

Choosing the right GitOps tool for your Amazon EKS cluster

AWS Prescriptive Guidance



Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Prescriptive Guidance: Choosing the right GitOps tool for your Amazon EKS cluster

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Introduction	1
Targeted business outcomes	2
Seamless integration with Amazon EKS	2
Scalability and performance	2
Security and compliance	2
Ease of use and learning curve	3
Community and network support	3
Multi-cluster management capabilities	3
Observability and monitoring	4
Flexibility and customization	4
Continuous delivery and progressive deployment support	4
Cost-effectiveness and resource utilization	5
GitOps tools for EKS clusters	6
Argo CD	6
GitOps support	6
Architecture	9
Flux	. 10
GitOps support	. 10
Architecture	. 12
Weave GitOps	. 13
GitOps support	. 13
Architecture	. 16
Jenkins X	. 17
GitOps support	. 17
Architecture	. 20
GitLab CI/CD	21
GitOps support	. 21
Spinnaker	24
GitOps support	. 24
Architecture	. 27
Rancher Fleet	. 28
GitOps support	. 29
Architecture	. 32
Codefresh	. 32

	GitOps support	32
	Pulumi	36
	GitOps support	36
Gi	tOps tools comparison	41
	Ease of use	41
	Kubernetes integration	41
	CI/CD capabilities	41
	GitOps purity	41
	Multi-cloud support	42
	Multi-cluster support	42
	Integration	42
	Community and support	42
	Enterprise features	42
	Flexibility and extensibility	43
	Scalability	43
	Infrastructure management	43
	Programming model and language support	43
Ar	go CD and Flux use cases	44
	General considerations	44
	Argo CD use cases	44
	Flux use cases	45
	Feature comparison	47
Be	st practices for choosing a GitOps tool	49
FΑ	Q	54
Re	sources	57
Do	cument history	58
Gl	ossary	59
	#	59
	A	60
	B	63
	C	65
	D	68
	E	72
	F	74
	G	76
	H	77

	78
	80
М	82
0	86
P	88
Q	91
R	
S	94
Γ	98
J	99
V	100
W	
Z	

Choosing the right GitOps tool for your Amazon EKS cluster

Pratap Kumar Nanda and Pradip Kumar Pandey, Amazon Web Services (AWS)

April 2025 (document history)

In the rapidly evolving landscape of cloud-native technologies, GitOps has emerged as a powerful methodology for managing and deploying applications and infrastructure. If you're using <u>Amazon EKS</u>), implementing GitOps principles can significantly enhance your deployment processes, improve reliability, and streamline operations. A variety of GitOps tools are available, and choosing the right one for your EKS cluster is a critical decision that can impact your team's efficiency and the overall success of your DevOps practices.

The selection of an appropriate GitOps tool for your Amazon EKS environment involves careful consideration of various factors, including your specific requirements, team expertise, scalability needs, and integration capabilities with existing AWS services. Each tool comes with its own set of features, strengths, and potential limitations, so it's essential to align your choice with your organization's goals and operational context.

This guide explores the key considerations in selecting GitOps tools for Amazon EKS, compares frequently used options, and provides insights to help you make an informed decision. It covers nine popular GitOps tools:

- Argo CD
- Flux
- Weave GitOps
- Jenkins X
- GitLab CI/CD
- Spinnaker
- Rancher Fleet
- Codefresh
- Pulumi

Targeted business outcomes

The following list discusses potential goals and outcomes when you choose a tool to implement GitOps principles in your development and operations processes.

Seamless integration with Amazon EKS

Your GitOps tool should integrate smoothly with Amazon EKS and provide compatibility with Amazon EKS-specific features and optimizations.

- Native Amazon EKS support: Look for tools that offer built-in support for Amazon EKS, including easy cluster connection and management.
- AWS service integration: Make sure that the tool can interact with other AWS services such as
 <u>AWS Identity and Access Management (IAM)</u>, <u>Amazon Elastic Container Registry (Amazon ECR)</u>,
 and Amazon CloudWatch.
- Amazon EKS add-on compatibility: Confirm that the tool supports <u>Amazon EKS add-ons</u> and can manage them effectively.

Scalability and performance

Your GitOps tool should be able to handle the scale of your Amazon EKS operations, from small clusters to large, multi-cluster environments.

- Resource efficiency: Evaluate the tool's resource consumption and its impact on cluster performance.
- Large-scale operations: Assess the tool's ability to manage numerous applications and clusters simultaneously.
- Performance under load: Consider how the tool performs during high-frequency updates and large-scale deployments.

Security and compliance

Security features and compliance capabilities are crucial, especially in regulated industries or when you handle sensitive data.

Access control: Look for robust role-based access control (RBAC) features that integrate with IAM.

Targeted business outcomes

- Secrets management: Evaluate how the tool handles sensitive information and integrates with AWS Secrets Manager or other solutions.
- Audit trails: Make sure that the tool provides comprehensive logging and auditing capabilities for compliance and troubleshooting.
- Security scanning: Consider tools that offer built-in security scanning for vulnerabilities in deployments.

Ease of use and learning curve

The tool should be user-friendly and align with your team's skills to ensure quick adoption and efficient usage.

- User interface: Assess the intuitiveness of both command line interface (CLI) and graphical user interface (GUI) features.
- Documentation quality: Look for comprehensive, up-to-date documentation and tutorials.
- Learning resources: Consider the availability of training materials, courses, and community resources.

Community and network support

A strong community and network can provide valuable resources, plugins, and long-term sustainability.

- Active development: Check the frequency of updates and the responsiveness of maintainers.
- Community size: Consider the size and activity of the user community for support and knowledge sharing.
- Third-party integrations: Evaluate the availability of plugins and integrations with other tools in your stack.

Multi-cluster management capabilities

If you have multiple EKS clusters, the ability to manage them efficiently is crucial.

• Centralized management: Look for features that allow managing multiple clusters from a single control plane.

- Cluster federation: Consider tools that support Kubernetes federation for multi-cluster applications.
- Environment parity: Assess how well the tool maintains consistency across different environments such as development, staging, and production.

Observability and monitoring

The tool should provide clear insights into the state of your deployments and cluster health.

- Deployment visibility: Look for features that offer clear views of deployment status and history.
- Integration with monitoring tools: Consider how well the tool integrates with popular monitoring solutions such as Prometheus and Grafana.
- Alerting capabilities: Assess the tool's ability to set up and manage alerts for deployment issues
 or drift.

Flexibility and customization

The ability to adapt the tool to your specific workflows and requirements is important for long-term satisfaction.

- Extensibility: Look for plugin architectures or APIs that enable you to extend the tool's functionality.
- Custom resource support: Confirm that the tool can handle custom Kubernetes resources effectively.
- Workflow customization: Assess how easily you can tailor the GitOps workflows to your team's needs.

Continuous delivery and progressive deployment support

Advanced deployment strategies are often crucial for minimizing risk and ensuring smooth updates.

- Canary deployments: Look for built-in support for canary releases.
- Blue/green deployments: Assess the tool's capabilities for blue/green deployment strategies.

• Rollback mechanisms: Ensure robust and easy-to-use rollback features for quick recovery from failed deployments.

Cost-effectiveness and resource utilization

Consider the overall cost of adopting and maintaining the tool, including both direct and indirect costs.

- Licensing costs: Compare open source options with commercial solutions, and consider support and enterprise features.
- Operational overhead: Assess the additional operational costs in terms of management and maintenance.
- Resource consumption: Evaluate the tool's efficiency in terms of compute and storage resources that would be required.

By carefully considering these outcomes and their aspects, you can make an informed decision on the most suitable GitOps tool for your EKS cluster, and make sure that the tool aligns with your organization's needs, capabilities, and long-term strategy.

GitOps tools for EKS clusters

There are several GitOps tools for Kubernetes that are currently available in the market. Here's a list of some of the most widely used options:

- Argo CD
- Flux
- Weave GitOps
- Jenkins X
- GitLab CI/CD
- Spinnaker
- Rancher Fleet
- Codefresh
- Pulumi

Follow the links to see detailed information about how these tools implement GitOps practices. Each tool has strengths and use cases. The choice depends on factors such as your specific requirements, existing infrastructure, team expertise, and desired features. It's important to evaluate these tools based on your organization's needs and the complexity of your Kubernetes environment.

Argo CD

Argo CD is a widely used GitOps continuous delivery (CD) tool for Kubernetes that complies with several key GitOps principles.

GitOps support

Area	Tool capabilities
Declarative configuration	Argo CD uses declarative configurations that are stored in Git repositories. The desired state of the application and infrastructure is defined in YAML files. These configurations describe

Argo CD 6

Area	Tool capabilities
	what should be deployed, not how to deploy them.
Version control system as the single source of truth	Git repositories serve as the single source of truth for the entire system. All changes to the application and infrastructure are made through Git. This ensures a complete audit trail and the ability to roll back to any previous state.
Automated synchronization	Argo CD continuously monitors the Git repository for changes. When changes are detected, it automatically synchronizes the actual state of the cluster with the desired state that's defined in Git. This ensures that the cluster always reflects the state that's described in the repository.
Kubernetes-native	Argo CD is designed specifically for Kubernete s environments. It leverages the declarative nature and custom resources in Kubernetes for managing applications.
Self-healing and drift detection	Argo CD regularly compares the live state of the cluster with the desired state in Git. If it detects any drift (differences between the actual and desired states), it can automatically correct these discrepancies.
Multi-cluster and multi-tenancy support	Argo CD can manage multiple Kubernetes clusters from a single instance. It supports multi-tenancy, so different teams can manage their applications independently.

