AWS Well-Architected Framework

End User Computing (EUC) Lens

End User Computing (EUC) Lens: AWS Well-Architected Framework

Table of Contents

End User Computing (EUC) Lens

Publication date: **September 18, 2025** (*Document revisions*)

This document describes the End User Computing (EUC) Lens for the AWS Well-Architected Framework, a collection of customer-proven best practices for designing and operating cloud-based virtual desktop and application streaming solutions. This document is intended for those in technology roles, such as chief technology officers (CTOs), architects, developers, and desktop administrators. After reading this document, you should understand typical customer scenarios, general design principles, and the best practices for deploying secure, operationally excellent, performant, cost-optimized, reliable, and sustainable EUC workloads.

Introduction

The Well-Architected End User Computing Lens can be used to guide the initial design of your EUC solution. This document can also be used to periodically evaluate your existing solution against best practices, to create a backlog of tasks to improve your workload, and to measure improvement over time.

Get started by familiarizing yourself with the scenarios that represent common EUC use cases and the design principles that offer high-level best practices. Refer to the pillars of the Well-Architected Framework to dive deep into pillar-specific design principles and best practices. The definitions section contains definitions of terms related to EUC workloads.

This document only covers details from the Well-Architected Framework that are specific to End User Compute. We recommend that you also consider the best practices from the AWS Well-Architected Framework when designing your architecture.

Custom lens availability

Custom lenses extend the best practice guidance provided by AWS Well-Architected Tool. AWS WA Tool allows you to create your own custom lenses, or to use lenses created by others that have been shared with you.

To determine if a custom lens is available for the lens described in this whitepaper, reach out to your Technical Account Manager (TAM), Solutions Architect (SA), or Support.

Definitions

Following is a list of definitions related to the AWS Well-Architected Framework and EUC workloads.

AWS Definitions

- EUC
  - **Amazon AppStream 2.0:** Secure, reliable, and scalable application streaming and low-cost virtual desktop service
  - **Amazon WorkSpaces Family:** Comprehensive, fully persistent, Virtual Desktop Infrastructure for most worker types
  - **Amazon WorkSpaces Core:** Virtual desktop infrastructure APIs for third-party VDI software
  - **Amazon WorkSpaces Secure Browser:** Secure, low-cost browser service for access to internal websites and Software as a Service apps
  - **Amazon WorkDocs:** Secure document sharing and content collaboration—connecting teams everywhere
  - **Amazon DCV:** Amazon DCV is a high-performance remote display protocol that provides secure remote desktop delivery and application streaming, avoiding the need for expensive dedicated workstations.
- Hardware
  - **Amazon WorkSpaces Thin Client:** Reduce costs, simplify logistics, and accelerate deployment using virtual desktops
- Storage
  - **Amazon FSx:** Launch, run, and scale feature-rich and highly performant file systems with just a few clicks
  - **Amazon S3:** Object storage built to retrieve any amount of data from anywhere
  - **Amazon EFS:** Share file data without provisioning storage
- Compute
  - **Amazon EC2:** Secure and resizable compute capacity for virtually any workload
- Cost

  - **Amazon WorkSpaces Family Pricing:** Pricing across the Amazon WorkSpaces Family services is designed to be flexible and cost-effective, allowing you to pay for the resources you need without over provisioning.
  - **Bring Your Own Windows Desktop Licenses (BYOL):** If your licensing agreement with Microsoft allows it, you can bring and deploy your Windows 10 or 11 desktop on your WorkSpaces.
  - **Cost Optimizer for Amazon AppStream 2.0:** Monitors your AppStream 2.0 app block builders and image
  - **Cost Optimizer for Amazon WorkSpaces:** Monitor Amazon WorkSpaces usage and optimize costs builders and notifies you and/or stops them when they are active for longer than specified thresholds.
- Managed directories for WorkSpaces
  - **AD Connector:** A directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in the cloud.
  - **AWS Managed Microsoft AD:** AWS Directory Service lets you run Microsoft Active Directory (AD) as a managed service.
  - **Simple AD:** Provides a subset of the features offered by AWS Managed Microsoft AD, including the ability to manage user accounts and group memberships, create and apply group policies, securely connect to Amazon EC2 instances, and provide Kerberos-based single sign-on (SSO).
  - **Cross Trust:** You can establish a trust relationship between your AWS Managed Microsoft AD directory and your on-premises domain.
- Protocols for Amazon WorkSpaces
  - **Amazon WSP (WorkSpaces Streaming Protocol):** Built using Amazon DCV technology, enabling high-performance remote access to Amazon WorkSpaces instances for a wide range of workloads and use cases.
  - **PCoIP (PC over IP):** Amazon WorkSpaces supports PCoIP when needed based on the type of devices your users will be accessing their WorkSpaces from, which operating system is on your WorkSpaces, what network conditions your users will be facing, and whether your users require bidirectional video support.
- Networking
  - **Internet gateway:** Horizontally scaled, redundant, and highly available VPC component that allows communication between your VPC and the internet. It supports IPv4 and IPv6 traffic.

- **NAT gateway:** Allow resources in private subnets to connect to the internet, other VPCs, or on-premises networks. These instances can communicate with services outside the VPC, but they cannot receive unsolicited connection requests.
- **Public subnets:** Subnet which has a direct route to an internet gateway. Resources in a public subnet can access the public internet.
- **Private subnets:** Subnet which does not have a direct route to an internet gateway. Resources in a private subnet require a NAT device to access the public internet.
- **Amazon Virtual Private Cloud (VPC):** Launch AWS resources in a logically isolated virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of AWS.
- **AWS Regions:** Each Region is designed to be isolated from the other Regions. This achieves the greatest possible fault tolerance and stability.
- **Availability Zone:** Each Region has multiple, isolated locations known as Availability Zones.
- **Amazon Route 53:** A reliable and cost-effective way to route end users to your Internet applications. As such, Amazon Route 53 is a highly available and scalable Domain Name System (DNS) web service that connects user requests to internet applications running on AWS or on-premises.
- **DHCP Option Sets in Amazon VPC:** Network devices in your VPC use Dynamic Host Configuration Protocol (DHCP). You can use DHCP option sets to control: The DNS servers, domain names, or Network Time Protocol (NTP) servers used by the devices in your VPC and whether DNS resolution is enabled in your VPC.

- Security
  - **AWS Identity and Access Management (IAM):** Helps an administrator securely control access to AWS resources.
  - **Security groups:** Controls the traffic that is allowed to reach and leave the resources that it is associated with.
- Monitoring
  - **Amazon CloudWatch:** Provides a reliable, scalable, and flexible monitoring solution that you can start using within minutes. You no longer need to set up, manage, and scale your own monitoring systems and infrastructure.
  - **Amazon EventBridge:** Serverless event bus to build event-driven applications at scale.

- **VPC Flow Logs:** Enables you to capture information about the IP traffic going to and from network interfaces in your VPC.

- Management
  - **AWS Management Console:** Everything you need to access and manage the AWS Cloud in one web interface
  - **AWS Command Line Interface (CLI):** A unified tool to manage your AWS services. You can control multiple AWS services from the command line and automate them through scripts.
  - **Amazon WorkSpaces API:** Provides detailed information about the actions, data types, parameters, and errors of the WorkSpaces service.
  - **Tag Editor:** Tags are key and value pairs that act as metadata for organizing your AWS resources.
  - **AWS Organizations:** An account management service that enables you to consolidate multiple AWS accounts into an organization that you create and centrally manage. AWS Organizations includes account management and consolidated billing capabilities that enable you to better meet the budgetary, security, and compliance needs of your business.
  - **End User Compute (EUC) Toolkit:** Offers a range of features to help manage EUC workloads at scale.
  - **Service control policies (SCPs):** A type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for the accounts in your organization. SCPs help you to verify that your accounts stay within your organization's access control guidelines.
- WorkSpaces
  - **Running mode:**
    - **AlwaysOn:** Use when paying a fixed monthly fee for unlimited usage of your WorkSpaces. This mode is best for users who use their WorkSpace full time as their primary desktop.
    - **AutoStop:** Use when paying for your WorkSpaces by the hour. With this mode, your WorkSpaces stop after a specified period of disconnection, and the state of apps and data is saved.
  - **WorkSpace bundles and images:**
    - **WorkSpace bundle:** A WorkSpace bundle is a combination of an operating system, and storage, compute, and software resources. When you launch a WorkSpace, you select the bundle that meets your needs. The default bundles available for WorkSpaces are called public bundles.

    - **Custom image:** If you have launched a Windows or Linux WorkSpace and have customized it, you can create a custom image from that WorkSpace. A custom image contains only the OS, software, and settings for the WorkSpace.
    - **Custom bundle:** After you create a custom image, you can build a custom bundle that combines the custom WorkSpace image and the underlying compute and storage configuration that you select. You can then specify this custom bundle when you launch new WorkSpaces to make sure that the new WorkSpaces have the same consistent configuration (hardware and software).
- AppStream 2.0
    - **Fleet types:**
        - **OnDemand:** Streaming instances run only when users are streaming applications and desktops.
        - **Always-On:** Streaming instances run constantly, even when no users are streaming applications and desktops.
        - **Elastic:** The pool of streaming instances is managed by AppStream 2.0. When your users select their application or desktop to launch, they will start streaming after the app block has been downloaded and mounted to a streaming instance.
    - **Images:** You can create Amazon AppStream 2.0 images that contain applications you can stream to your users and default system and application settings to enable your users to get started with those applications quickly.
    - **Image Builders:** Amazon AppStream 2.0 uses EC2 instances to stream applications. You launch instances from base images, called image builders, which AppStream 2.0 provides. To create your own custom image, you connect to an image builder instance, install and configure your applications for streaming, and then create your image by creating a snapshot of the image builder instance.

Partner software

- **WorkSpot:** A software partner that provides cloud-native virtual desktop infrastructure (VDI) turnkey solutions.
- **LeoStream:** A software partner that provides remote desktop access solutions supporting hosted desktop deployments.
- **VMWare:** A virtualization software provider.

- **ControlUp:** Provides support to IT teams when monitoring and troubleshooting virtual desktop systems. Offers real-time monitoring, troubleshooting, automation, and data analytics.
- **Nuvens:** A member of AWS' Partner Network (APN) that supports AWS' virtual desktop services, namely Amazon WorkSpaces Manager and AppStream 2.0. Our services support AWS' customers to provision, secure, and extract intelligence from end-point devices, end-user apps, and data on AWS.
- **LiquidWare:** Provides a bundle of solutions including ProfileUnity, FlexApp and Stratusphere UX that can be used to begin as on-premises VDI desktops and can provide a migration path to cloud-hosted or desktops as a service (DaaS), with a secure, high-quality work-from-anywhere desktop experience.
- **Lakeside Software:** Offers a suite of virtual solutions such as ProActiveIT, DEX, HelpDesk, Digital Workplace, and Systrack.

Industry definitions

- **Security Assertion Markup Language (SAML) 2.0:** A standard for exchanging authentication and authorization identities between security domains.
- **Pooled:** Creates a set (or pool) of virtual desktops. Users are connected to one of the machines and it is users' machine for the duration they are connected to it. Once the user disconnects, the machine becomes available to the pool again and a different user will be allocated to it.
- **Non-pooled (dedicated):** Provides each user with a persistent dedicated virtual machine. This approach offers individual isolation and customization options.
- **Federal Risk and Authorization Management Program (FedRAMP):** A US government-wide program that delivers a standard approach to the security assessment, authorization, and continuous monitoring for cloud products and services.

Design principles

Well-Architected design principles are a set of considerations used as the basis for a well-architected workload. We recommend that you follow these design principles for a successful EUC implementation.

Evaluate the scope of your EUC use cases

Begin your design process by making an inventory of the various EUC use cases in your organization. Most organizations will have multiple user personas that have unique requirements. For example, the different use cases within the same organization may require varying:

- Sets of applications
- Peripherals
- Levels of data persistence
- Dependencies on external systems or networks
- Support teams
- Cost concerns
- Availability requirements
- Security risk profiles

Enumerate as much of this data as possible and use this data to inform the EUC design process.

Based on your inventory of user personas and their requirements, select the most appropriate EUC service for each use case. Learn the fundamental aspects of the core AWS EUC services. For optimal efficiency in implementing diverse use cases within your organization, you may need to use multiple EUC services.

Engage your AWS account team and the AWS EUC specialist team for additional guidance during any stage of your EUC journey. For more information, see Operational excellence.

Isolate EUC resources and minimize permissions

EUC services typically have different admin teams and security risk profiles from other AWS workloads. This means deploying EUC services in isolation by segregating them at the account

boundary level. Consider any data sovereignty or regulatory compliance needs for your workloads (such as HIPAA, PCI, SOC, or FedRAMP) and use AWS guidance to build compliance-aligned workloads. Remove or block access to unneeded software. If required, control data egress using controls built into EUC services, such as client copy and paste, printer redirection, and upload and download functionality. Consider forwarding OS or application logs and using agents to validate security posture. Develop a vulnerability and patch management strategy to keep your instances and images secure and up to date with the latest security updates. For more information, see Security.

Design EUC solutions that maximize performance

Maximize client performance by deploying your EUC use cases near the user base. Similarly, deploy EUC use case dependencies (like directory services and file shares) near your EUC deployment to maximize application performance. Consider combining similar or overlapping use cases to reduce deployment and maintenance tasks. Consider separating use cases based on the different needs from your EUC use case inventory. For example, if use cases have different support teams or cost reporting needs, you may want to place them in different subnets, VPCs, or AWS accounts. When separating use cases, you still may be able to gain efficiencies by reusing images. Also, consider abstracting the applications from the images or creating a library of reusable scripts to deploy applications automatically. For more detail, see Performance efficiency.

Minimize EUC resources to optimize costs

Minimize resources needed to deliver your use cases, including instance and bundle types and fleet sizes. Review usage periodically to identify idle or underused resources (such as unused or over-provisioned instances, oversized fleets, and inefficient scaling policies). Deploy automated tools, such as the Cost Optimizer for Amazon WorkSpaces and the Cost Optimizer for Amazon AppStream 2.0,to help with this process. Use open-source OSes when use cases allow or bring your own OS licenses when available. For more detail, see Cost optimization.

Scenarios

AWS EUC services can be used to satisfy many use cases. Each of the services (Amazon WorkSpaces, Amazon AppStream 2.0, and WorkSpaces Secure Browser) are outlined in the following sections along with use cases that align with each service.

Scenarios

- Delivering desktops as a service
- Delivering native Windows and Linux applications
- Amazon delivering browser-based applications

Delivering desktops as a service

Delivering desktops as a service (DaaS) is a common approach to simplifying desktop delivery, reducing administrative overheads, and enabling flexible work models. DaaS offers flexibility when choosing an OS and underlying hardware specifications, enabling a wide variety of use cases with a pay-as-you-go pricing model and the freedom to alter the hardware specifications, if required. DaaS can also help improve your security posture by storing sensitive data on the virtual desktop instead of user devices or by simplifying access to network resources on-premises, in the cloud, or on the public Internet. In a corporate environment, DaaS can enable free seating by providing access to your personal virtual desktop from a desk using an inexpensive access device, like a thin client. Contrary to an application delivery scenario, a DaaS user can be given permission to install their own software and tooling on their personal virtual desktop.

With DaaS, applications can be pre-installed in your golden image or pushed to the virtual desktop using a desktop management suite. Users may also have access to a software center to install pre-packed applications in a self-service fashion. Application layering and application containerization can alternatively be used to run applications without the need for installation (for example, from a simple network share). DaaS can be combined with application virtualization to decouple applications with special hardware requirements from a virtual desktop with more common hardware specifications.

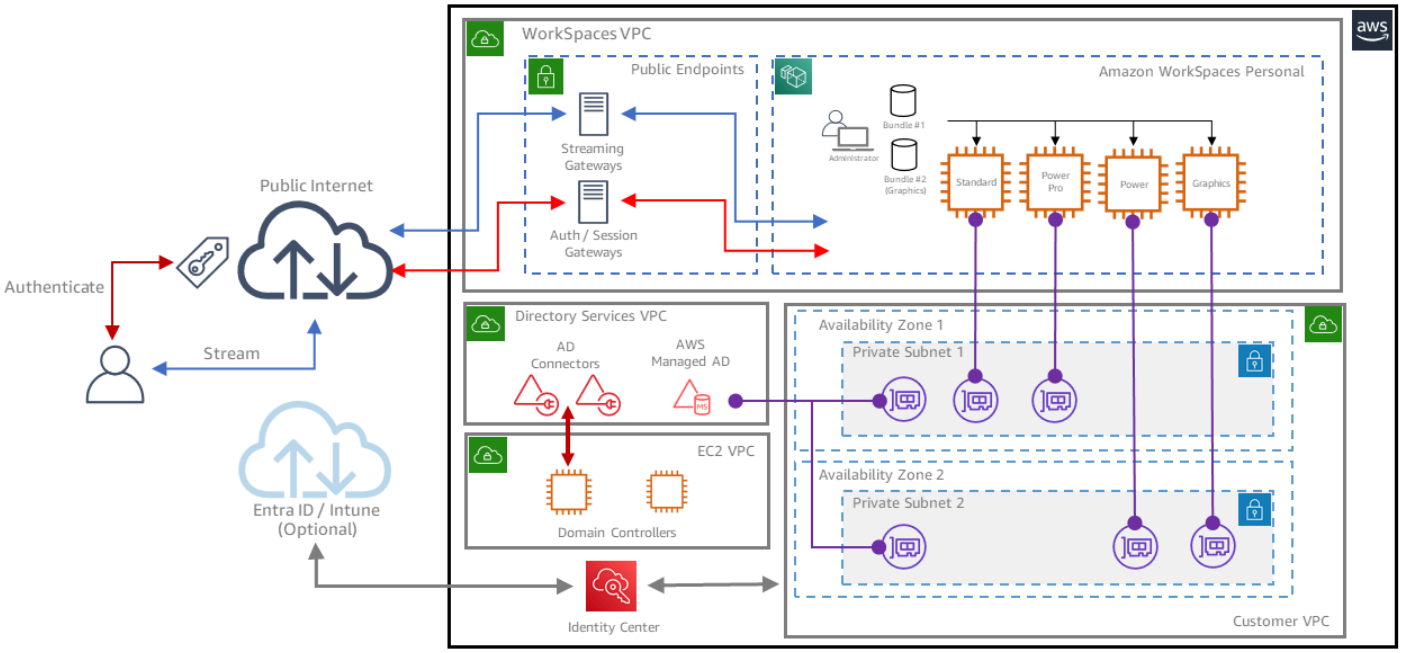

Common Amazon WorkSpaces Service deployment scenarios

The DaaS delivery model addresses many of the problems associated with providing desktops and applications to corporate users, customers and business partners. The following scenarios outline some key areas where a service-oriented approach to desktop delivery provides many advantages.

Scenario 1: Reducing costs and realizing the CapEx to OpEx shift

User scenario: We need to move away from upgrading, maintaining and periodically replacing the hardware required to our own EUC infrastructure. The costs of maintaining an on-premises data center, resilient infrastructure, and disaster recovery capabilities are growing, and we need to keep deploying more hardware to stay ahead of current technology trends.

With pay-as-you-go pricing, you can use services like Amazon WorkSpaces to move from a CapEx oriented model to a more OpEx oriented model, compared to an on-premises VDI environment. Planning demand can be difficult, and the ability to provision and decommission desktops when you choose means you can avoid sourcing new hardware and software licenses. It also helps reduce idle capacity that you might be forced to build out only for specific projects or seasonal business. You can save cost with VDI backend infrastructure, but you may also benefit from reduced cost or

an extended lifecycle of your users' devices, since this model may require less powerful hardware at the user's desk when all it does is access a desktop as a service.

Scenario 2: Protecting and using the value of existing investments

User scenario: I have an existing on-premises EUC infrastructure which is still licensed and needs to remain in place for some time. How can I reduce my desktop delivery overheads while embracing new cloud-based desktop delivery services?

If you have an existing desktop virtualization infrastructure on-premises and are looking to extend this into the cloud for peak usage or migrate part or all of this into the cloud, Amazon WorkSpaces Core can help you with this effort. Using APIs to the Amazon WorkSpaces services, Amazon WorkSpaces Core allows third-party desktop virtualization software such as VMware Horizon, Citrix DaaS, Workspot, and Leostream to be integrated with Amazon WorkSpaces.

Scenario 3: Reducing complexity

User scenario: The day-to-day overheads of managing, administering and supporting my on-premises EUC estate are considerable. How can we avoid some of the more challenging aspects of maintaining a complex desktop delivery infrastructure?

Standing up and maintaining an on-premises desktop virtualization environment can be challenging and requires knowledge and skilled staff in several different areas. A fully managed service like Amazon WorkSpaces takes a lot of this complexity away, leaving you with a significantly reduced number of tasks your IT staff has to deal with. The freed-up resources can be assigned to other projects where they can contribute more directly to your business goals.

			AWS End User Computing services	
Solution stack	On-premises VDI	DIY cloud VDI	Amazon WorkSpaces Core	Amazon WorkSpaces, AppStream 2.0, WorkSpaces Web
Image management	○	○	○	○
Directory services and policies	○	○	○	○
VDI control plane install and admin	○	○	Customer or partner	✓
Host admin	○	○	✓	✓
Storage admin	○	○	✓	✓
Load balancers install and admin	○	✓	✓	✓
Hypervisor install and admin	○	✓	✓	✓
Physical security	○	✓	✓	✓
Power / HVAC	○	✓	✓	✓
Rack and stack	○	✓	✓	✓

○ Customer or partner managed
✓ AWS managed

Scenario 4: Increased availability

User scenario: Adding and maintaining resilience for our desktop delivery infrastructure requires additional hardware and investment in expert staffing to maintain. How can we improve service delivery and reliability without significant additional capital investment?

The complexity of traditional self-managed desktop virtualization environments can impact the reliability and availability of the service you are providing to your users. Service interruptions and downtime may result in reduced user and customer satisfaction, productivity losses, and eventually financial losses. The Amazon WorkSpaces service comes with a service-level agreement (SLA) that commits to a monthly uptime percentage of at least 99.9%. In the event Amazon WorkSpaces does not meet the uptime commitment, you can receive a service credit.

Scenario 5: Cost optimization for business continuity

User scenario: The overheads of deploying and managing a secondary site for disaster recovery is prohibitive. How can we provide the level of continuity our business requires without adding cost and complexity?

Maintaining a standby deployment for your desktop virtualization environment can be a costly and labor-intensive effort. Amazon WorkSpaces comes with built-in features such as multi-Region replication that keep your users online and productive during disruptive events, all while

optimizing the cost of building a redundant, cloud-based virtual desktop environment across multiple Regions and reducing the administrative burden of maintaining such an environment.

Scenario 6: Improved security

User scenario: Our data and intellectual property are difficult to control, and it is hard to meet our industry compliance requirements. However, our users need the flexibility to access their data while on the move, using laptops, desktops, and mobile devices.

Using a service like Amazon WorkSpaces can give you better control over where your users process and store their data. For example, you may decide to allow access to certain backend services (like file servers or databases) from your Amazon WorkSpaces only. Or you may block users from copying data to their local devices. You might also consider placing an Amazon WorkSpaces Thin Client at the user's desk, which does not even have any local storage to store data. With Amazon WorkSpaces, both the execution of applications and the storage of data your users need can be centralized, regardless of employee location or the devices used for access.

Scenario 7: Enhanced flexibility to support the modern workforce

User scenario: It's costly and time-consuming to manage the replacement of old hardware, accommodating new devices, or providing compute capacity and the tools required to allow seasonal workers to staff key events.

With a service like Amazon WorkSpaces, you can change the compute type (bundle) and storage size as required throughout its lifecycle with minimal effort, helping you avoid the physical hardware upgrades and administrative and management overheads associated with a legacy laptop and PC estate.

When your business has to deal with seasonal peaks, the ability to provision and decommission new desktops with a service like Amazon WorkSpaces avoids the need to purchase new hardware and licenses that would otherwise sit unused in the cabinet during off-peak periods.

Scenario 8: Supporting mergers, acquisitions, and divestitures

User scenario: Our business is thriving and regularly acquires new businesses or divests legacy business units as our portfolio evolves. Onboarding employees from new companies and enabling them to quickly become productive by providing them with our applications and services takes significant time. Can we simplify this process?

If you are acquiring a new business or divesting a part of your business, you can often face the challenge of diverse and incoherent IT environments while having to provide access to certain

applications and data across the merged businesses or retaining access to certain applications and data for employees of a divested business. A service like Amazon WorkSpaces can be used temporarily or permanently to provide granular access to the specific required applications and data in an isolated environment.

Scenario 9: Simplification of the desktop refresh cycle

User scenario: We are planning a desktop refresh to roll out a new OS, and we need to replace a significant proportion of our laptops and desktops to meet the device requirements of this upgrade. Is there a better approach that avoids the need for hardware refreshes or upgrades when the devices are out of support or when operating system and applications demand more resources?

During a desktop refresh with a new OS, for example when moving from Microsoft Windows 10 to Microsoft Windows 11, you can use services like Amazon WorkSpaces to temporarily set up additional WorkSpaces for evaluation and testing of the new environment until you're ready to roll out the new version.

Scenario 10: Supporting a bring-your-own-device (BYOD) strategy

User scenario: How can we offer employees, contractors, or business partners the option of using their own devices to access our applications and data? This would simplify our access model and reduce the complexity of issuing, managing and replacing company owned hardware.

If you would like to offer your users the option to bring their own device, the Amazon WorkSpaces services can support this model. The WorkSpaces compute instance is hosted in the cloud, your data can remain in the cloud, and users can use a native client or browser to access these resources remotely, using a secure, high performance remoting protocol.

Scenario 11: Supporting compute-intensive workloads

User scenario: We have a number of contractors who occasionally need access to high performance, graphics-enabled machines to complete development work on our behalf. How can we avoid purchasing costly, high-end hardware to enable this sporadic use case?

Some of your workloads may have above-normal hardware requirements. Running CPU, RAM, or GPU-intensive workloads on the end-point can be extremely expensive and inefficient due to the cost of the device itself, as the device may sit under-utilized most of the time. Amazon WorkSpaces can provision these resources on demand, if and when they are required.

Scenario 12: Supporting cloud migrations and data center exit

User scenario: We have migrated several of our on-premises services to the cloud and are finding that some of our desktop applications are performing more slowly than they were before the migration. We believe that the latency between the desktops and the new cloud environment is causing these problems.

When migrating data and backend applications to the cloud, think about proximity requirements. For example, consider where the client applications are located that process this data or interact with the backend applications, both pre- and post-migration. The new location of data and backend applications may introduce increased network latency or reduced network throughput, which in turn may require bringing your client applications closer to the data and backend. When migrating to the AWS Cloud, services like Amazon WorkSpaces can bring your client applications closer to the data and backend applications now residing in the AWS Cloud.

Delivering native Windows and Linux applications

Application virtualization for Windows and Linux

Applications are at the core of employee productivity, providing the tools to create and manipulate the data that are critical to every business, enabling essential interaction with customers and business partners, and providing a solution that facilitates seamless communication and collaboration across an organization.

Maintaining an application estate, including upgrading, patching, and testing to maintain security and deliver value, can be a time-consuming and costly process to manage. The traditional tools used to perform this essential application maintenance are increasingly challenged to meet the agile needs of the modern enterprise.

For example, legacy application delivery techniques that push applications to every endpoint can be cumbersome and difficult to manage, increasing the time taken to realize the value of your investment in key business applications. At even moderate scale, to deliver an increasing number of unique user personas, the number of application combinations increases, and the test matrices required to verify application compatibility, and the reliability of your application set can become complex. If the application lifecycle management process subsequently becomes unpredictable, an increase in support overheads can occur, and users can become frustrated and less productive.

Simplifying and accelerating the lifecycle management of what can be a complex application estate promotes business agility and can be a key differentiator, verifying that employees have access to the most effective tools to maintain their productivity.

Key advantages of application virtualization

Application virtualization decouples the delivery and execution of each application from the endpoint device, whether this is a PC or laptop, a thin client device, a mobile device, or a virtual desktop. The endpoint device can receive applications which were developed for the same or a totally different OS system, a Windows application for example, can be delivered discretely to a macOS or Linux machine or a mobile OS system such as iOS or Android. As the applications are executed in, and delivered from a centralized location, such as the AWS Cloud, the data they create or consume can also be secured centrally, with policy-based controls which help prevent data leakage onto the endpoint device or into the wider internet landscape.

Virtual application delivery is achieved by installing the applications required by each user persona onto one or more centralized Windows or Linux images, hosted in your data center or Cloud of choice. These applications are then assigned and delivered individually, or in groups, to specific users using a remoting protocol which sends a pixel-based video stream over an encrypted data channel between the centralized server and the endpoint device.

Application images are created, updated and version controlled centrally. Each image can contain multiple applications or be configured to deliver a unique application with specific performance requirements such as high CPU, memory, or complex graphical applications requiring a GPU. Each discrete application image can be individually updated and re-versioned and the updates made available to users of that image immediately, they just need to logout and back in to receive the updated applications. Should a problem with the new image occur, the old image version can be immediately redeployed and accessed with a similar logoff/logon process.

Application virtualization has become one of the de facto methods of reducing the complexity of delivering and maintaining a complex application estate and can significantly accelerate the deployment of a variety of simple or complex application delivery scenarios. The flexibility of being able to access virtual applications from a location of your choice with a reliable network connection opens up a number of new delivery scenarios which allow employees to be more productive wherever they are located.

Amazon AppStream 2.0 delivers a service which accommodates the application delivery models and the associated advantages which are mentioned in the preceding section. The following diagram illustrates a typical deployment of the AppStream 2.0 service.

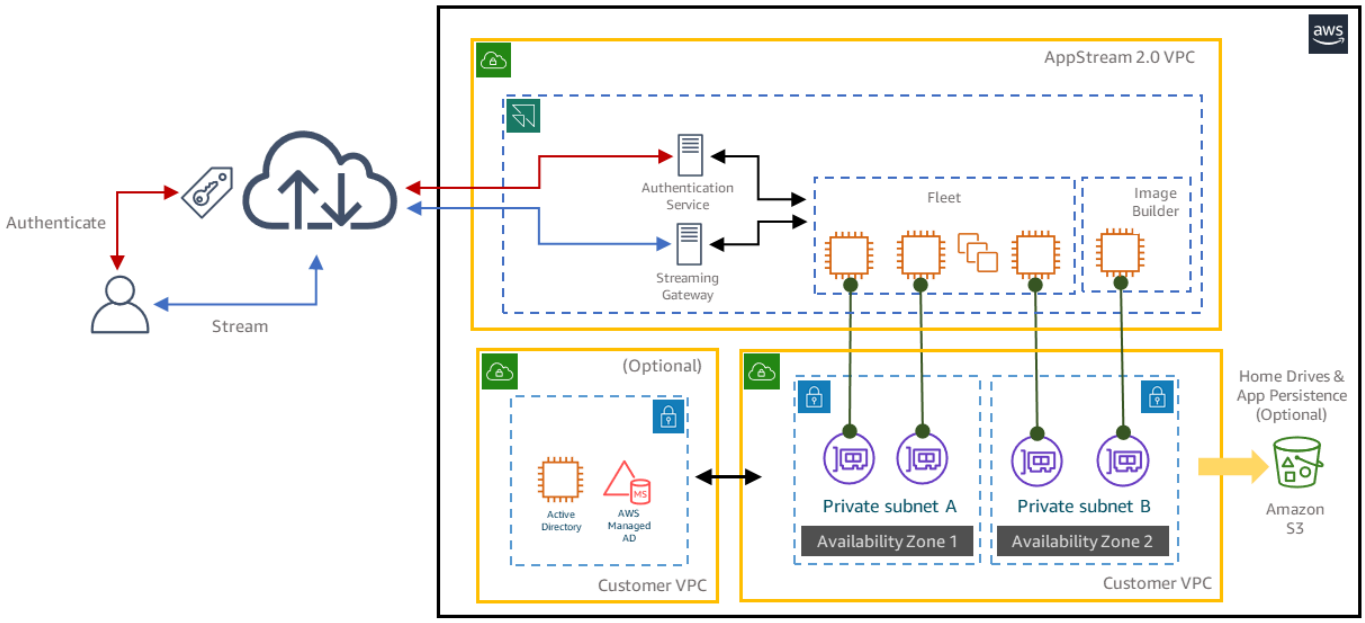

Prior to subscribing to the AppStream 2.0 service, the customer must create their own AWS landing zone and VPC, which will typically deploy subnets across multiple resilient Availability Zones. It is from these subnets that the application machine instances will communicate with external services.

