aws

The CI/CD litmus test: Is your pipeline fully CI/CD?

AWS Prescriptive Guidance

AWS Prescriptive Guidance: The CI/CD litmus test: Is your pipeline fully CI/CD?

Table of Contents

The CI/CD litmus test: Is your pipeline fully CI/CD?

Steven Guggenheimer and Ananya Koduri, Amazon Web Services (AWS)

August 2023 ([document history](#))

Is your pipeline automated? It's a simple question, but many organizations approach the answer too simply. The answer is much more complicated than a *yes* or *no*.

Innovations in technology are happening constantly, and sometimes it can be hard for organizations to keep up. Is this new thing a fad, or is it the next big thing? Should I overhaul my current practices, or should I wait? Often, by the time it becomes clear that something is indeed the next big thing, you can find yourself playing catchup. *Continuous integration and continuous delivery* (CI/CD) is here to stay, but it wasn't always that way. Many people took a long time to be convinced, and some people still need more convincing.

CI/CD is the process of automating the source, build, test, staging, and production stages of the software release process, and it is commonly described as a pipeline. Today, the cost savings and speed of CI/CD automations have convinced most organizations of its value. But transitioning to this new approach is not an easy task. You need to make sure your staff has the right training, you need to upgrade some resources, and then you need to test, test, test. There's a lot to do. In most cases, you want to make these changes gradually to help your organization adapt.

The purpose of this document is to define what it means to have a fully CI/CD process. It provides a tool to evaluate your own processes and presents a path forward for processes that aren't there yet. This path forward is rarely an overnight conversion. These processes are complex and depend upon many factors, including current employee skill sets and current infrastructure demands. We recommend you prioritize and make small, incremental changes.

Objectives

The following are the potential benefits of implementing the recommendations in this guide:

- **Efficiency** – A fully CI/CD deployment process can reduce complexity, workloads, and countless hours spent debugging, performing manual processes, and maintaining. For more information, see [Benefits of continuous delivery](#). According to a [TechAhead blog post](#), the implementation of the CI/CD process can lead to an estimated 20% savings in time, effort, and resources.

- **Cost reduction** – According to a [Forbes Insight report](#), "Three out of four executives agree that the amount of time, money, and resources spent on ongoing maintenance and management —versus new project development or new initiatives—is affecting the overall competitiveness of their organization." The shorter the development cycle, the higher the chances are that your organization can meet ambitious time-to-market goals and grab the right opportunities at the right time.

- **Speed** – Typically, a fully CI/CD pipeline is able to release software changes to customers within a few hours. Especially in cases with quick fault isolations and small patch pushes, the CI/CD pipeline helps improve the mean time to recovery (MTTR). For more information, see [Reducing MTTR](#).

- **Security** – Fully CI/CD pipelines also secure the release process by reducing the possible entry points for attacks and reducing the risk of human error. The security gains that come with fully automated CI/CD pipelines help avoid the costly consequences of data breaches, service outages, and more.

- **Reduced attrition** – Developers are more satisfied when they can spend more time creating great features and less time mired in an endless cycle of maintenance and debugging. For organizations, this means acquiring and retaining top talent for longer periods of time.

- **Superior quality code** – Developers release code into a shared repository in small batches, which enables them to conduct [parallel testing](#) (BrowserStack blog post). Rather than working in isolation, they share their builds with the team frequently, and they collaborate to identify critical bugs. This provides support for developers, which helps prevent bad code from making it to production. Support from developer peers contributes to high-quality releases and drives organizational growth.

- **Maintenance** – Maintenance and updates are a crucial part of making a great product. However, don't take the system down during peak traffic times. You can use a CI/CD pipeline to perform maintenance during low-usage hours that minimize downtime and performance impacts.

Understanding CI/CD

Continuous integration and continuous delivery (CI/CD) is the process of automating the software release lifecycle. In some cases, the *D* in CI/CD can also mean *deployment*. The difference between *continuous delivery* and *continuous deployment* occurs when you release a change to the production environment. With continuous delivery, a manual approval is required before promoting changes to production. Continuous deployment features an uninterrupted flow through the entirety of the pipeline, and no explicit approvals are required. Because this strategy discusses general CI/CD concepts, the recommendations and information provided are applicable to both the continuous delivery and continuous deployment approaches.

CI/CD automates much or all of the manual processes traditionally required to get new code from a commit into production. A CI/CD pipeline encompasses the source, build, test, staging, and production stages. In each stage, the CI/CD pipelines provisions any infrastructure that is needed to deploy or test the code. By using a CI/CD pipeline, development teams can make changes to code that are then automatically tested and pushed to deployment.

Let's review the basic CI/CD process before discussing some of the ways that you can, knowingly or unknowingly, deviate from being fully CI/CD. The following diagram shows the CI/CD stages and activities in each stage.

CONTINUOUS INTEGRATION (CI)	CONTINUOUS DELIVERY OR CONTINUOUS DEPLOYMENT (CD)			
SOURCE	**BUILD**	**TEST**	**STAGING**	**PRODUCTION**
Code repository	*Development environment*	*Development environment*	*Preproduction environment*	*Production environment*
Feature branches are committed to the main branch of a code repository.	Resources are provisioned in a development environment.	Multiple tests are performed, such as unit, integration, acceptance, and performance tests.	Resources are provisioned in a pre-production environment, and final unit and integration tests are performed.	The code is released to the production environment.

About continuous integration

Continuous integration occurs in a code repository, such as a Git repository in AWS CodeCommit. You treat a single, main branch as the source of truth for the code base, and you create short-lived branches for feature development. You integrate a feature branch into the main branch when you're ready to deploy the feature to upper environments. Feature branches are never deployed directly to upper environments. For more information, see Trunk-based approach in this guide.

Continuous integration process

1. The developer creates a new branch from the main branch.

2. The developer makes changes and builds and tests locally.

3. When the changes are ready, the developer creates a pull request (GitHub documentation) with the main branch as the destination.

4. The code is reviewed.

5. When the code is approved, it is merged into the main branch.

About continuous delivery

Continuous delivery occurs in isolated environments, such as development environments and production environments. The actions that occur in each environment can vary. Often, one of the first stages is used to make updates to the pipeline itself before proceeding. The end result of the deployment is that each environment is updated with the latest changes. The number of development environments for building and testing also varies, but we recommend you use at least two. In the pipeline, each environment is updated in order of its significance, ending with the most important environment, the production environment.

Continuous delivery process

The continuous delivery portion of the pipeline initiates by pulling the code from the main branch of the source repository and passing it to the build stage. The infrastructure as code (IaC) document for the repository outlines the tasks that are performed in each stage. Although using an IaC document is not mandatory, an IaC service or tool, such as AWS CloudFormation or AWS Cloud Development Kit (AWS CDK), is strongly recommended. The most common steps include:

1. Unit tests

2. Code build

3. Resource provisioning

4. Integration tests

If any errors occur or any tests fail at any stage in the pipeline, the current stage rolls back to its previous state, and the pipeline is terminated. Subsequent changes must start in the code repository and go through the fully CI/CD process.

Tests for CI/CD pipelines

The two types of automated tests that are commonly referred to in deployment pipelines are *unit tests* and *integration tests*. However, there are many types of tests that you can run on a code base and the development environment. The [AWS Deployment Pipeline Reference Architecture](#) defines the following types of tests:

- **Unit test** – These tests build and run application code to verify that it is performing according to expectations. They simulate all external dependencies that are used in the code base. Examples of unit test tools include [JUnit](#), [Jest](#), and [pytest](#).

- **Integration test** – These tests verify that the application satisfies technical requirements by testing against a provisioned test environment. Examples of integration test tools include [Cucumber](#), [vRest NG](#), and [integ-tests](#) (for AWS CDK).

- **Acceptance test** – These tests verify that the application satisfies user requirements by testing against a provisioned test environment. Examples of acceptance test tools include [Cypress](#) and [Selenium](#).

- **Synthetic test** – These tests run continuously in the background to generate traffic and verify that the system is healthy. Examples of synthetic test tools include [Amazon CloudWatch Synthetics](#) and [Dynatrace Synthetic Monitoring](#).

