

Hybrid Cloud with AWS

November 2020



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. 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.

© 2020 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Introduction 1
- Hybrid Cloud Use Cases 2
 - Application Migration to the Cloud 2
 - Cloud Services On-premises 2
 - Data Center Extension 3
 - Edge Computing 4
 - ISV and Software Compatibility 5
- Hybrid Architecture Tenets 5
- Operations and Management Framework for Hybrid Cloud with AWS 6
 - Hybrid Cloud Infrastructure 7
 - Core Services 7
 - Unified Hybrid Cloud Management 9
 - Example: Dropbox’s Hybrid Cloud Architecture 13
- AWS Hybrid Cloud Solutions 14
 - AWS Outposts 14
 - VMware Cloud on AWS 14
- Conclusion 15
- Contributors 15
- Document Revisions 15

Abstract

Businesses and organizations often have a mix of cloud, on-premises, and edge computing infrastructure. A hybrid cloud with Amazon Web Services (AWS) delivers IT resources like compute, storage, databases, and more, through an integration of AWS Cloud services with on-premises and edge infrastructure. Hybrid cloud architectures help organizations integrate their on-premises and cloud operations to support a broad spectrum of hybrid cloud use cases, often using a common set of services, tools, and interfaces. This whitepaper outlines major use cases, defines hybrid architectural tenets, describes an implementation framework, and provides AWS solutions to guide you in creating a hybrid cloud strategy, and to design and implement a hybrid cloud environment with AWS.

Introduction

Many businesses and organizations now adopt cloud computing as a key aspect of their technology strategy. They are moving their workloads to the AWS Cloud for greater agility, cost savings, performance, availability, resiliency, and scalability. While most applications can be easily migrated, some applications need to be re-architected or modernized before they can be moved to the cloud. There are a subset of applications that must remain on-premises due to low-latency, local data processing, high data transfer costs, or data residency requirements. This leads many organizations to seek [hybrid cloud](#) architectures to integrate their on-premises and cloud operations to support a broad spectrum of use cases.

Considerations for building a [hybrid cloud with AWS](#) include:

- **Creating a hybrid cloud strategy:** A hybrid cloud strategy provides guidelines that govern consumption of cloud and on-premises resources to support your business objectives. This whitepaper describes common use cases for building a hybrid cloud, such as ongoing migration to the cloud, ensuring business continuity during disasters, extending cloud infrastructure on-premises to support low-latency applications, or expanding international footprint on AWS. These use cases help you identify and define your business objectives for building a hybrid cloud, and provide guidelines for workload placements on the hybrid cloud.
- **Creating a technical strategy:** A technical strategy for the hybrid cloud identifies the guiding tenets of the hybrid cloud architecture, and defines an implementation framework. This whitepaper outlines common tenets of a hybrid cloud architecture to help you define guiding principles for a planned hybrid cloud implementation. An example tenet is having a consistent set of interfaces for resource provisioning and management across the hybrid cloud infrastructure.

The operations and management framework described in this whitepaper helps architects and systems integrators to identify the building blocks, best practices, and AWS services needed to create a technical strategy to implement a hybrid cloud with AWS.

Dropbox has built a hybrid cloud with AWS with a strategy to leverage the scale, agility, innovation, and global footprint provided by AWS. They have built a consistent set of tools and interfaces for provisioning and managing the hybrid infrastructure to ensure developer productivity. Dropbox's hybrid cloud with AWS powers the infrastructure that serves more than 500 million worldwide customers. Similarly, a large US-based insurance company has experienced higher business agility from adopting a cloud-first strategy. They consume Software-as-a-Service (SaaS) solutions first, then leverage AWS Managed Services and infrastructure to deliver cloud-native solutions, and only consume on-premises solutions when they have hybrid cloud use cases such as low-latency processing of data.

When you build a hybrid cloud strategy that best meets your business needs, consider tradeoffs such as the possibility of added operational overhead or reduced agility with on-premises infrastructure.

Hybrid Cloud Use Cases

These use cases help you identify and define your business objectives for building a hybrid cloud, such as ongoing migration to the cloud, ensuring business continuity during disasters, extending cloud infrastructure on-premises to support low-latency applications, or expanding your international footprint on AWS.

Application Migration to the Cloud

Large migrations from on-premises datacenters to AWS may involve thousands of applications and can take several years. Customers require a consistent operational environment across their hybrid cloud while they migrate their applications, to ensure business continuity. [Johnson & Johnson](#) and [Hess Corporation](#) created a hybrid cloud environment to support their migration to AWS.