Streaming and authentication traffic between the user and the application machine instance flows through a private network interface managed by the service. Each application machine instance also has a customer-facing network interface allowing connectivity to corporate resources on-premises through a Direct Connect or VPN connection or hosted locally in the cloud.

Authentication can be configured using local pooled users, which are managed by the AppStream service, or through a SAML 2.0 IdP and connectivity to the customers' existing active directory infrastructure.

An administrator connects to the Image Builder to manually configure a new or existing server instance (Windows Server 2022 or earlier supported versions) with the desired application set and desktop customizations. This process can also be automated using a variety of tools. From this image, hundreds or even thousands of users can be quickly provided with application services.

Once the image is created using the image builder, a fleet of application machine instances is defined that determine the hardware profile, active directory membership, timeout settings for users of the fleet, how many instances will be made available, and how the fleet will scale up and down.

The final step in providing visibility of the service to your end users is to create an application stack. This construct defines the naming convention for the application set, which applications will be visible to various users and groups, home drive and application persistence settings, and policy controls to limit access to clipboard, file transfer, and printing. Amazon AppStream 2.0 offers built-in storage for home drives and application settings in an Amazon S3 storage location, but other file sharing and profile management solutions such as Google Drive, OneDrive, WorkDocs, or using the AWS Windows FSx file service and FSLogix for profile management are possible.

Common Amazon AppStream 2.0 deployment scenarios

Several example problem-solutions are presented here that our typical customers encounter with their desktops. You can benefit by understanding how other customers implement Amazon AppStream 2.0 workloads to solve their VDI issues.

Scenario 1: Improve application delivery reliability

User scenario: A customer was struggling to maintain a fleet of 250 laptop devices using legacy application delivery technologies. They were experiencing configuration drift as many push installations failed or took several attempts to complete. The service desk became overwhelmed by new support calls when a new OS or application update was deployed, and users are becoming frustrated.

Amazon AppStream 2.0 provides a system to deliver a common set of applications to thousands of users from a single image. By maintaining a single image and using robust version control, a single set of applications can be delivered, updated, rolled out, and quickly rolled back if required. This customer installed the Amazon AppStream 2.0 client and delivered the required applications virtually to every laptop without needing to install local copies on every device. The applications delivered from Amazon AppStream 2.0 appeared as natively integrated into the Windows user interface, making the change almost seamless to the end users.

Scenario 2: Accelerate delivery of application updates

User scenario: We have an application that requires daily updates to deliver up to the minute features and capabilities. Our current deployment tools are struggling to reliably deliver these

updates to thousands of computers in the required timescales. We sometimes encounter issues with the new updates which can take hours to remediate.

Amazon AppStream 2.0 was selected here, as it facilitates the automated update of a centralized application image, allowing new application updates to be quickly installed and pushed out to thousands of users in a short time frame. Rollback is also quick and seamless, by reverting to a known good image if problems are detected in an application release.

Scenario 3: Simplify collaboration with business partners

User scenario: My customer would like to use application virtualization to make an internally developed application available to their key business partners, but we have no control over the management of the endpoint devices used by these external companies.

Amazon AppStream 2.0 can deliver applications into a standard HTML5 browser interface, removing the need to install client software on unmanaged devices. In this case, the customer was able to maintain complete control over the application set being delivered, allowing their business partners to securely access their services from a simple browser interface.

Scenario 4: Accelerate application delivery during acquisitions

User scenario: Customer has recently acquired a new business and needs to deliver some key business applications to the new employees. There is currently no infrastructure in place between the two organizations to allow authentication or corporate access to the parent company resources.

The customer can fast-track access to parent company assets from their new acquisition using Amazon AppStream 2.0, which provides remote access to the required applications over the internet while also being able to mandate the strict authentication requirements of the parent company.

Scenario 5: Deliver installable applications through SaaS

User scenario: Customer is an ISV who has developed a unique application to design products in a niche technology area. They would like to adopt a cost-effective way to deliver their application to thousands of customers without the costs of maintaining their own infrastructure or redeveloping their application as a SaaS offering.

As an ISV, being able to minimize the cost of delivering their applications using the service-oriented approach of Amazon AppStream 2.0 allowed this customer to maintain their competitive

advantage and to maximize their own revenues. As Amazon AppStream 2.0 application machine instances are charged at an hourly rate (or by the second for Elastic instances), it is simple to calculate the delivery costs for different service levels based on increasing hardware capabilities.

Scenario 6: Efficiently deliver applications during seasonal events

User scenario: My business needs to provide application access to several thousand customers during a seasonal event that we run couple of times every year. We currently stand-up dedicated hardware for these events, but this is costly, and the infrastructure is underutilized for the remainder of the year.

Amazon AppStream 2.0 fleets can be configured to scale up or down based on several criteria, such as the number of required instances or on a time schedule. The ability to scale and only pay for what you use was a compelling factor in this customer's adoption of the service.

Scenario 7: Provide remote application delivery to enable remote working

User scenario: It would improve our hiring process, employee engagement, and retention if we could offer the ability to work from home on a periodic basis and provide access to key productivity apps to users wherever they are geographically located. However, we don't issue laptops or mobile devices to many employees.

Amazon AppStream 2.0 provides secure remote internet access by default. All that is required to access your application set is a supported endpoint device such as a desktop, laptop, or mobile device. If the user cannot install a local Amazon AppStream 2.0 client, access is possible using an HTML5 browser. Following the adoption of Amazon AppStream 2.0, this customer can now offer remote working to their key staff, allowing them to be productive if travel to the office is disrupted or if personal circumstances mean that office attendance is difficult.

Scenario 8: Cross system application delivery by pixel streaming

User scenario: We would like to be able to deliver a new Windows application to all our users, but they are running a mixture of Windows, Linux, and macOS devices.

Amazon AppStream 2.0 can deliver Windows or Linux applications, virtually, to a diverse endpoint OS system combination such as macOS, Linux, Windows, or mobile devices while maintaining the native user experience of the application. This customer was able to standardize the delivery of a key Windows application to their supported OS systems.

Scenario 9: Test new applications rapidly and safely

User scenario: I need to test and upgrade my Windows application estate to the latest OS version, but the upgrade process is always time-consuming and support intensive.

Amazon AppStream 2.0 offers access to several OS systems versions in the Windows server and Amazon Linux line-up. You can select a later operating system version and perform exhaustive testing of your application set before migrating to this new version in a methodical fashion. No additional infrastructure needs to be deployed, and the test environment can be quickly decommissioned after testing is complete.

Scenario 10: Deliver new applications to older user devices

User scenario: My organization has an aggressive policy to meet regional environmental targets, but we are running a significant number of older application servers and endpoint devices that have green credentials which are difficult to incorporate into our plans.

Amazon AppStream 2.0 is delivered as a service using the AWS Cloud, which invests heavily in environmentally friendly data centers, providing businesses with a cost-effective way of realizing regional environmental targets with minimal effort. Thin or zero client devices can be deployed to access AppStream services, removing the power, heat, and recycling overheads of traditional endpoint devices.

Scenario 11: Help prevent data leakage when using remote applications

User scenario: As a financial, insurance, development, or design business, we need to deliver some key business applications to users across the globe but are concerned about leakage of our critical business intellectual property.

Accessing applications using Amazon AppStream 2.0 moves the execution of applications into a centralized and secure AWS Cloud environment which can be more tightly controlled than a distributed application delivery approach. Amazon Appstream 2.0 virtual channels which control clipboard, local file access, and printing can also be disabled, significantly reducing the likelihood of data leakage to uncontrolled endpoint devices and locations. As the applications are streamed to remote users, only delta changes in the application display are streamed over an encrypted virtual channel, securing information exchanged between client and server.

Scenario 12: Reduce latency for client-side applications

User scenario: We are an existing AWS customer with significant investment in AWS services, and we are looking for a way to localize access to strategic data lakes and AI/ML systems. How do we optimize the performance of the client-side applications that use these services?

Amazon AppStream 2.0 runs as an AWS service in many Regions across the globe. Accessing client applications using Amazon AppStream 2.0 verifies that they run near your other AWS services, maximizing network performance between the client and the data being manipulated.

Scenario 13: Deliver applications to users efficiently when training

User scenario: We are a training company who needs to increase our reach and reduce our service costs by offering training courses that can be attended remotely.

Amazon AppStream 2.0 runs in 15 AWS Regions, which means that deploying a training application with global reach is simply a matter of deploying Amazon AppStream 2.0 from multiple Regions. As Amazon AppStream 2.0 instance charges are based on hourly usage and scale across a wide range of instance types, you can deliver a range of different training courses with a predictable baseline cost.

Scenario 14: Deliver graphics applications efficiently

User scenario: We are an engineering company with significant requirements in terms of graphics processing for our development applications. How can Amazon AppStream 2.0 deliver the performance we need to take advantage of application virtualization?

Amazon AppStream 2.0 offers several graphics instance types offering NVidia T4 GPUs with up to 64Gb memory, which is adequate for a wide range of engineering application needs. Furthermore, the Amazon DCV protocol used by Amazon AppStream 2.0 to deliver remote access to your applications is highly optimized to deliver the best performance for highly graphical applications over a wide range of network conditions.

Amazon delivering browser-based applications

An increasing number of applications are being delivered to users using a web browser. What was once an exceptional use case for an infrequently used application is now a mainstream delivery mechanism for a vast array of applications being used within AWS customer organizations.

It can be challenging to deliver web-based applications securely, efficiently, and at a low cost. Amazon WorkSpaces Secure Browser provides this capability in the form of a low cost, fully managed, Linux-based service designed to facilitate secure browser access to customers' internal websites and software as a service (SaaS) applications from existing web browsers. This is achieved without the administrative burden of appliances, managing infrastructure, specialized client software, or VPN connections.

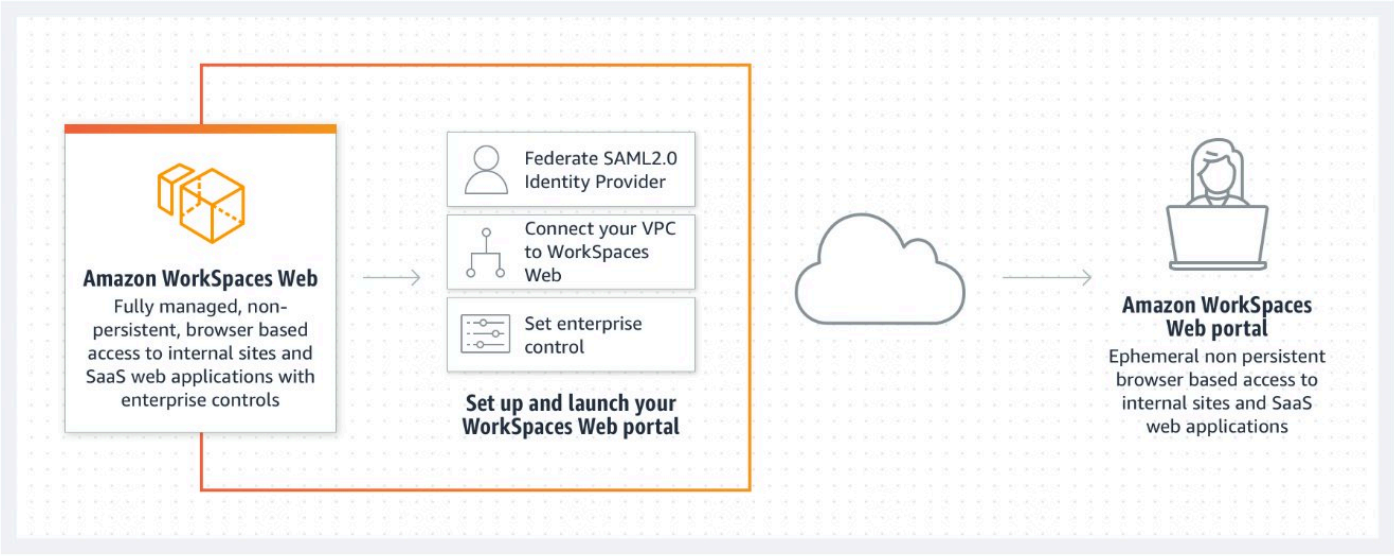

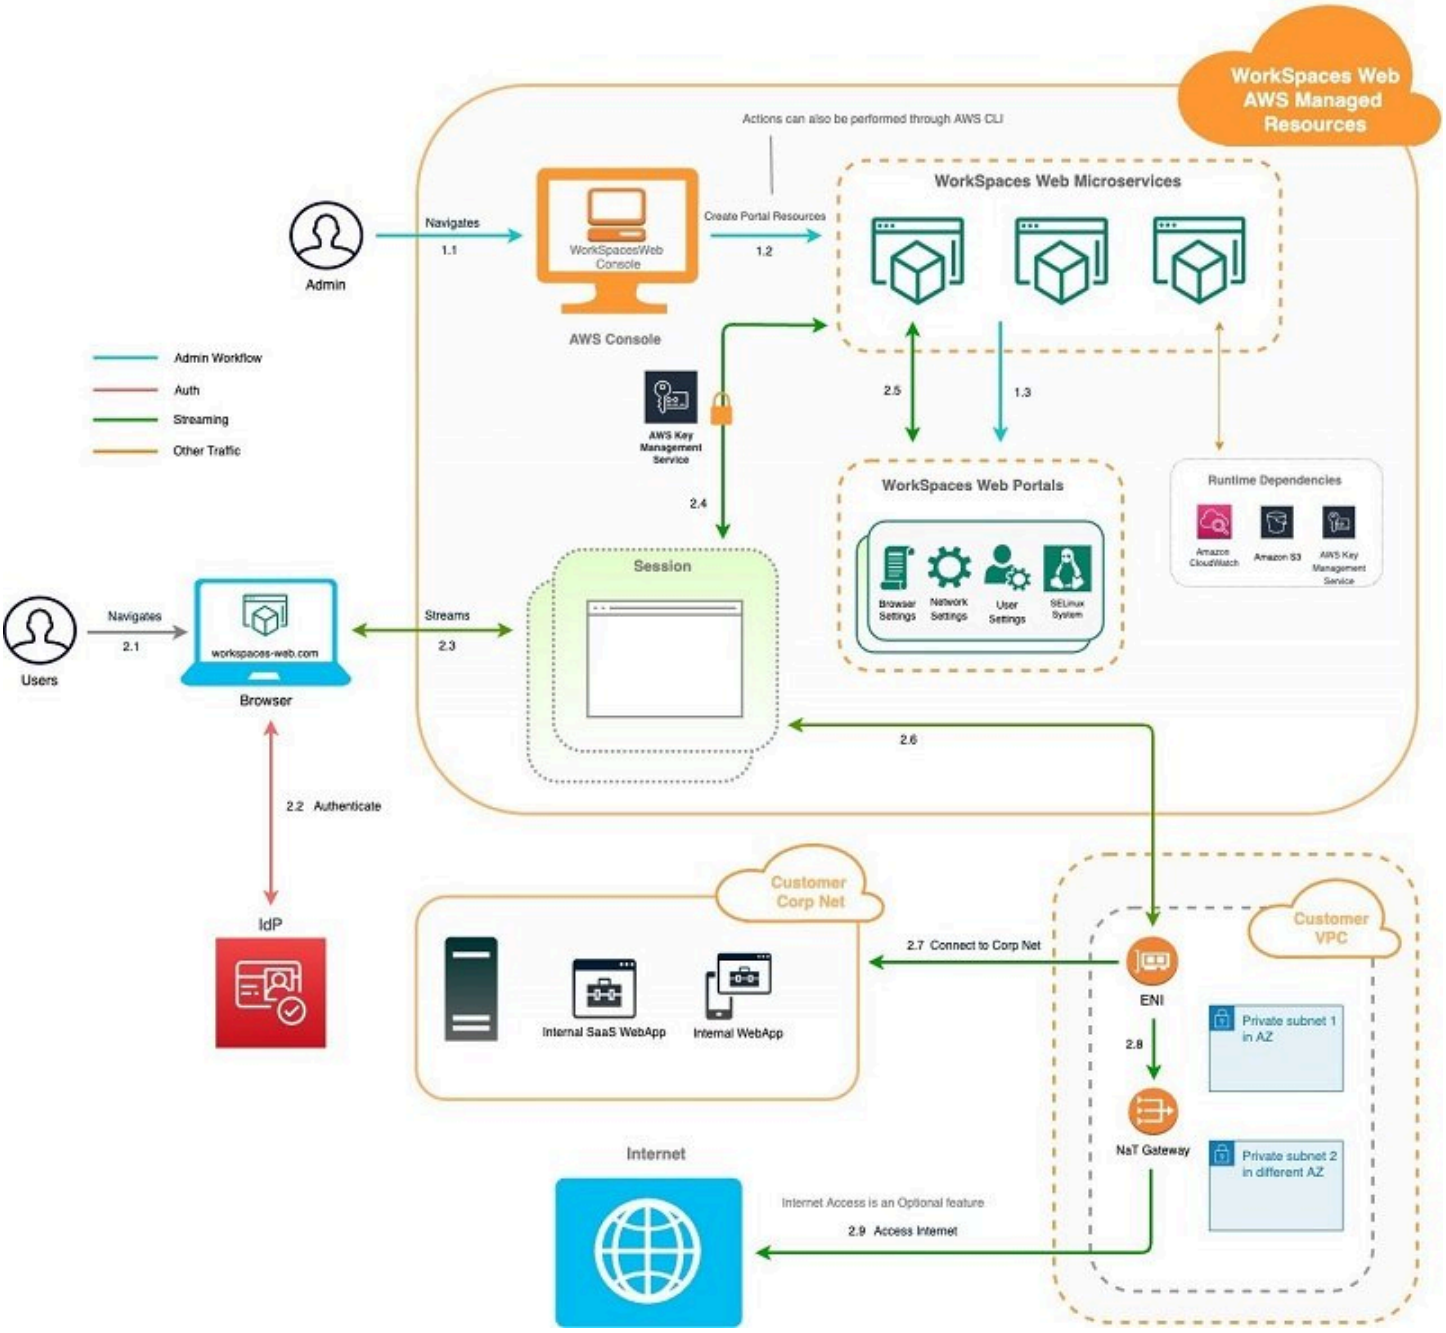

Amazon WorkSpaces Secure Browser is underpinned by a non-persistent system delivering a Chromium-based browser to users each time they authenticate and access the service. This verifies that each user is using an updated browser that is secure and has access to the resources that the user needs to fulfill their role. At the end of a user's session, they disconnect from the browser, and their browser session is terminated, avoiding sharing of data that may be cached within their browser.

The service automatically scales to satisfy the demand from users across a working day and week, and it is a cost-efficient service for delivering secure web applications to users. This is due to both the use of resources only when required by users and the managed service, which removes the need for organizations to operate and maintain a system that delivers a browser to users.

WorkSpaces Secure Browser provides compatibility with a broad range of web sites due to the use of a standard web browser. This compatibility verifies that web sites do not need to be re-developed to support WorkSpaces Secure Browser.

While WorkSpaces Secure Browser uses Chromium web browsers running on AWS resources with connectivity to customer VPCs, you should also provide connectivity to the web sites that require user access permissions. To do so, use the AWS networking service portfolio to provide connectivity to an organization's required intranet or internet web sites.

Users access WorkSpaces Secure Browser by authenticating against a SAML 2.0-compatible identity provider of the customer's choosing. Once authenticated, the user launches a WorkSpaces Secure Browser session and initiates the streaming of pixels to the user's endpoint device. The user's endpoint device must be a supported device running a browser and could be either BYOD or device owned and supplied by the user's employer.

Common WorkSpaces Secure Browser deployment scenarios

Several example problem solutions are presented here that our typical customers encounter with providing access to web applications. You can benefit by understanding how other customers implement WorkSpaces Secure Browser workloads to solve browser-based issues.

Scenario 1: Isolate access to specific browse applications

User scenario: We need to isolate the client component of a browser-based application so that users can access it from anywhere, remove possibilities of data leakage, and avoid the need and cost of delivering this on a full desktop.

WorkSpaces Secure Browser is a fully managed service that provides a resilient and scalable system to access web-based services. AWS manages the browser instances, Chrome policies can be used to secure the browser, and network controls, such as security groups, can be used to limit traffic flow. You can also configure WorkSpaces Secure Browser to obfuscate sensitive data such as credit card or social security numbers to assist with accommodating compliance initiatives such as PCI.

Scenario 2: Provide simplified access to web applications

User scenario: We have several workers who need a simple process to access their time sheets and other scheduling tools. These tools are all web-based, and the users will have no access to a computer due to the challenging nature of their work environment.

WorkSpaces Secure Browser can be accessed from a thin client device and provides a kiosk mode experience which only provides them with access to a limited application set. The thin client device can quickly be deployed, secured, and installed in areas where it may be challenging to install traditional computers.

Scenario 3: Provide safe access to external websites

User scenario: We would like to provide some users with the ability to access a wider range of external web sites from some devices but are concerned about the propagation of viruses or malware inside of our organization.

WorkSpaces Secure Browser is hosted by AWS in a secure environment, and each user's browsing instance is destroyed at the end of their session. Standard Chrome policies can be used to lock down each browser instance, and you can fully isolate the browsing environment and optionally have network traffic pass through a web proxy to further limit access to untrusted external sites. This solution significantly reduces the possibility of external browsing activities from compromising your production networks.

Operational excellence

Operational excellence discusses the foundational questions that need to be addressed in order to develop, run and support EUC workloads effectively, gaining insight into your operations, and to continuously improve supporting processes and procedures to deliver business value. For more information, see the Operational Excellence Pillar whitepaper.

Focus areas

Design principles

- **Design for the cloud:** You may be delivering a net new EUC service, migrating from an existing on-premises solution, or a hybrid of both, but regardless of approach, when you use cloud computing you can face the challenges and gain the advantages of the wide range of services on offer. Train your technical architects, delivery teams, and operations personnel so that they are ready to embrace this range of services to deliver cost efficiency, scalability, unparalleled resilience, and the global reach that AWS provides.
- **Embrace automation:** The opportunities to automate the provisioning of desktop and application delivery services increase significantly when moving to an AWS Cloud-hosted EUC service. Automation accelerates your time to deliver value. This is because it can cover selection of a variety of compute, software bundles, and APIs in order to streamline configuration, deployment, delivery, maintenance, and ongoing support for an EUC environment.
- **Define the service levels you need to provide**: The way your EUC workload is designed, built, and managed should be predicated upon business metrics that define the value of each service being offered. Some user personas, for example, may require higher resilience, better performance, and closer monitoring than users who require access to simple productivity applications. Service levels should reflect these requirements.

- **Understand your compliance requirements:** Many business verticals must adhere to stringent security and compliance standards and controls, which affect how desktops and applications must be deployed, managed, and monitored. Conditional access to services, secure storage of user data, or maintaining robust logging processes, for example, will all dictate how an EUC workload should be designed. For additional guidance on accommodating compliance requirements, see Security.
- **Understand your unique user personas:** When building new or migrating existing EUC workloads to AWS, it is critical to understand how each persona will need to interact with the services provided. The compute resources required, the desktop OS types you will need to support, the applications that each group requires, and the peripheral devices used must all be understood to design and build a solution that offers a superior user experience and reduces operational overheads by simplifying your delivery, maintenance, monitoring, and support processes.
- **Collaborate between key technology teams:** When designing and delivering EUC workloads, significant expertise is required across compute, storage, networking, security, and desktop or application delivery to effectively deliver successful outcomes. From project initiation, assemble cross functional teams and call upon expertise from every relevant technology discipline. Continue to involve these teams after the new EUC solution has been deployed to verify that evolution of the service continues to deliver incremental business value.
- **Implement robust change control processes:** EUC relies on many underlying technology stacks to deliver a reliable and performant service. For example, a change made to networking, authentication services, database services, or perimeter security can result in disruption for desktop and application users. To deliver service levels expected by your business, design a change control process that provides visibility and input into changes to supporting services to the EUC team. Follow the generic principle of making frequent, small, reversible changes.
- **Frequently refine operational processes:** The implementation of robust logging, alerting, and event management is important to all aspects of delivering an EUC service. Logging and monitoring help you retrospectively learn from events such as system outages, performance issues, and scaling activities to avoid future problems, improve user experience, and fine tune the environment to save costs.
- **Minimize duplicates and false positives:** When monitoring your EUC workload, take steps to reduce the number of events that are propagated to support teams for analysis and remediation. Over time, fine tune which level of skill is required to respond to each event type and route these to the most appropriate resources. Avoiding event fatigue helps verify that the most critical issues are surfaced and dealt with more efficiently.

- **Deploy systems health dashboards:** Performance dashboards simplify the reporting of overall systems health and allow support teams to quickly drill down into specific areas of concern. Creating useful dashboards helps your team spot service issues quickly and implement remediation processes. To improve user engagement, take a proactive approach to identifying and surfacing issues before they cause significant disruption. This also verifies that the provision of desktop and application delivery services is meeting the service levels required by the business.
- **Create and maintain service documentation:** Maintain design documentation from the initial deployment as the EUC workload evolves. This makes it much easier to hand over maintenance and support to new team members or partners and provides a solid architectural reference as the workload grows to incorporate new features and capabilities.
- **Provide adequate training and exposure to industry events:** Train architects, delivery teams, and support personnel according to their role in the organization. When you verify that each team has the right level of training to carry out their day-to-day activities, your teams can carry out deployment, maintenance, and support of the environment according to the relevant best practices. For example, have business and architect level employees attend industry specific events such as AWS re:Invent to verify that they have the necessary information on new products and features to evolve your EUC workloads. This training helps your teams take advantage of cost savings, resilience improvements, and other new services that can positively impact your business.

Organization

There are many organizational considerations that need to be made when striving for operational excellence while delivering AWS End User Computing solutions, these are discussed in detail in the overall AWS Well Architected Framework.

The following questions should be considered when evaluating organizational impact on the deployment of AWS EUC services.

EUCOPS01: Has a project board been convened which includes stakeholders and represent atives from key areas of the business?

A new deployment of cloud based EUC services requires a detailed understanding of how existing services are used and the business objectives which are met by delivering those services. A project

board containing authoritative and knowledgeable business representatives will benefit from historical knowledge and accelerate project delivery.

EUCOPS02: Have you convened a technical project team that represents the key technology areas which will integrate with AWS EUC services?

AWS EUC services integrate closely with many ancillary technology stacks. Verifying that the project team has representation from each of these areas can increase the chances of project success.

EUCOPS03: What business or technology problems will be solved by adopting the selected AWS EUC technologies?

Deploying a new service or technology should be the result of identifying a compelling business requirement or problem to be solved that derives value for the business. When requirements are clearly defined, quantifying success criteria and measuring success is more straightforward.

Many of the scenarios discussed in this lens map to specific business challenges which are addressed by specific EUC technologies.

EUCOPS04: If migrating to AWS EUC services from an existing platform, have you prioritized a list of mandatory features and capabilities?

If you are migrating from an existing EUC system, you may have used specific, sometimes mandatory functionality that delivers a focused benefit for some user personas. Identifying these features and capabilities is essential to maintain feature parity or save costs if these features have not been delivering the expected benefits.

Best practices

- EUCOPS01-BP01 Build a project team which includes executive level sponsors and relevant business and technical communities
- EUCOPS02-BP01 Engage technical stakeholders from all disciplines that affect your EUC services

- EUCOPS02-BP02 Build a matrix of all internal and external stakeholders who may be affected by changes to the way you deliver EUC services
- EUCOPS03-BP01 Identify the business goals and success criteria for your EUC project
- EUCOPS04-BP01 Identify the key capabilities and features that deliver business value and drive project success

EUCOPS01-BP01 Build a project team which includes executive level sponsors and relevant business and technical communities

When starting your AWS EUC project be sure to convene a project board which has sponsorship from a significant influencer or senior decision maker and buy in from both business and technology stakeholders, this will make sure that the views of the business and technology teams are considered when delivering the new service. Each invested project board member will have their own view of the governance and organizational challenges likely to arise throughout the project lifecycle, and their input is likely to reduce risk and enhance the chances of project success.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Many frameworks exist that define a structured project approach, including advice on project board formulation. The Prince II methodology is a common example that embraces this approach.

Implementing a structured approach or following a proven project management framework will make sure that project requirements, key timelines and success criteria are well documented, and that day-to-day tracking of project activities and progress towards key milestones is in effect.

EUCOPS02-BP01 Engage technical stakeholders from all disciplines that affect your EUC services

The deployment of AWS EUC services typically requires integration with many diverse technology areas. Build a project team which includes experts from multiple technology disciplines to identify resourcing issues early on, understand key technological blockers affecting deployment of the AWS EUC project, and accommodate and manage key processes and procedures used by each team to deliver successful outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Draw upon experiences from across all of your technical teams to make sure that relevant insights are considered as you navigate through the design, deployment and support of a new EUC deployment.

EUCOPS02-BP02 Build a matrix of all internal and external stakeholders who may be affected by changes to the way you deliver EUC services

To provide an initial starting point for building a project team, collate a list of all the technology areas which are directly or indirectly affected by an AWS EUC migration. Additionally, create a corresponding list of each service owner. This matrix helps you oversee a successful implementation or migration.

Level of risk exposed if this best practice is not established: High

Implementation guidance

EUC services in general, require a good understanding of many technology disciplines to design and deploy secure, reliable, and scalable solutions. The team that manages the AWS EUC services must have skills across general cloud principles, compute, storage, networking, applications, and security at a minimum to implement services that deliver against business requirements and internal and external SLAs while maintaining a first class user experience.

While technical skills in the EUC discipline are key, incorporate the teams responsible for maintaining other key processes, such as problem, incident, and capacity management and change control, into decision making processes. After a deployment or migration, these teams will be responsible for ongoing support and user experience.

Encourage participating teams to interact with each other in order to exchange opinions and ideas. Fostering collaboration and gathering diverse opinions typically results in a better overall solution, improving service and support.

Provide these teams sufficient resources not only to manage and maintain the planned infrastructure but also to perform continual service development. Plan to accommodate new features and functionality that meet evolving business needs. If your technical teams lack the AWS expertise necessary to deploy or migrate to Amazon WorkSpaces, Amazon AppStream 2.0, or

Amazon WorkSpaces Secure Browser, consider engaging with AWS Professional Services or one of our partners.

EUCOPS03-BP01 Identify the business goals and success criteria for your EUC project

Verify that the deployment of the selected AWS EUC services addresses the needs of both internal and external customers. For example, does the feature set of the selected AWS EUC technology stack meet the requirements of all user personas identified as part of the proposed project? Both employees, business partners, and external customers could be affected.

While adopting AWS EUC services to take advantage of a service oriented, pay as you go cost model has financial benefits, be sure that you understand the technological impact of this approach for both internal and external customers.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Most organizations have personas that benefit from a deployment of or migration to a cloud-hosted EUC service. While the business problems being solved may be clear, such as cost reduction, increased agility and resilience, or global reach, it is also important to evaluate whether the AWS EUC services being deployed meet the requirements of each user persona.

Create a matrix of user personas that captures each unique set of hardware and software requirements, accessibility options, and access requirements, highlighting where technology limitations may need to be considered to accommodate other benefits.

EUCOPS04-BP01 Identify the key capabilities and features that deliver business value and drive project success

While AWS EUC services may offer feature parity with your incumbent vendor, achieving this parity may require additional engineering effort to integrate with existing operational or support systems.

