

# Confluence Data Center on the AWS Cloud

## Quick Start Reference Deployment

*March 2017*

*Atlassian*

*AWS Quick Start Reference Team*

### Contents

Overview .....	2
Costs and Licenses .....	2
Architecture .....	2
Prerequisites .....	4
Specialized Knowledge .....	4
Technical Requirements .....	4
Deployment Options .....	4
Deployment Steps .....	4
Step 1. Prepare Your AWS Account .....	4
Step 2. Launch the Quick Start.....	5
Step 3. Configure Confluence Data Center .....	11
Step 4. Add instances to the Confluence Data Center cluster .....	15
Step 5. Enable collaborative editing with Synchrony .....	15
FAQ.....	18
Additional Resources .....	19
Send Us Feedback .....	20
Document Revisions .....	20

This Quick Start deployment guide was created by Amazon Web Services (AWS) in partnership with Atlassian.

## Overview

This Quick Start reference deployment guide provides step-by-step instructions for deploying Confluence Data Center on the Amazon Web Services (AWS) Cloud. [Quick Starts](#) are automated reference deployments that use AWS CloudFormation templates to launch, configure, and run the AWS compute, network, storage, and other services required to deploy a specific workload on AWS. Confluence Data Center is a team collaboration software from Atlassian. Confluence Data Center is a self managed solution that gives you high availability, performance at scale, and disaster recovery for uninterrupted access to Confluence for all your teams.

This Quick Start is for users who want to deploy Confluence Data Center in a supported configuration in the AWS Cloud, following AWS best practices.

## Costs and Licenses

You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using the Quick Start.

The AWS CloudFormation template for this Quick Start includes configuration parameters that you can customize. Some of these settings, such as instance type, will affect the cost of deployment. For cost estimates, see the pricing pages for each AWS service you will be using or the [AWS Simple Monthly Calculator](#). Prices are subject to change.

A Confluence Data Center license is required to deploy this Quick Start guide. Information on obtaining a Confluence Data Center license can be found in the [Atlassian pricing page](#). For evaluation, you can also request an evaluation license.

## Architecture

Deploying this Quick Start for a [new virtual private cloud \(VPC\)](#) with **default parameters** builds the following Confluence Data Center environment in the AWS Cloud.

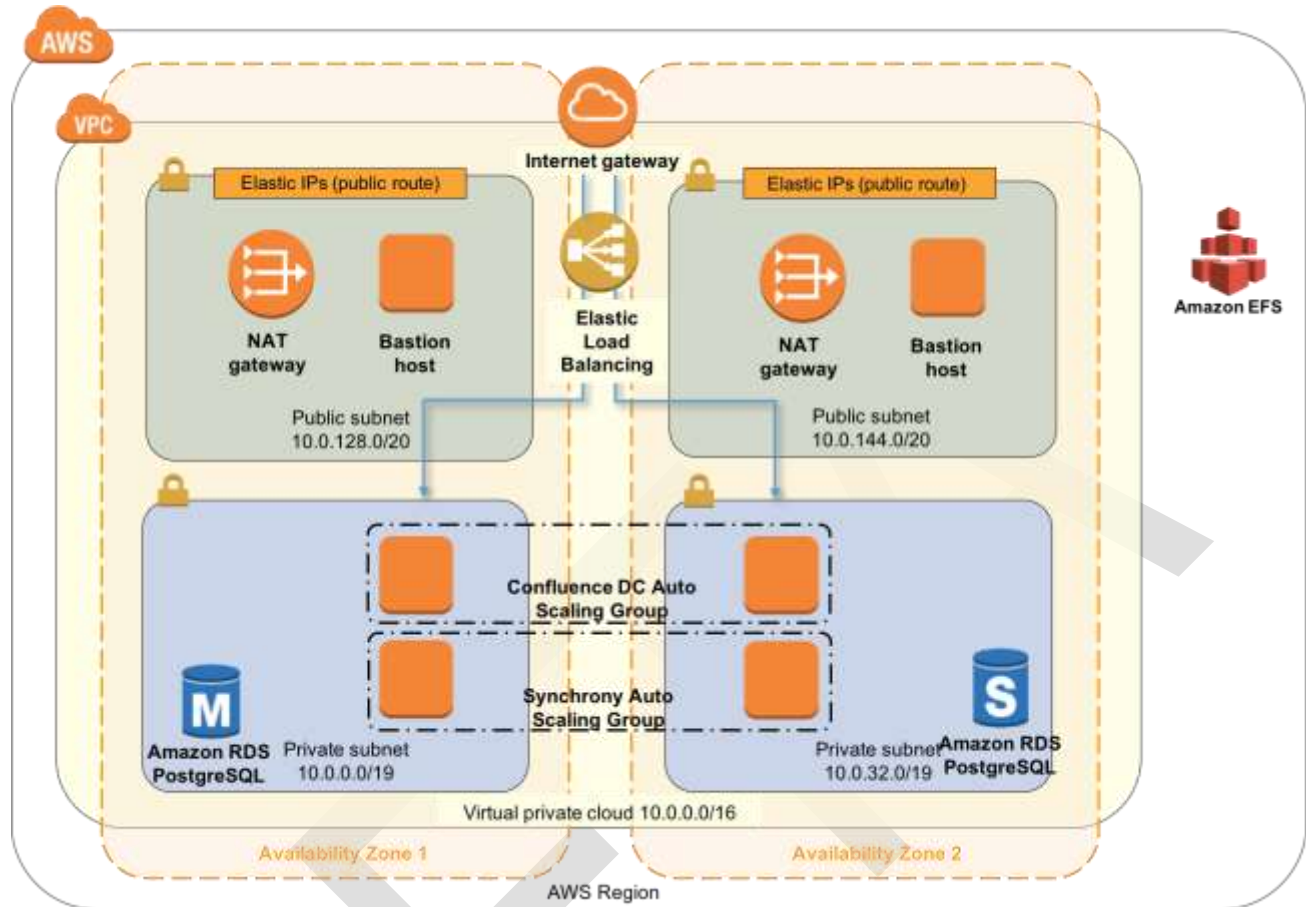


Figure 1: Quick Start Confluence Data Center architecture on AWS

The Quick Start sets up the following:

- **Amazon RDS PostgreSQL:** Confluence Data Center requires a supported external database. Amazon Relational Database Service (Amazon RDS) PostgreSQL in a Multi-AZ configuration allows failover in the event the master node fails.
- **Amazon Elastic File System:** Confluence Data Center uses a shared file system to store artifacts like attachments, avatars, icons, import/export files and plugins in a common location that is accessible to all Confluence nodes. The Quick Start architecture implements the shared file system using the highly available Amazon Elastic File System (Amazon EFS) service.
- **Confluence Auto Scaling group:** The Confluence Data Center product is installed on EC2 instances in an Auto Scaling group. The instances are based on Amazon Linux and use an Atlassian-provided AMI. The scaling metric is CPU utilization.
- **Synchrony Auto Scaling group:** Synchrony is required for the collaborative editing feature. Synchrony is installed on EC2 instances in an Auto Scaling group. The instances are based on Amazon Linux and use an Atlassian-provided AMI. The scaling metric is CPU utilization.

## Prerequisites

### Specialized Knowledge

Before you deploy this Quick Start, we recommend that you become familiar with the following AWS services. (If you are new to AWS, see [Getting Started with AWS](#).)

- [Amazon VPC](#)
- [Amazon EC2](#)
- [Amazon EBS](#)
- [Amazon EFS](#)
- [PostgreSQL on Amazon RDS](#)
- 

### Technical Requirements

This Quick Start requires an Atlassian account and a license to use Confluence Data Center.

## Deployment Options

This Quick Start provides two deployment options:

- **Deployment of Confluence Data Center into a new VPC** (end-to-end deployment) builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, bastion hosts, and other infrastructure components, and then deploys Confluence Data Center into this new VPC.
- **Deployment of Confluence Data Center into an existing VPC** provisions Confluence Data Center in your existing AWS infrastructure.

