UNITED STATES OF AMERICA
BEFORE THE
FEDERAL ENERGY REGULATORY COMMISSION

	)	
Reliability Technical Conference	)	**Docket No. AD18-11-000**
	)	

TECHNICAL CONFERENCE

Written Statement on Behalf of Microsoft Corporation[1]

July 31, 2018

Thank you for the opportunity to address the Federal Energy Regulatory Commission ("FERC" or the "Commission") on policy issues related to the reliability of the Bulk Power System. Microsoft offers comprehensive cloud computing services, including servers, storage, databases, networking, software, analytics and more, that are available to Microsoft customers via a common, internet-based cloud infrastructure and platform.[2] One of the primary benefits of cloud computing is the concept of a shared, common infrastructure across numerous customers simultaneously, which leads to economies of scale. This concept is called "multi-tenancy." A multi-tenant cloud platform means that multiple customer applications and data are stored on the same physical hardware. Microsoft uses logical isolation techniques[3] to separate cloud tenants and create an environment where customers can access and manage only their own cloud-based resources.[4]

---

[1] This statement is presented by Matt Rathbun, Chief Security Officer, Azure Global at Microsoft.
[2] The National Institute of Standards and Technology ("NIST") defines cloud computing as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." U.S. Dept. of Commerce, Nat'l Inst. of Standards and Tech., *The NIST Definition of Cloud Computing*, at 2 (Sept. 2011), http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf.
[3] "Logical isolation" is a configuration that prevents two sets of devices which share a physical network infrastructure from being able to communicate with each other.
[4] 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. Microsoft personnel have

As a cloud service provider,[5] Microsoft does not perform any bulk electric system ("BES") functions that would subject it to registration under the North American Electric Reliability Corporation's ("NERC") functional model. The only assets that Microsoft controls are the hardware and software that underlie the Microsoft cloud services offerings, which can be used by Registered Entities for a variety of purposes, including to run cloud-based BES Cyber Systems. Because (i) cloud service providers do not own or operate elements of the BES, and (ii) cloud service providers such as Microsoft cannot control how their customers use the cloud services, Microsoft construes existing guidance to mean that (i) cloud service providers are not required to register with NERC for their ownership and operation of the cloud, and that, (ii) the onus of meeting Critical Infrastructure Protection ("CIP") Reliability Standards ("CIP Standards") is on Registered Entities if they elect to use cloud services.[6] However, as discussed more fully below, Registered Entities that use cloud services can effectively meet the security assurances sought by the CIP Standards if NERC provides clarity on Registered Entities' use of the cloud.

---

very limited access to customers' cloud-based resources, which is outlined in the contracts for services. Access to customer resources is only permitted for Microsoft to operate its commercial cloud services. When access is required, it is done so via just-in-time access using temporary credentials, and all actions by Microsoft personnel are logged and audited. Controls for the protection of customer secrets are audited on a regular basis as part of existing independent third-party audits. Customers also have several options for encrypting their data in the cloud, including keeping encryption keys in hardware security modules that are FIPS 140-2 Level 2 validated.

[5] Cloud computing is the delivery of computing services—servers, storage, databases, networking, software, analytics, and more—over the internet ("the cloud"). Companies offering these computing services are called cloud service providers and typically charge for cloud computing services based on usage.

[6] As Microsoft pointed out in its comments submitted in the Commission's Notice of Proposed Rulemaking "Cyber Security Incident Reporting Reliability Standards," Docket Nos. RM18-2-000 and AD17-9-000, to date, neither the Commission nor NERC has provided any clear guidance on the extent to which Registered Entities that use commercial cloud-based BES Cyber Systems must report incidents, or attempted incidents, relating to their use of cloud infrastructure. Registered Entities that use cloud services should be responsible for ensuring their own compliance with the reporting requirements set forth in the CIP Standards with respect to their management, configuration, and operation of their cloud-based BES Cyber Assets, rather than placing the onus on a commercial cloud service provider that operates a multi-tenant environment. Moreover, the Commission should clarify what constitutes an "attempted" incident, especially with respect to a Registered Entity's cloud-based BES Cyber Systems. *See Cyber Security Incident Reporting Reliability Standards*, Docket Nos. RM18-2-000 and AD17-9-000, Comments of Microsoft Corporation, at 5-6 (filed Feb. 26, 2018).

The following discussion addresses the questions raised for Panel IV: **Addressing the Evolving Cybersecurity Threat**. The first portion addresses the questions posed in Paragraph a. of the Supplemental Notice of the Technical Conference ("Supplemental Notice"):

- **How are current trends in cyber threats and vulnerabilities affecting the behavior of grid owners and operators?**
- **How can grid operators be better prepared to protect their systems from these threats?**
- **How do you recommend organizations mitigate cyber risks?**
- **How can the Critical Infrastructure Protection Reliability Standards (CIP Standards) be improved to assist responsible entities in addressing emerging cyber threats?**
- **What information-sharing practices are required?**
- **How are best practices developed, applied, and improved?**

<u>Cyberattacks as a Growing Threat to the Bulk Power System</u>

As the Supplemental Notice points out, there is a widespread understanding among policymakers and industry that cyber-attacks are a persistent and growing threat to the reliable or resilient operation of the Bulk Power System. A report by McAfee and the Center for Strategic and International Studies concluded that the potential cost of cybercrime to the global community is now upwards of $600 billion.[7] The cost of a data breach, which according to the Ponemon Institute averages about $3.62 million,[8] is expected to exceed $150 million by 2020.[9]

NERC includes CIP Standards that are directed at addressing the risk posed by cyber-attacks. The focus of these CIP Standards is to protect the utility industry from efforts to harm the reliable operation of the BES. As noted in a 2016 report prepared by the Idaho National

---

[7] https://www.mcafee.com/enterprise/en-us/assets/reports/restricted/economic-impact-cybercrime.pdf at p.6.
[8] https://www.ibm.com/security/data-breach.
[9] *See* The Best VPN, *Cyber Security Statistics*, https://thebestvpn.com/cyber-security-statistics-2018/ (last updated Feb. 27, 2018).

Laboratory, "[t]he modern electric grid is dependent upon cyber-physical systems, 'engineered systems that are built from, and depend upon, the seamless integration of computational algorithms and physical components,' to generate, move, and distribute electricity efficiently."[10] Technological improvements in the bulk power industry have caused the U.S. BES to be increasingly vulnerable to intrusions from cyberspace. As explained below, a significant opportunity to help enhance the security of the BES is by encouraging the use of technologies provided by Microsoft and other companies, including cloud computing services.

Cloud Computing Services Can Enhance Cyber Security Protection

Electric utilities and other entities subject to NERC registration can choose to install their own on-site hardware, software and related assets, or can purchase those services from a cloud service provider, such as Microsoft. By relying on cloud service providers such as Microsoft, electric utility industry participants can deploy their cyber systems to cloud service platforms that provide sophisticated protection against cyber threats and vulnerabilities, including anti-malware, web application firewall, intrusion detection systems, "Denial of Service" protection systems, and similar protections not provided at most on-site cyber facilities. Microsoft has implemented sophisticated big-data analytics and machine learning algorithms to analyze vast amounts of data from the monitoring and diagnostics infrastructure deployed throughout Microsoft Azure[11], including fully–automated, real-time alerting capabilities. Utility industry customers can use the Azure Security Center to help monitor access to their resources, gain

[10] Idaho National Laboratory, *Cyber Threat and Vulnerability Analysis of the U.S. Electric Sector*, Mission Support Center Analysis Report, at 5 (Aug. 2016), https://www.energy.gov/sites/prod/files/2017/01/f34/Cyber%20Threat%20and%20Vulnerability%20Analysis%20of%20the%20U.S.%20Electric%20Sector.pdf.

[11] Microsoft offers cloud services through three cloud service models: Infrastructure as a Service ("IaaS"), Platform as a Service ("PaaS"), and Software as a Service ("SaaS"). Although other Microsoft cloud offerings may be available and useful to Registered Entities, this Written Statement refers only to Microsoft Azure for clarity and ease of discussion.

insight into the security state of their resources, find and remediate vulnerabilities, limit exposure to threats, and detect and respond quickly to cyberattacks. By using cloud services, Registered Entities can take advantage of the significant investments that cloud service providers make in cyber security, including research and development.[12]

The security of the physical location where data is stored must also be protected to maintain reliable operation of the BES. Cloud data centers, like those hosted by Microsoft, contain many security features built into the physical features of their cloud platform offering. Indeed, in addition to the "virtual" security offered through encryption, operational controls, and other software security measures, Microsoft's Azure data centers also include industry standard and best practices for physical security measures. These physical security processes and practices may far exceed typical utility industry on-premises approaches.

Microsoft also leverages third-party reviews, in addition to internal controls, to maintain the security baseline for its cloud offerings.[13] For example, Microsoft Azure undergoes a rigorous third-party audit through the Federal Risk and Authorization Management Program ("FedRAMP"). In connection with its commitment to cloud adoption, the Government

[12] It would be very difficult and economically inefficient for any Registered Entity to duplicate this level of effort with respect to its own cyber assets. Security and privacy are built into the Microsoft Azure platform. Security Development Lifecycle (SDL) is a software development process that secures software and addresses security compliance requirements while reducing development cost. Operational Security Assurance (OSA) is a framework that incorporates deep awareness of the cybersecurity threat landscape and the experience of running hundreds of thousands of servers in data centers around the world. These provide secure operations throughout the lifecycle of cloud-based services, in a manner that cannot be economically replicated by stand-alone utility cyber assets.

[13] Microsoft Azure has the broadest compliance coverage in the cloud industry, including key independent certifications and attestations such as ISO 27001, ISO 27017, ISO 27018, ISO 22301, ISO 20000-1, ISO 9001, Service Organization Controls (SOC) 1/2/3, Payment Card Industry (PCI) Data Security Standard (DSS) Level 1, HITRUST Alliance, Cloud Security Alliance ("CSA") Security, Trust & Assurance Registry ("STAR") Certification, CSA STAR Attestation, and FedRAMP, which is described above. In terms of U.S. government focused compliance coverage, Azure Government (Microsoft's Azure offering for the Government community) has: FedRAMP High Provisional Authorization to Operate (P-ATO) issued by the Joint Authorization Board (JAB); Department of Defense (DoD) – Defense Information Systems Agency – Security Requirements Guide – "Level 4 Provisional Authorizations"; and NIST Federal Information Processing Standard 140-2 "Level 2" certification for cryptographic module validation.

developed and manages FedRAMP as a government-wide program providing a standardized approach to security assessment, authorization, and continuous monitoring for cloud services.[14] The FedRAMP authorization is based on NIST controls and involves risk-based assessment and authorization with mandatory provisions for continuous monitoring. Federal agencies rely on FedRAMP for assurances that a cloud security provider's security controls are operating effectively.