You may want to leverage your on-premises investments in VMware while taking advantage of the agility and scalability offered by the AWS Cloud. AWS has partnered with VMware to enable you to migrate and run your VMware vSphere workloads on AWS, and leverage native AWS services for your on-premises environments through [VMware Cloud on AWS](#). [Stagecoach Group](#) has adopted VMware Cloud on AWS to accelerate their migration to AWS.

Cloud Services On-premises

Some applications have data residency, high data transfer costs, local data processing, or low-latency requirements. These applications must be deployed on-premises or close to the end users/systems. Customers want to seamlessly integrate these applications with their cloud deployments in a hybrid cloud environment to ensure operational consistency.

Similarly, customers who primarily operate on AWS may need to deploy applications on-premises for local data processing or low-latency needs. These customers want to continue leveraging existing cloud skill sets and tools that they have invested in for these on-premises deployments.

To support these use cases, [AWS Outposts](#) provides a consistent hybrid cloud solution that brings the same AWS infrastructure, services, APIs, management tools, support, and operating model that customers are familiar with, in [AWS Regions](#), to virtually any data center, co-location space, or on-premises facility. [VMware Cloud on AWS](#) provides an integrated cloud offering jointly developed by AWS and VMware. Additionally, using [Amazon Relational Database Service](#) (Amazon RDS) on [VMware](#), you can deploy

managed databases in on-premises VMware environments using the same Amazon RDS technology they use in the cloud.

Data Center Extension

With [data center extension](#), the AWS Cloud is an extension of your on-premises infrastructure to support applications that need to run on-premises. There are four broad use-cases:

Cloud Bursting

Cloud bursting is an application deployment model in which the application primarily runs in an on-premises infrastructure, and when the demand for capacity increases, AWS resources are utilized. Customers like [FuseFX](#), [Pacific Life Insurance](#), and Dropbox burst compute and storage resources to AWS from on-premises. There are two main reasons to use cloud bursting:

- **Bursting for compute resources:** You consume burst compute capacity on AWS through [Amazon Elastic Compute Cloud](#) (Amazon EC2) and [managed container services](#) of the [Amazon Elastic Container Service \(Amazon ECS\)](#), [Amazon Elastic Kubernetes Service](#) (Amazon EKS), and [AWS Fargate](#).
- **Bursting for storage:** In addition to integrating applications with [Amazon Simple Storage Service](#) (Amazon S3) APIs, [AWS Storage Gateway](#) enables on-premises workloads to use AWS Cloud storage. Capabilities such as [File Gateway](#), [Tape Gateway](#) and [Volume Gateway](#) help enable cloud bursting capabilities for block and file storage.

Backup and Disaster Recovery

Customers like [Scripps Network Interactive](#) implement a hybrid infrastructure with AWS for their application disaster recovery needs for applications that reside on-premises. AWS services like [Amazon S3 APIs](#), AWS Storage Gateway, [AWS Backup](#), [AWS DataSync](#) and [AWS Transfer for SFTP](#) enable you to implement a disaster recovery strategy with AWS for your data hosted on-premises.

Distributed data processing

Customers often deploy applications across on-premises data centers and AWS, with functionality split between the infrastructures. Most commonly, low-latency or local data processing components reside on-premises and other functionality, including asynchronous processing, archiving, compliance, business analytics processing, or machine learning-based predictions reside on AWS. AWS services like [AWS Storage Gateway](#), [AWS Backup](#), [AWS DataSync](#), [AWS Transfer Family](#), [Amazon Kinesis Data Firehose](#) and [Amazon Managed Streaming for Apache Kafka](#) (Amazon MSK) enable you to import data into AWS for data processing needs. When the data is imported, you

can leverage AWS services like [AWS Analytics](#), [AWS Machine Learning](#), [AWS Serverless](#), [AWS Containers](#), and more to process the data. Distributed data processing using AWS Services enables you to leverage AWS innovations in these areas, while meeting the requirements of low-latency or local data processing for the application.

Geographic expansion

You may need to deploy applications closer to your end users for compliance, data sovereignty, low-latency, or local data processing needs. Deploying physical infrastructure in new geographical areas can become prohibitively expensive, or constrained by legal requirements and local laws. Customers like Dropbox leverage [AWS global infrastructure](#)¹ as an extension to their existing infrastructure to deploy their applications and make them available globally.

