Addressing Data Residency Requirements with AWS

Comply with local legislation, regulations, contractual commitments with end users, or corporate policies, by using AWS infrastructure and services to store or process data in a specific location

Navigating the Complexities of Data Residency

Many organizations in both the public and private sectors have data residency requirements, driven by a variety of factors. This eBook explores the topic in depth. We define what data residency means and outline typical situations in which it may apply.

We then set out how you can map out your own data residency requirements, and explore your options for meeting them with AWS, using AWS Regions, AWS Local Zones or AWS Outposts.

What is ‘Data Residency’?

The term ‘data residency’ gets used interchangeably with various others, and can mean slightly different things to different organizations. Simply put, data residency is a requirement to store or process data in a particular geographical location.

We will first uncover:

- The main drivers of data residency requirements
- The myth about data residency and information security

Typical Data Residency Drivers

There are various reasons why organizations are subject to, or impose data residency requirements

The first is to meet **legal and regulatory demands**, including data locality laws. For example, digital payment service providers operating in the United Arab Emirates are required to store all user and transaction data within the country's borders.[1] A similar situation exists in India, where payment system providers must store all data related to their payment systems within India's national borders.[2,3] The French government has made it a requirement that data produced by local and national government be stored and processed in France.[4] And Australia requires organizations that store or process healthcare records, to do so within the country.[5] Publicly held companies subject to the control of the Capital Markets Board in Turkey, must keep their primary and secondary IT systems within the country.[6,7] There are many more examples like this around the world.

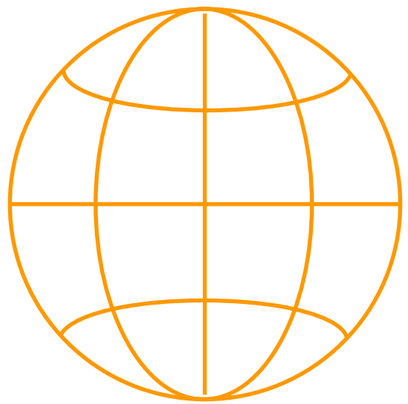

For other organizations, the main driver may be **different business and operating models**. A business might wish to carry out the majority of its activities in a particular country, so that it falls under that nation's financial rules. This can require the company to store and/or process some or all of its data within that country's borders.

An organization might have **contractual requirements** to store data in a particular country. An independent software vendor (ISV), for example, may have to agree to hold certain customers' data in a specific country, to meet those clients' own data residency requirements.

The fourth driver could be **corporate policy**. A business or public sector body might mandate that certain data must be stored or processed in a specified location. This may be partly or wholly driven by one of the factors above.

Data Residency and Information Security

As organizations attempt to tackle information security challenges, such as cybersecurity threats, it results in the implementation of data residency requirements, to try to safeguard data.

In their annual CIO Survey in 2020, Harvey Nash and KPMG found that 41% of organizations have experienced an increase in cybersecurity incidents, and that security is now the top technology investment priority.[8]

However, if security is the main driver of your data residency requirements, assess whether mandating data residency genuinely protects the data against the threats it faces.

Remote threats

Most vulnerabilities are exploited remotely over the internet. Any system connected to the internet is at risk, and requires appropriate logical security safeguards.

Manual human process failures

Human process failure often contributes to cybersecurity events. For example, the complexities of manually keeping interconnected business systems fully patched mean these can often remain on old software versions, months after vulnerabilities are discovered.

Errors or malicious behavior

A large number of data compromises are the result of unintentional errors or malicious individual actions.

Access credentials get lost or mismanaged. Phishing or social engineering attacks result in individuals handing over login information. Both enable attackers to access systems under the cloak of authorized accounts.

Equally, a malicious individual could leak information.

Customers can benefit from AWS security solutions

AWS provides exceptional security capabilities, and doing so at every level of the stack is critical to our operations and offerings. This includes the physical security of cloud data center infrastructure as well as the security of data.

AWS is architected to be the most flexible and secure cloud environment. The core infrastructure is designed to meet the most stringent business security requirements, so that customers can run their business in an automated driven environment with confidence.

AWS monitor infrastructure 24/7 to ensure the confidentiality and integrity of data, and have developed Outposts' services to satisfy the security needs of banks, the military and other high-sensitive organizations. AWS is the solution to secure workloads and applications, providing a variety of deeply integrated and tested services, that can be combined to automate tasks in innovative ways.