- **Performance test** – These tests simulate production capacity. They determine if the application meets performance requirements and compare metrics to past performance. Examples of performance test tools include [Apache JMeter](#), [Locust](#), and [Gatling](#).

- **Resilience test** – Also known as *chaos testing*, these tests inject failures into environments in order to identify risk areas. Periods when the failures are injected are then compared to periods without the failures. Examples of resilience test tools include [AWS Fault Injection Service](#) and [Gremlin](#).

- **Static application security test (SAST)** – These tests analyze code for security violations, such as SQL injection or cross-site scripting (XSS). Examples of SAST tools include Amazon CodeGuru, SonarQube, and Checkmarx.

- **Dynamic application security test (DAST)** – These tests are also known as *penetration testing* or *pen testing*. They identify vulnerabilities, such as SQL injection or XSS in a provisioned test environment. Examples of DAST tools include Zed Attack Proxy (ZAP) and HCL AppScan. For more information, see Penetration Testing.

Not all fully CI/CD pipelines run all of these tests. However, at a minimum, a pipeline should run unit tests and SAST tests on the code base as well as integration and acceptance tests on a test environment.

Metrics for CI/CD pipelines

According to the AWS Deployment Pipeline Reference Architecture, you should, at a minimum, track the following four metrics for CI/CD pipelines:

- **Lead time** – The average amount of time it takes for a single commit to get all the way into production. We recommend targeting a lead time between 1 hour and 1 day, as appropriate for your use case.

- **Deployment frequency** – The number of production deployments within a given period of time. We recommend targeting deployment frequencies between multiple times each day to twice each week, as appropriate for your use case.

- **Mean time between failure (MTBF)** – The average amount of time between the start of a successful pipeline and the start of a failed pipeline. We recommending targeting an MTBF that is as high as possible. For more information, see Increasing MTBF.

- **Mean time to recovery (MTTR)** – The average amount of time between the start of a failed pipeline and the start of the next successful pipeline. We recommending targeting an MTTR that is as low as possible. For more information, see Reducing MTTR.

These metrics help teams track their progress toward becoming fully CI/CD. Teams should have open discussions with the organization's stakeholders regarding what the optimal goals should be. Situations and needs vary greatly from organization to organization, and even from team to team.

It's very important to remember that rapid, drastic change usually increases the risk of problems arising. Set goals to aim for small, incremental improvements. A common optimal lead time for

fully CI/CD pipelines is less than 3 hours. A team that starts with a lead time of 5.2 days should target a reduction of one day every few weeks. After this team reaches a lead time of one day or less, they can stay there for several months and move to a more aggressive lead time only if the team and organization stakeholders deem it necessary.

How fully CI/CD processes are different

CI/CD pipelines use a modern *trunk-based workflow*, in which developers merge small, frequent updates into a main branch (or *trunk*) that is built and tested through the CD portion of the CI/CD pipeline. This workflow has replaced the *Gitflow workflow*, in which development and release branches are separated by a release schedule. In many organizations, Gitflow is still a popular method of version control and deployment. However, it is now considered legacy, and it can be challenging to integrate into a CI/CD pipeline.

For many organizations, the transition from a Gitflow workflow to a trunk-based workflow has been incomplete, and the result is that they get stuck somewhere along the way and never fully migrate to CI/CD. Somehow, their pipelines end up clinging to certain remnants of the legacy workflow, stuck in a transitional state between past and present. Review the differences in the Git workflows, and then learn how using a legacy workflow can affect the following:

- Environment integrity

- Releases

- Security

To make it easier to identify the remnants of a legacy Git workflow in a modern configuration, let's compare Gitflow to the modern, trunk-based approach.

Gitflow approach

The following image shows a Gitflow workflow. The Gitflow approach uses multiple branches to track several different versions of the code at the same time. You schedule releases of updates to an application for some point in the future while developers still work on the current version of the code. Trunk-based repositories can use feature flags to accomplish this, but it's built in to Gitflow by default.

One result of the Gitflow approach is that the application environments are usually out of sync. In a standard Gitflow implementation, development environments reflect the current state of the code while the preproduction and production environments remain frozen on the state of the code base from the most recent release.

This complicates things when a defect appears in the production environment because the code base that the developers work in cannot be merged into production without exposing unreleased features. The way Gitflow handles this situation is by using a hotfix. A hotfix branch is created from the release branch and then deployed straight into the upper environments. The hotfix branch is then merged into the development branch in order to keep the code current.

Trunk-based approach

The following image shows a trunk-based workflow. In a trunk-based workflow, developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially. Unit and integration tests occur between each environment.

Using this workflow, all environments are operating the same code base. There is no need for a hotfix branch for the upper environments because you can implement changes in the main branch without exposing unreleased features. The main branch is always assumed to be stable, free of

defects, and ready to release. This helps you integrate it as a source for a CI/CD pipeline, which can automatically test and deploy your code base through all environments in your pipeline.

Environment integrity benefits of a trunk-based approach

As many developers know, one change in code can sometimes create a [butterfly effect](#) (American Scientist article), where a small deviation that is seemingly unrelated sets off a chain reaction that causes unexpected results. Developers must then fully investigate to discover the root cause.

When scientists conduct an experiment, they separate the test subjects into two groups: the experimental group and the control group. The intention is to make the experimental group and the control group completely identical except for the thing being tested in the experiment. When something happens in the experimental group that doesn't happen in the control group, the only cause can be the thing being tested.

Think of the changes in a deployment as the experimental group, and think of each environment as separate control groups. The results of testing in a lower environment are only reliable when the controls are the same as in an upper environment. The more the environments deviate, the greater the chance of discovering defects in the upper environments. In other words, if the code changes are going to fail in production, we'd much rather them fail in beta first so that they never get to production. This is why every effort should be made to keep each environment, from the lowest test environment to production itself, in sync. This is called *environment integrity*.

The goal of any fully CI/CD process is to discover issues as early as possible. Preserving the environment integrity by using a trunk-based approach can virtually eliminate the need for hotfixes. In a trunk-based workflow, it's rare for an issue to first appear in the production environment.

In a Gitflow approach, after a hotfix is deployed directly to upper environments, it is then added to the development branch. This preserves the fix for future releases. However, the hotfix was developed and tested directly off of the current state of the application. Even if the hotfix works perfectly in production, there's a possibility that problems will arise when it interacts with the newer features in the development branch. Because deploying a hotfix for a hotfix is not typically desirable, this leads to developers spending extra time trying to retrofit the hotfix into the development environment. In many cases, this can lead to significant technical debt and reduce the overall stability of the development environment.

When a failure occurs in an environment, all changes are rolled back so that the environment is returned to its previous state. Any change to a code base should start the pipeline over again from

the very first stage. When an issue does arise in the production environment, the fix should go through the entire pipeline as well. The extra time it takes to go through the lower environments is usually negligible compared to the problems that are avoided by using this approach. Because the whole purpose of the lower environments is to catch mistakes before they reach production, bypassing these environments through a Gitflow approach is an inefficient and unnecessary risk.

Release benefits of a trunk-based approach

One of the things that often makes a hotfix necessary is that, in a legacy workflow, the state of the application that developers are working on might contain several unreleased features that are not yet live in production. The production environment and the development environment only become in sync when a scheduled release occurs, and then they immediately begin to diverge again until the next scheduled release.

Having scheduled releases is possible within a fully CI/CD process. You can delay the release of code to production by using feature flags. However, a fully CI/CD process allows more flexibility by making scheduled releases unnecessary. After all, *continuous* is a key word in CI/CD, and that suggests that changes are released as they become ready. Avoid maintaining a separate release environment that is almost always out of sync with the lower test environments.

If a pipeline is not fully CI/CD, the divergence between upper and lower environments usually occurs at the branch level. Developers work in a development branch and maintain a separate release branch that is updated only when it is time for a scheduled release. As the release branch and the development branch diverge, other complications can arise.

