

Structure 3.5 Release Notes

i 30th of January, 2017

Structure 3.5 is a major release that adds grouping by issue links, filtering by JQL or S-JQL in Structure Gadget and bidirectional links extender. This release also has quite a few other updates, important security fixes, performance and stability improvements.

! **Important for Structure 2.x users!** If you currently use Structure 2.11.2 or an earlier version, it is very important that you read [Structure 3.0 Release Notes](#) before upgrading.

[Download the latest Structure and Extensions](#)
[Structure 3 Demo Server](#)
[Structure 3 Quick Start Guide](#)

1. Version Highlights

Structure 3.5 is a major update in the Structure 3 series. It contains several new features:

- Grouping by issue links
- Specifying JQL or S-JQL as a filter in Structure dashboard gadget
- Bidirectional issue link extender
- Default level limit for extenders

This version also contains a critical security patch.

Upgrade is required for all JIRA instances running Structure 3.0 and using JIRA 7.x. If you are using a previous version of JIRA, please make sure to install a version with a security patch – either [Structure 3.3.4](#) for JIRA 6.4.x or [Structure 3.2.2](#) for JIRA 6.3.x.

! You need to have a Structure license with active maintenance (expiring not earlier than January 30th, 2017) to run use the new version or the patches.

2. Changes in Detail

2.1. Grouping by Issue Links

There's a new grouper that can arrange a list of issues based on the incoming links of a certain type.

For example, if you have a list of Features and there are links between Requirements and Features, you can have the Features placed under the Requirement that is linked to it. If there are two incoming links from different Requirements, the Feature will be placed under both parent items.

2.2. Improved Filtering in Structure Gadget

Before Structure 3.5, the content shown in a Structure dashboard gadget could be filtered with a JIRA's saved filter. In this version we made it possible to also filter by a text query, a JQL query or an S-JQL query.

Support for S-JQL is especially important here, because it allows you to see a part of a big structure based on a hierarchical query. For example, you can display only "critical" issues and all their children, or a contents of a specific folder.

Documentation: [Structure Gadget](#), [S-JQL Cookbook](#)

2.3. Bidirectional Issue Link Extender

The screenshot shows the Structure dashboard interface. At the top, there is a toolbar with icons for edit, delete, copy, paste, and close. A dropdown menu is open, showing 'Automation' and 'Group by...'. The 'Group by...' dropdown is currently set to 'link', and 'Issue Links...' is highlighted. Below the dropdown, the main content area displays a 'Summary' section for 'Demo Structure'. A blue bar indicates the grouping rule: 'Group each issue under issue that justifies it'. The structure is organized into two requirements: 'Requirement A' and 'Requirement B'. Under 'Requirement A', there are three features: 'Feature 1', 'Feature 2', and 'Feature AB3'. Under 'Requirement B', there are three features: 'Feature B1', 'Feature B2', and 'Feature AB3'. A red text overlay on the right side of the screenshot reads: 'Requirements "justify" features and features are grouped by the justifying requirements.'

Sometimes a link type does not have a "direction". An example is the standard JIRA's link type, "relates to". However, technically, there are two directions, but they are named identically. This was making it hard to build a hierarchy of related issues, because you would have to add a separate extender for each direction.

Starting with Structure 3.5, you can select both directions in the Links Extender, and it will produce a clean hierarchical view of the linked issues, regardless of the directions of the links.

The newly introduced Issue Links Grouper also support bidirectional operation.

2.4. Default Level Limit for Extenders

All extenders have "Levels" setting, which govern on what levels down the hierarchy the extender is applied.

Without a limit, the extender is applied at all levels under the parent that contains the extender. This sometimes led to problems where the extender was not configured properly and it resulted in structures 100 levels deep or more.

In Structure 3.5, the default setting for all extenders is to apply to 10 levels under. Should you need to increase that number, just edit the extender settings.

2.5. Security Patch

We have addressed two medium-to-critical security issues, affecting all Structure versions starting with 3.0.

Details of the issues are not publicly disclosed at this time, but you can find more information in the restricted Structure Security google group: <http://groups.google.com/group/structure-security-list> (please mention your Structure SEN number or license ID when requesting access).

2.6. Notable Fixes and Improvements

The following issues have been addressed in Structure 3.5:

- JQL queries containing structure() function should execute faster.
- Backlog template now offers more fields to group by.
- Fixed: Structure does not react to issue updates after upgrading Universal Plugin Manager plugin.
- Fixed: Structure requiring unlimited commercial license in case one of the application licenses for JIRA is an evaluation license.

The new version also contains other bug fixes and improvements.

3. Supported Versions

Structure 3.5 and all extensions support JIRA versions from 7.0 to 7.3.x. All editions of JIRA (JIRA Core, JIRA Software, JIRA Service Desk) are supported. JIRA Data Center is supported.



Structure 3.5 is the last version of Structure to support JIRA 7.0. Structure 3.6 will most likely support JIRA versions starting from 7.1.

This release is backwards-compatible with Structure 3.4, so all Structure extensions – [Structure.Pages](#), [Structure.Testy](#), as well as [Colors](#) plugin and other plugins that integrate with Structure 3.4 – should work with the new version.