Area	Tool capabilities
Application definition	Applications in Argo CD are defined by using the Application CRD (custom resource definitio n). This allows for a Kubernetes-native way of defining what should be deployed and how.
Separation of deployment and release	Argo CD separates the deployment of code from its release to users. This is achieved through various deployment strategies such as blue/green or canary deployments.
Observability and auditability	Argo CD provides a web UI and CLI for observing the state of applications and clusters. All actions are logged to provide a clear audit trail of changes and deployments.
Security and RBAC	Argo CD integrates with Kubernetes role-base d access control (RBAC). It supports single sign-on integration for authentication and authorization.
Pluggable architecture	Argo CD supports various source control management systems, Helm charts, Kustomize , and other Kubernetes manifest formats. This flexibility allows it to fit into diverse environments and workflows.
Continuous delivery (CD)	Although Argo CD focuses on continuous delivery, it can be integrated with continuous integration (CI) tools to create a complete CI/CD pipeline.

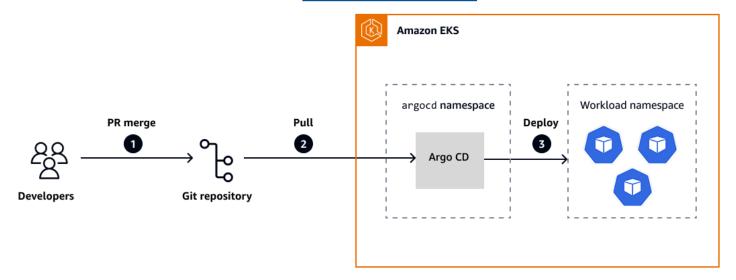
By adhering to these GitOps principles, Argo CD provides a robust, scalable, and secure way to manage Kubernetes deployments. It ensures that the operational state of your system is always in sync with the desired state that's defined in your Git repository, and promotes consistency, reliability, and ease of management in complex Kubernetes environments.

For scenarios and requirements that Argo CD can address, see <u>Argo CD use cases</u> later in this guide. For a comparison between Argo CD and Flux, see <u>Feature comparison</u> later in this guide.

For additional information, see the Argo CD documentation.

Architecture

The following diagram illustrates a GitOps-driven CD workflow that uses Argo CD within an EKS cluster. For detailed information, see the Argo CD documentation.



where:

- Step 1: Pull request (PR) merge. A developer commits changes to Kubernetes manifests or Helm charts that are stored in a Git repository. When the PR has been reviewed and merged into the main branch, the desired state of the application is updated in source control.
- **Step 2: Repository sync**. Argo CD runs within a dedicated namespace (argocd) in the EKS cluster and continuously monitors the configured Git repository. When it detects changes, it pulls the latest updates to reconcile the declared state.
- Step 3: Deployment to target namespace. Argo CD compares the desired state from Git with the live state in the cluster. It then applies the necessary changes to the target workload namespace so that the application is deployed or updated accordingly. This includes managing Kubernetes resources such as deployments, services, ConfigMaps, and secrets to maintain cluster consistency with the Git source of truth.

Architecture 9

Flux

Flux is another tool for Kubernetes that implements GitOps principles in a unique way.

GitOps support

Area	Tool capabilities
Git as the single source of truth	Flux uses Git repositories as the definitive source for defining the desired state of the system. All configuration for applications and infrastructure is stored in Git.
Declarative configuration	Flux works with declarative descriptions of the desired state of your cluster. These descriptions are typically Kubernetes manifests, Helm charts, or Kustomize overlays.
Automated aynchronization	Flux continuously monitors the Git repositor y for changes. When it detects changes, it automatically applies them to the cluster.
Kubernetes-native	Flux is built as a set of Kubernetes controlle rs and custom resources. It uses the extension mechanisms in Kubernetes to provide GitOps capabilities.
Pull-based deployment model	Unlike traditional push-based CI/CD systems, Flux uses a pull-based model. The cluster pulls the desired state from Git instead of using an external system to push changes.
Continuous reconciliation	Flux constantly compares the actual state of the cluster with the desired state in Git. It automatically corrects any drift that's detected between these states.
Multi-tenancy	Flux supports multi-tenancy through its concepts of <i>Kustomizations</i> and <i>HelmReleases</i> .

Flux 10

Area	Tool capabilities
	Different teams can manage their own parts of the configuration independently.
Progressive delivery	Flux supports advanced deployment strategie s, such as canary releases and A/B testing, through its Flagger component.
Helm integration	Flux includes native support for Helm, so you can easily manage Helm releases through GitOps.
Image update automation	Flux can automatically update container images in Git when new versions are available in the container registry.
Kustomize support	You can use the native support provided by Flux for Kustomize to customize and patch Kubernetes manifests.
Security and RBAC	Flux integrates with Kubernetes RBAC for access control. It supports secrets managemen t through various backends.
Observability	Flux provides status information and metrics about reconciliation and operations. It integrates with monitoring tools for enhanced observability.
Event-driven architecture	Flux uses an event-driven approach to implement reconciliations and updates.
Extensibility	The tool is designed to be extensible, so you can add custom controllers and resources.
Cross-cluster synchronization	Flux supports the management of multiple clusters from a single set of repositories.

Area	Tool capabilities
Dependency management	It allows for defining dependencies between different parts of your system and ensures the correct order of operations.
Webhook receivers	You can configure Flux to receive webhooks from Git providers or other systems to start immediate reconciliation.

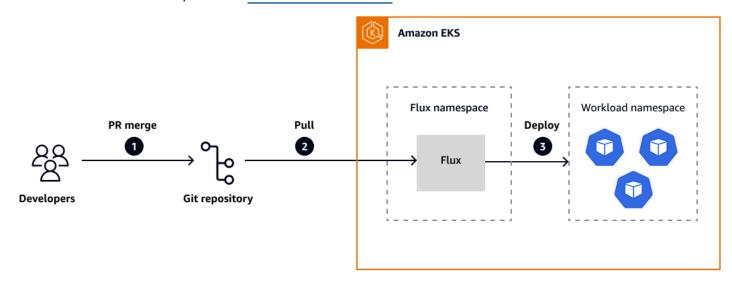
By implementing these GitOps principles, Flux provides a robust and flexible system for managing Kubernetes clusters and applications. It ensures that your infrastructure and applications are always in sync with your Git repositories, and provides consistency, reliability, and ease of management in complex Kubernetes environments. The tool's Kubernetes-native approach and focus on automation make it particularly well-suited for cloud-native environments.

For scenarios and requirements that Flux can address, see <u>Flux use cases</u> later in this guide. For a comparison between Argo CD and Flux, see <u>Feature comparison</u> later in this guide.

For additional information, see the Flux documentation.

Architecture

The following diagram illustrates a GitOps-driven CD workflow that uses Flux within an EKS cluster. For detailed information, see the Flux documentation.



where:

Architecture 12

- Step 1: Pull request (PR) merge. A developer commits changes to Kubernetes manifests or Helm charts that are stored in a Git repository. When the PR has been reviewed and merged into the main branch, the desired state of the application is updated in source control.
- Step 2: Repository sync. Flux runs within a dedicated namespace in the EKS cluster and continuously monitors the configured Git repository. When it detects changes, it pulls the latest updates to reconcile the declared state.
- Step 3: Deployment to target namespace. Flux compares the desired state from Git with the live state in the cluster. It then applies the necessary changes to the target workload namespace so that the application is deployed or updated accordingly.

Weave GitOps

Weave GitOps was developed by Weaveworks, which is the company that introduced the term *GitOps*. This tool provides a comprehensive GitOps solution that builds upon the core GitOps principles.

GitOps support

Area	Tool capabilities
Git as the single source of truth	Weave GitOps uses Git repositories as the authoritative source for defining the desired state of the system. All configurations, including application manifests, infrastructure definitions, and policies, are stored in Git.
Declarative configuration	The system relies on declarative descriptions of the entire system state. These descriptions ons are typically Kubernetes manifests, Helm charts, or other declarative formats.
Automated synchronization	Weave GitOps continuously monitors Git repositories for changes. When it detects changes, it automatically applies them to the target environment.