In addition to environments being out of sync, as developers work on the development branch and become accustomed to an application state that is far ahead of what is in production, they must readjust to the state of production every time an issue arises there. The state of the development branch could be many features ahead of production. When developers work in that branch every day, it's difficult to remember what is and isn't released to production. This adds risk that new bugs will be introduced while in the process of fixing other bugs. This result is a seemingly endless cycle of fixes that extend timelines and delay feature releases for weeks, months, or even years.

Security benefits of a trunk-based approach

A fully CI/CD process provides a fully automated single source of truth approach to deployment. The pipeline has a single point of entry. Software updates enter the pipeline at the beginning and are passed as-is from one environment to the next. If an issue is discovered at any stage in the

pipeline, the code changes that fix it must go through the same process and start at the first stage. Reducing the points of entry in a pipeline also reduces the possible ways that vulnerabilities can be introduced into the pipeline.

Additionally, because the entry point is the furthest possible point from the production environment, this drastically reduces the likelihood of vulnerabilities reaching production. If you implement a manual approval process in a fully CI/CD pipeline, you can still allow for go or no-go decision making about whether changes are promoted to the next environment. The decision maker is not necessarily the same person who deploys changes. This separates the responsibilities for the deployer of code changes and the approver of those changes. It also makes it more feasible for a less technical organization leader to perform the role of approver.

Finally, the single point of entry helps you limit write access to the production environment's user interface (UI) console to a few or even zero users. By reducing the number of users who can make manual changes in the console, you reduce the risk of security events. The ability to manually manage the console in the production environment is far more necessary in legacy workflows than in a CI/CD automated approach. These manual changes are more difficult to track, review, and test. They're usually performed to save time, but in the long run, they add significant technical debt to the project.

Console security issues aren't necessarily caused by bad actors. Many of the issues that occur in the console are accidental. Accidental security exposure is very common, and it has led to the rise of the zero-trust security model. This model posits, in part, that security accidents are less likely when even internal staff has as little access as possible, also known as *least-privilege permissions*. Preserving the integrity of the production environment by restricting all processes to an automated pipeline practically eliminates the risk of console-related security issues.

Litmus test for CI/CD pipelines

In chemistry, litmus paper is a thin strip of paper treated with a special red or blue dye that is used to determine the acidity of a substance. An acid turns blue litmus paper red, a base turns red litmus paper blue, and neutral substances don't affect the paper's color at all.

The way litmus paper determines acidity is by measuring a substance's pH level. If a pH level is higher than 8, it's acidic; if it's below 5, it's basic; and if it's between 5 and 8, it's neutral. Similarly, the CI/CD litmus test helps you measure the CI/CD level of your pipeline.

To test whether your pipeline is fully CI/CD

1. Start with a score of 0.

2. Answer each of the following questions, and add 1 to your score for each time you answer *yes*:

   - Do our repositories each have exactly one main branch that is used to deploy to environments?

   - Do we commit code to the main branch frequently and avoid having long-running feature branches?

   - Does our pipeline have a single point of entry? In other words, does our pipeline pull the code from each repository exactly one time?

   - Do we have more than one deployment environment?

   - When the pipeline is not running, are our upper and lower environments generally in sync?

   - Do we run tests on code before deploying?

   - Do we run tests on an environment before promoting to the next environment?

   - Does our pipeline do a full rollback and exit after a failure?

   - Does our pipeline restart from the first step when recovering from a failure?

   - Do we follow the same process to fix bugs in production that we do to release features to production?

   - Do we use some form of infrastructure as code (IaC) templates to deploy code?

3. Answer each of the following questions, and add 1 to your score for each time you answer *no*:

   - Do we ever deploy directly into a deployment environment from branches other than the main branch?

   - Do we ever deploy directly from any branch to an upper or production environment?

- Do we often find bugs in upper environments that were not present in lower environments?

- Do we ever bypass lower environments during a deployment?

- Do we wait until a scheduled release time to deploy to production?

- Do we regularly make updates in the production environment's console?

- Are there any manual deployment steps that must be done in the production environment's console in order to complete deployment?

- Does more than one person have write access to the production environment?

- Do more than five people have write access to the production environment?

4. Divide your score by 2. This is your pipeline's CI/CD score.

5. Compare your pipeline's CI/CD score to the following table to determine your pipeline's CI/CD level.

CI/CD score	CI/CD level
9.5 or above	Fully CI/CD
8–9	Mostly CI/CD
5–7	Neutral
Below 5	Not CI/CD

If you scored below 8, we recommend that you set a goal to move gradually toward the next level. When that goal is achieved, then product stakeholders should assess if and when a new goal should be set. The intention of this exercise is not necessarily to advocate for a change to your pipeline, but rather to bring awareness of what a fully CI/CD deployment process looks like and where your pipelines currently sit on that spectrum.

Best practices for CI/CD pipelines

The following are best practices for fully CI/CD pipelines:

- **Secure the production environment** – Because it's possible to accomplish virtually everything necessary for account and environment maintenance by using IaC, it's important to make every effort to secure the production environment by limiting console and programmatic access. We recommend limiting access to only a few, or even zero, users. When you deploy IaC through AWS CloudFormation, the user needs limited permissions. Most permissions are assigned to the CloudFormation service through a service role. For more information, see [Service role](#) in the CloudFormation documentation and [Implementing policies for least-privilege permissions for AWS CloudFormation](#).

- **Create separate accounts for each environment** – By dedicating a separate account to each environment, you can simplify the deployment process and create fine-grained access controls at the account level. When multiple environments share resources, it reduces the integrity of the environment as an isolated unit. It's best to keep environments synchronized and distinct. This is even more important for the production environment because everything in that account should be treated as a production resource.

- **Restrict personally identifiable information (PII) to the production environment** – Both for security and protection from liability risks, secure PII as much as possible. When possible in lower environments, use anonymized or sample data instead of copying potentially sensitive data from the production environment.

- **Review code in repositories** – A fully CI/CD process reduces the entry points for a pipeline to a single point, and that single point should be secured. For this reason, it's recommended that you require multiple code reviews before merging feature branches into the main branch. These code reviews can be conducted by any qualified team member, but at least one senior member should review. The code should be tested rigorously by the reviewer. After all, the best way to fix issues in a pipeline is to avoid introducing them into it. Also, it's important to resolve all comments made by any reviewer before merging. This resolution could simply be an explanation of why no changes are needed, but addressing all comments is an important extra check to help prevent the introduction of issues into the pipeline.

- **Make small and frequent merges** – In order to take full advantage of continuous integration, it's a good idea to push local changes into the pipeline continuously as well. After all, it's much more beneficial for the development environments to stay in sync if the local environments keep up with them too.

For more best practices for CI/CD pipelines, see [Summary of best practices](#) in *Practicing Continuous Integration and Continuous Delivery on AWS*.

FAQ

What are some key indicators that my deployment process is not fully CI/CD?

The most common indicator is when there are multiple repository branches that represent separate environments in a pipeline. Repositories in a fully CI/CD process use a trunk-based workflow, in which one branch acts as the single source of truth for that repository's deployments. For more information, see Trunk-based approach. Other indicators include manual deployment steps other than simple go or no-go decisions, the use of hotfixes, and scheduled releases.

What if I want to use a fully CI/CD process but still want to schedule releases of certain features for specific points in time?

This is usually done with feature flags. In this process, deployments are still made continuously, but certain features are hidden by using conditional closures in the code until it is time to release them.

What if some steps in my deployment process can't be automated?

One of the goals of a fully CI/CD pipeline is to minimize the need for manual processes, but there are certainly potential use cases where manual processes may be necessary. In fact, read-only processes, such as consulting application logs, can often be done in production environments with minimal risk. However, it is strongly recommended that you treat manual write actions in production as an absolute last resort.

What if my technical staff is more comfortable with legacy workflows than with a fully CI/CD process?

It's common for technical staff to be resistant to major changes, especially when something that used to be a best practice is replaced by something newer. Technology moves fast, and improvements are constantly being discovered. Although a certain degree of skepticism is a good quality for a technical staff to have, it's just as important for them to be open to change. Don't

move too fast with skeptical staff because they need to manage changes to the system before they are implemented. The key is to prevent skeptics from remaining static forever.

What if my environments are in multiple accounts? Can I still use a fully CI/CD process?