Aside from security benefits, cloud computing provides important reliability and resiliency benefits to NERC CIP workloads. Utility systems, including control systems, are built upon an approach of recovery from loss of one or more critical assets. This approach includes n-1 failures for grid assets covered by Contingency Analysis as part of the Energy Management System, and also for the computing infrastructure running the Energy Management System. In other words, the typical configuration is designed to recover from failure of a primary server. Cloud platforms are designed inherently to anticipate failure of each asset and to recover quickly and efficiently from failure in the underlying platform services, including compute, storage, and networking. Resiliency and availability are thus better assured by automated failure recovery of the equipment. For example, Azure Storage maintains three copies of customer data across separate fault domains in the primary region. Customers can also enable geo-redundant storage, which maintains three additional copies of customer data also across separate fault domains in the paired region. Azure Storage provides six complete replicas of customer data kept in two paired regions that are located at least 400 miles apart. In addition, Azure takes advantage of a computing model of high availability that runs multiple instances of an application, so that if any one fails, the rest pick up the load, and better assure availability and resiliency.

---

[14] *See* Steven VanRoekel, Federal Chief Information Officer, Memorandum for Chief Information Officers, at 1-2 (Dec. 8, 2011), available at https://www.fedramp.gov/assets/resources/documents/FedRAMP_Policy_Memo.pdf.

<u>Guidance from NERC and the Commission is Needed Regarding the Use of Cloud Computing by Registered Entities</u>

As explained previously, cyberattacks on the utility industry are on the rise, participation of new entrants to the power grid are making operation of the power system more complex, and the threat landscape is evolving at an accelerated rate. Both large and small Registered Entities may lack the staff, expertise, or financial resources to meet this evolving threat. At present, however, there is essentially no guidance from NERC regarding the appropriate security framework for a Registered Entity's use of cloud services. Registered Entities are subject to audits of CIP compliance.[15] A finding of non-compliance with CIP Standards may result in a significant penalty.[16] To date, however, NERC has not addressed whether and how a Registered Entity may pass a CIP Standards audit with workloads hosted in the cloud. As a result, Registered Entities that would otherwise benefit from cloud adoption currently face significant uncertainty arising from the potential for an auditor to disagree with the Registered Entity's approach to cloud deployment. Microsoft recommends that the Commission require NERC to address this issue as soon as possible.

The CIP Standards could better help Registered Entities approach security in the cloud if NERC were to explicitly endorse the use of cloud service providers with a FedRAMP authorization. FedRAMP offers an appropriate baseline to provide the assurances sought by the NERC CIP Standards in the cloud environment. NERC could rely on a cloud service provider's FedRAMP authorization to address NERC CIP audit requirements that apply to a cloud service provider. Each portion of the NERC CIP Standards for which the cloud service provider would be responsible under an audit maps to a NIST-based control. A cloud service provider's existing

---

[15] *See* 18 C.F.R. § 39.7 (2018) (Enforcement of Reliability Standards).
[16] *See id.*

FedRAMP authorization, which is based on NIST control evidence, provides assurance that controls equivalent to the applicable NERC CIP Standards have already been examined and approved through the rigorous auditing process that has been adopted by the U.S. Government. Under this approach, the cloud service provider's contributions would be deemed to have been met under the NERC CIP Standards if the cloud service provider obtains and maintains its FedRAMP status.[17]

Adopting the FedRAMP framework to demonstrate compliance with CIP Standards in the cloud would not replace or alter Registered Entities' NERC CIP compliance obligations. Regardless of the method a Registered Entity chooses to deploy its BES Cyber Systems technology workloads, a Registered Entity would still be required to exercise prudence and due diligence. However, adopting the FedRAMP framework would enable NERC Registered Entities to take advantage of the benefits of the cloud, while also ensuring the intended security controls captured by the NERC CIP Standards would in each case be met or exceeded. Such an approach is consistent with NERC's general risk-based approach to monitoring and assessment. Further, in adopting this approach, FERC would also promote the U.S. Government's goal of standardizing the cloud security framework. Even if FERC were to determine that FedRAMP is not an appropriate vehicle to address the security of NERC workloads in the cloud, further action is needed to provide clarity to industry. CIP Standards need to be amended to address the treatment of cloud environments for Registered Entities to be able to fully realize the benefits of the cloud.

[17] Microsoft has analyzed NERC CIP Standards requirements relative to NIST SP 800-53 controls and has concluded that a FedRAMP "Moderate" level audit provides extensive coverage for NERC CIP requirements. This approach is explained in more detail in the Microsoft Azure white paper "NERC CIP Standards and Cloud Computing," which is attached as Appendix A.

- **How could cloud computing, virtualization, and other technologies be deployed securely to help manage the emerging grid?**

In addition to security benefits, cloud computing provides important reliability and resiliency benefits to NERC CIP workloads. As discussed above, cloud platforms are designed inherently to anticipate failure and to recover from failure in the underlying platform services, including compute, storage, and networking. For example, customers can choose active geo-replication for Azure SQL Database, as well as variety of replication options for Azure Storage, including geo-redundant storage. Azure Storage maintains three copies of customer data across separate fault domains in the primary region. Customers can also enable geo-redundant storage, which maintains three additional copies of customer data also across separate fault domains in the paired region.

In addition to the inherent benefits of deployments directly to the cloud, supervisory control and data acquisition ("SCADA") systems can be deployed in the cloud for backup to primary control systems that might be on premise. These cloud-based SCADA systems can be available for grid black-start in the event of catastrophic compromise of the primary control system. These cloud-based SCADA deployments accrue minimal or no cost until used, so they can provide grid operators backup operations capability or insurance.

- **The Commission engages with other agencies and industry in mitigating the risk posed by cyber threats – including promoting information sharing, identifying and assessing threats, sharing lessons learned and best practices. How can we improve these efforts?**

Hyperscale cloud service providers[18] have significant information regarding cyber threats. Microsoft does not currently share threat intelligence with the Electronic Information

[18] According to Synergy Research Group, in 2017 there were 24 companies worldwide that are "hyperscale," accounting for 68% of the global cloud services market. *See* HPCwire, *Hyperscalers Emerging From 'Hype Phase'* (Apr. 12, 2017), https://www.hpcwire.com/2017/04/12/hyperscalers-emerging-hype-phase/.

Sharing and Analysis Center ("E-ISAC"), a NERC program that gathers and analyzes security data, shares appropriate data with stakeholders, coordinates incident management, and communicates mitigation strategies to stakeholders. Integration of Microsoft threat intelligence with E-ISAC is one possible avenue for improving threat visibility and assessment. Cloud deployments also lend themselves to widespread but controlled information dissemination. Automation of threat intelligence information sharing via cloud platforms can leverage deep cloud security models for controlling access, while providing immediate notification broadly across the industry.

- **How can cyber incident response plans be improved to address the evolving cyber threat landscape?**
- **For example, when a cyber system is compromised, antimalware software may not identify the system as compromised, and the only indicator may be the system's abnormal behavior.**

Employee training should remain the first step toward preventing cyber incidents, but Registered Entities that use cloud computing benefit from additional technical support. Detailed insight based on the latest threat telemetry and security intelligence is a prerequisite for addressing the evolving cyber threat landscape and offering advanced threat protection. By including threat telemetry on firewalls and other security applications within the cloud, Microsoft Azure customers receive automatic reports and analytics of threats to their cloud-based data and applications, thereby enhancing the customers' security intelligence. In addition, Microsoft Azure offers a mature security incident response plan that addresses how Microsoft Azure investigates, manages, and responds to security incidents. The Microsoft Azure security incident management program enumerates the steps, owners, and timelines for assessing and remediating threats using a standard operating procedure that contains a framework for evaluating the effectiveness of the program. Because of its economies of scale, Microsoft Azure can offer customers a team of dedicated security professionals to respond to suspected security events on a real-time basis.

- **When considering the emerging cyber threats to industrial control systems, what strengths and weaknesses in the body of CIP Reliability Standards are revealed?**
- **What role can the voluntary development, application, and sharing of best practices play?**

CIP Standards provide helpful guidance for some key considerations for cyber security, and the compliance framework ensures significant focus is applied to cybersecurity. However, this focus is largely an annual or bi-annual process, with the primary goal being compliance. In addition, NERC CIP Standards are industry-driven through often lengthy stakeholder proceedings, and, as a result, lag the state of the computing industry. For example, virtualization is now a fourteen-year-old technology, but NERC is only now addressing virtualization for incorporation in the next release of the CIP Standards. Accordingly, the CIP Standards in certain regards have not kept pace with the evolving technology landscape. Similarly, cyber threats are continuously evolving, and current CIP Standards do not completely embrace the breadth of information, the evolution of services, and the breadth of tools to identify and mitigate evolving cyber threats. Hyperscale cloud service providers such as Microsoft offer information, services, and tools to help Registered Entities address these evolving threats.

Voluntary sharing of best practices can help the collective industry address the evolving threat landscape. Addressing cyber security is very expensive. Sharing best practices can socialize some of this cost and improve the overall industry cyber security posture. It can also aid in industry mitigation. Whether this information is vetted and distributed via the E-ISAC, control systems providers, or directly among the Registered Entities themselves, the information exchange would be industry specific and could help bring all participants to a higher level of preparedness, whether at the top or the bottom of the 1400 NERC CIP registered entity list.

Appendix A

Microsoft Azure
NERC CIP Standards and Cloud Computing

Microsoft Azure

NERC CIP Standards and Cloud Computing

Microsoft

Abstract

Microsoft makes two different cloud environments available to electric utilities and other registered entities: Azure and Azure Government. Both provide a multi-tenant cloud services platform that registered entities 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 and Azure Government use 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 preventing customers from accessing one another's data or applications. This paper addresses common security and isolation concerns pertinent to the electric power industry. It also discusses compliance considerations for data and workloads deployed to Azure or Azure Government that are subject to NERC CIP standards.

Both Azure and Azure Government have the same comprehensive security controls in place, as well as the same Microsoft commitment on the safeguarding of customer data. Azure Government provides additional controls regarding US Government specific background screening requirements, including maintaining US persons for Azure Government operations. Moreover, Azure Government is only available in the United States to US-based registered entities.

Both Azure and Azure Government are suitable for registered entities deploying certain workloads subject to NERC CIP standards enforcement.

September 2017

https://aka.ms/AzureNERC

Acknowledgments

Author: Stevan Vidich

Reviewers: Larry Cochrane, Garima Jain, Matt Rathbun, Adam Soh, Matti Neustadt Storie

Contents

Introduction

This paper is intended for electric power utilities and [registered entities](#) considering cloud adoption for data and workloads subject to the North American Electric Reliability Corporation (NERC) Critical Infrastructure Protection (CIP) standards. It addresses compliance considerations involved with a NERC audit, and it provides a technical description of logical isolation measures implemented in Microsoft Azure and Azure Government to address tenant separation concerns. Specifically, this paper covers multi-tenancy and virtualization in place for Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) cloud service models to address common customer concerns when moving workloads from an on-premises environment to the cloud.

