

Search Indexing

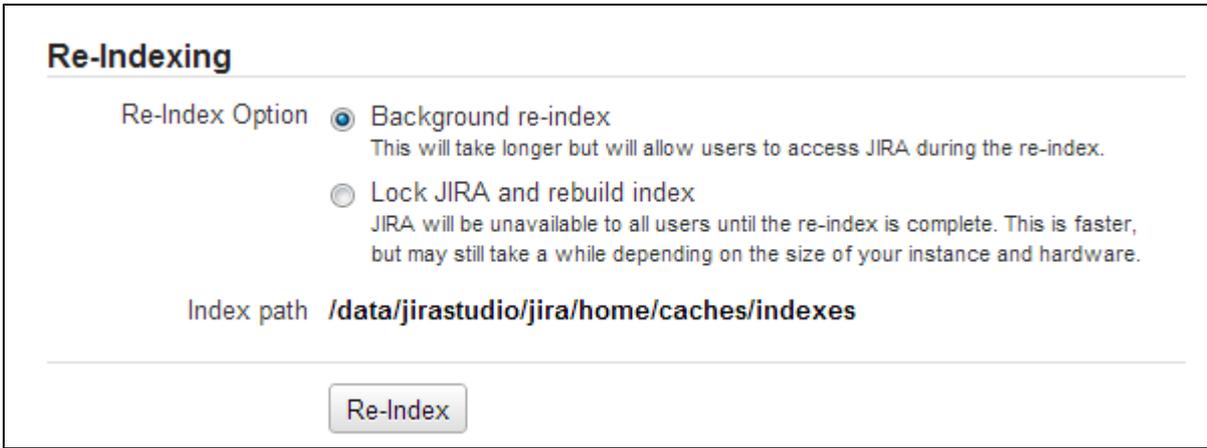
In order to provide fast searching, JIRA creates an index of the text entered into issue fields. This index is stored on the file system, and updated whenever issue text is added or modified. It is sometimes necessary to regenerate this index manually; for instance if issues have been manually entered into the database, or the index has been lost or corrupted.

See [Re-Indexing after Major Configuration Changes](#) for more information on when you should re-index.

Re-indexing JIRA

1. Log in as a user with the '**JIRA Administrators**' global permission.
2. Choose  **> System**. Select **Advanced > Indexing** to open the Indexing page.
 **Keyboard shortcut: 'g' + 'g' + type 'index'**
3. This page allows you to choose one of the following two re-indexing options:
 - **Background re-index** — This will re-index all issues in the background.
 - **Lock JIRA and rebuild index** — This will delete and rebuild all indices, including the comment and change history indices.

Screenshot: [Re-indexing JIRA](#)



Re-Indexing

Re-Index Option Background re-index
This will take longer but will allow users to access JIRA during the re-index.

Lock JIRA and rebuild index
JIRA will be unavailable to all users until the re-index is complete. This is faster, but may still take a while depending on the size of your instance and hardware.

Index path `/data/jirastudio/jira/home/caches/indexes`

Re-Index

Which re-indexing option should I use?

The **Background re-index** option should be used in the majority of circumstances, particularly following changes to the configuration. It will generally take significantly longer to perform than the **Lock JIRA and rebuild index** option, but it allows JIRA to remain usable while it is being done. There will however be a performance impact on JIRA as a whole. *We recommend that you perform this option during a low usage period.* The actual impact of running the **Background re-index** option will depend upon the customer's particular hardware and software installation as well as how many issues are in the system.

The **Lock JIRA and rebuild index** should be used when:

- the indices are corrupt, which may be caused by a system or disk failure or
- it is more important to have the re-index completed quickly than to have JIRA continuously available. The **Lock JIRA and rebuild index** option may be in the order of twice as fast as a background re-index.

The following table summarises the differences between the two options:

Background re-index	Lock JIRA and rebuild index
Slower to complete.	Faster to complete (may be up to twice as fast).
JIRA can be used by users during re-index.	JIRA cannot be used by users during re-index.

Can be cancelled at any time.	Cannot be cancelled once started.
-------------------------------	-----------------------------------

i NFS storage for JIRA indexes is not supported. See this page for details: [Supported Platforms](#).

Backing up and recovering your index

Enabling index recovery will cause a snapshot of the indexes to be taken periodically. This allows you to recover your index quickly, rather than rebuilding the index, if there is a failure. This is particularly useful if you have a large JIRA installation and you cannot afford for it to be offline for long. If you have a small JIRA instance, it may not be worth enabling index recovery, as it rebuilding the index won't take much time.

Whether a full index rebuild is faster than recovering from a snapshot depends on a number of factors, including how recent the snapshot being recovered was taken. Large and complex installations should test this process on a development/testing server before relying on it in production.

To enable index recovery:

1. Navigate to the **Indexing** page (as [described above](#)).
2. Click **Edit Settings** to enable index recovery and choose the frequency of snapshots.
 - Snapshots are stored in the `<yourjirahome>/export/indexsnapshots` directory.

To recover an index:

1. Navigate to the **Indexing** page (as [described above](#)).
2. Enter the name of the previously saved index in **File name** and click **Recover**.
 - JIRA will not be available during the recovery of the index.
 - If changes were made to the configuration that required a re-index after the snapshot was taken, then you will need to do a background re-index after the recovery. Note, JIRA will be available after the recovery.

Additional information

- JIRA will retain the last three snapshots at any time (in `<yourjirahome>/export/indexsnapshots`). Older snapshots will be automatically deleted. Note, snapshots may occupy considerable disk space and may need to be moved to offline storage or deleted as appropriate.
- The snapshot process is a relatively lightweight process and does not place much of a load on the system.
- The process of taking a snapshot will require temporary disk space equivalent to the index size. The resulting snapshots will each be about 25% the size of the index.
- All issues will be re-indexed appropriately during the recovery, including issues that were added, updated or deleted after the snapshot was taken.
- You can use the index recovery process to bring your index up to date, if you need to restore your JIRA database. The index snapshot must pre-date the database backup being restored.

Re-indexing a single project

If you have made a configuration change that affects a single project, you can re-index just that project. See [Re-indexing after Major Configuration Changes](#) for more information on when you should re-index.

To re-index a single project:

1. Navigate to the desired project and click the **Administration** tab.
2. Click **Actions > Re-index project** to start re-indexing the project.