The Quick Start also lets you configure additional settings such as CIDR blocks, instance types, and Confluence Data Center settings, as discussed later in this guide.

## Deployment Steps

### Step 1. Prepare Your AWS Account

1. If you don't already have an AWS account, create one at <https://aws.amazon.com> by following the on-screen instructions.
2. Use the region selector in the navigation bar to choose the AWS Region where you want to deploy *<software>* on AWS.
3. Create a [key pair](#) in your preferred region.
4. If necessary, [request a service limit increase](#) for the Amazon EC2 *<type>* instance type. You might need to do this if you already have an existing deployment that uses this instance type, and you think you might exceed the [default limit](#) with this reference deployment.

## Step 2. Launch the Quick Start

**Note** You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using this Quick Start. For full details, see the pricing pages for each AWS service you will be using in this Quick Start or the [AWS Simple Monthly Calculator](#). Prices are subject to change.

Launch

Launch

1. Choose one of the following options to launch the AWS CloudFormation template into your AWS account. For help choosing an option, see [deployment options](#) earlier in this guide.

<p><b>Option 1</b></p> <p>Deploy Confluence Data Center into a new VPC on AWS</p>	<p><b>Option 2</b></p> <p>Deploy Confluence Data Center into an existing VPC</p>
-----------------------------------------------------------------------------------	----------------------------------------------------------------------------------

**Important** If you're deploying Confluence Data Center into an existing VPC, make sure that your VPC has two private subnets in different Availability Zones for the database instances. These subnets require NAT gateways or NAT instances in their route tables, to allow the instances to download packages and software without exposing them to the Internet. You'll also need the domain name option configured in the DHCP options as explained in the Amazon VPC documentation. You'll be prompted for your VPC settings when you launch the Quick Start.

Each deployment takes about 20 minutes to complete.

2. Check the region that's displayed in the upper-right corner of the navigation bar, and change it if necessary. This is where the network infrastructure for **<software>** will be built. The template is launched in the US West (Oregon) Region by default.
3. On the **Select Template** page, keep the default setting for the template URL, and then choose **Next**.
4. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters, choose **Next**.

In the following tables, parameters are listed by category and described separately for the two deployment options:

- [Parameters for deploying <software> into a new VPC](#)

- [Parameters for deploying <software> into an existing VPC](#)
- **Option 1: Parameters for deploying Confluence Data Center into a new VPC**

[View template](#)

*VPC Network Configuration:*

Parameter label (name)	Default	Description
<b>Availability Zones</b> (AvailabilityZones)	<i>Requires input</i>	The list of Availability Zones to use for the subnets in the VPC. The Quick Start uses two Availability Zones from your list and preserves the logical order you specify.
<b>VPC CIDR</b> (VPCCIDR)	10.0.0.0/16	CIDR block for the VPC.
<b>Private Subnet 1 CIDR</b> (PrivateSubnet1CIDR)	10.0.0.0/19	CIDR block for the private subnet located in Availability Zone 1.
<b>Private Subnet 2 CIDR</b> (PrivateSubnet2CIDR)	10.0.32.0/19	CIDR block for the private subnet located in Availability Zone 2.
<b>Public Subnet 1 CIDR</b> (PublicSubnet1CIDR)	10.0.128.0/20	CIDR block for the public (DMZ) subnet located in Availability Zone 1.
<b>Public Subnet 2 CIDR</b> (PublicSubnet2CIDR)	10.0.144.0/20	CIDR block for the public (DMZ) subnet located in Availability Zone 2.
<b>Permitted IP range</b> (AccessCIDR)	<i>Requires input</i>	The CIDR IP range that is permitted to access <software>. We recommend that you set this value to a trusted IP range. For example, you might want to grant only your corporate network access to the software.
<b>Assign public IP</b> (AssociatePublicIpAddress)	true	Controls whether the EC2 instances are assigned a public IP address.
<b>SSL certificate name</b> (SSLCertificateName)	<i>Optional</i>	The name of your server certificate to use for HTTPS. Leave this parameter blank if you don't want to set up HTTPS at this time.

### Confluence Setup:

Parameter label (name)	Default	Description
<b>Version of Confluence</b> (ConfluenceVersion)	<i>Requires input</i>	The version of Confluence to install. This Quick Start supports Confluence Data Center version 6.1.0 and later. For a list of versions, see the <a href="#">Atlassian documentation</a> .
<b>Confluence cluster node instance type</b> (ClusterNodeInstanceType)	c3.xlarge	EC2 instance type for the Confluence Data Center nodes.
<b>Minimum number of cluster nodes</b> (ClusterNodeMin)	1	Minimum number of Confluence Data Center nodes in the Auto Scaling group.

Parameter label (name)	Default	Description
<b>Maximum number of cluster nodes</b> (ClusterNodeMax)	1	Maximum number of Confluence Data Center nodes in the Auto Scaling group. When you first launch the Quick Start, leave the default value of 1 node unchanged. In <a href="#">step 3</a> , after configuring Confluence Data Center for multinode clustering, you can change this parameter setting to the desired number of cluster nodes.
<b>Minimum number of Synchrony nodes</b> (ClusterNodeMin)	0	Minimum number of Synchrony nodes in the Auto Scaling group.
<b>Maximum number of Synchrony nodes</b> (ClusterNodeMax)	1	Maximum number of Synchrony nodes in the Auto Scaling group. When you first launch the Quick Start, leave the default value of 0 node unchanged. In <a href="#">step 3</a> , after configuring Confluence Data Center for multinode clustering, you can change this parameter setting to the desired number of Synchrony nodes, and then enable collaborative editing.

**Database:**

Parameter label (name)	Default	Description
<b>RDS instance class</b> (DBInstanceClass)	db.m4.large	EC2 instance type for the Amazon RDS database.
<b>Master password</b> (DBMasterUserPassword)	<i>Requires input</i>	Password for the master ("postgres") account. This password should be 8-128 alphanumeric characters.
<b>Confluence database password</b> (DBPassword)	<i>Requires input</i>	Password for the Confluence database user account, with a maximum length of 128 alphanumeric characters.
<b>Database storage</b> (DBStorage)	10	The storage size, in GiB, to allocate to the database. This value should be 100-6144, if you've selected Provisioned IOPS for the database storage type.
<b>Database storage type</b> (DBStorageType)	General Purpose (SSD)	Database storage type. For more information about storage types, see the <a href="#">AWS documentation</a> .
<b>RDS Provisioned IOPS</b> (DBIops)	—	IOPS for database storage. This value is used only when the <b>Database storage type</b> parameter is set to Provisioned IOPS. Allowed range is 1,000-30,000. The ratio of IOPS to allocated storage must be between 3 and 10.
<b>Enable RDS Multi-AZ deployment</b> (DBMultiAZ)	true	If <b>true</b> , the Quick Start deploys the Confluence instances in two Availability Zones for high availability. If high availability isn't a concern, you can set this parameter to <b>false</b> .

**Advanced :**

Parameter label (name)	Default	Description
<b>Catalina options</b>	<i>Optional</i>	Java options that are passed to the JVM that runs Confluence.



Parameter label (name)	Default	Description
(CatalinaOpts)		
<b>Start the collectd service</b> (StartCollectd)	false	Set this parameter to <b>true</b> to enable monitoring with the collectd.conf configuration file.
<b>Database snapshot ID to restore</b> (DBSnapshotId)	<i>Optional</i>	Amazon RDS snapshot ID of an existing database backup to restore. This parameter must be used with the <b>Home volume snapshot ID to restore</b> parameter. Leave this value blank for a new instance.