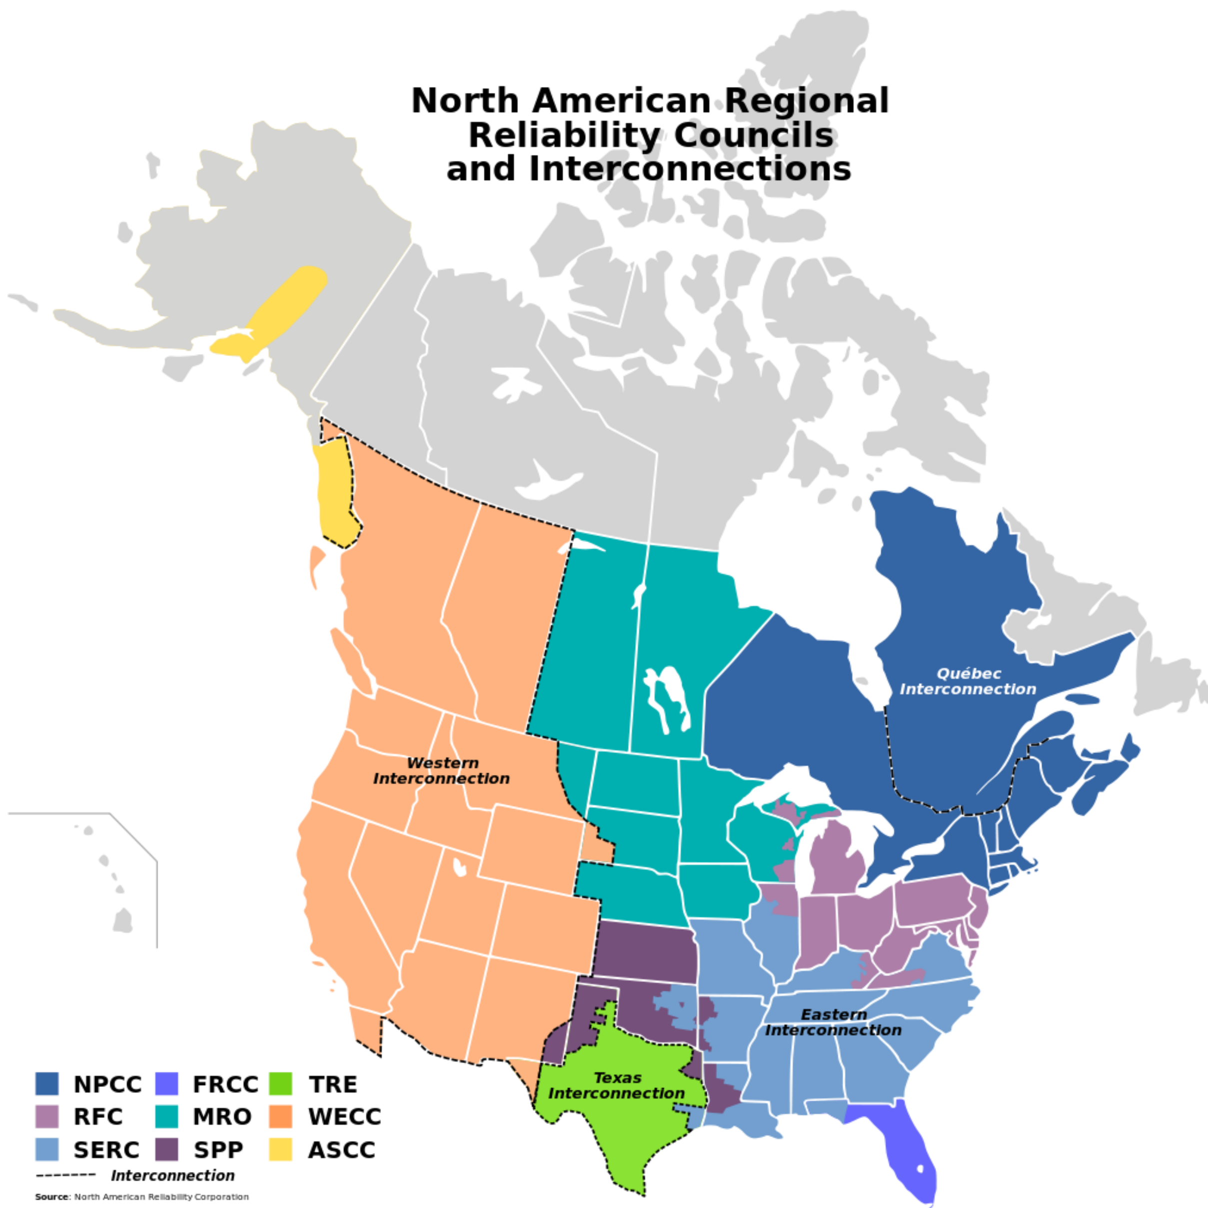

Figure 1: Regional reliability organizations and interconnections under NERC authority (Source: NERC)

NERC is a nonprofit regulatory authority whose mission is to ensure the reliability of the North American bulk power system. NERC is subject to oversight by the Federal Energy Regulatory Commission (FERC) and governmental authorities in Canada. In 2006, FERC granted the Electric Reliability Organization (ERO) designation to NERC in accordance with the Energy Policy Act of 2005 (U.S. Public Law 109-58). NERC has jurisdiction over users, owners, and operators of the bulk power system that serves more than 334 million people. Figure 1 shows regional reliability organizations and interconnections under NERC authority.

NERC develops and enforces reliability standards known as NERC CIP standards. In the United States, FERC approved the first set of CIP standards in 2007 and has continued to do so with every new revision. In Canada, the Federal, Provincial, and Territorial Monitoring and Enforcement Sub-group (MESG) develops provincial summaries for making CIP standards enforceable in Canadian jurisdictions.

Azure and Azure Government

Azure provides core infrastructure and virtualization technologies and services such as compute, storage, and networking that are designed with stringent controls to meet customer data separation requirements and help enable secure connection to customer on-premises environments. Most Azure services enable customers to specify the Region where their Customer Data will be stored. Microsoft may replicate Customer Data to other Regions within the same Geo for data resiliency but Microsoft will not replicate Customer Data outside the chosen Geo (e.g., United States).

Microsoft provides two different cloud environments to registered entities to deploy their applications and data: Azure and Azure Government. Azure is generally available in more than 50 Regions around the world; however, for registered entities subject to NERC CIP standards, the Geos of most interest are United States and Canada. As shown in Figure 2, Azure is available in 8 Regions located in Virginia, Iowa, Illinois, Texas, California, Washington, and Wyoming in the United States.

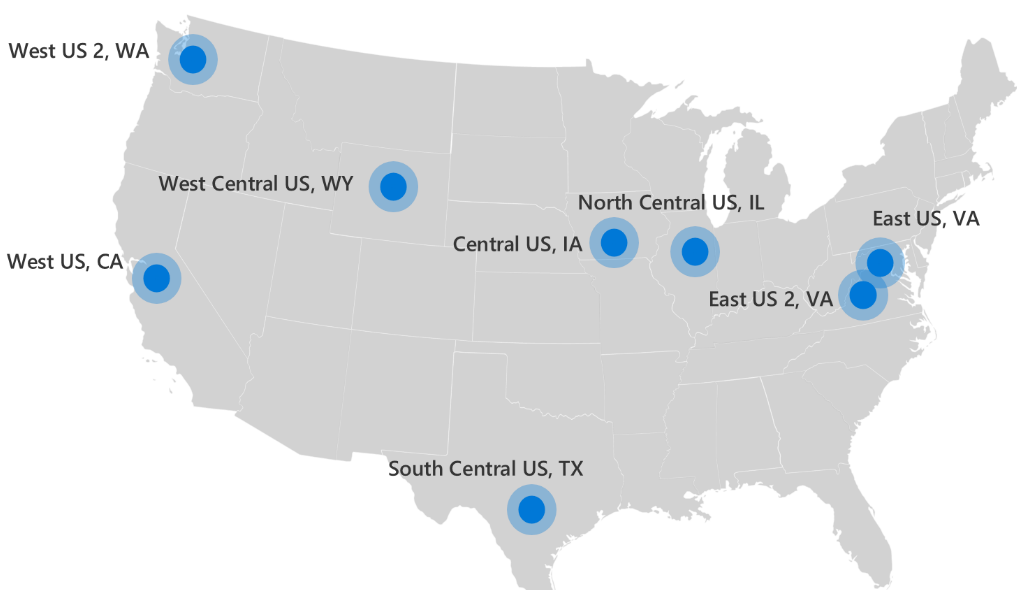

Figure 2: Azure cloud US locations

Moreover, Azure is available in two regions in Canada (Toronto and Quebec City). As shown in Figure 3, Azure Government is only available in the United States to US-based registered entities with four Regions located in Virginia, Iowa, Arizona, and Texas.

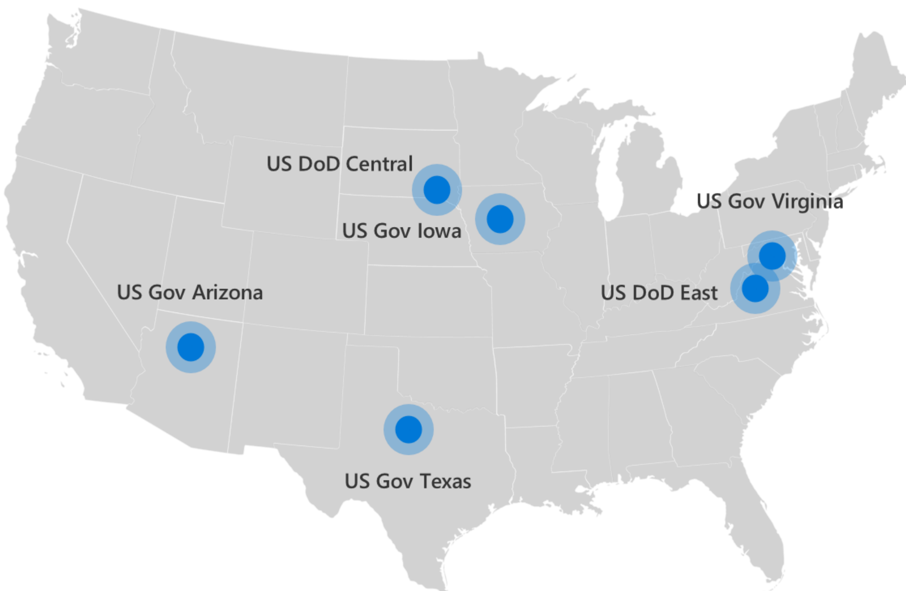

Figure 3: Azure Government cloud locations

Both Azure and Azure Government have the same strong security controls in place to provide robust assurances to customers about the safeguarding of customer data and applications. They offer a variety of services in a multi-tenant cloud environment that leverages virtualization technologies to provide scale and resource utilization, as well as superior data separation and isolation in a shared environment. This design helps ensure that electric utility customers and authorized administrators can use Azure and Azure Government efficiently and keep their data and workloads isolated from other tenants. Both cloud environments provide the same data redundancy for Azure Storage by maintaining three copies of Customer Data across separate fault domains in the primary Region. Customers can also enable geo-redundant storage, which maintains three additional copies of Customer Data also across separate fault domains in the paired Region. At any given time, Azure Storage provides 6 healthy replicas of Customer Data kept in two paired Regions that are located at least 400 miles apart.

