
Data Residency

AWS Policy Perspectives

July 2018



[Policy Perspectives]



© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.



Contents

Introduction.....	1
Why Data Residency Does Not Provide Better Security	2
Why the Cloud Does Not Impact Compelled Access Risk.....	4
Limiting Compelled Access.....	5
Why Unauthorized Access Risk is Lower in the Cloud.....	7
Mitigating Unauthorized Access.....	7
Hyperscale Cloud: A Transformational Approach to Security	9
CSP Responsibility: Native Security in the Cloud.....	11
Customer Responsibility: Secure Architecture Approach	11
Roles for Data Protection.....	12
Aligning Security Policy, Digital Transformation, and Economic Growth	14
Commercial and Public Sector Challenges with Data Residency	14
Considerations in Establishing Data Residency Policies.....	17



Introduction

In today's complex computing environment, public sector organizations continue to have legitimate concerns about the security of their data. As a result, some governments have determined that mandating data residency – the requirement that all customer content processed and stored in an IT system remain within a specific country's borders – provides an extra layer of security. Data residency reflects a combination of issues primarily associated with perceived (and in some cases real) security risks around third-party access to data, including foreign law enforcement agencies. Public sector customers want the assurance that their data is protected from unwanted access from not only nefarious attackers, but also other governments.

A position of stringent data residency sometimes restricts the use of large-scale, multi-national cloud service providers (CSPs), often called "hyperscale" CSPs. General cybersecurity concerns, as well as concerns about the potential overreach of government surveillance by some countries, have contributed to a continued focus on keeping data in-country. However, this restriction is counter-productive to the objective of effectively securing public sector data. As discussed below, a hyperscale CSP, which may be located out-of-country, provides their entire customer base with the ability to achieve high levels of data protection through protections of their own platform and with turn-key tooling for their customers. They do this while at the same time preserving nation-state regulatory sovereignty.

Hyperscale cloud services represent a transformational disruption in technology because of the high degree of efficiency, agility, and innovation to provide world-class security to support their customers. Hyperscale CSPs architect, operate, and maintain offerings to enable customers across multiple sectors (commercial, public, regulated) to address some of the most prevalent vulnerabilities and security risks. Customers rely on a hyperscale CSP's offerings to embody security practices that are dynamic and responsive to real-time threats, dramatically improving every customer's security posture. CSPs, for their part, have all the right incentives to maintain world-class cybersecurity as they would face substantial long-term consequences - including impacts associated with a system compromise, loss of customer trust, and brand damage. In other words, best-of-breed security is mandatory for a successful CSP. Security must be fully integrated into the design, development, and operations of hyperscale cloud services.

This paper will address the following:

- Real and perceived security risks expressed by governments when they demand in-country data residency.
- Commercial, public sector, and economic impact of in-country data residency policies with a focus on government data
- Considerations for governments to evaluate before enforcing requirements that can unintentionally limit public sector digital transformation goals leading to increased cybersecurity risk.



Why Data Residency Does Not Provide Better Security

Ownership and the geographic placement of data have become a major topic for cybersecurity and cloud policy initiatives around the globe. Historically, command and control over sensitive enterprise data meant housing the information locally on-premises or in physically accessible contractor-owned facilities within a country. Having full ownership of the “stack,” all the way from the building floor and walls to the software on the servers, made people feel comfortable that their data was as secure as possible. This rationale still exists for many governments.

As technology has evolved, three fundamental realities have disrupted the traditional “full stack control” model:

1. Most Vulnerabilities are Exploited

Remotely. The physical location of data has little to no impact on threats propagated over the Internet. Internet-connected systems expose an organization to a broad threat space, all of which are

Regardless of the physical location, if IT systems are in any way connected to the Internet (or other multi-party networks), even indirectly, they are at considerable risk.

propagated from any location. For instance, the recent Petya ransomware affected health care services, debilitating their operations and ability to perform patient care. This was a result of malware affecting their on-premises data center spread through the Internet. Despite a massive amount of effort to secure interconnected systems via firewalls and other anti-intrusion devices, experience has shown that perimeter security is a very small part of a protected system. Regardless of the physical location, if IT systems are in any way connected to the Internet (or other multi-party networks), even indirectly, they are at considerable risk and susceptible to a wide range of logical access threats.

2. Manual Processes Present Risk of Human Error. Human process failure plays a role in root cause failure (if not the entire cause) of most cybersecurity. A common example is a failure to patch vulnerable systems with published software updates for many months prior to an exploit. The manual process of updating systems with the latest patches is difficult and is not feasible to do regularly without automation.

3. Insider Threats Prevail as a Significant Risk. The vast majority of major data compromises have occurred either through unintentional errors or intentional malicious behavior by individuals using authorized accounts that had the right to access the data. The high-profile breaches of the last few years were largely attributed to poor cyber hygiene practices. The most common authorized account threat scenarios include:

- Inadvertent: credentials that are lost or mismanaged such that an attacker can act within a system as a valid user.
- Social engineering: phishing attacks and social engineering attacks that trick users or administrators into disclosing credentials to attackers.
- Malicious: classic insider threat – bad actors within the organization with nefarious intent.



Physical location of the data has no bearing on any of the above listed realities.

In today's climate, risk management is an even more formidable task when considering mobile technology and considering the interrelationships between external and internal entities. Any system connected to the Internet, directly or indirectly, presents a credible attack vector, without regard for the physical location of the infrastructure or system. As technology continues to advance and change customer threat vulnerabilities and vectors, governments must re-evaluate how they are modeling their strategies and risk tolerance. Real world examples have shown that storing data on your own servers, in your own datacenter, in your own country, is by no means an adequate basis for securing your data.