You can also deploy workloads in AWS Outposts in countries where AWS does not have an AWS Region yet. See the [AWS Outposts](#) section of this whitepaper.

Edge Computing

You may have [edge computing](#) needs at facilities like factories, mines, ships and windmills. AWS provides edge computing with [AWS Snowball](#), [AWS IoT Greengrass](#), and [AWS Wavelength](#).

With AWS Snowball edge computing, customers operating in disconnected, harsh, or air-gapped environments can pre-process information before transferring data to the cloud for durable retention and more advanced analysis. They can perform sophisticated analytics, machine learning, and run fully disconnected applications for traditional IT workloads on Amazon EC2 compute resources.

With AWS IoT Greengrass, you can enable devices and equipment to respond to local events in near real-time by acting locally on the data they generate. [AWS Lambda](#) functions deployed on AWS IoT Greengrass Core devices can use local device resources like cameras, serial ports, or GPUs, so device applications can quickly access and process local data. Devices can stay operational and function seamlessly, even with intermittent connectivity to the cloud.

You can reduce the cost of running edge applications by using AWS IoT Greengrass to filter locally before transmitting to the cloud. AWS Wavelength is an AWS Infrastructure offering optimized for mobile edge computing applications. Wavelength Zones are AWS infrastructure deployments that embed AWS compute and storage services within communications service provider (CSP) data centers at the edge of the 5G network, so application traffic from 5G devices can reach application servers running in Wavelength Zones without leaving the telecommunications network.

ISV and Software Compatibility

If you want run the same independent software vendor (ISV) software that you run on-premises in a hybrid or distributed model, you can use the [AWS Marketplace](#), a curated digital catalog, to find, buy, deploy, and manage third-party software on AWS.

AWS has [built the most complete and proven approach](#) for rapidly migrating tens to thousands of applications to the AWS Cloud to help you leverage your existing on-premises ISV software investments.

Recently, AWS launched the AWS Outpost [Service Ready](#) program, which offers products that integrate with AWS Outposts deployments. You can discover [products on this page](#) that are tested on AWS Outposts, and follow AWS security and architecture best practices. AWS Competency Partners are ready to help AWS customers migrate and deploy their applications to AWS Outposts. [AWS Partners](#) validated through the AWS Service Ready Program offer products tested to integrate with AWS Outposts deployments.

Hybrid Architecture Tenets

Depending on the identified use cases for a hybrid cloud implementation, enterprises must identify one or more tenets for their hybrid cloud architecture. These are some common tenets for a hybrid cloud architecture:

- **Operational consistency:** Hybrid cloud customers need operational consistency across the hybrid cloud in the form of a consistent set of interfaces and APIs for resource provisioning, monitoring and managing resources on the hybrid cloud.
- **Simple to control, manage, and secure:** Hybrid cloud customers want to manage hybrid cloud resources in a simple, consistent, and secure manner, similar to what they get with AWS APIs today.
- **Build once, deploy anywhere:** Hybrid cloud customers want to develop once and deploy workloads to cloud, on-premises, and edge environments in an agile and consistent fashion using a common set of development and management APIs, while getting consistent performance across the environment. In essence, the hybrid infrastructure must support the workload requirements, independent of where they are deployed.
- **Enterprise-class application Service Level Agreements (SLAs):** Hybrid cloud customers need their infrastructure to be highly reliable and available, similar to what they get with [AWS today](#).
- **Existing skill sets and tools:** While deploying a hybrid cloud, customers often want to leverage organizational skill sets and tools that they have already invested in.

Operations and Management Framework for Hybrid Cloud with AWS

The operations and management framework detailed here identifies the building blocks for architecting and implementing a hybrid cloud environment with AWS. This framework helps you identify the components and the corresponding considerations for building a hybrid cloud with AWS. This section also identifies AWS services and solutions to address the needs for each building block.




Figure 1- Operations and management framework for hybrid cloud with AWS

[AWS Outposts](#) vertically integrates across the layers of this framework by providing a hybrid cloud solution that brings [AWS infrastructure](#), [security](#), [services](#), [APIs](#), [management tools](#), [support](#) and [operating model](#) to data centers, co-location spaces, or on-premises facilities. AWS Outposts eliminates the undifferentiated heavy-lifting associated with building software and systems to integrate infrastructure in a hybrid cloud environment. It provides security, performance, and operational consistencies across the hybrid environment, while addressing the needs of running applications seamlessly on-premises or the AWS Cloud.