Yes, in fact, it's recommended to use a separate account for each environment. For more information about a pipeline that activates stages in different accounts, see Create a pipeline in CodePipeline that uses resources from another AWS account.

Next steps

Use the Litmus test for CI/CD pipelines section to evaluate the DevOps processes in your organization. Determine whether your processes are fully CI/CD. If they aren't, decide whether they need improvement to take full advantage of the benefits of CI/CD deployments.

How do you know when you're finished? Well, the answer is that a lot of organizations never actually finish. They stop somewhere along the way, at a suitable spot for their use case. Though a fully CI/CD pipeline is the best-case scenario, it depends heavily on the organizational situation and the stakeholders behind the decision. Stakeholders must decide which stage of the CI/CD implementation works best for their use case and how best to map the progression to the next phases.

For more information about designing and building CI/CD pipelines, see Resources.

Resources

AWS documentation and references

- [AWS Deployment Pipeline Reference Architecture](#)
- [What is Continuous Delivery?](#)
- [Practicing Continuous Integration and Continuous Delivery on AWS](#) (AWS Whitepaper)
- [Set Up a CI/CD Pipeline on AWS](#) (AWS Hands-on Tutorial)
- [Create a pipeline in AWS Regions that don't support AWS CodePipeline](#) (AWS Prescriptive Guidance)
- [Use AWS CodeCommit and AWS CodePipeline to deploy a CI/CD pipeline in multiple AWS accounts](#) (AWS Prescriptive Guidance)
- [Deployment Pipelines Reference Architecture and Reference Implementations](#) (AWS blog post)

Services and tools

- [The CI/CD Litmus Test](#)
- [AWS Cloud Development Kit (AWS CDK)](#)
- [AWS CloudFormation](#)
- [AWS CodeCommit](#)
- [AWS CodePipeline](#)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an RSS feed.

Change	Description	Date
Initial publication	—	August 25, 2023

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

Numbers

7 Rs

Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:

- Refactor/re-architect – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.

- Replatform (lift and reshape) – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.

- Repurchase (drop and shop) – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.

- Rehost (lift and shift) – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.

- Relocate (hypervisor-level lift and shift) – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. This migration scenario is specific to VMware Cloud on AWS, which supports virtual machine (VM) compatibility and workload portability between your on-premises environment and AWS. You can use the VMware Cloud Foundation technologies from your on-premises data centers when you migrate your infrastructure to VMware Cloud on AWS. Example: Relocate the hypervisor hosting your Oracle database to VMware Cloud on AWS.

- Retain (revisit) – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later

time, and legacy applications that you want to retain, because there's no business justification for migrating them.

- Retire – Decommission or remove applications that are no longer needed in your source environment.

A

ABAC

See attribute-based access control.

abstracted services

See managed services.

ACID

See atomicity, consistency, isolation, durability.

active-active migration

A database migration method in which the source and target databases are kept in sync (by using a bidirectional replication tool or dual write operations), and both databases handle transactions from connecting applications during migration. This method supports migration in small, controlled batches instead of requiring a one-time cutover. It's more flexible but requires more work than active-passive migration.

active-passive migration

A database migration method in which in which the source and target databases are kept in sync, but only the source database handles transactions from connecting applications while data is replicated to the target database. The target database doesn't accept any transactions during migration.

aggregate function

A SQL function that operates on a group of rows and calculates a single return value for the group. Examples of aggregate functions include SUM and MAX.

AI

See artificial intelligence.

AIOps

See [artificial intelligence operations](#).

anonymization

The process of permanently deleting personal information in a dataset. Anonymization can help protect personal privacy. Anonymized data is no longer considered to be personal data.

anti-pattern

A frequently used solution for a recurring issue where the solution is counter-productive, ineffective, or less effective than an alternative.

application control

A security approach that allows the use of only approved applications in order to help protect a system from malware.

application portfolio

A collection of detailed information about each application used by an organization, including the cost to build and maintain the application, and its business value. This information is key to [the portfolio discovery and analysis process](#) and helps identify and prioritize the applications to be migrated, modernized, and optimized.

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see [What is Artificial Intelligence?](#)

artificial intelligence operations (AIOps)

The process of using machine learning techniques to solve operational problems, reduce operational incidents and human intervention, and increase service quality. For more information about how AIOps is used in the AWS migration strategy, see the [operations integration guide](#).

asymmetric encryption

An encryption algorithm that uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted.

atomicity, consistency, isolation, durability (ACID)

A set of software properties that guarantee the data validity and operational reliability of a database, even in the case of errors, power failures, or other problems.

attribute-based access control (ABAC)

The practice of creating fine-grained permissions based on user attributes, such as department, job role, and team name. For more information, see ABAC for AWS in the AWS Identity and Access Management (IAM) documentation.

authoritative data source

A location where you store the primary version of data, which is considered to be the most reliable source of information. You can copy data from the authoritative data source to other locations for the purposes of processing or modifying the data, such as anonymizing, redacting, or pseudonymizing it.

Availability Zone

A distinct location within an AWS Region that is insulated from failures in other Availability Zones and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

AWS Cloud Adoption Framework (AWS CAF)

A framework of guidelines and best practices from AWS to help organizations develop an efficient and effective plan to move successfully to the cloud. AWS CAF organizes guidance into six focus areas called perspectives: business, people, governance, platform, security, and operations. The business, people, and governance perspectives focus on business skills and processes; the platform, security, and operations perspectives focus on technical skills and processes. For example, the people perspective targets stakeholders who handle human resources (HR), staffing functions, and people management. For this perspective, AWS CAF provides guidance for people development, training, and communications to help ready the organization for successful cloud adoption. For more information, see the AWS CAF website and the AWS CAF whitepaper.

AWS Workload Qualification Framework (AWS WQF)

A tool that evaluates database migration workloads, recommends migration strategies, and provides work estimates. AWS WQF is included with AWS Schema Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.

B

bad bot

>   A bot that is intended to disrupt or cause harm to individuals or organizations.

BCP

>   See business continuity planning.

behavior graph

>   A unified, interactive view of resource behavior and interactions over time. You can use a
>   behavior graph with Amazon Detective to examine failed logon attempts, suspicious API
>   calls, and similar actions. For more information, see Data in a behavior graph in the Detective
>   documentation.

big-endian system

>   A system that stores the most significant byte first. See also endianness.

binary classification

>   A process that predicts a binary outcome (one of two possible classes). For example, your ML
>   model might need to predict problems such as "Is this email spam or not spam?" or "Is this
>   product a book or a car?"

bloom filter

>   A probabilistic, memory-efficient data structure that is used to test whether an element is a
>   member of a set.

blue/green deployment

>   A deployment strategy where you create two separate but identical environments. You run the
>   current application version in one environment (blue) and the new application version in the
>   other environment (green). This strategy helps you quickly roll back with minimal impact.

bot

>   A software application that runs automated tasks over the internet and simulates human
>   activity or interaction. Some bots are useful or beneficial, such as web crawlers that index
>   information on the internet. Some other bots, known as *bad bots*, are intended to disrupt or
>   cause harm to individuals or organizations.

botnet

Networks of bots that are infected by malware and are under the control of a single party, known as a *bot herder* or *bot operator*. Botnets are the best-known mechanism to scale bots and their impact.

branch

A contained area of a code repository. The first branch created in a repository is the *main branch*. You can create a new branch from an existing branch, and you can then develop features or fix bugs in the new branch. A branch you create to build a feature is commonly referred to as a *feature branch*. When the feature is ready for release, you merge the feature branch back into the main branch. For more information, see About branches (GitHub documentation).

break-glass access

In exceptional circumstances and through an approved process, a quick means for a user to gain access to an AWS account that they don't typically have permissions to access. For more information, see the Implement break-glass procedures indicator in the AWS Well-Architected guidance.

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

buffer cache

The memory area where the most frequently accessed data is stored.

business capability

What a business does to generate value (for example, sales, customer service, or marketing). Microservices architectures and development decisions can be driven by business capabilities. For more information, see the Organized around business capabilities section of the Running containerized microservices on AWS whitepaper.

business continuity planning (BCP)

A plan that addresses the potential impact of a disruptive event, such as a large-scale migration, on operations and enables a business to resume operations quickly.