For example, a high-profile breach of a U.S. government agency affecting more than 20 million federal employees occurred in an on-premises environment as a result of compromised user credentials. These credentials were compromised and used over the wire from various locations - bypassing all protections the on-premises environment offered. The U.S. government agency breach is a good example of threats emanating over the Internet without geographic bounds.

This issue applies to more than just Internet-facing systems. Systems that do not have a direct connection to the Internet give users access via Virtual Private Network (VPN) connections from laptops, home computers, or mobile devices. Breaches do not require physical access to a server but instead exploit a lack of effectively implemented logical security controls. This demonstrates that data residency requirements have little relevance protecting information from today's most prevalent threats. Geographic locality requirements, therefore, have little relevance to protecting information from today's threats. Instead, the best mechanisms to protect, detect, respond, and recover is to use the transformational security a hyperscale CSP offers through modernization and automation. Hyperscale CSPs, like AWS, invest in and reflect technical and operational security best practices because it is core to their operations and offerings. Customers benefit when they leverage a CSP like AWS's infrastructure and cloud offerings.

Both Gartner¹ and IDC², two leading IT research organizations, concluded that the security posture of major CSPs is equal to or better than the best enterprise data centers and that security should no longer be considered a primary inhibitor to the adoption of public cloud services. In fact, enterprises actually benefit from the security native in the cloud.

1 <http://www.gartner.com/smarterwithgartner/is-the-cloud-secure/>

2 Pete Lindstrom, "Assessing the Risk: Yes, the Cloud Can Be More Secure Than Your On-Premises Environment," International Data Corporation (July 2015).



Why the Cloud Does Not Impact Compelled Access Risk

For some governments, data residency requirements are intended to mitigate risks related to another entity's access to their data. This section aims to address the perceived risk of an entity's ability to "compel access" to a sovereign entity's data when that data is stored within a hyperscale CSP environment. The concept of compelled disclosure or compelled access refers to access rights to data by governments or their agents using lawful means. The result of these types of access rights can mean that companies are subject to applicable laws and regulations at the national, provincial, and sector levels in any given country, and may be compelled to grant access or hand over data under such laws or regulations. The perceived concern is that compelled disclosure may potentially leave the data owner with no ability to prevent access to its data by an entity purporting to invoke applicable law. However, a sovereign nation's lawful access to data is not a cloud-specific issue.

Owning the physical system, either directly or through an outsourced contract, does not reduce the risk of compelled access because there are already other legal mechanisms in place that give governments in one jurisdiction the means to request access to data stored in another jurisdiction. For example, Mutual Legal Assistance Treaties (MLATs)³ and Letters Rogatory⁴ have been in place to govern a sovereign nation's requests for data long before the emergence of cloud technology.

As compared to a traditional on-premises environment, law enforcement must generally overcome more barriers when attempting to compel a CSP to disclose another customer's data. Law enforcement cannot search or seize data stored in a CSP's servers without abiding by the legal frameworks supporting a narrowly targeted set of law enforcement purposes. Further, CSPs can challenge requests that are overbroad, exceed the requestor's authority, or do not fully comply with applicable law.

More importantly, CSPs like AWS are fully committed to providing affected customers with notice of data requests, enabling the customer to engage with authorities and/or take further appropriate action to prevent against improper disclosure of its data. It is important to recognize that this complex challenge is not unique to the U.S. government or U.S.-based companies, because any multi-national company is subject to applicable laws and regulations at the national, provincial, and sector levels in any given country regardless of the location of data.

³ Mutual Legal Assistance Treaties (MLATs) generally allow for the exchange of evidence and information in criminal and related matters. <https://www.state.gov/j/inl/rls/nrcrpt/2012/vol2/184110.htm>

⁴ Letters Rogatory are requests from courts in one country to the courts of another country requesting the performance of an act which, if done without the sanction of the foreign court, could constitute a violation of that country's sovereignty. Letters Rogatory may be used to effect service of process or to obtain evidence if permitted by the laws of the foreign country. <https://travel.state.gov/content/travel/en/legal/travel-legal-considerations/international-judicial-assistance/obtaining-evidence/Preparation-Letters-Rogatory.html>



Limiting Compelled Access

Since the 20th century, many countries have had legal mechanisms in place to enable access to information stored abroad in response to appropriate lawful requests for information relating to criminal investigations and prosecutions. For example, a company doing business in Country X could be subject to a legal request for information even if the content is stored in Country Y under established bi-lateral and multi-lateral legal frameworks. In most instances, the recognized legal mechanism is a Mutual Legal Assistance Treaty (MLAT).

In addition to bi-lateral country MLATs, there are also key regional MLATs such as the Inter-American MLAT, the EU-US MLAT, and the ASEAN MLAT. In the absence of an MLAT, countries can obtain Letters Rogatory to seek assistance from foreign governments. Each jurisdiction's law will contain criteria that must be satisfied in order for the relevant law enforcement body to make a valid request. For example, the government agency seeking access may need to obtain a court order or warrant showing it has a valid reason for requesting access to content. While legitimate mechanisms, these legal instruments were not intended to address law enforcement access to data in a digital world.

In an effort to align asynchronous laws with modern technology, the U.S. passed the Clarifying Lawful Overseas Use of Data (CLOUD) Act in March 2018. The CLOUD Act provides a third legal international mechanism to acquire data stored overseas through direct requests issued to the service provider.⁵ The CLOUD Act sets forth procedures for the U.S. to enter into Executive Agreements with other countries. These Executive Agreements seek to remove legal restrictions on certain foreign nations' ability to seek data directly from U.S. providers, provided that the U.S. has determined that the foreign nation's laws adequately protect privacy and civil liberties. Under the CLOUD Act, CSPs have the right to resist disclosing information if doing so would conflict with another country's laws. The MLAT, Letters Rogatory, and Executive Agreements under the CLOUD Act all provide reciprocal international legal mechanisms for law enforcement access to data stored overseas.