Planning for the proof of concept (PoC) or pilot phase of a migration project is an opportunity to document acceptance criteria and define a list of the features and functionality currently being delivered by your existing vendor. When moving into pilot and production, these documents help you verify that all mandatory functionality has been tested and can be successfully delivered.

Sacrificing features in order to take advantage of a more flexible cost and delivery model for some user personas may be an acceptable approach. It may be feasible to deliver a high percentage of your existing use cases, covering most of your user population using built-in functionality and leaving just a small amount of engineering or the adoption of third-party solutions to accommodate the remaining use cases.

Level of risk exposed if this best practice is not established: High

Implementation guidance

A PoC or pilot may reveal that it is not possible to deliver all of the features and functionality offered by the incumbent EUC system. It may still be possible to roll out a significant part of the project and reduce costs or realize many of the other benefits of cloud delivery while investigating ways to bridge functionality gaps in areas where feature parity cannot be maintained.

Following are some examples of where it may be possible to dispense with bundled EUC functionality which is either no longer required or has been deprecated by new and improved capabilities:

- Many existing EUC or VDI system features, such as those that optimize compute resources, network bandwidth, or audio and video delivery, may no longer be required, as compute, network, and media capabilities have vastly improved over time.
- Accessing your desktop or application resources from an HTML5 browser as standard may be a significant change for the user experience, but standardization may offer operational and support savings in the medium to longer term.
- Deploying WorkSpaces with Ubuntu for developers may reduce development costs for a large population of users, moving away gradually from an incumbent, more costly EUC solution.
- Using a vendor-supplied profile management solution may now be less functional and performant than using a standard Microsoft solution such as FSLogix.
- It may be possible to dispense with complex legacy remote access solutions that have evolved over time in favor of the pervasive and secure capabilities available with the current generation of Amazon WorkSpaces and AppStream remoting protocols.

Prepare

To prepare for operational excellence, you must understand your workloads and their expected behaviors. You can then design them to provide insight to their status and build the procedures to support them.

EUCOPS05: Which tools will you use to monitor the end-to-end health of your AWS EUC service and other dependent services?

Deciding on the tools you use to provide insight and observability into your AWS EUC deployment is key to maintaining acceptable levels of performance and availability. Using well-known, incumbent tools while cloud alternatives are reviewed and deployed should be a part of your design and deployment process.

EUCOPS06: How do teams create, maintain and support your EUC service deployment?

While training is essential to verify competency across all aspects of AWS EUC service delivery, you should also provide a discrete training environment that won't impact production services.

EUCOPS07: How do you document the processes and procedures used to create, maintain, and support new AWS EUC service deployments?

Create and agree upon specific processes to write and continually update service documentation as your AWS EUC services evolve over time.

EUCOPS08: Do you have an existing, formalized change control process?

Closely control both direct changes to the AWS EUC service itself and indirect changes that affect dependent infrastructure. Formally review any activities that mightaffect the performance,

availability, or cost structure usinga cross-disciplinary panel of experts and test all changes before they are deployed into production.

EUCOPS09: How do you provide service and support readiness in production?

Each workload, desktop, or application service delivered on your AWS EUC system will have a specific architecture, a cadence for updates and patching, metrics for performance monitoring, and availability and resilience requirements. Train and provide appropriate permissions to a dedicated team with deep insight into delivering each workload to verify that each is delivered successfully, promoting user satisfaction and engagement.

EUCOPS10: What is the rollout and training plan for EUC service users?

While operational training is essential to provide the skills to build and maintain each service, user training is equally important to promote employee engagement and to reduce support calls. Project success or failure can pivot upon successful adoption of the services provided, and successful business outcomes can only be realized if the service users can efficiently use its capabilities.

Best practices

- EUCOPS05-BP01 Identify monitoring tools to provide the expected levels of insight into operational performance
- EUCOPS05-BP02 Store and regularly analyze log files to detect anomalous activity and behaviors
- EUCOPS06-BP01 Deploy test, development, and pre-production environments to reduce risk to production services
- EUCOPS07-BP01 Formalize the mandatory creation and maintenance of all EUC service-related documentation
- EUCOPS08-BP01 Adopt a mandatory and formalized process for managing any changes to EUC services and dependent infrastructure
- EUCOPS09-BP01 Maintain an up to date matrix of all EUC service owners and quick access links to the support plans for each service

- EUCOPS09-BP02 Allocate training time so your teams can build and maintain their skills to deploy and manage your AWS EUC environment
- EUCOPS10-BP01 Encourage user participation during service development and rollout to maximize engagement and project success

EUCOPS05-BP01 Identify monitoring tools to provide the expected levels of insight into operational performance

While existing, familiar tools can be used to monitor an AWS EUC deployment, there are many AWS services, such as automatic Amazon CloudWatch dashboards for Amazon WorkSpaces and Amazon AppStream, AWS CloudTrail for API call monitoring, and Amazon Kinesis for log propagation and centralized log storage.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Implement proactive monitoring of the health of all aspects of an AWS EUC deployment to quickly identify and remediate problems that affect the user population, their productivity, and any impact this may have on the business.

For both Amazon WorkSpaces and Amazon AppStream 2.0, it is important to monitor both the service itself in addition to any external service dependencies. Consider the following monitoring tools:

Amazon WorkSpaces

Amazon CloudWatch provides an automatic dashboard which gives an overview of overall service health, including:

- Available or unhealthy WorkSpaces
- Session launch times
- Connection success and failure
- Session latency
- Users connected, disconnected, stopped, or in maintenance

Additional metrics, such as instance specific CPU, memory, and disk performance can also be viewed. Develop custom CloudWatch widgets to fine tune the monitoring of specific groups of WorkSpaces.

Amazon AppStream 2.0

Amazon CloudWatch provides an automatic dashboard which gives an overview of overall service health, including fleet capacity and utilization.

CloudWatch alarms can be configured to send alerts when specific thresholds are met.

Each WorkSpaces and Amazon AppStream 2.0 instance exposes a network interface in the customers managed VPC which can be addressed by third party monitoring tools for traditional management.

As AppStream instances are ephemeral, logs required for compliance or historical monitoring, such as event logs, can be harvested at user logoff or shutdown using session scripts or in real time using services such as Amazon Kinesis.

External dependencies

Monitoring should also be in place for:

- Internet connectivity (user to Amazon WorkSpaces or Amazon AppStream 2.0 service)
- Amazon networking
- Active directory
- RADIUS (or other MFA provider)
- Microsoft PKI (If certificate-based authentication is in use)
- SAML 2.0 Identity Provider (IdP) availability (If SAML 2.0 authentication is in use)
- Private certificate authority (if certificate-based authentication is in use)
- User data repositories (like file shares and profile stores)
- Application web tiers
- Application database tiers
- Application licensing servers
- Web proxies
- Anti-virus infrastructure

If these services are hosted on Amazon EC2, Amazon CloudWatch can be used to monitor key health metrics and alert when service degradation is detected.

For services still hosted on-premises, Amazon CloudWatch agents can be installed which send key metrics to Amazon CloudWatch.

Log propagation

For centralized gathering of log files for troubleshooting and retrospective analysis, Amazon Kinesis agents can be deployed on WorkSpaces or AppStream 2.0 to deliver real-time propagation of OS and application-level logs to a central location.

For Amazon AppStream 2.0, propagating instance log files in real time to a centralized location is essential if you need to store logs for compliance purposes, as AppStream instances are destroyed at the end of each session. For more detail, see Using the Kinesis Agent to store AppStream 2.0 Windows event logs.

AWS Health dashboard

The AWS Health dashboard provides insight into the health and availability of AWS services running across regions. Individual regional services can be filtered in the web page or added to an RSS feed reader for additional visibility.

EUCOPS05-BP02 Store and regularly analyze log files to detect anomalous activity and behaviors

Maintaining a central store of log data and performance metrics is frequently a mandatory requirement if specific compliance standards need to be maintained. Even in the absence of compliance requirements, maintaining a central store of data facilitates a better understanding of service scaling, performance, and security enables analysis, which improves root cause analysis and drives incremental service improvement.

Review the available data sources that provide key insight into the usage of your EUC environment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Extracting performance data and log files from both WorkSpaces and AppStream 2.0 and storing it centrally is essential if you need to adhere to specific industry compliance standards or if you

want to perform retrospective analysis of data for troubleshooting purposes, root cause analysis, or predicting service scalability and requirements.

Amazon CloudWatch can be used to capture specific metrics and store the data longer term in Amazon S3. Amazon Kinesis agents can also be installed on WorkSpaces or AppStream 2.0 instances to propagate system logs in real time to a centralized location. For more detail, see Using Amazon Kinesis Agents to Store AppStream Event Logs.

EUCOPS06-BP01 Deploy test, development, and pre-production environments to reduce risk to production services

Training and testing should be undertaken in isolated environments, with little or no connectivity to production services.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Separate the EUC environments used for training, testing, and development from production services to reduce the risk of non-production activities affecting normal business operations. Create discrete test environments and use them for activities that could disrupt production services, cause outages or performance degradation, or compromise security.

When testing and staging new releases or updates to production systems, these should be undertaken in a separate environment that matches the current production deployment as closely as possible. This reduces the likelihood of issues arising from disparities between test, staging and production.

AWS EUC services are Regional in nature and delivered on an AWS account by account basis. Multiple Regions and accounts can be deployed to separate training, testing, and production environments.

Multiple AWS accounts can also be deployed to separate production workloads for scalability reasons, helping to avoid having all resources in the same place or where service separation is necessary to align with compliance or security requirements.

AWS Control Tower can be used to streamline the management and governance of multiple AWS accounts.

Unlike on-premises infrastructure, Amazon WorkSpaces and AppStream 2.0 environments can be deployed using automated processes and only attract costs while in use.

AWS CloudFormation templates can be used to deploy new AWS services such as WorkSpaces and AppStream 2.0 to reduce the likelihood of human error and reduce configuration drift.

AWS Systems Manager Runbooks can be used to automate some aspects of WorkSpaces deployment. For more detail, see SSM Runbooks for Amazon WorkSpaces.

EUCOPS07-BP01 Formalize the mandatory creation and maintenance of all EUC service-related documentation

Maintain a library of documentation related to the business requirements, architectural design, service delivery, and support of your EUC deployment.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Create deployment and operations guides and keep them updated over time verify that all processes used to install, administer, update, and maintain the AWS EUC environment are captured. This documentation provides an effective method of communicating how the environment should be managed to new administration or support staff and external partners, when required.

As iterative operational testing takes place, use lessons learned from failover and DR testing to evolve the deployment and operations guides and capture relevant changes that were needed in order to successfully complete testing.

While this documentation does not provide an exhaustive list of all aspects of a deployment that should be captured, gather as much detail of the end to end service configuration and the subsequent management processes as possible.

For each of the following topics, if a manual installation was performed, capture the specifics of how and why you configured each setting. If the deployment was automated, document the methods used (like AWS CloudFormation or Terraform), and call out the specifics of how and why each configuration decision was made.

Infrastructure build

How were the landing zone and your WorkSpaces or AppStream 2.0 environments created, which options were configured for each service, and why? CloudFormation templates can be used to reliably and repeatably build the baseline infrastructure and the rationale behind the

CloudFormation template creation. Deployment and rollback processes can be captured and documented.

Active Directory or RADIUS integration

Your Active Directory and RADIUS deployment and maintenance should be part of a separate operations guide chapter. For WorkSpaces and AppStream 2.0, capture the specifics of how you integrated Active Directory and RADIUS for the respective service. For WorkSpaces, document which directory integration method was used, and capture the manual steps used to deploy or details of the CloudFormation templates used to automate this process.

SAML 2.0 or certificate-based authentication

How was your SAML 2.0 IdP configured with respect to integration with Amazon WorkSpaces or AppStream 2.0? Which SAML attributes were used to drive AppStream application entitlements?

How will you monitor and manage the certificates used to build a chain of trust between AWS IAM and your SAML provider?

For certificate-based authentication, capture the installation choices taken and the integration points with Microsoft Certificate Services.

How will you manage certificates and expiry for integration between CBA and Microsoft Certificate Services?

Image management

Document the process followed to create each of your custom images. Which updates and hotfixes were applied, which applications were installed, how were they configured, and which registry or file system changes were required?

How were your applications installed and deployed (for example, did you use local images, App-V, MSIX, AppVolumes, network share, or third party isolation products?).

For AppStream 2.0, did you use session scripts? Document the scripts deployed and what each script does.

For WorkSpaces BYOL deployments, document the process followed to extract your Windows 10 or 11 images, sanitize them, and import into Amazon WorkSpaces.

How should image updates be managed, which version control and naming conventions will be applied, and how will you roll back to a known good configuration if required?

Client deployment

Which clients are required to access Workspaces or AppStream 2.0 (for example, Windows, macOS, or web), which user groups require each client type, and how should it be installed? How will clients be updated?

EUCOPS08-BP01 Adopt a mandatory and formalized process for managing any changes to EUC services and dependent infrastructure

Create a new process or integrate with existing processes that track all changes that can affect the stability and security of your EUC deployment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Track all changes which might directly or indirectly affect the performance and availability of Amazon WorkSpaces or Amazon AppStream 2.0 services. Implement a change control process that documents each service update with a robust risk assessment and back-out plan and involves technology stakeholders from all relevant areas. This process can reduce the risk of service outages or degradation.

Both WorkSpaces and Amazon AppStream 2.0 have key dependencies on many external services. If changes to any of these services is required, a representative from the AWS EUC team should be part of the change control team to review and quantify the risk of the change.

The service dependencies for Amazon WorkSpaces and Amazon AppStream 2.0 include, but are not restricted to:

- AWS networking
- Active Directory and connectors
- RADIUS
- Microsoft PKI (if certificate-based authentication is in use)
- Third Party PKI services that may be used to allocate public certificates for related services
- AWS KMS if used for encryption of WorkSpaces images
- SAML 2.0 IdP availability (if SAML 2.0 authentication is in use)
- Private certificate authority (if certificate-based authentication is in use)
- User data repositories (like file shares or profile stores)

- Application web tiers
- Application database tiers
- Application licensing servers
- Web proxies
- Firewalls and other related security infrastructure
- Anti-malware infrastructure
- Thin client management infrastructure

Amazon WorkSpaces and Amazon AppStream 2.0 also use other AWS services, such as Amazon EBS and Amazon S3 for storage, so you should understand any changes being made to these systems.

EUCOPS09-BP01 Maintain an up to date matrix of all EUC service owners and quick access links to the support plans for each service

Amazon WorkSpaces and Amazon AppStream 2.0, although easier to implement and administer than traditional on-premises alternatives, still require specific knowledge to deploy, manage, and support. To simplify the process of routing issues to the right owners, you should be able to quickly identify the teams who are responsible for implementation and support along with clear support plans for each application being delivered, expediting time to resolution.

Each application delivered by WorkSpaces or AppStream 2.0 should have a formalized support plan with designated business and technical owners who are responsible for and involved in the deployment, maintenance, and support of each application and its dependent technology stacks.

Each application set should have its own designated level of criticality, with associated SLAs that are clearly understood by the support teams involved. For disaster recovery purposes, the business should be able to identify relevant RTO and RPO parameters which each service should be engineered to accommodate so that critical business services can be delivered even under the most challenging circumstances.

If you are delivering WorkSpaces or AppStream 2.0 across multiple AWS Regions, verify that a support and escalation mechanism exists that documents the transfer of responsibility between regions when required. This documentation is important to sustain support efforts across time zones, maximizing service continuity.

Note: Your business RPO and RTO requirements may be more aggressive than the service can provide, and discrete groups of users may have different RPO and RTO requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Create a process to quickly identify roles and responsibilities for each application stack so that support teams can quickly identify the resources that need to be employed and address any issues in service delivery.

Resources

- WorkSpaces Service Level Agreement
- AppStream 2.0 Service Level Agreement

EUCOPS09-BP02 Allocate training time so your teams can build and maintain their skills to deploy and manage your AWS EUC environment

Training and enablement are key to maintaining a reliable, successfully-evolving EUC deployment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Provide targeted training on the AWS Cloud, Amazon WorkSpaces, and AppStream 2.0 to verify that architects, administrators, and support personnel all have the relevant skills to design, deploy, and maintain the AWS EUC environment. Give this training through authorized training courses and professional accreditations and create a training environment that can be used for evaluation and self-instruction, augmenting official coursework.

The core tenets are the same for WorkSpaces and AppStream 2.0 as they are for delivering a remotely accessed, centralized, and virtualized desktop and application delivery service, either on-premises or from an alternate cloud vendor. Skills in these areas are transferrable to deploying and managing AWS EUC services. It is essential for your deployment teams to have a good understanding of compute, networking, storage, virtualization, and application delivery, at a minimum.

Technical teams may need to be prepared in different ways depending on the nature of the adoption of Amazon WorkSpaces and AppStream 2.0 services:

Greenfield a net new deployment with no prior cloud or EUC skills

Teams need to be trained, and they should iteratively maintain their skills in AWS core competencies such as cloud delivery, compute, networking, and storage, in addition to specific training and exposure to Amazon WorkSpaces and AppStream 2.0. Focus on understanding the core tenets of cloud delivery such as reducing costs, increasing scalability and resilience, and taking advantage of the global reach of AWS Cloud services. This may be an area where AWS Professional Services or an AWS Partner may be able to add significant value until your technical teams are familiar with the technologies involved.

A net new deployment with existing EUC skills, but no prior cloud infrastructure skills

Teams need to be trained, and they should iteratively maintain their skills in AWS core competencies such as cloud delivery, compute, networking, and storage. Focus on understanding the core tenets of cloud delivery such as reducing costs, increasing scalability and resilience, and taking advantage of the global reach of AWS Cloud services.

Teams should still be trained on and exposed to Amazon WorkSpaces and AppStream 2.0, but technical teams with prior experience deploying and managing EUC solutions will find this relatively straightforward.

Migration from an on-premises deployment of an existing vendors EUC solution

Teams need to be trained, and they should iteratively maintain their skills in AWS core competencies such as cloud delivery, compute, networking, and storage. Focus on understanding the core tenets of cloud delivery such as reducing costs, increasing scalability and resilience, and taking advantage of the global reach of AWS Cloud services.

Teams should still be trained on and exposed to Amazon WorkSpaces and AppStream 2.0, but technical teams with prior experience deploying and managing EUC solutions will find this relatively straightforward.

Pay particular attention on the training and preparation needed to accommodate the differences between the incumbent solution and the way AWS EUC services are deployed and managed. Differences in image lifecycle management, application delivery, user access and peripheral support will be key.

Migration from an existing cloud or hybrid deployment of EUC services

Technical teams with existing expertise deploying cloud solutions from other vendors will have transferrable skills that shortcut training requirements. While AWS Cloud and EUC service training will still be required, the time to absorb and apply this knowledge will require less time and effort.

Pay particular attention on the training and preparation needed to accommodate the differences between the incumbent cloud and EUC solutions and the way AWS Cloud and EUC services are deployed and managed.

While Amazon WorkSpaces and AppStream 2.0 deliver standard Windows desktops and applications, which are created, managed, and maintained in a similar way to many other EUC and VDI systems, there are a few specific differences that need to be considered:-

Amazon WorkSpaces and Amazon AppStream 2.0 service specifics

Amazon WorkSpaces and Amazon AppStream 2.0 are fully managed services, meaning that there is no customer access to the control plane. While this reduces control plane hardware requirements and simplifies deployment, there are some specific differences that need to be considered:

- **Connectivity**: User connectivity to each of the services is typically through an AWS-controlled point of presence on the public internet. Both streaming authentication and streaming traffic are delivered in this fashion. For Amazon AppStream 2.0, streaming traffic can be routed to a customer-configured VPC endpoint.
  - AppStream 2.0 Interface VPC Endpoints
- **Compute instances**: Amazon WorkSpaces and Amazon AppStream 2.0 instances are a specifically engineered version of equivalent EC2 instances. As a result, the storage and networking configuration is subtly different.
- **Instance availability**: Customers already familiar with the AWS Cloud and Amazon EC2 may be accustomed to a large selection of available instance types. While Amazon WorkSpaces and Amazon AppStream 2.0 offer a range of compute instances to deliver most typical EUC use cases, these are only a subset of the instance types available on EC2.
- **Cost management**: Minimizing cost is a key consideration for most customers when adopting AWS EUC services. All personnel involved in deploying, managing, and maintaining the environment need to adopt a mindset that active resources add to the solution costs. For example, optimizing the running mode of WorkSpaces (Always-On or AutoStop), and managing the scale up and down policies and running mode for AppStream 2.0 (Always-On or On-Demand) verifies that you are managing costs effectively.

Both WorkSpaces and AppStream 2.0 have cost optimizers that can be used to reduce costs by automating the shutdown or running mode of compute resources:

- Cost Optimizer for Amazon WorkSpaces

- Cost Optimizer for Amazon AppStream 2.0
- Cost Optimization for AppStream 2.0 Fleets

Amazon WorkSpaces and AppStream 2.0 targeted training

While a basic knowledge of AWS services, such as deploying VPCs, subnets, networking, and storage, is required to deploy AWS EUC services, the following training, specific to AWS EUC services, is also available:

- An Introduction to AWS End User Computing
- Amazon WorkSpaces Primer
- Amazon WorkSpaces Deep Dive
- Amazon AppStream Primer

EUCOPS10-BP01 Encourage user participation during service development and rollout to maximize engagement and project success

Encourage users to participate in online and in-person training for any new service to promote trust in the new service, increase employee engagement, and reduce support overheads.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Following the design and initial implementation of a new AWS EUC service, implement a structured plan to train and prepare users impacted by the introduction of the new service prior to production launch.

Provide timely and effective user training to avoid overwhelming support teams during rollout, identify usability issues, and promote employee engagement. If users are well-educated in the use of new systems, the chances of project success are enhanced.

Many approaches can be taken to provide users with the knowledge they need to adapt to a new environment, including:

- **Rollout communication**: As the date for production rollout becomes imminent, keep users up to date with plans and changes to build trust. Use collaboration tools such as Microsoft

Teams, Cisco Webex, Zoom, or a series of lunch and learn activities, for example, to communicate planned changes to your user population.

- **Key users**: After a solution is designed and initially deployed, initiate a pilot production phase. During this period, provide access limited key user access to the new systems to test the new services, verify that desktops and applications are delivering the expected functionality and performance, and check that peripheral devices and user data access are working as expected.
- You can also use your key users when a full production rollout is underway to assist with deskside or departmental support, as they will already be well versed in your new AWS EUC services.
- **Face to face**: If there will be significant changes to the way desktop and applications are delivered, or if the functionality of familiar applications will change due to upgrades or replacement, plan formalized face to face training for groups of users to provide them the opportunity to use the new services in a controlled environment. Verify that experience trainers or key users are in attendance and provide space to ask relevant questions.
- **Floor walking**: For office-based employees, provide floor walkers for the first few days of a deployment who can react and respond to user questions. This process is a great way to build trust and engagement with the user population.
- **Online or web-based**: Develop web-based training materials that allow users to consume training at a more convenient time if your user working patterns are unpredictable. Online training is also a good way to augment face to face training, as users can reinforce their skills and return to the courses when needed. Issuing certificates and small incentives to complete training courses is also a good way of building employee engagement and confidence in a new desktop and application deployment.
- **Frequently asked questions (FAQ)**: An FAQ is a well-established and highly successful way of allowing users to help themselves, reducing support calls and wait times for users with more complex issues. Gather a comprehensive list of the most common support questions at the pilot phase as key users start to use the new systems and add to it as you identify and resolve new common issues. Deliver the FAQ as a web page to allow the information to evolve quickly and be immediately available for consumption.
- **Chatbot:** With the evolution of AI/ML and generative AI, delivering an interactive user support capability may be possible for larger deployments. There are many services like Slack, Microsoft Teams, or Amazon Q that can accelerate the delivery of online assistance.

People consume and retain information in different ways. Offering them varied and complementary ways to build their knowledge of new and improved desktop and application delivery services contributes to a more engaged workforce and a successful deployment.

Operate

Observability helps you focus on meaningful data and understand your workload's interactions and output. By concentrating on essential insights and removing unnecessary data, you maintain a straightforward approach to understanding workload performance. It is essential not only to collect data but also to interpret it correctly. Define clear baselines, set appropriate alert thresholds, and actively monitor for any deviations. A shift in a key metric, especially when correlated with other data, can pinpoint specific problem areas.

EUCOPS11: What are the health metrics you need to monitor for your EUC environment?

Agreeing on the metrics which are key to identifying service health help to verify adherence to service level agreements which have been agreed upon by the business, third parties or customers.

EUCOPS12: How do you identify and deal with variations in service availability and performance that exceed agreed baselines?

Following the identification of service degradation, quickly identifying the root cause and best method of remediation will be critical. Different user personas and application sets will likely have different priorities that must be dealt with based on their impact on the business.

EUCOPS13: Do you have a mechanism to regularly review key metrics?

To identify trends that can be used to improve service delivery and maintain a consistent user experience, establish an agreed cadence for reviewing data that identifies how your AWS EUC services are performing over time.

EUCOPS14: Are you gathering sufficient logging data to identify cross-platform or service failures that may have a common root cause?

While AWS EUC-focused logs can expose issues specific to these services, access to data from other contingent services (like networking, authentication, storage, or backend services) provides a broader view of the cause and effect of cross-service outages.

Best practices

- EUCOPS11-BP01 Create EUC health metrics that allow you to meet your operational goals
- EUCOPS12-BP01 Deploy alerting mechanisms that quickly identify anomalous metrics
- EUCOPS12-BP02 Define and maintain an alerting chain of command that quickly communicates issues in real time
- EUCOPS13-BP01 Perform regular service reviews to identify significant trends in performance, scalability, and availability
- EUCOPS14-BP01 Ingest log file data from multiple data sources to correlate key problem identifiers and trends

EUCOPS11-BP01 Create EUC health metrics that allow you to meet your operational goals

Spend time reviewing the available metrics which provide quick and insightful information into the health of your end-to-end EUC deployment.

Level of risk exposed if this best practice is not established: High

Implementation Guidance

While the tools and processes required to monitor AWS EUC service health are discussed earlier, from an operational perspective there are key metrics which, at a minimum, should be gathered to build a baseline for systems health across the tiers of an Amazon WorkSpaces or AppStream 2.0 deployment.

The following guidance discusses both the service specific metrics which should be gathered in addition to the monitoring other key services which contribute to AWS EUC service:

Amazon WorkSpaces and AppStream 2.0 Service or Instance metrics

Insight into both service level and instance-based performance metrics are key to identifying availability problems, performance problems or trends and to provide data for retrospective problem analysis. Consider gathering the following data, at a minimum, in order to maximize service efficiency and performance:

Amazon WorkSpaces: Amazon CloudWatch provides an automatic dashboard which gives an overview of overall service health for Amazon WorkSpaces, including:

Service metrics:

- Available or unhealthy WorkSpaces
- Session launch times
- Connection success or failure
- Session latency
- Users connected, disconnected, stopped, or in maintenance

Instance metrics:

- In-session latency
- Network health
- CPU usage
- Memory usage
- Root or user volume space usage

Custom dashboards can also be created which use these metrics to focus on a specific subset of your WorkSpaces.

- Monitor WorkSpaces Personal
- Monitor your WorkSpaces using CloudWatch metrics
- Creating Custom CloudWatch dashboards for Amazon WorkSpaces

CloudWatch Alarms can also be configured to send alerts when specific thresholds are met. For more information, see Creating CloudWatch Alarms.

Amazon AppStream 2.0: Amazon CloudWatch provides an automatic dashboard which gives an overview of overall Amazon AppStream 2.0 service health, including:

Service metrics:

- Fleet capacity or utilization
- Insufficient capacity errors
- Average actual capacity
- Average available capacity
- Average desired capacity
- Average in use capacity
- Average pending capacity

For multi-session AppStream deployments, additional performance metrics can be viewed for each instance or session, these metrics will also be available for single session fleets over time.

Instance metrics:

- Instance CPU utilization
- Instance memory utilization
- `PagingFileUtilizationInstance`
- Instance disk utilization

Session metrics:

- Session CPU utilization
- Session memory utilization

For more information, see Viewing Instance and Session Performance Metrics Using the Console.

CloudWatch Alarms can also be configured to send alarms when specific thresholds are met.

- Using Amazon CloudWatch alarms
- Monitoring Amazon AppStream 2.0 Resources

Other key areas to monitor

The following services, and associated metrics, at a minimum, should be monitored in order to understand the end to end health and performance of AWS EUC services.

Networking

With any cloud hosted desktop and application delivery service such as Amazon WorkSpaces or Amazon AppStream 2.0, users are connecting from a remote location, across a variety of network types, to a service running in a cloud data center. Once they are connected and logged in, they are dependent upon a number of backend services which are also connected to the AWS EUC service using a variety of devices which each have their own performance characteristics. Each part of the connection process and subsequent interaction with backend services should ideally, be monitored.

User endpoint device to AWS EUC service

The following articles discuss the latency and bandwidth requirements for Amazon WorkSpaces and Amazon AppStream 2.0 and tools that can be used to validate service performance:

- Client network requirements for WorkSpaces Personal
- AppStream Latency: Bandwidth Recommendations
- Measuring Client to AWS EUC region latency
- Visualizing AppStream 2.0 session latency metrics using AWS Lambda, Amazon Kinesis Data Stream and Amazon OpenSearch Service
- CloudWatch Internet Monitor
- Utilizing CloudWatch Internet Monitor with Amazon WorkSpaces Personal

AWS EUC compute instance to backend services:

Consider deploying third party tools which proactively monitor client to server operations such as network flow between WorkSpaces, AppStream 2.0 and supporting databases, data feeds, web servers and file or print services. These data points can be used to accurately determine service degradation or trends which might identify the need to scale supporting infrastructure service up or down.

AWS EUC compute instance to externally hosted services:

While there are no simple ways to individually gather the performance of compute instance to external service metrics, many third-party cloud providers provide API's which can be leveraged to determine service status. Both Microsoft and Google for example, expose API's that can be used to query individual cloud service availability. It should be possible to architect a centrally hosted

solution which pools key external resources and uses the metrics gathered to align with internal service availability

Backend service availability:

Consider using network analysis tools which can identify the reachability of key services using ICMP, TCP or application layer health probes. For Amazon WorkSpaces and Amazon AppStream which are dependent on low latency and available bandwidth, built in client-side network health tools will identify and notify the end user of performance degradation. In general, the ability to identify performance baselines for network packet flow is crucial. This applies to various supporting network infrastructures, including AWS-specific connections through Direct Connect, as well as connections to and from third-party cloud infrastructures.

Storage:

As discussed previously, both Amazon WorkSpaces and AppStream provide metrics which can trigger alarms based on certain thresholds such as if disk space is running low, but these do not include storage performance metrics. As part of your scalability testing during adoption of AWS EUC services, consider testing disk performance if your application workload is particularly disk i/o intensive. Some WorkSpaces and AppStream instance types are 'EBS Optimized' offering scalable disk throughput, for both services, GPU enabled instances offer the highest throughput and additional instance storage. The DISKSPD utility from Microsoft can be used to create synthetic disk i/o profiles for testing purposes.

- Amazon EBS Optimized Volumes and Instance Types
- Instance store temporary block storage for EC2 instances
- Microsoft: Use DISKSPD to test workload storage performance

If specific issues arise that require deeper insight into Amazon WorkSpaces or Amazon AppStream 2.0 storage performance, consider using Windows Task Manager or Performance Monitor, or iostat/iotop for Linux instances, to better understand disk i/o performance.

Active Directory

Active Directory performance is key to the user experience of Amazon WorkSpaces and AppStream 2.0 users as it directly affects the logon process. A badly performing Active Directory infrastructure may add significant logon time as Group Policies and Logon Scripts are processed. If AWS Managed Microsoft AD is being used, CloudWatch can be used to provide insights into directory performance. For EC2 hosted domain controllers, CloudWatch can also be used to gather most of

the metrics required to identify service degradation. For on-premises AD controllers, CloudWatch agents can be installed to centralize the collection of appropriate metrics, such as CPU, Memory, disk I/O and network utilization.