Hybrid Cloud Infrastructure

Physical infrastructure deployed on-premises and in AWS Regions provide the infrastructural foundation for a hybrid cloud. The network interconnecting these infrastructures enables traffic exchange within the hybrid environment.

On-premises and AWS Infrastructure

The customer infrastructure includes compute servers, storage nodes, networking devices, and edge computing devices. This infrastructure is hosted in customer-owned or leased facilities, manufacturing/retail facilities, or in spaces near end-users.

[AWS global infrastructure](#) consists of more than 24 geographical regions and 77 Availability Zones as of September 2020. AWS infrastructure provides a global edge network (currently 216 [points of presence](#)) to AWS customers for accelerating content delivery, domain name services, global load balancing, and security.

For on-premises infrastructure, you can deploy AWS hardware through AWS Outposts. AWS also provides edge computing infrastructure with [AWS Local Zones](#), AWS Wavelength, AWS Snowball Edge and AWS IoT Greengrass.

Network

The network interconnecting on-premises infrastructure with AWS can be through dedicated physical connections, VPN, or over the internet.

With [AWS Direct Connect](#), you can establish a private virtual interface from your on-premises network directly to your Amazon VPC. This provides an elastic, simple, and consistent network experience that can also increase bandwidth throughput. With [AWS site-to-site virtual private network \(VPN\)](#), you can create an IPsec VPN connection between your Amazon VPC and your on-premises network over the internet. Additionally, some applications, especially those leveraging IoT technologies, use the public internet to exchange traffic with AWS resources such as [AWS service endpoints](#) and public EC2 instances.

Core Services

We have identified three core services for a hybrid cloud implementation:

Device and Fleet Management Service

Device and [fleet management](#) service provides two main functions in a hybrid cloud:

- Host management of on-premises physical devices such as compute, networking, and storage devices. This includes device, configuration, software, metrics, and inventory management of the physical infrastructure.

- Device and fleet management service also provides the functionality and management interfaces to provision, manage and monitor infrastructure for the host devices. This includes management interfaces (such as create, delete, update, and read) for physical or virtual compute, storage and networking resources.

For the AWS physical infrastructure, all fleet management functions are managed by AWS on behalf of customers, including host management. [AWS APIs](#) provide capability for management and monitoring of AWS resources. [VMware vSphere](#) and [OpenStack](#) are examples of software that provide fleet management functions of host management, and interfaces for managing virtual infrastructure for on-premises compute environments.

For host management of on-premises infrastructure, you can manage servers in on-premises data center with [AWS Systems Manager](#). Systems Manager provides several [features](#) like remote command execution, patch management, inventory management, state management, and automation to help with host management functions. [AWS OpsWorks](#) provides a configuration management system using [Chef](#) and [Puppet](#) to automate how servers are configured, deployed, and managed in on-premises environments.

Metrics and Logging

Unified monitoring capability across the hybrid cloud simplifies operations and provides consistent health monitoring, alerting, logging, and auditing capabilities. A few major components for this service include:

- **Metrics and alerting:** Continuous monitoring of infrastructure, service, and application metrics provides the basis for secure, performant, reliable and cost-optimized operational practices. As a best practice, capturing of metrics from all sources in the hybrid environment must be at a unified repository.

[Amazon CloudWatch](#) provides a central repository for metrics collection, monitoring, alerting, and dashboarding. CloudWatch agents are deployed on EC2 instances, on-premises servers, and virtual machines, which export metrics on CPU, processes, memory, storage, and networking. [CloudWatch custom metrics](#) allow collection, storage, and monitoring of metrics from applications and infrastructure.

- **Auditing, logging and traceability:** Continuous collection, monitoring, and retaining logs related to management/control, application, and data-plane activities provides detective controls and auditing capabilities to identify security threats, to troubleshoot incidents, and for event correlation. As a best practice, all logs must be stored in a central repository for troubleshooting and further analytics processing.

[AWS CloudTrail](#) is a service that enables governance, compliance, operational auditing, and risk auditing of AWS accounts. With CloudTrail, customers can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. [Amazon CloudWatch Logs](#) enables you to centralize logs from all systems, applications, and AWS services. Use CloudWatch Logs to monitor, store, and access log files from EC2 instances, CloudTrail, Route 53, and custom sources.