### Amazon EC2 Configuration:

Parameter label (name)	Default	Description
<b>Key Name</b> (KeyPairName)	<i>Requires input</i>	Public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
<b>NAT Instance Type</b> (NATInstanceType)	t2.small	EC2 instance type for NAT instances. This parameter is used only if your selected AWS Region doesn't support NAT gateways.

### AWS Quick Start Configuration:

Parameter label (name)	Default	Description
<b>Quick Start S3 Bucket Name</b> (QSS3BucketName)	quickstart-reference	S3 bucket where the Quick Start templates and scripts are installed. Use this parameter to specify the S3 bucket name you've created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers, lowercase letters, uppercase letters, and hyphens, but should not start or end with a hyphen.
<b>Quick Start S3 Key Prefix</b> (QSS3KeyPrefix)	atlassian/Confluence/latest	The <a href="#">S3 key name prefix</a> used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes, but should not start or end with a forward slash (which is automatically added).

- **Option 2: Parameters for deploying Confluence Data Center into an existing VPC**

[View template](#)

*Confluence setup(Cluster Nodes)*

Parameter label (name)	Default	Description
------------------------	---------	-------------



Parameter label (name)	Default	Description
<b>Confluence cluster node instance type</b> (ClusterNodeInstanceType)	c3.xlarge	EC2 instance type for the Confluence Data Center nodes.
<b>Minimum number of cluster nodes</b> (ClusterNodeMin)	1	Minimum number of Confluence Data Center nodes in the Auto Scaling group.
<b>Maximum number of cluster nodes</b> (ClusterNodeMax)	1	Maximum number of Confluence Data Center nodes in the Auto Scaling group. When you first launch the Quick Start, leave the default value of 1 node unchanged. In <a href="#">step 3</a> , after configuring Confluence Data Center for multinode clustering, you can change this parameter setting to the desired number of cluster nodes.
<b>Minimum number of Synchrony nodes</b> (ClusterNodeMin)	0	Minimum number of Synchrony nodes in the Auto Scaling group.
<b>Maximum number of Synchrony nodes</b> (ClusterNodeMax)	1	Maximum number of Synchrony nodes in the Auto Scaling group. When you first launch the Quick Start, leave the default value of 0 node unchanged. In <a href="#">step 3</a> , after configuring Confluence Data Center for multinode clustering, you can change this parameter setting to the desired number of Synchrony nodes, and then enable collaborative editing.

*Database:*

Parameter label (name)	Default	Description
<b>Database instance class</b> (DBInstanceClass)	db.m4.large	EC2 instance type for the Amazon RDS database.
<b>Master password</b> (DBMasterUserPassword)	<i>Requires input</i>	Password for the master ("postgres") account. This password should be 8-128 alphanumeric characters.
<b>Confluence database password</b> (DBPassword)	<i>Requires input</i>	Password for the Confluence database user account, with a maximum length of 128 alphanumeric characters.
<b>Database storage</b> (DBStorage)	10	The storage size, in GiB, to allocate to the database. This value should be 100-6144, if you've selected Provisioned IOPS for the database storage type.
<b>Database storage type</b> (DBStorageType)	General Purpose (SSD)	Database storage type. For more information about storage types, see the <a href="#">AWS documentation</a> .
<b>RDS Provisioned IOPS</b> (DBIops)	—	IOPS for database storage. This value is used only when the <b>Database storage type</b> parameter is set to Provisioned IOPS. Allowed range is 1,000-30,000. The ratio of IOPS to allocated storage must be between 3 and 10.
<b>DBMultiAZ</b>	true	If <b>true</b> , the Quick Start deploys the Confluence instances in two Availability Zones for high availability. If high availability

Parameter label (name)	Default	Description
(DBMultiAZ)		isn't a concern, you can set this parameter to <b>false</b> .

### Networking:

Parameter label (name)	Default	Description
<b>VPC</b> (VPCID)	<i>Requires input</i>	ID of your existing VPC (e.g., vpc-0343606e).
<b>External subnets</b>	<i>Requires input</i>	Subnets (two or more) where your user-facing load balancer will be deployed. <b>MUST</b> be within the selected VPC.
<b>Internal subnets</b>	<i>Requires input</i>	Subnets (two or more) where your cluster nodes and other internal infrastructure will be deployed. <b>MUST</b> be within the selected VPC. Specify the ExternalSubnets again here if you wish to deploy the whole stack into the same subnets.
<b>Assign public IP</b> (AssociatePublicIpAddress)	true	Controls whether the EC2 instances are assigned a public IP address.
<b>Permitted IP range</b> (AccessCIDR)	<i>Requires input</i>	The CIDR IP range that is permitted to access <b>&lt;software&gt;</b> . We recommend that you set this value to a trusted IP range. For example, you might want to grant only your corporate network access to the software.
<b>Key Pair Name</b> (KeyPairName)	<i>Requires input</i>	Public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
<b>SSL certificate name</b> (SSLCertificateName)	<i>Optional</i>	The name of your server certificate to use for HTTPS. Leave this parameter blank if you don't want to set up HTTPS at this time.

### AWS Quick Start Configuration:

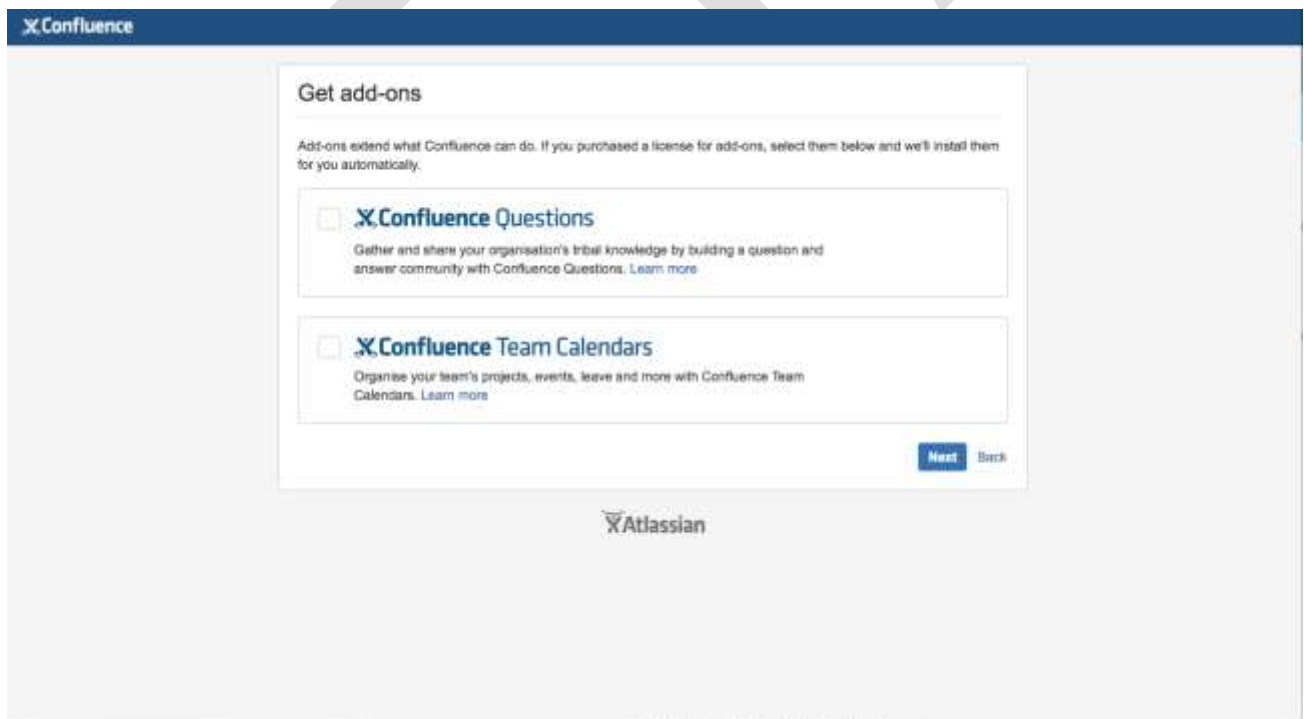
Parameter label (name)	Default	Description
<b>Quick Start S3 Bucket Name</b> (QSS3BucketName)	quickstart-reference	S3 bucket where the Quick Start templates and scripts are installed. Use this parameter to specify the S3 bucket name you've created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers, lowercase letters, uppercase letters, and hyphens, but should not start or end with a hyphen.
<b>Quick Start S3 Key Prefix</b> (QSS3KeyPrefix)	atlassian/Confluence/latest	The <a href="#">S3 key name prefix</a> used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes, but should not start or end with a forward slash (which is automatically added).