Weave GitOps 13

Area	Tool capabilities
Kubernetes-native architecture	Weave GitOps is built as a set of Kubernete s controllers and custom resources. It uses the extension mechanisms in Kubernetes to provide GitOps capabilities.
Continuous reconciliation	This tool constantly compares the actual state of the cluster with the desired state that's defined in Git. It automatically corrects any drift detected between these states.
Multi-cluster management	Weave GitOps supports the management of multiple Kubernetes clusters from a single control plane. It enables consistent application deployment across different environments.
Policy as code	Weave GitOps incorporates the concept of policy as code for enforcing security and compliance rules. Policies are version-c ontrolled alongside application code and infrastructure definitions.
Progressive delivery	This tool supports advanced deployment strategies such as canary releases and blue/ green deployments. It integrates with Flagger for automated, progressive delivery.
Observability and dashboards	Weave GitOps provides built-in dashboards for monitoring the state of applications and clusters. It offers insights into reconciliation processes and cluster health.
Secure by design	The tool implements security best practices , including RBAC integration and secrets management. It supports various authentic ation methods and integrates with enterprise identity providers.

Area	Tool capabilities
Extensibility and integration	The tool is designed to work with a wide range of cloud-native tools. It supports popular tools such as Flux, Helm, and Kustomize.
Self-service developer platforms	Weave GitOps enables the creation of self- service platforms for developers. It provides templates and guardrails for application deployment.
GitOps automation	The tool automates many aspects of the GitOps workflow, including pull request generation for updates.
Continuous delivery pipelines	It integrates with CI/CD systems to create end- to-end delivery pipelines.
Audit and compliance	Weave DevOps provides a complete audit trail of all changes and actions. It helps you meet compliance requirements through version control and automated processes.
Scalability	The tool is designed to scale from small projects to large, enterprise-grade deploymen ts.
Team collaboration	Weave GitOps facilitates collaboration between development and operations teams through Git-based workflows.
GitOps as a service	This tool offers GitOps as a managed service, which simplifies adoption and management.
Hybrid and multi-cloud support	Weave GitOps enables consistent managemen t across different cloud providers and onpremises environments.

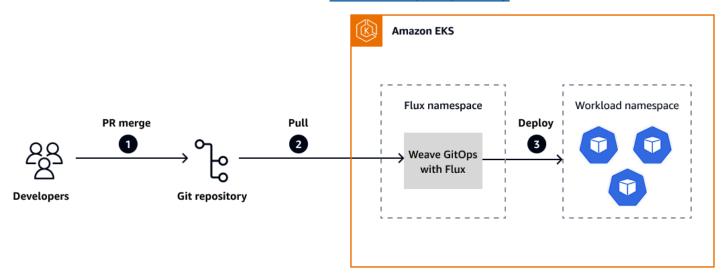
Area	Tool capabilities
Continuous security	The tool integrates security scanning and policy enforcement throughout the deployment process.

Weave GitOps implements these principles to provide a comprehensive GitOps solution that goes beyond basic deployment automation. It aims to create a complete operational model for cloudnative applications that focuses on security, scalability, and ease of use. By adhering to these GitOps principles, Weave GitOps helps organizations achieve consistent, auditable, and efficient management of their Kubernetes environments across multiple clusters and cloud providers.

For more information, see the Weave GitOps documentation.

Architecture

The following diagram illustrates a GitOps-driven CD workflow that uses Weave GitOps within an EKS cluster. For detailed information, see the Weave GitOps repository.



where:

• Step 1: Pull request (PR) merge. A developer commits changes to Kubernetes manifests or Helm charts that are stored in a Git repository. When the PR has been reviewed and merged into the main branch, the desired state of the application is updated in source control.

Architecture 16

- **Step 2: Repository sync**. Weave GitOps runs within the Flux namespace in the EKS cluster and continuously monitors the configured Git repository. When it detects changes, it pulls the latest updates to reconcile the declared state.
- Step 3: Deployment to target namespace. Weave GitOps compares the desired state from Git with the live state in the cluster. It then applies the necessary changes to the target workload namespace so that the application is deployed or updated accordingly.

Jenkins X

Jenkins X is a cloud-native, open-source CI/CD platform that implements GitOps principles for Kubernetes environments. Although Jenkins X isn't exclusively a GitOps tool such as Argo CD or Flux, it incorporates GitOps practices into its workflows.

GitOps support

Area	Tool capabilities
Git-centric workflow	Jenkins X uses Git repositories as the primary source of truth for both application code and configuration. All changes to applications and infrastructure are made through Git.
Environment as code (EaC)	Environments (such as staging and production) are defined as code in Git repositories. This allows for version control and review of environment configurations.
Automated CI/CD pipelines	Jenkins X automatically sets up CI/CD pipelines for projects. These pipelines are defined as code (pipeline as code) and stored in Git.
Kubernetes-native	Jenkins X is built specifically for Kubernetes environments. It uses Kubernetes resources and custom resource definitions (CRDs).

Jenkins X 17

Area	Tool capabilities
Preview environments	Jenkins X automatically creates temporary environments for pull requests. It enables easy review and testing of changes before merges.
Promotion between environments	Jenkins X uses a GitOps approach to promote applications between environments (for example, from staging to production). Promotions are handled by using pull requests to ensure proper review and approval processes.
Helm chart management	Jenkins X uses Helm charts to package and deploy applications. Charts are version-c ontrolled in Git repositories.
Automated versioning	Jenkins X automatically manages versioning of applications and releases. It uses semantic versioning and generates release notes.
ChatOps integration	Jenkins X supports ChatOps for common operations. This aligns with the GitOps principles of automation and collaboration.
Extensibility	This tool provides a plugin system for extending functionality. It allows integration with various cloud-native tools.
Infrastructure as code (IaC)	Jenkins X supports Terraform, AWS CloudForm ation, AWS Cloud Development Kit (AWS CDK), and other IaC tools to define and manage infrastructure. Infrastructure definitions are version-controlled alongside application code.
Automated rollbacks	Jenkins X supports automated rollbacks if issues are detected after deployment.

Area	Tool capabilities
Secrets management	The tool integrates with external secrets management solutions to handle sensitive information securely.
Observability	Jenkins X provides integration with monitorin g and logging tools for observability.
Multi-cloud support	Jenkins X is designed to work across different cloud providers and on-premises environme nts.
Team collaboration	This tool encourages collaboration through Git-based workflows and pull requests.
Continuous feedback	The tool provides quick feedback on changes through automated testing and preview environments.
DevOps best practices	Jenkins X implements DevOps best practices by default, including GitOps principles.
Declarative configuration	The tool uses declarative configurations for defining applications and environments.
Automated upgrades	Jenkins X provides tools for automating upgrades of the Jenkins X platform itself.

Jenkins X implements these GitOps principles to create a comprehensive CI/CD solution for Kubernetes. It aims to automate and streamline the entire software delivery process, from code commit to production deployment, while adhering to GitOps practices. By doing so, it helps teams achieve faster, more reliable, and more consistent deployments in cloud-native environments.

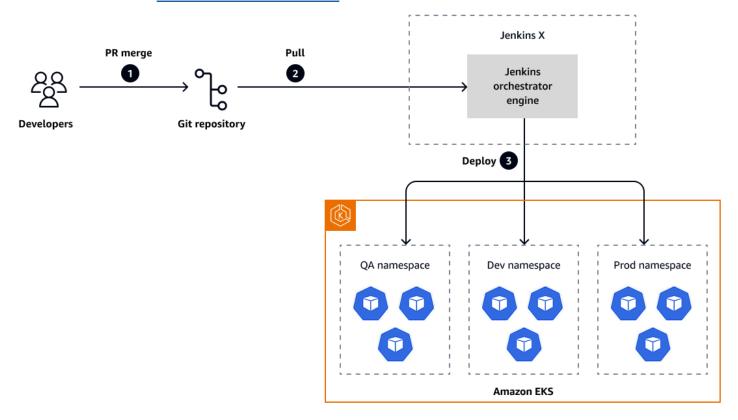
The key difference between Jenkins X and tools such as Argo CD or Flux is that Jenkins X provides a more comprehensive CI/CD solution, including build automation and pipeline management, while still incorporating GitOps principles for deployment and environment management. This makes it

particularly suitable for teams who need an all-in-one solution that covers both CI and CD aspects within a single GitOps framework.

For more information, see the Jenkins X documentation.

Architecture

The following diagram illustrates a GitOps-driven CD workflow that uses Jenkins X. For detailed information, see the Jenkins X documentation.



where:

- Step 1: Pull request (PR) merge. A developer creates a pull request that includes changes to Kubernetes manifests, Helm charts, or application code that's stored in a Git repository. After review and approval, the PR is merged into the main branch and updates the desired state in source control.
- **Step 2: Repository sync**. Jenkins X automatically triggers a CI/CD pipeline when it detects the change. The pipeline builds, tests, and promotes the application through different environments (for example, staging and production) by using GitOps principles.
- **Step 3: Deployment to target namespaces**. Jenkins X updates the environment repositories (staging and production) with the new application versions. The cluster automatically reconciles

Architecture 20

the changes by pulling the latest manifests from Git and deploying the application to the appropriate namespaces.

GitLab CI/CD

GitLab CI/CD is an integrated part of the GitLab platform that provides continuous integration, delivery, and deployment capabilities. Although GitLab CI/CD isn't exclusively a GitOps tool, you can configure it to implement GitOps principles, especially when you use it for Kubernetes deployments.