4. Installation and Upgrade

4.1. Installing Structure

If your JIRA server does not have Structure yet, the installation is simple:

1. Download and install Structure add-on, either from Atlassian Marketplace or from [Download](#) page. Pick the correct version based on your JIRA version!
2. When Add-on Manager reports about successful installation, click Get Started to visit a page with important guidance for the JIRA administrator. You may want to also check out the user's Get Started page, available under "Structure" top-level menu.
3. Monitor `catalina.out` or `jira-application.log` for log messages from Structure.

4.2. Upgrading Structure

Extend with Linked Issues

Runs As Admin

Link Type

Link Direction parent issue relates to → sub-issue relates to ← relates to

Extend Levels

Extender does not consider groups on the current nested levels.

Allow changes via Structure



- If you're upgrading from version 2.11.2 or earlier, please read [Structure 3.0.0 Release Notes](#).
- If you have installed Structure 3 beta or release candidate, please read [Structure 3.0.0 Release Notes](#).

Upgrade procedure from versions 3.x.x is simple:

1. Consider backing up JIRA data. Use **Administration | System | Backup System**. Starting from version 3.0.0 Structure data can be backed up together with JIRA data. (If you have a large instance and have proper backup strategy in place, you may skip this step.)
2. Install the new version of the plugin.
3. Upgrade [Structure.Testy](#) and [Structure.Pages](#) add-ons if you're using them.
4. Monitor `catalina.out` or `jira-application.log` for warnings or errors.

5. Enterprise Deployment Notes

Structure 3.5 is a fairly significant update. Some of the improvements are directly related to performance, stability and important Structure functionality.

5.1. Security Patches

Please note that the plugin contains security patches. It is important to upgrade to avoid potential data exposure. For the sake of safety, for time being we share the details only in the closed [Structure Security Google Group](#).

If you cannot upgrade to Structure 3.5 because you have JIRA 6, please check out [Structure 3.3.4](#) and [Structure 3.2.2](#).

5.2. Plugin Start-Up Sequence

We have made further improvements to the plugin start-up sequence, where it needs to orchestrate the components start-up with JIRA start-up. Over time, this has proven to be tricky, but our current solution should be very robust. The improvements in Structure 3.5 concern the case where the UPM (Universal Plugin Manager) plugin is updated. This causes the whole plugin system to restart, and it is yet another case with different JIRA behavior.

Although we are pretty sure in the current solution, keep an eye out at the start of JIRA or when UPM is upgraded, watch the logs, and if there are any warnings or errors, please let us know.

5.3. Automation Engine Locking

We have made improvements to the Automation Engine, which should reduce the number of cases where you see timeouts in the logs ("cannot acquire lock in 30000ms") for various structures.

Please do not confuse this with other locks! JIRA has a lot of locks and many have 30 second timeouts. These locks in particular don't block other JIRA actions, but come into play when multiple users open the same structure (or when a single user opens the same structure in multiple browser windows). The locks protect JIRA server from doing the same work in parallel. If there is a timeout, it means that some structure is generated for a long time – for more than 30 seconds – so either Automation engine has to crunch a lot of issues, or the system is slow generally, or there is some configuration problem. In any case, you can inspect the structure, finding it by its ID, which is given in the same warning in the logs.

However, this change is in a critical and very complex part of the product, so there's a small risk of locking issues, related to the aforementioned warnings in the logs, structure lock-downs (when a structure is not displayed, the users just see empty content), or redundant calculations on the server resulting in higher server load.

5.4. Asynchronous Index Writing

Structure keeps an index of "which structures contain which issues". This index is not critical to operation and is used only sometimes – for example, when selecting which structure to display on the issue page. The index also ignores the generated content; it only works for issues added manually or by synchronizers.

In some configurations, updating that index introduced unwanted latency, so we moved its updates to an asynchronous job, executed by a background thread. While this change is not likely to cause any trouble, the index updates may be delayed on a heavily loaded system and that may result in an incorrect structure being picked to show on an issue page.

5.5. Testing on Staging Environment

It is recommended to try Structure 3.5 on a staging environment before upgrading. The suggested tests are:

- Plugin lifecycle.
 - Installing and uninstalling. Disabling and enabling.
 - Upgrading UPM, upgrading JIRA Agile plugin.
 - Restarting JIRA. (There are no Data Center-specific changes in this version.)
- Try making changes to structures in rapid bursts. Watch out for JVM memory statistics and CPU consumption. See if there is a lot of activity even after the bursts stop.
- Check out generators that work with issue links – the Extender and the Grouper. Try using unidirectional and bidirectional configurations on large structures.
- Check Structure Gadget.
 - Check if it works on JIRA dashboard.

- If you have Confluence instance connected and Structure gadget is used there, check that it works in your configuration.
- Check "Export" button, exporting to Excel and to Printable page. Try on large structure and under a user account without administrative privileges.
- Automation locking.
 - Try opening a structure with automation from multiple browsers. You can build a temporary structure with considerable work for generators – for example, inserting 5,000 issues and using Issue Links extender.
 - Try using transformations (such as filters) on a structure, along with opening that structure in other browsers.
- Automation load testing.
 - Try opening user structures that contain considerable amount of issues, 10,000 or more. Click grid header to make Structure sort structure by some field. Do that in multiple browser tabs and using different columns.
 - Try creating a new structure and populate it with the help of an Inserter (use Automation | + | Insert | JQL) and raise the issue limit to 10,000. Repeat this several times with different structures.
 - Watch log files for errors and warnings.
- Automation stress testing.
 - Emulate peak number of users opening the most popular structure.



Should you have any questions on Enterprise Deployment, let us know at support@almworks.com.