Laws governing access to data stored abroad by law enforcement agencies in support of investigating serious crimes, such as terrorism, were not written with modern technology in mind. This resulted in cases where technology companies complying with a judicial warrant under one country's laws also faced the risk of violating another country's laws prohibiting disclosure. The CLOUD Act provides a new framework for challenging law enforcement requests when there are executive agreements in place between the U.S. and another country, and it also confirms, under principles of comity between nations, the right of service providers to resist disclosure of any data if doing so would conflict with another country's laws, even in the absence of an executive agreement. It also enables cloud service providers to disclose data to governments issuing orders or warrants for information based on sufficient facts demonstrating probable cause that a serious crime has occurred and that the information sought is directly related to that crime.

⁵ The CLOUD Act applies to both U.S. and foreign companies operating in the United States that provide "electronic communications services" and/or "remote computing services," such as businesses that offer email, electronic messaging, or cloud storage services to the public.



The national laws of a country generally apply to all companies with operations in that country, irrespective of where the company is incorporated or whether the information is stored in the cloud, an on-site data center,

Restricting CSPs to one jurisdiction does not better insulate data from governmental access

An [independent legal analysis](#) across early government cloud adopters assessed the country-specific laws governing law enforcement access to cloud-based data stored abroad. This study evaluated ten international jurisdictions- Australia, Canada, Denmark, France, Germany, Ireland, Japan, Spain, UK and U.S.- and found that restricting CSPs to one jurisdiction does not better insulate data from governmental access.

or in physical records. As nations continue to digitalize and advance towards modern information-based societies, lawful compelled access regimes in support of investigations for high crimes impacting national security, such as terrorism, have also been evolving. The enactment of the CLOUD Act is another framework aiming to strengthen legal due process for law enforcement request in this modern context.

The reality is that such compelled access occurs in a very limited number of cases, and generally only where there is an extreme need for

information (e.g., to prevent terror- related events). To mitigate even this low risk, organizations can practice due diligence and craft their own protections with available cloud services. In AWS, mitigations such as encryption of data at rest and in transit, data decomposition and distribution, and tokenization strategies can be employed for a fraction of the resource burden versus an on-premises solution.

AWS is vigilant about protecting our customer content, regardless of where a request for content comes from or who the customer is. AWS will not disclose customer content unless required to do so to comply with a legally valid and binding order, such as a subpoena or a court order. AWS carefully examines each request to authenticate its accuracy and verify that it complies with applicable law. AWS will challenge requests that are overbroad, exceed the requestor's authority, or do not fully comply with applicable law. Unless prohibited by law, AWS also attempts to re-direct the request directly to the customer, providing the customer with an opportunity to take action against the request. Additional information can be found in our latest transparency report and our Amazon Law Enforcement Guidelines.⁶

6 http://d0.awsstatic.com/certifications/Amazon_LawEnforcement_Guidelines.pdf



Why Unauthorized Access Risk is Lower in the Cloud

For some governments, data residency requirements are intended to mitigate risks related to another entity's access to their data. This section aims to address the perceived increase in unauthorized access risk when using a hyperscale CSP. Unauthorized access is the more common threat attempted by adversaries trying to gain access to customer data using various means. Unauthorized access can include third party access concerns, including the possibility of insider threats or outside bad actors.

Data residency requirements fail to address the common avenues attackers use to gain access. Exploiting these vectors almost always results from a failure in basic cyber hygiene disciplines, such as system inventory management, configuration management, data encryption, and privileged access management.

Mitigating Unauthorized Access

Preventing unauthorized access requires practicing proper security hygiene and implementing robust preventive and detective capabilities. For example, systems should be designed to limit the “blast radius” of any intrusion so that one compromised node has minimal impact on any other node in the enterprise. Hyperscale CSPs, such as AWS, provide a full security tooling environment to enable customers to maintain encrypted communications and implement tampering protections to mitigate the risk of unauthorized access. AWS does not have visibility into, or knowledge of, the contents of a customer account, including whether or not that content includes any personal information. AWS customers are empowered to use various techniques such as encryption,⁷ tokenization, data decomposition, and cyber deception to render content unintelligible to AWS or other parties seeking access to its content.

- Encryption - Appropriately encrypting data can make the data unreadable. This means storing encrypted data in the cloud, regardless of location, can provide adequate protection against the vast majority of exfiltration threats. It is crucial that the encryption keys for the data are carefully managed to ensure strong protections are maintained against any intercepting party. AWS provides services that can deliver these capabilities at an enterprise level with AWS CloudHSM or AWS Key Management Service (KMS).⁸ The amount of control that customers wish to have over the encryption method, storage of cryptographic keys, and management of cryptographic keys used with their data is up to the customer.⁹

7 AWS allows customers to use their own encryption mechanisms for nearly all AWS services, including Amazon S3, Amazon EBS, Amazon DynamoDB, and Amazon EC2. IPSec tunnels to VPC are also encrypted. Amazon S3 also offers Server Side Encryption as an option for customers. Customers may also use third-party encryption technologies.

8 The AWS CloudHSM (Hardware Security Module) service allows you to protect your encryption keys within HSMs designed and validated to government standards (FIPS 140-2 Level 3) for secure key management including robust tamper protection. AWS KMS, validated at FIPS 140-2 Level 2, provides a similar service, but one that is more scalable and more deeply integrated with a wide range of AWS services such that protections are provided automatically based on simple changes in service configuration. Using either service, you can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. For more details see <https://aws.amazon.com/cloudhsm/> and <https://aws.amazon.com/kms/>.