GitOps support

Area	Tool capabilities
Git as the single source of truth	GitLab CI/CD uses Git repositories to store both application code and infrastructure configurations. All changes to the system are made through Git, which ensures a complete history and audit trail.
Declarative configuration	GitLab CI/CD pipelines are defined in a .gitlab- ci.yml file, which is a declarative configura tion stored in the Git repository. Kubernetes manifests, Helm charts, or other infrastructure as code (IaC) files can be stored in the same repository to define the desired state of the infrastructure.
Automated pipelines	GitLab CI/CD automatically triggers pipelines when changes are pushed to the repositor y. These pipelines can include stages for building, testing, and deploying applications.
Kubernetes integration	GitLab CI/CD provides native Kubernete s integration and supports GitOps-style deployments to Kubernetes clusters. It can

GitLab CI/CD 21

Area	Tool capabilities
	automatically create and manage Kubernetes resources based on the configuration in Git.
Environment management	GitLab CI/CD supports the definition of multiple environments (such as staging and production) as code. Deployments to these environments can be automated or might require manual approval, in compliance with GitOps practices.
Review applications	GitLab can automatically create temporary environments for merge requests, similar to preview environments in other GitOps tools. This supports easy review and testing of changes before merges.
Continuous deployment	GitLab CI/CD can be configured to automatic ally deploy changes to Kubernetes clusters when changes are merged to specific branches.
IaC	GitLab CI/CD supports integration with tools such as Terraform and AWS CloudFormation to manage infrastructure as code. Infrastructure definitions can be version-controlled alongside application code.
Observability and monitoring	GitLab CI/CD provides built-in monitoring and observability features, including integration with Prometheus and Grafana.
Security scanning	GitLab CI/CD includes built-in security scanning tools that can be integrated into the CI/CD pipeline to enforce security as part of the GitOps workflow.

Area	Tool capabilities
Container registry	GitLab CI/CD includes a built-in container registry for seamless integration of container image management in the GitOps workflow.
Auto DevOps	The Auto DevOps feature in GitLab CI/CD can automatically configure CI/CD pipelines that follow GitOps principles for Kubernetes deployments.
Approval workflows	GitLab CI/CD supports approval processes for deployments, which provide controlled promotions between environments.
Secrets management	GitLab CI/CD provides features to securely manage and use secrets within CI/CD pipelines .
Versioning and releases	GitLab CI/CD supports automatic versioning and release management as part of the CI/CD process.
Rollbacks	GitLab CI/CD enables easy rollbacks to previous versions if issues are detected after deployment.
Audit logs	GitLab CI/CD provides comprehensive audit logs for all actions to support the traceability aspect of GitOps.
Multi-project pipelines	GitLab CI/CD supports complex GitOps workflows that span multiple projects or repositories.
ChatOps	GitLab CI/CD supports ChatOps integrations, which provide collaboration and operations through chat interfaces.

Area	Tool capabilities
Kubernetes cluster management	GitLab CI/CD provides features for managing Kubernetes clusters directly from the GitLab interface.

Although GitLab CI/CD is not exclusively designed for GitOps, it can be used effectively to implement GitOps practices, especially for teams that already use GitLab as their primary development platform. Its integrated approach, which combines source control, CI/CD, and Kubernetes management, makes it a powerful tool for implementing GitOps workflows.

The key difference between GitLab CI/CD and dedicated GitOps tools such as Argo CD or Flux is that GitLab provides a more comprehensive platform that includes source control management, issue tracking, and other development tools along with its CI/CD capabilities. This makes it particularly suitable for teams that need an all-in-one solution that can implement GitOps practices within a broader development system.

For more information about GitLab CI/CD and its architecture, see the <u>GitLab CI/CD</u> documentation.

Spinnaker

Although Spinnaker isn't exclusively designed as a GitOps tool, you can configure it to implement GitOps principles, especially when you use it for cloud-native and Kubernetes deployments.

GitOps support

Area	Tool capabilities
Declarative configuration	Spinnaker uses declarative pipeline definitio ns, which are typically stored as JSON or YAML files. These pipeline definitions can be version-controlled in Git repositories, in alignment with GitOps practices.
laC	Spinnaker supports the definition of infrastru cture and deployment configurations as code.

Spinnaker 24

Area	Tool capabilities
	These definitions can be stored in Git repositor ies, and can serve as the single source of truth.
Multi-cloud deployments	Spinnaker is designed to work across multiple cloud providers and Kubernetes clusters. It enables consistent GitOps practices across diverse environments.
Pipeline as code	Spinnaker pipelines can be defined as code and stored in Git repositories. This allows for version control and the review of deployment processes.
Automated deployments	You can configure Spinnaker to automatic ally start deployments based on changes in Git repositories. The tool supports continuous deployment practices that are central to GitOps.
Immutable infrastructure	Spinnaker promotes the use of immutable infrastructure, which is a key concept in GitOps. It encourages the deployment of new instances instead of modifying existing instances.
Rollbacks and versioning	Spinnaker provides robust rollback capabilities and quick reversion to previous known good states. It supports versioning of deploymen ts, in alignment with GitOps principles of traceability.
Approval workflows	Spinnaker includes manual judgment stages in pipelines to support controlled promotions between environments. This supports GitOps practices of separation between deployments and releases.

Area	Tool capabilities
Canary and blue/green deployments	Spinnaker supports advanced deployment strategies that align with GitOps practices for safe and controlled releases.
Integration with version control systems	Spinnaker can integrate with various Git providers to start pipelines based on repositor y events.
Kubernetes integration	Spinnaker provides native support for Kubernetes and supports GitOps-style management of Kubernetes resources.
Artifact management	Spinnaker supports artifact management and versioning, which are crucial for maintaining a GitOps workflow.
Observability and monitoring	Spinnaker offers integration with monitorin g tools to support the observability aspect of GitOps.
Audit trail	Spinnaker provides detailed deployment logs and history, which support the auditability principle of GitOps.
Role-based access control (RBAC)	This tool implements RBAC for fine-grained control over who can perform which actions, in alignment with GitOps security practices.
Templating and parameterization	Spinnaker supports templating in pipeline definitions to enable reusable and parameter ized deployments.
Environment promotion	Spinnaker facilitates the promotion of applications between environments (for example, from staging to production) in a controlled manner.

Area	Tool capabilities
Integration with CI tools	Spinnaker can integrate with various continuous integration (CI) tools to provide a complete CI/CD pipeline that adheres to GitOps principles.
Custom stages and extensions	This tool supports custom stages and extensions, so teams can implement GitOps workflows that are tailored to their needs.
Centralized management	Spinnaker provides a centralized platform for managing deployments across multiple environments and cloud providers.

Although Spinnaker isn't primarily marketed as a GitOps tool, its flexibility and robust feature set make it capable of implementing GitOps workflows, especially in complex, multi-cloud environments. The key difference between Spinnaker and dedicated GitOps tools such as Argo CD or Flux is that Spinnaker offers a more comprehensive continuous delivery platform with advanced deployment strategies and multi-cloud support.

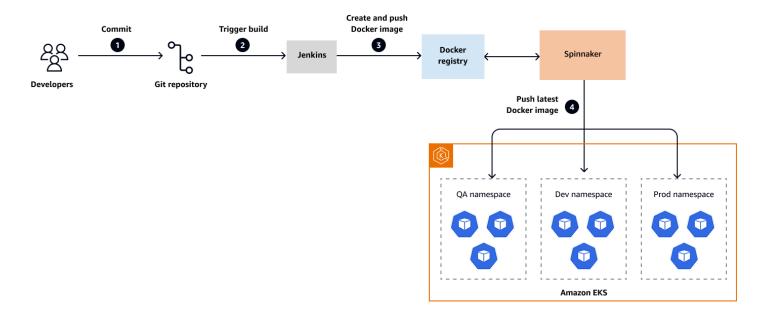
Spinnaker's strength lies in its ability to handle complex deployment scenarios across various cloud providers and its support for advanced deployment strategies. When Spinnaker is configured properly, it can implement GitOps principles effectively. This makes it a powerful tool for organizations who want to adopt GitOps practices in diverse and complex environments.

For more information, see the Spinnaker documentation.

Architecture

The following diagram illustrates a GitOps-driven CD workflow that uses Spinnaker and Jenkins X. For detailed information, see the Spinnaker documentation.

Architecture 27



where:

- **Step 1: Code commit.** Developers commit application code changes to a Git repository. These changes could include updates to the application itself, Dockerfiles, or Kubernetes manifests.
- Step 2: Jenkins build and image creation. Jenkins is automatically triggered by the Git repository through a webhook or polling. Jenkins builds the application, creates a Docker image, and pushes the built image to a configured Docker registry (such as Amazon ECR or Docker Hub).
- Step 3: Spinnaker image monitoring and pipeline trigger. Spinnaker continuously monitors the Docker registry for new images. When a new image version is detected, Spinnaker automatically triggers a pipeline to start the deployment process.
- Step 4: Deployment to target namespaces. Spinnaker deploys the new Docker image to Amazon EKS. Based on pipeline configurations, the image is deployed to target namespaces in the cluster. Spinnaker ensures that the latest application version is deployed while adhering to defined deployment strategies such as blue/green or canary deployments.

