

Practicing Continuous Integration and Continuous Delivery on AWS

Accelerating Software Delivery with DevOps

June 2017



Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

The Challenge of Software Delivery	1
What is Continuous Integration and Continuous Delivery/Deployment?	2
Continuous Integration	2
Continuous Delivery and Deployment	2
Continuous Delivery Is Not Continuous Deployment	3
Benefits of Continuous Delivery	3
Automate the Software Release Process	4
Improve Developer Productivity	4
Improve Code Quality	4
Deliver Updates Faster	4
Implementing Continuous Integration and Continuous Delivery	5
A Pathway to Continuous Integration/Continuous Delivery	5
Teams	9
Testing Stages in Continuous Integration and Continuous Delivery	11
Building the Pipeline	14
Pipeline Integration with AWS CodeBuild	22
Pipeline Integration with Jenkins	23
Deployment Methods	24
All at Once (In-Place Deployment)	25
Rolling Deployment	26
Immutable and Blue/Green Deployment	26
Database Schema Changes	27
Summary of Best Practices	27
Conclusion	29
Further Reading	29
Contributors	30

Abstract

This paper explains the features and benefits of using continuous integration, continuous delivery (CI/CD), and Amazon Web Services (AWS) tooling in your software development environment. Continuous integration and continuous delivery are best practices and a vital part of a DevOps initiative.

The Challenge of Software Delivery

Enterprises today face the challenge of rapidly changing competitive landscapes, evolving security requirements, and performance scalability. Enterprises must bridge the gap between operations stability and rapid feature development. Continuous integration and continuous delivery (CI/CD) is a practice that enables rapid software changes while maintaining system stability and security.

Amazon realized early on that the business needs of delivering features for Amazon.com retail customers, Amazon subsidiaries, and Amazon Web Services (AWS) would require new and innovative ways of delivering software. At the scale of a company like Amazon, thousands of independent software teams must be able to work in parallel to deliver software quickly, securely, reliably, and with zero tolerance for outages.

By learning how to deliver software at high velocity, Amazon and other forward-thinking organizations pioneered DevOps.¹ DevOps is the combination of cultural philosophies, practices, and tools that increases an organization's ability to deliver applications and services at high velocity. Using DevOps principles, organizations can evolve and improve products at a faster pace than organizations that use traditional software development and infrastructure management processes. This speed enables organizations to better serve their customers and compete more effectively in the market.

Some of these principles, such as two-pizza teams² and microservices/service oriented architecture (SOA), are out of scope of this whitepaper. In this whitepaper we discuss the CI/CD capability that Amazon has built and continuously improved. CI/CD is key to delivering software features rapidly and reliably.

AWS now offers these CI/CD capabilities as a set of developer services: [AWS CodeStar](#),³ [AWS CodeCommit](#),⁴ [AWS CodePipeline](#),⁵ [AWS CodeBuild](#),⁶ and [AWS CodeDeploy](#).⁷ Developers and IT operations professionals practicing DevOps can use these services to rapidly and safely securely deliver software. Together, they help you securely store and apply version control to your application's source code. You can use AWS CodeStar to rapidly orchestrate an end-to-end software release workflow using these services. For an existing environment, AWS CodePipeline has the flexibility to integrate each service

independently with your existing tools. These are highly available, easily integrated services that can be accessed through the AWS Management Console, AWS APIs, and AWS SDKs like any other AWS service.

What is Continuous Integration and Continuous Delivery/Deployment?

This section discusses the practices of continuous integration and continuous delivery. We explain the difference between continuous delivery and continuous deployment.

Continuous Integration

Continuous integration (CI) is a software development practice where developers regularly merge their code changes into a central repository, after which automated builds and tests are run. CI most often refers to the build or integration stage of the software release process and requires both an automation component (e.g., a CI or build service) and a cultural component (e.g., learning to integrate frequently). The key goals of CI are to find and address bugs more quickly, improve software quality, and reduce the time it takes to validate and release new software updates.

Continuous integration focuses on smaller commits and smaller code changes to integrate. A developer commits code at regular intervals, at minimum once a day. The developer pulls code from the code repository to ensure the code on the local host is merged before pushing to the build server. At this stage the build server runs the various tests and either accepts or rejects the code commit.

The basic challenges of implementing CI include more frequent commits to the common codebase, maintaining a single source code repository, automating builds, and automating testing. Additional challenges include testing in similar environments to production, providing visibility of the process to the team, and allowing developers to easily obtain any version of the application.

Continuous Delivery and Deployment

Continuous delivery (CD) is a software development practice where code changes are automatically built, tested, and prepared for production release. It

expands on continuous integration by deploying all code changes to a testing environment, a production environment, or both after the build stage has been completed. Continuous delivery can be fully automated with a workflow process or partially automated with manual steps at critical points. When continuous delivery is properly implemented, developers always have a deployment-ready build artifact that has passed through a standardized test process.

With continuous deployment, revisions are deployed to a production environment automatically without explicit approval from a developer, making the entire software release process automated. This, in turn, allows for a continuous customer feedback loop early in the product lifecycle.

Continuous Delivery Is Not Continuous Deployment

One misconception about continuous delivery is that it means every change committed is applied to production immediately after passing automated tests. However, the point of continuous delivery is not to apply every change to production immediately, but to ensure that every change is ready to go to production.

Before deploying a change to production, you can implement a decision process to ensure that the production deployment is authorized and audited. This decision can be made by a person and then executed by the tooling.

Using continuous delivery the decision to go live becomes a business decision, not a technical one. The technical validation happens on every commit.

Rolling out a change to production is not a disruptive event. Deployment doesn't require the technical team to stop working on the next set of changes, and it doesn't need a project plan, handover documentation, or a maintenance window. Deployment becomes a repeatable process that has been carried out and proven multiple times in testing environments.⁸

Benefits of Continuous Delivery

CD provides numerous benefits for your software development team including automating the process, improving developer productivity, improving code quality, and delivering updates to your customers faster.

Automate the Software Release Process

CD provides a method for your team to check in code that is automatically built, tested, and prepared for release to production so that your software delivery is efficient, resilient, rapid, and secure.

Improve Developer Productivity

CD practices help your team's productivity by freeing developers from manual tasks, untangling complex dependencies, and returning focus to delivering new features in software. Instead of integrating their code with other parts of the business and spending cycles on how to deploy this code to a platform, developers can focus on coding logic that delivers the features you need.

Improve Code Quality

CD can help you discover and address bugs early in the delivery process before they grow into larger problems later. Your team can easily perform additional types of code tests because the entire process has been automated. With the discipline of more testing more frequently, teams can iterate faster with immediate feedback on the impact of changes. This enables teams to drive quality code with a high assurance of stability and security. Developers will know through immediate feedback whether the new code works and whether any breaking changes or bugs were introduced. Mistakes caught early on in the development process are the easiest to fix.

Deliver Updates Faster

CD helps your team deliver updates to customers quickly and frequently. When CI/CD is implemented, the velocity of the entire team, including the release of features and bug fixes, is increased. Enterprises can respond faster to market changes, security challenges, customer needs, and cost pressures. For example, if a new security feature is required, your team can implement CI/CD with automated testing to introduce the fix quickly and reliably to production systems with high confidence. What used to take weeks and months can now be done in days or even hours.