9 AWS encryption options are detailed via the following links: 1) [Securing Data at Rest with Encryption](#), 2) [Protecting Data Using Encryption in Amazon S3](#), 3) [AWS Key Management Service Cryptographic Details](#), and 4) [Overview of AWS Security Processes](#).



- Tokenization – Tokenization is a process that allows you to define a sequence of data to represent an otherwise sensitive piece of information (e.g., a token to represent a customer’s credit card number). A token is meaningless on its own and cannot be mapped back to the data it represents without use of the tokenization system. Token vaults can be constructed in VPCs to store sensitive information in encrypted form while sharing tokens out to approved services for transmitting obfuscated data. In addition, AWS has a number of partners that specialize in providing tokenization services that integrate with popular databases and other storage services.
- Data Decomposition – This is a process that reduces data sets into unrecognizable elements that have no significance on their own.¹⁰ These elements or fragments are then stored in a distributed fashion so that any compromise in one node would yield only an insignificant data fragment. A particular advantage of this technique is it requires a threat actor to compromise all nodes, obtain all fragments, and know the algorithm (or fragmentation scheme) to piece together the data in a coherent way.
- Cyber Deception Defense – Cyber deception architectures and solutions can be a key component for mitigating advanced adversaries. Deception solutions can use highly sophisticated traps and decoys to present an attacker with the perception that they have infiltrated the system while in reality diverting them to a highly controlled environment. Intelligence about the attacker is gathered in order to mitigate future threats and the attack is neutralized.

A related concern to unauthorized access is third party access to customer content and the adequacy of access control measures to prevent unauthorized access by CSP personnel. Third party access to AWS systems are allocated based on least privilege, approved by an authorized individual prior to access provisioning, and supervised by an AWS employee. Duties and areas of responsibility (for example, access request and approval, change management request and approval, etc.) must be segregated across different individuals to reduce opportunities for an unauthorized or unintentional modification or misuse of AWS systems. AWS personnel with a business need to access the management plane are required to first use multi-factor authentication, distinct from their normal corporate Amazon credentials, to gain access to purpose-built administrative hosts. These administrative hosts are systems that are specifically designed, built, configured, and hardened to protect the management plane. All such access is logged and audited. When an employee no longer has a business need to access the management plane, the privileges and access to these hosts and relevant systems are revoked. AWS has implemented a session lock out policy that is systematically enforced. The session lock is retained until established identification and authentication procedures are performed.

AWS also monitors for unauthorized remote management and expeditiously disconnects or disables unauthorized remote access once it is detected. All remote administrative access attempts are logged, and the logs are reviewed, not just by humans for suspicious activity, but also by automated machine-learning systems built by the AWS security team to detect unusual access patterns that may indicate unauthorized attempts to access data. If suspicious activity is detected, the incident response procedures are initiated. Further, AWS has established formal

¹⁰ An array of research is available on data decomposition techniques. One such report that was reviewed for this document is Data protection by means of fragmentation in various different distributed storage systems - a survey, Kapusta and Memmi, June 20, 2017.



policies and procedures to delineate standards for logical access to the AWS infrastructure and hosts. The policies also identify functional responsibilities for the administration of logical access and security. Unless prohibited by law, AWS requires that all employees undergo a background investigation commensurate with their position and level of access.

Finally, customer virtual instances are solely controlled by the customer who has full root access or administrative control over accounts, services, and applications. AWS personnel do not have the ability to log into customer instances.

Hyperscale Cloud: A Transformational Approach to Security

Leading hyperscale CSPs, like AWS, offer customers the opportunity to build adaptive and highly resilient security for their workloads. Restricting operations to specific in-country requirements would inhibit service innovation and hinder the ability to compensate for threats, such as ones that target availability. Another detrimental byproduct of in-country geographic constraints is that threat actors can gain targeting accuracy knowing the data must reside within specific areas. Hyperscale CSPs have available offerings and supporting architectures to offer both defense in depth¹¹ and defense in breadth¹² capabilities. This is due to security mechanisms being intrinsic to the design and operation of hyperscale CSP offerings.

An unintended by-product of in-country data residency requirements is that threat actors can gain better accuracy in targeting systems knowing the data resides in specific locations.

The following six items reflect the core security attributes that are an integral part of a hyperscale CSP like AWS:

1. Deep integration of security and compliance (seldom accomplished in traditional systems) means that security directly benefits from compliance because security controls are continuously monitored and updated.
2. Economies of scale apply not only to technology, but also security personnel and processes, resulting in unprecedented return on investment as compared to traditional systems.
3. The CSP takes on a major portion of the security “surface area,” executing with professional focus and skill beyond almost any customer on earth. As a result, customers can refocus their security professionals and resources on a much smaller part of the challenge such as application security.
4. The cloud provides visibility, homogeneity, and automation never seen before in traditional systems, all of which massively benefit security. This includes significantly deep auditing and logging capabilities that, for example, can record API calls that log actions taken by a CSP that may affect the customer’s account.

11 Defense in depth is the practice of implementing multiple layers of security controls to provide independence and redundancy. If one layer of controls fails the subsequent layer is available to mitigate further incursion against an asset.

12 Defense in breadth is the approach of using multidisciplinary activities to provide numerous protection mechanisms at each identified layer of defense. Generally, this means more automation and more varied security controls at each layer.



5. CSPs operate as a sort of “system container” that provides far more insight into system behavior and functioning, including security operations, providing customers with a new layer of “defense in depth.”
6. With easy and cheap access to massive amounts of storage and processing capacity, AWS customers “use the cloud to secure the cloud”, i.e., they run big data analytics on security data and log data, which provides more insight into their security posture and results in a much faster remediation of issues.