Azure Government is a US government community cloud that is physically separated from the Azure cloud. It provides additional assurances regarding US Government specific background screening requirements, including US person verification for Azure Government operation personnel with potential access to customer data. Azure Government can also support customers subject to certain export controls laws and regulations. **Both Azure and Azure Government are suitable for registered entities deploying certain workloads subject to NERC CIP standards enforcement.**

All Azure and Azure Government employees in the United States are subject to the Microsoft cloud background check every two years, as discussed in the Background Screening section. Azure Government personnel are additionally subject to the verification of US persons, as well as the National Agency Check with Law and Credit (NACLC) that involves fingerprint background checks against an FBI database.

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, and FedRAMP Moderate Provisional Authorization to Operate (P-ATO) issued by the Joint Authorization Board (JAB).

In terms of US Government focused compliance coverage, Azure Government has

- **FedRAMP High** P-ATO issued by the JAB
- Department of Defense (**DoD**) Defense Information Systems Agency (DISA) Security Requirements Guide (SRG) **Level 4** Provisional Authorizations (**Level 5** Provisional Authorization is available in the DoD Region)
- **FIPS 140-2 Level 2** certification for cryptographic module validation
- Contractual amendments available to support FBI Criminal Justice Information Services (**CJIS**) and Internal Revenue Service **IRS 1075** requirements
- Contractual amendment to support Directorate of Defense Trade Controls (DDTC) International Traffic in Arms Regulations (**ITAR**) requirements
- Support for **NIST SP 800-171** guidance for the protection of Controlled Unclassified Information (CUI) and DoD Defense Federal Acquisition Regulation Supplement (**DFARS**) Clause 252.204-7012

More information on the scope of services covered by these compliance offerings can be found from the Overview of Microsoft Azure Compliance or at the Microsoft Trust Center.

Nuclear electric utility customers may also be subject to the Department of Energy (DoE) / National Nuclear Security Administration (NNSA) **10 CFR Part 810** export control requirements. DoE 10 CFR Part 810 (final rule) became effective on 25 March 2015, and, among other things, it controls the export of unclassified nuclear technology and assistance. § 810.7 (b) states that specific DoE authorization is required for providing or transferring sensitive nuclear technology to any foreign entity. Export is the transfer of protected technology or information to a foreign destination or foreign person irrespective of the destination, whereas Deemed Export represents the transmission of protected technology and information to a foreign person inside the United States. Azure Government is designed to be able to meet specific controls that restrict access to information and systems to US persons. This commitment is not applied in Azure so customers deploying in Azure should consider whether additional technical measures, such as data encryption, should be taken to address DoE 10 CFR Part 810 requirements. **Nuclear utility customers are wholly responsible for ensuring their own compliance with all applicable laws and regulations. The forgoing is not legal advice, and customers should consult their legal advisors for any questions regarding regulatory compliance.**

Classifying NERC CIP data and workloads

Customers operating Bulk Electric Systems are wholly responsible for ensuring their own compliance with NERC CIP standards. Neither Azure nor Azure Government constitutes a Bulk Electric System (BES) or BES Cyber Asset. As stated by NERC, CIP standards apply to the BES:

- Generally, 100kV and above, but with some exceptions, primarily for radial lines

- 20MVA and above generating units, 75MVA and above generating plants, with some exceptions for wholly behind-the-meter generation
- Includes Control Centers that monitor and control the BES

As stated by NERC, CIP standards do not apply to distribution (i.e., non-BES) with several exceptions, primarily Under Frequency Load Shedding (UFLS), Under Voltage Load Shedding (UVLS), Blackstart Resources (generation), and Cranking Paths.

To assess the suitability of NERC CIP standards data and workloads for cloud deployment, registered entities should consult with their own compliance officers and NERC auditors. Below are some key BES-related definitions provided by NERC in the current set of CIP standards and NERC's Glossary of Terms:

- **Cyber Asset:** Programmable electronic devices, including the hardware, software, and data *in those devices*.
- **BES Cyber Asset (BCA):** A Cyber Asset that if rendered unavailable, degraded, or misused would, within 15 minutes of its required operation, misoperation, or non-operation, adversely impact one or more Facilities, systems, or equipment, which, if destroyed, degraded, or otherwise rendered unavailable when needed, would affect the reliable operation of the Bulk Electric System. Redundancy of affected Facilities, systems, and equipment shall not be considered when determining adverse impact. Each BES Cyber Asset is included in one or more BES Cyber Systems.
- **BES Cyber System (BCS):** One or more BES Cyber Assets logically grouped by a responsible entity to perform one or more reliability tasks for a functional entity.
  - Components of the BCS also include "glue" infrastructure components (e.g., networking infrastructure) necessary for the system to perform its reliability tasks, such as network switches
  - Tremendous flexibility is built into the definition – BCS could be the entire control system, or a subset based on function (HMI, server, database, FEP, etc.).
- **Electronic Security Perimeter (ESP):** The logical border surrounding a network to which BES Cyber Systems are connected using a routable protocol.
- **Protected Cyber Asset (PCA):** One or more Cyber Assets connected using a routable protocol within or on an Electronic Security Perimeter that is not part of the highest impact BES Cyber System within the same Electronic Security Perimeter. The impact rating of Protected Cyber Assets is equal to the highest rated BES Cyber System in the same ESP.
- **Electronic Access Point (EAP):** A Cyber Asset interface on an Electronic Security Perimeter that allows routable communication between Cyber Assets outside an Electronic Security Perimeter and Cyber Assets inside an Electronic Security Perimeter.
- **Electronic Access Control or Monitoring Systems (EACMS):** Cyber Assets that perform electronic access control or electronic access monitoring of the Electronic Security Perimeter(s) or BES Cyber Systems. This includes intermediate Systems.
- **Control Center:** One or more facilities *hosting operating personnel that monitor and control* the Bulk Electric System (BES) in real-time to perform the reliability tasks, including their associated data centers, of: 1) a Reliability Coordinator, 2) a Balancing Authority, 3) a Transmission Operator for transmission Facilities *at two or more locations*, or 4) a Generator Operator for generation Facilities at two or more locations.

  - Includes rooms and equipment where power system operators sit, as well as rooms and equipment containing the "back office" servers, databases, telecommunications equipment, etc.
  - They may all be in the same room or be in different buildings or in different cities.

As stated by NERC, BES Cyber Assets perform real-time functions of monitoring or controlling the BES. There is heavy emphasis in the current definition on physical assets within the Electronic Security Perimeter (e.g., the very specific term *"in those devices"* referring to BES Cyber Assets), and no provisions for key cloud concepts such as virtualization, logical isolation, and multi-tenancy. **To properly accommodate BES Cyber Assets and Protected Cyber Assets in cloud computing, existing definitions in NERC CIP standards would need to be revised or augmented with a cloud implementation guide.** However, NERC has acknowledged that there are many workloads that deal with CIP sensitive data and do not fall under the 15-minute rule. More detailed guidance was provided by NERC in November 2016 at the Emerging Technology Roundtable on Cloud Computing.

Depending registered entity's implementation, some of the following workloads may or may not be considered a BCS or placed within the ESP:

- Storing all transmission substation data in a cloud based HIS, including substation equipment status, P&C settings, and substation topology.
- Transmission network planning using a cloud-based application and cloud-based storage.
- Transmission demand forecasting using a cloud-based Machine Learning algorithm.
- Contingency Analysis conducted in the cloud to reduce the risk of outages.
- Utility asset management and predictive maintenance for transmission assets.
- Geospatial asset location information.
- Common Information Model (CIM) modeling and existing CIM network model.
- Streaming of operational phasor data to the cloud for storage and analytics.
- Black-Scholes pricing models for bulk generation energy trading.

These workloads require careful assessment that takes into consideration individual utility needs. Another class of data not subject to the 15-minute rule is the BES Cyber System Information (BCSI) provided that proper security controls are in place to safeguard BCSI. The following definition is provided by NERC:

BES Cyber System Information (BCSI): Information about the BES Cyber System that could be used to gain unauthorized access or pose a security threat to the BES Cyber System. BES Cyber System Information does not include individual pieces of information that by themselves do not pose a threat or could not be used to allow unauthorized access to BES Cyber Systems, such as, but not limited to, device names, individual IP addresses without context, ESP names, or policy statements. Examples of BES Cyber System Information may include, but are not limited to, security procedures or security information about BES Cyber Systems, Physical Access Control Systems, and Electronic Access Control or Monitoring Systems that is not publicly available and could be used to allow unauthorized access or unauthorized distribution; collections of network addresses; and network topology of the BES Cyber System.

NERC CIP compliance requirements can be addressed during a NERC audit and in line with the shared responsibility model for cloud computing. We believe that the robust logical isolation capability in a multi-tenant cloud combined with commitments on data location and personnel background screening allow for using Azure or Azure Government cloud services in a manner compliant with NERC CIP standards. Microsoft has developed a Cloud Implementation Guide for NERC Audits, and is prepared to assist registered entities with NERC audits by furnishing Azure or Azure Government audit documentation and control implementation details in support of customer's NERC audit requirements.

Compliance considerations for NERC CIP standards

NIST SP 800-145 defines the following cloud service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). The shared responsibility model for cloud computing is depicted in Figure 4. With on-premises deployment in their own datacenter, customers assume the responsibility for all layers in the stack. As workloads get migrated to the cloud, Microsoft assumes progressively more responsibility depending on the cloud service model. For example, with the IaaS model, Microsoft's responsibility ends at the Hypervisor layer, and customers are responsible for all layers above the virtualization layer, including maintaining the base operating system in guest Virtual Machines. With finished cloud services in the SaaS model such as Microsoft Office 365 or Dynamics 365, Microsoft assumes responsibility for all layers in the stack; however, customers are still responsible for administering the service, including granting proper access rights to end users.