Identity, Security and Access Management

Establishing a unified identity and access management solution is key to providing secure and consistent access to services in a hybrid cloud environment. As a best practice, a single Identity Provider (IdP), which manages identity information for principals while providing authentication services to resources on the hybrid cloud, must be instituted.

[AWS Directory Services](#) provide multiple ways to set up and run directories like [Amazon Cloud Directory](#), [Amazon Cognito](#), and [Microsoft AD](#) to serve as the IdP for the hybrid cloud.

[AWS Identity and Access Management](#) (AWS IAM) and [Amazon Cognito Identity Pools](#) enable identity federation through integration with IdPs supporting Security Assertion Markup Language (SAML) or Open-ID Connect (OIDC) to obtain temporary, limited-privilege AWS credentials to manage and access resources on AWS and in a hybrid cloud deployment with AWS Outposts.

Finally, [AWS Single Sign-On \(SSO\)](#) enables you to integrate services in the [Unified hybrid cloud management layer](#) to the same IdPs as used to manage access to AWS resources and services.

Unified Hybrid Cloud Management

The hybrid cloud management layer provides a unified set of interfaces to consume hybrid cloud services like compute, storage, networking, databases, analytics, and others. These interfaces provide capabilities for provisioning, editing, deleting, monitoring, and operating resources and services on the hybrid cloud. This section describes design practices, components, and AWS services that address the needs of building a unified hybrid cloud management layer in support of hybrid cloud services.

[AWS Outposts](#) natively provides unified hybrid cloud management through the use of the same APIs and management tools across on-premises and AWS infrastructure. AWS Outposts supports [several AWS services](#), including compute, storage, networking, and higher-level services allowing consistent operations across the hybrid cloud, eliminating the need to build and manage custom software.

Compute Services

Compute services in the hybrid cloud provide the interfaces to manage compute (instances, containers, functions) resources. Figure 2 provides an example customer software implementation of a unified management interface for compute services.

In this example, a hybrid cloud user authenticates with the Identity, security, and access management service of the hybrid cloud to gain authorization to the management interfaces of the compute service. The compute service provides a unified provisioning, monitoring, and operating interface for the user. Internally, the compute service interacts with the core fleet or device management layer for on-premises infrastructure management, [AWS EC2 APIs](#) for EC2 management, as well as the core services of metrics and logging services for metrics and logging needs and identity, security and access management for gaining access authorization to on-premises and AWS resources through their respective APIs.