With the speed of innovation and increasing scale, the cloud security story will only get better. For example, in just the past year AWS added powerful security capabilities such as Amazon GuardDuty¹³, a managed threat detection offering that continuously monitors for malicious or unauthorized behavior; Amazon Macie¹⁴, an offering that uses machine learning to protect sensitive data; and AWS CloudHSM 2.0¹⁵, a fully managed offering that uses FIPS 140-2 Level 3¹⁶ validated hardware automatically deployed in a highly available and redundant multi-availability zone cluster that enables customers to easily generate, manage, and use their own encryption keys in the AWS Cloud while providing AWS with zero access to master keys or core encryption operations.

Encryption should be considered a core service because it can act as a means to protect data in the event other capabilities fail. It adds an additional layer of security and assurance for the confidentiality and integrity of data in transit and at rest. The combination of AWS Key Management Service (KMS) and AWS CloudHSM are the centerpiece of a rigorous encryption solution.¹⁷ Hyperscale CSPs like AWS offer ubiquitous encryption which can be out of reach for on premise operations. For example, AWS Key Management Service (KMS), FIPS 140-2 Level 2 validation, offers a Bring Your Own Keys (BYOK) option that enables customers to use their own locally generated and stored key materials with AWS services. Customers can meet specific security and compliance requirements around highly sensitive workloads with this capability as they can retain and manage their key material outside of AWS.

13 <https://aws.amazon.com/guardduty/>

14 <https://aws.amazon.com/maciek/>

15 <https://aws.amazon.com/cloudhsm/>

16 FIPS 140-2, Security Requirements for Cryptographic Modules cover 11 areas related to the design and implementation of a cryptographic module.

17 https://d1.awsstatic.com/whitepapers/compliance/AWS_Logical_Separation_Handbook.pdf



CSP Responsibility: Native Security in the Cloud

The AWS infrastructure is custom-built for the cloud, with all elements designed to intercommunicate well and present the smallest attack surface possible. In addition, the physical security controls present in our data centers have been designed to be some of the most stringent in the world. AWS architecture has been reviewed and validated against dozens of international compliance frameworks.¹⁸ We use independent third-party assessors and auditors to evaluate and attest our adherence to these regimes, and provide customers with access to the resulting reports and supporting evidence. In order to meet such a large range of security requirements, AWS builds its data centers and architecture to scale and advance with the pace of innovation. This approach has led to AWS being trusted by governments, military organizations, global banks, healthcare institutions, and other high-sensitivity organizations.

At AWS, our unique environment has been an impetus to build many of our own security tools. These tools automate a broad swath of routine tasks allowing our security experts to focus on critical aspects of safeguarding the environment. Our tooling results in security requirements that are baked-in and adhered to throughout the system development lifecycle. Common security concerns are remediated in the initial system development phases allowing our security experts to focus on mitigating advanced and complex threats at the production level.

Our security teams monitor the infrastructure all day, every day, and are well-connected to all major security watchdog groups and vendors to immediately identify potential threats. They are doing this at massive scale, which is something that sets the AWS security organization apart. By using complex algorithms to scan across millions of active customer accounts running virtually every conceivable type of workload, we can see issues that may only occur once in a billion operations multiple times a day. When we remediate the issue, we do so for the entire platform. That kind of visibility and response simply isn't achievable for the vast majority of organizations running on-premises data centers. The value that comes from focused expertise and massive scale explains why Gartner and IDC have determined that public cloud infrastructure as a service (IaaS) workloads will experience fewer security incidents than those in traditional data centers. Gartner's research estimates at least 60% reduction in security incidents.¹⁹

Customer Responsibility: Secure Architecture Approach

The security capabilities that are native to hyperscale cloud providers like AWS empower customers to create unique architectures for mitigating access risks. On-premise and similar facilities lack the homogeneity, economies of scale, visibility, and automation that can bring major security advancements. These advancements are necessary to construct highly secure systems that can counter the evolving threats seen both externally and internally. On-premises facilities struggle to employ these new operating concepts due to the resource requirements for network refactoring and new system procurement, as well as the human labor required due to the lack of software-defined infrastructure. Hyperscale CSPs build a level of agility and adaptability into

18 See <https://aws.amazon.com/compliance>

19 <https://www.gartner.com/smarterwithgartner/is-the-cloud-secure/>



their infrastructure in order to organically implement these security advancements. This means that customers can use new advancements more readily since they are natively integrated into the CSP offerings, allowing customers to craft systems using unique architectures such as micro-segmentation, polymorphic²⁰ designs, and multi-level deception networks.

For example, taking a closer look at micro-segmentation-based design on AWS, a customer can use a wide array of technology including Amazon Virtual Private Cloud (Amazon VPC), AWS Identity and Access Management (IAM), Security Groups, Network Access Control Lists, numerous encryption and logging services, along with AWS Certificate Manager to form the basis for constructing a Zero Trust Model²¹ (ZTM) network. In concept, the ZTM can provide a distinct advantage for threat mitigation and monitoring performance. Organizations have a clear need to implement a ZTM or similar security segmentation design to counter today's threats, but it is extremely difficult and expensive to construct this type of architecture in traditional enterprise environments. Moving to a public cloud provider gives organizations the opportunity to implement ZTM and similar concepts without the significant cost and resource burden associated with the physical network retrofit/build.

Roles for Data Protection

There are five important basic concepts regarding data ownership and management in the shared responsibility model:

1. Customers continue to own their data.
2. Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
3. Customers can download or delete their data whenever they like.
4. Customers can “crypto-delete” their data by deleting the master encryption keys that are required to decrypt the data keys, which are, in turn, required to decrypt the data.
5. Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit and at rest.