For more information, see Performance tuning for Active Directory Services.

SAML 2.0

SAML integration with AWS EUC services is typically provided by external providers such as Azure AD, Okta or Ping Identity. These systems usually provide an API which can be used to extract service level heath metrics for propagation into an existing SIEM system. Azure Monitor or the Okta System Log API, for example, can be used to understand availability and performance.

Certificate-based authentication (CBA)

If end-to-end single sign-on is required for Amazon WorkSpaces or AppStream 2.0 deployments which are integrated with SAML, CBA can be used to emulate a virtual smart card login for each user. While falling back to a standard AD username and password login is possible if CBA is unavailable, if you do not elect to use this option it will be essential to implement monitoring for CBA to avoid login failures. The AWS Private Certificate Authority is a resilient service by default and presents operational metrics through CloudWatch:

- Monitor AWS Private CA with CloudWatch Events

As certificate-based authentication relies upon a private certificate authority (PCA) which in turn requires a Microsoft Certificate Service infrastructure, refer to the following documentation to understand which key metrics should be monitored:

- Microsoft: Securing PKI: Appendix A: Events to Monitor

Network file services

Amazon AppStream 2.0 and WorkSpaces are typically integrated with backend network file services which provide storage for user data and user profiles. These repositories are typically critical to employee productivity and should form part of end-to-end service monitoring. If Amazon FSx for Windows is being used for backend storage, a comprehensive CloudWatch dashboard is available which exposes system performance. If traditional Windows file servers are being used in EC2 or on-premises, Microsoft provides direction on how to use the SMB performance metrics to gather the relevant performance statistics.

- Monitoring with Amazon CloudWatch
- Performance tuning for SMB file servers

RADIUS

If RADIUS is being used with Amazon WorkSpaces, the documentation for the RADIUS provider in use should be consulted as these can be Windows or Linux based and will expose performance metrics in different ways.

Application web tiers

Availability and performance of web tiers that support the applications being delivered from AWS EUC services is typically controlled by load balancers than can also execute L2, L4 or L7 health probes to ascertain service health and optionally perform auto-scaling if required. Refer to your web server vendors documentation for information on monitoring your specific web tiers.

Application database tiers

Availability and performance of database tiers that support the applications being delivered from AWS EUC services is also key to end-to-end service health. Refer to your web server vendors documentation for information on monitoring your specific web tiers.

Application licensing servers

Monitoring license server availability and performance is critical as failure of these servers can result in complete denial of service for a specific application tier. Please refer to your license server or application vendors documentation for information on monitoring these components.

Web proxies or app firewalls

Web proxy and app firewall tiers are typically load balanced and auto scaled for resilience and scalability, but monitoring these is important as failure of this tier can result in users being denied Internet access, the impact of which can be significant. Please refer to your web proxy vendor documentation for information on monitoring these components.

Anti-virus infrastructure

While anti-virus and anti-malware products are unlikely to cause systems outage, from a security perspective, being sure that Amazon WorkSpaces and AppStream 2.0 instances are being

effectively protected can avoid wider service outage due to intrusion and malign interference from external bad actors. Furthermore, understanding and minimizing the impact of anti-virus and anti-malware scans, is key.

WorkSpaces and AppStream 2.0 instance metrics

Amazon WorkSpaces and AppStream 2.0 compute instances are standard Windows Client/Server, or Linux instance types. They each have a network interface exposed to a customer managed VPC and can be managed and monitored in the same way as traditional desktops.

Amazon CloudWatch can be used to extract instance specific metrics such as CPU, Memory, Disk or Network utilization, and existing third party tools can be used to extract similar information.

Be aware that as AppStream 2.0 is a non-persistent application and desktop delivery service, instances are terminated and destroyed when the last user session is ended (consider single session versus multi-session), this needs to be considered when gathering performance statistics or system logs.

There are a number of utilities created by AWS employees that assist in the gathering and presentation of AWS EUC instance metrics, the EUC Toolkit for example, can be used for this purpose, the PowerShell code for this utility can also be downloaded and used as a reference for building your own PowerShell management utilities.

- Monitor your WorkSpaces using CloudWatch metrics
- Monitoring and Reporting for Amazon AppStream 2.0
- Monitoring Amazon WorkSpaces Secure Browser
- Use the EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces

In summary, AWS EUC deployments are dependent on the reliability and performance of both the Amazon WorkSpaces or AppStream 2.0 services themselves and also many external systems, taking a holistic approach to management of each component of the end to end deployment is key to maintaining end user engagement and productivity.

EUCOPS12-BP01 Deploy alerting mechanisms that quickly identify anomalous metrics

AWS EUC services provide access to desktops and applications which can be highly variable in their resource requirements over time. Weekly, monthly, quarterly, and year-end activities can

cause spikes in resource consumption that might result in unnecessary alerts and a degraded user experience.

Level of risk exposed if this best practice is not established: High

Implementation guidance

The design and pilot phases of an AWS EUC project should identify resource requirements for each application set over a typical business cycle. Identify the peak activity levels to verify that the compute instance types selected for both Amazon WorkSpaces and AppStream 2.0 can deliver performance that maintains a good user experience and improves productivity.

Third party tools from vendors such as ControlUp, Nuvens, LiquidWare, Lakeside Software, and Aternity can be used to collect resource usage trends and build baselines for key applications. Some of these can be found on the AWS Marketplace.

AWS and the AWS Partner Network offer many services and automation capabilities you can use to automatically and elastically scale backend application services or to provide increased compute capabilities during periods of heavy utilization**.**

EUCOPS12-BP02 Define and maintain an alerting chain of command that quickly communicates issues in real time

As important as gathering relevant service metrics and alerts is expediting the propagation of those alerts to the right teams, individuals, or automated processes. This propagation helps you quickly surface and remediate associated issues.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Accelerate the awareness of key events, and check that the notification to initiate the appropriate process of remediation is quickly followed.

There are several ways to verify that both operations and support teams in addition to internal and external users and management teams are appraised of service health:

- **Health dashboards**: Build a set of centralized service health dashboards that are tailored to provide the right information to the right people from operations, support, users, or

management. Dashboards help your teams quickly identify and track issues to resolution. User-level dashboards promote transparency, reduce support calls, and increase user engagement as new production services are introduced.

- **Effective communication**: Develop a communications protocol to effectively communicate about extended outages as they are identified to internal and external customers. Keeping customers informed, specifically around outage timelines, is key to building trust and engagement.
- **Effective routing**: Automate the process of prioritizing and effectively routing the right alerts to the right teams at the right time, which increases operational efficiency and contributes to an improved user experience and higher productivity.
- Consider the following factors when identifying roles and responsibilities for event response, escalation, and propagation:
  - **Roles and responsibilities**: Define clear lines of responsibility for escalation, problem resolution, and root cause analysis.
  - **Incident assignment**: Identify alert categories so that specific events can be directed to the team most appropriate to resolve. For example, first, second, and third lines of support.
  - **RPO and RTO requirements**: Involve the business in understanding and calculating RTO and RPO requirements to prioritize problem tracking and remediation accordingly.
  - **Cost of outages**: Quantify the cost to the business of specific categories of outage and use this data to inform the escalation and notification process. It may be pertinent to revise support processes to involve more skilled support teams to react to specific event types that have higher business impact.
  - **Tool**s: Map out a matrix that identifies the right tools, metrics, and notification processes to verify that critical events are surfaced and distributed effectively to the appropriate teams and individuals.
  - **Alert fatigue**: Filter out duplicate alerts and false positives, as they can lead to alert fatigue and loss of focus on important issues.
  - **Geographic reporting**: For multi-Region deployments, dynamically adjust notification distribution lists to accommodate support in applicable time zones and geographic areas.

EUCOPS13-BP01 Perform regular service reviews to identify significant trends in performance, scalability, and availability

Perform regular reviews of service performance and capabilities to maintain visibility of key issues and focus on service improvement and readiness.

Level of risk exposed if this best practice is not established: High

Implementation guidance

While real-time monitoring and alerting is essential in meeting business and technical SLAs with internal and external customers, performing periodic review of logfile and monitoring data can help to identify problem trends and to put in place remediation steps to avoid future outages.

Along with incumbent monitoring tools, you can use Amazon CloudWatch and Amazon Kinesis to centrally store data to use for retrospective performance and systems health analysis.

EUCOPS14-BP01 Ingest log file data from multiple data sources to correlate key problem identifiers and trends

Identify and implement mechanisms to maintain a centralized source of EUC service data that can be used for root cause analysis of cross service issues.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Store logs and metrics from AWS EUC services and their dependent services in a centralized location to allow analysis tools to build a picture of cross system failures that are affecting AWS EUC reliability. For example, expiry of a critical SSL certificate on a load balancer or remote access tier may be the root cause of login degradation or other failures at the AWS EUC service tier.

Use Amazon CloudWatch to gather metrics and logs, which are stored for subsequent analysis, to identify problems or trends that have occurred over time.

Amazon Kinesis agents can be installed onto Amazon WorkSpaces or AppStream 2.0 images to export log file data in real time to a centralized location for retrospective analysis.

For larger environments, consider creating a data lake of key data from system logs, performance monitoring, and security tools across all service lines. Develop analysis capabilities using Amazon AI/ML tools to generate a more holistic insight into end to end systems health and scalability.

> **Note**
>
> Review CloudWatch and Kinesis data retention policies and service charges to verify that data availability and costs are within EUC project guidelines.

Evolve

End User Computing solutions evolve over time. New features are added or deprecated, end user requirements change, and business requirements evolve to maintain competitive advantage. This section discusses how change should be embraced and incorporated into the day to day activities which govern your EUC deployment.

EUCOPS15: How do you track and manage changes to the solution design?

Define internal processes that propagate any changes to AWS EUC services or supporting infrastructure into a version-controlled suite of operational documentation.

EUCOPS16: How do you control, test, and validate changes to the AWS EUC environment and deploy changes repeatably and reliably?

Select a toolset and processes that foster predictable delivery and maintenance of AWS EUC services to improve user satisfaction, reduce support overhead, and achieve positive business outcomes.

EUCOPS17: How do you make data-driven, iterative improvements across all aspects of your AWS EUC deployment?

Perform periodic review of service effectiveness based on analysis of archived log and performance data and support calls. Combine this data with training and exposure to product roadmaps and feature updates in order to plan service improvements that add incremental value to your AWS EUC service delivery.

Best practices

- EUCOPS15-BP01 Update your solution design documentation over time, and use version control to track changes
- EUCOPS16-BP01 Implement automated processes to verify that service updates can be repeatably deployed, updated, and rolled back

- EUCOPS17-BP01 Provide time and resources for your teams to keep up to date with changes and feature updates

EUCOPS15-BP01 Update your solution design documentation over time, and use version control to track changes

Keep key architectural designs, operations handbooks, and support guides up to date, maintaining a library of reference material that can be used by new personnel, partners, or other support teams.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

For both Amazon WorkSpaces and AppStream 2.0, each service should have been deployed based upon a design that resulted from the collected input of key technology and business stakeholders. Evolving the solution design should be managed in a similarly inclusive fashion. Agree and sign off on all changes to the initial design through a project board before updating the solution. This approach verifies that invested parties have validated the key metrics required to deliver the new service and that the updated solution meets the requirements of both technical and business stakeholders.

Design documentation should be maintained as continually updated documents that represent the state of the AWS EUC service deployments over time. It should capture the rationale for each design decision in addition to the technical and architectural solutions deployed to achieve each requirement. Maintain iterative versions of the design as changes are made so that you can see a historical view of the deployment.

A design document is an essential piece of knowledge collateral which is invaluable for training purposes, onboarding new technical team members, reviewing and implementing changes to the infrastructure, and when working with partners to integrate new technologies or handover support to new teams.

EUCOPS16-BP01 Implement automated processes to verify that service updates can be repeatably deployed, updated, and rolled back

Selecting an automation toolset and defining the processes that facilitate repeatable, predictable delivery and maintenance of AWS EUC services is key to achieving simplified administration, reduced support overheads, end user satisfaction and positive business outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Although Amazon WorkSpaces and Amazon AppStream 2.0 are fully managed services, there are a number of touch points when maintenance the associated infrastructure and the desktop and applications delivered by the service requires periodic updates.

Every AWS EUC environment is unique. Verify that you understand each area of your AWS EUC deployment that may need to be updated over time and develop a formalized plan on how each of these areas need to be managed.

The following questions and discussions can provide you steps for improvement.

What updates are required?

- **Amazon WorkSpaces**: For WorkSpaces, the custom bundles created to deliver persistent desktops to users will require updates over time in the form of operating system patches, hotfixes, and security and application updates. Once your WorkSpaces have been deployed, they must be individually managed, as each WorkSpace can be uniquely changed and reconfigured by its assigned user if they are given the appropriate rights. The customer is responsible for making these changes. Amazon WorkSpaces have a regular automatic maintenance schedule which keeps the WorkSpaces specific agents aligned with the service control plane. For detail on the maintenance process for Always-On and AutoStop WorkSpaces, see Maintenance in WorkSpaces Personal.
- **Amazon AppStream 2.0**: For AppStream 2.0, each private image used to deploy a non-persistent desktop or application experience will periodically require updates in the form of operating system patches, hotfixes, and security and application updates. As AppStream 2.0 instances are deployed from a common image, only the private image for each fleet version needs to be updated. New instances launched when users log in will automatically inherit the changes made to the private image. The customer is responsible for making these changes.

- Maintenance of the agent software installed on each image can be automated or controlled by the customer if specific versions are required. For more information on the processes of maintaining agent versions for each image, see:
  - Update Management in Amazon AppStream 2.0 PDF RSS Focus mode
  - Manage AppStream 2.0 Agent Versions
- Amazon AppStream 2.0 also offers an application delivery option called elastic fleets that you can use to quickly deploy and manage portable applications. For more information, see Applications Manager.

How do you manage updates?

Creating and delivering updates that are consistent and repeatable is the best way of reducing problems and frustration for users when configuration changes are made to the workloads delivered by AWS EUC services. You can use software deployment tools to build new software packages, perform unit and interoperability testing, and roll out or roll back changes without touching each desktop or application server individually. This form of automation drastically reduces the chance of human error and configuration drift across large desktop and application estates, saving on support costs, reducing downtime, and maximizing productivity.

- **WorkSpaces**: Workspaces provides a management console and a corresponding API, which can be used to create and configure new WorkSpace bundles. Once created from a custom bundle, each WorkSpace is persistent but decoupled from the custom image and requires discrete management and maintenance.
- To update existing WorkSpaces, use the customer-facing network interface attached to each WorkSpace to integrate with software deployment toolsets such as AWS Systems Manager or existing on-premises tools such as Microsoft Endpoint Configuration Manager (MECM), Puppet Enterprise, or Ansible.
  - Software deployment to Amazon WorkSpaces using AWS Systems Manager
  - Automatically create customized Amazon WorkSpaces Windows images
- **AppStream 2.0**: AppStream 2.0 provides a management console and a corresponding API, which can be used to automate the delivery of an image builder that updates each version of a private image. As the image builder has a network interface in a customer-managed VPC, traditional software distribution tools and automation frameworks can also be used to push updates to this instance from where a new version of an image is created and assigned to fleets.

- AppStream 2.0 also offers an automated option called Managed Image Updates, which automates and simplifies the process of updating AppStream agent software and OS patches. For more information, see the following:
  - Administer Your Amazon AppStream 2.0 Images
  - Automatically create customized AppStream 2.0 Windows images
  - Automate the creation of AppStream 2.0 resources using AWS CloudFormation

How will you test and validate updates?

- **WorkSpaces**: Before production rollout, any OS or application updates need to be tested on a WorkSpace created from the same custom bundle as the WorkSpace group being updated. Several custom bundles may exist with different application combinations that need to be independently tested. Once testing is complete, you can roll out changes to each WorkSpace created from the initial custom bundle, either manually or using automation tools such as Microsoft WSUS or Microsoft MECM (SCCM).
- If WorkSpace users have been given full administrative access to their desktop, it is possible that they may have updated their WorkSpace OS or application independently, making the process of applying consistent, reliable updates across the WorkSpace estate challenging. Unless strictly necessary, we don't recommend allowing users to update their own WorkSpaces.
- Should an update fail, a snapshot of the two WorkSpaces storage volumes is taken every 12 hours, which may provide a recovery position. WorkSpaces can be rebuilt or recovered. For more information, see Rebuild a WorkSpace in WorkSpaces Personal.

For more flexible backup and recovery options, consider using traditional backup and recovery tools and techniques, or consider AWS Backup.

- **AppStream 2.0**: As AppStream 2.0 delivers tens, hundreds, or thousands of instances from a common private image, testing can be done by creating a single instance test or development fleet from a new version of an image, allowing administrators to fully test changes before assigning the image to a production fleet and making it available to users.
- Multiple fleets or fleet versions can be created from a private image, allowing administrators to roll forward or roll back to a known operational state if problems occur. Multiple versions of an image can also be deployed in parallel if extended user testing is preferred.

How do you track changes and manage releases?

Track specific changes to your AWS EUC environments by date and time to maintain a paper trail that can be used to pinpoint and cross reference if a specific change was responsible for a positive or negative change in functionality or performance. For example, creating a retrospective back-out or remediation plan in the event of an issue that occurs days or weeks after a change is made to the environment will be much easier if comprehensive change management is observed.

For both Amazon WorkSpaces and Amazon AppStream 2.0 specifically, adopt a version numbering scheme and capturing a log of changes made to each custom bundle or private image to trace issues back to a specific image version, if required.

You can use AWS CloudTrail to log API calls used to make changes to both Amazon WorkSpaces and AppStream 2.0 environments.

- Logging AppStream 2.0 API calls with AWS CloudTrail
- Logging WorkSpaces API Calls by Using CloudTrail

Automating changes to Amazon WorkSpaces and Amazon AppStream 2.0

By using automation, you can avoid many of the configuration drift and image consistency problems seen with manual deployments. The following articles provide some options for automating the creation and management of AWS EUC services.

- Automating the provisioning of AWS WorkSpaces
- Automatically create customized AppStream 2.0 Windows images
- Best practices for automating your AWS End User Computing deployments
- Amazon WorkSpaces and AppStream 2.0 Terraform Resources
- Deploying and Managing Amazon WorkSpaces applications with Ansible
- DXC Technology creates DevSecOps and CI/CD for mainframe and Java using Amazon AppStream 2.0
- Announcing the Amazon WorkSpaces dynamic inventory plugin for Ansible®
- Terraform resources for AWS WorkSpaces
- Automation of infrastructure and application deployment for Amazon AppStream 2.0 with Terraform

EUCOPS17-BP01 Provide time and resources for your teams to keep up to date with changes and feature updates

Provide ongoing training to keep key personnel up to date with changes that are occurring in the industry and specifically in their domain of expertise.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Take advantage of new and improved capabilities of Amazon WorkSpaces and AppStream 2.0 services and deploy new updates to deliver incremental features and performance for your desktops and applications delivered by those services. Staying up to date is key to deliver business outcomes that provide a competitive advantage.

Perform periodic reviews of new service capabilities and improvements in desktop and application delivery to maximize your investment.

You can continually improve service for Amazon WorkSpaces and AppStream 2.0 in many ways, including:

- **New features**: Identify, test, and implement new service features that provide added value in order to deliver incremental end user and customer benefits, sometimes at zero cost. Follow the AWS Desktop and Application Streaming Blog to stay updated on the latest developments in this area.

  Patching: Regularly apply operating system patches, hotfixes, and important updates to both the operating system and applications to avoid known issues and vulnerabilities, which helps you avoid costly service outages.
- **Performance**: Periodically review performance logs to scale your environment cost effectively and maintain a compelling user or customer experience.
- **Service review**: Periodically review all support tickets related to the AWS EUC deployment to understand root cause and identify problem trends, helping you avoid future outages and associated costs.
- **Change management**: Involve your AWS EUC team in the change board for all dependent technology areas (such as compute, storage, networking, and security). Provide visibility of changes in other technology domains to guide and inform improvements in the delivery of desktops and applications.

- **Industry awareness**: Attend key EUC industry events to identify new industry trends and partners who provide added value. The opportunity to attend industry events is also an opportunity to meet other users of AWS EUC services and learn from their valuable experiences.
- **Expert roundtables**: Promote the use of and participation in regular expert roundtables where technology teams can present improvements and advances across diverse areas of expertise. This helps the EUC team identify where they can apply improvements in other areas to improve AWS EUC service delivery.

Key AWS services

EUC:

- Amazon AppStream 2.0
- Amazon WorkSpaces Family
- Amazon WorkSpaces Secure Browser

Storage:

- Amazon FSx
- Amazon S3
- Amazon Elastic Block Store

Managed directories for WorkSpaces:

- AD Connector
- AWS Managed Microsoft AD
- Simple AD

Resources

- AWS Compliance
- Amazon Web Services: Risk and Compliance
- AWS Cloud Security
- Engage with AWS Partners

- AWS Professional Services
- AWS Training and Certification
- AWS Certification
- Amazon WorkSpaces – Getting Started
- Best Practices for Deploying WorkSpaces
- Amazon AppStream 2.0 – Getting Started
- Best Practices for Deploying Amazon AppStream 2.0
- Using Kinesis Agent for Microsoft Windows to store AppStream 2.0 Windows Event Logs
- AWS Health Dashboard
- Best Practices to Automate your Amazon End User Computing Deployments
- Manage directories for WorkSpaces Personal
- Amazon WorkSpaces: Multi-Region Resilience for WorkSpaces Personal
- Enable and Administer Home Folders for Your AppStream 2.0 Users
- How Application Settings Persistence Works
- Amazon AppStream 2.0 Now Supports Copying Images Across AWS Regions
- The AWS Fault Injection Service
- Amazon WorkSpaces Service Level Agreement
- Amazon AppStream 2.0 Service Level Agreement
- AppStream 2.0 Interface VPC Endpoints
- Cost Optimizer for Amazon WorkSpaces
- Cost Optimizer for Amazon AppStream 2.0
- Optimizing costs using Amazon AppStream 2.0 fleet options
- Introduction to AWS End User Computing
- Amazon WorkSpaces Primer
- Amazon WorkSpaces Deep Dive
- Amazon AppStream 2.0 Primer
- Monitor WorkSpaces Personal
- Creating custom Amazon CloudWatch dashboards and widgets for Amazon WorkSpaces
- Using Amazon CloudWatch alarms
- Client network requirements for WorkSpaces Personal

- Amazon AppStream 2.0: Bandwidth requirements
- Measuring Client to AWS EUC region latency
- Visualizing AppStream 2.0 session latency
- Amazon CloudWatch: Using Internet Monitor
- Amazon EBS-optimized instance types
- Instance store temporary block storage for EC2 instances
- Microsoft: Use DISKSPD to test workload storage performance
- Performance Tuning for Active Directory Services
- Monitor AWS Private CA with CloudWatch Events
- Securing PKI: Appendix A: Events to Monitor
- Amazon FSx for Windows File Server: Monitoring with Amazon CloudWatch
- Performance tuning for SMB file servers
- Monitor your WorkSpaces using CloudWatch metrics
- Amazon AppStream 2.0: Monitoring and Reporting
- Monitoring Amazon WorkSpaces Secure Browser
- Use the EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces
- Ultimate KPI Tree Guide: How to Build Killer KPI Trees in 7 Steps
- AWS Backup
- AWS Storage Gateway
- Maintenance in WorkSpaces Personal
- Update Management in Amazon AppStream 2.0
- Manage AppStream 2.0 Agent Versions
- AppStream 2.0 App Blocks
- Software deployment to Amazon WorkSpaces using AWS Systems Manager
- Update Management in Amazon AppStream 2.0
- Rebuild a WorkSpace in WorkSpaces Personal
- Restore a WorkSpace in WorkSpaces Personal
- Logging AppStream 2.0 API calls with AWS CloudTrail
- Logging WorkSpaces API calls with AWS CloudTrail
- Automate provisioning of Amazon WorkSpaces using AWS Lambda

- Automating the creation of AppStream 2.0 Windows Images
- Best practices for automating your AWS End User Computing deployments
- Deploying and Managing Amazon WorkSpaces applications with Ansible
- DXC Technology creates DevSecOps and CI/CD for mainframe and Java using Amazon AppStream 2.0
- AWS Desktop and Application Streaming Blog
- AWS Sustainability
- AWS Training and Certification

Security

Using managed AWS End User Computing (EUC) services helps reduce the scope of your infrastructure management tasks, letting you focus on mitigating common security risks with less overhead.

The security pillar includes the ability to protect information, systems, and assets while delivering business value. This section provides in-depth, best-practice guidance for architecting secure EUC workloads on AWS.

Focus areas

Design principles

There are several principles that can help strengthen the security of EUC workload in addition to the overall Well-Architected Framework security design principles:

- **Implement lifecycle management for AWS End User Computing instances and applications delivered by End User Computing services:** End to end lifecycle management should be adopted for the software used to deliver applications to your end users. Lifecycle management includes considering the operating system, middleware, runtime environments and patches associated with this software to make sure that components are patched and updated to the most recent release. It also includes the ongoing management of user identities used

to access End User Computing services to make sure that accounts are regularly validated and where used, that certificates are validated and renewed.

- **Design EUC solutions that respect data classification and restrict access to data:** EUC solutions should respect the classification of data and restrict the ability for users to access data with classifications that they are not entitled to access.
- **Design for continuous monitoring of end user sessions:** Implement continuous monitoring of user sessions to determine the user experience that users are receiving while using AWS End User Computing services. Make sure that key operating system performance metrics covering processor, memory, disk and network are monitored. Review the performance of user sessions on a regular basis to determine if any anomalous metric data is being generated due to systems being compromised.
- **Limit access to AWS EUC services to approved and adherent devices:** Access from unknown or unsupported device systems should be restricted to the absolute minimum to allow users to connect to and use the service with the required level of functionality for their role. Adherent device access should be used, when possible, where devices are assessed for compliance against a set of criteria before being allowed to access the requested service.

Security foundations

EUCSEC01: How do you identify and categorize risk profiles for users accessing your EUC services?

The organization implementing EUC services may have different users with different risk profiles associated with them. Use the following best practices to achieve the appropriate security posture for the users accessing the services.

EUCSEC02: How do you maintain security and compliance for users accessing your EUC environment?

The organization implementing EUC services may have external regulatory frameworks, standards, and internal policies that are in scope for the EUC environment. When implementing the

configuration of EUC services, it is important that these are fully considered for all types of users and stakeholders accessing your environment.

EUCSEC03: How do you implement least privilege access to applications?

The principle of least privilege maintains that a user or entity should only have access to the specific data, resources and applications needed to complete a required task

Best practices

- EUCSEC01-BP01 Identify discrete groups of users that require access and implement security controls appropriate for their risk profiles
- EUCSEC02-BP01 Identify external stakeholders and their security or regulatory compliance requirements
- EUCSEC03-BP01 Restrict user permissions to the minimum required to perform their role

EUCSEC01-BP01 Identify discrete groups of users that require access and implement security controls appropriate for their risk profiles

When modelling user access to computing systems, it is important to consider different risk profiles associated with discrete groups of users. For example, internal employees and external contractors will have different risk profiles associated with them. Because of their risk profile, different security controls should be applied to the groups of users.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Create a group to manage users associated with each risk profile. If different discrete sets of users will be interacting with the AWS EUC services, take a risk-based approach to determine the risk profile associated with each group. The groups being considered here are broader than other AWS services, as you need to consider users across multiple lines of business each with their own discrete risk profiles in addition to the standard administrators, developers, and operators.

Based on the risk profile, implement different security controls to mitigate residual risks within the groups of users. A matrix can be used to assess the risks associated with users. For example, in a scenario where four groups of internal and external users will be accessing the EUC services,

a 2x2 matrix can be created that captures the type of users on one axis (for example, internal or external) and the risk profile of the group of users on the other (for example, high or low risk). By populating the matrix with the different groups, you can determine the appropriate risk posture and apply the appropriate level of security controls for the user group, such as enforcing multi-factor authentication. An example matrix is shown in the following figure for groups of internal and external users that will access a computing service.

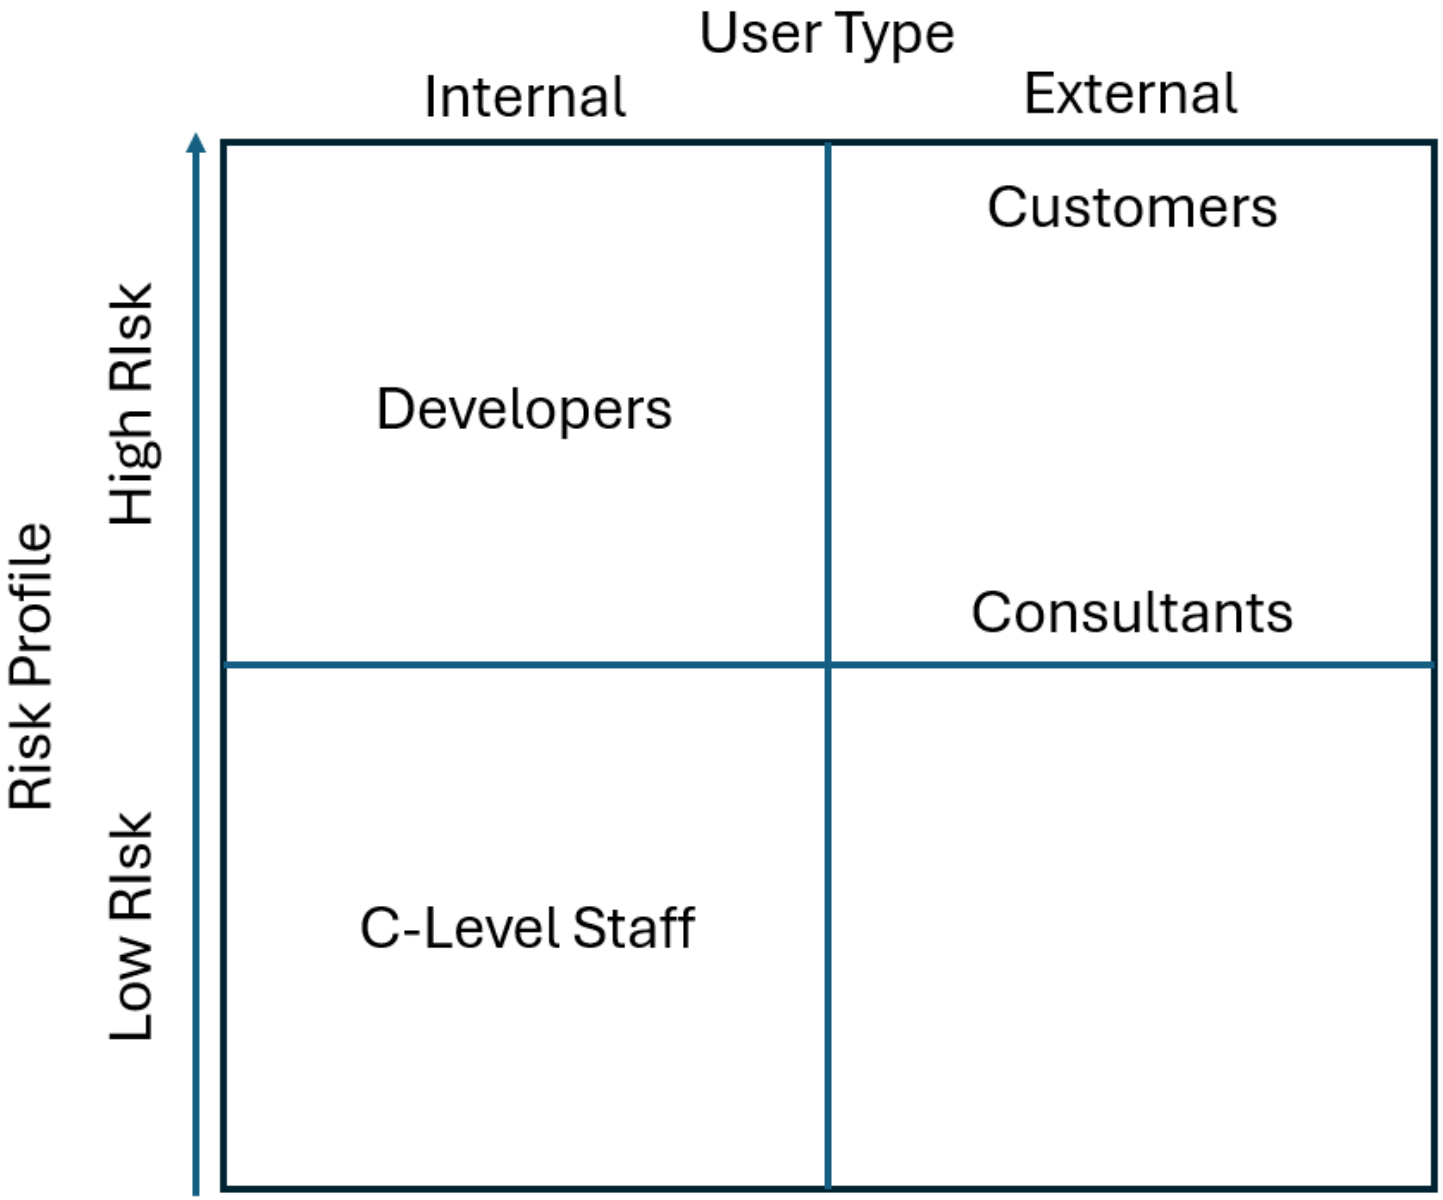

EUCSEC02-BP01 Identify external stakeholders and their security or regulatory compliance requirements

When creating and configuring an end user computing environment, verify that the regulatory requirements for the users of your environment are met. Consider the broader regulatory frameworks and their associated requirements in relation to accessibility that may be in scope for users with specific accessibility requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Determine internal and external policies that are applicable to your environment. To help identify stakeholders external to the organization, consider the following groups of potential sources of policy:

- Government
- Legal (for example, employment law, health and safety regulation, financial regulation, or accessibility)
- Industry (for example, financial services regulators)

By considering each of these groups, you can assess the different potential sources of regulatory compliance for relevance against the applications being delivered, as well as the data they process and visualize.