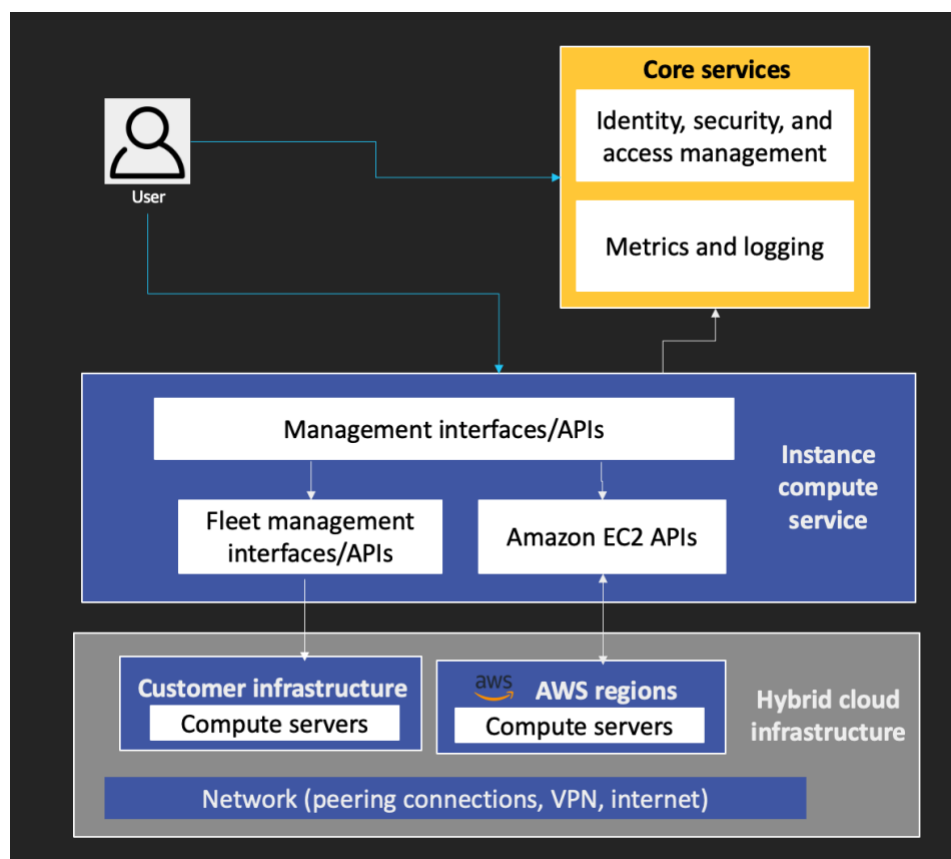


Figure 2 – Example of compute service on a hybrid cloud

A consistent mechanism for managing guest operating systems on AWS instances and in on-premises virtual machines across the hybrid cloud provides seamless operations for activities such as software/patch management and policy enforcement. You can manage on-premises servers and AWS EC2 instances with [AWS Systems Manager](#). Systems Manager provides several [features](#), such as remote command execution,

patch management, inventory management, state management, and automation, to help with host management functions.

Storage Services

Outside of providing core storage services for block, file, and objects, hybrid cloud use-cases often require moving data between on-premises data centers and AWS. These use cases are cloud bursting for storage, disaster recovery (data replication and backups), distributed data processing (for analytics processing on AWS), or geographic expansion (move data closer to customers). Data movement is required for files, block storage, transactional data in databases, and streaming data.

- **Files:** The [File Gateway](#) interface, [AWS DataSync](#), [AWS Transfer for SFTP](#), [Amazon EFS](#), [Amazon FSx for Lustre](#), and [Amazon FSx for Windows File Server](#) are used for integrating files between the environments and enabling cloud bursting, disaster recovery, and application migration use-cases.
- **Block storage:** [Volume Gateway](#) provides an on-premises iSCSI interface to provide S3-based storage in AWS (in gateway-cache mode). [AWS Storage Gateway](#) can be used for cloud bursting, storage extension, migration, or backups of block stores.
- **Transactional data:** [AWS Database Migration Service](#) (AWS DMS) provides integration between on-premises SQL/NoSQL databases and AWS-based databases (on EC2 or AWS RDS, DynamoDB, DocumentDB) by providing migration and synchronization between the databases. Additionally, DMS can be used to migrate data from on-premises databases to S3 directly for analytics workflow integration.
- **Streaming data:** Streaming records from on-premises data centers and AWS sources can be collected and analyzed in managed stream stores on AWS, including [Amazon Kinesis](#) and [Amazon MSK](#).

The AWS Outposts service also provides on-premises storage with EBS. [S3 on AWS Outposts](#) enables customers to store object data on premises using the S3 API.

Networking and Security Services

Networking and security services enable you to create and manage networks for applications and secure them on the hybrid cloud. A few major components are discussed here:

- **Virtual Networking:** Virtual networking enables you to provision logically isolated sections of the infrastructure where they can launch resources. You can define your own IP addressing, subnets, routing policies, securities, and gateways in the virtual network based on the application requirements. On the hybrid cloud, these virtual networks extend between AWS and on-premises infrastructures, allowing applications to function across the environment.

[Amazon VPC](#) enables you to create virtual networks in AWS Regions. [AWS Direct Connect](#) private virtual interfaces (VIFs), transit VIFs, and [site-to-site VPN](#) provide mechanisms to extend the virtual network between Amazon VPC and on-premises networks.

- **Load balancing:** In a hybrid environment, load balancers are used to distribute traffic to targets across on-premises and public cloud environments. Load balancers abstract the location of the physical resources in the hybrid cloud by presenting a unified front end for an application service. [AWS Application Load Balancer](#) and [Network Load Balancer](#), deployed in AWS regions, support targets in AWS regions as well as on-premises. They also support containers as targets deployed across the hybrid environment. Application load balancing on AWS Outposts is fully managed, operates in a single subnet, and scales automatically up to the capacity available on the Outposts rack to meet varying levels of application load without manual intervention.
- **Unified DNS:** As a best practice, internal DNS resolutions for applications and services deployed in virtual networks on the hybrid cloud must be unified across the infrastructure. [AWS Route53 resolver and conditional forwarding rules](#)² provide a mechanism to unify DNS resolutions across on-premises DNS servers and resolvers hosted on AWS.

For public DNS resolutions, internet traffic is routed to the front-ends of web applications deployed on the hybrid cloud through public DNS services like Amazon Route53. In the hybrid cloud, the application front ends (implemented using load balancers or API endpoints on instances) reside either on-premises, in AWS Regions, or split across the infrastructure. AWS Route53 features routing mechanisms for active-backup and active-active hybrid environments.

- **Infrastructure Security:** Infrastructure security in a hybrid cloud must be applied to all layers of the technology stack across both on-premises and AWS environments. These layers include security at the edge network, perimeter, load balancers, network devices, host and guest operating systems, applications, virtual networks, subnets, and compute. Customers require a common set of security policies that they can apply to AWS or on-premises infrastructure. AWS provides tools like [AWS Web Application Firewall](#) (WAF), [AWS Shield](#), [VPC Security Groups](#), and [VPC Network ACLs](#) to enforce security boundaries.

Example: Dropbox's Hybrid Cloud Architecture

Dropbox built a hybrid cloud with AWS, enabling the use cases of distributed data processing, cloud bursting, disaster recovery, and geographical expansion.

Figure 3 provides a view of the technical architecture for Dropbox's hybrid software stack built with AWS.

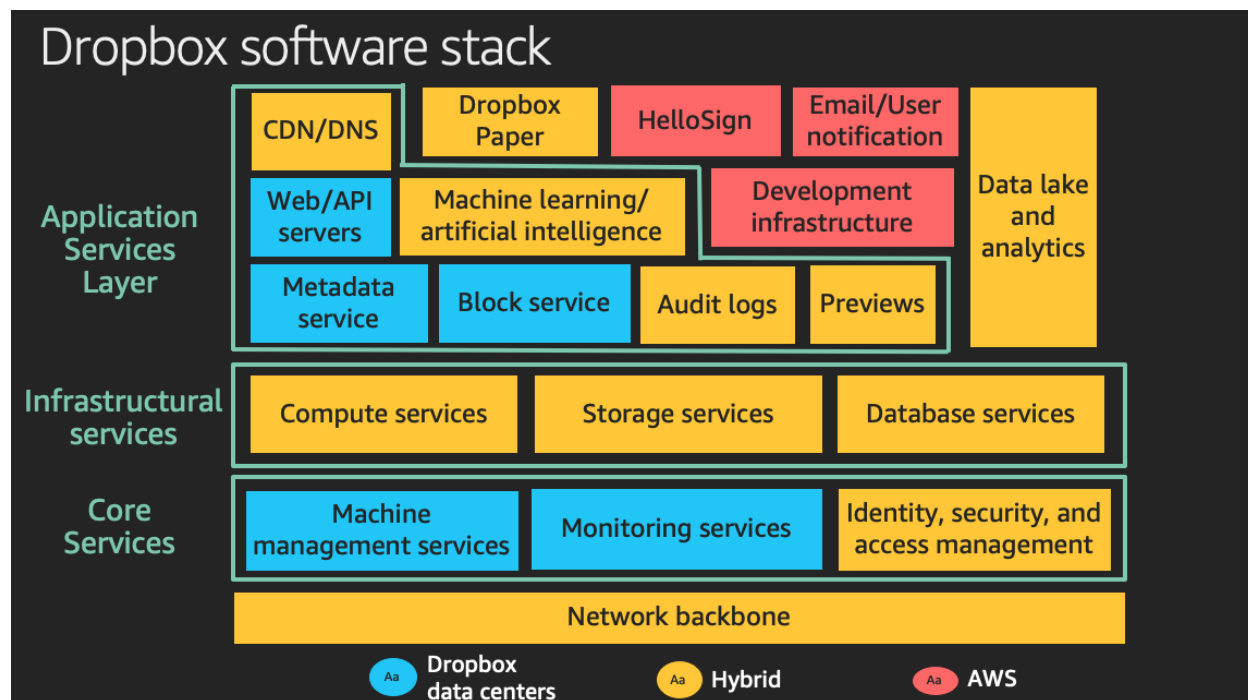


Figure 3 – Dropbox's hybrid software stack

Dropbox's machine management service provides host management capabilities for on-premises infrastructure, as well as a common set of tools for systems management for infrastructure in both AWS and on-premises, such as device inventory, software management, installation, patching, remote command execution, service discovery, and more. Additionally, Dropbox's monitoring services provide unified metrics and logging service for the hybrid cloud.