Implementing Continuous Integration and Continuous Delivery

In this section we discuss the ways in which you can begin to implement a CI/CD model in your organization. This whitepaper doesn't discuss how an organization with a mature DevOps and cloud transformation model builds and uses a CI/CD pipeline. To help you on your DevOps journey, AWS has a number of [certified DevOps Partners](#)⁹ who can provide resources and tooling. For more information on preparing for a move to the AWS Cloud, see the [AWS Cloud Transformation Maturity whitepaper](#).¹⁰

A Pathway to Continuous Integration/Continuous Delivery

CI/CD can be pictured as a pipeline (see Figure 1), where new code is submitted on one end, tested over a series of stages (source, build, staging, and production), and then published as production-ready code. If your organization is new to CI/CD it can approach this pipeline in an iterative fashion. This means that you should start small, and iterate at each stage so that you can understand and develop your code in a way that will help your organization grow.




Figure 1: CI/CD pipeline

Each stage of the CI/CD pipeline is structured as a logical unit in the delivery process. In addition, each stage acts as a gate that vets a certain aspect of the code. As the code progresses through the pipeline, the assumption is that the quality of the code is higher in the later stages because more aspects of it continue to be verified. Problems uncovered in an early stage stop the code from progressing through the pipeline. Results from the tests are immediately sent to the team, and all further builds and releases are stopped if software does not pass the stage.

These stages are suggestions. You can adapt the stages based on your business need. Some stages can be repeated for multiple types of testing, security, and performance. Depending on the complexity of your project and the structure of your teams, some stages can be repeated several times at different levels. For example, the end product of one team can become a dependency in the project of the next team. This means that the first team's end product is subsequently staged as an artifact in the next team's project.

The presence of a CI/CD pipeline will have a large impact on maturing the capabilities of your organization. The organization should start with small steps and not try to build a fully mature pipeline, with multiple environments, many testing phases, and automation in all stages at the start. Keep in mind that even organizations that have highly mature CI/CD environments still need to continuously improve their pipelines.

Building a CI/CD-enabled organization is a journey, and there are many destinations along the way. In the next section we discuss a possible pathway that your organization could take, starting with continuous integration through the levels of continuous delivery.

Continuous Integration

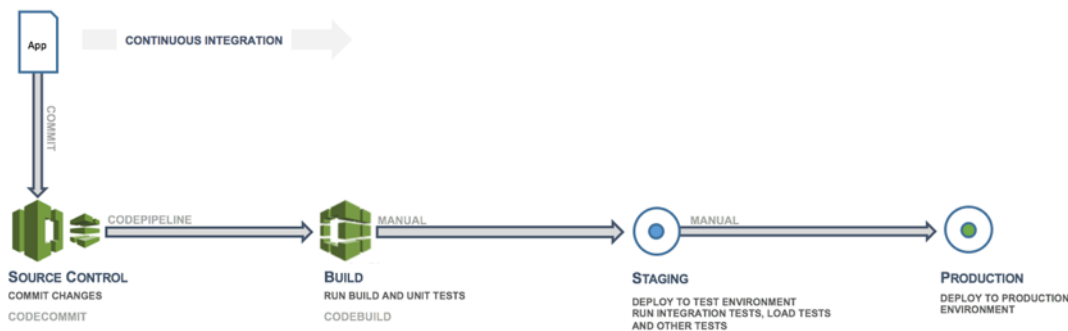


Figure 2: Continuous integration—source and build

The first phase in the CI/CD journey is to develop maturity in continuous integration. You should make sure that all of the developers regularly commit their code to a central repository (such as one hosted in CodeCommit or GitHub) and merge all changes to a release branch for the application. No developer should be holding code in isolation. If a feature branch is needed for a certain period of time, it should be kept up to date by merging from upstream as

often as possible. Frequent commits and merges with complete units of work are recommended for the team to develop discipline and are encouraged by the process. A developer who merges code early and often will likely have fewer integration issues down the road.

You should also encourage developers to create unit tests as early as possible for their applications and to run these tests before pushing the code to the central repository. Errors caught early in the software development process are the cheapest and easiest to fix.

When the code is pushed to a branch in a source code repository, a workflow engine monitoring that branch will send a command to a builder tool to build the code and run the unit tests in a controlled environment. The build process should be sized appropriately to handle all activities, including pushes and tests that might happen during the commit stage, for fast feedback. Other quality checks, such as unit test coverage, style check, and static analysis, can happen at this stage as well. Finally, the builder tool creates one or more binary builds and other artifacts, like images, stylesheets, and documents, for the application.

Continuous Delivery: Creating a Staging Environment

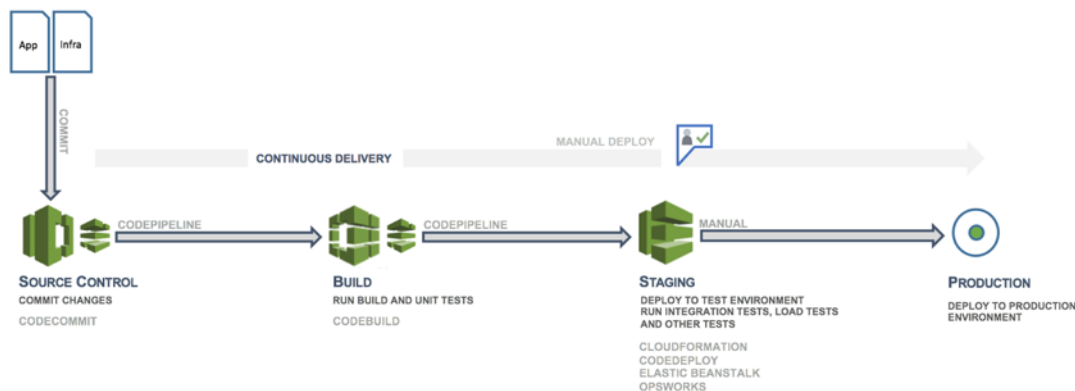


Figure 3: Continuous delivery—staging

Continuous delivery (CD) is the next phase and entails deploying the application code in a staging environment, which is a replica of the production stack, and running more functional tests. The staging environment could be a static environment premade for testing, or you could provision and configure a

dynamic environment with committed infrastructure and configuration code for testing and deploying the application code.

Continuous Delivery: Creating a Production Environment

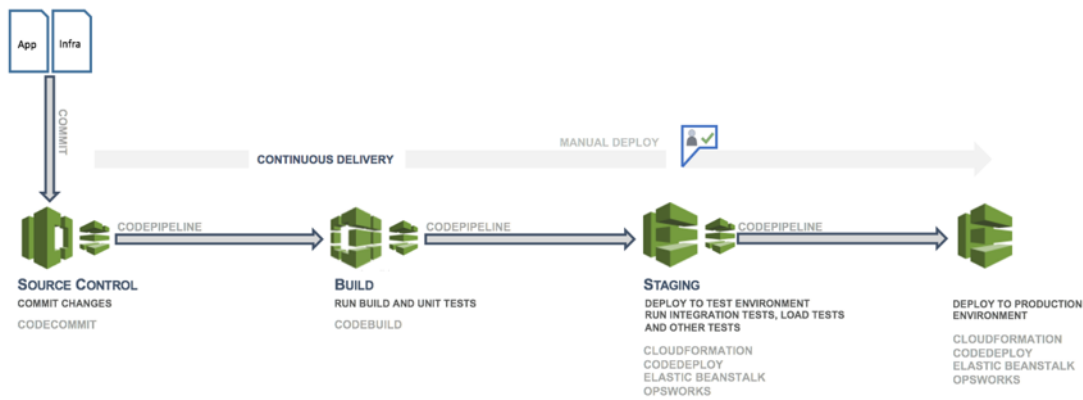


Figure 4: Continuous delivery—production

After the staging environment is built using infrastructure as code (IaC), a production environment can be built very quickly in the same way.