Find out more about the benefits of AWS security on our dedicated AWS Cloud Security page. We also highlight some of the specific security benefits of AWS Outposts on page 10 of this eBook.

Data That may Have Residency Requirements

Today, organizations around the world store enormous amounts of data, so it's helpful to know where to start looking when you're carrying out your own data residency assessment.

There are two main categories of information that typically fall under data residency requirements.

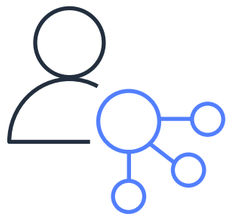

Personally identifiable data

Virtually every organization stores or processes at least some personally identifiable data, so this should be the first area you consider.

This might be **financial information**, such as an individual's transaction history, or other personally identifiable details.

Healthcare data is another example, where organizations store information about patients and their medical histories.

Public sector organizations, and their commercial partners, also hold large amounts of data about individuals, created through their use of public services. This **citizen information** could include addresses, dates of birth, financial data, details of someone's children or dependents, and their history of interactions with that organization.

National data

National data is often considered sensitive, and consequently subject to data residency requirements in certain jurisdictions.

Examples include **geospatial** information, such as maps and seismic data, which may be collected by either public sector bodies or private companies. Data associated with the **military**, including intelligence, operations and technology, is another example.

Similarly, data about **critical national infrastructure and resources** may need to be stored within the country. This could include information about the design and operation of power-generation facilities, utility and communications networks, and transportation infrastructure.

Understand Your Own Data Residency Requirements

Since every organization will hold different data, and be subject to different laws, regulations, contracts and policies, there's no one-size-fits-all checklist that will give a comprehensive understanding of your specific data residency requirements. Speak to AWS if you have any queries on this process.

To establish what these requirements could be, we recommend going through our **five-step process:**

1. Create (or update) your data catalog

Firstly, you need a complete and up-to-date data catalog, showing the data your organization is acquiring, producing, processing, storing and sharing.

3. Apply the regulations to your data

You now know what data you're storing and processing, and the applicable laws, regulations, policies or contractual requirements. Put these insights together, by updating your data catalog with data residency metadata. For each asset subject to residency requirements, set out where it needs to be stored and the reason for this need, such as the specific regulation you must comply with.

5. Stay up-to-date

Requirements will evolve over time, as laws, regulations and other factors change. Keep your residency metadata up to date by periodically reviewing it with the appropriate stakeholders.

2. Understand your data residency obligations

Work with your legal and compliance teams to map out what data residency requirements apply to your organization and its data. Depending on the nature of your operations, these requirements may only apply to certain subsets of data.

4. Communicate the requirements

Ensure everyone in your organization is aware that certain data is subject to residency requirements, and that they know where to find the rules. Also make clear who they should speak to for guidance, if they need to work with that data in any way.

Meet Your Data Residency Needs With AWS

Once you establish you have data residency requirements to meet, you have four broad technology options:

- Run your own infrastructure in your chosen location
- Use the AWS Cloud
- Use AWS Local Zones
- Bring cloud benefits to your own facility, with AWS Outposts

1. Buy and run your own technology in your chosen location

The traditional method of buying hardware and installing it on-premises, in a colocation space or in a data center, will be familiar to anyone working in IT.

From a data residency perspective, this approach provides the means to keep data in a location where the organization can own the infrastructure, the facility and even the land on which it's located.

However, widespread adoption of cloud computing has highlighted the drawbacks of working in this conventional way:

- High capital costs to buy and periodically refresh infrastructure
- Lack of flexibility
- High maintenance and management overhead
- Greater security responsibilities
- No consistency between on-premises and the cloud for development or operations teams

These challenges, coupled with the cloud-first strategies being adopted in many organizations, mean IT teams are now looking for alternative ways to meet their data residency requirements. This is where AWS can help.

2. AWS Regions: The public cloud

AWS has data centers across the Americas, EMEA and Asia-Pacific, clustered into 'Regions'. Many organizations with data residency requirements find they can meet them by designing their workload to store and process data in a selected AWS Region.

This will be the most secure, scalable and flexible way for you to run these workloads, with access to all the AWS services available in that Region. In addition, you'll benefit from the physical and operational security in place at all AWS locations.

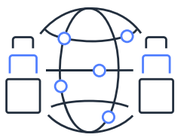

3. AWS Local Zones: Cloud capabilities where no AWS Region exists

AWS Local Zones deploys AWS compute, storage, database and other select services closer to large population, industry and IT centers where no AWS Region exists today. While its primary purpose is to run latency-sensitive portions of applications close to end users, Local Zones can also be used to meet your data residency requirements.

Discover more about AWS Local Zones.

4. AWS Outposts: Bringing cloud benefits to your own facility

There may be no AWS Region or Local Zone in the location where you need to store or process data. Or there may be other reasons you can't run a particular workload in your local AWS Region or Local Zone. In these situations, you can bring the benefits of AWS into your own data center, colocation space or on-premises facility, using AWS Outposts.

Outposts is a fully managed service that extends the AWS cloud into virtually any facility. Using the same infrastructure found in AWS Regions, run AWS compute, storage, database and other services locally, while seamlessly accessing the full range of AWS services available in your local Region.

This gives you the best of both worlds when it comes to meeting data residency requirements: full control over the location of your data, coupled with additional benefits, including:

Eliminate upfront capital costs

Outposts offers pricing models that reduce or even eliminate the upfront hardware cost, and spread it over a three-year term.

Flexible configurations and space to grow

Outposts can be purchased in a variety of off-the-shelf or custom configurations to meet your workload requirements. Increase the compute and storage capacity in your Outposts as your needs evolve, by upgrading your configuration in the future.

Discover more about AWS Outposts.

Consistent experience for developers and operations teams

With the same hardware infrastructure, services, APIs and tools on both Outposts and in the AWS Region, your developers and operations teams enjoy a truly consistent hybrid experience across on-premises and the cloud.

Reduce management overheads

Free up resources from low-level maintenance. AWS takes care of Outposts installation, monitoring, maintenance, software patches and upgrades, giving your teams more time for high-value work.

Benefit from AWS security

As part of the managed service, AWS protects your Outposts infrastructure, similar to how it secures its cloud infrastructure. This uses the AWS Nitro Chip, which is built into every Outposts host.

The Nitro System continuously monitors, protects and verifies your Outposts hardware and firmware, thereby removing this burden from your in-house teams, who can focus on securing the workloads on Outposts, and the physical security of the facility (if applicable).

Virtualization resources are offloaded onto dedicated hardware and software, which minimizes the attack surface. And the Nitro System's security model is locked down, preventing administrative access, including by Amazon employees. This eliminates the possibility of human error or malicious tampering.

Your data is encrypted at rest in the Nitro System in every host. The encryption key is wrapped to an external key, which is stored in a removable device, present in every host in the Outpost. Destroying the device is equivalent to destroying the data. When AWS removes the Outposts rack from your premises, you must destroy the removable device, thereby ensuring your data does not leave the premises.

In addition, the AWS Outposts hardware is installed in an enclosed rack with a lockable door and tamper-detection, to protect against physical attacks.

Discover more about AWS Nitro System.

AWS Services on Outposts: At-a-Glance

Implement familiar AWS services on Outposts in your on-premises facility, colocation space or data center.

Outposts offers a wide, and ever-growing, range of AWS services. Discover the current range below (as of November 2020).

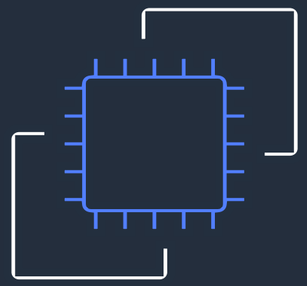

Compute

Select from a range of EC2 instances that use Intel® Xeon® Scalable processor technology, with or without local instance storage:

- General purpose
- Compute optimized
- Memory optimized
- Graphics optimized
- I/O optimized

Storage

- Amazon EBS gp2
- Amazon S3 on Outposts

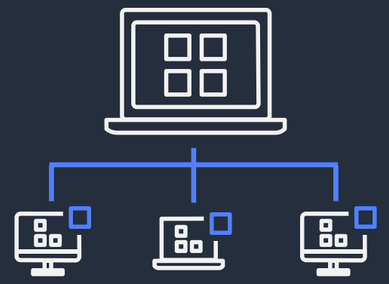

Networking

- VPC extension: Seamlessly extend your existing Amazon VPC to your Outpost
- Local gateway: Connect your Outpost resources to your on-premises networks
- Application Load Balancer (ALB): Automatically distribute incoming HTTP(S) traffic across multiple targets on your Outposts

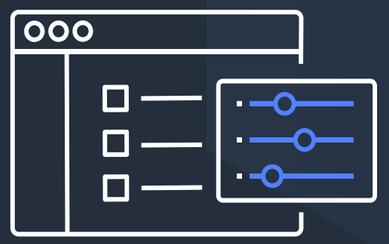

Containers

- Amazon ECS: Run a highly scalable, high-performance container orchestration service that supports Docker containers
- Amazon EKS: Run Kubernetes on AWS without the need to install and operate your own Kubernetes control plane

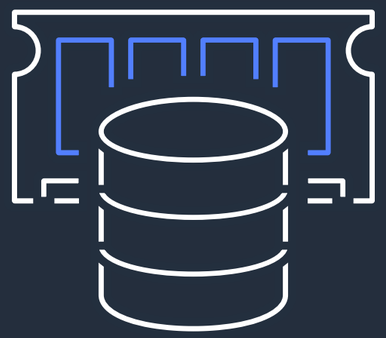

Databases

- Amazon RDS: Set up, operate and scale a relational database on Outposts
- ElastiCache: A fully managed in-memory data store, where you can power real-time applications with sub-millisecond latency. ElastiCache enables you to set up, run, and scale open-Source compatible in-memory data stores on Outposts

Data analytics

- Amazon EMR: Set up, deploy, manage and scale Apache Hadoop, Apache Hive, Apache Spark and Presto clusters on Outposts

Plus: Access all the services available in your local AWS Region

Outposts is an extension of your local AWS Region: extend your Amazon VPC on-premises and connect to services available in the Region, all from within your private VPC environment.

And as new versions of AWS services become available in the cloud, the services on your Outposts are automatically upgraded.

Case Study: Transforming Banking Services

A large bank, offering corporate, investment, and personal banking services wanted to modernize their current workflows by delivering digital banking services including e-wallet and mobile payments to their customers, while keeping customer data secure and resident in their country.

The challenge

Providing their customers with a much more modernized approach to everyday banking, produced a handful of challenges that had to be mitigated. These included: transforming the overall user experience while still complying with data residency requirements, ensuring enhanced security on customer's personal data, protection against remote cyber threats and malicious behavior and, ensuring that they are able to access the benefits of the cloud while running workloads on-premises.

The solution

By using AWS Outposts, they were able to transform their customer experience by developing new automated applications. Build and deploy modern containerized workloads while storing and processing data on-premises. Access multiple AWS regions located across the country to enable them to meet data residency and other business requirements. And finally, protect their customers identifiable data using safe-guarded AWS security infrastructure.

Business outcome

Using AWS cloud based services allowed for an accelerated digital transformation within the organization. Which, increased the developer and infrastructure teams productivity by leveraging the same infrastructure services, APIs and tools across on-premises and the cloud.

Next Steps

Many organizations in both the public and private sectors have data residency requirements, driven by legal and regulatory demands. With AWS Outposts you can bring all the advantages of the cloud to your facility, while storing and processing data in your chosen location.

1. Engage

Reach out to your account team or fill out our contact form. Alternatively, go into the AWS Management Console.

2. Choose

Select your size and then order the Outpost rack configuration that best suits. Custom configuration is available.

3. Install and Launch

AWS will install and deliver your configuration. Use standard AWS APIs or Management Console to launch and run AWS resources locally.

Learn more
https://aws.amazon.com/outposts

1 https://www.dlapiperdataprotection.com/index.html?t=law&c=AE
2 https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=43574
3 https://rbidocs.rbi.org.in/rdocs/Content/PDFs/OVER13062020_FF0CC640BA0CC434C83A1E847B4FB3FD6.PDF (Section 6.4.9)
4 http://www2.itif.org/2017-cross-border-data-flows.pdf
5 https://www.legislation.gov.au/Details/C2019C00337 (Clause 77)
6 http://www.erdem-erdem.av.tr/publications/newsletter/management-of-information-systems/
7 https://www.mondaq.com/turkey/leasing/960124/localization-of-data-storage-through-cloud-computing-systems-
8 https://home.kpmg/content/dam/kpmg/xx/pdf/2020/09/harvey-nash-kpmg-cio-survey-2020-executive-summary.pdf