Further up the hybrid stack, Dropbox's compute, storage, and database hybrid services leverage the core services, and provide the management interfaces for provisioning, managing, and operating infrastructure on the hybrid cloud.

Using these infrastructure and core services, Dropbox's developers are provided with a unified set of interfaces and tools to provision, operate, and manage infrastructure and services on the hybrid cloud with AWS. Dropbox has implemented distributed data processing in the hybrid cloud to leverage AWS innovation for several workloads such as data lake and analytics, machine learning infrastructure, document previews, audit logging, and more. Dropbox leverages the scale and agility provided by AWS through

deployments of [Dropbox Paper](#) and [Dropbox HelloSign](#) products on AWS. Additionally, Dropbox has deployed their application services in AWS Regions in Frankfurt, Tokyo, and Sydney, using their hybrid architecture for geographic expansion. Details of Dropbox's hybrid cloud strategy and implementation are described in [Dropbox's re:Invent presentation](#).

AWS Hybrid Cloud Solutions

In addition to providing services and solutions to implement the building blocks identified in the [Operations and Management Framework for Hybrid Cloud with AWS](#) section of this whitepaper, AWS also provides AWS Outposts and [VMware Cloud on AWS](#) to provide a fully integrated hybrid cloud solution.

AWS Outposts

[AWS Outposts](#) brings native AWS services, infrastructure, and operating models to virtually any data center, co-location space, or on-premises facility. You can use the same APIs, the same tools, the same hardware, and the same functionality across on-premises infrastructure and the cloud to deliver a consistent hybrid experience. Outposts can support workloads that must remain on-premises due to low latency or local data processing needs.

AWS Outposts come in two variants:

- [VMware Cloud on AWS Outposts](#) enables you to use the same VMware control plane and APIs you use to run your infrastructure,
- The AWS native variant of AWS Outposts enables you to use the same APIs and control plane you use to run in the AWS cloud, but on-premises.

AWS Outposts infrastructure is fully managed, maintained, and supported by AWS to deliver access to the latest AWS services.

VMware Cloud on AWS

[VMware Cloud on AWS](#) is an integrated cloud offering jointly developed by AWS and VMware, delivering a highly scalable, secure, and innovative service that enables organizations to seamlessly migrate and extend their on-premises [VMware vSphere](#)-based environments to the AWS Cloud running on EC2 [bare-metal infrastructure](#). VMware Cloud on AWS is ideal for enterprise IT infrastructure and operations organizations looking to migrate their on-premises vSphere-based workloads to the public cloud, consolidate and extend their data center capacities, and optimize, simplify and modernize their disaster recovery solutions.

With VMware Cloud on AWS, you can simplify your Hybrid IT operations by using the same VMware Cloud Foundation technologies, including vSphere, [vSAN](#), [NSX](#), and [vCenter Server](#) across your on-premises data centers and on the AWS Cloud without having to purchase any new or custom hardware, rewrite applications, or modify your operating models. The service automatically provisions infrastructure and provides full VM compatibility and workload portability between your on-premises environments and the AWS Cloud.

Conclusion

AWS offers the broadest and deepest set of services to address a broad spectrum of use cases to help you build a hybrid cloud with AWS. AWS provides:

- A highly integrated solution for VMware workloads in the cloud and on-premises
- The ability to bring the same infrastructure and services on-premises for applications with data residency, low latency or local data processing requirements with AWS Outposts
- The most services to help you implement the building blocks identified in the Operations and Management framework for hybrid cloud with AWS
- The largest and most reliable global footprint to help you meet reliability and global coverage requirements

Contributors

Contributors to this document include:

- Anuj Dewangan, Principal Solutions Architect, Amazon Web Services
- Tom Laszewski, PE Transformation Strategist, Amazon Web Services
- Rob Chen, Head of Infrastructure Solution Marketing, Amazon Web Services
- Andrei Savine, WW Enterprise Transformation Architect, Amazon Web Services

Document Revisions

Date	Description
November 2020	First publication

Notes

¹ <https://aws.amazon.com/about-aws/global-infrastructure/>

² <https://aws.amazon.com/blogs/security/simplify-dns-management-in-a-multiaccount-environment-with-route-53-resolver/>