C

CAF

> See [AWS Cloud Adoption Framework](#).

canary deployment

> The slow and incremental release of a version to end users. When you are confident, you deploy the new version and replace the current version in its entirety.

CCoE

> See [Cloud Center of Excellence](#).

CDC

> See [change data capture](#).

change data capture (CDC)

> The process of tracking changes to a data source, such as a database table, and recording metadata about the change. You can use CDC for various purposes, such as auditing or replicating changes in a target system to maintain synchronization.

chaos engineering

> Intentionally introducing failures or disruptive events to test a system's resilience. You can use [AWS Fault Injection Service (AWS FIS)](#) to perform experiments that stress your AWS workloads and evaluate their response.

CI/CD

> See [continuous integration and continuous delivery](#).

classification

> A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

client-side encryption

> Encryption of data locally, before the target AWS service receives it.

Cloud Center of Excellence (CCoE)

A multi-disciplinary team that drives cloud adoption efforts across an organization, including developing cloud best practices, mobilizing resources, establishing migration timelines, and leading the organization through large-scale transformations. For more information, see the CCoE posts on the AWS Cloud Enterprise Strategy Blog.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to edge computing technology.

cloud operating model

In an IT organization, the operating model that is used to build, mature, and optimize one or more cloud environments. For more information, see Building your Cloud Operating Model.

cloud stages of adoption

The four phases that organizations typically go through when they migrate to the AWS Cloud:

- Project – Running a few cloud-related projects for proof of concept and learning purposes

- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)

- Migration – Migrating individual applications

- Re-invention – Optimizing products and services, and innovating in the cloud

These stages were defined by Stephen Orban in the blog post The Journey Toward Cloud-First & the Stages of Adoption on the AWS Cloud Enterprise Strategy blog. For information about how they relate to the AWS migration strategy, see the migration readiness guide.

CMDB

See configuration management database.

code repository

A location where source code and other assets, such as documentation, samples, and scripts, are stored and updated through version control processes. Common cloud repositories include GitHub or AWS CodeCommit. Each version of the code is called a *branch*. In a microservice structure, each repository is devoted to a single piece of functionality. A single CI/CD pipeline can use multiple repositories.

cold cache

A buffer cache that is empty, not well populated, or contains stale or irrelevant data. This affects performance because the database instance must read from the main memory or disk, which is slower than reading from the buffer cache.

cold data

Data that is rarely accessed and is typically historical. When querying this kind of data, slow queries are typically acceptable. Moving this data to lower-performing and less expensive storage tiers or classes can reduce costs.

computer vision (CV)

A field of AI that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, AWS Panorama offers devices that add CV to on-premises camera networks, and Amazon SageMaker provides image processing algorithms for CV.

configuration drift

For a workload, a configuration change from the expected state. It might cause the workload to become noncompliant, and it's typically gradual and unintentional.

configuration management database (CMDB)

A repository that stores and manages information about a database and its IT environment, including both hardware and software components and their configurations. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in an AWS account and Region, or across an organization, by using a YAML template. For more information, see Conformance packs in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

The process of automating the source, build, test, staging, and production stages of the software release process. CI/CD is commonly described as a pipeline. CI/CD can help you automate processes, improve productivity, improve code quality, and deliver faster. For more information, see Benefits of continuous delivery. CD can also stand for *continuous deployment*. For more information, see Continuous Delivery vs. Continuous Deployment.

CV

See [computer vision](#).

D

data at rest

Data that is stationary in your network, such as data that is in storage.

data classification

A process for identifying and categorizing the data in your network based on its criticality and sensitivity. It is a critical component of any cybersecurity risk management strategy because it helps you determine the appropriate protection and retention controls for the data. Data classification is a component of the security pillar in the AWS Well-Architected Framework. For more information, see [Data classification](#).

data drift

A meaningful variation between the production data and the data that was used to train an ML model, or a meaningful change in the input data over time. Data drift can reduce the overall quality, accuracy, and fairness in ML model predictions.

data in transit

Data that is actively moving through your network, such as between network resources.

data mesh

An architectural framework that provides distributed, decentralized data ownership with centralized management and governance.

data minimization

The principle of collecting and processing only the data that is strictly necessary. Practicing data minimization in the AWS Cloud can reduce privacy risks, costs, and your analytics carbon footprint.

data perimeter

A set of preventive guardrails in your AWS environment that help make sure that only trusted identities are accessing trusted resources from expected networks. For more information, see [Building a data perimeter on AWS](#).

data preprocessing

To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

data provenance

The process of tracking the origin and history of data throughout its lifecycle, such as how the data was generated, transmitted, and stored.

data subject

An individual whose data is being collected and processed.

data warehouse

A data management system that supports business intelligence, such as analytics. Data warehouses commonly contain large amounts of historical data, and they are typically used for queries and analysis.

database definition language (DDL)

Statements or commands for creating or modifying the structure of tables and objects in a database.

database manipulation language (DML)

Statements or commands for modifying (inserting, updating, and deleting) information in a database.

DDL

See database definition language.

deep ensemble

To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

defense-in-depth

An information security approach in which a series of security mechanisms and controls are thoughtfully layered throughout a computer network to protect the confidentiality, integrity, and availability of the network and the data within. When you adopt this strategy on AWS, you add multiple controls at different layers of the AWS Organizations structure to help secure resources. For example, a defense-in-depth approach might combine multi-factor authentication, network segmentation, and encryption.

delegated administrator

In AWS Organizations, a compatible service can register an AWS member account to administer the organization's accounts and manage permissions for that service. This account is called the *delegated administrator* for that service. For more information and a list of compatible services, see Services that work with AWS Organizations in the AWS Organizations documentation.

deployment

The process of making an application, new features, or code fixes available in the target environment. Deployment involves implementing changes in a code base and then building and running that code base in the application's environments.

development environment

See environment.

detective control

A security control that is designed to detect, log, and alert after an event has occurred. These controls are a second line of defense, alerting you to security events that bypassed the preventative controls in place. For more information, see Detective controls in *Implementing security controls on AWS*.

development value stream mapping (DVSM)

A process used to identify and prioritize constraints that adversely affect speed and quality in a software development lifecycle. DVSM extends the value stream mapping process originally designed for lean manufacturing practices. It focuses on the steps and teams required to create and move value through the software development process.

digital twin

A virtual representation of a real-world system, such as a building, factory, industrial equipment, or production line. Digital twins support predictive maintenance, remote monitoring, and production optimization.

dimension table

In a star schema, a smaller table that contains data attributes about quantitative data in a fact table. Dimension table attributes are typically text fields or discrete numbers that behave like text. These attributes are commonly used for query constraining, filtering, and result set labeling.

disaster

An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location. These events can be natural disasters, technical failures, or the result of human actions, such as unintentional misconfiguration or a malware attack.

disaster recovery (DR)

The strategy and process you use to minimize downtime and data loss caused by a disaster. For more information, see Disaster Recovery of Workloads on AWS: Recovery in the Cloud in the AWS Well-Architected Framework.

DML

See database manipulation language.

domain-driven design

An approach to developing a complex software system by connecting its components to evolving domains, or core business goals, that each component serves. This concept was introduced by Eric Evans in his book, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Boston: Addison-Wesley Professional, 2003). For information about how you can use domain-driven design with the strangler fig pattern, see Modernizing legacy Microsoft ASP.NET (ASMX) web services incrementally by using containers and Amazon API Gateway.

DR

See disaster recovery.

drift detection

Tracking deviations from a baselined configuration. For example, you can use AWS CloudFormation to detect drift in system resources, or you can use AWS Control Tower to detect changes in your landing zone that might affect compliance with governance requirements.

DVSM

See development value stream mapping.

E

EDA

See [exploratory data analysis](#).