Data protection measures are most effectively applied after defining data handling roles to determine appropriate stakeholder roles and responsibilities. Most data protection schemes differentiate between the data controller (also referred to as “user”) and data processor and levy obligations based on those distinct roles. For example, under the EU General Data Protection Regulation, the data controller is responsible for implementing appropriate technical and organizational measures to protect data against accidental or unlawful destruction or accidental loss, alteration, unauthorized disclosure, or access. Where processing is carried out by a data processor

20 In simple terms, polymorphic design enables the creation of moving targets making it harder for adversaries execute successful attacks.

21 The concept that was originally coined by Forrester Research. It proposes that no entity on the network is trusted. The objective is to enforce secure access to all resources whether internal or external. This means that an organization must understand and classify its data and map how that data, particularly sensitive data, flows between storage, processing, transit and consumers. Then once the data is understood an organization can implement the ZTM mechanisms that enforce and automate absolute least privilege, end-to-end encryption, and full traffic inspection.



on the data controller's behalf, the data controller is also responsible for choosing a processor that provides sufficient technical and organizational measures governing the processing to be carried out. These distinctions help to delineate responsibilities between outsource providers and their customers.

As a provider of self-service infrastructure that is completely under the customers' control – including with respect to how and whether the data is processed – AWS provides the infrastructure services for customers who want to upload and process content on AWS. AWS does not have visibility into or knowledge of what customers are uploading onto its network, including whether or not that content includes any personal data. AWS customers are also empowered and encouraged to use encryption to render content unintelligible for AWS and any third party seeking to access data.

AWS services are content agnostic in that they offer the same high level of security to all customers, regardless of the type or geographical region of content being processed or stored. In other words, AWS adopts the same high security bar across all of our offerings. This means that we take the highest classification level of data traversing and stored in our commercial cloud and apply those same levels of protections to all of our offerings and for all of our customers. These offerings are then queued for certification against the security and compliance high bar, which translates to customers benefiting from elevated levels of protection for customer data processed and stored in the cloud. The AWS Cloud has been certified against numerous regulated industry (healthcare, financial, etc.), national (e.g. U.S. FedRAMP, Germany C5, Australia IRAP), and global accreditations (e.g. ISO 27001,²² ISO 27018,²³ Payment Card Industry (PCI) Data Security Standard (DSS),²⁴ and Service Organization Controls (SOC)²⁵), which test

Free flow of non-personal data proposed as the de facto EU and Trans-Pacific Regions

The EU Commission recently published a draft Regulation on the free flow of data **prohibiting national data localization rules in EU Member States** and recognizing the principle of free movement of non-personal data within the EU. This proposal establishes cross-border data flow as the de facto standard, placing the onus on Member States to provide public safety justification for imposing data localization requirements. While in the early stages of deliberation, this proposition recognizes the economic and security advantages of cross-border data flows, which are outweighing the considerations for enforcing data residency policies. Further, in early 2018, the **Comprehensive and Progressive Trans-Pacific Partnership Agreement** established among 11 countries also supports **cross-border data flows** and does not require companies to establish in-country computing facilities as a condition of doing business in that country.

22 ISO 27001/27002 is a widely-adopted global security standard that sets out requirements and best practices for a systematic approach to managing company and customer information that's based on periodic risk assessments appropriate to ever-changing threat scenarios.

23 ISO 27018 is a code of practice that focuses on protection of personal data in the cloud. It is based on ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls applicable to public cloud Personally Identifiable Information (PII). It also provides a set of additional controls and associated guidance intended to address public cloud PII protection requirements not addressed by the existing ISO 27002 control set.

24 The Payment Card Industry Data Security Standard (also known as PCI DSS) is a proprietary information security standard administered by the PCI Security Standards Council (<https://www.pcisecuritystandards.org/>), which was founded by American Express, Discover Financial Services, JCB International, MasterCard Worldwide and Visa Inc. PCI DSS applies to all entities that store, process or transmit cardholder data (CHD) and/or sensitive authentication data (SAD) including merchants, processors, acquirers, issuers, and service providers.

25 Service Organization Controls reports (SOC 1, 2, 3) are intended to meet a broad range of financial auditing requirements for U.S. and international auditing bodies. The audit for this report is conducted in accordance with the International Standards for Assurance Engagements No. 3402 (ISAE 3402) and the American Institute of Certified Public Accountants (AICPA): AT 801 (formerly SSAE 16).



and validate the security of our systems against the most rigorous standards.

Aligning Security Policy, Digital Transformation, and Economic Growth

Policies must evolve to meet the changing realities of technology and the world it helps to create. Otherwise, governments will continue to lag behind in upgrading their operations, servicing their citizens, and adopting the most modern and secure solutions. This section describes how AWS addresses the security objectives underlying data residency requirements to abate policy maker concerns. It also explores the economic and IT modernization challenges associated with data residency and policy considerations to advance secure public sector cloud adoption.

Commercial and Public Sector Challenges with Data Residency

Governments must consider how their national policies work to advance or impede economic growth and workforce development opportunities that are empowered by hyperscale cloud services. There can be significant negative impacts to implementing data residency requirements, such as:

- **Adverse effect on local business multi-national commercial expansion efforts** - As businesses grow and expand outside regional operations, it is vital that they have access to resources that have a global reach. Restricting access to hyperscale CSP services severely limits the level of user experience that a business can provide to its global customer base.
- **Limited geo-redundancy options compared to global CSP regions** - For governments and businesses, ensuring redundancy in the event of operational failure due to a disaster or other circumstances is vital for stability. Having clustered operations in only one country exposes the organization to a level of risk that can far outweigh data access concerns.
- **Expensive cost structures necessary to accommodate stringent requirements** - Single tenant or community built “cloud” environments require a level of pricing for operational sustainability that can actually take away from procuring the additional capabilities needed for achieving defense in depth.