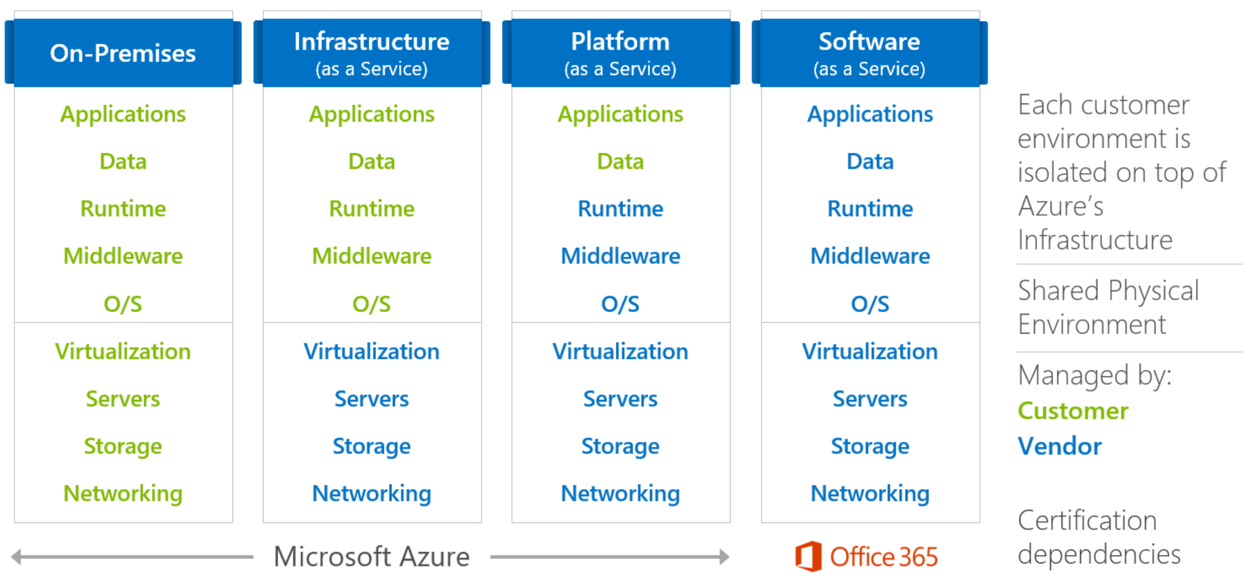

Figure 4: Shared responsibility model in cloud computing

The concept of shared responsibility extends also to certification dependencies and compliance obligations. When registered entities deploy applications to Azure or Azure Government, they take certification dependencies on Microsoft. Customers are ultimately responsible for meeting their NERC CIP compliance obligations; however, they inherit security controls from the underlying cloud platform, and can count on Microsoft for audit assistance.

Both Azure and Azure Government are audited extensively by independent third-party auditors, and some of these audits can be leveraged by registered entities when assessing their NERC CIP compliance obligations. In discussion with NERC regulators, the following independent third-party audits were identified as relevant and potentially useful to registered entities:

1. Cloud Security Alliance (CSA) STAR program
2. AICPA SOC 2 Type 2 attestation
3. US Government FedRAMP authorization

Microsoft already supports all three of these compliance audits and has the respective certifications and attestations in place that are available to customers.

Cloud Security Alliance

The Cloud Security Alliance (CSA) is a nonprofit organization led by a broad coalition of industry practitioners, corporations, and other important stakeholders. It is dedicated to defining best practices to help ensure a more secure cloud computing environment, and to helping potential cloud customers make informed decisions when transitioning their IT operations to the cloud.

In 2013, the CSA and the British Standards Institution launched the Security, Trust & Assurance Registry (STAR), a free, publicly accessible registry in which Cloud Service Providers (CSPs) can publish their CSA-related assessments. CSA STAR is based on the following components:

- Cloud Controls Matrix (CCM): a controls framework covering fundamental security principles across 16 domains (see Table 1) to help cloud customers assess the overall security risk of a CSP.
- The Consensus Assessments Initiative Questionnaire (CAIQ): a set of nearly 300 questions based on the CCM that a customer or cloud auditor may want to ask of CSPs to assess their compliance with CSA best practices.

Table 1: Cloud Control Matrix domains

CSA CCM 16 control areas	
Application and interface security	Human resources
Audit assurance and compliance	Identity and access management
Business continuity management and operational resilience	Infrastructure and virtualization security
Change control and configuration management	Interoperability and portability
Data security and information lifecycle management	Mobile security
Datacenter security	Security incident management, e-discovery, and cloud forensics
Encryption and key management	Supply chain management, transparency and accountability
Governance and risk management	Threat and vulnerability management

CSA CCM version 3.0.1 contains control mappings to NERC CIP v3, ISO 27001, SOC 2, FedRAMP, PCI DSS and many more standards. Even though the current set of CIP standards are not represented in the matrix, it is still a very useful tool for electric utility customers to assess how NERC CIP requirements map to established standards and audits that are applicable to cloud service providers (Figure 5). CSA keeps updating the CCM, so future updates may reflect the current set of NERC CIP standards.

CSA STAR consists of three levels of assurance aligned with the control objectives in the CCM:

- Level 1: STAR Self-Assessment
- Level 2: STAR Certification, STAR Attestation, and C-STAR Assessment
- Level 3: STAR Continuous Monitoring, which is still under development by CSA

Whereas the STAR Self-Assessment can be submitted directly by a CSP using either the CCM or CAIQ, Level 2 entries such as STAR Certification and STAR Attestation require rigorous, independent assessments by accredited auditing firms. The C-STAR Assessment is specific to the Greater China market, and it harmonizes CSA best practices with specific Chinese national standards.

CCMv3.0.1 CLOUD CONTROLS MATRIX VERSION 3.0.1

Control Domain	CCM V3.0 Control ID	Updated Control Specification	FedRAMP Security Controls (Final Release, Jan 2012) --MODERATE IMPACT LEVEL--	ISO/IEC 27001:2013	NERC CIP	NIST SP800-53 R3
Identity & Access Management *User Access Revocation*	IAM-11	Timely de-provisioning (revocation or modification) of user access to data and organizationally-owned or managed (physical and virtual) applications, infrastructure systems, and network components, shall be implemented as per established policies and procedures and based on user's change in status (e.g., termination of employment or other business relationship, job change or transfer). Upon request, provider shall inform customer (tenant) of these changes, especially if customer (tenant) data is used as part the service and/or customer (tenant) has some shared responsibility over implementation of control.	NIST SP 800-53 R3 AC-2 NIST SP 800-53 R3 AC-2 (1) NIST SP 800-53 R3 AC-2 (2) NIST SP 800-53 R3 AC-2 (3) NIST SP 800-53 R3 AC-2 (4) NIST SP 800-53 R3 AC-2 (7) NIST SP 800-53 R3 PS-4 NIST SP 800-53 R3 PS-5 NIST SP 800-53 R3 SC-30	Annex A A.9.2.6 A.9.1.1 A.9.2.1, A.9.2.2 A.9.2.3	CIP-004-3 R2.2.3 CIP-007-3 - R5.1.3 -R5.2.1 - R5.2.3	AC-2 PS-4 PS-5
Identity & Access Management *User ID Credentials*	IAM-12	Internal corporate or customer (tenant) user account credentials shall be restricted as per the following, ensuring appropriate identity, entitlement, and access management and in accordance with established policies and procedures: • Identity trust verification and service-to-service application (API) and information processing interoperability (e.g., SSO and Federation) • Account credential lifecycle management from instantiation through revocation • Account credential and/or identity store minimization or re-use when feasible • Adherence to industry acceptable and/or regulatory compliant authentication, authorization, and accounting (AAA) rules (e.g., strong/multi-factor, expireable, non-shared authentication secrets)	NIST SP 800-53 R3 AC-1 NIST SP 800-53 R3 AC-2 NIST SP 800-53 R3 AC-3 NIST SP 800-53 R3 AC-11 NIST SP 800-53 R3 AC-11 (1) NIST SP 800-53 R3 AU-2 NIST SP 800-53 R3 AU-2 (3) NIST SP 800-53 R3 AU-2 (4) NIST SP 800-53 R3 AU-11 NIST SP 800-53 R3 IA-1 NIST SP 800-53 R3 IA-2 NIST SP 800-53 R3 IA-2 (1) NIST SP 800-53 R3 IA-2 (2) NIST SP 800-53 R3 IA-2 (3) NIST SP 800-53 R3 IA-2 (8) NIST SP 800-53 R3 IA-5 NIST SP 800-53 R3 IA-5 (1) NIST SP 800-53 R3 IA-5 (2) NIST SP 800-53 R3 IA-5 (3) NIST SP 800-53 R3 IA-5 (6) NIST SP 800-53 R3 IA-5 (7) NIST SP 800-53 R3 IA-6 NIST SP 800-53 R3 IA-8 NIST SP 800-53 R3 SC-10	A.9.2.6 A.9.1.1 A.9.2.1, A.9.2.2 A.9.2.4 A.9.2.5 A.9.4.2	CIP-004-3 R2.2.3 CIP-007-3 - R5.2 - R5.3.1 - R5.3.2 - R5.3.3	AC-1 AC-2 AC-3 AC-11 AU-2 AU-11 IA-1 IA-2 IA-5 IA-6 IA-8 SC-10
Identity & Access Management *Utility Programs Access*	IAM-13	Utility programs capable of potentially overriding system, object, network, virtual machine, and application controls shall be restricted.	NIST SP 800-53 R3 AC-6 NIST SP 800-53 R3 AC-6 (1) NIST SP 800-53 R3 AC-6 (2) NIST SP 800-53 R3 CM-7 NIST SP 800-53 R3 CM-7 (1)	A.9.1.2 Deleted A.9.4.4	CIP-007-3 - R2.1 - R2.2 - R2.3	AC-5 AC-6 CM-7 SC-3 SC-19
Infrastructure & Virtualization Security *Audit Logging / Intrusion Detection*	IVS-01	Higher levels of assurance are required for protection, retention, and lifecyle management of audit logs, adhering to applicable legal, statutory or regulatory compliance obligations and providing unique user access	NIST SP 800-53 R3 AU-1 NIST SP 800-53 R3 AU-2 NIST SP 800-53 R3 AU-2 (3) NIST SP 800-53 R3 AU-2 (4)	A.12.4.1 A.12.4.1 A.12.4.2, A.12.4.3 A.12.4.3	CIP-007-3 - R6.5	AU-1 AU-2 AU-3 AU-4

CSA CCM V3.0.1 | CCM Change Log

Figure 5: Spreadsheet extract showing CCM mapping to NERC CIP v3, ISO 27001:2013, and NIST SP 800-53 R3 (Source: CSA)

Azure has completed the STAR Self-Assessment based on both CCM and CAIQ. Moreover, **Azure and Azure Government have [STAR Certification and STAR Attestation] produced by independent auditing firms** for the services listed as in-scope on the Trust Center, as shown in Figure 6.

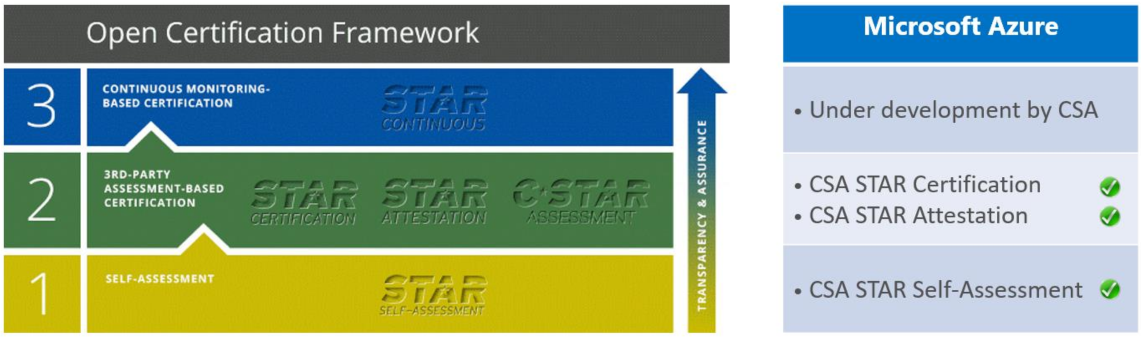

Figure 6: CSA Open Certification Framework (Source: CSA)

AICPA SOC 2 Type 2 Attestation

The American Institute of Certified Public Accountants (AICPA) has established three Service Organization Control (SOC) reporting options (SOC 1, SOC 2, and SOC 3) to assist CPAs with examining and reporting on a service organization's controls (Table 2).

