



Truly Consistent Hybrid Cloud

AWS Outposts

From startups to the largest enterprises and leading government agencies, the cloud is used for almost every kind of workload. This includes anything from running enterprise applications to creating new and innovative applications. But some use cases cannot yet be fully migrated to the cloud. Perhaps they need to work with large local datasets, share data with on-premises applications with single-digit-millisecond latency, or meet data residency requirements. In these use cases, constantly passing data to and from the cloud is too slow, too resource-intensive or not permitted. To meet these needs, customers are asking for cloud services to be deployed on premises for a truly consistent hybrid experience. AWS Outposts enables AWS compute (utilizing the same Intel® Xeon® Scalable processor technology in Amazon EC2 instances), storage, database and other services to run locally.

How AWS Outposts helps any organization anywhere

"AWS services can run locally on AWS Outposts. Services available in the local AWS Region can work with Outposts too, and tools hosted in the local AWS Region can be used to manage workloads running on Outposts."



On-premises infrastructure has traditionally required customers to install and manage their own hardware. However, organizations are becoming increasingly familiar with the benefits of running workloads in the cloud, where both hardware installation and management are taken care of on their behalf. When you add in all the other advantages of cloud computing (lower cost, increased flexibility), it's clear why there's a strong desire to extend these to on premises.

What does this look like in practice? It can mean deploying cloud infrastructure and applications in locations where they never existed before, such as at an organization's own data center or colocation facility. When cloud applications are deployed in closer proximity to end-users and dependent systems, they can better meet low-latency requirements. Organizations can use cloud infrastructure and services to achieve ultra-low-latency performance, process local data in specific geographic locations, or modernize legacy applications while keeping them on premises.

There are further advantages of extending the cloud to on-premises environments. By using a common set of tools, services and APIs across the cloud and on premises, applications currently running in the cloud can be extended to anywhere where the customer needs it, without the need for costly replatforming or refactoring. Equally, by using the same tools to deploy and manage workloads in both environments, operational overhead can be reduced.

True hybrid computing does bring challenges, of course. How can your organization access infrastructure that serves your unique use cases? How can you integrate edge resources with on-premises and cloud infrastructure, so that development and operational work is not duplicated? There is a need for infrastructure services that give your organization:

- **Computing resources** that meet your use cases where they are needed most for lower latency, data residency requirements or local data processing needs
- **Consistency of infrastructure, service availability, and performance** across all deployments (wherever they are situated) so that developers can build an application once and deploy it flexibly
- **Simplified infrastructure deployment and management** that minimizes operational risk, maintenance downtime, and time and resources spent

AWS Outposts extends AWS infrastructure and services on premises

AWS Outposts is a fully managed service that extends cloud computing capabilities to customers' locations wherever they need it, including data centers, colocation spaces, and other facilities. AWS services can run locally on AWS Outposts. Services available in the local AWS Region can work with Outposts too, and tools hosted in the local AWS Region can be used to manage workloads running on Outposts.

This solution brief discusses key use cases driving the need for hybrid computing. It also explores how Outposts extends customers' familiar AWS experience to new locations, both from a technical standpoint and in a range of real-world settings with real-world benefits. Read on to learn how AWS Outposts can meet your organization's needs and how to get started with the service.

Use cases for AWS Outposts



AWS Outposts is designed to serve technical use cases with these needs...

Low-latency compute

Low-latency applications typically fall into one of two categories: interactive applications and latency-sensitive business applications. Firstly, interactive applications include media-editing tools or online multiplayer gaming, where services need to be deployed close to end users, within a latency target, such as 20 milliseconds.

Secondly, latency-sensitive business applications are where the workload needs to be connected to another device or system on the local network. Examples include health management systems (HMS) and manufacturing process control systems. These often need to respond to other connected systems with single-digit-millisecond latency. For example, a smart factory's quality-control system must identify a defective product and instruct a connected system to remove it from the line, so that faulty goods don't reach customers.

Complementing the low latency delivered by AWS Outposts instances due to their location on the enterprise edge, Intel® Xeon® Scalable processors (designed with Intel® Mesh Architecture) deliver efficient and scalable, low-latency data flow across both edge and cloud environments.

When the nearest public cloud servers are not close enough to satisfy these needs, organizations need to be able to run applications at the network edge. In these cases, there is a need to extend cloud infrastructure to new locations.

Local data processing

Many organizations want to bring cloud services to their data centers, factories, and other premises. They envision a hybrid architecture, with important control applications running on-premises, and other resources, such as data lakes and machine learning workloads, in the cloud region.