Continuous Deployment

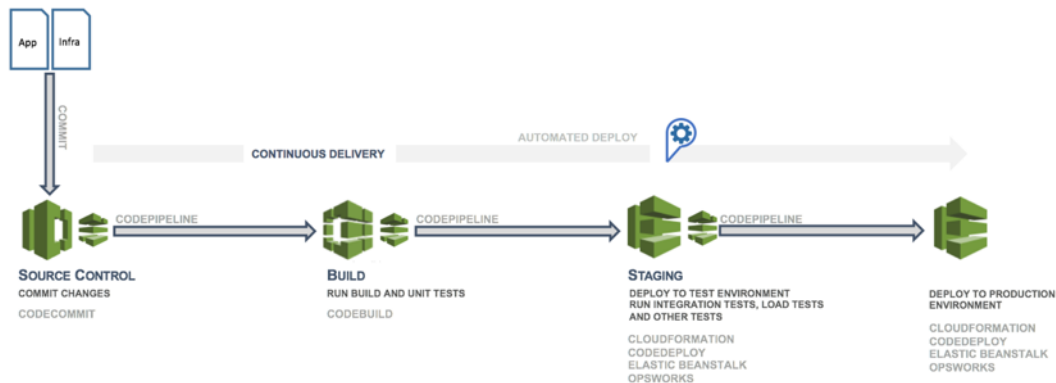


Figure 5: Continuous deployment

The final phase in the CI/CD pipeline is continuous deployment, which may include full automation of the entire software release process including deployment to the production environment. In a fully mature CI/CD

environment, the path to the production environment is fully automated, which allows code to be deployed with high confidence.

Maturity and Beyond

As your organization matures, it will continue to develop the CI/CD model to include more of the following improvements:

- More staging environments for specific performance, compliance, security, and user interface (UI) tests
- Unit tests of infrastructure and configuration code along with the application code
- Integration with other systems and processes such as code review, issue tracking, and event notification
- Integration with database schema migration (if applicable)
- Additional steps for auditing and business approval

Even the most mature organizations that have complex multi-environment CI/CD pipelines continue to look for improvements. DevOps is a journey, not a destination. Feedback about the pipeline is continuously collected, and improvements in speed, scale, security, and reliability are achieved as a collaboration between the different parts of the development teams.

Teams

AWS recommends organizing three developer teams for implementing a CI/CD environment: an application team, an infrastructure team, and a tools team (see Figure 6). This organization represents a set of best practices that have been developed and applied in fast-moving startups, large enterprise organizations, and in Amazon itself. The teams should be no larger than groups that two pizzas can feed, or about 10-12 people. This follows the communication rule that meaningful conversations hit limits as group sizes increase and lines of communication multiply.¹¹

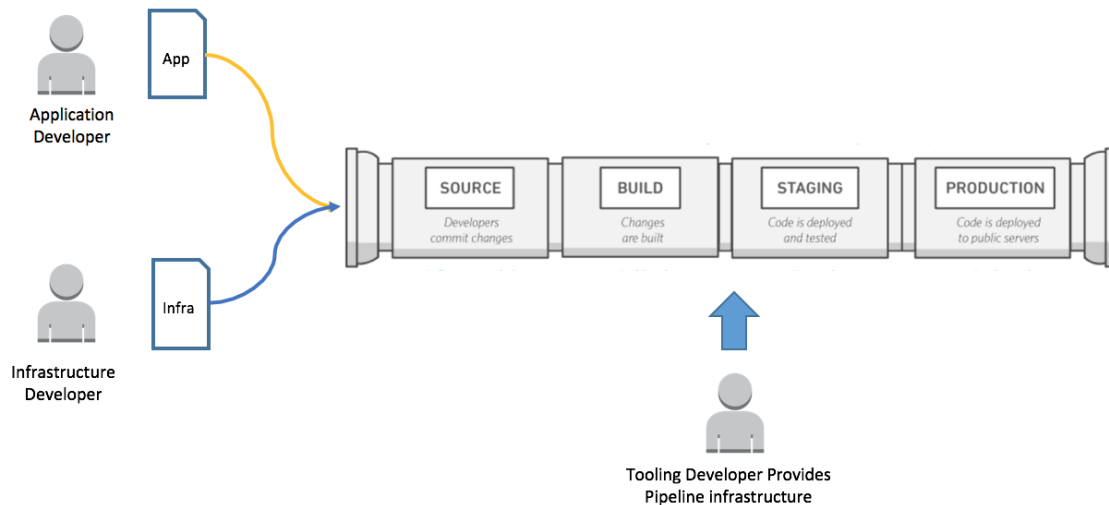


Figure 6: Application, infrastructure, and tools teams

Application Team

The application team creates the application. Application developers own the backlog, stories, and unit tests, and they develop features based on a specified application target. This team's organizational goal is to minimize the time these developers spend on non-core application tasks.

In addition to having functional programming skills in the application language, the application team should have platform skills and an understanding of system configuration. This will enable them to focus solely on developing features and hardening the application.

Infrastructure Team

The infrastructure team writes the code that both creates and configures the infrastructure needed to run the application. This team might use native AWS tools, such as AWS CloudFormation, or generic tools, such as Chef, Puppet, or Ansible. The infrastructure team is responsible for specifying what resources are needed, and it works closely with the application team. The infrastructure team might consist of only one or two people for a small application.

The team should have skills in infrastructure provisioning methods, such as AWS CloudFormation or HashiCorp Terraform. The team should also develop configuration automation skills with tools such as Chef, Ansible, Puppet, or Salt.

Tools Team

The tools team builds and manages the CI/CD pipeline. They are responsible for the infrastructure and tools that make up the pipeline. They are not part of the two-pizza team, however they create a tool that is used by the application and infrastructure teams in the organization. The organization needs to continuously mature its tools team, so that the tools team stays one step ahead of the maturing application and infrastructure teams.

The tools team must be skilled in building and integrating all parts of the CI/CD pipeline. This includes building source control repositories, workflow engines, build environments, testing frameworks, and artifact repositories. This team may choose to implement software such as AWS CodeStar, AWS CodePipeline, AWS CodeCommit, AWS CodeDeploy, and AWS CodeBuild, along with Jenkins, GitHub, Artifactory, TeamCenter, and other similar tools. Some organizations might call this a DevOps team, but we discourage this. Instead, we encourage thinking of DevOps as the sum of the people, processes, and tools in software delivery.

Testing Stages in Continuous Integration and Continuous Delivery

The three CI/CD teams should incorporate testing into the software development lifecycle at the different stages of the CI/CD pipeline. Overall, testing should start as early as possible. The testing pyramid, shown below, is a concept provided by Mike Cohn in *Succeeding with Agile*.¹² It shows the various software tests in relation to their cost and speed at which they run.

