Structure 4.2 Release Notes

27th of August, 2017

Structure 4.2 adds major Formula language improvements, custom reporting period for Work Logged column, and a way to see large values that don't fit the column width. Other updates and improvements are also included.

Download the latest Structure and Extensions
Structure Demo Server

1. Version Highlights

- Formula Column language improvements.
- Filtering work logs by a custom period in the Work Logged column.
- Large texts or other values that don't fit the column width can now be displayed fully.

2. Changes in Detail

2.1. Formula Language Improvements

There are a number of Expr language improvements delivered with this version.

First, we're adding **aggregate functions**, which are able to calculate a total value over sub-issues. In this version, we include only `SUM{}` aggregate function, but more aggregate functions will follow soon. But already with `SUM{}` only you can calculate a number of interesting metrics, such as WSJF or percentage of bugs in a structure. Here's a simple example of a formula that gives the number of bugs: `SUM{ IF (type=Bug;1) }`

Secondly, we're adding **local variables**, which is really helpful if the formula contains some expression multiple times.

Lastly, now you can write **comments** in a formula, which should make it is easier to read the formula later and make changes.

Remember that with Formula Column, it is possible to define a metric and share it with the team via Views.

Documentation: Expr Language, Aggregate Function Reference

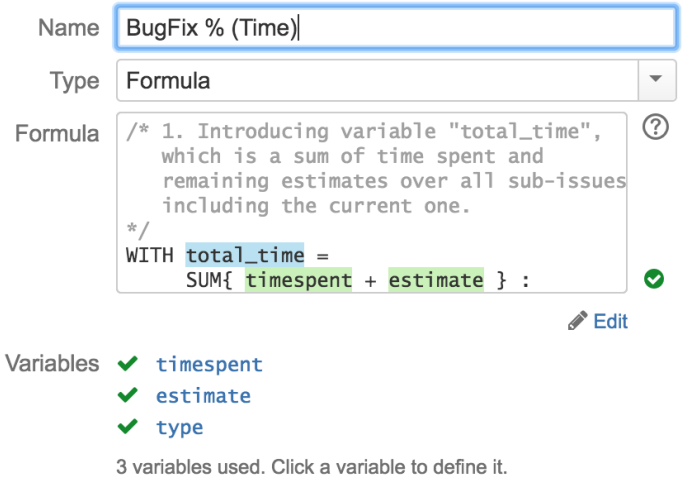

2.2. Custom Period in the Work Logged Column

Work Logged column displays total hours logged for an issue, with some additional filters. One of the filters selects only the work logs filed for a specific time period.

With Structure 4.2, it is possible to set an arbitrary period for Work Logged column (in addition to already available predefined periods like Today or This Year).

This addition option allows you to build task-based timesheet reports.

Documentation: Work Logged Column

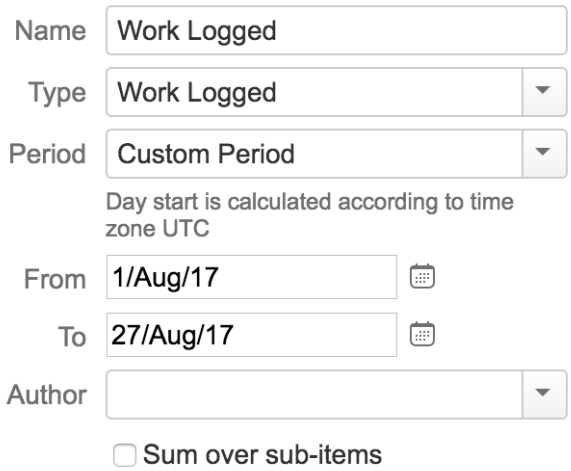

2.3. Full Content Hover Box

When a field value is too large to fit into the cell, only a part of it is shown by Structure.

With version 4.2, it is possible to view the full content by hovering mouse pointer over the "more" sign (three vertical dots). This sign appears near the right edge of a cell if there's more to show.

This can come in handy if you do not want to click every issue to view their description in the issue details panel.

Documentation: Displaying Full Cell Content

2.4. Notable fixes and improvements

- The process of restoring Structure from backup now displays a progress bar and can be cancelled.
- Fixed: Changing filter text in the panel header does not result in change in filter results.
- Fixed: License installed on one node is not applied on another node in the cluster (JIRA Data Center).
- Fixed: Count leaves column has incorrect values when exported to Excel.
- Fixed: Values in Formula Cells with Duration format are missing when exported.
- Fixed: Potential data corruption during backup because of runtime exception.
- Fixed: Sub-task synchronizer doesn't add freshly created sub-tasks.
- Fixed: Time Tracking panel disappears when switching between issues on Internet Explorer.

The version also includes performance, reliability and other internal improvements.

3. Supported Versions

Structure 4.2 and all extensions support JIRA versions from 7.2 to 7.4.x. All editions of JIRA (JIRA Core, JIRA Software, JIRA Service Desk) are supported. JIRA Data Center is supported.

⚠️ Structure.Pages 1.2 or earlier are not fully compatible with Structure 4.2. If you're using Structure.Pages, make sure to also upgrade it to version 1.3 or later.