5. On the **Options** page, you can [specify tags](#) (key-value pairs) for resources in your stack and [set advanced options](#). When you're done, choose **Next**.
6. On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the check box to acknowledge that the template will create IAM resources.
7. Choose **Create** to deploy the stack.
8. Monitor the status of the stack. When the status is **CREATE\_COMPLETE**, the Confluence Data Center cluster is ready.
9. Use the URLs displayed in the **Outputs** tab for the stack to view the resources that were created.

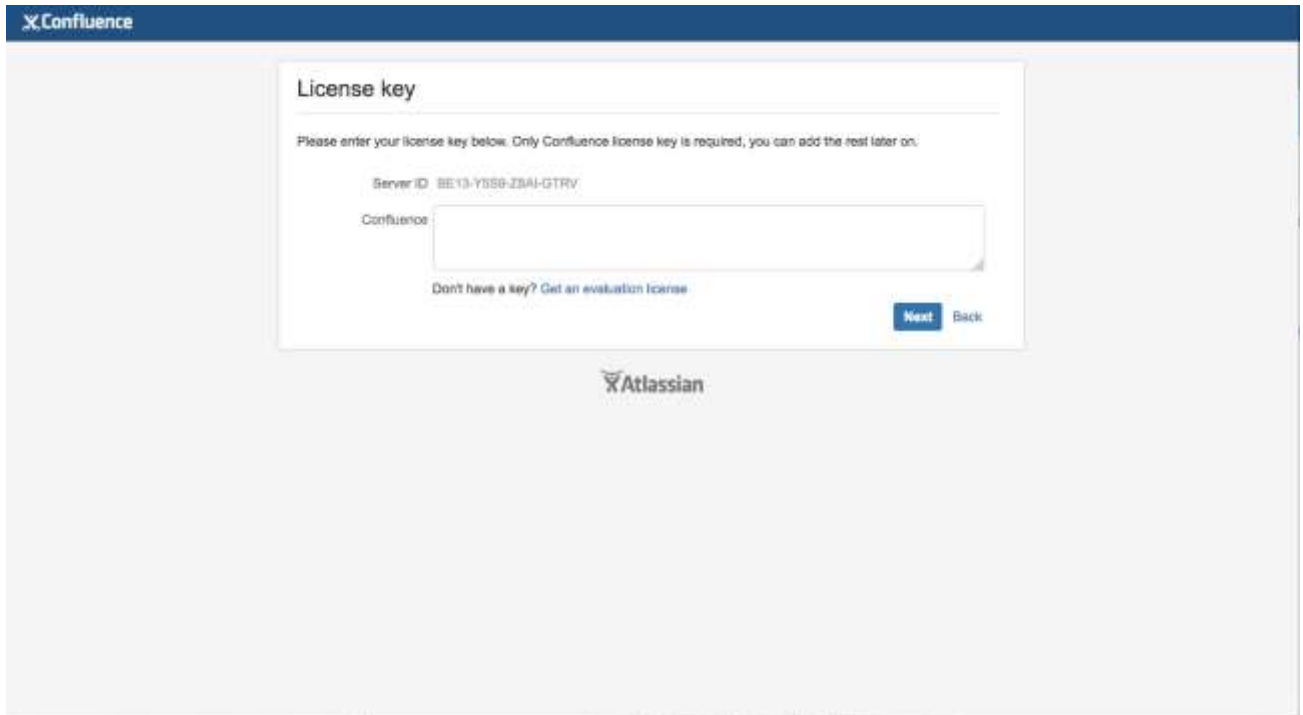
### Step 3. Configure Confluence Data Center

When you launch the Quick Start, it deploys a single Confluence node (Auto Scaling group of min=1 and max=1).

1. Click the URL displayed in the **Outputs** tab, of the Cloudformation stack, to go to the Confluence setup screen.
2. In the **Get add-ons** page of Confluence setup, click on **Next**. You can enable add-ons after the setup if you wish.



3. On the **License Key** page, enter a valid Confluence Data Center license, and click **Next**. If you don't have a valid license for Confluence Data Center, click **Get an evaluation license**. You'll be taken to [my.atlassian.com](http://my.atlassian.com) where you can generate a Data Center evaluation license. You can't use a Confluence Server license with this QuickStart.

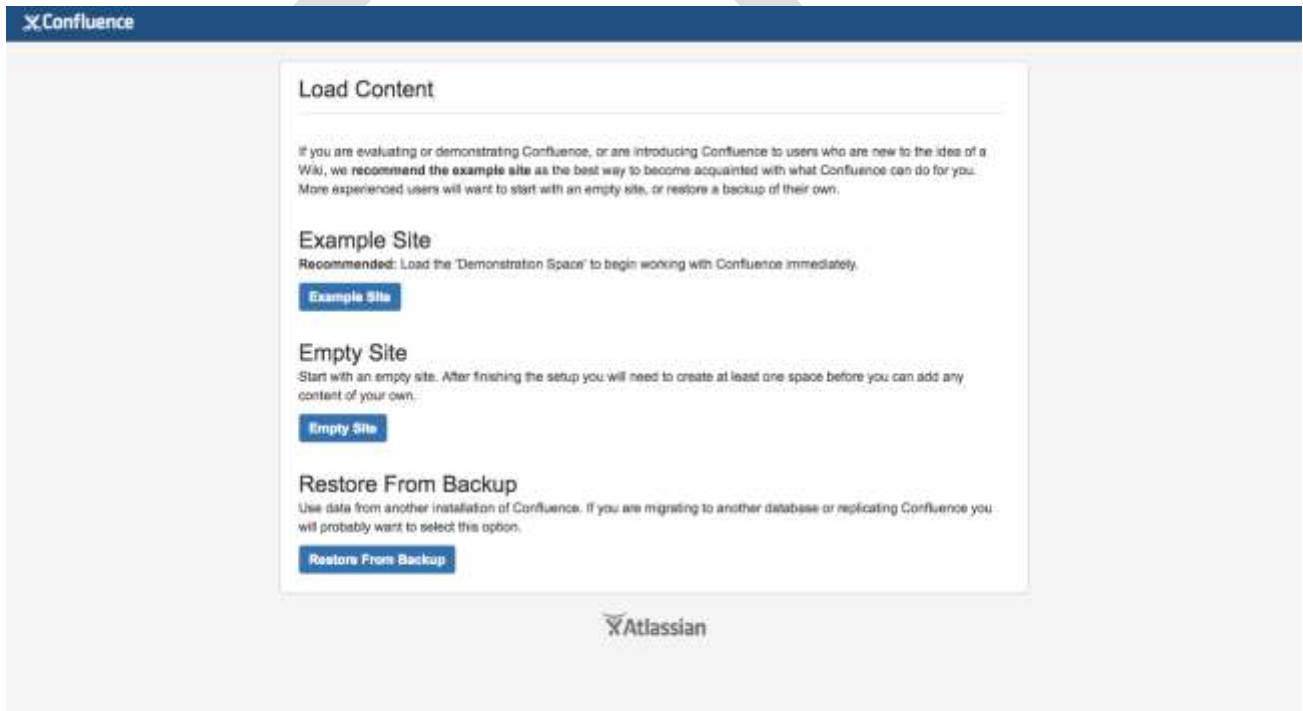


The screenshot shows the 'License key' page in the Confluence installation wizard. At the top left, there is a blue header with the Confluence logo and the word 'Confluence'. The main content area is white and contains the following elements:

- License key**: A heading followed by a text input field.
- Instructions**: A paragraph stating, "Please enter your license key below. Only Confluence license key is required, you can add the rest later on."
- Server ID**: A label followed by the value "BE13-Y559-ZBAI-GTRV".
- Confluence**: A label followed by a large text input field for the license key.
- Link**: A text link that says "Don't have a key? Get an evaluation license".
- Buttons**: Two buttons at the bottom right, "Next" (highlighted in blue) and "Back".

