable hardware from major vendors such as Dell EMC, HP, Lenovo, and Cisco. Azure Stack is manufactured, configured, and deployed by the hardware vendor, and can be ruggedized and security-hardened to meet a broad range of environmental and compliance standards, including the ability to withstand transport by aircraft, ship, or truck, and deployment into colocation, mobile, or modular data centers. Azure Stack can be used in exploration, construction, agriculture, oil and gas, manufacturing, disaster response, government and military efforts in hospitable or the most extreme conditions and remote locations. Azure Stack allows customers the full autonomy

to monitor, manage and provision their own private cloud resources while meeting their connectivity, compliance, and ruggedization requirements.

Machine Learning model training

[Artificial Intelligence](#) (AI) holds tremendous potential for governments. [Machine Learning](#) (ML) is a data science technique that allows computers to learn to use existing data, without being explicitly programmed, to forecast future behaviors, outcomes, and trends. Moreover, ML [techniques](#) can discover patterns, anomalies, and predictions that can help governments in their missions. As technical barriers continue to fall, decision-makers face the opportunity to develop and explore transformative AI applications. There are five main vectors that can make it easier, faster, and cheaper to adopt ML:

- Unsupervised learning
- Reducing need for training data
- Accelerated learning
- Transparency of outcome
- Deploying closer to where data lives

In the following sections, we expand on areas that can help government agencies with some of the above vectors.

IoT analytics

In recent years, we have been witnessing massive proliferation of Internet of Things (IoT) devices and sensors. In almost all cases, these sensors gather signals and data from the environments and conditions they're designed for. The spectrum of capabilities for IoT sensors expands from measuring the level of moisture in soil all the way to gathering intelligence at 18,000 feet altitude. The high number of use cases imposes the necessity of applying data-analysis tools and procedures to realize value from the huge volumes of gathered data by IoT devices.

Governments are increasingly employing IoT devices for their missions, which could include maintenance predictions, borders monitoring, weather stations, smart meters, and field operations. In many cases, the data is often analyzed and inferred from where it's gathered. The main challenges of IoT analytics are: (1) large amount of data from independent sources, (2) analytics at the edge and often in disconnected scenarios, and (3) data and analysis aggregation.

With innovative solutions such as [IoT Hub](#) and [Data Box Edge](#), Azure services are well positioned to help governments with these challenges.

Precision Agriculture with Farm Beats

Agriculture plays a vital role in most economies worldwide. In the US, over 70 per cent of the rural households depend on agriculture as it contributes about 17% to the total GDP and provides employment to over 60% of the population. In project [Farm Beats](#), we gather a lot of data from farms that we couldn't get before, and then by applying AI and ML algorithms we are able to turn this data into actionable insights for farmers. We call this technique data-driven farming. What we mean by data-driven farming is the ability to map every farm and overlay it with data. For example, what is the soil moisture level 6-inches below soil, what is the soil temperature 6 inches below soil, etc. These maps can then enable techniques, such as Precision Agriculture, which has been shown to improve yield, reduce costs, and benefit the environment. Despite the fact the Precision Agriculture as a technique

was proposed more than 30 years ago, it hasn't taken off. The biggest reason is the inability to capture a lot of data from farms to accurately represent the conditions in the farm. Our goal as part of the Farm Beats project is to be able to accurately construct these precision maps at a fraction of the cost.

Unleashing the power of analytics with Synthetic Data

Synthetic Data is data that is artificially created rather than being generated by actual events. It is often created with the help of computer algorithms and it is used for a wide range of activities, including usage as test data for new products and tools, as well as for ML models validation and improvements. Synthetic Data can meet very specific needs or conditions that are not available in existing real data. For governments, the nature of Synthetic Data removes many barriers and helps data scientists with privacy concerns, accelerated learning, and data volume reduction needed for the same outcome. The main benefits of Synthetic Data are:

- Overcoming restrictions: Real data may have usage constraints due to privacy rules or other regulations. Synthetic Data can replicate all important statistical properties of real data without exposing real data.
- Scarcity: Providing data where real data does not exist for a given event.
- Precision: Synthetic Data is perfectly labeled.
- Quality: The quality of Synthetic Data can be precisely measured to fit the mission conditions.

Synthetic Data can exist in several forms, including text, audio, video, and hybrid.

Knowledge mining

The exponential growth of unstructured data gathering in recent years has created many analytical problems for government agencies. This problem intensifies when data sets come from diverse sources such as text, audio, video, imaging, etc. [Knowledge mining](#) is the process of discovering useful knowledge from a collection of diverse data sources. This widely used data mining technique is a process that includes data preparation and selection, data cleansing, incorporation of prior knowledge on data sets, and interpretation of accurate solutions from the observed results. This process has proven to be very useful for large volumes of data in different government agencies.