EUCSEC03-BP01 Restrict user permissions to the minimum required to perform their role

To implement the principle of least privilege when configuring AWS EUC services, define appropriate access controls for role categories like users, service desk users, first-level administrators, second-level administrators, and accounts used for automation.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Limit the use of administrator permissions:** Users should not be granted local administrator access to Amazon WorkSpaces or AppStream 2.0 instances unless it is required for them to

undertake their role. Use tools and products that provide the ability to temporarily provide elevated rights in preference to granting users long term administrative access.

- **Do not provide all support staff with administrator permissions:** Grant service desk users should the minimal set of access permissions to allow them to perform their function. This can vary among organizations, but service desk users should not be granted full access to the Amazon WorkSpaces and AppStream 2.0 services.
- **Use administrative toolsets and automation to avoid the need to provide administrator permissions:** Administrators providing first-level support for the users consuming the AWS EUC service can use the enhanced administrative toolset that AWS offers in the form of the EUC toolkit. For more detail on the EUC Toolkit, see Use EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces.
- **Audit and monitor privileged or sensitive operations:** Log any privileged or sensitive operations associated with the management of AWS EUC services. These logs can then be used to generate alerts as required.
- **Use temporary elevated access for privileged or sensitive operations:** When users occasionally require elevated or privileged access to support or operate the environment, provide a way for them to gain temporary elevated access. For an example of temporary elevated access to AWS IAM Identity Center, see Temporary elevated access for AWS accounts.
- **Restrict the allocation and use of IAM permissions providing service access:** Administrators providing second or third-level support that use the AWS Management console require IAM permissions. Grant the minimal set of permissions to administrative users providing an enhanced level of support to users using Amazon WorkSpaces and AppStream 2.0 for them to fulfill their role.
- **Restrict the scope of access for service accounts:** Restrict permissions for service accounts for Amazon WorkSpaces (with Active Directory Connector) and Amazon AppStream 2.0 (with domain-joined fleets) to only allow them to create computer objects within their designated Organizational Unit (OU). For implementing service accounts, see Amazon AppStream 2.0 Active Directory Administration and AD Connector prerequisites.

Identity and access management

EUCSEC04: How do you separate end user systems to meet your organization's policies?

When implementing EUC services, your organization may have requirements to separate compute devices accessed directly by end users from others used for infrastructure applications.

EUCSEC05: How do you manage application entitlements in your EUC environments?

Users should be entitled to access individual applications rather than provided access to all applications on end user systems. Apply this in a consistent way so that there is a minimal chance of operational failure or accidental granting of full access to all applications.

EUCSEC06: How do you authenticate and authorize access to your end user applications?

Strong and consistent authentication and authorization are key to the secure operation of an end user system to help prevent unauthorized access. Authentication using multiple factors may be a requirement, and the authentication system in use should satisfy this requirement.

Best practices

- EUCSEC04-BP01 Separate end user systems between different groups of users when required to satisfy policy or regulatory requirements
- EUCSEC05-BP01 Evaluate applications and data access requirements and implement entitlements accordingly
- EUCSEC06-BP01 Rely on a centralized authentication system that satisfies security requirements for your EUC environment
- EUCSEC06 BP02 Strengthen SAML federation to reduce security risks

EUCSEC04-BP01 Separate end user systems between different groups of users when required to satisfy policy or regulatory requirements

Many organizations have security requirements that mandate the segregation of systems accessed and interacted with by end users from servers that perform an infrastructure or application hosting function. Regardless of whether there is a specific security requirement, end user systems should be segregated from each other. This is for multiple reasons including reducing the risk of unintended access and exposure to unsafe software.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Use distinct AWS accounts to separate EUC services from other AWS workloads:** Separate AWS EUC workloads deployed to an AWS account from application and infrastructure servers and services that are consumed by the EUC workloads using different AWS accounts. You can use AWS Control Tower and AWS Organizations are two services to implement and manage a multi-account structure in your AWS environment. Create an AWS account for EUC workloads and use other accounts for infrastructure and application services. For more detail, see SEC01-BP01 Separate workloads using accounts.
- **Use IAM roles with AppStream 2.0 to enable access to AWS services**: To access AWS services from an AppStream 2.0 instance, use an IAM role and verify that the IAM policy attached to it is scoped to the specific services required. This approach avoids the need for users in AppStream 2.0 sessions to have access with additional credentials. If groups of users require differing levels of access to other AWS services, consider creating an additional role for each set of permissions. To help determine the least privilege policies based on the needed access, analyze user access with AWS IAM Access Analyzer. For further detail, see Use IAM Access Analyzer policy generation to grant fine-grained permissions for your AWS CloudFormation service roles.
- **Restrict access to only authorized applications**: By default, AppStream 2.0 allows users or applications to start programs on the instance, beyond what is specified in the image application catalog. This is useful when your application relies on another application as part of a workflow, but it may be undesirable for the user to be able to start that dependent application directly. For example, an application starts the browser to provide help instructions from an application vendor's website, but the ability for the user to start the browser directly must be blocked.

  In some situations, it can be desirable to control which applications can be launched on streaming instances. Microsoft AppLocker is application control software that uses explicit control policies to enable, or disable, the applications a user can run. An alternative to Microsoft AppLocker is FSLogix Application Masking which is available with Windows desktop and server operating systems. The use of application entitlements with AppStream 2.0 can restrict the ability of users to launch only authorized applications, but this control by itself does not prevent the launch of other applications on AppStream 2.0 instances. To achieve this, we recommend the two preceding approaches AppLocker or FSLogix.
- **Secure access to the S3 buckets used by Amazon AppStream 2.0:** Review, maintain, and update S3 bucket policies as appropriate. These reviews should verify that restricted access is in place to protect S3 buckets that are created and used to persist user data for both home folders and

application settings persistence when enabled. This blocks non-AppStream 2.0 administrators from accessing the data. Use S3 bucket policies and IAM policies together. For more information, see IAM Policies and Bucket Policies and ACLs! Oh, My! (Controlling Access to S3 Resources).

EUCSEC05-BP01 Evaluate applications and data access requirements and implement entitlements accordingly

Assess the types of users, their associated risk profile, and the access that each group of users requires to understand the access permissions each group of users require. Map the requirements to security groups, such as Active Directory security groups and the required permissions granted to these groups. Continually maintain the users associated with the group to verify ongoing appropriate access to applications and data.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Automate user entitlements**: Use a provisioning and entitlement system that automates the addition and removal of users from groups that provide role-based permissions access. Automation creates consistency in the approach for handling permissions.
- **Use templates for user creation:** Use templates when creating user accounts to avoid manual configuration of user groups and settings that may lead to overly permissive access.
- **Review user entitlements regularly**: Review user entitlements regularly to verify that they are aligned with each user's current role and access requirements to fulfill the role. Consider a regular cadence, such as a quarterly or monthly review.

EUCSEC06-BP01 Rely on a centralized authentication system that satisfies security requirements for your EUC environment

Evaluate your organization's security policies to determine the requirements that authentication systems need to provide for end users accessing EUC services.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Use authentication providers configured for best practices**: Consider the following regarding authentication of users accessing AWS EUC services using Microsoft Active Directory or an authentication provider:
  - Use a strong password policy.
  - Use multi-factor authentication (MFA) to provide additional protection to end users in your environment. For Amazon WorkSpaces and AppStream 2.0 environments integrated with a SAML IdP, enable MFA in the IdP. For Amazon WorkSpaces Personal where a SAML IdP is not in use, implement a RADIUS server to provide the MFA capability.
  - Consider adding password expiration policy to require users to change their passwords regularly.
  - When using a SAML identity provider (IdP), consider enabling advanced features like geo-restrictions and conditional access.
  - Using a corporate managed (and HR linked) identity provider improves security by automatically propagating role and permission changes to the EUC environment. It also promotes the best practice of managing access based on user lifecycle.

- **Users should be authenticated and authorized to access EUC services**: Use an authentication system, such as a SAML 2.0 IdP or Microsoft Active Directory, to authenticate users prior to them accessing an AWS EUC service. Verifying authenticating or authorization checks that only entitled users can access the applications and data accessible from Amazon WorkSpaces and AppStream 2.0 instances.
- **Manage user entitlements using groups where possible**: Use groups within Active Directory or your authentication provider instead of granting access to individual users. This approach simplifies the administration process and helps you perform access reviews and updates more efficiently.

EUCSEC06 BP02 Strengthen SAML federation to reduce security risks

To help prevent an opportunity for SAML assertions to be misused when using Certificate Based Authentication by incorrectly associating with Active Directory user security objects, strong mapping should be used.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Use strong mapping between SAML IdP and Active Directory. You can use certificate-based authentication (CBA) with Amazon WorkSpaces, which you can use to remove the user password prompt when using a SAML 2.0 identity provider. To establish a strong mapping between Active Directory users and SAML assertions, ObjectSid must be configured within the SAML assertion. CBA will fail if the attribute does not match the Active Directory security identifier (SID) for the user in the `SAML_Subject NameID`. For more detail, see Certificate-based authentication.

Detection

EUCSEC07: How do you monitor and track access to EUC environment?

Logging successful and unsuccessful access attempts to EUC instances and applications is a key practice for the secure delivery of end user applications. The logs can be used for troubleshooting purposes but also to audit access patterns and identify malicious behaviours.

EUCSEC08: How do you protect EUC endpoints from malicious software?

The ongoing logging of user activity and the associated analysis of those logs helps you detect unexpected behavior as anomalies that deviate away from the standard baseline behaviors.

EUCSEC09: How do you detect unexpected or unwanted configuration changes to end user applications?

Many organizations have security policies that require ongoing adherence to the base line security posture determined for systems delivering applications and desktops to users. They may also require that the security posture is not deviated away over time due to configuration drift.

Best practices

- EUCSEC07-BP01 Monitor user access to EUC instances and aggregate logs in central location

- EUCSEC08-BP01 Install endpoint protection software on instances to detect unexpected behavior
- EUCSEC09-BP01 Verify that your instances are configured as expected

EUCSEC07-BP01 Monitor user access to EUC instances and aggregate logs in central location

Record events generated when users access systems to a central system to log attempted and successful user authentication, as well as access to applications. Prior to implementation, consider that the use of a central system for security events may be subject to local regulation and legal framework.

Events logged in the central system should include the following data attributes:

- Timestamps
- User ID
- IP address
- Outcome of access attempt (success or failure)

Additional attributes or metadata may be required for compliance reasons. Evaluate any applicable regulatory and organizational security policy requirements to determine the complete set of attributes to record.

For completeness, consider all possible sources of events, including:

- Service-emitted events and logs (for example, Amazon WorkSpaces EventBridge events and AppStream 2.0 usage reports)
- Data plane logs collected through agents installed onto Amazon WorkSpaces or AppStream 2.0 instances

For Windows instances, use events recorded in the Windows security log alongside a log management system to collect and aggregate data from various sources, such as other text-based logs, network devices, and security applications. This integration provides a deeper insight into potential security issues so that your organization can address them.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Use agents on Amazon AppStream 2.0 and Amazon WorkSpaces instances to aggregate security logs. If instance security logs need to be captured from AppStream 2.0 instances, then event forwarding agents such as Amazon CloudWatch, Amazon Kinesis Agent for Windows, or Telegraf can be used to forward relevant events into the central security logging system.

For WorkSpaces, these agents can be pre-installed into a WorkSpaces custom bundle to make sure a logging capability is available before users attempt to access WorkSpaces. For AppStream 2.0, these agents need to be installed into the Image Builder for On-Demand and Always-On fleets.

EUCSEC08-BP01 Install endpoint protection software on instances to detect unexpected behavior

Endpoint protection software can provide the capability to detect anomalous behavior on end user computing services.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Configure security software for Amazon AppStream 2.0:** If you choose to install security software (for example, anti-virus or behavioral anomaly detection) on your image, we recommend that you do not enable automatic updates for the software. Otherwise, the software may attempt to update itself with the latest definition or configuration files or other updates during user sessions, which can affect performance. In addition, updates made to the software will not persist beyond the current user session. To verify that your fleet instances have the latest updates, we recommend that you do either of the following:
  - Update your image builder and create a new image on a regular basis (for example, by using the Image Assistant CLI operations).
  - Use security software that delegates scanning, detection, or other operations to an continuously updated external server.
  - For more detail, see Administer Your Amazon AppStream 2.0 Images and Best Practices for Deploying Amazon AppStream 2.0.
- **Configure security software for Amazon WorkSpaces:** Security software can adversely affect the operation of Amazon WorkSpaces if it is not configured to consider the requirements of the service. For details on the configuration elements that are required to be considered as

exclusions for anti-malware scanning, see Required configuration and service components for WorkSpaces Personal. The configuration of endpoint security software should verify that the status of the agents deployed on Amazon WorkSpaces is centralized to provide a consolidated view of the status of the deployed Amazon WorkSpaces.

EUCSEC09-BP01 Verify that your instances are configured as expected

Unexpected configuration changes to end user systems can help you identify possible threat actors. Users should not need to reconfigure applications or operating systems for the daily use of their application portfolio.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Automate configuration management tools to verify compliance. Use centralized control and enforcement of configuration settings applied to Amazon WorkSpaces and AppStream 2.0 instances to verify that configuration settings align with the desired configuration of instances.

- For Active Directory domain-joined Windows instances, use Group Policy Objects (GPOs) to apply a known configuration to instances.
- For Amazon WorkSpaces Linux instances, consider configuration management tools such as Ansible, Chef, and Puppet to apply a known configuration.
- For Amazon AppStream 2.0 On-Demand and Always-On fleets, apply desired configuration settings to the instance used to create the Image for the associated fleet.

After deployment, you can audit Amazon WorkSpaces Personal instances to determine if the expected and desired configuration of instances is in effect or whether this has been overridden or tampered with. Configuration management tools such as Ansible, Chef, and Puppet can help with this, as can PowerShell Desired State Configuration.

Infrastructure protection

EUCSEC10: How do you implement network protection in your EUC environment?

Network protection is required between infrastructure, application servers, and end user facing systems to help you adhere to your organization's security policies and reduce risks.

EUCSEC11: How do you scan for vulnerabilities and perform patch management for your EUC instances?

Organizations commonly implement security policies requiring robust and continuous protection of systems that provide end user computing services. As new exploits emerge and vulnerabilities are discovered regularly, maintaining strong security is an ongoing process. You should routinely assess and update your systems with security patches to maintain an effective defence against evolving threats.

EUCSEC12: How do you prevent user access to non-essential software binaries present on systems that cannot be uninstalled?

Many organizations have security policies that require that users in end user computing systems only have access to the software and applications they need to fulfil their role.

Best practices

- EUCSEC10-BP01 Implement network separation for AWS EUC instances
- EUCSEC10-BP02 Restrict access to open ports on instances to reduce risks
- EUCSEC11-BP01 Perform vulnerability scanning on EUC instances
- EUCSEC12-BP01 Allow user access to only the software binaries needed to perform their job

EUCSEC10-BP01 Implement network separation for AWS EUC instances

Separating end user systems from infrastructure, application servers, and data at the network level verifies that you can enforce minimal access between systems to help prevent unauthorized access to data and applications.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Enforce network separation between user instances and other services. EUC instances provided by Amazon WorkSpaces or AppStream 2.0 usually have network connectivity to other workloads in the same network subnet. The use of security groups within VPCs can restrict lateral movement and are recommended for implementation. For defense-in-depth, non-end user instances such as application servers, authentication providers, and other infrastructure services should reside on subnets different to those where user instances reside.

You can apply security controls to the non-end user instances at various points using AWS capabilities, such as separate AWS accounts and VPCs, VPC endpoints, proxy servers, and network firewalls. Review network security best practices for WorkSpaces and AppStream 2.0 to improve security posture in your EUC environment.

EUCSEC10-BP02 Restrict access to open ports on instances to reduce risks

Restrict use of network ports on end user systems to reduce the potential exposure surface of these systems. Block network ports that aren't required for the operation and support of end user systems using host-based or network firewalls.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Implement networking security controls on Amazon EUC instances. AWS provides several services and capabilities that can help you secure AWS EUC instances for Amazon WorkSpaces and AppStream 2.0. In addition to these services, consider OS capabilities and additional software to provide the required level of security.

For AWS networking, the following services and features should be evaluated:

- Network ACLs
- Security groups
- AWS Network Firewall
- NAT Gateway

Consider these services to create a baseline of network security. Additionally, review and explore best practices for VPC and networking in WorkSpaces, as well as best practices for deploying AppStream 2.0, as you evaluate your network security.

In addition to AWS security capabilities and services, when users require access to the Internet from browsers installed in Amazon WorkSpaces or AppStream 2.0 instances, consider using a web proxy to log web site access and implement restrictions on where users can browse.

In Amazon WorkSpaces and AppStream 2.0 instances, consider existing OS software to harden the instances. For example, you can use host-based firewalls available within the operating system to restrict accessible ports in your instances. In addition, consider endpoint protection software to identify and mitigate security risks that may be introduced into the environment using software local to the instances. For detail on the ports required by Amazon WorkSpaces and AppStream 2.0, see the following:

- List of ports required by Amazon AppStream 2.0
- List of ports required for Amazon WorkSpaces

EUCSEC11-BP01 Perform vulnerability scanning on EUC instances

The frequent release of patches for vulnerabilities in operating systems and applications means that you should patch them on a frequent basis to address potential risks.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Perform frequent vulnerability scanning and patch instances accordingly:** Software patching is critical for the security and performance of compute resources. Frequent patching is a best practice in the security pillar of the Well-Architected Framework.
- **Regularly patch Amazon AppStream 2.0 images:** As part of the AWS Shared Responsibility Model, customers are responsible for patching and securing their AppStream 2.0 images. When an image is built and deployed, there are five categories of software that require patching in your AppStream 2.0 image:

- **Applications and dependencies:** Customers are responsible for patching the applications and dependencies in images.

- **Operating system:** Customers are responsible for installing and maintaining updates for Linux and Windows.
- **Software components:** These are drivers, agents, and other software required for AppStream 2.0 operation (for example, the Amazon CloudWatch agent). AppStream 2.0 periodically releases new base images that contain new agents and drivers. Customers can recreate their images using the latest base image to bring the software components to the latest baseline.
- **AppStream 2.0 agent**: Customers can choose to consistently use the latest agent version in the Image Assistant. With this option, streaming instances that are launched from the image automatically use the latest version of the agent.
- **Clients**: Where the Amazon AppStream 2.0 client is in use, this should also be updated upon the release of each new version.

- **Regularly patch Amazon WorkSpaces Personal instances:** Amazon WorkSpaces Personal instances need to be scanned for vulnerabilities and patched regularly post-deployment. Use configuration management tools or patch management tools to satisfy the requirement for ongoing assessment and deployment of patches. The Amazon WorkSpaces client should also be updated upon the release of a new version.

EUCSEC12-BP01 Allow user access to only the software binaries needed to perform their job

Users should only have access to the software binaries required for them to perform their role. Access to additional software that could introduce risks should be blocked.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Implement security controls to restrict access to software binaries. Permissions applied to software binaries present on Amazon WorkSpaces or AppStream 2.0 instances should restrict the ability for users to run the programs and applications that they require to fulfill their role. Evaluate other software binaries present in the image to verify that the default permissions applied in the file system do not permit users to run them.

System hardening should also be considered to further secure the operating system image. For reference, consider the Center for Internet Security (CIS) AWS End User Compute Services Benchmark. You can apply your chosen security settings by incorporating them into the image

pre-deployment, post-deployment using scripts, or for Windows instances by using Group Policy Objects (GPOs). In addition, for Windows instances, consider FSLogix or AppLocker to restrict access to specific software binaries.

Data protection

EUCSEC13: How do you address data residency requirements in your EUC environment?

Data residency requirements relate to the physical or geographic location where an organization chooses to store or process regulated data. End user systems that process regulated data are subject to data residency requirements and must be considered when choosing where to implement an end user system.

EUCSEC14: How do you protect data processed and stored in EUC instances?

Many organizations maintain security policies requiring protection of both data at rest (information stored in persistent storage) and data in transit (information moving between source and destination systems).

EUCSEC15: How do you reduce the risk of data loss in your EUC environment?

The provisioning of an end-user system to deliver applications marks the beginning of its lifecycle. Consider how you want to address ongoing management, operation, and configuration. These aspects must be addressed throughout the system's lifecycle to maintain security, reliability, availability, performance, and operational continuity.

Best practices

- EUCSEC13-BP01 Align your compliance of data storage with policies and regulatory requirements
- EUCSEC14-BP01 Encrypt disk volumes to protect data at rest
- EUCSEC14-BP02 Encrypt data in transit in your EUC environment

- EUCSEC14-BP03 Limit egress channels available to users to only the required set of channels to perform their role
- EUCSEC15-BP01 Encourage users to store data on long-term storage services

EUCSEC13-BP01 Align your compliance of data storage with policies and regulatory requirements

The storage of data accessible by users and applications on end user systems should align with and comply with the data residency requirements for the data and the organization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Data should be stored and accessed in compliance with the in-scope policies and regulatory requirements. The location of data and the applications accessing data should align with the compliance framework and requirements for the respective organization. To achieve this, consider your AWS Region for compliance against the data sovereignty requirements for the application and data. Additionally, consider data permissions to verify compliance and enforce least privilege access. Keep latency between end user devices and the data they need to access in consideration when choosing the location of the EUC environment but also adhere to data residency requirements.

EUCSEC14-BP01 Encrypt disk volumes to protect data at rest

Protect security, integrity, and availability of data at rest to make sure it is reliably accessible when needed.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Encrypt Amazon WorkSpaces Personal disk volumes. Each Amazon WorkSpace Personal instance is provisioned with a root volume (C: drive for Windows WorkSpaces Personal, root file system for Amazon Linux WorkSpaces Personal) and a user volume (D: drive for Windows WorkSpaces Personal, /home for Amazon Linux WorkSpaces Personal). The encrypted WorkSpaces feature encrypts one or both volumes. For WorkSpaces Personal instances used by users (rather than for creating custom images), it is a best practice for these to be encrypted. For more details, see Encrypted WorkSpaces in WorkSpaces Personal.

EUCSEC14-BP02 Encrypt data in transit in your EUC environment

Use encryption to protect data confidentiality while in transit inside your EUC environment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Use AWS EUC streaming protocols to encrypt streaming data in transit. Amazon WorkSpaces and Amazon AppStream 2.0 provide data encryption of pixel streaming traffic between instances and end user devices by default. Evaluate the default levels of encryption to verify that they provide sufficient protection in terms of key length and cipher suites and satisfy the requirements of the organization. For further details regarding the encryption used for Amazon AppStream, see Data Protection in Amazon AppStream 2.0 , and for Amazon WorkSpaces, see Data Protection in Amazon WorkSpaces.

EUCSEC14-BP03 Limit egress channels available to users to only the required set of channels to perform their role

End user systems can provide multiple channels for users to export and access data. Evaluate these channels to determine their suitability in the specific use case being delivered. Block channels not required for specific use cases.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Encrypt streaming and control data traffic using strong ciphers**: To protect data confidentiality, WorkSpaces using PCoIP are encrypted using an AES 128-bit cipher by default. For encryption up to AES 256-bit see Data protection in Amazon WorkSpaces. Evaluate your security requirements and use stronger ciphers where necessary. For Windows, you can implement this using the Group Policy template, and for Linux WorkSpaces, the appropriate configuration file needs to be edited to increase the default level of encryption. For example, for PCoIP Amazon Linux 2 WorkSpaces, edit the /etc/pcoip-agent/pcoip-agent.conf file.
- WorkSpaces using the Amazon DCV protocol have streaming and control data in-transit encrypted using DTLS 1.3 encryption for UDP traffic and TLS 1.3 encryption for TCP traffic with AES-256 ciphers. For details of the implementation, see Data protection in Amazon WorkSpaces.
- **Restrict data access to required functionality within Amazon WorkSpaces**: To protect data on the endpoint used to connect to an Amazon WorkSpaces session and the WorkSpace itself,

enable data exportation features only when needed and allowed to users. For example, Amazon WorkSpaces can block copying in-session clipboard contents to the endpoint, copying of files between client and WorkSpace, and block printers attached to the endpoint from being mapped into the session. The blocking of these capabilities can remove these potential data exportation vectors from the Amazon WorkSpaces service.

- The implementation of these controls is through Group Policy on Windows WorkSpaces, editing the /etc/pcoip-agent/pcoip-agent.conf file on Amazon Linux 2 WorkSpaces using PCoIP, or editing the /etc/wsp/wsp.conf file on Ubuntu Amazon WorkSpaces using Amazon DCV. For details on how to configure clipboard and other settings on Windows WorkSpaces, see Manage your Windows WorkSpaces in WorkSpaces Personal.
- **Restrict data access to required functionality within Amazon AppStream 2.0**: To protect data on the endpoint used to connect to an Amazon AppStream 2.0 session and the AppStream 2.0 instance itself, implement controls to close potential inbound or outbound channels that are not required by the users connecting to the service. The service has controls to configure the clipboard, file transfer, printing to a local device, and file system redirection. You can configure each of these options on an AppStream 2.0 stack and disable them when not required. For details on configuring data access restrictions with Amazon AppStream 2.0, see Create an Amazon AppStream 2.0 Fleet and Stack.

EUCSEC15-BP01 Encourage users to store data on long-term storage services

Educate users to avoid storing critical data directly on EUC systems without also saving that data to an approved, long-term storage solution that is regularly backed up. This practice verifies that data remains visible, accessible, and protected against loss in the event of an EUC instance failure or lifecycle event such as termination or rebuild.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Verify that EUC services are not used for long-term data storage. EUC services such as Amazon WorkSpaces and Amazon AppStream 2.0 are optimized for application delivery and user productivity, rather than as primary long-term data storage solutions. Amazon AppStream 2.0 streaming instances are non-persistent, meaning data stored locally during a session is lost when the instance is recycled or terminated. Amazon WorkSpaces provides persistent root and user volumes, which are well-suited for user profiles, application settings, and day-to-day productivity

tasks. However, for storing critical or long-term data, organizations should use dedicated storage services that offer centralized management, data durability, and backup capabilities.

To reduce the risk of data loss, encourage users to save important data to approved, persistent storage services that align with the organization's data protection and recovery requirements. Recommended options include Amazon FSx for Windows File Server, which integrates with WorkSpaces to provide durable, backed-up home directories, as well as Amazon S3 or other enterprise-grade cloud storage services. Implement clear governance policies and user education to verify that data is consistently stored in systems designed for long-term retention and resilience.

Incident response

There are no incident response practices unique to this lens. For information on Security, refer to the Security Pillar - AWS Well-Architected Framework. For rapid recovery in response to a security incident affecting an EUC service consider the best practices documented within the reliability pillar of the EUC lens.

Application security

There are no application security practices unique to this lens. For information on Security, refer to the Security Pillar - AWS Well-Architected Framework.

Key AWS services

- Amazon WorkSpaces
- Amazon AppStream 2.0
- Amazon Virtual Private Cloud
- AWS Identity and Access Management
- Amazon CloudWatch
- Amazon EventBridge
- Amazon S3
- Amazon FSx for Windows Server
- Amazon Kinesis Agent for Microsoft Windows
- AWS Directory Service

Resources

Related documents:

- Identity and access management for WorkSpaces
- Using AWS Managed Policies and Linked Roles to Manage Administrator Access to AppStream 2.0 Resources
- AppStream 2.0 Active Directory Administration
- Getting started with AD Connector
- Using an IAM Role to Grant Permissions to Applications and Scripts Running on AppStream 2.0 Streaming Instances
- Amazon AppStream 2.0: Manage Application Entitlements
- Administer Your Amazon AppStream 2.0 Images
- Required configuration and service components for WorkSpaces Personal
- Tutorial: Creating and Streaming from Interface VPC Endpoints
- IP address and port requirements for WorkSpaces Personal
- Encrypted WorkSpaces in WorkSpaces Personal
- Data protection in Amazon WorkSpaces
- Manage your Windows WorkSpaces in WorkSpaces Personal
- Create an Amazon AppStream 2.0 Fleet and Stack
- Disaster Recovery considerations with Amazon AppStream 2.0
- Cross-Region redirection for Amazon WorkSpaces
- Multi-Region Resilience for Amazon WorkSpaces
- Announcing cross-region data replication for Amazon WorkSpaces
- Business Continuity and Disaster Recovery with Amazon WorkSpaces
- Building a multi-region disaster recovery environment for Amazon WorkSpaces
- Getting started with GuardDuty
- Administer Your Amazon AppStream 2.0 Images
- Certificate-based authentication with WorkSpaces
- IAM Policies and Bucket Policies and ACLs! Oh, My! (Controlling Access to S3 Resources)
- What is AWS Control Tower?
- Temporary elevated access for AWS accounts

- Use IAM Access Analyzer policy generation to grant fine-grained permissions for your AWS CloudFormation service roles
- Best Practices for VPCs and Networking in Amazon WorkSpaces Deployments
- Best Practices for Deploying Amazon AppStream 2.0

Related partner solutions:

- Microsoft AppLocker
- Microsoft FSLogix

Reliability

The reliability pillar for end-user computing encompasses the ability of your user-facing workloads and services to perform their intended functions consistently and correctly. This pillar provides in-depth, best-practice guidance for implementing reliable and resilient end-user computing environments on AWS.

The reliability pillar for end-user computing provides guidance to help customers apply best practices in the design, delivery, and maintenance of their AWS-based end-user computing infrastructure and services.

By following the guidance in this whitepaper, customers can build and operate highly reliable and resilient end-user computing solutions on AWS, delivering consistent and dependable experiences for their end-users. You can find prescriptive guidance on implementation in this Reliability Pillar whitepaper.