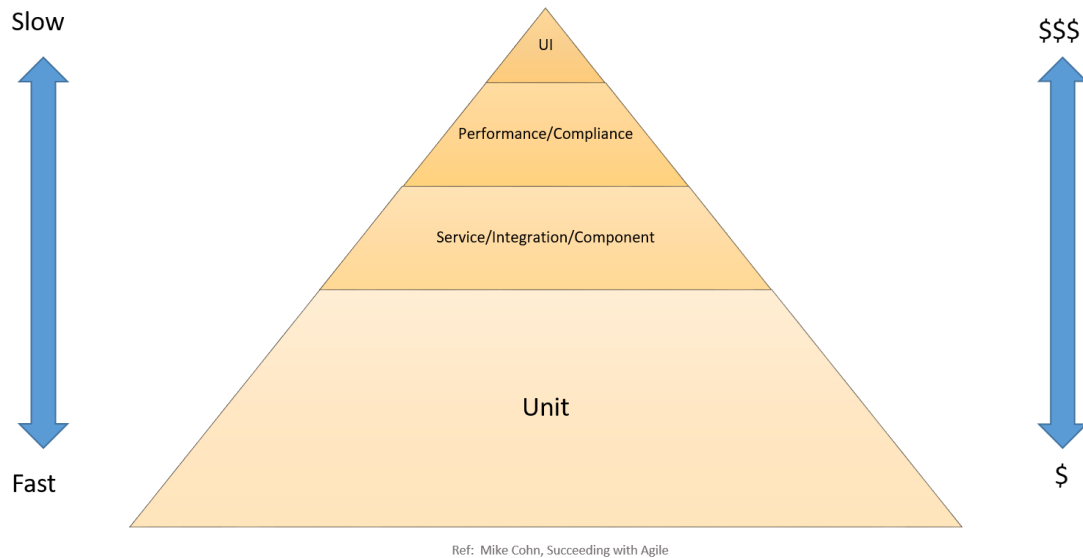


Figure 7: CI/CD testing pyramid

Unit tests are on the bottom of the pyramid. They are both the fastest to run and the least expensive. Therefore, unit tests should make up the bulk of your testing strategy. A good rule of thumb is about 70 percent. Unit tests should have near-complete code coverage because bugs caught in this phase can be fixed quickly and cheaply.

Service, component, and integration tests are above unit tests on the pyramid. These tests require detailed environments and, therefore, are more costly in infrastructure requirements and slower to run. Performance and compliance tests are the next level. They require production-quality environments and are more expensive yet. UI and user acceptance tests are at the top of the pyramid and require production-quality environments as well.

All of these tests are part of a complete strategy to assure high-quality software. However, for speed of development, emphasis is on the number of tests and the coverage in the bottom half of the pyramid.

In the following sections we discuss the CI/CD stages.

Setting up the Source

At the beginning of the project it's essential to set up a source where you can store your raw code and configuration and schema changes. In the source stage,

choose a source code repository such as one hosted in GitHub or AWS CodeCommit.

Setting Up and Executing Builds

Build automation is essential to the CI process. When setting up build automation, the first task is to choose the right build tool. There are many build tools, such as Ant, Maven, and Gradle for Java; Make for C/C++; Grunt for JavaScript; and Rake for Ruby. The build tool that will work best for you will depend on the programming language of your project and the skill set of your team. After you choose the build tool, all the dependencies need to be clearly defined in the build scripts, along with the build steps. It's also a best practice to version the final build artifacts, which makes it easier to deploy and to keep track of issues.

In the build stage, the build tools will take as input any change to the source code repository, build the software, and run the following types of tests:

Unit Testing – Tests a specific section of code to ensure the code does what it is expected to do. The unit testing is performed by software developers during the development phase. At this stage, a static code analysis, data flow analysis, code coverage, and other software verification processes can be applied.

Static Code Analysis – This test is performed without actually executing the application after the build and unit testing. This analysis can help to find coding errors and security holes, and it also can ensure conformance to coding guidelines.

Staging

In the staging phase, full environments are created that mirror the eventual production environment. The following tests are performed:

Integration Testing – Verifies the interfaces between components against software design. Integration testing is an iterative process and facilitates building robust interfaces and system integrity.

Component Testing – Tests message passing between various components and their outcomes. A key goal of this testing could be idempotency in component testing. Tests can include extremely large data volumes, or edge situations and abnormal inputs.

System Testing – Tests the system end-to-end and verifies if the software satisfies the business requirement. This might include testing the UI, API, backend logic, and end state.

Performance Testing – Determines the responsiveness and stability of a system as it performs under a particular workload. Performance testing also is used to investigate, measure, validate, or verify other quality attributes of the system, such as scalability, reliability, and resource usage. Types of performance tests might include load tests, stress tests, and spike tests. Performance tests are used for benchmarking against predefined criteria.

Compliance Testing – Checks whether the code change complies with the requirements of a nonfunctional specification and/or regulations. It determines if you are implementing and meeting the defined standards.

User Acceptance Testing – Validates the end-to-end business flow. This testing is executed by an end user in a staging environment and confirms whether the system meets the requirements of the requirement specification. Typically, customers employ alpha and beta testing methodologies at this stage.

Production

Finally, after passing the previous tests, the staging phase is repeated in a production environment. In this phase, a final **Canary test** can be completed by deploying the new code only on a small subset of servers or even one server, or one Region before deploying code to the entire production environment. Specifics on how to safely deploy to production are covered in the Deployment Methods section.

Next, we discuss building the pipeline to incorporate these stages and tests.

Building the Pipeline

This section discusses building the pipeline. Start by establishing a pipeline with just the components needed for CI and then transition later to a continuous delivery pipeline with more components and stages. This section also discusses how you can consider using AWS Lambda functions and manual approvals for large projects, plan for multiple teams, branches, and Regions.

Starting with a Minimum Viable Pipeline for Continuous Integration

Your organization's journey toward continuous delivery begins with a minimum viable pipeline (MVP). As discussed in [Continuous Delivery Maturity](#), teams can start with a very simple process, such as implementing a pipeline that performs a code style check or a single unit test without deployment.

A key component is a continuous delivery orchestration tool. To help you build this pipeline, Amazon developed [AWS CodeStar](#).