Some of these on-premises applications need to process large datasets that are created or stored at an edge location. Large datasets cannot be efficiently moved to and from the cloud for processing, due to bandwidth limitations or low-latency requirements. For example, an autonomous vehicle testing facility generates large amounts of telematics data that can be mined for insights and used to train artificial intelligence models. The manufacturer may want to perform initial processing of this data on-premises, before sending it to a data lake for deeper, long-term analysis. The manufacturer needs edge infrastructure that can store and process large local datasets and integrate seamlessly with the cloud region – without the need for a complex hybrid-cloud architecture.

Data residency

Organizations may need to meet data residency requirements for several reasons. Some are driven by regulations mandating that sensitive data, such as personal or financial information, is stored and processed inside a specific municipality, state, country or region. Equally, some may have contractual obligations specifying where applications are deployed.

Some organizations may have existing public cloud applications they want to shift to another location, while others are new to cloud. They need to control where their workloads run and where their compute and storage resources reside, and to use local operational tooling, perhaps for monitoring and stability. They also need low-friction movement of workloads between public cloud and on premises, and vice versa, so they can adapt to changing requirements.

Modernization

Some organizations have a long-term goal of modernizing legacy on-premises applications, such as those running on ageing mainframes, and moving them to the cloud. This can involve splitting large applications into smaller, cloud-suitable ones.

Whatever form the modernization takes, it can present numerous challenges. Applications can be complex, part of a delicate ecosystem, or there may be interdependencies between components that require them to be on the same network.

Having access to consistent infrastructure, services, tools, and performance across on-premises IT and the cloud, eases the modernization journey for all involved. It enables organizations to rebuild applications once, keep them on the local network initially, and easily deploy them in the cloud and at other edge locations when ready. As Intel Architecture and its supporting software ecosystem are the most pervasive and rigorously tested compute environment, you can have greater confidence that your modernized applications with Outposts will behave consistently and be performant as you migrate between on-premises, the edge and cloud.

Real-world applications for AWS Outposts

The technical use cases for AWS Outposts underpin opportunities for organizations in a wide range of industries to innovate new products, services, experiences, and efficiencies.



Healthcare & life sciences

Healthcare providers can use AWS Outposts to easily apply AWS analytics and machine learning services to health management systems that have low-latency processing requirements. Data can be stored locally on Outposts to enable rapid retrieval of medical information. In life science fields, lower latencies can deliver faster results exactly where scientists are conducting their experiments.



Telecommunications

Service providers can use cloud services and tools to orchestrate, update, scale, and manage the lifecycle of Virtual Network Functions (VNFs) across cloud, on-premises, and edge. Virtual network functionality and new network services can be deployed at any network operations center.



Media and entertainment

Artists and producers can access the latest GPU innovations on-premises for graphics processing, audio and video rendering, and for running other media applications. Live and real-time event streaming applications or interactive experiences requiring low latency, can be run at edge locations, close to end users.



Multiplayer gaming & iGaming

Real-time multiplayer gaming is a huge market where Outposts can offer ultra-low-latency performance (seamless interactions between players) and improved graphics processing at locations close to players. In the world of iGaming (casino-style betting games and online sports bets), Outposts' single-digit-millisecond performance can mean the difference between winning and losing when it comes to a user's experience.



Financial services

Financial institutions can build next-generation trading and exchange platforms that serve all participants at low latency, using AWS services and programming interfaces. Banking, payments-processing, and risk-management services can be delivered to customers and financial clients from in-country locations, to meet data-residency requirements.



Retail

With AWS Outposts, retailers can leverage AWS database, container, and analytics services to enable retail innovations such as connected store experiences, run point-of-sale systems to process in-person transactions locally, and deliver consistent and reliable operations at each retail location.



Manufacturing

With AWS Outposts, manufacturers can use AWS services to run control systems (e.g. MES and SCADA) close to factory-floor equipment and centralize their operation. On-premises applications can be seamlessly integrated with services running in the AWS Region, for centralized operations.

AWS Outposts meets your workload demands wherever they are needed

AWS Outposts meets the above use cases by extending AWS infrastructure, services, APIs, and tools to virtually any location. With AWS Outposts, the cloud is brought as close to the endpoint as needed, to meet low-latency performance needs, store and process data locally or in a specific geographic location, or modernize applications in place. AWS Outposts uses the same Intel-processor-powered hardware and delivers the same experience as in an AWS Region, but in a location of the customer's choosing. AWS also manages, monitors, and updates Outposts infrastructure and services, just like in the cloud.

AWS Outposts delivers:



A truly consistent hybrid experience

AWS Outposts offers a consistent experience with the same hardware infrastructure, services, APIs, management, and operations on-premises as in the cloud. Unlike other hybrid solutions that require use of different APIs, manual software updates, and purchase of third-party hardware and support, Outposts enables developers and IT operations teams to achieve the same pace of innovation across different environments.



Fully managed infrastructure

The customer chooses where the hardware goes. AWS delivers, installs, monitors, patches, and updates it. Organizations can reduce the time, resources, operational risk, and maintenance downtime required for managing their IT. Infrastructure and services running locally on Outposts are automatically upgraded, and with in-built real-time health monitoring customers get peace of mind and can set alerts on capacity issues before workloads run out of compute, storage, network capacity, or GPU.



Data storage and processing on-premises

AWS Outposts enables customers to securely store and process data that needs to remain on-premises, or where there is no AWS Region today. This can help address storage requirements of companies in highly regulated industries, or with specific data-residency requirements. Local data processing needs can include processing large datasets that can't be easily moved or transcoding, filtering and caching of data that must be done on premises. With Outposts, customers have access to Amazon Block Store (EBS) for high performance block storage and object storage with Amazon Simple Storage Service (S3).



AWS Services on-premises

Customers can run Amazon Cloud services such as Amazon EC2, Amazon S3 and EBS, container-based services such as Amazon Elastic Container Service (ECS), Amazon Elastic Kubernetes Service (EKS), Amazon Load Balancer (ALB) and database services such as Amazon Relational Database Service (RDS), and Amazon EMR for big data workloads, locally on AWS Outposts. Organizations can also access the full range of AWS services available in the Region to build, manage, and scale on-premises applications.

Outposts is a way to provide AWS infrastructure and services all over the world, where organizations need them, in ways that serve organizations' technical use cases. Where ultra-low-latency performance is required, applications can be deployed on AWS Outposts on the local network and with AWS services (such as EC2, EBS, EKS, and RDS) running locally. For local data-processing, AWS Outposts provides local compute, storage and AWS services, as well as consistent hybrid-cloud infrastructure and tools, making it easy to move data to the cloud when ready. Organizations with data-residency needs can move applications and data between AWS Regions and Outposts with minimal friction. And AWS Outposts make application-modernization projects faster, providing a helpful staging area that organizations can use to incrementally convert applications in-place. Modernized applications can subsequently be relaunched in the AWS Region via a simple lift-and-shift migration.

How AWS Outposts works

An Outpost is a pool of AWS compute and storage capacity deployed at a customer site and fully managed by AWS. With an Outpost installed customers can create subnets on an Outpost and specify them when creating AWS resources such as EC2 instances, EBS volumes, ECS clusters, and RDS instances. Instances in Outpost subnets communicate with other instances in the AWS Region using private IP addresses, all within the same VPC.



Defining AWS Outposts

- **Outpost site** – The customer-managed physical buildings where AWS will install your Outpost. A site must meet the facility, networking, and power requirements for your Outpost.
- **Outpost configurations** – Combination of Amazon EC2 compute capacity, Amazon EBS storage capacity, and network switches. Each configuration has unique power, cooling, and weight support requirements.
- **Outpost capacity** – Compute and storage resources available on the Outpost. You can view and manage the capacity for your Outpost from the AWS Outposts console.
- **Outpost equipment** – Physical hardware that provides access to the AWS Outposts service, including racks, servers, switches, and cabling owned and managed by AWS.
- **Service link** – Network route that enables communication between your Outpost and its associated AWS Region. Each Outpost is an extension of an Availability Zone and its associated Region.
- **Local gateway** – A logical interconnect virtual router that enables communication between your Outpost and your on-premises network.

AWS Outposts specifications

AWS Outposts are delivered as an industry-standard 42U rack that is 80 inches (203.2cm) tall, 24 inches (60.96cm) wide, and 48 inches (121.92cm) deep - custom designed by Intel and AWS. Outposts can be installed in virtually any data center, colocation space, or on-premises facility that can maintain a consistent connection to an AWS Region. Inside are hosts, switches, a network patch panel, a power shelf, and blank panels. It has redundant active components, including network switches and hot spare hosts.

Outposts support a range of EC2 instance types (powered by Intel® Xeon® Scalable processors), that are categorized to your workload need: general purpose, compute optimized, memory optimized, graphics optimized or I/O optimized.



GET STARTED

AWS Outposts enables organizations from a range of industries to meet key edge computing challenges by creating new hybrid infrastructure deployments that bring the cloud, and their applications, closer to endpoints where they are needed. Contact us through your account team or our online form: <https://aws.amazon.com/contact-us>

Alternatively, select the right configuration for you via the AWS Management Console:

1. Log into the AWS Management Console to create a site
2. Select an Outposts configuration
3. Place your order

Learn more at aws.amazon.com/outposts