Rancher Fleet

Rancher Fleet is a GitOps-at-scale solution that is designed specifically for managing multiple Kubernetes clusters. It adheres closely to GitOps principles while focusing on scalability and multi-cluster management.

Rancher Fleet 28

Area	Tool capabilities
Git as the single source of truth	Fleet uses Git repositories as the authorita tive source for defining the desired state of applications and resources across multiple clusters. All configurations, including Kubernetes manifests, Helm charts, and custom resources, are stored in Git.
Declarative configuration	Fleet works with declarative descriptions of the desired state for applications and resources. These can be raw Kubernetes YAML, Helm charts, Kustomize files, or Fleet-specific custom resources.
Automated synchronization	Fleet continuously monitors Git repositories for changes. It automatically applies changes to the target clusters when it detects differenc es between the Git state and the cluster state.
Multi-cluster management	Fleet is designed specifically to manage deployments across multiple Kubernetes clusters. It can handle thousands of clusters from a single control plane.
Kubernetes-native architecture	Fleet is built as a set of Kubernetes custom resources and controllers. It uses the extension mechanisms in Kubernetes for GitOps operations.
Continuous reconciliation	Fleet constantly compares the actual state of clusters with the desired state that's defined in Git. It automatically corrects any drift detected between these states.

Area	Tool capabilities
Cluster grouping and targeting	Fleet enables you to group clusters and target deployments to specific groups or individua l clusters. It supports consistent application deployment across different environments and cluster types.
Layered configurations	Fleet supports layered configurations, which provide base configurations with environme nt-specific overlays. This aligns with GitOps practices of managing multiple environments efficiently.
Helm integration	Fleet provides native support for Helm charts and provides easy management of complex applications. It can version and manage Helm releases through GitOps workflows.
Custom resource definitions (CRDs)	Fleet uses custom resources such as GitRepo and Bundle to define deployments. These CRDs provide a Kubernetes-native way to define GitOps workflows.
Security and RBAC	Fleet integrates with Kubernetes RBAC for access control. It supports secure management of sensitive information and credentials.
Observability	Fleet provides status information about the synchronization state of clusters and applicati ons. It offers insights into GitOps processes across the fleet of clusters.
Scalability	Fleet is designed to scale to manage thousands of clusters efficiently. It supports large-scale GitOps operations in enterprise environments.

Area	Tool capabilities
Dependency management	You can define dependencies between different resources and applications. Fleet ensures that the correct order of operations is followed in complex deployments.
Customization and extensibility	Fleet supports custom scripts and lifecycle hooks for advanced customization of deployments. It allows for integration with existing tools and workflows.
Offline and air-gapped support	Fleet can operate in environments that have limited or no internet connectivity. It supports GitOps workflows in high-security or regulated environments.
Progressive rollouts	Fleet supports staged rollouts across clusters, which allow for controlled and gradual deployment strategies.
Unified management interface	Fleet provides a single interface for managing GitOps workflows across all clusters. It simplifies operations in complex, multi-cluster environments.
Integration with other Rancher tools	Fleet integrates with other Rancher tools to provide a comprehensive Kubernetes management solution.
Audit trail and compliance	Fleet maintains a clear audit trail of all changes and deployments. It helps you meet compliance requirements through version-c ontrolled, Git-based operations.

Rancher Fleet implements these GitOps principles with a strong focus on scalability and multicluster management. Its design is particularly suited for organizations that manage large numbers of Kubernetes clusters across different environments, data centers, or cloud providers.

The key differentiator of Fleet is its ability to handle GitOps at scale. This feature makes it especially valuable for large enterprises or managed service providers that manage numerous clusters. Tools such as Argo CD or Flux are often used for individual cluster management, whereas Fleet is designed to manage GitOps across a large fleet of clusters.

By adhering to these GitOps principles, Rancher Fleet provides a solution for organizations that want to implement consistent, scalable, and automated management of applications and resources across a diverse and large-scale Kubernetes environment.

For more information, see the Fleet documentation.

Architecture

For architecture and workflow information, see the Fleet repository.

Codefresh

Codefresh is a modern CI/CD platform that supports GitOps principles, particularly for Kubernetes deployments. Codefresh offers a comprehensive set of CI/CD features, and its GitOps capabilities are notable.

GitOps support

Area	Tool capabilities
Git as the single source of truth	Codefresh uses Git repositories as the authoritative source for application code, infrastructure definitions, and pipeline configurations. All changes to the system are made through Git, which ensures a complete history and audit trail.
Declarative configuration	Codefresh supports declarative pipeline definitions by using YAML files that are stored in Git. Kubernetes manifests, Helm charts, AWS CloudFormation templates, and other IaC files can be version-controlled in the same repositories.

Architecture 32

Area	Tool capabilities
GitOps dashboard	Codefresh provides a dedicated GitOps dashboard for visualizing and managing GitOps workflows. It offers a clear view of the synchronization status between Git and cluster states.
Automated synchronization	Codefresh continuously monitors Git repositor ies for changes. It automatically starts pipelines to apply changes to the target environments when it detects differences.
Kubernetes integration	Codefresh offers deep integration with Kubernetes to support GitOps-style deployments across multiple clusters. It supports various Kubernetes resources and custom resource definitions (CRDs).
Environment management	You can define and manage multiple environments (such as development, staging, and production) as code. Codefresh supports promotion between environments by using GitOps practices.
Argo CD integration	Codefresh integrates with Argo CD for enhanced GitOps capabilities. It combines its CI capabilities with the CD strengths of Argo CD to provide a complete GitOps solution.
Helm support	Codefresh supports Helm charts, and provides easy management of complex applications through GitOps. It also offers Helm chart versioning and promotion.

Area	Tool capabilities
Progressive delivery	Codefresh supports advanced deploymen t strategies such as canary and blue/green deployments. You can implement and manage these strategies through GitOps workflows.
Rollbacks and versioning	Codefresh enables easy rollbacks to previous versions if issues are detected after deployment. It maintains deployment versioning for traceability.
Approval workflows	Codefresh supports manual and automated approval processes for deployments. It enables controlled promotions between environments, in compliance with GitOps practices.
IaC	Codefresh supports integration with IaC tools such as AWS CloudFormation and Terraform . It enables version control of infrastructure definitions alongside application code.
Observability and monitoring	Codefresh provides built-in monitoring and observability features. It also offers integrations with external monitoring tools for enhanced visibility.
Security scanning	Codefresh includes security scanning capabilit ies that can be integrated into GitOps workflows. Security checks are part of the automated deployment process.
Audit trails	Codefresh maintains comprehensive audit logs for all actions and changes. It supports the traceability and compliance aspects of GitOps.

Area	Tool capabilities
RBAC and access control	Codefresh implements role-based access control (RBAC) for fine-grained permissions management. This helps ensure secure GitOps operations across teams and environments.
GitOps automation	Codefresh offers features to automate various aspects of GitOps workflows, including pull request (PR) creation and merging.
Multi-cloud and hybrid deployments	Codefresh supports GitOps workflows across multiple cloud providers and on-premises environments.
Templating and parameterization	Codefresh supports templates in pipeline and deployment configurations. This enables reusable and parameterized GitOps workflows.
Integrated image management	Codefresh provides built-in container image management capabilities. It integrates image builds and deployments into GitOps workflows.
GitOps for secrets management	Codefresh offers secure ways to manage secrets within GitOps workflows. It integrates with external secrets management solutions.
Collaboration features	Codefresh provides features for team collaboration within GitOps processes. These features include commenting, notifications, and shared dashboards.

The Codefresh approach to GitOps is notable for its integration of CI/CD capabilities with GitOps practices. It aims to provide a comprehensive platform that covers the entire software delivery lifecycle while adhering to GitOps principles.

The key differentiator of Codefresh in the GitOps area is its unified platform approach, which combines CI capabilities with CD and GitOps features. This makes it particularly suitable for teams that want an all-in-one solution that can handle complex CI/CD scenarios while implementing GitOps practices.

Codefresh offers a platform for organizations that want to adopt GitOps methodologies within a broader CI/CD context, especially when working with Kubernetes and cloud-native technologies.

For more information, see the Codefresh documentation.

Pulumi

Pulumi is an IaC platform that isn't exclusively designed for GitOps. However, it can be used effectively to implement GitOps principles, especially for cloud infrastructure and Kubernetes deployments.

GitOps support

Area	Tool capabilities
IaC	Pulumi lets you define your infrastructure by using general-purpose programming languages such as Python, TypeScript, and Go. This code-based approach aligns with the GitOps emphasis on versioned, declarative configurations.
Git as the single source of truth	Infrastructure code in Pulumi can be stored and version-controlled in Git repositories. This ensures that Git serves as the single source of truth for infrastructure definitions.
Declarative desired state	Although Pulumi uses programming languages, it still describes the desired state of infrastructure declaratively. The code defines what the infrastructure should look like, not the step-by-step process to create it.