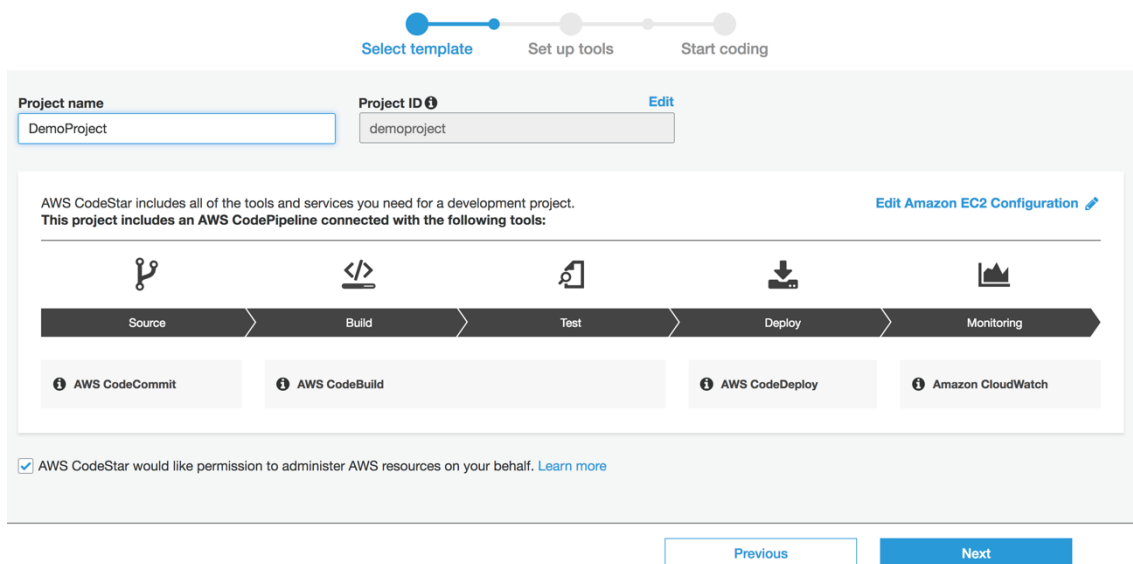


Figure 8: AWS CodeStar Setup Page

AWS CodeStar uses AWS CodePipeline, AWS CodeBuild, AWS CodeCommit, and AWS CodeDeploy with an integrated setup process, tools, templates, and dashboard. AWS CodeStar provides everything you need to quickly develop, build, and deploy applications on AWS. This allows you to start releasing code faster. Customers who are already familiar with the AWS Management Console and seek a higher level of control can manually configure their developer tools of choice and can provision individual AWS services as needed.

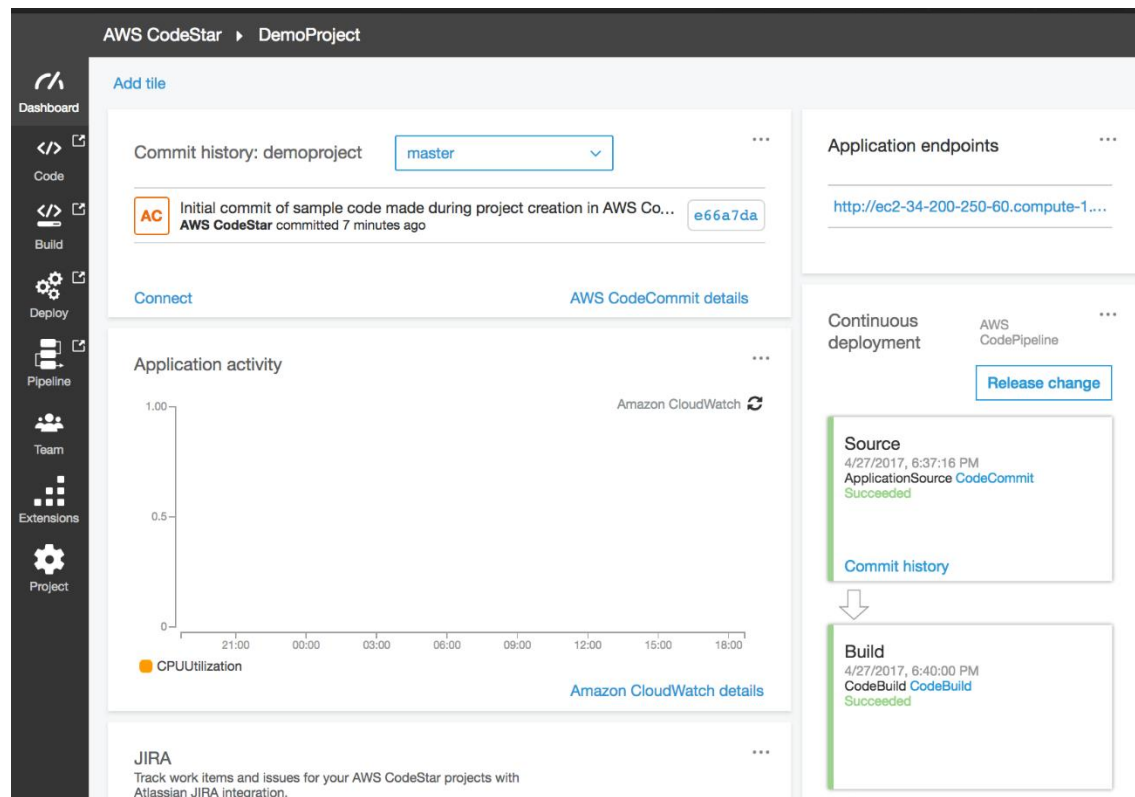


Figure 9: AWS CodeStar Dashboard

AWS CodePipeline is a CI/CD service that can be used through AWS CodeStar or through the AWS Management Console for fast and reliable application and infrastructure updates. AWS CodePipeline builds, tests, and deploys your code every time there is a code change, based on the release process models you define. This enables you to rapidly and reliably deliver features and updates. You can easily build out an end-to-end solution by using our pre-built plugins for popular third-party services like GitHub or by integrating your own custom plugins into any stage of your release process. With AWS CodePipeline, you only pay for what you use. There are no upfront fees or long-term commitments.

The steps of AWS CodeStar and AWS CodePipeline map directly to the [source, build, staging, and production CI/CD stages](#). While continuous delivery is desirable, you could start out with a simple two-step pipeline that checks the source repository and performs a build action:

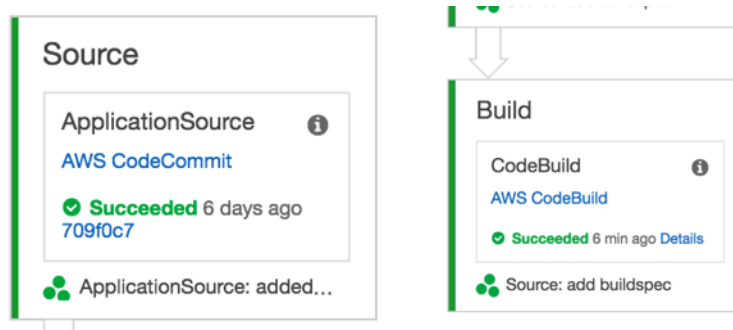


Figure 10: AWS CodePipeline—source and build stages

For AWS CodePipeline, the source stage can accept inputs from GitHub, AWS CodeCommit, and Amazon Simple Storage Service (S3). Automating the build process is a critical first step for implementing continuous delivery and moving toward continuous deployment. Eliminating human involvement in producing build artifacts removes the burden from your team, minimizes errors introduced by manual packaging, and allows you to start packaging consumable artifacts more often.

AWS CodePipeline works seamlessly with AWS CodeBuild, a fully managed build service, so you can easily set up a build step within your pipeline that packages your code and runs unit tests. With AWS CodeBuild, you don't need to provision, manage, or scale your own build servers. AWS CodeBuild scales continuously and processes multiple builds concurrently so your builds are not left waiting in a queue. AWS CodePipeline also integrates with build servers such as Jenkins, Solano CI, and TeamCity.

For example, in the build stage below, three actions (unit testing, code style checks, and code metrics collection) run in parallel. Using AWS CodeBuild, these steps can be added as new projects without any further effort in building or installing build servers to handle the load.