Focus areas

- Design principles
- Definitions
- Foundations
- Workload architecture
- Change management
- Failure management
- Key AWS services
- Resources

Design principles

- **Automatically recover from failure:** By monitoring Amazon WorkSpaces and AppStream for key performance indicators (KPIs), you can run automation when a threshold is breached. These KPIs should be a measure of business value, not just technical aspects of the service operation. This allows for automatic notification and tracking of failures, as well as automated recovery processes that remediate the failure. For example, you can use Amazon CloudWatch to set alarms and Amazon EventBridge to route these events to AWS Lambda or AWS Systems Manager, which can run recovery actions automatically.

- **Test recovery procedures:** Enable the simulation of various failure modes or the replication of historical failure scenarios. This approach facilitates the identification and exploration of potential failure pathways, allowing for preemptive testing and remediation. Addressing common and previously experienced vulnerabilities proactively helps you reduce the risk of encountering actual failures in production environments. Regularly test Amazon WorkSpaces and AppStream recovery using AWS Fault Injection Service or other DR tools to validate recovery procedures.
- **Scale horizontally to increase aggregate workload availability:** Replace one large resource with multiple small resources to reduce the impact of a single failure on the overall workload. Distribute requests across multiple, smaller resources to verify that they don't share a common point of failure. For Amazon WorkSpaces, this means deploying instances across multiple Availability Zones and using service-specific load balancing to distribute user connections, enhancing availability and fault tolerance.
- **Stop guessing capacity:** A common cause of failure in on-premises workloads is resource saturation. This occurs when the demand placed on a workload exceeds its capacity. In the cloud, you can monitor demand and workload utilization and adjust the number of WorkSpaces to maintain the optimal level to satisfy demand without over-provisioning or under-provisioning.

Definitions

Foundations:

- Foundational requirements are those whose scope extends beyond a single workload or project. Before architecting any system, foundational requirements that influence reliability should be in place.
- In an on-premises environment, these requirements can cause long lead times due to dependencies and therefore must be incorporated during initial planning. With AWS, however, most of these foundational requirements are already incorporated or can be addressed as needed. The cloud is designed to be nearly limitless, so it's the responsibility of AWS to satisfy the requirement for sufficient networking and compute capacity, leaving you free to change the end user computing environment on demand.

Amazon WorkSpaces architecture:

- A reliable Amazon WorkSpaces deployment starts with upfront design decisions for both software and infrastructure. Your architecture choices will impact your WorkSpaces behavior

across the six Well-Architected pillars. For reliability, there are specific patterns you must follow to provide high availability and improve fault tolerance.

- Use Amazon Virtual Private Cloud (VPC) for network isolation and security.

Change management:

- Changes to your Amazon WorkSpaces environment must be anticipated and accommodated to achieve reliable operation. Changes include those imposed on your WorkSpaces, such as spikes in demand, as well as those from within, such as feature deployments and security patches.
- Use AWS CloudFormation, Service Catalog, or AWS Systems Manager to manage and automate changes to your WorkSpaces environment, creating consistency and reducing the risk of errors.

Failure management:

- Amazon WorkSpaces are deployed in a specific AWS Region with built-in redundancies to protect against component failures. This provides high availability and minimizes downtime.
- There is the potential for failures to impact your WorkSpaces environment. Therefore, you must take steps to implement resiliency if you need your WorkSpaces to be reliable.
- Spread awareness amongst the people designing, implementing, and operating your Amazon WorkSpaces about business objectives and the required reliability goals to achieve them. Leaders or system owners must provide training and guidance to verify that individuals understand and can design for the reliability requirements pertinent to their roles.

Foundations

EUCREL01: How do you increase resilience and minimize impact of failure in your EUC environment?

An Amazon WorkSpaces environment can achieve increased resilience using the native deployment pattern within the managed services that dictates that a minimum of two Availability Zones (AZs) are used. Where increased resilience is required at the regional rather than zonal level, multi-region Amazon WorkSpaces environments can be deployed in separate regions.

For Amazon AppStream 2.0, increased resilience can be achieved by deploying fleets across a minimum of two AZs and using three AZs where possible. Where increased resilience is required at the regional rather than zonal level, multi-region Amazon AppStream 2.0 environments can be deployed in separate regions. This is achieved by copying images between regions and establishing separate fleets in multiple regions.

Best practices

- EUCREL01-BP01 Add redundancy and remove single points of failure in your environment

EUCREL01-BP01 Add redundancy and remove single points of failure in your environment

The principle of assuming that failures will occur represents a paradigm shift in the approach to designing Amazon WorkSpaces and AppStream 2.0 environments. By adopting this mindset, organizations can prioritize resilience and develop strategies that minimize the impact of failures, thereby reducing downtime and mitigating potential business disruptions.

Level of risk exposed if this best practice is not established: High

Implementation guidance

When designing an Amazon WorkSpaces or AppStream 2.0 environment, the approach should prioritize resilience and minimize the impact of failures by assuming that failures will occur and implementing robust strategies.

Implement redundancy at every layer of your architecture. This includes network paths, storage, and virtual desktops. Use multiple instances of Amazon WorkSpaces or AppStream 2.0 so that if one fails, others can take over seamlessly. For AppStream 2.0, use automatic scaling to match the number of running instances to user demand, keeping performance consistent even during usage spikes.

Regularly test your failure recovery procedures. Use AWS tools such as AWS Fault Injection Service to simulate different failure scenarios and validate your recovery strategies.

Implement robust data backup and disaster recovery plans. Regularly back up user data and configurations, and verify that you have a tested recovery plan in place to restore operations quickly in case of a failure.

Set up comprehensive monitoring using Amazon CloudWatch to keep track of the performance and health of your WorkSpaces and AppStream 2.0 environments. Create alarms and automated responses to detect and remediate detected issues promptly.

Continuously review and improve your architecture and operational procedures. Learn from historical incidents and update your strategies to help prevent future occurrences.

Workload architecture

EUCREL02: How do you minimize impact of Regional disruptions in your EUC environment?

Explore the resilience strategies provided by Amazon WorkSpaces Multi-Region to support operational continuity and minimize service interruptions, especially during events with a low Recovery Time Objective (RTO).

EUCREL03: How do you minimize impact of networking disruptions in your EUC environment?

Explore the benefits of configuring redundant networking components for Amazon WorkSpaces and AppStream 2.0 to enhance resilience and support continuous connectivity. Additionally, create redundancy to allow for authentication and authorization remain operational even if certain components or resources experience disruptions.

EUCREL04: How do you protect data processed and stored in EUC instances from loss?

Implement data replication and backup strategies, including automated solutions like Amazon WorkSpaces Automated Snapshots, to strengthen resilience and protect data within Amazon WorkSpaces environments.

EUCREL05: How do you monitor availability and respond to availability drops in your EUC environment?

Use Amazon CloudWatch to maintain resilience for Amazon WorkSpaces and AppStream 2.0. Consider automated remediation workflows that can help you minimize downtime and address critical issues.

EUCREL06: How do you test disaster recovery plans for your EUC environment on AWS?

Promote preparedness and resilience planning in your Amazon WorkSpaces and AppStream 2.0 environments. Document and understand the steps that are essential to verify readiness for disruptive events or data loss situations.

Best practices

- EUCREL02-BP01 Use multiple regions for your EUC environment to minimize downtime
- EUCREL03-BP01 Add redundancy to networking connections
- EUCREL04-BP01 Establish data integrity with replication and backup strategies
- EUCREL05-BP01 Monitor and automate remediation for Amazon WorkSpaces and AppStream 2.0
- EUCREL06-BP01 Plan for disaster recovery of EUC through testing and procedures

EUCREL02-BP01 Use multiple regions for your EUC environment to minimize downtime

Amazon WorkSpaces Multi-Region Resilience offers cost-effective, easy-to-manage operational continuity solutions that keep your users online and productive . Organizations should proactively design their environment to anticipate failure and plan for a fast recovery.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Implement Amazon WorkSpaces Multi-Region Resilience to enable cost-effective and smoothly managed operational continuity. This approach verifies that users remain online and productive with minimal recovery time through standby WorkSpaces in alternative AWS Regions during disruptive events. Additionally, regularly test your multi-Region setup to verify its effectiveness in supporting operational continuity. Conduct failover drills and simulations to validate the RTO

and identify any potential areas of improvement in your resilience strategy. By using Multi-Region Resilience, you can minimize service interruptions and provide uninterrupted access to Amazon WorkSpaces for your users, even during disruptive events.

EUCREL03-BP01 Add redundancy to networking connections

Use a redundant networking architecture for Amazon WorkSpaces and AppStream 2.0, incorporating multiple Active Directory (AD) Controllers, AD connectors, DNS servers, gateways, VPNs, or AWS Direct Connect links. This approach supports continuous connectivity by providing alternative pathways for network traffic, reducing the risk of service disruptions due to network incidents and enhancing overall system resilience. Redundant networking helps mitigate the impact of network failures and supports uninterrupted access to WorkSpaces environments.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Enhance the resilience of Amazon WorkSpaces and AppStream 2.0 by configuring redundant networking components such as VPN connections or AWS Direct Connect links. This setup provides alternative paths for network traffic, mitigating the impact of network incidents and supporting continuous access to WorkSpaces environments. Verify that you have multiple AD controllers and connectors across multiple Availability Zones.

Additionally, regularly monitor and test the redundant networking setup to check its effectiveness in maintaining continuous connectivity. Conduct failover tests and simulations to validate the redundancy configuration and identify any potential areas for improvement. By implementing redundant networking architecture, you can strengthen the resilience of your EUC environment and reduce the risk of downtime caused by network disruptions.

EUCREL04-BP01 Establish data integrity with replication and backup strategies

Implement data replication and backup strategies to safeguard user data and configurations. Use automated backup solutions such as Amazon WorkSpaces Automated Snapshots to create regular backups of WorkSpaces volumes. Store data backups securely and verify that you can promptly restore backups in the event of data loss or corruption.

Level of risk exposed if this best practice is not established: High

Implementation guidance

To fortify resilience and safeguard data within Amazon WorkSpaces environments, adopt a comprehensive approach to data replication and backup strategies. Use automated solutions to regularly capture backups of user data. Store backups are stored securely, and check that you can readily access them for prompt restoration in the event of data loss or corruption.

Additionally, establish a backup retention policy to determine how long backups are retained, and verify your compliance with regulatory requirements. Regularly test the effectiveness of your backup and restoration processes and identify any potential areas for improvement proactively. By implementing robust data protection practices, you can strengthen the resilience of your WorkSpaces infrastructure and protect valuable user data and configurations, supporting operational continuity and reducing the risk of data loss.

EUCREL05-BP01 Monitor and automate remediation for Amazon WorkSpaces and AppStream 2.0

Implement comprehensive monitoring and alerting for Amazon WorkSpaces and AppStream 2.0 environments. Monitor key metrics such as instance health, network connectivity, and user activity using Amazon CloudWatch. Set up automated remediation workflows to respond to critical issues, such as instance failures or resource constraints, and automatically initiate recovery actions to minimize service interruptions.

Level of risk exposed if this best practice is not established: High

Implementation guidance

To strengthen resilience in EUC environments, establish robust monitoring using Amazon CloudWatch to detect anomalies in instance health, network connectivity, and user activity. Set up CloudWatch alarms to proactively identify potential areas of improvement and run automated remediation workflows. Configure these workflows to automatically address critical events, such as instance failures or resource constraints, by initiating recovery actions such as instance restarts or scaling adjustments.

Thoroughly test and validate automated remediation processes to minimize service interruptions and maintain continuous operations. Additionally, regularly review and refine your monitoring and automation strategies to align with evolving workload demands and infrastructure changes, supporting ongoing resilience in your EUC environments.

EUCREL06-BP01 Plan for disaster recovery of EUC through testing and procedures

Develop and regularly test disaster recovery plans for Amazon WorkSpaces and AppStream 2.0 deployments. Document procedures for restoring user data and configurations in the event of disruptive incidents or data loss. Conduct periodic disaster recovery drills to assess the effectiveness of recovery procedures and verify that you are ready to respond to unexpected incidents.

Level of risk exposed if this best practice is not established: High

Implementation guidance

To improve readiness for significant incidents or data loss incidents in Amazon WorkSpaces and AppStream 2.0 deployments, organizations must prioritize the development and regular testing of disaster recovery plans.

Document procedures for restoring user data and configurations, including backup and restoration processes, and verify that this documentation is quickly accessible to relevant personnel. Additionally, conduct periodic disaster recovery drills to simulate real-world scenarios and validate the effectiveness of recovery procedures. Use these drills to identify areas for improvement in the disaster recovery plan and take proactive measures to address them.

By investing in proactive disaster recovery planning and testing, organizations can mitigate the impact of unexpected events, provide business continuity, and protect valuable data and resources in their Amazon WorkSpaces and AppStream 2.0 environments. These best practices help organizations strengthen the resilience and availability of their EUC environments, minimize the impact of potential incidents, and support continuous access to virtual desktop resources for users.

Change management

EUCREL07: How do you track changes and implement reversible deployments in your change management process?

Maintain thorough documentation for changes in your Amazon WorkSpaces and AppStream 2.0 environments.

EUCREL08: How do you test andd verify changes in your EUC environment before pushing to end users?

Evaluate changes in a controlled environment before deploying them in the production environment for Amazon WorkSpaces and AppStream 2.0.

EUCREL09: How do you test rollback plans in your EUC change management process?

Create rollback plans for changes in Amazon WorkSpaces and AppStream 2.0, and test these procedures to maintain continuous operations and facilitate timely recovery from issues or failures.

EUCREL10: How do you communicate and coordinate changes with EUC environment stakeholders?

Notify stakeholders about changes, and coordinate with other teams to support continuous operations for users in your EUC environments.

EUCREL11: How do you perform post-change assessments in your EUC environment?

When you perform post-change evaluations in Amazon WorkSpaces and AppStream 2.0, use metrics and user feedback to inform future change management.

Best practices

- EUCREL07-BP01 Document changes for transparency and traceability
- EUCREL08-BP01 Test and validate changes to promote reliable deployment
- EUCREL09-BP01 Implement and test rollback plan for every change you make in EUC environments
- EUCREL10-BP01 Implement communication plans with EUC environment stakeholders

- EUCREL11-BP01 Implement post-change assessment to evaluate impact and optimize performance

EUCREL07-BP01 Document changes for transparency and traceability

Maintain comprehensive documentation of all modifications made to the EUC environment. Document the details of each change, including the rationale, implementation steps, and anticipated outcomes. Documentation helps promote transparency and traceability and provides a reference for future troubleshooting or auditing purposes.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Thoroughly document all changes made to the Amazon EUC environment to promote transparency and shared responsibility. This documentation serves as a valuable resource for troubleshooting and auditing purposes, providing insights into the history of changes and facilitating the identification of potential areas of improvement.

Additionally, consider establishing version control mechanisms and documentation guidelines to maintain consistency and facilitate collaboration among team members. By prioritizing comprehensive change documentation, organizations can enhance transparency, foster shared responsibility, and streamline troubleshooting and auditing processes in their WorkSpaces and AppStream 2.0 environments.

EUCREL08-BP01 Test and validate changes to promote reliable deployment

Thoroughly test and validate changes in an isolated environment before implementing them in the production environment for Amazon WorkSpaces and AppStream 2.0. By conducting testing in a controlled environment, organizations can assess the impact of changes on WorkSpaces and application availability under conditions that closely resemble real-world usage scenarios. This approach evaluates whether the changes achieve the desired outcomes before deployment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Before implementing changes in the production environment for Amazon WorkSpaces and AppStream 2.0, test and validate the changes in a controlled testing environment. Test changes under conditions resembling real-world usage scenarios to assess their impact on availability, performance, and functionality.

Validate that changes achieve desired outcomes and identify any potential areas for improvement or dependencies. Additionally, consider implementing automated testing frameworks and deployment pipelines to streamline the testing and validation process and maintain consistency in testing procedures across different environments.

By prioritizing comprehensive testing and validation, organizations can reduce the risk of service interruptions and issues during deployment, supporting smooth transitions and maintaining the stability and performance of their EUC environments.

EUCREL09-BP01 Implement and test rollback plan for every change you make in EUC environments

Develop rollback plans for changes in Amazon WorkSpaces and AppStream 2.0 to anticipate and address potential failures and their impacts to system stability or resilience. By establishing these plans, businesses can proactively address unforeseen situations that may arise during implementation, creating a smooth transition back to previous configurations if needed. Test rollback procedures beforehand to gain insights into their effectiveness and identify any potential areas of improvement. This proactive approach minimizes service interruptions and facilitates prompt recovery from incidents or failures, supporting continuous service delivery and reducing impact on users.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Develop rollback plans for changes that could potentially impact WorkSpaces and AppStream 2.0 resiliency or stability. Define procedures for reverting to the previous state if a change causes unexpected incidents or service disruptions. Test and validate rollback plansto verify their effectiveness and minimize the time required to restore service in such situations in the event of a rollback.

Additionally, consider implementing automated rollback mechanisms where feasible to streamline the recovery process and reduce manual intervention. By prioritizing rollback planning and testing,

organizations can enhance their ability to respond effectively to unexpected challenges and maintain continuous operations in their EUC environments.

EUCREL10-BP01 Implement communication plans with EUC environment stakeholders

Include WorkSpaces and AppStream 2.0 users, administrators, and support teams in your communications. Provide advance notice of scheduled maintenance windows or change activities to minimize potential service interruptions for users. Coordinate with other teams or departments as necessary to implement changes smoothly and with minimal impact on related systems or services.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Effective communication and coordination are critical components in minimizing service interruptions for end users during periods of change. Establish clear communication channels.

Begin by identifying all stakeholders involved, including users, administrators, support teams, and any other relevant parties. Develop a comprehensive communication plan outlining how and when stakeholders will be informed of changes. Use various channels such as email, in- system announcements, or dedicated communication systems to facilitate timely notifications.

Transparency is key, so communicate changes clearly, and provide rationale, expected impact, and necessary instructions. Additionally, establish a feedback mechanism for stakeholders to express concerns or ask questions. Keep stakeholders informed with regular updates on implementation progress and timelines.

By prioritizing communication and coordination, you can minimize disruptions and support a smooth transition process for all involved.

EUCREL11-BP01 Implement post-change assessment to evaluate impact and optimize performance

After implementing changes, conduct a post-change assessment to assess their impact on the resiliency and performance of your EUC services. Monitor key metrics and user feedback to identify any potential areas for optimization or adjustment that may have resulted from the changes. Use

post-change evaluations to inform continuous improvement efforts and refine future change management processes.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

In the post-change evaluation process, establish clear objectives to focus on while you assesse the impact of your changes on resiliency and performance. Continually monitor key metrics such as latency, resource utilization, and system stability to evaluate the effects of changes on your EUC environments accurately.

Additionally, gather feedback from users regarding their experience with the implemented changes to gain valuable insights into usability and functionality. Comparing post-change metrics with baseline data helps identify significant deviations or improvements, facilitating a thorough assessment of change impact.

Ultimately, the insights gained from post-change evaluations should drive continuous improvement in change management processes. By using evaluation results to refine procedures and enhance the resilience and performance of EUC environments over time, organizations can effectively adapt to evolving needs and challenges.

By following these change management practices, organizations can effectively manage changes to their EUC environments, maintain service availability and resiliency, and reduce the risk of service interruptions or incidents that could impact users' access to virtual desktop resources.

Failure management

EUCREL12: How do you address EUC specific issues and requirements in your incident response plan?

Develop and document incident response plans specifically tailored to your Amazon WorkSpaces and AppStream 2.0 environments.

Best practices

- EUCREL12-BP01 Develop an EUC-specific incident response plan that improves reliability in your environment

EUCREL12-BP01 Develop an EUC-specific incident response plan that improves reliability in your environment

When developing incident response plans for Amazon WorkSpaces and AppStream 2.0, it's important to address their unique characteristics such as the session-based nature of AppStream 2.0 and the persistent data in WorkSpaces. Plans should include strategies for handling issues with scaling, session failures, and network dependencies like VPCs or AWS Direct Connect. Active Directory integration is crucial for both services, so steps for troubleshooting authentication failures or AD synchronization must be detailed. The plan should also account for Region-specific outages, using cross-Region backups or failover mechanisms for user data and application availability.

Additionally, document user connectivity issues and regular backups to provide seamless recovery and data protection. Verify that the incident response plans are comprehensive, covering procedures for responding to various types of incidents or events specific to WorkSpaces and AppStream 2.0. Collaborate with key stakeholders in the process to gather insights into potential scenarios and verify alignment with organizational goals. Regularly review and refine these plans to incorporate lessons learned and evolving requirements, maintaining their effectiveness and relevance over time.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

When developing incident response plans for Amazon WorkSpaces and AppStream 2.0, customize them to suit the specific features and challenges posed by these cloud services. Verify that these plans are thorough, encompassing procedures for addressing various incidents or situations specific to WorkSpaces and AppStream 2.0.

Collaborate with key stakeholders to gather valuable insights and align plans with organizational objectives. Document comprehensive procedures, clearly define roles and responsibilities, establish effective communication channels, prioritize incidents based on severity, and set response timelines.

Regularly review and update these plans to incorporate lessons learned and evolving requirements, improving their ongoing effectiveness and relevance. This structured and inclusive approach fosters readiness to respond swiftly and effectively, bolstering overall system reliability and resilience.

Key AWS services

EUC:

- Amazon AppStream 2.0
- Amazon WorkSpaces

Storage:

- Amazon FSx
- Amazon S3

Managed Directories for WorkSpaces:

- AD Connector
- AWS Managed Microsoft AD
- Simple AD
- Cross Trust

Resources

Related documents:

- Amazon WorkSpaces Documentation
- AWS Sign-In endpoints and quotas
- Advancing business continuity with Amazon WorkSpaces Multi-Region Resilience
- Building for business continuity with Amazon WorkSpaces and AWS Directory Services
- Creating custom Amazon CloudWatch dashboards and widgets for Amazon WorkSpaces
- VPC design
- Integration with Microsoft Active Directory
- Introduction to AWS End User Computing Services
- Amazon WorkSpaces Primer
- Amazon WorkSpaces Deep Dive
- Amazon AppStream 2.0 Primer

Related videos:

- AWS - Workspaces

Performance efficiency

The performance efficiency pillar includes the ability to use cloud resources efficiently to meet performance requirements, and to maintain that efficiency as demand changes and technologies evolve. This section provides in-depth, best practice guidance for architecting performant EUC workloads on AWS.

Focus areas

Design principles

- **Minimize latency**: EUC workloads are sensitive to latency. For best performance, minimize the latency between end users and EUC services, as well as between EUC instances and dependencies.
- **Monitor performance metrics**: Use performance metrics to understand the behavior of both individual instances and the holistic health of your EUC environment. Adjust configurations to meet evolving performance requirements.
- **Consider mechanical sympathy**: Understand the design goals of AWS EUC services and features and align them with your workload goals. For further information related to mechanical sympathy, see Consider Mechanical Sympathy.

Architecture selection

EUCPERF01: How do you choose AWS Regions and Availability Zones for your EUC deployments?

Selecting the most appropriate Regions or Availability Zones to deploy your EUC services will be a critical factor to consider to provide the best performance for your end users, partners, and customers.

EUCPERF02: What are the external considerations that affect your choice of regions for EUC deployment?

A well-performing AWS EUC architecture will consider the location of the users accessing the services and the latency to key service endpoints in each Region. Consider the proximity of user data such as home drives, user profile stores, databases, and data feeds to the users to design an efficient data flow. For further information related to tradeoffs to consider in relation to latency as well as how to determine the latency between user locations and the location of AWS EUC services, see EUC latency tradeoffs and How to check latency to the closest AWS Region.

EUCPERF03: How do you improve performance of EUC backend services to meet overall performance goals?

A typical EUC deployment uses many backend services which are deployed, managed and supported by the business. AWS offers a range of managed services which offer resilience and scalability, and which augment the performance of your desktop and application delivery tiers.

Best practices

- EUCPERF01-BP01 Check Regional support for the required EUC services
- EUCPERF01-BP02 Consider the requirements of your Availability Zones when architecting your AWS EUC services

- EUCPERF01-BP03 Consider disaster recovery (DR) requirements when architecting your AWS EUC solution
- EUCPERF02-BP01 Identify geographic distribution of end users and design to minimize latency
- EUCPERF02-BP02 Scale your EUC environment to accommodate the required number of end users
- EUCPERF02-BP03 Evaluate external data sources that your environment integrates with, and assess its impact on performance
- EUCPERF03-BP01 Consider modernization of backend services to use managed services from AWS for best performance

EUCPERF01-BP01 Check Regional support for the required EUC services

Not all AWS regions support EUC services such as AppStream 2.0, WorkSpaces and WorkSpaces Secure Browser.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Check to see if the relevant AWS EUC service is available in your most proximal Region. If the required service is not available in this Region, check to be sure that you can deliver the required performance from the Region closest to you or with lowest latency. For information on EUC Regional support, see:

- WorkSpaces Regional Support
- AppStream 2.0 Regional Support
- WorkSpaces Secure Browser Regional Support

The WorkSpaces Connection Health Checker details the latency between a specific endpoint device and the WorkSpaces service running in each available Region. This data is also a good indicator of latency for WorkSpaces Secure Browser and AppStream 2.0 if they are running in the same Region.

EUCPERF01-BP02 Consider the requirements of your Availability Zones when architecting your AWS EUC services

Within each Region, only select Availability Zones support each AWS EUC service. This is important if you are architecting solutions with extreme performance or security requirements that demand that applications and desktops reside on the same subnet as the user data they need to access.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

For the WorkSpaces service line, explore the Availability Zone information.

- Amazon WorkSpaces Availability Zone Support
- Amazon WorkSpaces Secure Browser

For AppStream 2.0, selecting a subnet when creating a new fleet automatically checks if the associated Availability Zone can support the requested requirements, which are based on several criteria such as instance type and availability.

EUCPERF01-BP03 Consider disaster recovery (DR) requirements when architecting your AWS EUC solution

Will a secondary Region support the latency that is acceptable to support the selected AWS EUC service in a DR scenario, or can you accept degraded performance and relaxed service level agreements to continue to do business?

Level of risk exposed if this best practice is not established: Low

Implementation guidance

For WorkSpaces, the use of cross-Region redirection or Multi-Region Resilience allows the manual or partially automated process of using alternate regions to support your WorkSpaces users in the event of a serious outage.

For AppStream 2.0, the master images created in one Region can be copied to a secondary Region to enable the configuration of identical regional deployment for DR purposes.

Review each of these DR features to be sure that they offer adequate performance and capabilities depending on the Region that is selected for the purpose.

You should also replicate user data and other critical backend services in each Region to provide localized access if similar levels of performance are expected in a DR scenario.

For more detail on Cross-Region redirection and Multi-Region Resilience, see Business continuity for WorkSpaces Personal.

EUCPERF02-BP01 Identify geographic distribution of end users and design to minimize latency

When migrating to or implementing AWS EUC services, consider the location of each group of users with respect to the service endpoints for AWS WorkSpaces, AppStream 2.0, or WorkSpaces Secure Browser. You should deliver services from the Region with the lowest latency to most users.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Capture the location of each user group, and calculate the average latency between each group and their most proximal AWS Region that supports the required AWS EUC service. Due to Regional network routing and capabilities, it is possible the most proximal AWS Region does not necessarily offer the lowest latency.

If you must deploy AWS EUC services in a non-optimal Region (which is sometimes necessary to access other AWS services which have already been deployed), then be sure that you test your application to verify that they offer acceptable performance at the latency levels being experienced.

For an example of how latency might affect the user experience, see EUC latency trade-offs.

EUCPERF02-BP02 Scale your EUC environment to accommodate the required number of end users

The number of users accessing the selected AWS EUC service should not affect the performance of the service itself, as AWS provides both scale and resilience for the components that affect authentication and streaming of user sessions. Many supporting components, however, need to be scaled to support the user numbers you intend to deploy.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Understand the backend requirements for your deployment and scale them accordingly. For example, a WorkSpaces compute instance with 2 vCPU and 4Gb of RAM may offer acceptable performance to run a targeted application set, but if access to user data or an application database backend is compromised by server performance or network constraints, then the user may complain that WorkSpaces is performing badly. Ideally, perform end to end testing for each application set using scalability testing tools to be sure that they will deliver acceptable performance in production as the services scale.

EUCPERF02-BP03 Evaluate external data sources that your environment integrates with, and assess its impact on performance

The location of user data and the services used to deliver access to this data are key to providing the best performance for consumers of an AWS EUC deployment. Latency incurred while accessing data sources may incur additional delays and contribute to end user frustration and lack of engagement, as well as increased support calls.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Define a data architecture that describes how data is managed, from collection through transformation, distribution, and consumption. This informs the EUC architects where to place key application and desktop delivery services and where optimization may be required to avoid performance degradation.

If migrating from an existing on-premises EUC architecture, you may need to deploy AWS Direct Connect or AWS Site-to-Site VPN connections to provide access between AWS and your on-premises infrastructure. For best practices related to networking for Amazon WorkSpaces and descriptions for how and when to use Direct Connect and VPN connections, see Best Practices for VPCs and Networking in Amazon WorkSpaces Deployments.

Be sure to architect network solutions with low enough latency and sufficient bandwidth to support appropriate data access between desktops, applications, and any on-premises data sources.

If your AWS EUC solution integrates with services offered by other cloud providers, such as email, collaboration tools, or SaaS applications, be sure to size internet connections or private networks accordingly to avoid high latency and bandwidth constraints.

EUCPERF03-BP01 Consider modernization of backend services to use managed services from AWS for best performance

By using AWS EUC services, you are already taking advantage of the reduced infrastructure and management overheads of maintaining your own environment. Taking the same approach to other backend services which support the EUC deployment can further increase operational efficiency.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Review the backend services required for your AWS EUC deployment, and determine if transitioning to managed service equivalents from AWS might improve performance, simplify cost modelling, and reduce the administrative and support overheads of delivering these in-house. Examples include:

- **Amazon FSx for Windows File Server**: Resilient, high performance file shares, user data storage and profile management.
- **Amazon RDS**: A range of high-performance managed SQL services.
- **Amazon CloudWatch**: Insight into operational metrics and alerting to maintain performance and efficiency.
- **AWS CloudTrail**: Records logs that provide insight into activities undertaken within an AWS account.

Reduce the overhead of managing your own infrastructure, and invest the time saved to perform continual service improvement, increasing the performance efficiency of your EUC deployment.

Compute and hardware

EUCPERF04: How do you select the most appropriate AWS compute solution?

Following the sizing of the compute requirements for your EUC solution, you can start planning the most appropriate compute instance types to deliver each workload efficiently and cost effectively.

Best practices

- EUCPERF04-BP01 Evaluate available instance types (AppStream) and hardware bundles (WorkSpaces)
- EUCPERF04-BP02 Identify all user types, and deploy required fleet types and instance types as needed
- EUCPERF04-BP03 Determine the running mode and size of hardware bundles needed to support each user type's applications

EUCPERF04-BP01 Evaluate available instance types (AppStream) and hardware bundles (WorkSpaces)

AppStream 2.0 groups instances into families, such as General Purpose (stream.standard). Within each family, there are different instance sizes, such as stream.standard.medium and stream.standard.large. Each size has a different number of vCPUs and memory. Graphics optimized families include instances with one or more GPUs. For more information on the Graphics G4 (stream.graphics.g4dn), Graphics G5 (stream.graphics.g5), and Memory Optimized (stream.memory.z1d) families, see Amazon EC2 Instance Types.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

WorkSpaces bundle selection begins with determining if your workload requires a GPU. If it does, evaluate the Graphics G4 and Graphics G5 families. If it does not require a GPU, evaluate the General Purpose, Compute Optimized, and Memory Optimized families. In addition to large amounts of memory, stream.memory.z1d instances offer the highest CPU clock rates of the AppStream 2.0 instance family.