At the bottom center of the page is the Atlassian logo.

4. In the **Load Content** page, click on **Example Site**

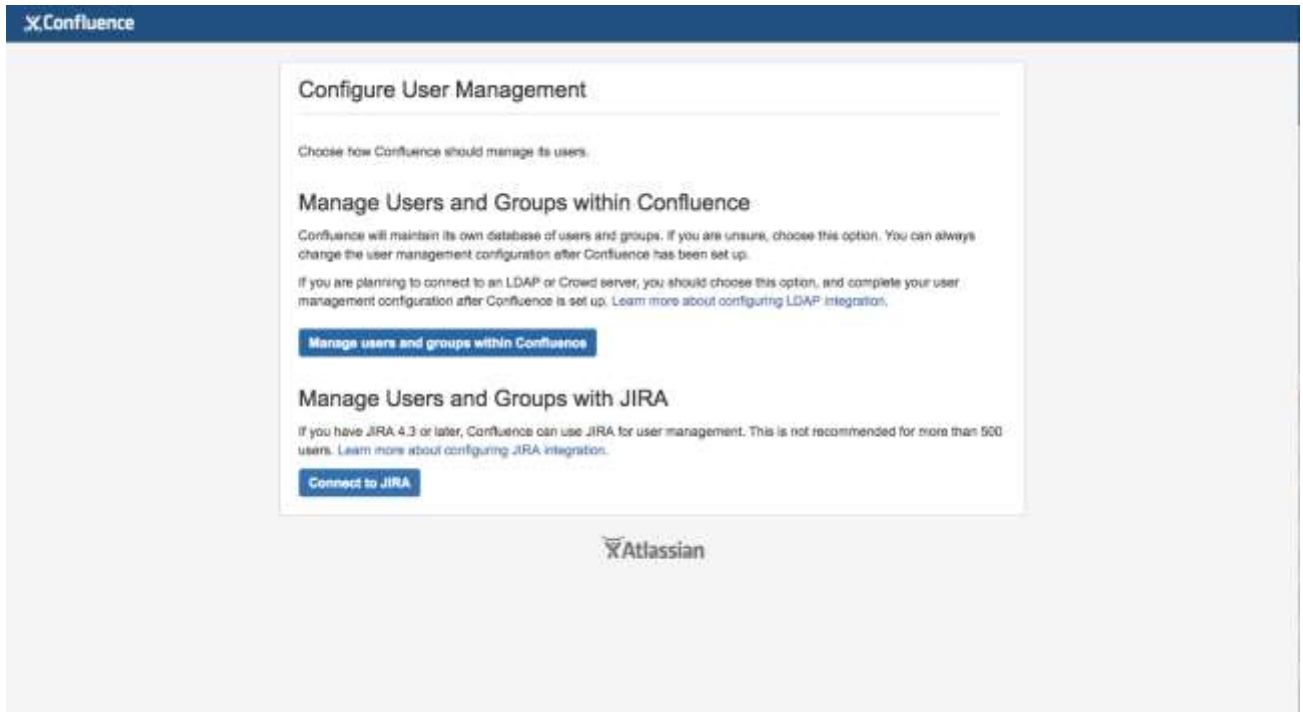


The screenshot shows the 'Load Content' page in the Confluence installation wizard. At the top left, there is a blue header with the Confluence logo and the word 'Confluence'. The main content area is white and contains the following elements:

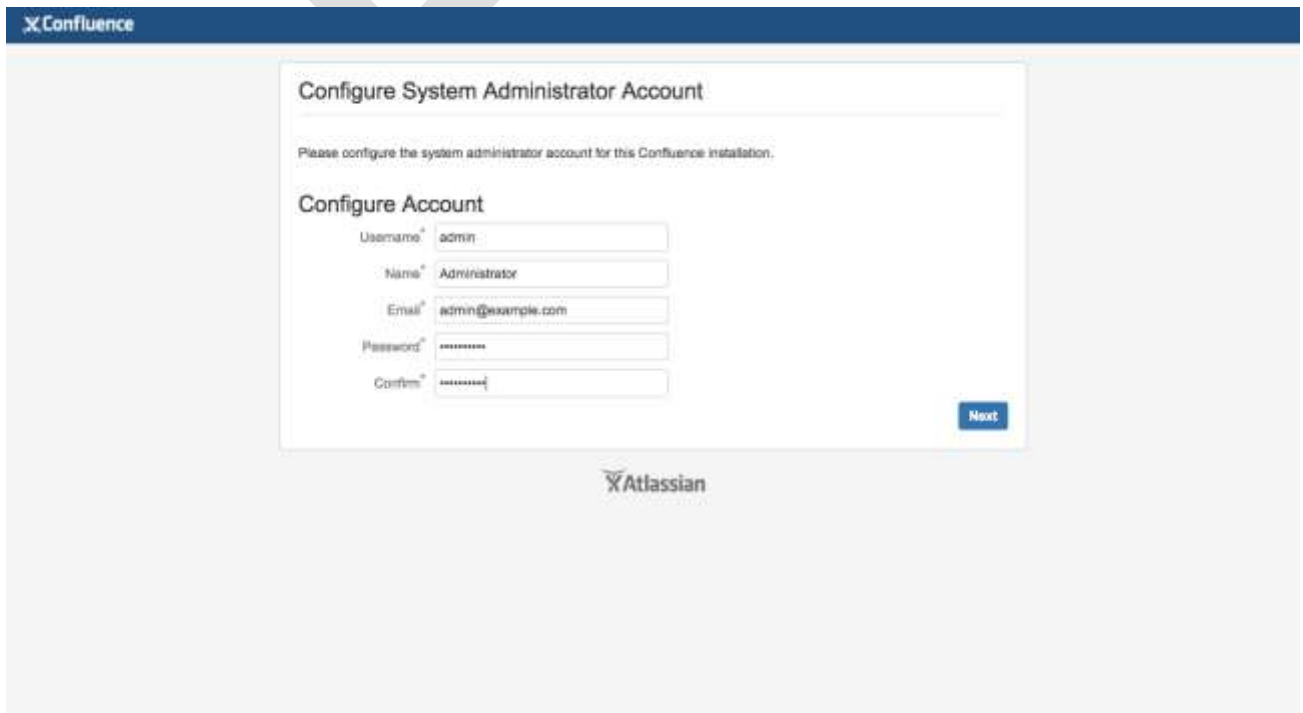
- Load Content**: A heading followed by a text input field.
- Text**: A paragraph explaining that for evaluation or demonstration, the 'example site' is recommended as the best way to become acquainted with Confluence. It also mentions that more experienced users might prefer an empty site or a backup.
- Example Site**: A heading followed by a paragraph: "Recommended: Load the 'Demonstration Space' to begin working with Confluence immediately." Below this is a blue button labeled "Example Site".
- Empty Site**: A heading followed by a paragraph: "Start with an empty site. After finishing the setup you will need to create at least one space before you can add any content of your own." Below this is a blue button labeled "Empty Site".
- Restore From Backup**: A heading followed by a paragraph: "Use data from another installation of Confluence. If you are migrating to another database or replicating Confluence you will probably want to select this option." Below this is a blue button labeled "Restore From Backup".

At the bottom center of the page is the Atlassian logo.

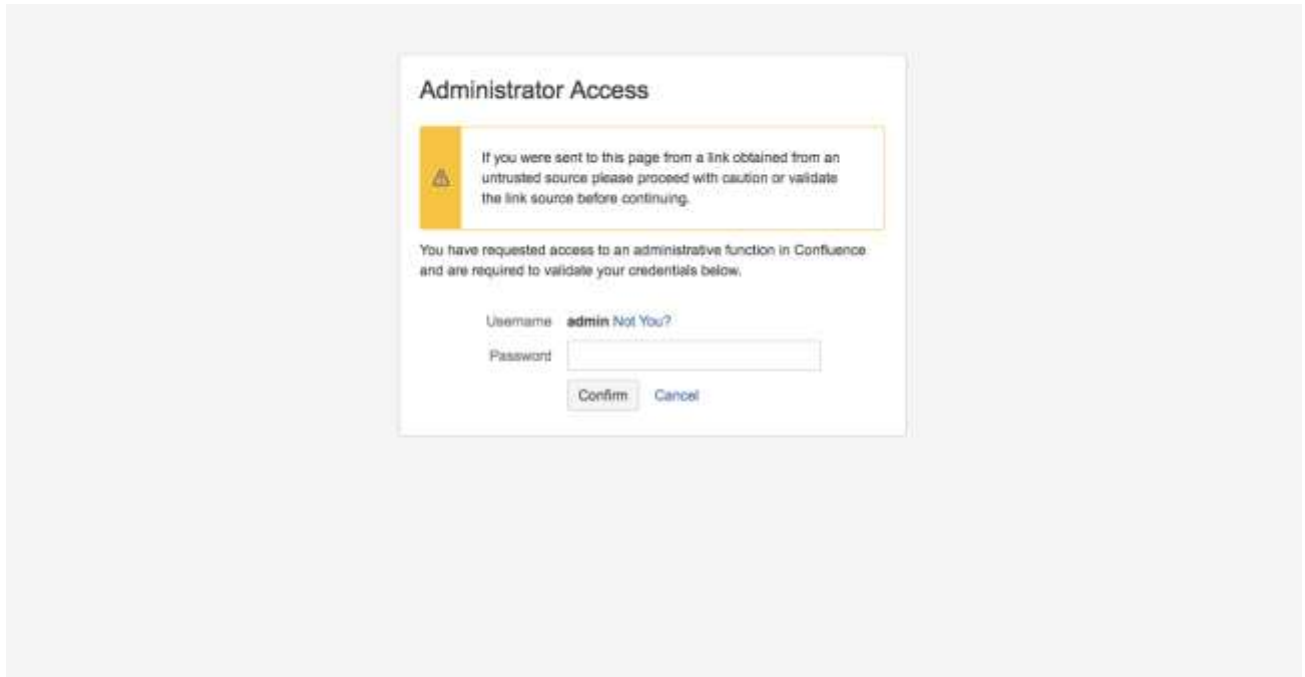
5. In the **Configure User management** page, choose **Manage users and groups within Confluence**



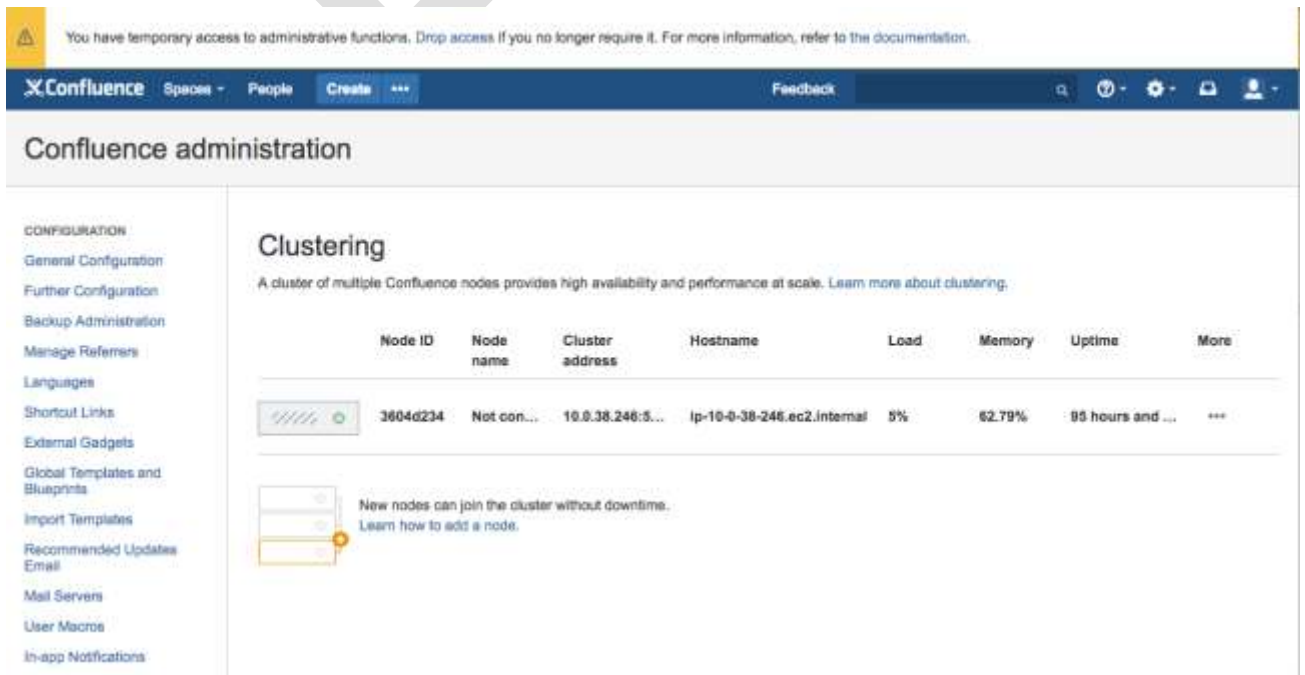
6. To set up Confluence Data Center, you need to create an Administrator account and password. The Administrator account has full access to all data in Confluence, so we highly recommend that you choose a strong password for this account. Enter the Administrator's user details in the **Configure System Administrator Account**, and click **Next**.



7. You should now see the **Setup Successful** page. Click the **Further configuration** hyperlink to go directly to the Confluence administration console, and log in with the admin user account created in the previous step.




8. Go to Clustering in the administration console sidebar. You should see a page similar to that illustrated in Figure 5, which shows that the node is ready for clustering.