Cloud technology is the enabler for commercial and public sector advancements, and the extent to which governments promote or oppose the principle of cross-border data flows will impact the strength of their local economies as well as their global marketplace competitiveness.

Commercial Impact

Enabling the free flow of data across borders has significant net positive impact on the global economy. Recent studies by various research organizations emphasize this impact, and go further to highlight the cost to establishing barriers to data flows. A February 2016 report by McKinsey Global Institute estimated that cross-border data flows contributed nearly \$2.8 trillion to the global economy in 2014²⁶ through its enablement of

26 <http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/digital-globalization-the-new-era-of-global-flows>



the flow of goods, services, and other resources. The reports estimate that this figure could reach \$11 trillion by 2025. Governments that require localization of data and limit cross-border economic flows pay a high price. The European Centre for International Political Economy (ECIPE), an independent policy think tank, issued a study on the economic impact of data localization requirements that discriminate against foreign suppliers in seven jurisdictions: Brazil, China, EU, India, Indonesia, South Korea and Vietnam.²⁷ Their research concluded that unilateral restrictions on cross-border data flow and access to foreign markets negatively impacts economic growth and recovery because it limits access to competitive pricing, job growth in many services and goods sectors, and investment opportunities. The study noted that data residency requirements not only impact data flow, but also a broader set of commercial expansion opportunities that rely on cross-border data flows.

A similar study by the World Bank studied six developing countries and the EU 28 Member States, and found that data localization requirements can reduce the GDP by up to 1.7 percent, investments up to 4.2 percent, and exports by 1.7 percent.²⁸ This impact is most felt by small-scale businesses and start-ups. Through the use of cloud, for example, individuals and small to medium enterprises (SMEs) are able to access IT resources at a cost and scale once accessible only to entities with far greater capitalization. SMEs are primary drivers for new job creation. Cloud computing lowers the barriers for business creation and market access, enabling more start-ups to form, ultimately creating more jobs. However, according to the European Commission, technology companies like a CSP can face significant costs to adapt to various national laws leading to the costs of selling online outweighing the benefits. Most recently, in May 2017, the Information Technology and Innovation Foundation, a non-partisan research institute, independently arrived at similar findings.²⁹

27 European Centre for International Political Economy (ECIPE): “The Costs of Data Localization: A Friendly Fire on Economic Recovery,” http://www2.itif.org/2015-cross-border-data-flows.pdf?_ga=1.8208626.1580578791.1473954628.

28 <http://documents.worldbank.org/curated/en/961621467994698644/pdf/102724-WDR-WDR2016Overview-ENGLISH-WebResBox-394840B-OUO-9.pdf>

29 Nigel Cory, “Cross-Border Data Flows: Where Are the Barriers, and What Do They Cost?” Information Technology and Innovation Foundation (May 2017) http://www2.itif.org/2017-cross-border-data-flows.pdf?_ga=2.243762501.1722557619.1508762047-1611916082.1508762047.



A key conclusion consistent across these studies is that prohibiting cross-border data flows in the form of data residency requirements can impact local and regional economic growth and competitiveness in the global market, with the greatest impact borne by SMEs. A secure system in the EU is no more or less secure than a similarly architected system in Latin America. Governments misunderstand that data protection does not generally depend on where the information is stored, but rather what measures are used to secure the data. Physical location generally has no relevance because data centers are almost always connected to broadly accessible networks, and thus real security depends on the technical, operational, and managerial practices and processes implemented by the CSP and the customer.³⁰

Costs of Exclusively Operating In-Country Data Centers

A 2015 study by an information security company evaluated how an in-country data center model is much more expensive compared to leveraging global CSPs. The study found that the cost of cloud services can increase substantially due to data localization, depending on the availability of alternative services. The study found that:

- If Brazil had enacted data localization as part of its “Internet Bill of Rights” in 2014, companies would have had to pay an average of 54 percent more to use cloud services (of all categories) from local cloud providers compared with the lowest worldwide price.
- If the European Union enacted data localization, companies would still have had to pay up to 36 percent more to use similar services provided by hyperscale CSPs. At the time of the study, some of the lowest-cost data centers were located in the European Union.²⁹

Public Sector Impact

Countries enforcing barriers to data flows can limit their citizens ability to take advantage of innovative services that improve their quality of life and government services delivery. For example, artificial intelligence and machine learning (AI/ML) applications require customized infrastructure for optimal functioning,³² and while global CSPs continue to expand their data center footprint, it is unrealistic to assume that data centers will be established in every country. Hence, as AI/ML is increasingly used to improve services, such as health care prognoses and weather forecasting for emergency preparedness, citizens in countries with strict data residency requirements will lag behind in accessing technological breakthroughs for citizen-related services.

30 Ibid p.4 Similar conclusions are independently drawn by this paper.

31 http://www2.itif.org/2017-cross-border-data-flows.pdf?_ga=2.51021357.566718019.1510350061-1611916082.1508762047

32 For example, systems with general-purpose GPU capabilities and Field Programmable Gate Arrays (FPGA).



There are also cascading socio-economic costs to limiting data flows, specifically on trade competitiveness and workforce development. As cloud technology becomes ubiquitous and more strongly tied to economic advancement, digital trade (and reducing the barriers to it) will become a higher priority for governments. Countries allowing free data flows will be at an advantage by accessing leading edge technology, which will in turn impact the modernization of commercial and public sector services, improve worker productivity, and accelerate local job and skills growth across sectors. Countries restricting data flows and digital trade will, in time, notice a competitive disadvantage. For instance, the full range of benefits associated with IoT to enable “smart” farming, manufacturing, or cities cannot be realized with restrictive policies that place limits on big data analytics, machine learning, or other features serviced by free yet secure movement of data.