Pulumi 36

Area	Tool capabilities
Automated synchronization	Pulumi can be integrated with CI/CD pipelines to automatically apply changes when code is updated in Git. This enables continuous deployment of infrastructure changes, which is a key GitOps principle.
Multi-cloud and Kubernetes support	Pulumi supports a wide range of cloud providers and Kubernetes, so you can follow GitOps practices across diverse environments. The tool enables consistent management of resources across different platforms.
State management	Pulumi manages the state of infrastructure, which can be stored remotely and securely. This state management is crucial for GitOps practices, and ensures consistency between the defined state and the actual state of your infrastructure.
Drift detection and reconciliation	Pulumi can detect differences between the desired state (in code) and the actual state of infrastructure. It reconciles these differences in alignment with the GitOps principle for continuous reconciliation.
Policy as code	You can define and enforce policies as code by using Pulumi CrossGuard. This enables version-controlled, GitOps-style management of compliance and security policies.
Secrets management	Pulumi provides secure ways to manage sensitive information within infrastructure code. It supports integration with external secrets management systems, which is crucial for GitOps security practices.

Area	Tool capabilities
Modular and reusable components	Pulumi supports the creation of reusable components and modules. This modularity aligns with GitOps practices for managing complex, multi-environment deployments.
Preview and plan	Pulumi offers the ability to preview changes before applying them. This supports the GitOps principle of safe, predictable changes to infrastructure.
Rollbacks and history	Pulumi maintains a history of deployments and supports rollbacks to previous states. This aligns with GitOps principles of traceability and reversibility.
Continuous delivery for infrastructure	Pulumi can be integrated into CI/CD pipelines for continuous delivery of infrastructure changes. It supports automated testing and validation of infrastructure code.
RBAC and access control	Pulumi provides role-based access control for managing who can make changes to infrastru cture. This supports GitOps security and governance practices.
Observability and logging	Pulumi offers logging and monitoring capabilities for infrastructure changes. These capabilities support the observability aspect of GitOps practices.
Integration with other tools	Pulumi can integrate with various tools in the cloud. This flexibility allows for comprehensive GitOps workflows.

Area	Tool capabilities
Environment management	Pulumi supports the management of multiple environments (development, staging, production) by using the same codebase with different configurations. This aligns with GitOps practices for consistent multi-environment management.
Dependency management	Pulumi handles dependencies between resources, and ensures the correct order of operations. This is crucial for complex GitOps deployments that involve interdependent components.
Custom resource providers	Pulumi lets you create custom providers to manage any API-driven service. This extends GitOps practices to a wide range of resources beyond standard cloud offerings.
Collaboration features	Pulumi supports team collaboration through shared state and access controls. This facilitat es GitOps workflows in team environments.

By using these Pulumi features, organizations can implement GitOps practices for their infrastructure, especially in scenarios where they need fine-grained control or complex logic, or want to manage a diverse set of cloud and on-premises resources within a single, consistent framework.

Pulumi's approach to GitOps is unique because it brings the power and flexibility of general-purpose programming languages to infrastructure management while adhering to GitOps principles. This can be particularly advantageous for teams that prefer to work with familiar programming languages and want to apply software engineering best practices to infrastructure management.

The key differentiator of Pulumi in GitOps is its use of standard programming languages to define infrastructure. Traditional GitOps tools often use YAML or domain-specific languages, whereas

Pulumi allows for more complex logic, better code reuse, and easier integration with existing development workflows.

For more information, see the Pulumi documentation.

GitOps tools comparison

Here's a comparison of the nine GitOps tools that were discussed in the previous sections. When you choose a tool, consider your specific requirements, existing infrastructure, team expertise, and desired level of control and customization.

Ease of use

- Argo CD, Flux, and Rancher Fleet are generally easier to set up.
- Spinnaker and Jenkins X have steeper learning curves.
- Weave GitOps might require more setup for advanced features.
- GitLab CI/CD and Codefresh offer integrated experiences.

Kubernetes integration

- Argo CD, Flux, and Rancher Fleet are very Kubernetes-centric.
- Jenkins X and Weave GitOps offer broader DevOps capabilities.
- The other tools support Kubernetes without an exclusive focus on it.

CI/CD capabilities

- Jenkins X, GitLab CI/CD, and Codefresh offer complete CI/CD solutions.
- Argo CD, Flux, and Weave GitOps focus more on the CD aspect of the workflow, and often require integration with separate CI tools.

GitOps purity

- Argo CD and Flux are tools that focus specifically on GitOps.
- The other tools incorporate GitOps principles to varying degrees.

Ease of use 41

Multi-cloud support

- Spinnaker and Pulumi excel in multi-cloud scenarios.
- The other tools can work across clouds but might require additional setup.

Multi-cluster support

- All tools support multi-cluster deployments.
- Argo CD and Weave GitOps have more advanced multi-cluster management features.

Integration

- Flux has strong Cloud Native Computing Foundation (CNCF) backing.
- Argo CD has a large and active community.
- Argo CD and Flux have strong Kubernetes integration.
- Jenkins X uses the broader Jenkins system.
- Weave GitOps is newer but is growing with strong commercial backing.
- GitLab CI/CD integrates tightly with GitLab.
- Rancher Fleet works well within the Rancher system.

Community and support

- Flux has strong CNCF backing.
- Argo CD, GitLab, and Spinnaker have large communities.
- Commercial support is available for most tools.

Enterprise features

- Weave GitOps and Jenkins X offer more enterprise-focused features by default.
- Argo CD and Flux have enterprise offerings or can be extended for enterprise use.

Multi-cloud support 42

Flexibility and extensibility

- · Flux is highly modular and extensible.
- Argo CD offers good customization options.
- Jenkins X is very extensible but might require more effort.
- Weave GitOps aims to provide a complete solution with less need for extensibility.

Scalability

- Spinnaker and GitLab CI/CD are known for enterprise scalability.
- Argo CD and Flux handle large-scale Kubernetes deployments well.

Infrastructure management

- Pulumi focuses on infrastructure management.
- Weave GitOps and Flux offer good IaC capabilities.

Programming model and language support

- In Pulumi, you can define infrastructure by using general-purpose programming languages such as Python, Go, TypeScript, C#, and Java. Pulumi's use of standard languages enables integration of infrastructure code with familiar development workflows, testing practices, and complex logic.
- Terraform uses HashiCorp Configuration Language (HCL).
- AWS CloudFormation uses JSON and YAML templates.
- Argo CD, Flux, Rancher Fleet, Weave GitOps, Spinnaker, and GitLab CI/CD primarily manage YAML or declarative configuration files.
- Jenkins X manages YAML and scripting-based pipelines but doesn't natively offer generalpurpose programming for IaC.

Flexibility and extensibility 43

Argo CD and Flux use cases

This section focuses on two tools, Argo CD and Flux, which provide pure GitOps functionality. In this context, *pure GitOps* refers to a model where a Git repository serves as the single source of truth for the desired state of applications and infrastructure. All changes are made through Git commits, and the system automatically synchronizes the live environment to match the state that's defined in the repository. No manual intervention is required outside of Git operations.

General considerations

- You might prefer to use Argo CD in environments where visual management and applicationcentric workflows are important.
- You might choose Flux if you require lighter-weight solutions, strong multi-tenancy, or deep integration with the broader Cloud Native Computing Foundations (CNCF) network.
- Argo CD often appeals to teams that are transitioning from traditional CI/CD to GitOps because
 of its intuitive UI.
- Flux is often favored in cloud-native environments where CLI-based workflows and IaC practices are already established.

Ultimately, the choice between Argo CD and Flux often depends on your specific organizational needs, existing tooling, and team preferences. Both tools are capable of handling most GitOps scenarios, so we recommend that you evaluate them based on your specific use cases and requirements.

Argo CD use cases

Visual management:

- When you need a user-friendly UI to manage deployments and visualize application states.
- For teams that prefer a graphical interface for monitoring and troubleshooting.

Application-centric approach:

• When you want to manage deployments at the application level instead of managing individual resources.

General considerations 44

• For organizations that structure their deployments around application concepts.

Multi-cluster management:

- When managing deployments across multiple clusters is a primary requirement.
- For complex, distributed environments with many clusters.

Rollback and sync waves:

- When you need fine-grained control over the deployment process, including sync waves and manual interventions.
- For scenarios that require complex rollback strategies.

Integration with existing tools:

 When you're already using other tools in the Argo project such as Argo Workflows and Argo Events.

Enterprise environments:

For large enterprises that need robust RBAC and single sign-on integration by default.

Flux use cases

Lightweight deployments:

- When you need a more lightweight, less resource-intensive GitOps solution.
- For edge computing or IoT scenarios where resources might be constrained.

Automated image updates:

- When automatic detection and deployment of new container images is a key requirement.
- For teams that focus on continuous deployment with frequent image updates.

Multi-tenancy:

Flux use cases 45

- When strong multi-tenancy support is needed, especially in shared cluster environments.
- For service providers or large organizations that have strict separations between teams or projects.

IaC:

- When managing both applications and infrastructure through the same GitOps workflow is important.
- For teams that are heavily invested in the IaC paradigm.

Helm integration:

- When extensive use of Helm charts is part of your deployment strategy.
- For environments that have complex Helm-based deployments.

CNCF project integration:

- When tight integration with other CNCF projects is important.
- For organizations that align with CNCF technologies and principles.