Build

CodeBuild

AWS CodeBuild

✓ Succeeded

12 min ago

Details

UnitTest

AWS CodeBuild

No executions yet

StyleChecker

AWS CodeBuild

No executions yet

CodeMetrics

AWS CodeBuild

No executions yet

Source: add buildspec

Figure 11: AWS CodePipeline — build functionality

The source and build stages shown in Figure 10, along with supporting processes and automation, support your team's transition toward a Continuous Integration. At this level of maturity, developers need to regularly pay attention to build and test results. They need to grow and maintain a healthy unit test base as well. This, in turn, will bolster the entire team's confidence in the CI/CD pipeline and further its adoption.

Continuous Delivery Pipeline

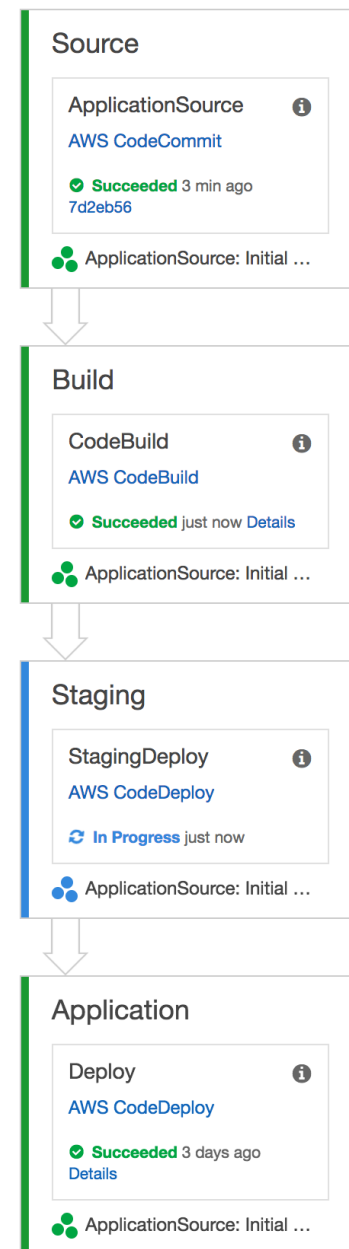
After the continuous integration pipeline has been implemented and supporting processes have been established, your teams can start transitioning toward the continuous delivery pipeline. This transition requires teams to automate both building and deploying applications.

A Continuous Delivery pipeline is characterized by the presence of Staging and Production steps, where the Production step is performed after a manual approval.

In the same manner as the Continuous Integration pipeline was built, your teams can gradually start building a Continuous Delivery pipeline by writing their deployment scripts.

Depending on the needs of an application, some of the deployment steps can be abstracted by existing AWS services. For example, AWS CodePipeline directly integrates with AWS CodeDeploy, a service that automates code deployments to Amazon EC2 instances and instances running on-premises, AWS OpsWorks, a configuration management service that helps you operate applications using Chef, and to AWS Elastic Beanstalk, a service for deploying and scaling web applications and services.

AWS has detailed [documentation](#) on how to implement and integrate AWS CodeDeploy with your infrastructure and pipeline.¹³



After your team successfully automates the deployment of the application, deployment stages can be expanded with various tests. For example:

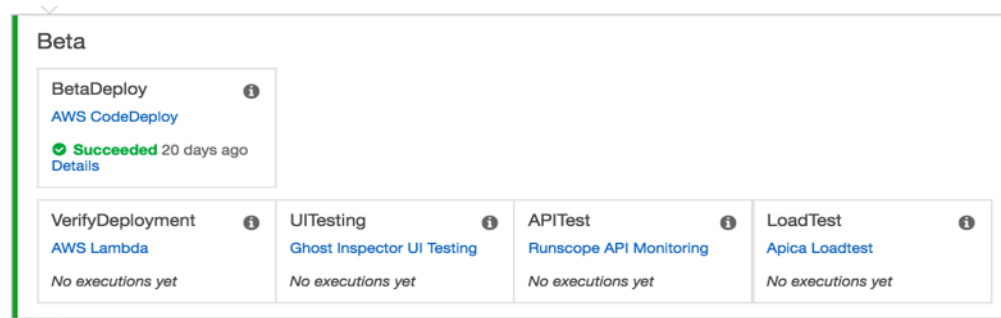


Figure 12: AWS CodePipeline—code tests in deployment stages

You can add other out-of-the-box integrations with services like Ghost Inspector, Runscope, Apica, and others.

Adding Lambda Actions

AWS CodeStar and AWS CodePipeline support [integration with AWS Lambda](#).¹⁴ This integration enables implementing a broad set of tasks, such as creating custom resources in your environment, integrating with third-party systems (such as Slack), and performing checks on your newly deployed environment.

Lambda functions can be used in CI/CD pipelines to do the following tasks:

- Roll out changes to your environment by applying or updating an AWS CloudFormation template.
- Create resources on demand in one stage of a pipeline using AWS CloudFormation and delete them in another stage.
- Deploy application versions with zero downtime in AWS Elastic Beanstalk with a Lambda function that swaps CNAME values.
- Deploy to Amazon EC2 Container Service (ECS) Docker instances.
- Back up resources before building or deploying by creating an AMI snapshot.
- Add integration with third-party products to your pipeline, such as posting messages to an IRC client.

Manual Approvals

Add an approval action to a stage in a pipeline at the point where you want the pipeline execution to stop so that someone with the required AWS Identity and Access Management (IAM) permissions can approve or reject the action.

If the action is approved, the pipeline execution resumes. If the action is rejected—or if no one approves or rejects the action within seven days of the pipeline reaching the action and stopping—the result is the same as an action failing, and the pipeline execution does not continue.




Figure 13: AWS CodeDeploy—manual approvals

Deploying Infrastructure Code Changes in a CI/CD Pipeline

AWS CodePipeline lets you select AWS CloudFormation as a deployment action in any stage of your pipeline. You can then choose the specific action you would like AWS CloudFormation to perform, such as creating or deleting stacks and creating or executing change sets. A stack is an AWS CloudFormation concept and represents a group of related AWS resources. While there are many ways of provisioning Infrastructure as Code, AWS CloudFormation is a comprehensive tool recommended by AWS as a scalable, complete solution that can describe the most comprehensive set of AWS resources as code. We recommend using AWS CloudFormation in an AWS CodePipeline project to [track infrastructure changes and tests](#).¹⁵

CI/CD for Serverless Applications

You can also use AWS CodeStar, AWS CodePipeline, AWS CodeBuild, and AWS CloudFormation to build CI/CD pipelines for serverless applications. These are applications based on AWS Lambda, which are composed of Lambda functions triggered by events. If you are a serverless application developer, you can use the combination of AWS CodePipeline, AWS CodeBuild, and AWS CloudFormation to automate the building, testing, and deployment of serverless applications that are expressed in templates built with the AWS Serverless Application Model. For more information, see the AWS Lambda documentation for [Automating Deployment of Lambda-based Applications](#).¹⁶

Pipelines for Multiple Teams, Branches, and Regions

For a large project, it's not uncommon for multiple project teams to work on different components. If multiple teams use a single code repository, it can be mapped so that each team has its own branch, plus there should be an integration or release branch for the final merge of the project. If a service-oriented or microservice architecture is used, each team could have its own code repository.

