Structure 3.4 Release Notes

12th of December, 2016
Structure 3.4 is a major release that adds quick transformations (a.k.a. quick filters), integration with Portfolio for JIRA and stable public API. It also has quite a few other updates, bug 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 3.4 is a major update in the Structure 3 series. It contains several new features:

- Quick Transformations
- Integration with Portfolio for JIRA
- Issue Details layout on the Project page
- Grouping by text fields
- Performance improvements
- Public API

Upgrade is recommended to all users.

2. Changes in Detail

2.1. Quick Transformations

Quick transformations, also known as "quick filters", are predefined transformations that may be configured by the structure owner and then quickly turned on, separately or in combination.

Previously, Structure had only two predefined quick transformations – Unresolved and Assigned to me. Now it is possible to define your own quick transformations and share them with the team.

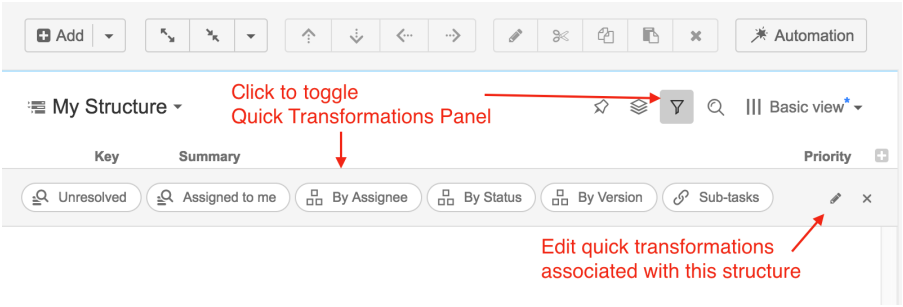

Documentation: Quick Transformations

2.2. Integration with Portfolio for JIRA

We have added support for Parent Link field in Portfolio for JIRA.

Now, when you have Portfolio installed, you get two new generators:

- An extender called **Child Issues (Portfolio)**, which adds "child" issues, as defined by Parent Link field, under their "parent" issues.
- A grouper called **Parent Issue (Portfolio)**, which goes the other direction – groups issues based on the value in their Parent Link field, which adds the parent issues to structure.

Both generators support updates, so you can change the Parent Link field by dragging and dropping issues.

This feature allows you to recreate in Structure the hierarchy you have in Portfolio and then build on top of it.

Supported Portfolio for JIRA versions: 2.1.5 or later.

2.3. Issue Details Layout on Project Page

Good news for those of you who prefer to use Project tab with Structure, rather than Structure Board – we have added "Grid + Details" layout to those available on the project page.

See Viewing Issue Details for documentation about this layout.

2.4. Grouping by Text Field

We have added a special generator that lets you group by a text attribute. A common example is grouping by the value of a custom text field.

To use this grouper, select **Automation | +**, then **Group**, and find **Text Attribute...** in the drop-down list. In the dialog that follows, select the desired attribute.

This grouper does not handle changes. You cannot move an issue from one group to another – you will need to update the field.

2.5. Performance Improvements

There are a few performance improvements in this version:

- We have replaced one of the core components, *RowManager*, that has been causing some trouble in the past. The new version should be much faster.
- There were significant fixes related to synchronizer performance.
- Backup creation with history is now much faster.
- We fixed some performance issues with structure() function used on Agile boards.
- Fixed performance issues caused by too many groups in the user directory.

2.6. Public API

We are finally releasing our stable API for Structure 3 series.

With the introduction of the new architecture earlier this year we had to make breaking changes to the API and took some time to work out the kinks and make it stable. Structure 3.4 ships with Structure API 16.0.0, which can be used in your custom plugins or from integration scripts.

We have documented the most important parts of the API but the documentation work still continues. You can expect that API coverage will increase with every new version.

Documentation: Structure Developer's Guide

2.7. Notable Fixes and Improvements

The following other issues have been addressed in Structure 3.4:

- Structure restore writes progress percent to the logs, so JIRA administrator can check it.
- Fixed: Compatibility issues with Zephyr plugin.
- Fixed: Incorrect sorting direction when adding a sorter, in certain cases.

There are also other bug fixes and improvements.

3. Supported Versions

Structure 3.4 and all extensions support JIRA versions from 7.0 to 7.2. All editions of JIRA (JIRA Core, JIRA Software, JIRA Service Desk) are supported. JIRA Data Center is supported.

Structure 3.3 was the last version of Structure to support JIRA 6. Structure 3.4 supports JIRA versions starting from 7.0.

3.1. Structure.Pages Upgrade Required

Please note that if you're using Structure.Pages, you will need to upgrade it to version 1.2.0.

See Structure.Pages 1.2 Release Notes for more information about updates in Structure.Pages.

3.2. Structure.Testy Upgrade Required

Please note that if you're using Structure.Testy, you will need to upgrade it to version 2.2.0.

Structure.Testy 2.2 includes new high-level API that makes it easier to integrate your scripts with Testy and update test statuses programmatically.

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

- 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

We now include an additional section in Release Notes that aims to address the concerns of deploying the upgrade at a large enterprise. In this section we will suggest how the changes in the new version may affect stability and performance of JIRA and provide ideas for testing the new version on a staging environment.

Structure 3.4 is a fairly serious update. A lot of code underwent refactoring, performance optimization and testing.

The following changes are quite important for the large-scale instances.

5.1. New implementation of RowManager component

We have rewritten the component that is central to the Structure's operation. The component's role is to store temporary rows as they are generated by Automation engine. The previous version was using "MapDB" library and offloading the stored values to disk. This turned out to be unreliable in some circumstances and under a heavy load.

The new version is much simpler and efficient, however, it does store this information in memory. The amount of data stored is not overly large, but current implementation never removes data from memory once it is stored. Therefore, on an active enough and large enough instance, this component may eventually grow to take a lot of memory resources or even cause an OutOfMemoryError.

Therefore, JIRA administrators on large instances are advised to keep an eye on the amount of Heap memory consumed by JIRA's Java process, and if it reaches what you'd consider to be a high threshold, you can clean up the memory taken by Structure with a simple operation: disable Structure plugin and then enable it back.

We are continuing our work on improving this component. Our next versions will contain improvements for this aspect.

5.2. Indexing on Data Center

We have identified a problem that can possibly manifest on JIRA Data Center cluster, when a node wants to recover indexes (from a downloaded index zip from another node), but it happens to be running a Structure synchronizer. As synchronizers used to take read lock on reindexing, they prevented full-stop reindex and index recovery from happening at the same time with synchronization. But on large instances a single synchronization job may execute for a long time – and this caused "Wait attempt timed out" error on the reindexing node.

The new locking mechanism, which is added in Structure 3.4, does not take read lock for more than a fragment of a second. While still ensuring that Structure's synchronizers and generators would execute based only on consistent query results, it allows JIRA full reindex or index recovery to happen at any time.

5.3. Improved Start-up Sequence

We have changed the way Structure plugin starts. Since we only support JIRA 7.0 or later with this version, we were to use some of the improvements Atlassian team prepared in JIRA 7 to increase reliability of the start sequence.

The related problem that happened to our customers in the past was that the plugin wouldn't start – it would wait for Active Objects component (provided by JIRA) to become available, but fail after 60 seconds of waiting. Usually, another attempt to manually start Structure would succeed.

With the new start-up sequence, the probability of such behavior should be less. Unfortunately, it's hard to tell precisely, because there's a dependency on JIRA subsystem. However, the plugin start-up has become much more predictable and repeatable.

5.4. Performance Improvements

Most of the performance improvements mentioned in the Release Notes above are quite important for large-scale instances.

5.5. Testing on Staging Environment

The following checks are suggested for Structure 3.4 on a staging environment with production data.

1. Start-up and shutdown.
    a. Try installing and uninstalling Structure several times. Watch the logs. On JDC, watch logs on every node.
    b. Try starting and stopping JIRA when Structure is installed. On JDC, try stopping/starting individual nodes and the whole cluster.
2. Index recovery on JIRA Data Center. Try forcing index recovery by manually delaying your node indexes. Before that, install synchronizers that are known to run for a long time – for example, a links synchronizer on all of your issues, assuming you have a lot of links in JIRA.
3. Automation load testing.
    a. 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.
    b. 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.
    c. Watch log files for errors and warnings.
4. Automation stress testing.
    a. Emulate peak number of users opening the most popular structure.
5. If you're using JIRA Software, try creating an Agile board that is based on `structure()` function (see S-JQL Cookbook). Make sure there is about 1,000 issues in the result. See how well Agile board loads.
6. Check **Administration | Structure | Maintenance** page. Try to run Structure backup and measure how much time does it take. Watch for errors in the logs.

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