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with [cloud computing](#), edge computing can reduce communication latency and improve response time.

encryption

A computing process that transforms plaintext data, which is human-readable, into ciphertext.

encryption key

A cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique.

endianness

The order in which bytes are stored in computer memory. Big-endian systems store the most significant byte first. Little-endian systems store the least significant byte first.

endpoint

See [service endpoint](#).

endpoint service

A service that you can host in a virtual private cloud (VPC) to share with other users. You can create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts or to AWS Identity and Access Management (IAM) principals. These accounts or principals can connect to your endpoint service privately by creating interface VPC endpoints. For more information, see [Create an endpoint service](#) in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

A system that automates and manages key business processes (such as accounting, [MES](#), and project management) for an enterprise.

envelope encryption

The process of encrypting an encryption key with another encryption key. For more information, see Envelope encryption in the AWS Key Management Service (AWS KMS) documentation.

environment

An instance of a running application. The following are common types of environments in cloud computing:

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.

- lower environments – All development environments for an application, such as those used for initial builds and tests.

- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.

- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

In agile methodologies, functional categories that help organize and prioritize your work. Epics provide a high-level description of requirements and implementation tasks. For example, AWS CAF security epics include identity and access management, detective controls, infrastructure security, data protection, and incident response. For more information about epics in the AWS migration strategy, see the program implementation guide.

ERP

See enterprise resource planning.

exploratory data analysis (EDA)

The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

F

fact table

The central table in a [star schema](). It stores quantitative data about business operations. Typically, a fact table contains two types of columns: those that contain measures and those that contain a foreign key to a dimension table.

fail fast

A philosophy that uses frequent and incremental testing to reduce the development lifecycle. It is a critical part of an agile approach.

fault isolation boundary

In the AWS Cloud, a boundary such as an Availability Zone, AWS Region, control plane, or data plane that limits the effect of a failure and helps improve the resilience of workloads. For more information, see [AWS Fault Isolation Boundaries]().

feature branch

See [branch]().

features

The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

How significant a feature is for a model's predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations (SHAP) and integrated gradients. For more information, see [Machine learning model interpretability with :AWS]().

feature transformation

To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to benefit from the data. For example, if you break down the "2021-05-27 00:15:37" date into "2021", "May", "Thu", and "15", you can help the learning algorithm learn nuanced patterns associated with different data components.

FGAC

See [fine-grained access control]().

fine-grained access control (FGAC)

The use of multiple conditions to allow or deny an access request.

flash-cut migration

A database migration method that uses continuous data replication through change data capture to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

G

geo blocking

See geographic restrictions.

geographic restrictions (geo blocking)

In Amazon CloudFront, an option to prevent users in specific countries from accessing content distributions. You can use an allow list or block list to specify approved and banned countries. For more information, see Restricting the geographic distribution of your content in the CloudFront documentation.

Gitflow workflow

An approach in which lower and upper environments use different branches in a source code repository. The Gitflow workflow is considered legacy, and the trunk-based workflow is the modern, preferred approach.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction of compatibility with existing infrastructure, also known as brownfield. If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts

for remediation. They are implemented by using AWS Config, AWS Security Hub, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See high availability.

heterogeneous database migration

Migrating your source database to a target database that uses a different database engine (for example, Oracle to Amazon Aurora). Heterogeneous migration is typically part of a re-architecting effort, and converting the schema can be a complex task. AWS provides AWS SCT that helps with schema conversions.

high availability (HA)

The ability of a workload to operate continuously, without intervention, in the event of challenges or disasters. HA systems are designed to automatically fail over, consistently deliver high-quality performance, and handle different loads and failures with minimal performance impact.

historian modernization

An approach used to modernize and upgrade operational technology (OT) systems to better serve the needs of the manufacturing industry. A *historian* is a type of database that is used to collect and store data from various sources in a factory.

homogeneous database migration

Migrating your source database to a target database that shares the same database engine (for example, Microsoft SQL Server to Amazon RDS for SQL Server). Homogeneous migration is typically part of a rehosting or replatforming effort. You can use native database utilities to migrate the schema.

hot data

Data that is frequently accessed, such as real-time data or recent translational data. This data typically requires a high-performance storage tier or class to provide fast query responses.

hotfix

An urgent fix for a critical issue in a production environment. Due to its urgency, a hotfix is usually made outside of the typical DevOps release workflow.

hypercare period

Immediately following cutover, the period of time when a migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this period is 1–4 days in length. At the end of the hypercare period, the migration team typically transfers responsibility for the applications to the cloud operations team.

I

IaC

See infrastructure as code.

identity-based policy

A policy attached to one or more IAM principals that defines their permissions within the AWS Cloud environment.

idle application

An application that has an average CPU and memory usage between 5 and 20 percent over a period of 90 days. In a migration project, it is common to retire these applications or retain them on premises.

IIoT

See industrial Internet of Things.

immutable infrastructure

A model that deploys new infrastructure for production workloads instead of updating, patching, or modifying the existing infrastructure. Immutable infrastructures are inherently more consistent, reliable, and predictable than mutable infrastructure. For more information, see the Deploy using immutable infrastructure best practice in the AWS Well-Architected Framework.

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The AWS Security Reference Architecture recommends

setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

A cutover strategy in which you migrate your application in small parts instead of performing a single, full cutover. For example, you might move only a few microservices or users to the new system initially. After you verify that everything is working properly, you can incrementally move additional microservices or users until you can decommission your legacy system. This strategy reduces the risks associated with large migrations.

Industry 4.0

A term that was introduced by Klaus Schwab in 2016 to refer to the modernization of manufacturing processes through advances in connectivity, real-time data, automation, analytics, and AI/ML.

infrastructure

All of the resources and assets contained within an application's environment.

infrastructure as code (IaC)

The process of provisioning and managing an application's infrastructure through a set of configuration files. IaC is designed to help you centralize infrastructure management, standardize resources, and scale quickly so that new environments are repeatable, reliable, and consistent.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see Building an industrial Internet of Things (IIoT) digital transformation strategy.

inspection VPC

In an AWS multi-account architecture, a centralized VPC that manages inspections of network traffic between VPCs (in the same or different AWS Regions), the internet, and on-premises networks. The AWS Security Reference Architecture recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

Internet of Things (IoT)

> The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see What is IoT?

interpretability

> A characteristic of a machine learning model that describes the degree to which a human can understand how the model's predictions depend on its inputs. For more information, see Machine learning model interpretability with AWS.

IoT

> See Internet of Things.

IT information library (ITIL)

> A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.

IT service management (ITSM)

> Activities associated with designing, implementing, managing, and supporting IT services for an organization. For information about integrating cloud operations with ITSM tools, see the operations integration guide.

ITIL

> See IT information library.

ITSM

> See IT service management.

L

label-based access control (LBAC)

> An implementation of mandatory access control (MAC) where the users and the data itself are each explicitly assigned a security label value. The intersection between the user security label and data security label determines which rows and columns can be seen by the user.

landing zone

A landing zone is a well-architected, multi-account AWS environment that is scalable and secure. This is a starting point from which your organizations can quickly launch and deploy workloads and applications with confidence in their security and infrastructure environment. For more information about landing zones, see Setting up a secure and scalable multi-account AWS environment.

large migration

A migration of 300 or more servers.

LBAC

See label-based access control.

least privilege

The security best practice of granting the minimum permissions required to perform a task. For more information, see Apply least-privilege permissions in the IAM documentation.

lift and shift

See 7 Rs.

little-endian system

A system that stores the least significant byte first. See also endianness.

lower environments

See environment.

M

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see Machine Learning.

main branch

See branch.

malware

> Software that is designed to compromise computer security or privacy. Malware might disrupt computer systems, leak sensitive information, or gain unauthorized access. Examples of malware include viruses, worms, ransomware, Trojan horses, spyware, and keyloggers.

managed services

> AWS services for which AWS operates the infrastructure layer, the operating system, and platforms, and you access the endpoints to store and retrieve data. Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB are examples of managed services. These are also known as *abstracted services*.

manufacturing execution system (MES)

> A software system for tracking, monitoring, documenting, and controlling production processes that convert raw materials to finished products on the shop floor.