In the first scenario, if a single pipeline is used it's possible that one team could affect the other teams' progress by blocking the pipeline. We recommend that you create specific pipelines for the team branches and another release pipeline for the final product delivery.

Pipeline Integration with AWS CodeBuild

AWS CodeBuild is designed to enable your organization to build a highly available build process with almost unlimited scale. AWS CodeBuild provides quickstart environments for a number of popular languages plus the ability to run any Docker container that you specify.

With the advantages of tight integration with AWS CodeCommit, AWS CodePipeline, and AWS CodeDeploy, as well as Git and CodePipeline Lambda actions, the CodeBuild tool is highly flexible.

Software can be built through the inclusion of a `buildspec.yml` file that identifies each of the build steps, including pre- and post- build actions, or specified actions through the CodeBuild tool.

You can view detailed history of each build using the CodeBuild dashboard. Events are stored as Amazon CloudWatch Logs log files.

Build						
Build ARN	arn:aws:codebuild:us-east-1:77ecd198b59:build/codebuild-demo:99a13e63-950d-4f00-85a9-f77ecd198b59					
Build project	codebuild-demo					
Source provider	AWS CodePipeline					
Repository						
Start time	7 minutes ago					

Build history						
Expand to view the build's details.						
<div>Stop</div> <div>Refresh</div>						
<div>Viewing 1 to 10 builds</div> <div>Builds per page 10</div>						
	Project	Build run	Submitter	Status	Duration	Completed
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	CodePipeline-C...	Succeeded	1 min, 13 secs	5 minutes ago
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	CodePipeline-C...	Succeeded	1 min, 49 secs	12 minutes ago
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	root	Succeeded	1 min, 51 secs	Nov 18, 2016 11...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	root	Failed	1 min, 40 secs	Nov 18, 2016 4:...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	root	Failed	1 min, 50 secs	Nov 18, 2016 4:...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	root	Succeeded	1 min, 51 secs	Nov 18, 2016 3:...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	CodePipeline-C...	Succeeded	1 min, 8 secs	Nov 17, 2016 10...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	root	Succeeded	1 min, 52 secs	Nov 17, 2016 10...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	CodePipeline-C...	Failed	47 secs	Nov 17, 2016 10...
<input type="checkbox"/>	codebuild-demo	codebuild-demo...	CodePipeline-C...	Failed	1 min, 5 secs	Nov 17, 2016 10...

Figure 14: CloudWatch Logs log files in AWS CodeBuild

Pipeline Integration with Jenkins

You can use the Jenkins build tool to create delivery pipelines. These pipelines use standard jobs that define steps for implementing continuous delivery stages.¹⁷ However, this approach might not be optimal for larger projects because the current state of the pipeline doesn't persist between Jenkins restarts, implementing manual approval is not straightforward, and tracking the state of a complex pipeline can be complicated.

Instead, we recommend that you implement continuous delivery with Jenkins by using the [AWS Code Pipeline Plugin](#).¹⁸ The Pipeline plugin allows complex workflows to be described using Groovy-like domain-specific language and can

be used to orchestrate complex pipelines. The Pipeline plugin's functionality can be enhanced by the use of satellite plugins such as Pipeline Stage View Plugin, which visualizes the current progress of stages defined in a pipeline, or Pipeline Multibranch Plugin, which groups builds from different branches.

We recommend that you store your pipeline configuration in *Jenkinsfile* and have it checked into a source code repository. This allows for tracking changes to pipeline code and becomes even more important when working with the Pipeline Multibranch Plugin. We also recommend that you divide your pipeline into stages. This logically groups the pipeline steps and also enables the Pipeline Stage View Plugin to visualize the current state of the pipeline.

Figure 15 shows a sample Jenkins pipeline, with four defined stages visualized by the Pipeline Stage View Plugin.

master - Stage View

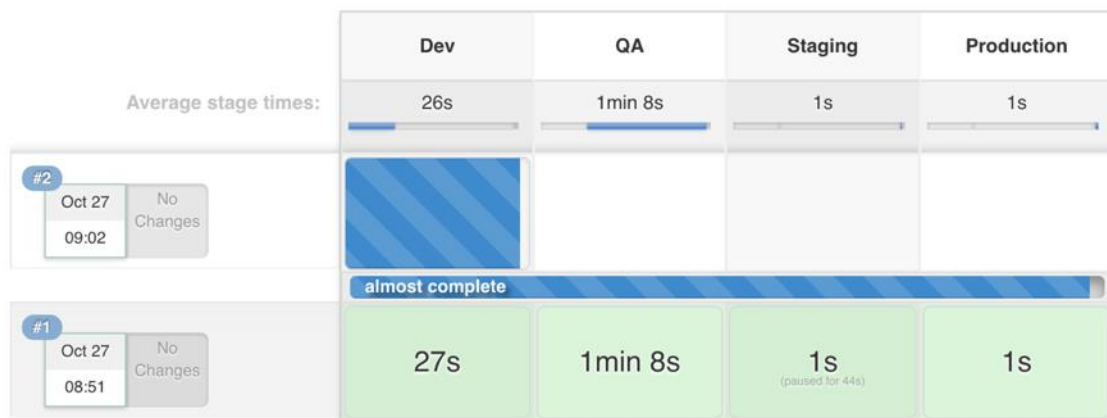
















Figure 15: Defined stages of Jenkins pipeline visualized by the Pipeline Stage View Plugin

Deployment Methods

You can consider multiple deployment strategies and variations for rolling out new versions of software in a Continuous Delivery process. This section discusses the most common deployment methods: all at once (deploy in place), rolling, immutable, and blue/green. We indicate which of these methods are supported by AWS CodeDeploy and AWS Elastic Beanstalk.

The following table summarizes the characteristics of each deployment method.

Method	Impact of Failed Deployment	Deploy Time	Zero Downtime	No DNS Change	Rollback Process	Code Deployed To
Deploy in place	Downtime		X	✓	Re-deploy	Existing instances
Rolling	Single batch out of service. Any successful batches prior to failure running new application version.	  †	✓	✓	Re-deploy	Existing instances
Rolling with additional batch (beanstalk)	Minimal if first batch fails, otherwise similar to rolling.	   †	✓	✓	Re-deploy	New & existing instances
Immutable	Minimal	   	✓	✓	Re-deploy	New instances
Blue/green	Minimal	   	✓	X	switch back to old environment	New instances

All at Once (In-Place Deployment)

All at once or in-place deployment is a method you can use to roll out new application code to an existing fleet of servers. This method replaces all the code in one deployment action. It requires downtime because all servers in the fleet are updated at once. There is no need to update existing DNS records. In case of a failed deployment, the only way to restore operations is to redeploy the code on all servers again.

In AWS Elastic Beanstalk this deployment is called [All at Once](#),¹⁹ and is available for single and load-balanced applications. In AWS CodeDeploy this deployment method is called [In-Place Deployment](#) with a deployment configuration of AllAtOnce.²⁰

Rolling Deployment

With rolling deployment, the fleet is divided into portions so that all of the fleet isn't upgraded at once. During the deployment process two software versions, new and old, are running on the same fleet. This method allows a zero-downtime update. If the deployment fails, only the updated portion of the fleet will be affected.