Modular architecture:

- When you need flexibility to use only specific components of the GitOps toolkit.
- For teams that want to build custom GitOps workflows by using modular components.

Progressive delivery:

 When advanced deployment strategies such as canary releases or A/B testing are core requirements.

Flux use cases 46

Feature comparison

Area	Argo CD	Flux	
Support for core GitOps principles	⊘	\odot	Yes
Architecture	End-to-end application for implementing Kubernetes GitOps workflows	Provides Kubernetes CRDs and controllers for GitOps	
Setup	Simple	Complex	
Helm support	⊘	⊘	Yes
Kustomize support	⊗	\odot	Yes
Integrated GUI	CLI and fully-featured web UI	CLI and optional lightweight web interface	
RBAC support	Granular control	Kubernetes native RBAC	
Multi-tenancy and multi-clu ster support	Excellent support for multi- cluster	Excellent support for multi- tenancy	
Single sign-on authentication	⊘	\odot	Yes
Sync automation	Ability to sync windows	Ability to set reconcillation intervals	

Feature comparison 47

Area	Argo CD	Flux	
Partial sync	⊘		۷o
Reconciliation process	Supports manual and automatic syncs. Several different strategies are available.	Supports manual and automatic syncs.	
Extentensibility	Supports custom lugins. Limited customization options.	Supports custom controller. Good extensibility and third- party integrations.	
Cummunity support	Large and active community.	Smaller but growing community.	
Scalability	Good scalability, but limited by the web UI's data-fetching rate. Community analysis suggests support for tens of thousands of applications.	Clear guides for horizontal and vertical scalability, up to tens of thousands of applications.	

Feature comparison 48

Best practices for choosing a GitOps tool

This section provides considerations, tips, and best practices for choosing a GitOps tool for your EKS cluster. The right choice depends on your specific context, requirements, and long-term strategy. It's often beneficial to conduct a proof of concept with your top choices before you make a final decision.

Assess your organization's needs and capabilities:

- Consider your team's current skill set and willingness to learn new tools.
- Evaluate the complexity of your Amazon EKS environment. (For example, are you using a single cluster or multiple clusters?)
- Determine your specific requirements for compliance, security, and scalability.

Best practice

Create a detailed requirements document that outlines required features and useful, but not required, capabilities.

Evaluate tool maturity and adoption:

- Research the maturity of potential GitOps tools and their adoption rates in the industry.
- Look for tools that have a proven track record in Amazon EKS environments.



Best practice

Prioritize tools that have been widely adopted and have a strong presence in the Cloud Native Computing Foundation (CNCF) network.

Consider integration with your existing toolchain:

- Assess how well the GitOps tool integrates with your current CI/CD pipeline, monitoring solutions and other operational tools.
- Look for native integrations with AWS services such as IAM, Amazon ECR, and CloudWatch.



Create a proof of concept to test integration capabilities before you make a final decision.

Evaluate security features:

- Prioritize tools that have robust role-based access control (RBAC) capabilities and integrate well with IAM.
- Look for features that support secure secrets management and policy enforcement.



Best practice

Choose a tool that supports GitOps-based security practices, including policy as code and automated compliance checks.

Assess scalability and performance:

- Consider how the tool performs with large numbers of applications and clusters.
- Evaluate its impact on cluster performance and resource consumption.



Best practice

Conduct performance testing with workloads that are similar to your production environment to make sure that the tool can handle your scale.

Consider multi-cluster and multi-environment support:

- If you have, or plan to have, multiple EKS clusters, prioritize tools that have strong multi-cluster management capabilities.
- Look for features that support consistent deployments across different environments (such as development, staging, and production).



Choose a tool that allows for centralized management of multiple clusters while maintaining environment-specific configurations.

Evaluate observability and monitoring capabilities:

- Look for tools that provide clear visibility into the state of your deployments and cluster health.
- Consider how well the tool integrates with your existing monitoring and logging solutions.



Best practice

Prioritize tools that offer customizable dashboards and alerting mechanisms for proactive issue detection.

Assess the learning curve and documentation:

- Evaluate the quality and comprehensiveness of the tool's documentation.
- Consider the availability of training resources and community support.



Best practice

Choose a tool that has well-maintained documentation, active community forums, and official training programs or certifications.

Consider cost and resource utilization:

- Evaluate both the direct costs (such as licensing and support) and indirect costs (such as operational overhead and training costs) of adopting the tool.
- Assess the tool's efficiency in terms of compute and storage resource consumption.



Perform a total cost of ownership (TCO) analysis that includes both short-term and longterm costs.

Evaluate flexibility and customization options:

- Look for tools that let you customize workflows to fit your specific needs.
- Consider the extensibility of the tool through plugins or APIs.



Best practice

Choose a tool that balances default functionality with the ability to customize for your unique requirements.

Assess continuous delivery and progressive deployment capabilities:

- Look for tools that support advanced deployment strategies such as canary releases and blue/ green deployments.
- Evaluate the ease of implementing and managing these strategies.



Best practice

Prioritize tools that offer built-in support for progressive delivery patterns to minimize risk in your deployments.

Consider vendor lock-in and portability:

- Assess the tool's dependencies on specific cloud providers or technologies.
- Consider the ease of migrating to a different tool in the future if needed.



Favor tools that use open standards and provide export capabilities for your GitOps configurations.

Evaluate community support and extensions:

- Look at the size and activity of the user community.
- Assess the availability of third-party integrations and plugins.



Best practice

Join community forums or user groups to get firsthand experiences from other users before you make a decision.

Consider compliance and audit requirements:

- Evaluate how well the tool supports your compliance needs, including audit trails and reporting.
- Look for features that help maintaining and demonstrate compliance.



Best practice

Choose a tool that provides comprehensive audit logs and supports the generation of compliance reports.

Assess rollback and disaster recovery capabilities:

- Evaluate the ease and reliability of rollback mechanisms.
- Consider how the tool supports disaster recovery scenarios.



Best practice

Test rollback and recovery processes thoroughly as part of your evaluation.

FAQ

Q: What are the most popular GitOps tools for Amazon EKS?

A: The most popular GitOps tools for Amazon EKS include Argo CD, Flux, Jenkins X, and GitLab CI/CD. Each tool has strengths, but Argo CD and Flux are particularly well-regarded for their Kubernetes-native approach and strong community support.

Q: How does GitOps improve EKS cluster management?

A: GitOps improves EKS cluster management by providing version control for infrastructure, automated deployments, improved security through declarative configurations, easier rollbacks, and better auditability. It also enhances collaboration and reduces human error in deployments.

Q: What key features should I look for in a GitOps tool for Amazon EKS?

A: Key features to look for include: seamless Amazon EKS integration, robust RBAC, multi-cluster support, observability features, support for progressive delivery strategies, scalability, and integration with AWS services such as IAM and Amazon ECR.

Q: How do I ensure security when implementing GitOps in Amazon EKS?

A: To ensure security, choose a tool that has strong RBAC integration with IAM, secure secrets management, support for encrypted Git repositories, and the ability to implement security policies as code. Also, verify that the tool provides comprehensive audit logs.

Q: Can GitOps tools handle multi-cluster Amazon EKS environments?

A: Yes, GitOps tools such as <u>Argo CD</u> and <u>Flux</u> have robust multi-cluster management capabilities. They allow you to manage multiple EKS clusters from a single control plane, which ensures consistency across environments.

Q: How do GitOps tools integrate with existing CI/CD pipelines?

A: GitOps tools typically integrate with existing CI/CD pipelines by acting as the deployment stage of the pipeline. They can be triggered by CI tools when changes are pushed to the Git repository, and they automate the deployment process to EKS clusters.

Q: What are the challenges of implementing GitOps in Amazon EKS?

A: Common challenges include managing secrets securely, ensuring proper access controls, handling stateful applications, managing drift between Git and cluster state, and adapting team workflows to the GitOps model.

Q: How do GitOps tools handle rollbacks in Amazon EKS?

A: GitOps tools typically handle rollbacks by reverting to a previous commit in the Git repository. This automatically triggers a deployment of the previous known good state, which results in fast and reliable rollbacks.

Q: Can GitOps tools manage Amazon EKS add-ons and other AWS resources?

A: Many GitOps tools can manage Amazon EKS add-ons and some AWS resources, especially when they're combined with IaC tools such as Terraform or AWS CloudFormation. However, the extent of this capability can vary; see the GitOps tools section for specific information about each tool.

Q: How do GitOps tools support compliance requirements in Amazon EKS?

A: GitOps tools support compliance by providing a clear audit trail of all changes, enforcing approval processes, implementing policy as code for automated compliance checks, and offering detailed logging and reporting features.

Q: What's the learning curve for implementing GitOps in Amazon EKS?

A: The learning curve can vary depending on the tool and your team's existing knowledge. Generally, teams that are familiar with Git, Kubernetes, and Amazon EKS will adapt more quickly than others. Most popular tools offer extensive documentation and training resources to ease adoption.

Q: How do GitOps tools handle secrets management in Amazon EKS?

A: GitOps tools typically integrate with external secrets management solutions such as AWS Secrets Manager or HashiCorp Vault. Some tools also offer built-in encryption for secrets that are stored in Git repositories.