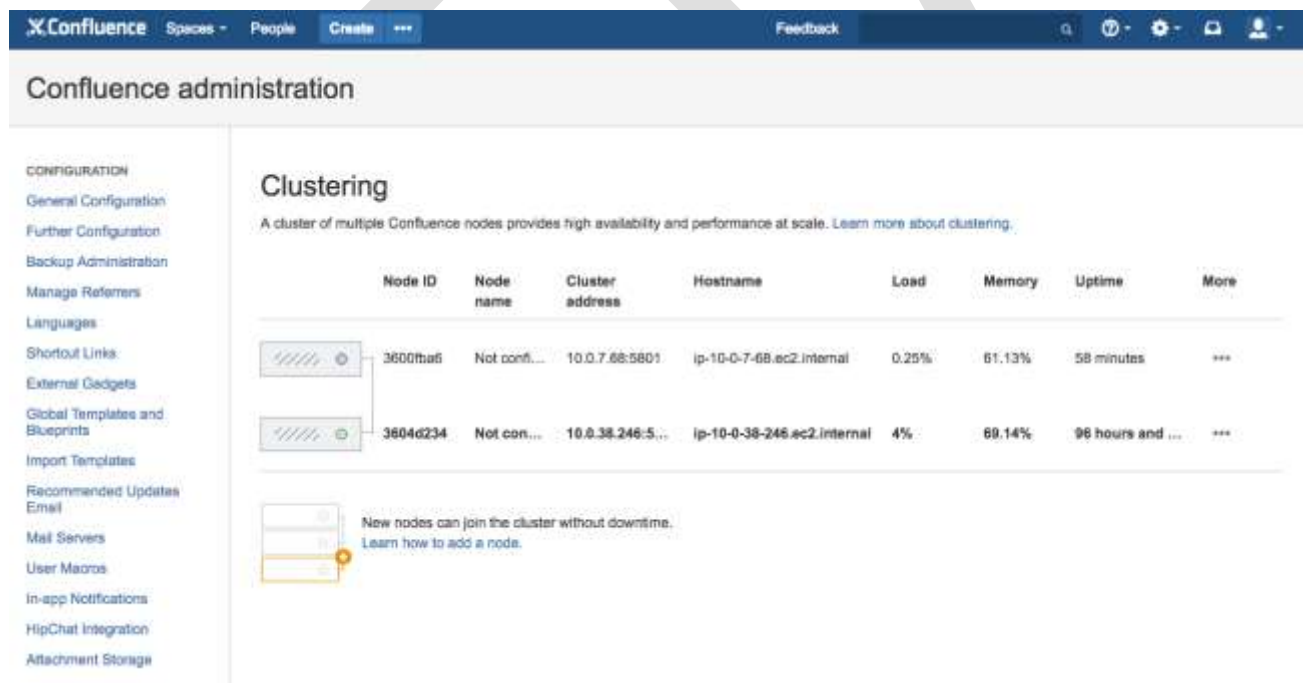


Your Confluence Data Center deployment is now in a state where you can add nodes that will

automatically cluster with your existing node.

## Step 4. Add instances to the Confluence Data Center cluster

1. Sign in to the AWS Management Console, use the region selector in the navigation bar to choose the AWS Region for your deployment, and open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation/>.
2. Choose the Confluence Data Center template. From the **Actions** list, choose **Update Stack**.
3. On the **Select Template** page, leave **Use current template** selected, and then choose **Next**.
4. On the **Specify Details** page, in the **Confluence Setup** section of **Parameters**, enter the value **2** for **Minimum number of cluster nodes**, and a desired value for **Maximum number of cluster nodes**, and click through to update the stack.
5. After the stack has finished updating, go to  > General Configuration > Clustering to confirm that the additional nodes have formed a cluster.



Confluence administration

**Clustering**

A cluster of multiple Confluence nodes provides high availability and performance at scale. [Learn more about clustering.](#)


Node ID	Node name	Cluster address	Hostname	Load	Memory	Uptime	More
3600baf6	Not confi...	10.0.7.68:5801	ip-10-0-7-68.ec2.internal	0.25%	61.13%	58 minutes	***
3604d234	Not con...	10.0.38.246:5...	ip-10-0-38-246.ec2.internal	4%	69.14%	96 hours and ...	***

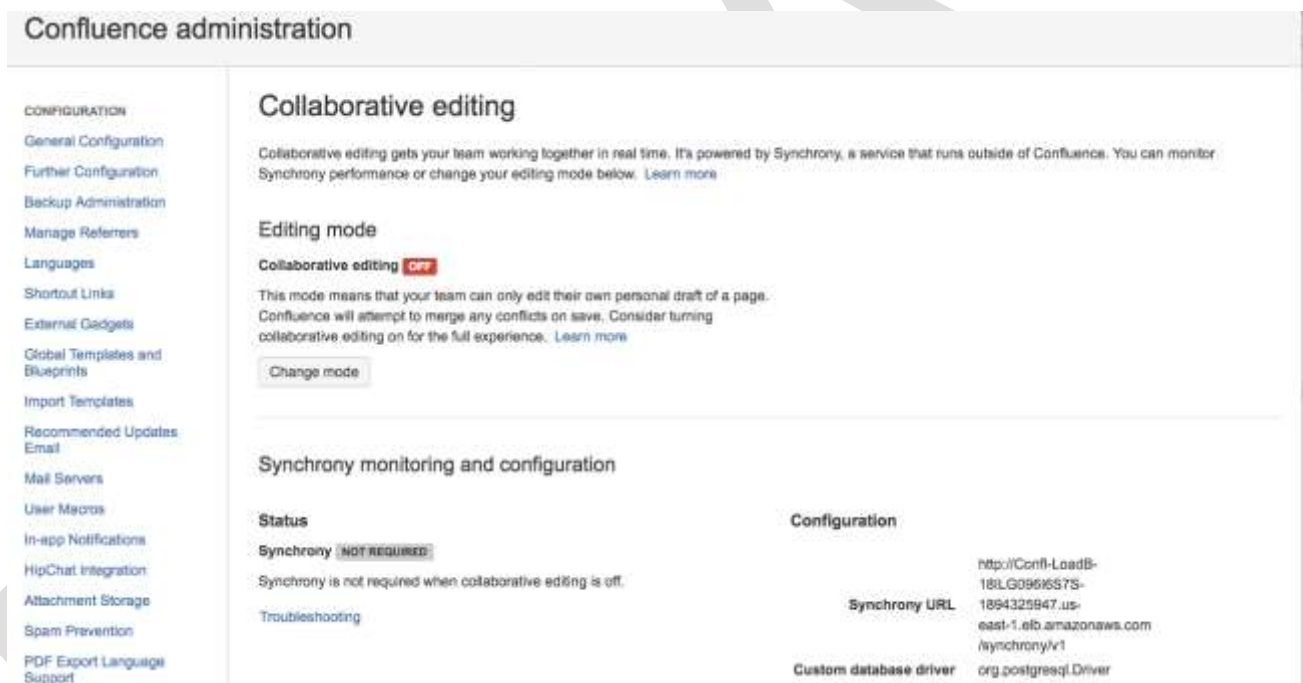
New nodes can join the cluster without downtime. [Learn how to add a node.](#)

## Step 5. Add nodes to the Synchrony cluster and enable collaborative editing

1. Sign in to the AWS Management Console, use the region selector in the navigation bar to choose the AWS Region for your deployment, and open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation/>.



2. Choose the Confluence Data Center template. From the **Actions** list, choose **Update Stack**.
3. On the **Select Template** page, leave **Use current template** selected, and then choose **Next**.
4. On the **Specify Details** page, in the **Confluence Setup** section of **Parameters**, enter the value **2** for **Minimum number of Synchrony nodes**, and a desired value for **Maximum number of Synchrony nodes**, and click through to update the stack.
5. After the CloudFormation stack has finished updating, in the EC2 console, confirm that the Synchrony nodes are in running state.
6. In Confluence Data Center, go to  > General Configuration > **Collaborative editing**, and verify collaborative editing is off.