A variation of the rolling deployment method, called canary release, involves deployment of the new software version on a very small percentage of servers at first. This way, you can observe how the software behaves in production on a few servers, while minimizing the impact of breaking changes. If there is an elevated rate of errors from a canary deployment, the software is rolled back. Otherwise, the percentage of servers with the new version is gradually increased.

AWS Elastic Beanstalk has followed the rolling deployment pattern with two deployment options, [rolling and rolling with additional batch](#).²¹ These options allow the application to first scale up before taking servers out of service, preserving full capability during the deployment. AWS CodeDeploy accomplishes this pattern as a variation of an in-place deployment with patterns like [OneAtATime](#) and [HalfAtATime](#).²²

Immutable and Blue/Green Deployment

The immutable pattern specifies a deployment of application code by starting an entirely new set of servers with a new configuration or version of application code. This pattern leverages the cloud capability that new server resources are created with simple API calls.

The blue/green deployment strategy is a type of immutable deployment which also requires creation of another environment. Once the new environment is up and passed all tests, traffic is shifted to this new deployment. Crucially the old environment, that is, the “blue” environment, is kept idle in case a rollback is needed.

AWS Elastic Beanstalk supports [immutable](#)²³ and [blue/green](#)²⁴ deployment patterns. AWS CodeDeploy also supports the [blue/green pattern](#).²⁵ For more

information on how AWS services accomplish these immutable patterns, see the [Blue/Green Deployments on AWS whitepaper](#).²⁶

Database Schema Changes

It's common for modern software to have a database layer. Typically, a relational database is used, which stores both data and the structure of the data. It's often necessary to modify the database in the continuous delivery process. Handling changes in a relational database requires special consideration, and it offers other challenges than the ones present when deploying application binaries. Usually, when you upgrade an application binary you stop the application, upgrade it, and then start it again. You don't really bother about the application state. That is handled outside of the application.

When upgrading databases, you do need to consider state because a database contains much state but comparatively little logic and structure.

The database schema before and after a change is applied should be seen as different versions of the database. You could use tools such as Liquibase and Flyway to manage the versions.

In general, those tools employ some variant of the following methods:

- Add a table to the database where a database version is stored.
- Keep track of database change commands and bunch them together in versioned change sets. In the case of Liquibase, these changes are stored in XML files. Flyway employs a slightly different method where the change sets are handled as separate SQL files or occasionally as separate Java classes for more complex transitions.
- When Liquibase is being asked to upgrade a database, it looks at the metadata table and determines which change sets to run in order to make the database up-to-date with the latest version.

Summary of Best Practices

The following are some best practice dos and don'ts for CI/CD.

Do:

- Treat your infrastructure as code
 - Use version control for your infrastructure code.
 - Make use of bug tracking/ticketing systems.
 - Have peers review changes before applying them.
 - Establish infrastructure code patterns/designs.
 - Test infrastructure changes like code changes.
- Put developers into integrated teams of no more than 12 self-sustaining members.
- Have all developers commit code to the main trunk frequently, with no long-running feature branches.
- Consistently adopt a build system such as Maven or Gradle across your organization and standardize builds.
- Have developers build unit tests toward 100% coverage of the code base.
- Ensure that unit tests are 70% of the overall testing in duration, number, and scope.
- Ensure that unit tests are up-to-date and not neglected. Unit test failures should be fixed, not bypassed.
- Treat your continuous delivery configuration as code.
- Establish role-based security controls (that is, who can do what and when).
 - Monitor/track every resource possible.
 - Alert on services, availability, and response times.
 - Capture, learn, and improve.
 - Share access with everyone on the team.
 - Plan metrics and monitoring into the lifecycle.
- Keep and track standard metrics.
 - Number of builds.
 - Number of deployments.

- Average time for changes to reach production.
- Average time from first pipeline stage to each stage.
- Number of changes reaching production.
- Average build time.
- Use multiple distinct pipelines for each branch and team.

Don't:

- Have long-running branches with large complicated merges.
- Have manual tests.
- Have manual approval processes, gates, code reviews, and security reviews.

Conclusion

Continuous Integration and Continuous Delivery provides an ideal scenario for your organization's application teams. Your developers simply push code to a repository. This code will be integrated, tested, deployed, tested again, merged with infrastructure, go through security and quality reviews, and be ready to deploy with extremely high confidence.

When CI/CD is used code quality is improved, and software updates are delivered quickly and with high confidence that there will be no breaking changes. The impact of any release can be correlated with data from production and operations. It can be used for planning the next cycle, too—a vital DevOps practice in your organization's cloud transformation.

Further Reading

For more information on the topics discussed in this whitepaper, read the following AWS whitepapers:

- [Overview of Deployment Options on AWS](#)
- [Blue/Green Deployments on AWS](#)
- [Jenkins on AWS](#)

- [Microservices on AWS](#)
- [Docker on AWS: Running Containers in the Cloud](#)

Contributors

The following individuals and organizations contributed to this document:

- David Stacy, Senior Consultant - DevOps, AWS Professional Services
- Mikhail Prudnikov, Senior Solutions Architect, AWS
- Asif Khan, Solutions Architect, AWS
- Xiang Shen, Senior Solutions Architect, AWS

Notes

- ¹ <https://aws.amazon.com/devops/what-is-devops/>
- ² <http://www.businessinsider.com/jeff-bezos-two-pizza-rule-for-productive-meetings-2013-10>
- ³ <https://aws.amazon.com/codestar/>
- ⁴ <https://aws.amazon.com/codecommit/>
- ⁵ <https://aws.amazon.com/codepipeline/>
- ⁶ <https://aws.amazon.com/codebuild/>
- ⁷ <https://aws.amazon.com/codedeploy/>
- ⁸ Kief Morris, Infrastructure as Code
- ⁹ <https://aws.amazon.com/devops/partner-solutions/>
- ¹⁰ <https://do.awsstatic.com/whitepapers/AWS-Cloud-Transformation-Maturity-Model.pdf>
- ¹¹ [The Two Pizza Rule is Jeff Bezos Secret to Productive Meetings](#)
- ¹² Mike Cohn, Succeeding with Agile
- ¹³ <http://docs.aws.amazon.com/codepipeline/latest/userguide/getting-started-w.html#getting-started-w-create-deployment>
- ¹⁴ <https://docs.aws.amazon.com/codepipeline/latest/userguide/how-to-lambda-integration.html>
- ¹⁵ <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/continuous-delivery-codepipeline.html>
- ¹⁶ <http://docs.aws.amazon.com/lambda/latest/dg/automating-deployment.html>
- ¹⁷ <https://jenkins.io/doc/book/pipeline/overview/>
- ¹⁸ <https://wiki.jenkins-ci.org/display/JENKINS/AWS+CodePipeline+Plugin>
- ¹⁹ <http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

20

<http://docs.aws.amazon.com/codedeploy/latest/userguide/deployments.html>

21 [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html)

[features.rolling-version-deploy.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html)

22 [http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-](http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html)
[configurations.html](http://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html)

23 [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html)
[updates-immutable.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html)

24 [http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMEswap.html)
[features.CNAMEswap.html](http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMEswap.html)

25

[http://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html#welcome-](http://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html#welcome-deployment-overview-blue-green)
[deployment-overview-blue-green](http://docs.aws.amazon.com/codedeploy/latest/userguide/welcome.html#welcome-deployment-overview-blue-green)

26 https://do.awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf