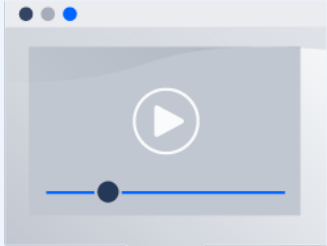


# Automation in Confluence Data Center

## Win back time and mitigate human error

As your organization's Confluence usage spreads and grows to enterprise scale, certain administrative tasks can feel daunting. While in some situations this means that these tasks become more complex, it's more often the case that they simply become more time-intensive. Instead of focusing your time on what's ahead, you're stuck putting out fires in a way that's tedious and prone to human error.

The appeal of automation has always revolved around two attributes: 1) its ability to speed up common tasks that allow you to be more efficient, and 2) improve the quality of your outputs. With the automation features that Confluence Data Center provides, you're able to achieve this and more, helping to win you back valuable time and fuel a more efficient way of working.



**Confluence and SharePoint: What's the difference?**

Every collaboration tool has different strengths, so how do you know which tools are best for your organization? [Watch our on-demand webinar](#) to learn what differentiates Confluence from file management systems like SharePoint.

---

## Taking advantage of availability controls in a clustered environment

One of the primary benefits of Data Center's clustered environment is the ability to use your load balancer to manage the use of your nodes. The two main ways that Confluence leverages automation in lockstep with Data Center's clustered environment are through the process of rebuild your index and setting up Synchrony - the engine that powers collaborative editing in Confluence.

 Reindex your cluster without downtime

 A simpler and safer Synchrony setup

# Set it and forget it: Hassle-free reindexing

A healthy search index is essential to the operation of your Confluence site. When something goes wrong, rebuilding your search index can be a major undertaking, especially in big, busy sites. If you run Confluence in a cluster, this is further compounded by the number of manual steps required to get the new index on to each node.

From Confluence Data Center 7.7 and on, we've reduced the number of steps you need to take, and removed the need for any manual file handling. When you rebuild the index on one node, Confluence will automatically propagate the new index file to every node in the cluster. Just sit back and watch it happen. A refreshed UI allows you to see exactly where the process is up to, from any node in the cluster.

Rebuilding the index is resource intensive, and it's important to minimize the performance impact on your team. A great option is to use your load balancer to temporarily direct traffic away from the node performing the reindex, leaving it to do its job in the background, while the other nodes in your cluster continue to serve requests.

To illustrate just how much time this wins back for you, [look at the process before automated re-indexing was an option](#) vs. the process now.

## Confluence 6.0-7.6

The previous experience required the following steps, each including roughly five sub-tasks for each:

- Preparation phase
- Cleaning the current index
- Rebuilding the index
- Copy the index file to a shared location
- Copy the index files from the shared home to the remaining nodes of the cluster
- Cleanup tasks
- Confirming content index is working on all nodes

## Confluence 7.7 and on

The automated experience includes only one task, and it's [accessible within the Confluence UI](#).\*

\*This method requires zero downtime, and you can also choose to remove the node performing the reindex from your load balancer, to further minimize any impact on users.

# Eliminate human error with a more straightforward Synchrony setup

Any admin will tell you that managing Synchrony is almost like managing a separate instance. The process has a large number of steps – not dissimilar to that of rebuilding your index – as well as files that need to be edited and options that can be overlooked, making it uniquely prone to human error.

The process itself includes the following steps:

- Provision your Synchrony nodes
- Create the Synchrony home directory
- Edit the start and stop scripts
- Add additional Synchrony nodes and configure your load balancer
- Start Confluence one node at a time

For those of you who want a simple setup with virtually no maintenance, Data Center admins can choose to let Confluence manage Synchrony for you. This means

Confluence will automatically launch a Synchrony process on the same node, and manage it for you with no manual setup is required. If you're happy with your current setup, or if you need to ensure the editor is highly available, [you can continue to run a standalone Synchrony cluster](#). The choice is yours.

---

While rebuilding your content index or setting up Synchrony are not everyday tasks, they do occur and can be extremely time-intensive. By taking these lengthy and detailed processes out of your hands, Confluence Data Center's automation features not only help you win back valuable time to be more efficient, but they ensure that these processes are performed correctly. When managing an instance at enterprise scale, automating tasks like these can make a big impact.

## Is it time to consider Confluence Data Center?

We're continuing to modernize Confluence Data Center so it can support your business objectives for as long as you need. If you're interested in learning more about what Data Center has to provide, you can test it out for free now.

[Try it now](#)