**Confluence administration**

**Collaborative editing**

Collaborative editing gets your team working together in real time. It's powered by Synchrony, a service that runs outside of Confluence. You can monitor Synchrony performance or change your editing mode below. [Learn more](#)

**Editing mode**

Collaborative editing **OFF**

This mode means that your team can only edit their own personal draft of a page. Confluence will attempt to merge any conflicts on save. Consider turning collaborative editing on for the full experience. [Learn more](#)

[Change mode](#)

**Synchrony monitoring and configuration**

**Status**

Synchrony **NOT REQUIRED**

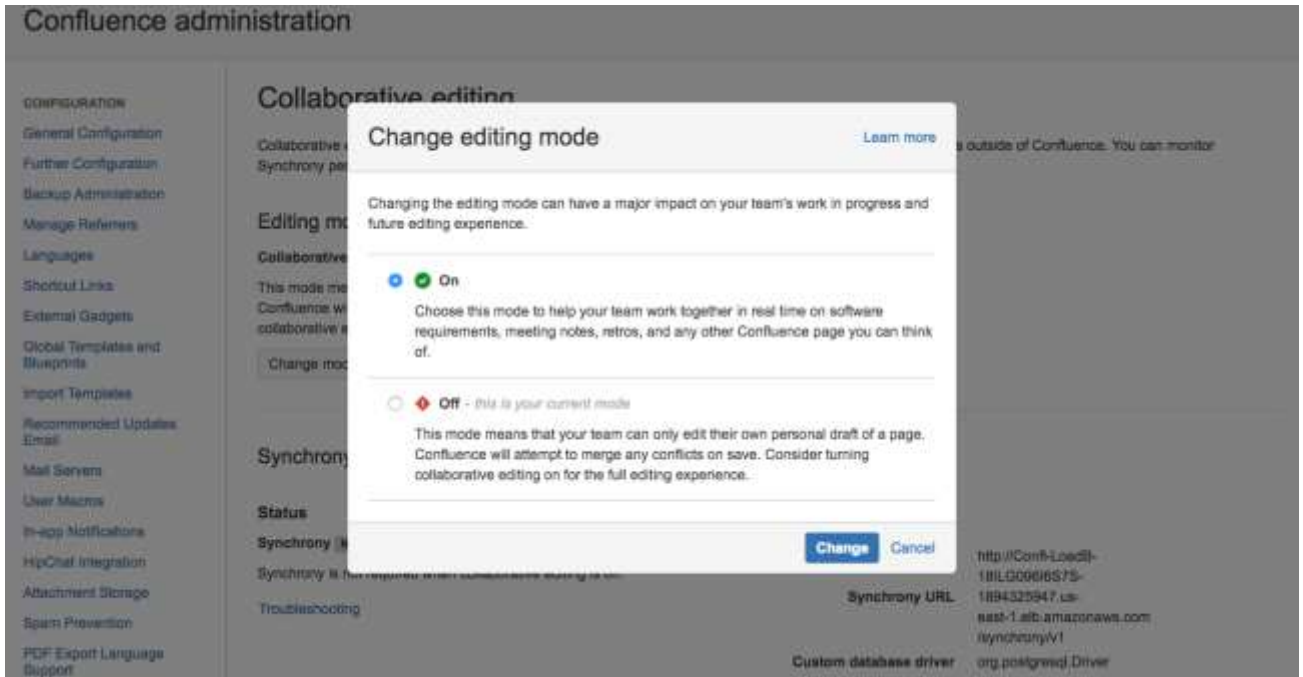
Synchrony is not required when collaborative editing is off.

[Troubleshooting](#)

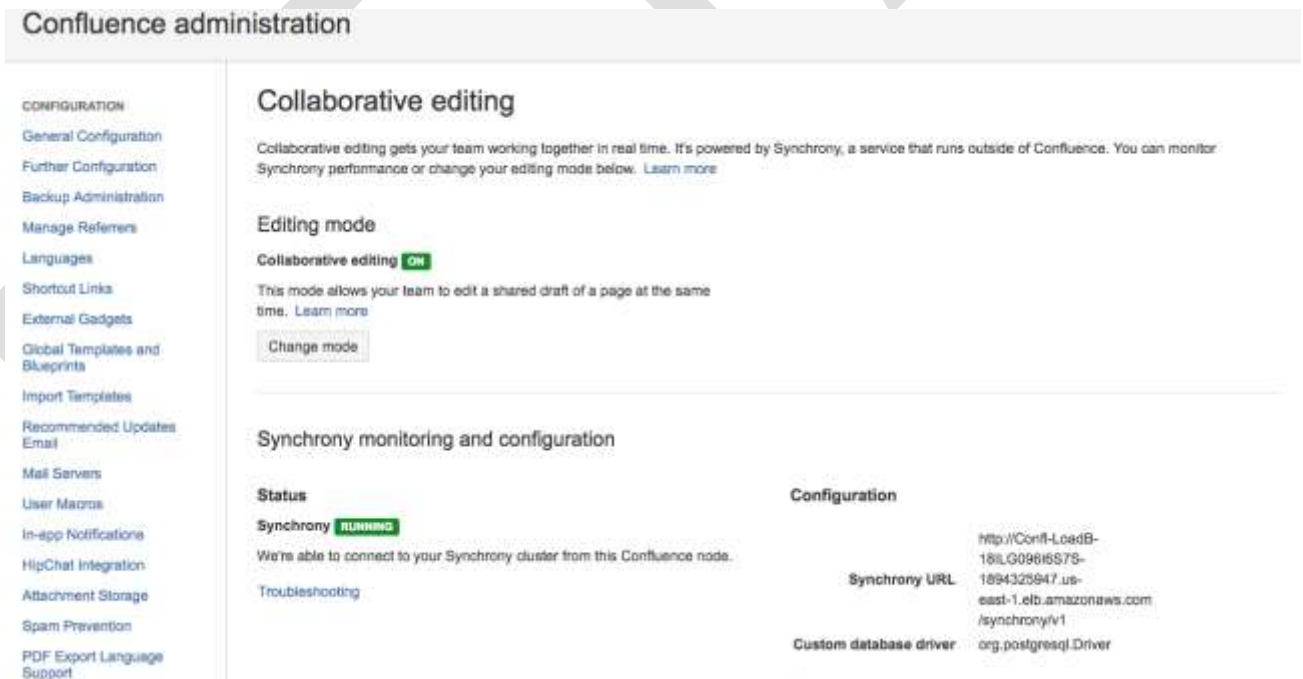
**Configuration**

<b>Synchrony URL</b>	http://Confl-LoadB-18LG0966S7S-1894325947.us-east-1.elb.amazonaws.com/synchrony/v1
<b>Custom database driver</b>	org.postgresql.Driver

7. Click the **Change Mode** button, then choose **On**, and click on **Change**.



- 8. Verify that Collaborative editing is on, and the status of Synchrony is “Running”



## FAQ

*Any tips or answers to anticipated questions. This could include the following troubleshooting information. If you don't have any other Q&A to add, change this heading to "Troubleshooting" and remove the Q/A headings below.*

**Q.** I encountered a `CREATE_FAILED` error when I launched the Quick Start. What should I do?

**A.** If AWS CloudFormation fails to create the stack, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained and the instance will be left running, so you can troubleshoot the issue. (You'll want to look at the log files in

`%ProgramFiles%\Amazon\EC2ConfigService` and `C:\cfn\log`.)

**Important** When you set **Rollback on failure** to **No**, you'll continue to incur AWS charges for this stack. Please make sure to delete the stack when you've finished troubleshooting.

For additional information, see [Troubleshooting AWS CloudFormation](#) on the AWS website or contact us on the [AWS Quick Start Discussion Forum](#).

## Additional Resources

### AWS services

- Amazon EC2  
<https://docs.aws.amazon.com/AWSEC2/latest/WindowsGuide/>
- AWS CloudFormation  
<https://aws.amazon.com/documentation/cloudformation/>
- Amazon VPC  
<https://aws.amazon.com/documentation/vpc/>

### Atlassian documentation

- Confluence Data Center  
<https://confluence.atlassian.com/display/DOC/Confluence+Data+Center>

### Quick Start reference deployments

- AWS Quick Start home page  
<https://aws.amazon.com/quickstart/>

## Send Us Feedback

We welcome your questions and comments. Please post your feedback on the [AWS Quick Start Discussion Forum](#).

You can visit our [GitHub repository](#) to download the templates and scripts for this Quick Start, and to share your customizations with others.

## Document Revisions

Date	Change	In sections
24 Jan 2017	Updated document with review comments from Atlassian	<a href="#">Links to revised sections</a>
9 Jan 2017	Initial publication	—

### 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. The software included with this paper is licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. A copy of the License is located at <http://aws.amazon.com/apache2.0/> or in the "license" file accompanying this file. This code is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.