MAP

> See Migration Acceleration Program.

mechanism

> A complete process in which you create a tool, drive adoption of the tool, and then inspect the results in order to make adjustments. A mechanism is a cycle that reinforces and improves itself as it operates. For more information, see Building mechanisms in the AWS Well-Architected Framework.

member account

> All AWS accounts other than the management account that are part of an organization in AWS Organizations. An account can be a member of only one organization at a time.

MES

> See manufacturing execution system.

Message Queuing Telemetry Transport (MQTT)

> A lightweight, machine-to-machine (M2M) communication protocol, based on the publish/ subscribe pattern, for resource-constrained IoT devices.

microservice

> A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include

microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see Integrating microservices by using AWS serverless services.

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed, and scaled to meet demand for specific functions of an application. For more information, see Implementing microservices on AWS.

Migration Acceleration Program (MAP)

An AWS program that provides consulting support, training, and services to help organizations build a strong operational foundation for moving to the cloud, and to help offset the initial cost of migrations. MAP includes a migration methodology for executing legacy migrations in a methodical way and a set of tools to automate and accelerate common migration scenarios.

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the AWS migration strategy.

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners, migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the discussion of migration factories and the Cloud Migration Factory guide in this content set.

migration metadata

The information about the application and server that is needed to complete the migration. Each migration pattern requires a different set of migration metadata. Examples of migration metadata include the target subnet, security group, and AWS account.

migration pattern

A repeatable migration task that details the migration strategy, the migration destination, and the migration application or service used. Example: Rehost migration to Amazon EC2 with AWS Application Migration Service.

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The MPA tool (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

The process of gaining insights about an organization's cloud readiness status, identifying strengths and weaknesses, and building an action plan to close identified gaps, using the AWS CAF. For more information, see the migration readiness guide. MRA is the first phase of the AWS migration strategy.

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the 7 Rs entry in this glossary and see Mobilize your organization to accelerate large-scale migrations.

ML

See machine learning.

modernization

Transforming an outdated (legacy or monolithic) application and its infrastructure into an agile, elastic, and highly available system in the cloud to reduce costs, gain efficiencies, and take advantage of innovations. For more information, see Strategy for modernizing applications in the AWS Cloud.

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and

milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see Evaluating modernization readiness for applications in the AWS Cloud.

monolithic applications (monoliths)

Applications that run as a single service with tightly coupled processes. Monolithic applications have several drawbacks. If one application feature experiences a spike in demand, the entire architecture must be scaled. Adding or improving a monolithic application's features also becomes more complex when the code base grows. To address these issues, you can use a microservices architecture. For more information, see Decomposing monoliths into microservices.

MPA

See Migration Portfolio Assessment.

MQTT

See Message Queuing Telemetry Transport.

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

mutable infrastructure

A model that updates and modifies the existing infrastructure for production workloads. For improved consistency, reliability, and predictability, the AWS Well-Architected Framework recommends the use of immutable infrastructure as a best practice.

O

OAC

See origin access control.

OAI

See origin access identity.

OCM

See organizational change management.

offline migration

> A migration method in which the source workload is taken down during the migration process. This method involves extended downtime and is typically used for small, non-critical workloads.

OI

> See operations integration.

OLA

> See operational-level agreement.

online migration

> A migration method in which the source workload is copied to the target system without being taken offline. Applications that are connected to the workload can continue to function during the migration. This method involves zero to minimal downtime and is typically used for critical production workloads.

OPC-UA

> See Open Process Communications - Unified Architecture.

Open Process Communications - Unified Architecture (OPC-UA)

> A machine-to-machine (M2M) communication protocol for industrial automation. OPC-UA provides an interoperability standard with data encryption, authentication, and authorization schemes.

operational-level agreement (OLA)

> An agreement that clarifies what functional IT groups promise to deliver to each other, to support a service-level agreement (SLA).

operational readiness review (ORR)

> A checklist of questions and associated best practices that help you understand, evaluate, prevent, or reduce the scope of incidents and possible failures. For more information, see Operational Readiness Reviews (ORR) in the AWS Well-Architected Framework.

operational technology (OT)

> Hardware and software systems that work with the physical environment to control industrial operations, equipment, and infrastructure. In manufacturing, the integration of OT and information technology (IT) systems is a key focus for Industry 4.0 transformations.

operations integration (OI)

The process of modernizing operations in the cloud, which involves readiness planning, automation, and integration. For more information, see the operations integration guide.

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the organization and tracks the activity in each account. For more information, see Creating a trail for an organization in the CloudTrail documentation.

organizational change management (OCM)

A framework for managing major, disruptive business transformations from a people, culture, and leadership perspective. OCM helps organizations prepare for, and transition to, new systems and strategies by accelerating change adoption, addressing transitional issues, and driving cultural and organizational changes. In the AWS migration strategy, this framework is called *people acceleration*, because of the speed of change required in cloud adoption projects. For more information, see the OCM guide.

origin access control (OAC)

In CloudFront, an enhanced option for restricting access to secure your Amazon Simple Storage Service (Amazon S3) content. OAC supports all S3 buckets in all AWS Regions, server-side encryption with AWS KMS (SSE-KMS), and dynamic PUT and DELETE requests to the S3 bucket.

origin access identity (OAI)

In CloudFront, an option for restricting access to secure your Amazon S3 content. When you use OAI, CloudFront creates a principal that Amazon S3 can authenticate with. Authenticated principals can access content in an S3 bucket only through a specific CloudFront distribution. See also OAC, which provides more granular and enhanced access control.

ORR

See operational readiness review.

OT

See operational technology.

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The AWS Security Reference Architecture recommends

setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

An IAM management policy that is attached to IAM principals to set the maximum permissions that the user or role can have. For more information, see Permissions boundaries in the IAM documentation.

personally identifiable information (PII)

Information that, when viewed directly or paired with other related data, can be used to reasonably infer the identity of an individual. Examples of PII include names, addresses, and contact information.

PII

See personally identifiable information.

playbook

A set of predefined steps that capture the work associated with migrations, such as delivering core operations functions in the cloud. A playbook can take the form of scripts, automated runbooks, or a summary of processes or steps required to operate your modernized environment.

PLC

See programmable logic controller.

PLM

See product lifecycle management.

policy

An object that can define permissions (see identity-based policy), specify access conditions (see resource-based policy), or define the maximum permissions for all accounts in an organization in AWS Organizations (see service control policy).

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements. For more information, see Enabling data persistence in microservices.

portfolio assessment

A process of discovering, analyzing, and prioritizing the application portfolio in order to plan the migration. For more information, see Evaluating migration readiness.

predicate

A query condition that returns `true` or `false`, commonly located in a WHERE clause.

predicate pushdown

A database query optimization technique that filters the data in the query before transfer. This reduces the amount of data that must be retrieved and processed from the relational database, and it improves query performance.

preventative control

A security control that is designed to prevent an event from occurring. These controls are a first line of defense to help prevent unauthorized access or unwanted changes to your network. For more information, see Preventative controls in *Implementing security controls on AWS*.

principal

An entity in AWS that can perform actions and access resources. This entity is typically a root user for an AWS account, an IAM role, or a user. For more information, see *Principal* in Roles terms and concepts in the IAM documentation.

Privacy by Design

An approach in system engineering that takes privacy into account throughout the whole engineering process.

private hosted zones

A container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs. For more information, see Working with private hosted zones in the Route 53 documentation.

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

product lifecycle management (PLM)

The management of data and processes for a product throughout its entire lifecycle, from design, development, and launch, through growth and maturity, to decline and removal.

production environment

See [environment](#).

programmable logic controller (PLC)

In manufacturing, a highly reliable, adaptable computer that monitors machines and automates manufacturing processes.

pseudonymization

The process of replacing personal identifiers in a dataset with placeholder values. Pseudonymization can help protect personal privacy. Pseudonymized data is still considered to be personal data.

publish/subscribe (pub/sub)

A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based [MES](#), a microservice can publish event messages to a channel that other microservices can subscribe to. The system can add new microservices without changing the publishing service.

Q

query plan

A series of steps, like instructions, that are used to access the data in a SQL relational database system.

query plan regression

When a database service optimizer chooses a less optimal plan than it did before a given change to the database environment. This can be caused by changes to statistics, constraints, environment settings, query parameter bindings, and updates to the database engine.