With respect to other add-ons and custom integrations, this release is backwards-compatible with Structure 3.4–4.1. Structure.Testy extension, Colors plugin, integrations with our partner add-ons 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. ⚠️ If you have Structure.Pages installed, make sure you've upgraded to version 1.3 or later, both on JIRA and on Confluence side.
4. 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.

Upgrade procedure from versions 3.x-4.1 is simple:

1. Consider backing up JIRA data. Use **Administration | System | Backup System**. (If you have a large instance and have proper backup strategy in place, you may skip this step.)
2. Back up Structure data. Use **Administration | Structure | Backup Structure** menu item. If you have a lot of structures and a large JIRA, consider turning off "**Backup History**" option to avoid long backup process.
3. Install the new version of the plugin.
4. Monitor `catalina.out` or `jira-application.log` for warnings or errors.

5. Enterprise Deployment Notes

Structure 4.2 has a number of changes that are particularly important for large installations and JIRA Data Center instances.

5.1. Events, Re-indexing and Experimental Feature

Over last few months, we've been working on tackling a race condition that may happen on a clustered JIRA. In short, Structure has its own means of inter-node communications, which is used to notify about issue changes. These notifications are important because changes made on one node may cause synchronizers or automation rules to be executed on another node. These rules or synchronizers often access JIRA Lucene index, to run a query, for example.

The problem is a race between our notification subsystem and JIRA's own notifications between nodes, which cause issue to be re-indexed on all nodes. If JIRA re-indexing is too late and our recalculation happens first, the users working with that node may see outdated or inconsistent data.

The solution to this problem required a lot of changes to Structure's event subsystem. Since it's a core component right in the middle of Structure's architecture, we approached its rollout carefully. Currently, the new event distribution approach for Data Center is an experimental "dark feature", which means that you need to turn it on explicitly, if needed.

To activate the updated "index monitoring" based event distribution subsystem on JIRA Data Center:

1. Verify that you need it. The cause for activating it could be the users complaining about not seeing the effects of actions of other users, or if your testing reveals this problem.
2. Set "`structure.delegatingItemTracker.enableReindexMonitor`" system property to "`true`".
3. Disable and then enable Structure. (Make sure there's enough time for disabling to propagate through all nodes.)
4. Verify: a new, undeletable system custom field called "**Structure Index Monitor**" should appear. This field will not have effect on issues, and it will disappear on its own if you turn this dark feature off.
5. Note that JIRA will tell you that a Full Re-index is required (since a new field added). Feel free to ignore this request, unless you also have other reasons for re-indexing.

In one of the upcoming versions of Structure, we'll enable this dark feature by default.

Takeaways:

- Current versions of Structure, including 4.2, in certain cases may show outdated data to some users on JIRA Data Center.
- The solution is currently an experimental feature, which can be enabled through a system property.
- There's an automatically created system custom field if this feature is enabled. It should not interfere with JIRA configuration.
- Even if the experimental feature is not enabled, there have been certain changes in a core subsystem in Structure.

5.2. Cancelling Structure Restore

Structure restore may take a long time, if there's a lot of Structure data (especially, history records). It may take an hour or more to restore Structure from backup, unless you decide to skip restoring history data.

With Structure 4.2, you can cancel the restore process if it takes too long. However, please be aware that after that you'll have some of the Structure data restored and some of it missing. You may need to communicate that fact to the users or adjust the access to Structure.

As an alternative, you can try Merge Structure operations, which works like Structure Restore but allows you to pick which structures and other data to restore.

5.3. Formulas and Attribute Subsystem

Attribute subsystem in Structure is responsible for delivering column values to the grid in Structure Widget, and for a few other things. In Structure 4.2 there were a few internal improvements in the Attributes subsystem, which may affect performance and functionality. (Actually, some of the changes should increase performance.)

One of the big features that we're working on over last several releases is Formula Column. This is an important addition to Structure that brings capability to define and run arbitrary metrics over structures of issues. Structure 4.2 introduces important additional functionality to formulas.

The expansion of formula feature may invite users to create new automated structures and running a formula with aggregation over those structures. On large JIRA instances this may lead to increased server load, as formulas with aggregation may need to scan a lot of issues to calculate a total.

Previously, we've added a hard limit on a structure's size (100,000 rows – can be adjusted). Calculating a total over 100,000 issues would require JIRA to retrieve that many issues from the database, which is not going to be quick and would add a load on JIRA and the database. It's a good idea to advise the users to calculate metrics on such large structures only if there's a real need. In most cases, structures can be made smaller. When in doubt, feel free to reach out to support@almworks.com.

5.4. Testing on Staging Environment

You can try load testing and stress testing Structure 4.2 and JIRA on a staging environment before upgrading.

It makes sense to test the following parts:

- Open large structures and add multiple predefined columns from Calculated section. (See the description above about Formula Column feature.) Do that from a number of client computers using different user accounts.
- Open the same structure in different browsers, having logged in under different users. Make sure the users go to different nodes in cluster. (Use direct node's addresses if necessary.) Make updates to a test structure and test issues in one browser and watch if changes are picked up (without refreshing the page) on another browser. (You might need to switch to the other browser to see the changes.)

The usual load and stress testing can also be applied.

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