Table 2: Service Organization Control (SOC) reporting options (Source: AICPA)

SOC Summary Chart (With Suggested Guidance):			
	SOC 1 Report	**SOC 2 Report**	**SOC 3 Report**
Kind of controls addressed by the report	Controls likely to be relevant to user entities financial statements	Controls over the security, availability and processing integrity of a system and the confidentiality and privacy of information processed by the system	Controls over the security, availability and processing integrity of a system, and the confidentiality and privacy of information processed by the system
Standard under which the engagement is performed and other related guidance	SSAE No. 16, Reporting on Controls at a Service Organization AICPA Guide. Service Organizations, Applying SSAE No. 16 (SOC 1℠)	AT 101, Attestation Engagements AICPA Guide, Reporting on Controls at a Service Organization (SOC 2℠)	AT 101, Attestation Engagements AICPA Technical Practice Aid, Trust Services Principles, Criteria and Illustrations
Content of report	Description of service organization's system CPA's opinion on fairness of presentation of the description, suitability of design and in a type 2 report, the operating effectiveness of controls A type 2 report includes a description of the CPA's tests of controls and results	Description of service organization's system CPA's opinion on the fairness of presentation of the description, suitability of design and in a type 2 report, the operating effectiveness of controls A type 2 report includes a description of the CPA's tests of controls and results	An unaudited system description used to delineate the boundaries of the system CPA's opinion on whether the entity maintained effective controls over its system

A SOC 2 Type 2 is a restricted use report intended to report on controls relevant to Security, Availability, Confidentiality, Processing Integrity, and Privacy system attributes. SOC 2 engagements are conducted in accordance with the Trust Services Principles and Criteria, as well as the requirements stated in the AICPA AT Section 101 standard. A Type 2 audit includes auditor's opinion on the control effectiveness to achieve the related control objectives during the specified monitoring period. Azure SOC 2 Type 2 attestation is based on a rigorous independent third-party audit conducted by a reputable CPA firm. For information about the in-scope services included in the Azure SOC attestations, see the Trust Center.

Customers can download the latest Azure SOC 2 Type 2 attestation report from the Service Trust Portal. Moreover, **Microsoft is prepared to assist registered entities with their NERC CIP compliance obligations by furnishing additional information, including control implementation details.**

US Government FedRAMP Authorization

The US Federal Risk and Authorization Management Program (FedRAMP) was established in December 2011 to provide a standardized approach for assessing, monitoring, and authorizing cloud service providers. It became operational in June 2012, and it is mandatory for certain US federal procurement programs.

Cloud Service Providers (CSPs) desiring to sell services to a federal agency requiring FedRAMP can take three paths to demonstrate FedRAMP compliance: 1) earn a Provisional Authorization to Operate (P-ATO) from the Joint Authorization Board (JAB); 2) receive an Authorization to Operate (ATO) from a federal agency; or 3) work independently to develop a CSP Supplied Package that meets program requirements. Each of these paths requires a stringent technical review by the FedRAMP Program Management Office (PMO) and an assessment by an independent third-party assessor organization (3PAO) that is accredited by the program.

FedRAMP is based on the National Institute of Standards and Technology (NIST) SP 800-53 Rev 4 standard, augmented by FedRAMP controls and enhancements. FedRAMP authorizations are granted at three impact levels based on the NIST FIPS 199 guidelines—Low, Moderate, and High. These levels rank the impact that the loss of confidentiality, integrity, or availability could have on an organization—Low (limited effect), Moderate (serious adverse effect), and High (severe or catastrophic effect). Consequently, the higher the impact level the more extensive the FedRAMP controls set is. Table 3 shows the number of FedRAMP controls and enhancements for Low, Moderate, and High baselines.

Table 3: FedRAMP control baselines

FedRAMP control baseline	Low	Moderate	High
Total number of controls and enhancements	125	325	421

Azure maintains a P-ATO at the Moderate impact level for the in-scope services listed at the Trust Center, and was the first public cloud with IaaS and PaaS services to receive a P-ATO. The JAB has also granted a FedRAMP High P-ATO to Azure Government for the in-scope services listed at the Trust Center, the highest bar for FedRAMP authorization. Once a P-ATO is granted, a CSP still requires an authorization—an ATO—from any government agency it works with. In the case of Azure and Azure

Government, a government agency can leverage the respective P-ATOs in its own security authorization process and rely on it as the basis for issuing an agency ATO that also meets FedRAMP requirements.

The JAB is the primary governance and decision-making body for FedRAMP. Representatives from the Department of Defense (DoD), Department of Homeland Security (DHS), and General Services Administration (GSA) serve on the board. The board grants a P-ATO to CSPs that have demonstrated FedRAMP compliance. It's important to note that FedRAMP is not a point-in-time certification but an assessment and authorization program that also comes with provisions for continuous monitoring mandated by DHS. A CSP is required to furnish a variety of evidence to demonstrate continuous compliance, including vulnerability scans, penetration test results, plan of actions and milestones, etc. FedRAMP is one of the most rigorous and demanding audits that a CSP can undergo.

A comparison between the FedRAMP Moderate control set and NERC CIP requirements reveals that FedRAMP control baseline encompasses all NERC CIP requirements. Microsoft has developed a Cloud Implementation Guide for NERC Audits (available to customers under a non-disclosure agreement) that includes control mapping between the current set of NERC CIP standards and FedRAMP control set (NIST 800-53 Rev 4). Figure 7 shows the current NERC CIP standards and FedRAMP control families.

NERC CIP standards

Reliability Standards

Standard Number	Title
⊟ **(CIP) Critical Infrastructure Protection** (82)	
⊟ Subject to Enforcement (11)	
CIP-002-5.1	Cyber Security — BES Cyber System Categorization
CIP-003-6	Cyber Security - Security Management Controls
CIP-004-6	Cyber Security - Personnel & Training
CIP-005-5	Cyber Security - Electronic Security Perimeter(s)
CIP-006-6	Cyber Security - Physical Security of BES Cyber Systems
CIP-007-6	Cyber Security - System Security Management
CIP-008-5	Cyber Security - Incident Reporting and Response Planning
CIP-009-6	Cyber Security - Recovery Plans for BES Cyber Systems
CIP-010-2	Cyber Security - Configuration Change Management and Vulnerability Assessments
CIP-011-2	Cyber Security - Information Protection
CIP-014-2	Physical Security

FedRAMP control set

ID	Family	Low	Moderate
AC	Access Control	11	18 (25)
AT	Awareness and Training	4	4 (1)
AU	Audit and Accountability	10	11 (8)
CA	Certification, Accreditation, and Security Assessment	7 (1)	8 (7)
CM	Configuration Management	8	11 (15)
CP	Contingency Planning	6	9 (15)
IA	Identification and Authentication	7 (8)	8 (19)
IR	Incident Response	7	9 (9)
MA	Maintenance	4	6 (5)
MP	Media Protection	4	7 (3)
PE	Physical and Environmental Protection	10	16 (4)
PL	Planning	3	4 (2)
PS	Personnel Security	8	8 (1)
RA	Risk Assessment	4	4 (6)
SA	System and Services Acquisition	6 (1)	9 (13)
SC	System and Communications Protection	10	20 (12)
SI	System and Information Integrity	6	12 (16)
	Totals (Controls and Enhancements):	125	325

Figure 7: NERC CIP standards and FedRAMP control set (Source: NERC and FedRAMP)

There are many valid reasons why an electric utility subject to NERC CIP compliance obligations might want to leverage an existing FedRAMP P-ATO or ATO when assessing the security posture of a cloud service provider:

- Reinventing the established NIST SP 800-53 standard and FedRAMP assessment and authorization program would be a significant undertaking.
- FedRAMP is already in place, and it is an adopted framework for US Government agencies when assessing cloud service providers.
- In the United States, FERC approves NERC CIP standards. As a US federal agency, FERC relies on FedRAMP when assessing cloud service providers for their own cloud computing needs. Given that NERC is interested in harmonizing CIP standards with cloud computing standards, the choice of FedRAMP as a compliance path for cloud service providers would seem logical.
- The program relies on an in-depth audit with mandatory provisions for continuous monitoring, and it provides strong assurances to customers that audited controls are operating effectively.

At the Emerging Technology Roundtable on Cloud Computing, NERC had requested that industry member utilities formally communicate to NERC the desire to **consider FedRAMP authorization as a possible compliance path for cloud vendors.** Note that existing NERC CIP compliance obligations would remain unchanged, and they would still be the responsibility of registered entities.

Background screening

Background screening requirements are in NERC CIP-004-6 under R2 (formal training), R3 (personnel risk assessments) and R4 (access authorization). Requirements are enforced on support and operations personnel with access to NERC CIP protected assets and data. Registered entities have written these requirements into their policies under the goals provided by NERC CIP standards. Some registered entities have written requirements for restriction on data access to US persons into their policies as well. Nuclear electric utility companies may additionally be subject to export control requirements mandated by the Department of Energy (DoE) under 10 CFR Part 810 and administered by the National Nuclear Security Administration (NNSA). Among other things, these requirements are in place to prevent the export of unclassified nuclear technology and assistance to foreign persons. The Nuclear Regulatory Commission (NRC) regulates the export and import of nuclear facilities and related equipment and materials under 10 CFR Part 110. The NRC does not regulate nuclear technology and assistance related to these items which are under the DoE jurisdiction. Consequently, NRC 10 CFR Part 110 regulations would not be applicable to Azure. **Nuclear utility customers are wholly responsible for ensuring their own compliance with all applicable laws and regulations. The forgoing is not legal advice, and customers should consult their legal advisors for any questions regarding regulatory compliance.**

All Azure and Azure Government employees in the United States are subject to the Microsoft cloud background check, as outlined in Table 4. Personnel with the ability to access customer data in Azure Government are additionally subject to the verification of US persons, as well as the National Agency Check with Law and Credit (NACLC) that involves fingerprint background checks against an FBI database.

Table 4: Background screening for Azure and Azure Government personnel

Applicable screening and background check	Environment	Frequency	Description
Microsoft cloud background check	Azure Azure Gov	Upon employment	• Education history (highest degree) • Employment history (7-yr history)
		Every 2 years	• Social Security Number search • Criminal history check • Office of Foreign Assets Control (OFAC) list • Bureau of Industry and Security (BIS) list • Directorate of Defense Trade Controls (DDTC) debarred list
National Agency Check with Law and Credit (NACLC)	Azure Gov	Every 5 years	• Ads fingerprint background check against FBI database
US persons	Azure Gov	Upon employment	• Verification of US person

Information security training and awareness is provided to all Azure and Azure Government engineering personnel on an ongoing basis to educate them about applicable policies, standards, and information security practices. All engineering staff are required to complete a computer-based training module when they join the team. In addition, all staff participate in mandatory security, compliance, and privacy training administered annually. Training is also covered by controls in many compliance certifications and attestations applicable to Azure and Azure Government.