WorkSpaces provides hardware bundles with different amounts of vCPUs and memory. Graphics.G4dn and GraphicsPro.G4dn bundles include GPUs.

For specifications and recommended uses cases, see Amazon WorkSpaces.

EUCPERF04-BP02 Identify all user types, and deploy required fleet types and instance types as needed

Not all end users necessarily require the same level of performance. Users who perform routine tasks such as data entry, document review, or customer service may need a low level of

performance, while content or video editors, investment and securities traders, or graphics users may require performant desktops. Other users may require moderate levels of performance as their workloads may be unpredictable.

It's important to have a high degree of familiarity with the applications that need to be delivered using Amazon AppStream 2.0 in terms of their compute resource requirements. By understanding core compute requirements such as the amount of memory, CPU, network bandwidth, latency, and disk space that applications require, you can determine the optimum fleet type and instance sizes required for the workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Determine the compute requirements for your applications.

- Assess your users' applications and tasks, and deploy a sufficient level of fleet types and instance types as are needed.
- Monitor the resulting user feedback to verify that performance meets their needs without overprovisioning their instance types.
- If performance or productivity suffers for various users, increase the performance of their instances. This can be achieved by using larger instances with more CPU or in the case of AppStream 2.0 using a different instance family that provides higher clock speed for CPU cores.

EUCPERF04-BP03 Determine the running mode and size of hardware bundles needed to support each user type's applications

It's important to have a high degree of familiarity with the applications that need to be delivered using Amazon WorkSpaces Personal in terms of their compute resource requirements and their usage pattern. By understanding core compute requirements such as the amount of memory, CPU, network bandwidth, latency, and disk space that applications require, you can more effectively determine the optimum WorkSpaces Personal bundle type. The optimal running mode required to support the workload is determined by understanding the pattern of usage of the application.

Implementation guidance

Determine the compute requirements for your applications.

- Assess your users' applications and tasks and deploy a sufficient level of performance as is needed.
- Monitor the resulting user feedback to verify that performance meets their needs without overprovisioning their hardware types.
- If performance or productivity suffers for various users, increase the size of their instances.
- For Personal WorkSpaces, establish the current or required pattern of usage of the applications or desktops being delivered. Select an Always-On running mode for user environments that are broadly used throughout each month (> 80 hours), select the Auto-Stop running mode where usage will be <80 hours per month. Alternatively, consider implementing the Cost Optimizer for Amazon WorkSpaces Solution to automatically select the optimum running mode for each instance.
- Enable self-service WorkSpace management capabilities for your users.

Data management

EUCPERF05: How do you choose storage solutions that maximize performance efficiency in your EUC environment?

Storage solutions can be divided into multiple types: block, object, file, and instance. Each type has different characteristics and is suitable for different use cases. For an overview, see the following resources:

- What is Block Storage?
- What is Object Storage?
- What is File Storage?
- What is Instance Storage?

You may also want to consider third-party cloud storage solutions such as Microsoft OneDrive for Business or Google Drive, which offer granular policy management and help you meet specific security and compliance requirements.

Best practices

- EUCPERF05-BP01 Understand your existing storage requirements, policies, and solutions

- EUCPERF05-BP02 Understand integrated storage capabilities (AppStream)
- EUCPERF05-BP03 Understand integrated storage capabilities (WorkSpaces)
- EUCPERF05-BP04 Use instance storage when available and appropriate
- EUCPERF05-BP05 Consider the benefits of additional AWS storage services

EUCPERF05-BP01 Understand your existing storage requirements, policies, and solutions

If your EUC workload already uses storage volumes, operations policies, and vendor solutions, make sure that you not only understand what products and services they are based on, but also identify the features, advantages, and benefits associated with each in your existing workload. Decide whether these are best suited to your applications and technical goals. Otherwise, develop a set of new functional requirements and solutions that will better address your requirements.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Begin by understanding the storage requirements for each use case. Relevant requirements include the following:

- Individual file size
- Total data size
- Average and peak IOPS
- Whether storage is per-user or shared
- File and folder permissions

Also, consider organizational policies and existing solutions (for example, if policy dictates that users store all files in a central repository).

EUCPERF05-BP02 Understand integrated storage capabilities (AppStream)

For persistent, per-user storage, AppStream 2.0 offers built-in connectors to Amazon S3 home folders, Google Drive for Google Workspace, and OneDrive for Business. For more information on these connectors, see Enable and Administer Persistent Storage for Your AppStream 2.0 Users.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use Amazon S3 home folders when you need a simple, fully-managed solution for persisting user files between sessions and users don't need to access their files from outside their AppStream 2.0 sessions. Use Google Drive for Google Workspaces or OneDrive for Business when you use Windows fleets and your users have a license for one of the services.

If the integrated storage features of Amazon AppStream 2.0 do not offer the capabilities you require, consider Amazon FSx for Windows File Server, Amazon FSx for NetApp ONTAP, or Amazon EC2 hosted file sharing. You can use these fully or partly-managed solutions to store user data or user profiles, such as FSLogix, close to your AWS EUC control plane.

EUCPERF05-BP03 Understand integrated storage capabilities (WorkSpaces)

Most existing workloads, either physical or virtual, will make use of integrated storage that provides the system drive and data drives. For virtualized desktops and servers, this will be virtual drives created from hyperconverged storage. Some workloads, if not already virtualized, may also have fast boot and data drives (like SSD or NVMe) or additional integrated storage in the form of internal hard drives or externally-connected hard drives that deliver large or faster storage for specific applications.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

If any of the workloads you are migrating to AWS EUC services have been configured with and require high performance or additional high-density storage, carefully review the AWS instance types that provide higher performance storage. The Graphics G4 instance types offer a local NVMe instance store which may meet your requirements.

This may also be an opportunity to review alternate networked AWS Storage solutions as they might provide the speed and density you require.

EUCPERF05-BP04 Use instance storage when available and appropriate

An instance store provides temporary block-level storage for your instance. This storage is located on disks that are physically attached to the host computer. Instance store is ideal for temporary

storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content.

For AppStream 2.0, the Graphics G4, Graphics G5, and Memory Optimized (stream.memory.z1d) instance families include NVMe instance storage volumes. For further information related to the instance storage volumes and initializing, see Instance store temporary block storage for EC2 instances.

For WorkSpaces, the graphics.g4dn and GraphicsPro.G4dn bundles provide NVMe instance storage volumes.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use the local instance store on instances that support it to optimize the performance of end user applications. When doing so, consider that the instance store is not backed up and should only be used to satisfy temporary storage requirements. See Local Instance Store for GPU-enabled Bundles for more information.

EUCPERF05-BP05 Consider the benefits of additional AWS storage services

As an alternative to internal storage, some workloads benefit from shared storage for collaboration or to enable persisting data in centralized locations. Using non-internal storage services delivers storage with customizable performance, which gives administrators more control for common storage attributes like IOPS, throughput, and volume size that directly impact performance and user experience.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Review additional storage services if any of the workloads you are migrating to AWS EUC services require tunable performance, larger volume sizes exceeding those provided by the EUC services, or granular control over throughput and IOPs, including Amazon FSx for Windows File Server, Amazon FSx for NetApp ONTAP , and Amazon EFS.

For more information, see Persistent storage for Amazon AppStream 2.0 Linux Fleets on Amazon Elastic File System and Connect Amazon FSx for NetApp ONTAP to Amazon AppStream 2.0 Linux instances.

Networking and content delivery

The optimal networking solution for a workload varies based on latency, throughput requirements, jitter, and bandwidth. Physical constraints, such as the location of users and data sources, will affect the performance of the solution.

EUCPERF06: How do you configure network connectivity in your EUC environment for best performance?

In an EUC architecture, there are two key networking configurations to consider:

- Connections from end users to their most proximal AWS EUC service connection point
- Connections from the EUC instances to any backend infrastructure and other services

Best practices

- EUCPERF06-BP01 Minimize latency between end users and EUC services
- EUCPERF06-BP02 Minimize latency between EUC instances and dependent services
- EUCPERF06-BP03 Make sure that EUC network configurations don't interfere with service management connections

EUCPERF06-BP01 Minimize latency between end users and EUC services

Like many other vendors, AWS EUC solutions deliver their services using a remote display protocol to stream the pixel information to the endpoint device, which is highly efficient and capable of tolerating a variety of network conditions. Low latency, low packet loss, and jitter are key to delivering the best service for end users.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Minimize latency between end user devices (like desktops, laptops, and thin clients) and the AWS EUC service endpoints by avoiding proxies, inspection appliances, and VPNs.

Determine whether there are any conditions which might introduce latency between your end users and the AWS EUC service endpoints. Test connectivity under various conditions to identify the maximum latency that can be tolerated by the application set being deployed, and verify that your network can scale to reliably deliver the number of users being deployed.

If end users will be working from home, try to establish a minimum level of network connectivity that should provide a good user experience. Most home broadband connections are more than capable of delivering low latency for home working, but problems with home networks can be difficult to diagnose.

Verify that endpoint devices can run the local client application (WorkSpaces or AppStream Client) that processes and displays the encrypted pixel stream which flows between the end user and the AWS EUC service connection points (streaming gateways). If the workload delivers collaboration tools such as Microsoft TEAMs, Zoom, or Webex, optimization capabilities will try to offload processing to the local endpoint device, which must be capable of handling this additional load.

EUCPERF06-BP02 Minimize latency between EUC instances and dependent services

In most cases, EUC users require connections to resources outside their EUC instances. Common dependencies include web or application servers, database servers, and storage services.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

When possible, deploy these dependencies in the same AWS Region and ideally the same Availability Zone. If the system of record must reside elsewhere, consider deploying caches or replicas. For example, if your Active Directory domain controllers are on your on-premises network, deploy replicas on Amazon EC2.

When connecting to Amazon S3, use gateway VPC endpoints. For more information on configuring gateway endpoints, see Gateway endpoints for Amazon S3.

EUCPERF06-BP03 Make sure that EUC network configurations don't interfere with service management connections

AppStream 2.0 instances use a dedicated management network interface (eth0) for streaming and service management connections.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Do not configure applications or the operating system to interfere with the connections listed in Amazon AppStream 2.0 Connections to Your VPC. If private network connectivity from AppStream 2.0 instances to resources outside your VPC is required, use a VPC-level solution such as AWS Site-to-Site VPN or AWS Transit Gateway. Do not use a client VPN on the AppStream 2.0 instance, as this is complex and error-prone to configure properly.

WorkSpaces instances use a dedicated management network interface (eth0) for streaming and service management connections.

Process and culture

When architecting an AWS EUC solution, there are principles and practices that you can adopt to help you better run an efficient, high-performing environment. This focus area offers best practices to help adopt a culture that delivers and maintains a level of performance that both meet and exceed the requirements of the business.

EUCPERF07: How do you test performance of your EUC environment?

Validate your workload performance against real-life testing conditions based on your end users' use cases, performance metrics, and constraints within their workloads.

EUCPERF08: How do you monitor performance and availability in your EUC environment?

Unless you can accurately measure your resource performance against your key performance indicators and functional requirements, you cannot be certain that the workload will meet your objectives consistently. Establish a set of baseline requirements. Once established, routinely measure performance to understand how your workload's performance may vary against external demands.

EUCPERF09: How do you evolve your workload to benefit from new feature releases?

New releases are inevitable, and updates can change the performance characteristics of a system. Understand what changes are being made, and develop strategies for taking advantage of performance increases and resource consumption due to those changes.

EUCPERF10: How do you manage ongoing and on-demand sizing changes in your EUC environment?

Consider the growth of performance needs over time. Also consider cyclical events such as end-of-month or end-of-quarter reporting, annual updates, and holiday demands.

Best practices

- EUCPERF07-BP01 Conduct realistic end-to-end testing aligned with organizational objectives
- EUCPERF08-BP01 Establish and monitor service metrics and KPIs
- EUCPERF08-BP02 Monitor Amazon AppStream 2.0 CloudWatch metrics
- EUCPERF08-BP03 Monitor Amazon WorkSpaces Personal CloudWatch metrics
- EUCPERF08-BP04 Monitor operating system metrics
- EUCPERF09-BP01 Follow AWS EUC news sources
- EUCPERF10-BP01 Align the instance type and instance size of a fleet with the workload
- EUCPERF10-BP02 Enable self-service WorkSpaces Personal management capabilities, and allow users to request changes by an administrator
- EUCPERF10-BP03 Install only the application features required by end users
- EUCPERF10-BP04 Remove caches, temporary data, log files, and unneeded files such as tutorials and sample data before creating an image
- EUCPERF10-BP05 Tune application performance where possible to optimize compute resource usage

EUCPERF07-BP01 Conduct realistic end-to-end testing aligned with organizational objectives

When planning to conduct testing, consider how your users interact with the EUC service and on an everyday basis. Create tests that align with the primary use of the service initially and expand to edge cases over time or in response to incidents to verify that they do not arise in future iterations of the service.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Conduct tests that align with the expected use of the service. Work backwards from organizational objectives to conduct realistic tests. For example, consider any use case where remote users process invoices in an accounting application. Key metrics may include the number of invoices that each user processes per hour and their accuracy. A realistic test would include experienced application users processing actual invoices, using representative client devices under typical network conditions.

EUCPERF08-BP01 Establish and monitor service metrics and KPIs

When using an AWS EUC service to deliver a service to your users, it's important to consider the service metrics that are key to the delivery of the service for your organization to verify that the service is operating at the required service levels.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Determine service metrics and KPIs for your service. Some examples of key measures to consider are:

- Service availability
- Mean time to repair (MTTR)
- First call resolution (FCR)
- SLA breach rate
- User and customer satisfaction (CSAT)
- Cost per contact
- Net promoter score
- Incident volume
- Problem resolution time

Consider how metrics available within the AWS EUC services outlined in the following sections can be used to support or determine your service metrics.

EUCPERF08-BP02 Monitor Amazon AppStream 2.0 CloudWatch metrics

Use Amazon CloudWatch to establish and monitor your AppStream 2.0 workload's performance against the KPIs established for your service. Use the Automatic dashboard in Amazon CloudWatch to monitor your fleet capacity over time or consider creating a custom dashboard tailored to your environment.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Measure your workload's performance across Amazon AppStream 2.0 fleets and fleet instances.

EUCPERF08-BP03 Monitor Amazon WorkSpaces Personal CloudWatch metrics

Use CloudWatch to establish and monitor your Amazon WorkSpaces workload's performance against these KPIs and requirements.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Monitor the Amazon WorkSpaces service and instances using Amazon CloudWatch. Use the guidance provided in the following articles to measure your workload's performance.

- Monitor your WorkSpaces using CloudWatch metrics
- Creating custom Amazon CloudWatch dashboards and widgets for Amazon WorkSpaces
- Monitor your WorkSpaces health using the CloudWatch automatic dashboard
- Utilizing CloudWatch Internet Monitor with Amazon WorkSpaces

EUCPERF08-BP04 Monitor operating system metrics

Operating systems can add significant variations in performance to your Workload depending on the compute, storage, and memory resources required. Test with all operating systems that are intended to be supported by your deployment.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Monitor the performance of instances delivering end user services.

- Use operating system metrics such as Windows Performance Counters for detailed insight into instance performance.
- Use the EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces.
- For ongoing monitoring and analysis, consider using the Amazon Kinesis Agent for Windows to monitor Windows Performance Counters for performance trend analysis of key system metrics.

EUCPERF09-BP01 Follow AWS EUC news sources

Many customers can benefit from keeping up with news from software publishers and partners in the end user computing domain. Stay updated on developments by following news feeds and social media.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Subscribe to AWS feeds and blogs to keep up to date. Follow the Desktop and Application Streaming blog and End User Computing What's New Feed.

EUCPERF10-BP01 Align the instance type and instance size of a fleet with the workload

As needed, user environments can be updated on a pre-determined schedule or in response to periodic changes in performance to satisfy a change in the anticipated demand for resources.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Determine the optimal instance family and size for your applications.

- The non-graphics instance families can utilize the same image across them. This provides image portability across these instance families and the instance sizes associated with them and allows varying requirements for compute resources to be catered for.

- Images created for a graphics instance family (for example, stream.graphics.g5) can only be associated with that family due to the specific GPU drivers for the associated GPU. Consequently, choose a graphics instance family carefully from the outset to avoid the need to create a new image for a different GPU family.

EUCPERF10-BP02 Enable self-service WorkSpaces Personal management capabilities, and allow users to request changes by an administrator

The WorkSpaces Personal self-service options allow users to ramp up or down instance performance over time.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Enable user self-service where possible to optimize processes.

- Identify the most flexible compute types for users to anticipate required changes in performance. Consider the following:
  - You can change the compute type from Graphics.g4dn to GraphicsPro.g4dn, or from GraphicsPro.g4dn to Graphics.g4dn.
  - However, you cannot change the compute type of Graphics.g4dn and GraphicsPro.g4dn to other types.
  - You cannot change the compute type of Graphics and GraphicsPro to another type.
- Consider these capabilities and limitations when initially configuring your users' environments.

EUCPERF10-BP03 Install only the application features required by end users

Some applications provide the ability to tailor an installation to remove features that are not required by users.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Do not install application features not required by users. Install the minimal set of features in applications that are required by users to perform their roles. This helps to reduce compute requirements and also helps to remove potential security risks that may arise that are associated with those features.

EUCPERF10-BP04 Remove caches, temporary data, log files, and unneeded files such as tutorials and sample data before creating an image

Remove non-required files that are installed, downloaded, or created by applications to optimize storage consumption.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Remove unneeded files from images to optimize storage consumption.

Unnecessary files included in an Amazon WorkSpaces golden image use space for each WorkSpace provisioned using that image. Similarly, for Amazon AppStream 2.0 where the image builder volume size is limited, removing unneeded files can provide additional storage space for other applications.

Consider data access patterns and whether data not included in an image can be downloaded when needed. For example, if 10% of users access an application library that can be downloaded when needed, omit the library from images.

EUCPERF10-BP05 Tune application performance where possible to optimize compute resource usage

To provide the optimal access to compute resource for your applications, consider tuning the performance of applications or software where possible to reduce their compute resource utilization.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

By reducing the compute resource utilization for software used to provide non-end user facing functionality, such as security agents, additional resources are made available to benefit the applications users interact with. The disabling of non-essential functionality within software can yield a performance benefit for end user software.

Key AWS services

EUC:

- Amazon AppStream 2.0
- Amazon WorkSpaces

Storage:

- Amazon FSx
- Amazon S3

Monitoring:

- Amazon CloudWatch

Resources

Related documents:

- **AppStream 2.0**
  - Create an Image That Always Uses the Latest Version of the AppStream 2.0 Agent
  - Fleet auto scaling strategies
  - Best practices for scaling policy design
- **EC2:**
  - Instance store volume limits for Amazon EC2 instances
- **EUC:**
  - Desktop and Application Streaming Blog

  - EC2 Instance Types
- **WorkSpaces:**
  - Monitor your WorkSpaces using CloudWatch metrics
  - Enable self-service WorkSpaces management capabilities for your users in WorkSpaces Personal

Related partner solutions:

- ControlUp
- eG Innovations
- Liquidware

Cost optimization

The cost optimization pillar of the EUC Well-Architected Lens includes the continuous process of refinement and improvement of an AWS EUC environment over its entire lifecycle. Cost optimization is a key effort, from the initial design of your first proof of concept, to the ongoing operation of production workloads. It enables building and operating cost-aware AWS EUC environments that minimize costs, maximize return on investment, and achieve business outcomes. Best practice focus areas include: Cloud Financial Management, expenditure and usage awareness, resource cost-effectiveness, resource demand and supply management, and optimization.

Focus areas

Design principles

The following design principles are additional considerations over and above the Well-Architected Framework cost optimization design principles.

- **Establish ownership of cost optimization for your EUC services:** EUC services are typically consumed by end-users, which is why consumption and cost is highly dependent on user behavior, usage patterns, self-service capabilities, application requirements, and other factors. Over and above knowledge of the cloud financial management, cost optimization for EUC services requires knowledge of the specific EUC services in use and how EUC services generally work to understand the levers they have to optimize cost. This includes Operating System choices, application workloads and their hardware requirements, licensing, storage requirements, desktop & application management, and other areas.

- **Govern self-service capabilities of EUC services:** With EUC services you have the option to provide your users with certain self-service capabilities, which gives them flexibility to adjust their environment (e.g. CPU, RAM, Disk) according to their requirements. However, some of these self-service capabilities have an impact on your cost. With choice comes responsibility. If you leave the choice to your end-users, you should include the changes made via self-service capabilities in your cost & usage reporting so you can react accordingly.
- **Evaluate cost when selecting AWS EUC services:** AWS offers different EUC services that lend themselves to different use cases. It is important to evaluate the application landscape, understand usage patterns, and understand hardware requirements to map the specific workload to a suitable EUC service. For example, in some cases a monthly billing model may be more cost-effective, while in other cases billing by the hour or even by the second may be the better choice. With different use cases and personas, it is common practice to make use of a mix of services & billing models as appropriate for the given use case. It is therefore important to - ideally upfront - gather data on usage, usage patterns and resource utilization to select the most appropriate service for a given workload.
- **Use existing licenses for cost optimization when appropriate:** Whilst we give customers choice on the Operating System they consume, most customers deploy Microsoft Windows Operating Systems. Depending on your existing licensing with Microsoft, you may be able to use existing M365 or RDS CAL licenses with certain AWS EUC services. It is highly recommended you assess your eligibility to bring your own licenses upfront, since this may allow you to reduce the cost of your AWS EUC service. Consult the Amazon WorkSpaces FAQs on Windows BYOL, the Amazon AppStream 2.0 FAQs on Pricing and Billing, and the Microsoft Licensing on AWS guide for further detail, and contact microsoft@amazon.com.

Practice Cloud Financial Management

EUCCOST01: How do you establish ownership of cost optimization for your EUC services?

AWS EUC services represent a workload that may require dedicated resources for cost optimization, since these services are typically consumed by end-users, and the actual consumption can vary significantly based on user behavior, seasonality, self-service capabilities, and other factors.

Best practices

- EUCCOST01-BP01 Evaluate EUC specific cost model awareness in your cloud business

- EUCCOST01-BP02 Increase awareness of the EUC cost model in your cloud business office to promote cost optimization

EUCCOST01-BP01 Evaluate EUC specific cost model awareness in your cloud business

You may have a cloud business office, Cloud Center of Excellence, or a FinOps team that is responsible for establishing and maintaining cost awareness across your organization. However, AWS EUC services often use many other services and may require a solid understanding of Microsoft licensing to optimize the cost. If you are using the ITIL framework, you may already have defined service owners for individual services who own the financials for their services.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Evaluate the required domain knowledge in your cloud business office. We recommend you evaluate the required EUC domain knowledge in your cloud business office to understand if the team is ready to support cost optimization for EUC services. These individuals need to be intimately familiar with the cost optimization levers specific to EUC services, such as Microsoft licensing, WorkSpaces running modes, WorkSpaces bundles, AppStream 2.0 Fleet types, and AppStream 2.0instances. Provide EUC-specific cloud financial training to them if there is a knowledge gap.

EUCCOST01-BP02 Increase awareness of the EUC cost model in your cloud business office to promote cost optimization

In case your cloud business office lacks the required knowledge in the EUC domain, consider additional training and enablement to close these gaps.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

If an existing team is tasked with cost optimization of AWS EUC services, train these individuals on the specifics of the EUC services so they can perform their duties. This includes training on the services, generic training on how a typical EUC environment (virtual or physical) operates, and specifically how to identify resource under- and over-provisioning.

Expenditure and usage awareness

EUCCOST02: How do you monitor cost and utilization of your EUC services?

While AWS Cost and Usage Reports include insights into AWS EUC services, some of the AWS EUC services offer additional reporting capabilities.

EUCCOST03: How do you govern usage of your EUC services?

Enabling your end users to provision or manage AWS EUC resources can help reduce the cost if the users are educated on the cost impact of the decisions they make. Determine the level of self-service capabilities you want to give your end users and assess the impact of this (best and worst case) on your cost. Depending on the outcome of this assessment, you may decide to run certain self-service capabilities through an approval workflow before implementing them. You can configure AppStream 2.0 to implement scaling policies that dynamically and automatically provision resources according to demand. AWS Cost and Usage Reports provide insights into EUC service usage and costs, and some AWS services offer additional reporting capabilities.

Best practices

- EUCCOST02-BP01 Monitor your EUC cost and usage proactively
- EUCCOST03-BP01 Determine the level of self-service capabilities to provide your users
- EUCCOST03-BP02 Use a self-service portal to request your ITSM

EUCCOST02-BP01 Monitor your EUC cost and usage proactively

AWS Cost and Usage Reports help you gain detailed insights onboth your AppStream 2.0 and your WorkSpaces service usage and cost. In addition, AppStream 2.0 offers separate Usage Reports with further detail. Amazon WorkSpaces comes with a WorkSpaces CloudWatch automatic dashboard that provides insight into the performance of your WorkSpaces resources and helps you identify performance issues. Amazon AppStream 2.0 Fleet Usage and Instance/Session Performance Metrics are available in the AppStream 2.0 Console and Amazon CloudWatch.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

For WorkSpaces, enable AWS Cost and Usage Reports with resource IDs to analyze and visualize your cost and usage. Resource IDs help you see the cost and usage data for an individual WorkSpace. Consider building an Build an enterprise cost and usage dashboard for Amazon WorkSpaces.

Furthermore, the Cloud Intelligence Dashboards section of AWS Well-Architected Labs explores how to build a CUDOS Dashboard that includes Amazon WorkSpaces cost and usage data. The Cost Optimizer for Amazon WorkSpaces referred to in EUCCOST-BP05 also generates basic usage reports in Amazon S3.

AppStream 2.0 also offers built-in usage reports. Enable AppStream 2.0 Usage Reports to gain valuable insights into your AppStream 2.0 usage. For details on visualizing your AppStream 2.0 usage, see Analyze your AppStream 2.0 usage reports using Amazon Athena and Quick Suite. If you are using Amazon AppStream 2.0 features such as Enable Application Settings Persistence for Your AppStream 2.0 Users or Enable and Administer Home Folders for Your AppStream 2.0 Users , include the underlying Amazon S3 buckets in your cost and usage monitoring.

EUCCOST03-BP01 Determine the level of self-service capabilities to provide your users

Amazon WorkSpaces offers self-service capabilities that you can enable for your users. Assess the impact of granting access to these self-service capabilities and selectively disable or enable them based on your requirements.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Evaluate the cost impact of enabling certain self-service WorkSpaces management capabilities for your users, and then select which of these self-service capabilities you want to provide to your users. For more information, see Enable self-service WorkSpaces management capabilities for your users in WorkSpaces Personal . Consider creating internal policies to govern which capabilities are allowed. Changing the compute type (bundle), increasing the root and user volume size, and changing the running mode may increase your cost. Instead of enabling these capabilities for your users, you may consider providing these capabilities through your IT service management so that changes requested by a user requires prior approval.

EUCCOST03-BP02 Use a self-service portal to request your ITSM

Instead of enabling self-service capabilities for your users, use a self-service portal to allow users to request resources, or enable workflow-based request for EUC resources with your ITSM. This gives you better control and limits the exposure to unforeseen cost increases.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Consider implementing a self-service portal to request WorkSpaces or changes to your existing WorkSpaces, like running mode, bundle, or storage. A self-service portal can allow your users to provision and terminate their EUC services as required. For an example for Amazon WorkSpaces, see Creating a self-service portal for Amazon WorkSpaces end users.

Additionally, consider using your ITSM solution to enable workflow-based requests for new WorkSpaces or changes to existing WorkSpaces, like running mode, bundle, or storage. For examples of integrating with ServiceNow, see How to enable self-service Amazon WorkSpaces by using Service Catalog Connector for ServiceNow and Managing Amazon WorkSpaces by integrating Service Catalog with ServiceNow.

Cost effective resources

EUCCOST04: How do you allocate cost to your business owners?

Since AWS EUC services are typically consumed by end users, we recommend allocating the cost to your business owners at an organizational, departmental, functional or other level. Tag your AWS EUC resources to apply this cost allocation per your requirements. For information on building a tagging strategy, see Best Practices for Tagging AWS Resources. Prepare your tagging strategy ahead of deploying your EUC resources to assist with cost allocation and reporting.

EUCCOST05: How do you evaluate cost and measure efficiency when selecting EUC services?

Usage patterns and hardware requirements for your applications may vary significantly, which is why a one-size-fits-all approach often does not work well for an EUC environment. It is

therefore important to evaluate your application landscape, understand the usage patterns of individual applications, and understand their hardware requirements. This understanding helps you compartmentalize your application landscape and select the most appropriate AWS EUC service for your applications based on their hardware requirements and usage pattern.

Best practices

- EUCCOST04-BP01 Tag your Amazon WorkSpaces and Amazon AppStream 2.0 resources
- EUCCOST05-BP01 Gather usage data and hardware requirements in your existing environment
- EUCCOST05-BP02 Select the most cost-effective service for your EUC workload
- EUCCOST05-BP03 Rightsize your EUC resources
- EUCCOST05-BP04 Choose an appropriate running mode for your EUC workload where applicable

EUCCOST04-BP01 Tag your Amazon WorkSpaces and Amazon AppStream 2.0 resources

Tagging your Amazon AppStream 2.0 resources or tagging WorkSpaces resources helps you allocate your cost to logical groups, such as departments or business entities.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Plan your tagging strategy before you start deploying your EUC resources. You may think of tagging EUC resources with information such as cost center, department, usernames, projects, location, or deployment types (like development, test, and production). The more dimensions you add with your tags, the easier it will be to report and break down the cost once you are in production.

If you already use tagging in your organization, implement a standardized approach for tagging that aligns with the approach being used by the rest of the organization, which results in a standardized format for the key value pairs being used for tags in the organization. Using Service control policies (SCPs) with AWS Organizations enforces tags to restrict resource creation unless they are correctly tagged.

EUCCOST05-BP01 Gather usage data and hardware requirements in your existing environment

Before selecting a service for your EUC workload, gather usage data in your existing EUC environment. Collect data in different areas, like usage patterns and resource utilization. Usage patterns portray how intensively your applications are being used (for example, hours per day and days per week). Resource utilization details how efficiently your compute resources are being used by these applications (like CPU, RAM, GPU, disk space, and disk IO). Both areas help you select the optimal service for a given application or set of applications. You can gather this data using OS or third-party tools.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

If you use a desktop virtualization Environment, your VDI solution may include reporting tools that can provide you with the required data. Tools like Citrix Director or VMware vRealize Operations Manager can be used for this.

Alternatively, you may use scripting to wrap application launches and log the usage of applications using these scripts in a file or database that you can use later to analyze the data. Your OS may include tools to visualize and log the resource utilization of your applications.

For example, Windows offers the Windows Performance Monitor to capture performance metrics over an elapsed period of time.

If you do not have any tools available to gather usage patterns, you can conduct a survey with a representative selection of users to understand their usage of your applications.

EUCCOST05-BP02 Select the most cost-effective service for your EUC workload

Invest time into planning your EUC deployment. A persistent Amazon WorkSpaces, for example, is a desktop as a service assigned to a named user. If this named user needs to run a certain resource-intensive application only occasionally, it is not recommended to over-provision the hardware resources for this WorkSpace to meet the application requirements, as these resources will be under-utilized most of the time. Instead, consider deploying this application to an Amazon AppStream 2.0 fleet, where you have a more granular choice of instance types and are charged for the actual usage only per hour or even per second.

The usage patterns and usage data collected help you govern your application landscape and select the most appropriate service and bundle and instance for each of your applications.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Amazon WorkSpaces offers a variety of different bundles to choose from, and each one has a different hardware configuration (vCPU and RAM), some of which supporting a GPU. In total, you have the choice between five non-GPU bundles and four GPU-enabled bundles.

With Amazon AppStream 2.0, you have a more granular choice from many non-GPU and GPU-enabled instance types. Review your application workloads and match them to the most appropriate service and bundle or instance type to avoid over-provisioning of resources.