For instance, captured data from the field often includes documents, pamphlets, letters, spreadsheets, propaganda, videos, and audio files across many disparate structured and unstructured formats. Buried within the data are [actionable insights](#) that can enhance effective and timely response to crisis and drive decisions. The objective of knowledge mining is to enable decisions that are better, faster, and more humane by implementing proven commercial algorithm-based technologies.

Scenarios for Confidential Computing

Security is a key driver accelerating the adoption of cloud computing, but it's also a major concern when customers are moving extremely sensitive IP and data to the cloud.

Microsoft Azure provides broad capabilities to secure data at rest and in transit, but sometimes the requirement is also to protect data from threats as it's being processed. Microsoft Azure [Confidential Computing](#) is designed to meet this scenario through the use of Trusted Execution Environments (TEEs) or encryption mechanisms to protect customer data while in use. TEEs are hardware or software implementations that safeguard data being processed from access outside the TEE. The hardware provides a protected container by securing a portion of the processor and memory. Only authorized

code is permitted to run and to access data, so code and data are protected against viewing and modification from outside of TEE.

TEEs can directly address these types of scenarios. For example, consider the scenario where data coming from a public or unclassified source needs to be matched with data from a highly sensitive source. Using Confidential Computing can enable that matching to occur in the public cloud while protecting the highly sensitive data from disclosure. This is a common circumstance in highly sensitive national security and law enforcement scenarios.

A second scenario involves data coming from multiple sources that needs to be analyzed together, even though none of the sources have the authority to see the data. Each individual provider encrypts the data they provide and only within the TEE is that data decrypted. As such, no external party and even none of the providers are able to see the combined data set. This is a particularly valuable capability for secondary use of healthcare data.

Frequently asked questions

This section addresses common customer questions related to Azure public, private, and hybrid cloud deployment models.

Data residency and data sovereignty

- **Data location:** How does Microsoft keep data within a specific country's boundaries? In what cases does data leave? What data attributes leave? **Answer:** Microsoft provides [strong customer commitments](#) regarding cloud services data residency and transfer policies:
 - **Data storage for regional services:** Most Azure services are deployed regionally and enable the customer to specify the region into which the service will be deployed, e.g., Europe. Microsoft will not store Customer Data outside the customer-specified Geo except for [Cloud Services](#), [Cognitive Services](#), [Azure Databricks](#), and Preview services as described on the [data location page](#). This commitment helps ensure that Customer Data stored in a given region will remain in the corresponding Geo and will not be moved to another Geo for the majority of regional services, including Storage, SQL Database, Virtual Machines, etc.
 - **Data storage for non-regional services:** Certain Azure services do not enable the customer to specify the region where the services will be deployed as described on the [data location page](#). For a complete list of non-regional services, customers should see [Services by Region](#).
- **Sovereign cloud deployment:** Why doesn't Microsoft deploy a sovereign, physically isolated cloud instance in every country that requests it? **Answer:** Physical isolation or "air gapping", as a strategy, is diametrically opposed to the strategy of hyperscale cloud. The value proposition of the cloud, rapid feature growth, resiliency, and cost-effective operation, break down when the cloud is fragmented and physically isolated. These strategic challenges compound with each additional sovereign cloud or fragmentation within a sovereign cloud.
- **Sovereign cloud customer options:** How can Microsoft support governments who need to operate cloud services completely in-country by local security-cleared personnel? What options does Microsoft have for cloud services operated entirely on-premises within customer owned datacenter where government employees exercise sole operational and data access control? **Answer:** Government customers can use [Azure Stack](#) to deploy a private cloud on-premises managed by the customer's own security-cleared, in-country personnel. Customers can run

many types of VM instances, App Services, Containers (including Cognitive Services containers), Functions, Azure Monitor, Key Vault, Event Hubs, and other services while using the same development tools, APIs, and management processes they use in Azure. With Azure Stack, customers have sole control of their data, including storage, processing, transmission, and remote access.