As discussed in the storage isolation section, it’s important to note that Azure and Azure Government personnel do not have persistent access to customer data by default. Access to customer data is not needed to operate Azure and Azure Government. Customers who require technical assistance can open a troubleshooting ticket and authorize access to their data if needed. All actions taken by support personnel are logged and audited.

Logical isolation considerations

A multi-tenant cloud platform implies that multiple customer applications and data are stored on the same physical hardware. Azure and Azure Government use 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.

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 or Azure Government 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 or Azure Government 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 or Azure Government 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 and Azure Government services such as compute, storage, and networking.

Compute isolation

- A customer cannot intercept VM traffic that belongs to another customer
- When VMs belonging to multiple customers are deployed on the same node, it is not possible for one VM to starve neighboring VMs of compute resources
- Customer VMs cannot launch denial of service attacks against other VMs
- Azure provides VM instances that are isolated to hardware dedicated to a single customer

Microsoft Azure and Azure Government compute platforms, which includes Web Roles, Worker Roles, and Virtual Machines, are 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. Every physical node in Azure and Azure Government has one or more virtual machines, called instances, scheduling them on physical CPU cores, assigning them dedicated RAM, and granting and controlling access to local disk and network I/O.

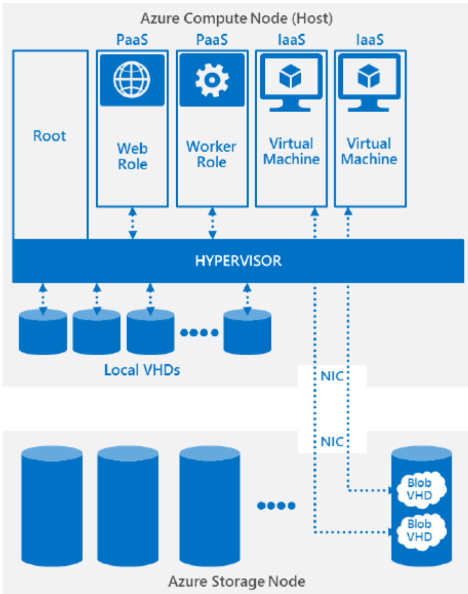

Figure 8: Isolation of Hypervisor, Root VM, and Guest VMs

On each Azure or Azure Government 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 (see Figure 8). Isolation of the Root VM from the Guest VMs and the Guest VMs from one another is a key concept in Azure and Azure Government security architecture.

Figure 9 shows a simplified view of a server's software architecture. The host partition (also called the root partition) runs the Server Core profile of Windows Server as the Host OS. The only difference between this diagram and a standard Hyper-V architecture diagram is the presence of the Azure Fabric Controller Host Agent in the host partition and the Guest Agents in the guest partitions (Russinovich, 2012). The Fabric Controller is the brain of the Azure and Azure Government compute platforms, and the Host Agent is its proxy, integrating servers into the platform so that the Fabric Controller can deploy, monitor, and manage the virtual machines that define Azure and Azure Government Cloud Services. By default, only PaaS roles have Guest Agents, which are the Fabric Controller's proxies for providing runtime support and monitoring the health of the roles. For IaaS VMs, customers can choose to install a VM agent that can be used to bootstrap VM Extensions offered by both Microsoft and partners for configuring, managing, and accelerating VMs.

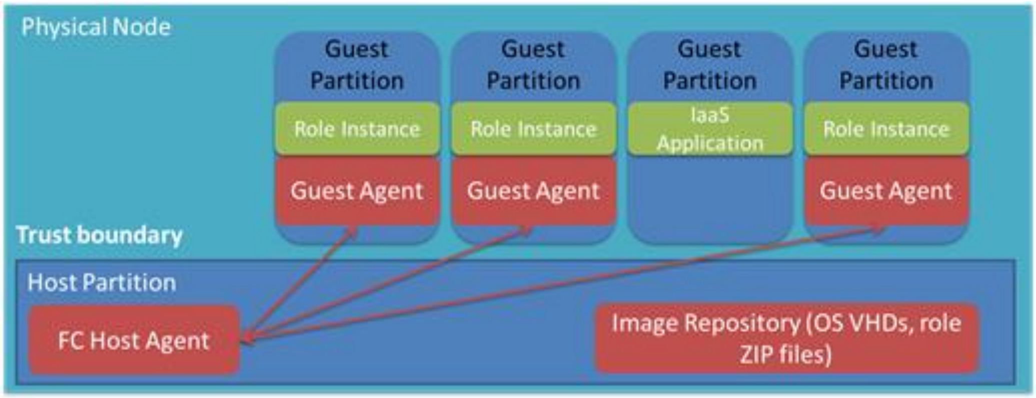

Figure 9: Host and Guest Agents deployed in Azure nodes

A critical boundary is the isolation of the Root VM from the Guest VMs and the Guest VMs from one another, managed by the Hypervisor and the Host OS. As discussed in the networking isolation section, the Hypervisor and the Host OS provide network packet filters that help assure untrusted virtual machines cannot generate spoofed traffic, cannot receive traffic not addressed to them, cannot direct traffic to protected infrastructure endpoints, and cannot send or receive inappropriate broadcast traffic (Kaufman and Venkatapathy, 2010).

As the central orchestrator of much the Azure Fabric, significant controls are in place to mitigate threats to Fabric Controllers, especially from potentially compromised Fabric Agents within customer applications. As part of the defense-in-depth design, Fabric Controllers are protected even against trusted components within Azure and Azure Government to prevent any compromise of those components from being leveraged to compromise the entire fabric.

Communication from the Fabric Controller to Fabric Agent is unidirectional – the Fabric Agent implements an SSL-protected service that is accessed by the Fabric Controller, and it's designed to only reply to requests. A Fabric Agent cannot initiate connections to the Fabric Controller or other privileged internal nodes. The Fabric Controller strongly parses all responses as though they were untrusted communications. In addition to cryptographically authenticating all incoming requests, Fabric Controller will only accept incoming requests from a limited set of IP addresses, which explicitly excludes requests from the Internet or from customer guest VMs. Finally, the main VLAN that interconnects untrusted customer nodes is completely separate from the VLANs that host Fabric Controllers and network and infrastructure devices. In this manner, the authentication interfaces for the Fabric Controllers and Devices have limited exposure to a compromised node that hosts customer VMs.

A common customer concern in a multi-tenant public cloud relates to the possibility of resource starvation when multiple VMs belonging to different customers share the same physical server. Fabric Controller ensures that a customer VM cannot consume more resources than the customer paid for, which in turn prevents a VM from starving neighboring VMs of compute resources. Also, it is not possible for a customer VM to launch denial of service attacks on neighboring VMs or other parts of the Azure infrastructure. Fabric Controller monitors for these types of attacks and removes offending VMs from Azure.

Azure GS-5, G5, DS15 v2, and D15 v2 VM instances allow customers to be deployed on hardware dedicated to a single customer. Although the primary purpose of these VMs is to support enterprisegrade applications that demand faster CPUs, better local disk performance, or have higher memory demands, they can also be used by customers who prefer VM instances that are isolated to hardware dedicated to a single customer.

A core principle of security is that by presenting a smaller attack surface to an adversary you stand a better chance of defending yourself. A compromised physical machine usually has only one way to get to other machines, i.e., through the network. Virtual machines (VMs) can have this same security model by limiting channels out of the VM and helping ensure there is only a narrow attack surface in your virtualization software. Keeping a compromised virtual machine isolated requires a great deal of rigor and accuracy in the VM monitor (i.e., Hypervisor), and all the software in the host that interacts with a VM. The use of sound security practices reduces the risk of compromise of these components and provides greater assurance that a virtual machine stays isolated.

Storage isolation

- A customer cannot read data belonging to another customer
- Data deleted by a customer cannot be accessed by another customer
- Data destruction and disposal follows NIST SP 800-88 R1
- Azure administrators do not have access to customer data by default

Microsoft Azure and Azure Government separate customer VM-based compute resources from storage as part of its fundamental design. The separation allows compute and storage to scale independently, making it easier to provide multi-tenancy and isolation. The resulting Flat Network Storage provides very high bandwidth network connectivity for storage clients, enabling Azure and Azure Government to support IaaS Virtual Machines where persistent disks for VMs are stored as durable network attached blobs in Azure Storage.

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. A key concept behind storage organization is a storage account, whereby a single customer may have many storage accounts. Storage Account Key controls all access to a storage account, and access to data in a specific account is granted only to entities having the secret key for that account. Storage keys are 512-bit long, and they are generated randomly when the storage account is created (or later at the request of the customer). A storage account may have two active keys at any given time to support key rollover. Customers can use Azure Active Directory and Role-Based Access Control to restrict access to the Storage Account Key only to specific users. Note that access to Azure SQL Database features standard SQL authentication, i.e., connection string with user name and password.

Azure and Azure Government drives, disks, and images are all Virtual Hard Drives (VHDs) stored as page blobs within customer's storage account (see Figure 10). A VHD can be uploaded into a storage account and used for either PaaS or IaaS; however, access from customer compute VMs will differ for PaaS versus IaaS. Drives are used by the PaaS roles (Worker role, Web role) to mount a VHD and assign a drive letter. They are implemented with a kernel mode driver that runs within the VM and communicates with storage using VM's virtual network adapter. Disks are attached to IaaS Virtual Machines to persist operating system as well as any additional data needed by the application. The code that communicates with storage is not within the VM; it must pass through the Hypervisor and then into the Host OS before access to storage can be granted.

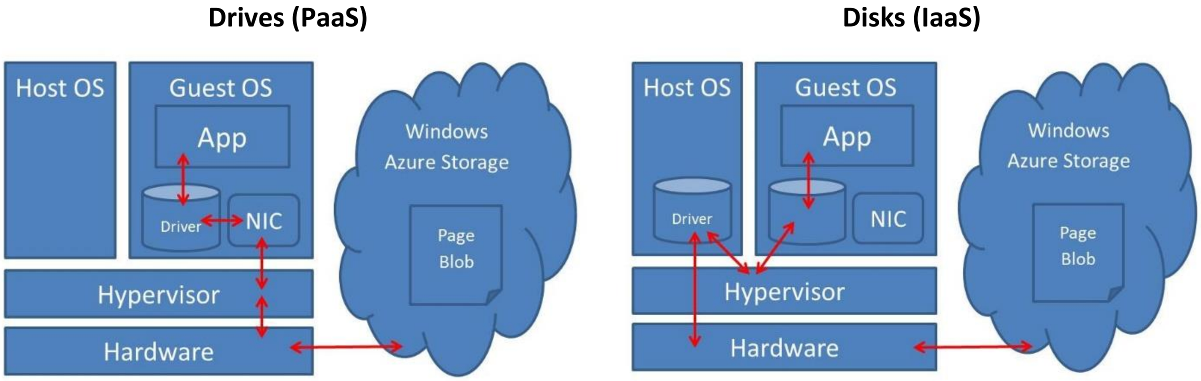

Figure 10: IaaS disks and PaaS drives are stored VHD files in Azure Blob storage

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 (Myers, 2014).