Q: Can GitOps tools work with both stateless and stateful applications in Amazon EKS?

A: Yes, GitOps tools can work with both stateless and stateful applications. However, managing stateful applications often requires additional considerations, such as handling persistent volumes and ensuring data consistency during updates.

Q: How do GitOps tools support canary or blue/green deployments in Amazon EKS?

A: Many GitOps tools offer built-in support for advanced deployment strategies. They can manage the gradual rollout of new versions, monitor for issues, and automatically roll back if problems are detected. All of these operations are defined as code in the Git repository.

Q: What's the difference between using a GitOps tool and using kubect1 apply with a CI/CD pipeline?

A: GitOps tools offer advantages over simple kubectl apply commands, including automated drift detection and reconciliation, improved security through pull-based deployments, better auditability, and more sophisticated deployment strategies. They also provide a more comprehensive approach to managing the entire cluster state.

Resources

The following resources provide official documentation, practical guides, case studies, and in-depth analyses that can help you make an informed decision when choosing a GitOps tool for your EKS cluster. They cover various aspects of GitOps, including implementation strategies, best practices, comparisons between different tools, and real-world experiences.

AWS resources:

- Amazon EKS documentation
- Automating Amazon EKS with GitOps (AWS blog post)
- Introduction to GitOps on EKS with Weaveworks (AWS workshop)
- Flux lab (Amazon EKS workshop)
- Argo CD lab (Amazon EKS workshop)

GitOps and tool documentation:

- GitOps Best Practices for Continuous Deployment and Progressive Security (on-demand DevOps.com webinar)
- · Kubernetes documentation
- Argo CD documentation
- Flux documentation
- Weave GitOps documentation
- Jenkins X documentation
- GitLab CI/CD documentation
- Spinnaker documentation
- Rancher Fleet documentation
- Codefresh documentation

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	_	April 30, 2025

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 onpremises 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.
 You migrate servers from an on-premises platform to a cloud service for the same platform.
 Example: Migrate a Microsoft Hyper-V application to 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.

59 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 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.

ΑI

See artificial intelligence.

AIOps

See artificial intelligence operations.

Ā 60

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 <u>operations</u> 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.

A 61

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 <u>ABAC for AWS</u> 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.

A 62

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.

B 63

botnet

Networks of <u>bots</u> that are infected by <u>malware</u> 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 <u>About branches</u> (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 <u>Implement break-glass procedures</u> 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 <u>Organized around business capabilities</u> section of the <u>Running containerized microservices on AWS</u> 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.

B 64



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 <u>AWS Fault Injection Service (AWS FIS)</u> 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.

C 65

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 <u>The Journey Toward Cloud-First</u> & 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 Bitbucket Cloud. 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.

C 66

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 <u>AI</u> that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, Amazon SageMaker AI 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 Deployment.

C 67

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 <u>star schema</u>, 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 <u>disaster</u>. For more information, see <u>Disaster Recovery of Workloads on AWS: Recovery in the Cloud</u> 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.

EDI

See electronic data interchange.

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with <u>cloud computing</u>, edge computing can reduce communication latency and improve response time.

electronic data interchange (EDI)

The automated exchange of business documents between organizations. For more information, see What is Electronic Data Interchange.

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 <u>service endpoint</u>.

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

E 72

information, see <u>Create an endpoint service</u> 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, <u>MES</u>, 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.

E 73

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 <u>star schema</u>. 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.

F 74

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.

few-shot prompting

Providing an <u>LLM</u> with a small number of examples that demonstrate the task and desired output before asking it to perform a similar task. This technique is an application of in-context learning, where models learn from examples (*shots*) that are embedded in prompts. Few-shot prompting can be effective for tasks that require specific formatting, reasoning, or domain knowledge. See also <u>zero-shot</u> prompting.

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 <u>change data</u> <u>capture</u> to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

FΜ

See foundation model.

foundation model (FM)

A large deep-learning neural network that has been training on massive datasets of generalized and unlabeled data. FMs are capable of performing a wide variety of general tasks, such as understanding language, generating text and images, and conversing in natural language. For more information, see What are Foundation Models.

F 75

G

generative AI

A subset of <u>AI</u> models that have been trained on large amounts of data and that can use a simple text prompt to create new content and artifacts, such as images, videos, text, and audio. For more information, see What is Generative AI.

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 <u>Restricting the geographic distribution of your content</u> 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 <u>trunk-based workflow</u> is the modern, preferred approach.

golden image

A snapshot of a system or software that is used as a template to deploy new instances of that system or software. For example, in manufacturing, a golden image can be used to provision software on multiple devices and helps improve speed, scalability, and productivity in device manufacturing operations.

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.

G 76

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.

Н

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 rearchitecting effort, and converting the schema can be a complex task. <u>AWS provides AWS SCT</u> 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.

holdout data

A portion of historical, labeled data that is withheld from a dataset that is used to train a machine learning model. You can use holdout data to evaluate the model performance by comparing the model predictions against the holdout data.

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.

77 T

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.

lloT

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 <u>mutable infrastructure</u>. For more information, see the <u>Deploy using immutable infrastructure</u> best practice in the AWS Well-Architected Framework.

78

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The <u>AWS Security Reference Architecture</u> 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 <u>Klaus Schwab</u> 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 <u>Building an industrial Internet of Things (IIoT) digital transformation strategy</u>.

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.

79

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.

L 80

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 language model (LLM)

A deep learning <u>AI</u> model that is pretrained on a vast amount of data. An LLM can perform multiple tasks, such as answering questions, summarizing documents, translating text into other languages, and completing sentences. For more information, see <u>What are LLMs</u>.

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.

LLM

See large language model.

lower environments

See environment.

81

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 <u>publish/</u> 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 <u>AWS migration strategy</u>.

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 <u>discussion of migration</u> 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 <u>migration readiness guide</u>. MRA is the first phase of the <u>AWS migration strategy</u>.

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the <u>7 Rs</u> 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.



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).

0 86

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 <u>Industry 4.0</u> transformations.

operations integration (OI)

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

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.

O 87

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 <u>AWS Security Reference Architecture</u> 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 <u>Permissions boundaries</u> 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.

P 88

PLM

See product lifecycle management.

policy

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

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 <u>Evaluating migration readiness</u>.

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 <u>Preventative controls</u> 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.

P 89

privacy by design

A system engineering approach that takes privacy into account through the whole development 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 <u>security control</u> 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 <u>Controls reference guide</u> in the AWS Control Tower documentation and see <u>Proactive controls</u> in <u>Implementing security controls on AWS</u>.

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.

prompt chaining

Using the output of one <u>LLM</u> prompt as the input for the next prompt to generate better responses. This technique is used to break down a complex task into subtasks, or to iteratively refine or expand a preliminary response. It helps improve the accuracy and relevance of a model's responses and allows for more granular, personalized results.

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.

P 90

publish/subscribe (pub/sub)

A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based <u>MES</u>, 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).

RAG

See Retrieval Augmented Generation.

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.

Q 91

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.

R 92

repurchase

See 7 Rs.

resiliency

An application's ability to resist or recover from disruptions. <u>High availability</u> and <u>disaster</u> 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 <u>Responsive controls</u> in *Implementing security controls on AWS*.

retain

See 7 Rs.

retire

See 7 Rs.

Retrieval Augmented Generation (RAG)

A <u>generative AI</u> technology in which an <u>LLM</u> references an authoritative data source that is outside of its training data sources before generating a response. For example, a RAG model might perform a semantic search of an organization's knowledge base or custom data. For more information, see What is RAG.

R 93

rotation

The process of periodically updating a <u>secret</u> 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 What's in a Secrets Manager secret? in the Secrets Manager documentation.

security by design

A system engineering approach that takes security into account through the whole development process.

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 <u>detective</u> or <u>responsive</u> 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 <u>Service control policies</u> 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 <u>AWS service endpoints</u> 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 <u>service-level indicator</u>. 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 <u>Shared responsibility model</u>.

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 <u>introduced by Martin Fowler</u> as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see <u>Modernizing legacy Microsoft ASP.NET (ASMX) web services incrementally by using containers and Amazon API Gateway</u>.

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 <u>Amazon CloudWatch Synthetics</u> to create these tests.

system prompt

A technique for providing context, instructions, or guidelines to an <u>LLM</u> to direct its behavior. System prompts help set context and establish rules for interactions with users.

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 <u>Tagging</u> 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 <u>What is a transit gateway</u> 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.

T 98

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 <u>Using AWS Organizations with other AWS services</u> 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 <u>Quantifying uncertainty in</u> 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.

U 99



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.

V 100

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 <u>immutable</u>.

Z

zero-day exploit

An attack, typically malware, that takes advantage of a <u>zero-day vulnerability</u>. 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.

zero-shot prompting

Providing an <u>LLM</u> with instructions for performing a task but no examples (*shots*) that can help guide it. The LLM must use its pre-trained knowledge to handle the task. The effectiveness of zero-shot prompting depends on the complexity of the task and the quality of the prompt. See also <u>few-shot prompting</u>.

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.

Z 101