R

RACI matrix

See responsible, accountable, consulted, informed (RACI).

ransomware

A malicious software that is designed to block access to a computer system or data until a payment is made.

RASCI matrix

See responsible, accountable, consulted, informed (RACI).

RCAC

See row and column access control.

read replica

A copy of a database that's used for read-only purposes. You can route queries to the read replica to reduce the load on your primary database.

re-architect

See 7 Rs.

recovery point objective (RPO)

The maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

recovery time objective (RTO)

The maximum acceptable delay between the interruption of service and restoration of service.

refactor

See 7 Rs.

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see Specify which AWS Regions your account can use.

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

rehost

See 7 Rs.

release

In a deployment process, the act of promoting changes to a production environment.

relocate

See 7 Rs.

replatform

See 7 Rs.

repurchase

See 7 Rs.

resiliency

An application's ability to resist or recover from disruptions. High availability and disaster recovery are common considerations when planning for resiliency in the AWS Cloud. For more information, see AWS Cloud Resilience.

resource-based policy

A policy attached to a resource, such as an Amazon S3 bucket, an endpoint, or an encryption key. This type of policy specifies which principals are allowed access, supported actions, and any other conditions that must be met.

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the

matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

A security control that is designed to drive remediation of adverse events or deviations from your security baseline. For more information, see Responsive controls in *Implementing security controls on AWS*.

retain

See 7 Rs.

retire

See 7 Rs.

rotation

The process of periodically updating a secret to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

The use of basic, flexible SQL expressions that have defined access rules. RCAC consists of row permissions and column masks.

RPO

See recovery point objective.

RTO

See recovery time objective.

runbook

A set of manual or automated procedures required to perform a specific task. These are typically built to streamline repetitive operations or procedures with high error rates.

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API

operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCADA

See [supervisory control and data acquisition](#).

SCP

See [service control policy](#).

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata. The secret value can be binary, a single string, or multiple strings. For more information, see [Secret](#) in the Secrets Manager documentation.

security control

A technical or administrative guardrail that prevents, detects, or reduces the ability of a threat actor to exploit a security vulnerability. There are four primary types of security controls: [preventative](#), [detective](#), [responsive](#), and [proactive](#).

security hardening

The process of reducing the attack surface to make it more resistant to attacks. This can include actions such as removing resources that are no longer needed, implementing the security best practice of granting least privilege, or deactivating unnecessary features in configuration files.

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers, networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

A predefined and programmed action that is designed to automatically respond to or remediate a security event. These automations serve as [detective](#) or [responsive](#) security controls that help you implement AWS security best practices. Examples of automated response actions include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

server-side encryption

Encryption of data at its destination, by the AWS service that receives it.

service control policy (SCP)

A policy that provides centralized control over permissions for all accounts in an organization in AWS Organizations. SCPs define guardrails or set limits on actions that an administrator can delegate to users or roles. You can use SCPs as allow lists or deny lists, to specify which services or actions are permitted or prohibited. For more information, see Service control policies in the AWS Organizations documentation.

service endpoint

The URL of the entry point for an AWS service. You can use the endpoint to connect programmatically to the target service. For more information, see AWS service endpoints in *AWS General Reference*.

service-level agreement (SLA)

An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

service-level indicator (SLI)

A measurement of a performance aspect of a service, such as its error rate, availability, or throughput.

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a service-level indicator.

shared responsibility model

A model describing the responsibility you share with AWS for cloud security and compliance. AWS is responsible for security *of* the cloud, whereas you are responsible for security *in* the cloud. For more information, see Shared responsibility model.

SIEM

See security information and event management system.

single point of failure (SPOF)

A failure in a single, critical component of an application that can disrupt the system.

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

A database organizational structure that uses one large fact table to store transactional or measured data and uses one or more smaller dimensional tables to store data attributes. This structure is designed for use in a [data warehouse](#) or for business intelligence purposes.

strangler fig pattern

An approach to modernizing monolithic systems by incrementally rewriting and replacing system functionality until the legacy system can be decommissioned. This pattern uses the analogy of a fig vine that grows into an established tree and eventually overcomes and replaces its host. The pattern was [introduced by Martin Fowler](#) as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see [Modernizing legacy Microsoft ASP.NET (ASMX) web services incrementally by using containers and Amazon API Gateway](#).

subnet

A range of IP addresses in your VPC. A subnet must reside in a single Availability Zone.

supervisory control and data acquisition (SCADA)

In manufacturing, a system that uses hardware and software to monitor physical assets and production operations.

symmetric encryption

An encryption algorithm that uses the same key to encrypt and decrypt the data.

synthetic testing

Testing a system in a way that simulates user interactions to detect potential issues or to monitor performance. You can use Amazon CloudWatch Synthetics to create these tests.

T

tags

Key-value pairs that act as metadata for organizing your AWS resources. Tags can help you manage, identify, organize, search for, and filter resources. For more information, see Tagging your AWS resources.

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

task list

A tool that is used to track progress through a runbook. A task list contains an overview of the runbook and a list of general tasks to be completed. For each general task, it includes the estimated amount of time required, the owner, and the progress.

test environment

See environment.

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that captures these patterns. You can then use the ML model to make predictions on new data for which you don't know the target.

transit gateway

A network transit hub that you can use to interconnect your VPCs and on-premises networks. For more information, see [What is a transit gateway](#) in the AWS Transit Gateway documentation.

trunk-based workflow

An approach in which developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially.

trusted access

Granting permissions to a service that you specify to perform tasks in your organization in AWS Organizations and in its accounts on your behalf. The trusted service creates a service-linked role in each account, when that role is needed, to perform management tasks for you. For more information, see [Using AWS Organizations with other AWS services](#) in the AWS Organizations documentation.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

two-pizza team

A small DevOps team that you can feed with two pizzas. A two-pizza team size ensures the best possible opportunity for collaboration in software development.

U

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data. For more information, see the [Quantifying uncertainty in deep learning systems](#) guide.

undifferentiated tasks

Also known as *heavy lifting*, work that is necessary to create and operate an application but that doesn't provide direct value to the end user or provide competitive advantage. Examples of undifferentiated tasks include procurement, maintenance, and capacity planning.

upper environments

See environment.

V

vacuuming

A database maintenance operation that involves cleaning up after incremental updates to reclaim storage and improve performance.

version control

Processes and tools that track changes, such as changes to source code in a repository.

VPC peering

A connection between two VPCs that allows you to route traffic by using private IP addresses. For more information, see What is VPC peering in the Amazon VPC documentation.

vulnerability

A software or hardware flaw that compromises the security of the system.

W

warm cache

A buffer cache that contains current, relevant data that is frequently accessed. The database instance can read from the buffer cache, which is faster than reading from the main memory or disk.

warm data

Data that is infrequently accessed. When querying this kind of data, moderately slow queries are typically acceptable.

window function

A SQL function that performs a calculation on a group of rows that relate in some way to the current record. Window functions are useful for processing tasks, such as calculating a moving average or accessing the value of rows based on the relative position of the current row.

workload

A collection of resources and code that delivers business value, such as a customer-facing application or backend process.

workstream

Functional groups in a migration project that are responsible for a specific set of tasks. Each workstream is independent but supports the other workstreams in the project. For example, the portfolio workstream is responsible for prioritizing applications, wave planning, and collecting migration metadata. The portfolio workstream delivers these assets to the migration workstream, which then migrates the servers and applications.

WORM

See write once, read many.

WQF

See AWS Workload Qualification Framework.

write once, read many (WORM)

A storage model that writes data a single time and prevents the data from being deleted or modified. Authorized users can read the data as many times as needed, but they cannot change it. This data storage infrastructure is considered immutable.

Z

zero-day exploit

An attack, typically malware, that takes advantage of a zero-day vulnerability.

zero-day vulnerability

An unmitigated flaw or vulnerability in a production system. Threat actors can use this type of vulnerability to attack the system. Developers frequently become aware of the vulnerability as a result of the attack.

zombie application

An application that has an average CPU and memory usage below 5 percent. In a migration
project, it is common to retire these applications.