- **Local jurisdiction:** Is Microsoft subject to local country jurisdiction based on the availability of Azure public cloud service? **Answer:** Yes, Microsoft must comply with applicable local laws; however, government requests for Customer Data must also comply with applicable laws. A subpoena or its local equivalent is required to request non-content data and a warrant, court order, or its local equivalent is required for content data. Every year, Microsoft rejects a number of law enforcement requests for Customer Data. Challenges to government requests can take many forms. In many of these cases, Microsoft simply informs the requesting government that it is unable to disclose the requested information and explains the reason for rejecting the request. Where appropriate, Microsoft challenges requests in court. Customers should review the [Law Enforcement Requests Report](#) that Microsoft publishes twice a year. For example, in the first half of 2018, Microsoft received 50 requests from law enforcement for accounts associated with enterprise cloud customers. In 32 cases, these requests were rejected, withdrawn, or law enforcement was successfully redirected to the customer. In 18 cases, Microsoft was compelled to provide responsive information: 10 of these cases required the disclosure of some customer content and in 8 of the cases Microsoft was compelled to disclose non-content information only.
- **Autarky:** Can Microsoft cloud operations be separated from the Internet or the rest of Microsoft cloud and connected solely to local government network? Are operations possible without external connections to a third party? **Answer:** Yes, depending on the cloud deployment model.
 - **Public Cloud:** Azure regional datacenters can be connected to local government network through dedicated private connections such as ExpressRoute. Independent operation without any connectivity to a third party such as Microsoft is not possible in public cloud.
 - **Private Cloud:** With Azure Stack, customers have full control over network connectivity and can operate Azure Stack in fully disconnected mode.
- **Data flow restrictions:** What provisions exist for approval and documentation of all data exchange between customer and Microsoft for local, in-country deployed cloud services? **Answer:** Options vary based on the cloud deployment model.
 - **Private cloud:** For private cloud deployment using Azure Stack, customers can control which data is exchanged with third parties. Azure Stack telemetry can be turned off based on customer preference and Azure Stack can be operated fully disconnected. Moreover, Azure Stack offers the [capacity-based billing model](#) in which no billing or consumption data leaves the customer's premises.
 - **Public cloud:** In Azure public cloud, customers can leverage [Network Watcher](#) to monitor network traffic associated with their workloads. For public cloud workloads, all billing data is generated through telemetry used exclusively for billing purposes and sent to Microsoft billing systems. Customers can [download and view](#) their billing and usage data; however, they cannot prevent this information from being sent to Microsoft. Microsoft engineers [do not have default access](#) to Customer Data. For customer-initiated support requests, [Azure Customer Lockbox](#) can be used to enable customers to Approve/Deny elevated requests for Customer Data access. Moreover, customers have control over data encryption at rest using customer-managed encryption keys.

- **Patching and maintenance for private cloud:** How can Microsoft support patching and other maintenance for Azure Stack private cloud deployment? **Answer:** Microsoft [update packages for Azure Stack](#) are released monthly. Government customers are sole operators of Azure Stack and they can import and install these update packages from the administrator portal.

Safeguarding of customer data

- **Microsoft network security:** What network controls and security does Microsoft use? Can customer requirements be considered? **Answer:** For insight into Azure infrastructure protection, customers should review Azure [network architecture](#), Azure [production network](#), and Azure [infrastructure monitoring](#). Customers deploying Azure applications should review Azure [network security overview](#) and [network security best practices](#). To provide feedback or requirements, customers should engage their Microsoft account teams.
- **Customer separation:** How does Microsoft logically or physically separate customers within its cloud environment? Is there an option for select customers to ensure complete physical separation? **Answer:** Azure uses [logical isolation](#) to segregate each customer's applications and data from those of others. This approach provides the scale and economic benefits of multi-tenant cloud services while rigorously enforcing controls designed to keep customers from accessing one another's data or applications. There is also an option to enforce physical compute isolation by provisioning Azure E64is v3, E64i v3, GS5, G5, DS15 v2, and D15 v2 Virtual Machine instances. These VMs allow customers to be deployed on [hardware dedicated to a single customer](#). Utilizing isolated VMs essentially guarantees that customer VM will be the only one running on that specific server node.
- **Data encryption at rest and in transit:** Does Microsoft enforce data encryption by default? Does Microsoft support customer-managed encryption keys? **Answer:** Yes, Azure Storage and Azure SQL Database encrypt data by default and support customer-managed keys. Azure [Storage Service Encryption for Data at Rest](#) helps ensure that data is automatically encrypted before persisting it to Azure Storage and decrypted before retrieval. Customers can [use their own encryption keys for Azure Storage](#) encryption at rest and manage their keys in Azure Key Vault. Storage Service Encryption is enabled by default for all new and existing storage accounts and cannot be disabled. When provisioning storage accounts, customers can enforce “[Secure transfer required](#)” option, which allows access only from secure connections. This option is enabled by default when creating a storage account in the Azure Management Portal. Azure SQL Database enforces [data encryption in transit](#) by default and provides [Transparent Data Encryption](#) (TDE) at rest by [default](#) allowing customers to use Azure Key Vault and [Bring Your Own Key](#) (BYOK) functionality to control key management tasks including key rotation, key deletion, permissions, etc.
- **Data encryption during processing:** Can Microsoft protect Customer Data while it is being processed in memory? **Answer:** Yes, Microsoft Azure [Confidential Computing](#) is designed to address this requirement through the use of Trusted Execution Environments (TEEs, aka enclaves) that protect customer data while in use. TEEs are hardware or software implementations that safeguard data being processed from access outside the TEE. The hardware provides a protected container by securing a portion of the processor and memory. Only authorized code is permitted to run and to access data, so code and data are protected against viewing and modification from outside of the TEE.
- **FIPS 140-2 validation:** Does Microsoft offer FIPS 140-2 Level 3 validated Hardware Security Modules in Azure? **Answer:** Yes, Azure provides [Dedicated HSMs](#) that have FIPS 140-2 Level 3

validation, as well as Common Criteria EAL4+ certification and conformance with eIDAS requirements.