Consider Amazon AppStream 2.0 with appropriate instance types for workloads that can be characterized as CPU-intensive or RAM-intensive or requires a GPU and that typically shows a lower utilization.

In a typical EUC environment, users are often using certain applications permanently over the course of a day and other applications only occasionally. For a CPU-intensive or RAM-intensive workload, or for applications requiring a GPU, Amazon AppStream 2.0 can be the more cost-effective solution, especially if the application is only used occasionally. If you have any usage data (usage patterns) on these applications, we recommend you review these and calculate a cost estimate of the usage on Amazon AppStream 2.0using these usage patterns. This helps you understand if provisioning the application on Amazon AppStream 2.0 will be more cost-effective than provisioning it on Amazon WorkSpaces if choosing a more powerful bundle.

Even the combined usage of a less powerful WorkSpaces instance for standard applications and AppStream 2.0 for more demanding workloads can come at a lower cost compared to a more powerful WorkSpaces bundle as the only service. If there isn't enough data to make a decisive decision, identify a mechanism to capture this data in your existing environment or perform a proof of concept (PoC) to capture this data.

If your users only need to access web-based applications, consider using Amazon WorkSpaces Secure Browser. Examples of web-based applications are Salesforce, SAP-Fiori, Confluence, or your intranet websites. WorkSpaces Secure Browser service is a low cost, fully-managed, Linux-based service designed to provide secure browser access to internal websites SaaS applications for up to 200 streaming hours.

If you need a persistent environment with users who require a high degree of flexibility in customizing their environment and installing their own applications, Amazon WorkSpaces Personal is your best option. As opposed to Amazon WorkSpaces Personal, Amazon AppStream 2.0 is not designed to allow users to install their own software due to the non-persistent nature of the AppStream 2.0 fleet.

EUCCOST05-BP03 Rightsize your EUC resources

Choosing the right Amazon WorkSpaces bundle or Amazon AppStream 2.0 instance type for your EUC workloads is important to operate your EUC environment in a cost-effective manner. The chosen configuration needs to support the hardware requirements of your applications, while at the same time avoiding over-provisioning resources.

Capture metrics in an existing reference environment (physical machines or virtual desktops) to understand how the existing resources are being used. This data helps you choose the right bundles and instance types with AWS EUC services. To capture these metrics, use tools like Microsoft Performance Monitor or third-party solutions like Liquidware Stratusphere UX and Control-Up DX solutions.

Once your workload is in production, continually monitor relevant metrics, helping you react to changing requirements by adjusting the bundle and instance type. Monitor your WorkSpaces health using the WorkSpaces CloudWatch automatic dashboard, which provides insight into the performance of your WorkSpaces resources and helps you identify performance issues. Amazon AppStream 2.0 fleet usage, instance, and session Performance Metrics are available in the AppStream 2.0 console and Amazon CloudWatch.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

AWS EUC services offer a variety of different bundles and instance types, including GPU-enabled choices. Assuming you have captured and analyzed your metrics in an existing reference environment, you can map your workloads to the most cost-effective Amazon WorkSpaces or AppStream 2.0 bundles and instance types. If you have use cases that require a GPU and are heavily utilized (high number of hours per month), consider using AppStream 2.0, which gives you a more granular choice of GPU-enabled instances. Use the AWS Pricing Calculator or the Amazon AppStream 2.0 Pricing tool to determine which of the two solutions is more cost-effective for your specific workload.

EUCCOST05-BP04 Choose an appropriate running mode for your EUC workload where applicable

Amazon WorkSpaces can be used with monthly and hourly pricing, while Amazon AppStream 2.0 supports Always-On, On-Demand, and Elastic fleets. Choosing an appropriate running mode can significantly impact the cost of your EUC services. Historical usage data (usage patterns) of a reference environment can help you assess which running mode to use for your EUC workloads.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

When you use Amazon WorkSpaces, you can choose between Always-On and On-Demand running modes, which translate into monthly and hourly billing respectively. For the non-GPU bundles, there is a breakeven point at roughly 80 hours of usage per month, at which point the Always-On WorkSpace will be more cost-effective. If your users use their WorkSpace for less than 80 hours per month, the On-Demand running mode is usually the more cost-effective model for non-GPU bundles.

You can deploy the Cost Optimizer for Amazon WorkSpaces to get reports with recommendations on which running mode to select for your WorkSpaces and automatically convert your WorkSpaces to the most cost-effective running mode. For the GPU bundles, the breakeven point varies from bundle to bundle. The Amazon WorkSpaces Pricing page helps you calculate the breakeven point for these bundles.

Amazon AppStream 2.0 offers three different fleet types: Always-On, On-Demand, and Elastic. Explore the fleet types to determine the right balance between cost-effective operation and desired user experience.

- With Always-On fleets, your fleet instances will constantly be running while the fleet is in a started state, and you'll be charged the respective instance fee per hour per instance in your fleet.
- On-Demand fleets have those fleet instances not in use in a stopped state, for which you'll be charged the lower stopped instance fee per hour per stopped instance in your fleet.
  - This can make a significant difference to your cost, especially when your fleet instances are higher-end instances.
  - However, using On-Demand fleets will prolong the logon time by up to 120 seconds.

- Both Always-On and On-Demand fleet instances are charged on one-hour increments, while Elastic Fleet instances are charged on one second increments, with a minimum of 15 minutes.
- As opposed to Always-On and On-Demand, Elastic fleets do not require you to manage scaling policies and provision buffer capacity, since the pool of Instances in an Elastic fleet is managed by AppStream 2.0.

Amazon AppStream 2.0 offers multi-session fleets, which allow multiple users to use a single AppStream 2.0 fleet instance. Depending on the user density you can achieve on a given instance, you may be able to further optimize your AppStream 2.0 costs compared to a single-session fleet. If you plan to use multi-session fleets, consider resource requirements, instance specifications, and user behavior. For specific guidance, see Multi-Session Recommendations .

Manage demand and supply resources

EUCCOST06: How do you optimize cost using existing licenses when appropriate?

You may already have existing license agreements with Microsoft in place. You can use these licenses with AWS EUC services to reduce your cost.

EUCCOST07: How do you track and identify idle resources to avoid unnecessary charges?

Amazon WorkSpaces can be used in AlwaysOn and AutoStop running mode, which correspond to monthly and hourly billing respectively. If deployed your WorkSpaces with monthly billing but are using them less than expected, switching the billing mode can reduce your cost.

With Amazon AppStream 2.0, you will likely use app block builders or image builders to generate app blocks and images. These resources are charged hourly or in one second increments with a 15-minute minimum if you keep them running.

Best practices

- EUCCOST06-BP01 Explore a bring your own license (BYOL) approach
- EUCCOST07-BP01 Use the available cost optimizers for Amazon WorkSpaces and Amazon AppStream 2.0

EUCCOST06-BP01 Explore a bring your own license (BYOL) approach

If you already have suitable license agreements with Microsoft in place for Operating Systems or Microsoft Remote Desktop Client Access Licenses, consider bringing these licenses for use with AWS EUC services to reduce the cost.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

If you meet the requirements stated in Amazon WorkSpaces FAQs , Amazon WorkSpaces allows you to Bring Your Own Windows desktop licenses in WorkSpaces (BYOL) for Windows 10 and 11. This can reduce your monthly or hourly WorkSpaces charges. When calculating the TCO, consider that BYOL requires a certain minimum commitment of WorkSpaces per AWS Region that you want to deploy in.

When using Amazon AppStream 2.0, you'll be charged a monthly user fee in the form of a Microsoft RDS SAL fee. For more information , see Amazon AppStream 2.0 pricing. If you have Microsoft License Mobility, you may be eligible to bring your own Microsoft RDS Client Access License (CAL) licenses and use them with Amazon AppStream 2.0. For users covered by your own licenses, you won't incur monthly AppStream 2.0 user fees.

EUCCOST07-BP01 Use the available cost optimizers for Amazon WorkSpaces and Amazon AppStream 2.0

Leverage available tools from AWS and partners to support you with cost monitoring and optimization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Amazon AppStream 2.0 uses app block and image builders that are charged hourly or in one second increments with a 15-minute minimum if you keep them running. You must explicitly stop them to stop the billing. The Cost Optimizer for Amazon AppStream 2.0 can monitor your AppStream 2.0 app block and image builders and notify you or stop them when they are active for longer than specified thresholds.

Third-party tools like the AppStream Optimizer by Cambrian Technologies use machine learning to optimize your AppStream 2.0 Fleets and achieve a better utilization. This helps reduce your cost by reducing idle capacity.

Deploy the Cost Optimizer for Amazon WorkSpaces to receive reports with recommendations on which running mode to select for your WorkSpaces and automatically convert your WorkSpaces to the most cost-effective running mode.

Optimize over time

EUCCOST08: How do you optimize your EUC service consumption over time?

Once you've gone into production with your AWS EUC services, continually monitor your environment to react to changing requirements. New operating systems or application releases may be more demanding and require you to choose another bundle or instance type. Organizational changes or changes in usage patterns can require adjustments to your scaling policies with a service like Amazon AppStream 2.0. Identify and terminate unused resources.

Best practices

- EUCCOST08-BP01 Monitor your Amazon WorkSpaces usage, and implement the Cost Optimizer for Amazon WorkSpaces
- EUCCOST08-BP02 Monitor your Amazon AppStream 2.0 fleet utilization, and optimize scaling policies and buffer capacity

EUCCOST08-BP01 Monitor your Amazon WorkSpaces usage, and implement the Cost Optimizer for Amazon WorkSpaces

The Cost Optimizer for Amazon WorkSpaces generates reports you can use to understand the usage of individual WorkSpaces. Based on these reports, identify underutilized WorkSpaces or WorkSpaces that are no longer in use so that you can assess whether to terminate them.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Deploy the Cost Optimizer for Amazon WorkSpaces, and perform regular reviews of your WorkSpaces usage reported by the Cost Optimizer for Amazon WorkSpaces. Based on your findings, decide which WorkSpaces to terminate, and initiate a conversation with owners of underutilized WorkSpaces to understand if these are still needed. Agree on how, when, and by whom any changes are to be applied.

EUCCOST08-BP02 Monitor your Amazon AppStream 2.0 fleet utilization, and optimize scaling policies and buffer capacity

Use Amazon CloudWatch to observe and monitor your Amazon AppStream 2.0 resources. Amazon AppStream 2.0 publishes several AppStream 2.0 Metrics and Dimensions to Amazon CloudWatch that you can visualize and use to check if you are overprovisioning buffer capacity or if you are running into capacity shortages at times. Use these metrics to adjust your AppStream 2.0 Fleet capacity and scaling policies to minimize idle capacity and reduce insufficient capacity errors where possible.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Create your own customized CloudWatch dashboards to visualize key AppStream 2.0 metrics for your AppStream 2.0 fleets. These dashboards can contain several widgets that display a view of selected metrics of a specific AppStream 2.0 fleet or across multiple AppStream 2.0 fleets. Review these dashboards on a regular basis.

Additionally, use the EUC Toolkit to review Amazon CloudWatch and OS-level metrics. This Toolkit also helps you manage large WorkSpaces and AppStream 2.0 deployments at scale. After review of the metrics, determine whether changes to the fleet capacity or scaling policies are required, and plan for how to implement those changes. For more information, see Use the EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces

Key AWS services

- Amazon AppStream 2.0
- Amazon WorkSpaces
- Amazon WorkSpaces Secure Browser

- Amazon WorkSpaces Core
- Amazon CloudWatch
- AWS Cost & Usage Reports
- AWS Organizations
- AWS Pricing Calculator

Resources

Related documents:

- Creating Cost and Usage Reports
- AWS Well-Architected Labs → Cloud Intelligence Dashboards
- AppStream 2.0 Usage Reports
- AppStream 2.0 Metrics and Dimensions
- Enable self-service WorkSpace management capabilities for your users
- Bring Your Own Windows desktop licenses in Amazon WorkSpaces
- Service Control Policies (SCPs) with AWS Organizations
- Tagging your AWS resources
- Tag resources in WorkSpaces Personal
- Tagging Your Amazon AppStream 2.0 Resources
- Amazon AppStream 2.0 Pricing
- Amazon WorkSpaces Pricing
- Amazon WorkSpaces FAQs → Windows Bring Your Own License (BYOL)
- Microsoft Windows Performance Monitor Overview
- Analyze your Amazon AppStream 2.0 usage reports using Amazon Athena and Quick Suite
- Build an enterprise cost and usage dashboard for Amazon WorkSpaces
- Creating a self-service portal for Amazon WorkSpaces end users
- How to enable self-service Amazon WorkSpaces by using Service Catalog Connector for ServiceNow
- Managing Amazon WorkSpaces by integrating Service Catalog with ServiceNow
- Use the EUC Toolkit to manage Amazon AppStream 2.0 and Amazon WorkSpaces
- Best Practices for Deploying Amazon WorkSpaces

- Best Practices for Deploying Amazon AppStream 2.0

Related AWS solutions:

- Cost Optimizer for Amazon WorkSpaces

Related partner solutions:

- Liquidware StratusphereTM UX
- ControlUp Real-Time DX solutions
- Cambrian AppStream Optimizer

Related videos:

- How can I reduce costs for my Amazon AppStream 2.0 Deployment on AWS
- Amazon AppStream 2.0 Fleet Auto Scaling
- Amazon AppStream 2.0 Fleet types - Always-On vs. On-Demand
- Amazon AppStream 2.0 Session Timeouts
- End User Computing Innovation Day: Cost optimization for the reinvented workplace

Related training:

- Introduction to AWS End User Computing Services
- Amazon WorkSpaces Primer
- Amazon WorkSpaces Deep Dive
- Amazon AppStream 2.0 Primer
- AWS Billing and Cost Management and Cost Management
- AWS Foundations: Cost Management

Related tools:

- Cost Optimizer for Amazon AppStream 2.0

Sustainability

The sustainability pillar provides design principles, operational guidance, best-practices, and improvement plans to meet sustainability targets for your AWS workloads.

End user computing (EUC) plays a significant role in supporting environmental, social, and governance (ESG) goals, ultimately contributing to sustainability. EUC in AWS includes Amazon WorkSpaces and Amazon AppStream 2.0 and allows employees to work from anywhere, anytime, and on many devices, thereby reducing the need for commuting and lowering carbon emissions. By promoting remote work, EUC also contributes to a more sustainable approach by reducing the need for physical office spaces and their associated energy consumption.

You are encouraged to use the prescriptive guidance in the Sustainability Pillar whitepaper that applies to most workloads. The guidance that follows is specific to EUC workloads.

Focus areas

- Design principles
- Region selection
- Alignment to demand
- Software and architecture
- Data management
- Hardware and services
- Process and culture
- Key AWS services
- Resources

Design principles

As part of the design of an End User Computing solution, there are some key principles to consider towards sustainability that you can embrace and adapt with other principles alignment such as cost, performance, and operational excellence:

- **Define your end user experience:** By offering EUC access to your users, you have a choice to offer different types of compute to align with each of your scenarios. In addition, based on your user personas, you can optimize the session delivery to fit each of your use case.

- **Manage with automation and tools:** With a robust pipeline and tools in place in the different environments (like development, sandbox, or production), you can optimize unused resources and align the load and capacity to your unique requirements. This principle mirrors operational excellence principles.
- **Control and manage user data:** Data has gravity, cost, and sustainability impacts. User profiles or application data are often a key consideration with an EUC deployment. By controlling the volume, backups, and retention policies, you can lower your sustainability impact.
- **Establish your device strategy:** EUC services provided by AWS are accessible through a few types of client devices. The strategy for your end-user devices will be highly important to meet your sustainability goals.

Region selection

There are no sustainability practices unique to the EUC Lens for Region selection. Refer to the Region selection and best practice SUS01-BP01 from the sustainability pillar of the Well-Architected Framework for more information.

Alignment to demand

EUCSUS01: How do you select the type of fleet and running mode for your end users?

With Amazon AppStream 2.0, you can select between Always-On or On-Demand fleet types. Similarly, Amazon WorkSpaces offers two running modes with AlwaysOn and AutoStop.

EUCSUS02: How do you select the bundle or instance family for your end users?

Selecting an appropriate bundle or instance family involves evaluating your application's compute, memory, storage, and network requirements.

EUCSUS03: How do you align session settings with efficient resource management?

Timeouts are crucial for an AWS EUC deployment, as timeouts enable efficient resource management by disconnecting idle sessions.

For example, consider two scenarios where disconnect timeouts have been configured for an AppStream 2.0 fleet at six hours and 15 minutes respectively. In the case of the 15 minute timeout, a session will be timed out and the instance terminated within 15 minutes.

In comparison, the six hour timeout will result in an instance running for six hours after the user disconnects before it is terminated, which incurs a larger carbon footprint.

Timeouts optimize costs and the carbon footprint for AppStream 2.0and WorkSpaces instances by helping to prevent unnecessary resource consumption and associated charges.

EUCSUS04: How do you align your scaling strategy with efficient use of resources?

Scaling in Amazon AppStream 2.0 helps create a seamless user experience by providing the necessary resources to handle fluctuating user demands, while also optimizing costs and sustainability by avoiding over-provisioning when demand is low.

Best practices

- EUCSUS01-BP01 Choose the appropriate fleet type
- EUCSUS01-BP02 Choose the appropriate running mode for your Amazon WorkSpaces
- EUCSUS02-BP01 Select the instance type or bundle to match software requirement and user personas
- EUCSUS03-BP01 Adapt your AppStream 2.0 fleet timeout
- EUCSUS03-BP02 Adapt the AutoStop timeout and idle disconnect timeout for Amazon DCV
- EUCSUS04-BP01 Implement a scaling methodology in AppStream 2.0

EUCSUS01-BP01 Choose the appropriate fleet type

By selecting Always-On instances in AppStream 2.0, your instances are constantly kept running and ready to receive user connection. With On-Demand, your instances will be provisioned based on your scaling policies, but instances start only when users initiate the connection. Elastic fleet is a fleet of instances managed by AWS directly, and you only pay when your user is launching a new session and there is no scaling management.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Encourage the usage of On-Demand fleet type. With On-Demand, streaming instances run only when users are streaming and therefore have a lower carbon footprint in comparison to Always-On fleets. The number of streaming instances will still require auto scaling rules. Once the user disconnects, the instance is terminated.

An additional option is to select a multi-session fleet according to the performance pillar to select the right instances type.

Elastic fleets offer a pool of streaming instances managed by AppStream 2.0 service. When you use Elastic fleets, an app block (also known as a virtual hard disk) will be downloaded and mounted from Amazon S3. You do not have to configure scaling policies, so you will not consume and reserve unnecessary resources. Elastic fleets do not support domain join, for further details see: Using Active Directory with AppStream 2.0.

EUCSUS01-BP02 Choose the appropriate running mode for your Amazon WorkSpaces

The running mode of a WorkSpace determines its immediate availability and how you pay for it (monthly or hourly). You can choose between the following running modes when you create the WorkSpace:

- **AlwaysOn:** You are paying a fixed monthly fee for unlimited usage of your WorkSpaces. This mode is best for users who use their WorkSpace full time as their primary desktop.
- **AutoStop:** You are paying for your WorkSpaces by the hour. With this mode, your WorkSpaces stop after a specified period of disconnection, and the state of apps and data is saved.

Level of risk exposed if this best practice is not established: High

Implementation guidance

AutoStop instances are stopped when users disconnect and therefore help lower the carbon footprint associated with WorkSpace instances in comparison to AlwaysOn instances. Below a certain threshold, which depends on the bundle selected, we recommend AutoStop mode.

Use Cost Optimizer for Amazon WorkSpaces to set the appropriate running mode of a WorkSpaces based on past usage and improve the sustainability position for WorkSpace environments.

EUCSUS02-BP01 Select the instance type or bundle to match software requirement and user personas

Consider the performance needs, cost implications, and any specific workload characteristics (for example, GPU requirements). Benchmark and test different instance types to find the best fit for your workload. Regularly review and adjust your instance type selection as your application's demands change over time.

Level of risk exposed if this best practice is not established: High

Implementation guidance

AppStream 2.0 offers eight instance families and a set of instance types per family. Explore these instance families and types to identify the appropriate requirement for each use case. For graphics workloads, use Graphics G4dn and Graphics G5. Once you have defined the instance family, you can benchmark at least two instances type to identify the best choice.

Amazon WorkSpaces offers nine bundles, from value to GraphicsPro.g4dn. Once you have selected applications and usage for each use case, identify the requirement in term of CPU, memory, and GPU for each of them.

EUCSUS03-BP01 Adapt your AppStream 2.0 fleet timeout

Configure timeouts for AppStream 2.0 fleets to minimize unnecessary resource consumption whilst also factoring in usability. Minimize resource consumption by verifying that instances are not consuming resources unnecessarily when users are not using them or unlikely to use them.

Usability is an important consideration when shortening timeouts. Setting them too low results in sessions being terminated too early with the risk of impacting user productivity, whereas setting them too high results in instances running without any users, which incurs a higher carbon footprint as well as higher costs.

Strike an appropriate balance in timeout durations to maintain user productivity while reducing resource consumption in periods of low usage.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

You can select a session duration to configure a maximum active session for a user, which defaults to 16 hours. Disconnect timeout and idle disconnect timeout determine when to log off an existing user session. By default, they are both configured at 15 minutes each. The default value can be reduced without disrupting the end user experience.

For example, you can set the idle disconnect timeout for five minutes. You can set timecout configurations in the fleet console.

EUCSUS03-BP02 Adapt the AutoStop timeout and idle disconnect timeout for Amazon DCV

The AutoStop timeout in WorkSpaces is only available with AutoStop. This is not applicable to AlwaysOn WorkSpaces. In WorkSpaces, you can configure how long a user can be inactive while connected to a WorkSpace before they are disconnected. Amazon DCV (Desktop Cloud Virtualization) is the remote display protocol used by Amazon WorkSpaces to stream pixels, keystrokes and mouse movements.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

By default, AutoStop time (in hours**)** is set to one hour, which means that the WorkSpace stops automatically an hour after the WorkSpace is disconnected. Keep the AutoStop time at the default value, as this is the lowest value offered.

EUCSUS04-BP01 Implement a scaling methodology in AppStream 2.0

Scaling policies improve resource utilization and cost management for application streaming workloads.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Either fleet type (On-Demand or Always-On) requires a methodology to verify that the appropriate number of instances are available when users initiate a connection.

A combination of step scaling, scheduled scaling, or target tracking scaling is recommended to match each fleet usage. To avoid extra consumption of instances, monitor your fleet usage and modify your scaling policies accordingly. The following resources describe in further detail the differences between the types of scaling and how to configure them to align with the pattern of usage for the applications being delivered. Keep in mind that the fleet type choice is only available during the fleet creation process.

- AppStream 2.0 Fleet Types
- Fleet Auto Scaling for Amazon AppStream 2.0
- Scaling Your Desktop Application Streams with Amazon AppStream 2.0
- Scale your Amazon AppStream 2.0 fleets
- Monitoring Amazon AppStream 2.0 Resources

Software and architecture

EUCSUS05: Do you use an automation pipeline to manage your images?

Using an automation pipeline for Amazon AppStream 2.0 image management promotes sustainability by optimizing resource utilization, and enabling efficient application delivery.

EUCSUS06: Do you have a mechanism to avoid unused active instances?

Identifying your unused workload will impact positively your cost and sustainability.

Best practices

- EUCSUS05-BP01 Optimize machine image creation, copying, and sharing to each environment (like development, testing, and production)
- EUCSUS06-BP01 Stop image builders and app block builders when not in use
- EUCSUS06-BP02 Implement the Cost Optimizer for Amazon WorkSpaces

EUCSUS05-BP01 Optimize machine image creation, copying, and sharing to each environment (like development, testing, and production)

Using automation with machine images facilitates scalability and elasticity, minimizing over-provisioning and associated energy consumption. Centralized management and compliance reporting further support sustainability initiatives. Overall, automation pipelines contribute to lower environmental impact and improved resource optimization.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use a dedicated and separate account to create your Amazon AppStream images to manage your changes and your image history. Push the image (copy or share) with other development or production AWS accounts. For more detail, see UpdateImagePermissions and UpdateWorkspaceImagePermission.

EUCSUS06-BP01 Stop image builders and app block builders when not in use

In AppStream 2.0, image builders and app block builders are two instances used only when creating your baseline image or application package. There is no requirement to keep them running.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

The Cost Optimizer for Amazon AppStream 2.0 monitors your AppStream 2.0 image builders, notifying you and halting them when they are active for longer than specified thresholds.

EUCSUS06-BP02 Implement the Cost Optimizer for Amazon WorkSpaces

The Cost Optimizer for Amazon WorkSpaces analyzes your Amazon WorkSpaces usage data and automatically converts the WorkSpace to the most cost-effective billing option (hourly or monthly), depending on your individual usage.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

This solution analyzes your Amazon WorkSpaces usage data and automatically converts WorkSpaces to the most cost-effective billing option (hourly or monthly). This verifies that the lowest carbon footprint and cost is associated with each individual WorkSpace instance based on the unique usage pattern for each user. This data provides the opportunity to identify usage of WorkSpaces and delete unused WorkSpaces through the definition of a rule.

Data management

EUCSUS07: How do you manage data within EUC services to promote efficient resource usage?

By implementing effective user profile management strategies in EUC services, organizations can provide a seamless and personalized experience for their users, maintain data integrity and persistence, simplify administration, improve regulatory compliance, and optimize overall performance.

Best practices

- EUCSUS07-BP01 Identify the volume and data requirement for your user profiles

EUCSUS07-BP01 Identify the volume and data requirement for your user profiles

Each user persona may require different volume and performance to align with your business case.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Limit user data to application data and mandatory user data profile. It is a best practice to monitor the storage usage of home folders, application settings persistence, or other storage solutions like FSLogix, OneDrive, and Google Drive. With FSLogix, enable de-duplication in FSx and VHD disk compaction. Fsx for Windows / Fsx on tap.

For more information, see:

- How Application Settings Persistence Works
- Use Amazon FSx for Windows File Server and FSLogix to Optimize Application Settings Persistence on Amazon AppStream 2.0

Hardware and services

EUCSUS08: What is your end user client device lifecyle strategy?

A device lifecycle strategy is a comprehensive plan for managing computing devices from procurement to retirement. It encompasses processes for deployment, maintenance, support, refresh cycles, and secure disposal.

Best practices

- EUCSUS08-BP01 Extend device lifecycle, and review a bring your own device (BYOD) strategy
- EUCSUS08-BP02 Migrate end users to a thin client or web-based client device

EUCSUS08-BP01 Extend device lifecycle, and review a bring your own device (BYOD) strategy

The strategy defines policies for device usage, security, and compliance while optimizing costs through efficient lifecycle management. Implementing a robust device lifecycle strategy fosters standardization, security, and productivity across an organization's device fleet.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

By using either AppStream 2.0 or WorkSpaces, you can extend your device lifecycle. The performance of the local device will not be affected, and a low-performance device can connect and stream an intensive application. Examples of this strategy include Windows laptops, Chromebooks, or other user-owned devices.

EUCSUS08-BP02 Migrate end users to a thin client or web-based client device

Thin client or web-based client devices can reduce investment and are aligned to the previous best practice.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

With WorkSpaces Thin Client, you can offer a device with a direct connection to AppStream 2.0 or WorkSpaces and WorkSpaces Secure Browser. The total lifecycle carbon emission for Amazon WorkSpaces Thin Client is 77kg CO2e as verified by the Carbon Trust. For more information, see Amazon WorkSpaces Thin Client has received Carbon Trust verification for the product's carbon footprint.

Process and culture

There are no process and culture practices specific to the EUC Lens. For more information, see the following best practices:

- SUS06-BP01 Adopt methods that can rapidly introduce sustainability improvements
- SUS06-BP02 Keep your workload up-to-date
- SUS06-BP03 Increase utilization of build environments

Key AWS services

- Amazon AppStream 2.0
- Amazon WorkSpaces
- Amazon S3
- Amazon FSx for Windows
- Amazon FSx for NetApp ONTAPv
- Business Analytics Service - Quick Suite
- Amazon Cloudwatch
- Amazon WorkSpaces Thin Client

Resources

Related documents:

- Manage the running mode in WorkSpaces Personal
- AppStream 2.0 Instance Families
- Bundle options for WorkSpaces Personal
- Create an AppStream 2.0 Fleet and Stack
- AppStream 2.0Fleet Types
- Fleet Auto Scaling for Amazon AppStream 2.0
- Scale your Amazon AppStream 2.0 fleets
- UpdateImagePermissions
- UpdateWorkspaceImagePermission
- How Application Settings Persistence Works
- Use Amazon FSx for Windows File Server and FSLogix to Optimize Application Settings Persistence on Amazon AppStream 2.0
- What to Consider when Selecting a Region for your Workloads
- Amazon WorkSpaces Thin Client has received Carbon Trust verification for the product's carbon footprint
- Github - Cost Optimizer for Amazon AppStream 2.0
- Cost Optimizer for Amazon WorkSpaces

Conclusion

The AWS Well-Architected Framework provides architectural best practices across the six pillars for designing and operating reliable, secure, efficient, and cost-effective systems in the cloud. The End User Computing lens provides additional insights that allow you to dive further into an existing or proposed workload. Using the framework in your architecture will help you produce efficient end-user systems, freeing you to focus on innovation that delivers more value to your end-users.

Whether you have just started designing a greenfield end user environment on AWS or you are looking to migrate an existing one, we hope that his paper has given you new perspectives or sparked new ideas. We encourage you to take advantage of the recommendations offered here as well as the knowledge and experience of our AWS Solution Architects. We would love to hear from you – especially about your success stories building end user environments on AWS. Contact your account team or use Contact Us through our website.

Contributors

The following individuals and organizations have contributed to this document:

- Arvind Raghunathan, Principal Operations Lead SA, Amazon Web Services
- Ben Trunnell, Senior Solutions Architect, Amazon Web Services
- Brian Sheppard, Principal Technical Account Manager (US-CIENG), AWS Enterprise Support, Amazon Web Services
- Bruce Ross, Well-Architected Lens Leader, Amazon Web Services
- Christian O'Donoghue, Senior Technical Account Manager (UKI), Enterprise Support, Amazon Web Services
- Daniel Arrieta Alvarez, Senior Partner Specialist Solutions Architect - LATAM, Amazon Web Services
- Daniel Matlock, Technical Account Manager, ADC, Amazon Web Services
- Dylan Barlett, Senior Developer Advocate, Amazon WorkSpaces, Amazon Web Services
- Ese Alofoje, Cloud Support Engineer II (EAP), Support Engineering, Amazon Web Services
- Grant Joslyn, Senior Solutions Architect, Amazon Web Services
- Jeremy Schiefer, Principal Security Enablement Solutions Architect, WWPS, Amazon Web Services
- Klaus Becker, Senior EMEA End User Computing Specialist Solutions Architect, Amazon Web Services
- Madhuri Srinivasan, Senior Technical Writer, Well- Architected, Amazon Web Services
- Mahmoud Matouk, Principal Security Lead SA, Amazon Web Services
- Manny Aragones, ESL Technical Account Manager (US-CIENG), Enterprise Support, Amazon Web Services
- Mark Homer, Senior EMEA End User Computing Specialist Solutions Architect, Amazon Web Services
- Matthew Wygant, Sr. TPM Guidance, Amazon Web Services
- Noah Jackson, Global End User Computing Solutions Architect, Amazon Web Services
- Peter David, Senior EMEA End User Computing Specialist Solutions Architect, Amazon Web Services
- Ryan Dsouza, Principal Guidance Lead SA, Amazon Web Services

- Sai Sanjay Garishe, Technical Account Manager, US-DNB-Games, Amazon Web Services
- Stewart Matzek, Senior Technical Writer, Well- Architected, Amazon Web Services
- Thomas Sagaspe, France End User Computing Specialist Solutions Architect, Amazon Web Services

Document revisions

Change	Description	Date
Initial release	Initial release of the EUC Lens.	September 18, 2025

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.