There is continued high demand for cloud computing skills in key areas like application security, cloud enterprise application development, enterprise cloud migration, and big data. The U.S. Bureau of Labor Statistics reports that forecasted demand for jobs in information security is expected to grow at a rate of 37% between the period of 2012-2022. To meet new job demand, governments will have to invest in provide educational and training opportunities for individuals in acquiring technological skills.

Limits on access to the kinds of sophisticated IT services provided by hyperscale CSPs will also lead to a perpetual gap in developing and maintaining a highly skilled, technically savvy workforce. This is because workforce aptitude is correlated with the technological sophistication of an organization, which in turn is based on the ability of the organization to access state-of-the-art technology. The effective use of modern technology demands a workforce with commensurate skills to operate that technology. Given the breadth and pace of innovation with cloud services, there is a known and widening skills gap. Governments, in particular, have fallen behind in the race for experts who are essential to modernizing applications while at the same time protecting public sector information and systems from highly sophisticated adversaries and breaches that increase in frequency and impact.

Considerations in Establishing Data Residency Policies

As discussed above, nation-state regulatory sovereignty over data can still be achieved while taking advantage of the cost and security benefits of hyperscale CSPs like AWS. The security measures deployed throughout AWS services, and verified through our third-party audits, provide a high level of assurance to prevent and address unlawful data access risk events.

We encourage governments to consider the following policies to meet the security objectives associated with data residency.

1. Develop policies and requirements that allow for the use of out-of-country data processing facilities if the data is processed and stored in a modern, highly secure, hyperscale cloud environment. Customers can also choose locations with data protection laws consistent with their own and where data transfer agreements are already in place.
2. Align national policies and regulatory requirements to the principle of free movement of data cross-border to effectively balance security, economic, and IT modernization goals.



3. Evaluate data transfer models, such as the US-EU Privacy Shield, and standardized contractual clauses, such as the EU Model Clauses, which have been approved by the EU data protection authorities and may be used in agreements between service providers and their customers to ensure that any personal data leaving the European Economic Area will be transferred in compliance with the General Data Protection Regulation (GDPR).³³ These types of data transfer agreements provide assurances that CSPs are safeguarding personal data responsibly as well as a pre-approved means to protect and support international data flow in a secure and compliant manner.

4. Ensure that CSPs and third-party contractors demonstrate robust security controls to address unauthorized third-party access to data, systems, and assets through internationally-recognized third party accreditations (e.g. ISO 27001, ISO 27018, SOC, PCI DSS, etc.).

5. Classify data and define data handling roles and responsibilities to determine appropriate data protection obligations for each party.

Governments should consider leveraging ISO 27018 for defining the roles of the data controller and processor. Governments can work with CSPs to adequately understand and apply data protection responsibilities for the controller versus processor for each of the cloud service models.

6. Ensure customer understanding and implementation of security services for encrypting data. AWS has pioneered encryption services that provide customers with the ability to fully control encryption keys. AWS provides customers with the option to encrypt data using their own keys that can be stored outside AWS or securely within the offerings, enabling them to control their keys and access to data while meeting strict security and compliance obligations.

The EU's General Data Protection Regulation, which became enforceable in May 2018, is intended to harmonize data protection laws throughout the European Union (EU) by applying a single data protection law that is binding throughout each member state. The GDPR does not require data residency laws within the EU, but rather supports legal frameworks in the form of data transfer models and standardized contractual clauses (i.e., the EU Model Clauses) to encourage trans-regional data flows.

Article 45 of the GDPR sets forth the principle that transfers of personal data to a third country or international organization may take place if the third country, territory or one or more specified sectors within that country, or the international organization in question ensures an adequate level of protection. To achieve this, governments may:

- Change their existing data protection law and engage in adequacy discussions with other countries. For example, New Zealand is in the process of achieving an adequacy decision by the EU Commission.
- Establish bilateral frameworks such as the US-EU Privacy Shield.

³³ The AWS GDPR Data Processing Addendum, which includes the EU Model Clauses, is now part of our online Service Terms. This means all AWS customers globally can rely on the terms of the AWS GDPR DPA whenever they use AWS services to process personal data under the GDPR. More information about the AWS approach to GDPR compliance is available here: <https://aws.amazon.com/compliance/gdpr-center/>.



7. Engage in bi-lateral and multi-lateral efforts to update the MLAT process so that it balances governmental needs to expeditiously obtain evidence necessary in investigations and prosecutions with an individual's right to privacy over electronic content they own. We support legislation that updates privacy and law enforcement access to electronic communications -- both domestically and internationally. We also encourage governments to review and update their national laws to address the roles, responsibilities, and mechanisms governing lawful access to data consistent with the principles of the MLAT process.

Conclusion

While governments may perceive a sense of increased security when imposing data residency requirements for data processed and stored in local IT facilities because they offer physical proximity and control, deeper evaluation shows that restricting IT services to the local jurisdiction-only does not provide better overall data security. From a risk-benefit perspective, hyperscale CSPs, like AWS, can better help manage cybersecurity risks while still minimizing the risk of foreign government access to data. Governments also need to consider the significant trade-offs associated with data residency requirements. Not only will governments that use restrictive data residency requirements forfeit access to some of the most secure computing environments on earth, but, beyond security, they will be forced to deal with a perpetual lag in access to cost-effective, state-of-the-art technology needed for their own digital transformation. We encourage governments to re-evaluate the security objectives that they actually achieve through data localization restrictions relative to the significant economic, IT modernization, and security opportunity costs. The security capabilities of hyperscale CSPs not only address top-of-mind concerns, but provide security at a bar higher than traditional on-premises or locally contracted facilities. Policy solutions, such as data transfer agreements and leveraging well-reputed international security accreditations, can serve as sufficient means to address data residency objectives while promoting public sector digital transformation goals.