- **Customer provided crypto:** Can customers bring their own cryptography or encryption hardware? **Answer:** If customers expect to use back-end services integrated with [Azure Key Vault](#) (e.g., Azure Storage, SQL Database, Data Lake Storage, Disk Encryption, etc.), then they need to use Microsoft provided cryptography and encryption hardware, e.g., Hardware Security Modules (HSMs). Customers can use their own HSMs deployed on-premises with their own crypto algorithms. Similarly, if customers are migrating their HSM applications from on-premises to Azure, they can use [Azure Dedicated HSM](#) and retain their crypto algorithms.
- **Access to customer data by Microsoft personnel:** How does Microsoft restrict access to customer data by Microsoft engineers? **Answer:** Microsoft engineers [do not have default access to Customer Data](#) in the cloud. Instead, they are granted access, under management oversight, only when necessary using the [restricted access workflow](#). For customer-initiated support requests, Azure [Customer Lockbox](#) provides customers with the capability to control how a Microsoft Engineer accesses their data. As part of the support workflow, a Microsoft Engineer may require elevated access to Customer Data. Azure Customer Lockbox puts the customer in charge of that decision by enabling the customer to Approve/Deny such elevated requests.

Operations

- **Code review:** What can Microsoft do to help ensure that no malicious code has been inserted into the services that customers use? Can customers review Microsoft code deployments? **Answer:** Microsoft has full control over all source code that comprises Azure services. The procedure for patching guest VMs differs greatly from traditional on-premises patching where patch verification is necessary following installation. In Azure, patches are not applied to guest VMs; instead, the VM is simply restarted and when the VM boots, it is guaranteed to boot from a known good image that Microsoft controls. There is no way to insert malicious code into the image or interfere with the boot process. PaaS VMs offer more advanced protection against persistent malware infections than traditional physical server solutions, which if compromised by an attacker can be difficult to clean, even after the vulnerability is corrected. With PaaS VMs, reimaging is a routine part of operations, and it can help clean out intrusions that have not even been detected. This makes it much more difficult for a compromise to persist. Customers cannot review Azure source code; however, online access to view source code is available for key products through the Microsoft [Government Security Program](#) (GSP).
- **DevOps personnel (cleared, nationality, etc.):** What controls or clearance levels does Microsoft have for the personnel that have DevOps access to cloud environments or physical access to data centers? **Answer:** Microsoft conducts background screening on operations personnel with access to production systems and physical data center infrastructure. Microsoft cloud background check includes verification of education and employment history upon hire, and additional checks conducted every two years thereafter (where permissible by law), including criminal history check, OFAC list, BIS denied persons list, and Office of Defense Trade Controls debarred parties list.
- **Data center site options:** Is Microsoft willing to deploy a data center to a specific physical location to meet more advanced security requirements? **Answer:** Customers should enquire with their Microsoft account team regarding options for data center locations.
- **Service availability guarantee:** How do we ensure that Microsoft (or particular government or other entity) can't turn off our cloud services? **Answer:** Customers should review the [Online](#)

[Services Terms](#) for contractual commitments Microsoft makes regarding service availability and use of online services.

- **Non-traditional cloud service needs:** What is the recommended approach for managing scenarios where Azure services are required in periodically internet free/disconnected environments? **Answer:** In addition to Azure Stack which is intended for on-premises deployment, a ruggedized and field-deployable version called [Tactical Azure Stack](#) is also available to address tactical edge deployments for limited or no connectivity, fully mobile requirements, harsh conditions requiring military specification solutions, etc.

Transparency and audit

- **Audit documentation:** Does Microsoft make all audit documentation readily available to customers to download and examine? **Answer:** Yes, Microsoft makes all independent third-party audit reports and other related documentation available to customers from the [Service Trust Portal](#).
- **Process auditability:** Does Microsoft make its processes, data flow, and documentation available to customers or regulators for audit? **Answer:** Yes, Microsoft offers a Regulator Right to Examine, which is a program Microsoft implemented to provide the Regulator with direct right to examine Azure, including the ability to conduct an on-site examination, to meet with Microsoft personnel and Microsoft external auditors, and to access any related information, records, reports, and documents.
- **Service documentation:** Can Microsoft provide in-depth documentation covering service architecture, software and hardware components, and data protocols? **Answer:** Yes, Microsoft provides extensive and in-depth [Azure online documentation](#) covering all these topics. For example, customers can review documentation on Azure [products](#), [global infrastructure](#), and [API reference](#).