

Confluence Datacenter Migration

- [Before you begin](#)
 - [Clustering requirements](#)
 - [Supported platforms](#)
 - [Terminology](#)
- [Set up Data Center](#)
 - [1. Upgrade Confluence Server](#)
 - [2. Apply Data Center license](#)
 - [3. Create a shared home directory](#)
 - [4. Start Confluence](#)
- [Add more Confluence nodes](#)
 - [5. Copy Confluence to second node](#)
 - [6. Configure load balancer](#)
 - [7. Start Confluence one node at a time](#)
 - [8. Test your Confluence cluster](#)
 - [9. Set up a Synchrony cluster \(optional\)](#)
- [Security](#)
- [Troubleshooting](#)

This page outlines the process for migrating an existing Confluence Server (non-clustered) site to Confluence Data Center (clustered).

If you're installing Confluence for the first time (you don't have any existing Confluence data to migrate), see [Installing Confluence Data Center](#).

If you're wanting to switch back to a non-clustered solution, see [Moving from Data Center to Server](#).

Your Confluence license determines the type of Confluence you have: Server or Data Center. Confluence will auto-detect the license type when you enter your license key, and automatically prompt you to begin the migration.

Not sure if you should upgrade from Confluence Server to Data Center? Learn more about the [benefits of Confluence Data Center](#).

Before you begin

Clustering requirements

To run Confluence in a cluster you must:

- Have a Data Center license (you can [purchase a Data Center license](#) or create an evaluation license at my.atlassian.com)
- Use a [supported](#) external database, operating system and Java version
- Use a load balancer with session affinity and WebSockets support in front of the Confluence cluster
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory)
- Use OAuth authentication if you have [application links](#) to other Atlassian products (such as Jira)

Supported platforms

See our [Supported Platforms](#) page for information on the database, Java, and operating systems you'll be able to use. These requirements are the same for Server and Data Center deployments. See [Confluence Data Center Technical Overview](#) for important hardware and infrastructure considerations.

We also have specific guides and deployment templates to help you running Confluence Data Center in [AWS](#) or [Azure](#). Check them out to find out what's required.

Terminology

In this guide we'll use the following terminology:

- Installation directory - The directory where you installed Confluence on a node.
- Local home directory - The home or data directory on each node (in non-clustered Confluence this is simply known as the home directory).
- Shared home directory - The directory you created that is accessible to all nodes in the cluster via the same path.
- Synchrony home directory - The directory where you configure and run Synchrony from (this may be on a confluence node, or on its own node)

At the end of the installation process, you'll have an installation and *local* home directory on each node, and a single *shared* home directory (a total of 5 directories in a two node cluster) for Confluence plus directories for Synchrony.

Set up Data Center

1. Upgrade Confluence Server

If you plan to upgrade Confluence as part of your migration to Data Center, you should upgrade your existing Confluence Server site as the first step.

2. Apply Data Center license

1. Go to [blocked URL](#) > **General Configuration** > **License Details**
2. Enter your new Confluence Data Center license key.
3. You'll be prompted to stop Confluence to begin the migration.

At this stage your home directory (configured in `confluence\WEB-INF\classes\confluence-init.properties`) should still be pointing to your existing (local) home directory.

3. Create a shared home directory

1. Create a directory that's accessible to all cluster nodes via the same path. The directory should be empty. This will be your **shared home** directory.
2. In your existing Confluence home directory, move the contents of `<confluence home>/shared-home` to the new shared home directory you just created.

To prevent confusion, we recommend deleting the empty `<confluence home>/shared-home` directory once you've moved its contents.

3. Move your attachments directory to the new shared home directory (skip this step if you currently store attachments in the database).

4. Start Confluence

The setup wizard will prompt you to complete the migration, by entering:

Configure cluster node

Cluster Name:

A name to identify the cluster.

Shared Home Directory:

A filepath to a shared directory. Should be accessible by all nodes in the cluster.

Interface:	Name	IPv6	IPv4
☒	eth0	fe80:0:0:0:f816:3eff:fe01:77a4%eth0	10.6.221.30
☐	virbr0		192.168.122.1

The network interface Confluence will use to communicate with other cluster nodes.

Join Configuration
☒ Use multicast
☐ Use TCP/IP
☐ Use AWS

Multicast Address:
☒ Automatically generate a multicast address

Create new cluster

- A name for your cluster
- The path to the shared home directory you created earlier
- The network interface Confluence will use to communicate between nodes
- A multicast address (automatically generated or enter your own) or the IP addresses of each cluster node
- How you want Confluence to discover cluster nodes:
 - Multicast - enter your own multicast address or automatically generate one.
 - TCP/IP - enter the IP address of each cluster node
 - AWS - enter your IAM Role or secret key, and region.

[AWS node discovery...](#)

Add more Confluence nodes

5. Copy Confluence to second node

To copy Confluence to the second node:

1. Shut down Confluence on node 1
2. Shut down your application server on node 2, or stop it automatically loading web applications
3. Copy the installation directory from node 1 to node 2
4. Copy the local home directory from node 1 to node 2
If the file path of the local home directory is not the same on nodes 1 and 2 you'll need to update the `<installation_directory>/confluence/WEB-INF/classes/confluence-init.properties` file on node 2 to point to the correct location.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

6. Configure load balancer

Configure your load balancer for Confluence. You can use the load balancer of your choice, but it needs to support session affinity and WebSockets.

You can verify that your load balancer is sending requests correctly to your existing Confluence server by accessing Confluence through the load balancer and creating a page, then checking that this page can be viewed/edited by another machine through the load balancer.

7. Start Confluence one node at a time

You must only start Confluence **one node at a time**. The first node must be up and available before starting the next one.

1. Start Confluence on node 1
2. Wait for Confluence to become available on node 1
3. Start Confluence on node 2
4. Wait for Confluence to become available on node 2.

The Cluster monitoring console ([blocked URL](#) > **General Configuration** > **Clustering**) shows information about the active cluster.

When the cluster is running properly, this page displays the details of each node, including system usage and uptime. Use the [blocked URL](#) menu to see more information about each node in the cluster.

[blocked URL](#)

8. Test your Confluence cluster

Remember, to test creating content you'll need to access Confluence via your load balancer. You can't create or edit pages when accessing a node directly.

A simple process to ensure your cluster is working correctly is:

1. Access a node via your load balancer, and create a new document on this node
2. Ensure the new document is visible by accessing it directly on a different node
3. Search for the new document on the original node, and ensure it appears
4. Search for the new document on another node, and ensure it appears

If Confluence detects more than one instance accessing the database, but not in a working cluster, it will shut itself down in a *cluster panic*. This can be fixed by [troubleshooting the network connectivity of the cluster](#).

9. Set up a Synchrony cluster (optional)

Synchrony is required for collaborative editing. You have two options for running Synchrony with a Data Center license:

- **managed by Confluence** (recommended)
This is the default setup. Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual steps are required.
- **Standalone Synchrony cluster** (managed by you)
You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required. See [Set up a Synchrony cluster for Confluence Data Center](#) for a step-by-step guide.

Head to [Administering Collaborative Editing](#) to find out more about collaborative editing.

Security

Ensure that only permitted cluster nodes are allowed to connect to the following ports through the use of a firewall and / or network segregation:

- 5801 - [Hazelcast](#) port for Confluence
- 5701 - [Hazelcast](#) port for Synchrony
- 25500 - Cluster base port for Synchrony
- 54327- Multicast port for Synchrony (only required if running Synchrony standalone cluster)

Troubleshooting

If you have problems with the above procedure, please see our [Cluster Troubleshooting guide](#).

If you're testing Confluence Data Center by running the cluster on a single machine, please refer to our developer instructions on [Starting a Confluence cluster on a single machine](#).