When a 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. At the primary location, a customer can immediately try to access the blob or entity, and they won't find it in their index, since both Azure and Azure Government provide strong consistency for the delete. So, the customer can verify directly that the data has been deleted.

A customer cannot read deleted data of another customers. 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.

For data disposal, 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 these 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."

A common customer concern is related to potential access to customer data by Azure and Azure Government administrators. With respect to access management, it's important to note that Azure and Azure Government personnel do not have persistent access to customer data by default. Access to customer data is not needed to operate Azure and Azure Government. Should customers request assistance (e.g., via a troubleshooting ticket), Azure support engineers are provided with just-in-time access using temporary credentials. Access to customer data is revoked as soon as the issue is resolved, and all actions taken by support personnel are logged and audited. Controls for the protection of customer secrets are audited on a regular basis as part of existing Azure and Azure Government audits, including SOC 2 Type 2, FedRAMP Moderate, and FedRAMP High. Customers also have several options for encrypting their data at rest, including keeping encryption keys in 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 or community 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. Virtual networks (VNET) in Azure and Azure Government help ensure that each customer's private network traffic is logically isolated from traffic belonging to other customers. A customer subscription can contain multiple logically isolated private networks, and include firewall, load-balancing, and network address translation (see Figure 11). These sub-divided networks generally fall into one of the two categories:

- **Deployment network:** Each deployment can be isolated from other deployments at the network level. Multiple VMs within a deployment can communicate with each other through private IP addresses.
- **Virtual network:** Each virtual network is isolated from other virtual networks. Multiple deployments inside the same subscription can be placed on the same virtual network, and then communicate with each other through private IP addresses.

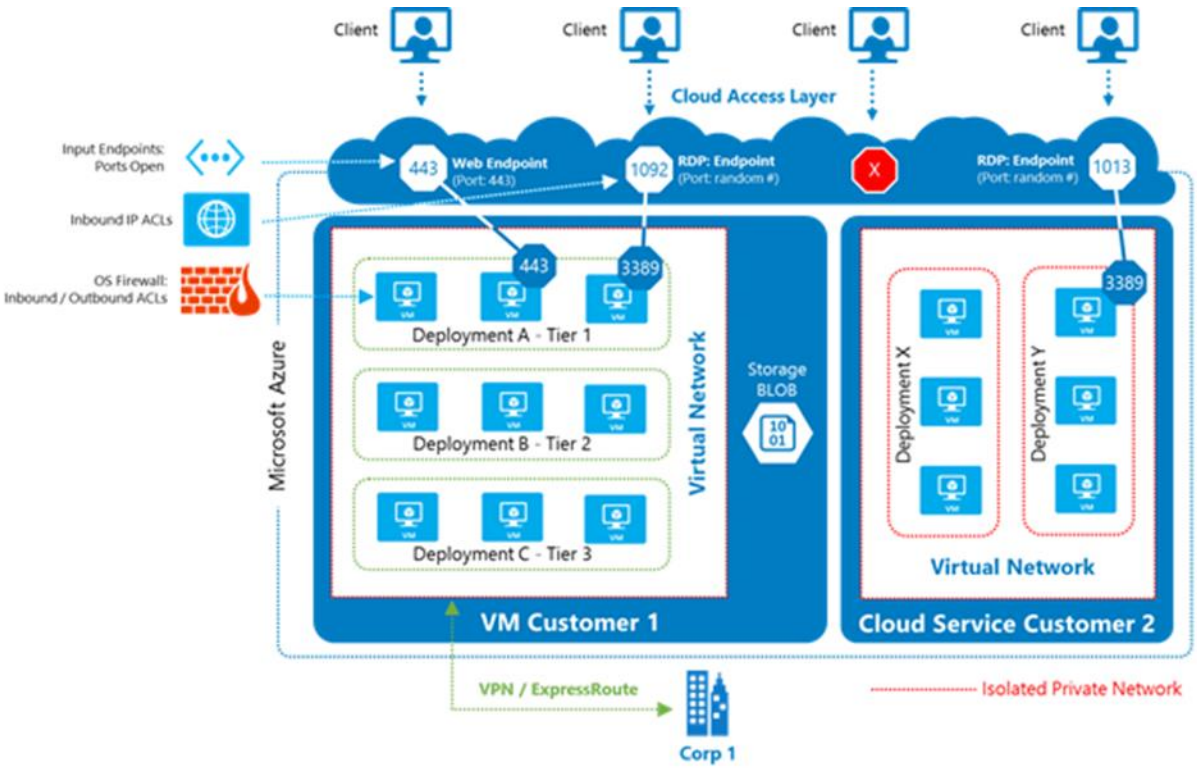

Figure 11: An example of a virtual network topology

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 and Azure Government provide 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.
    a) 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.
    b) VMs cannot capture any traffic on the network that is not destined to them.
- Customer VMs cannot send traffic to Azure or Azure Government 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 and Azure Government 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 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 isolation with connectivity to customer intranets

Virtual Networks provide 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 and Azure Government or other machines on a customer's own network (presumably behind the customer's datacenter firewall). With VNETs, 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 or Azure Government 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. The customer end of this tunnel can be implemented either by configuring on-premises routers with tunnel endpoint information, or with a software based relay (in the case where traffic load is too light to justify a hardware investment). Azure VMs connected to VNETs can be domain-joined to customer domain controllers in order to be managed consistently with the customer's on-premises resources.

Security considerations

Table 5 provides a summary of key security considerations for physically isolated on-premises deployments (e.g., bare metal, which may be familiar to electric utility customers) versus logically isolated cloud-based deployments (e.g., Azure and Azure Government). It's useful to review these considerations prior to examining risks identified to be specific to shared cloud environments.

Table 5: Key security considerations for physical versus logical isolation

Security Consideration	On-Premises with Bare Metal	Azure and Azure Government
Firewalls, networking	• Physical network enforcement (switches, etc.) • Physical host-based firewall can be manipulated by compromised application • 2 layers of enforcement	• Physical network enforcement (switches, etc.) • 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	• Large hardware attack surface exposed to complex workloads, enables firmware based advanced persistent threat (APT)	• 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	• Side channel attacks may be a factor, although reduced vs. shared hardware	• Side channel attacks assume control over VM placement across applications; may not be practical in large cloud service
Patching	• Varied effective patching policy applied across host systems • Highly varied/fragile updating for hardware & firmware	• Uniform patching policy applied across host and VMs
Security analytics	• Security analytics dependent on host-based security solutions, which assume host/security software has not been compromised	• Outside VM (hypervisor based) forensics/snapshot capability allows assessment of potentially compromised workloads
Security policy	• Security policy verification (patch scanning, vulnerability scanning, etc.) subject to tampering by compromised host • Inconsistent security policy applied across customer entities	• Outside VM verification of security policies • Possible to enforce uniform security policies across customer entities
Logging and monitoring	• Varied logging and security analytics solutions	• Common Azure platform logging and security analytics solutions • Most existing on-premises / varied logging and security analytics solutions also work
Malicious insider	• Persistent threat caused by system admins having elevated access rights typically for the duration of employment	• Greatly reduced threat because admins have zero standing access rights by default

Listed below are key risks that are unique to shared cloud environments that may need to be addressed when accommodating data and workloads subject to NERC CIP standards.

Exploitation of vulnerabilities in virtualization technologies, interfaces to external systems, APIs, and management systems

Compared to traditional on-premises hosted systems, both Azure and Azure Government provide 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 significantly better **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 and Azure Government that make it unlikely that such an attack would be successful:

- The standard Azure and Azure Government cryptographic libraries have been designed to resist such attacks by not having cache access patterns depend on the cryptographic keys being used.
- All Azure and Azure Government 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).

Zero standing access rights and Just-in-Time (JIT) access provisions eliminate the risks associated with traditional on-premises administrator 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.

Residual risk that may or may not be acceptable to NERC customers even with proper configuration of the virtual and physical environment to mitigate threats.

NERC regulated utility customers should identify and review all residual risks and consider whether Azure or Azure Government provide a suitable solution. Specific mitigation paths for identified risks can also be discussed during the customer onboarding process.

Summary

Microsoft Azure and Azure Government are multi-tenant cloud platforms available to electric power utilities and other registered entities. A multi-tenant cloud platform implies that multiple customer applications and data are stored on the same physical hardware. Azure and Azure Government use 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. Table 6 summarizes key considerations for cloud adoption. Both Azure and Azure Government are suitable for registered entities deploying certain workloads subject to NERC CIP standards enforcement.

Table 6: NERC CIP considerations for cloud computing

Requirement	Azure	Azure Gov
Data or workload subject to NERC CIP standards	✓	✓
Data must reside in continental United Status	✓	✓
Security controls mapped to CSA CCM v3.0.1	✓	✓
FedRAMP Moderate authorization (325 controls)	✓	✓
FedRAMP High authorization (421 controls)	✗	✓
Support for DoE export control requirements	✗	✓
Microsoft cloud background check	✓	✓
Background screening with fingerprint check	✗	✓
Require US persons for operational personnel	✗	✓

Both Azure and Azure Government have comprehensive security controls and compliance coverage to provide robust assurances to customers about the safeguarding of customer data and applications. Azure Government is a US government community cloud that is physically separated from the Azure cloud. It provides additional assurances regarding US Government specific background screening requirements, including US person verification for Azure Government operation personnel with potential access to customer data. Moreover, Azure Government is only available in the United States to US-based registered entities.

Nuclear electric utilities may also be subject to the DoE 10 CFR Part 810 export control requirements on unclassified nuclear technology and assistance. Azure Government is designed to meet specific controls regarding access to information and systems by US persons. This commitment is not applied in Azure so customers deploying in Azure should consider whether additional technical measures, such as data encryption, should be taken to help secure data that should not be disclosed to foreign persons.

Registered entities subject to NERC CIP compliance obligations can leverage existing audits when assessing the security posture of a cloud service provider, including Cloud Security Alliance STAR program, SOC 2 Type 2 attestation, and FedRAMP authorization. FedRAMP may be considered by NERC when harmonizing CIP standards and cloud computing or assessing cloud service providers for NERC CIP data and workloads. Customers contemplating a NERC audit should review Microsoft's Cloud Implementation Guide for NERC Audits available to customers under a non-disclosure agreement from the Service Trust Portal. They can also engage Microsoft for audit assistance, including furnishing Azure or Azure Government audit documentation and control implementation details in support of NERC audit requirements. Registered entities are ultimately responsible for meeting their NERC CIP compliance obligations.