

Structure 5.1 Release Notes

25th of October, 2018

Structure 5.1 adds wiki markup, an admin interface for dark features and several performance improvements and fixes.

Download the latest version of Structure and its Extensions
Try It: [Structure Sandbox Server](#) (no installation required)

1. Version Highlights

- Formula columns now support wiki markup
- New admin interface for enabling/disabling Structure dark features
- Inactive users are now flagged when grouping by Assignee
- Several performance improvements and fixes

Data Center customers: Due to a problem in the Atlassian Universal Plugin Manager, the previous Data Center approved version of Structure, Structure 5.0.1, was reported as incompatible. This problem does not occur with Structure 5.1.

2. Changes in Detail

2.1. Wiki Markup

Wiki markup is now supported within [Formula columns](#).

Using wiki markup, you can:

- Specify the text color within a column
- Highlight cells with background coloring
- Insert images
- Add emoticons

SAFe Structure

Due Date*

Key	Summary	Σ Story Points	Assignee	Task Status	TP	
SPR-5	SAFe Epic 5	42	M. Reynolds	Due Soon		
STMA-16	Team A Story 16	15	Unassigned	Due Soon		
STMA-4	Team A Story 4	15	C. Bacca (Inactive)	On Track!		
STMA-5	Team A Story 5	12	C. Bacca (Inactive)	On Track!		
SPR-4	SAFe Epic 4	22	M. Reynolds	Due Soon		
STMA-7	Team A Story 7	9	Jack Brown	OVERDUE		
STMA-14	Team A Story 14	13	Unassigned	OVERDUE		

Markup content can be included when exporting a structure to Excel or a printable page.

To learn more, see [Wiki Markup in Formula Columns](#).

2.2. Inactive User Flag

When grouping by Assignee (either through automation or transformation), the "Inactive" identifier is now visible beside the group name:

SAFe Structure						Basic view	
Key	Summary	Σ Story Points	Assignee	Pri	TP		
rlitch	▶ Bob	228		<div></div>			
demo	▶ Demo User	15		<div></div>			
admin	▶ M. Reynolds	159		<div></div>			
mary	▶ Mary (Inactive)	8		<div></div>			
✓ STMB-23	Sub-task 9	7	Mary (Inactive)	<div></div>			
✓ STMB-22	Sub-task 8	1	Mary (Inactive)	<div></div>			
nah	▶ Nah Duo	3		<div></div>			
✓ STMB-19	Sub-task 6	3	Nah Duo	<div></div>			

Showing 73 items 3 items with duplicates Info

2.3. Admin Interface for Dark Features

It is now possible to easily enable or disable dark features directly from a new Structure dark features interface.

Administration Search Jira admin

Applications

Projects

Issues

Add-ons

User management

Latest upgrade report

System

Structure

STRUCTURE ADMINISTRATION

Configuration

Defaults

Backup Structure

Restore Structure

Migrate Structure

Maintenance

License Details

Support

Setup Guide

Back to project: SAFe Team B

Structure Dark Features and Fine Tuning

⚠ Please be careful! Incorrect value of a property may lead to unpredictable results. Make sure you follow the instructions from support.

Property Key

Add Property

Key	Value	
com.almworks.jira.colors.dataVersion	1	
com.almworks.jira.structure.system.refreshedOnStart	false	

To access the interface, you must have Jira Administration permissions and enter the interface location directly into your browser: https://YOUR_JIRA_ADDRESS/secure/admin/StructureDarkFeatures.jspa

Notes

- Properties added with our admin interface are added to Application properties
- A Structure property gets its value from Application properties; if no value is found, the property gets its value from System properties

For more information, see [Advanced Configuration](#).

2.4. Notable Fixes and Improvements

- Group, project and role lookup is now available when editing structure permissions
- The #distinct modifier has been added to the JOIN function in formulas, to prevent values from being added multiple times
- Shared perspectives can be viewed with anonymous access - issues that require specific permissions will not be shown in the structure
- When adding generators, the "allow changes" option can be unchecked by default
- Fixed: Exporting to Excel issue when exporting a perspective created by Structure 4.x
- Fixed: Unable to move issues under a folder when using Rank sort generator
- Fixed: Clicking column header does not sort structure when user has view permission
- Fixed: Reordering sub-tasks on the issue page does not reorder issues in the structure
- Several additional fixes and performance improvements

3. Supported Versions

Structure 5.1 and all extensions support Jira versions from 7.2 or later. All editions of Jira (Jira Core, Jira Software, Jira Service Desk) are supported. Jira Data Center is supported.


With respect to other add-ons and custom integrations, this release is backwards-compatible with Structure 3.4–5.0. [Structure.Testy](#) extension, [Colors](#), [Structure.Pages](#), [Structure.Gantt](#) and integrations with third-party apps should continue working normally.

 Structure 5.1 will be the last version to support Jira 7.2-7.5. Structure 5.2 will support Jira 7.6+.

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 app, either from the Atlassian Marketplace or our [Download](#) page.
2. When Add-on Manager reports the 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 the "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 the Confluence side. If your Confluence version is not compatible with Structure Helper 1.3, you should stay with version 1.2 for Structure.Pages and Structure Helper apps, but please note that there are limitations to its compatibility with Structure 4.2 and higher, so a Confluence upgrade to version 6.1 or later is recommended.
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](#).

The upgrade procedure from versions 3.x–5.0 is simple:

1. Consider backing up Jira data. Use **Administration | System | Backup System**. (If you have a large instance and a 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 the **"Backup History"** option to avoid a long backup process.
3. Install the new version of the plugin.
4. Monitor `catalina.out` or `jira-application.log` for warnings or errors.

 We strongly recommend that you back up your data before upgrading. The introduction of manual adjustments required changes to our backup file format, which makes previous versions of Structure unable to restore data from backup files created by Structure 5.0 and later. For more information, see [Backup Format Version Change](#).

5. Enterprise Deployment Notes

Structure 5.1 has a number of changes that are important for large-scale Jira Server and Jira Data Center instances.

5.1. Structure Index Monitor Field

DATA CENTER

"Structure Index Monitor" is a system custom field that is automatically installed on Jira Data Center instances since Structure 4.6. It helps Structure capture events related to Lucene-based index subsystem. This is what it looks like on the Jira administrator's Custom Fields page:

Structure Index Monitor LOCKED System custom field, required for correct functioning of Structure add-on. Does not store any data.	Structure Index Monitor	Not configured for any context	
---	-------------------------	--------------------------------	---

This custom field does not store any data in the index, so it does not have a performance impact. Structure Index Monitor field cannot be removed, but it will disappear if Structure is uninstalled and will not affect Jira at all.

One application of this field is ensuring correct cache operation in Structure. When an issue is updated on one node, the indexes on all nodes must be updated to make sure JQL queries return the correct result everywhere. Structure caches also need to be checked on all nodes, and potentially a JQL query should be re-run to make sure that generated content is up to date. The cache update should not happen earlier than the index update, otherwise a JQL query would return an outdated result.

With Structure 5.1, we use the Structure Index Monitor field to help with the timing of cross-node updates. It eliminates a chance of a race condition that would cause Structure to show outdated content on a particular node for a long time after a relevant update has happened.

To test the rollout of this change, you can observe the behavior of the system under a constant stream of issue updates (one change per second, for example), happening on one node, while a user observes a structure with JQL-based automation on another node. At some point the stream of changes should stop, and a few seconds later the user should see the most up-to-date information in Structure.

This feature has been experimental for a while, and we're enabling it now as the default. It is still possible to turn it off and go back to the way events were propagated between nodes in Structure 5.0 and earlier, by using the dark features interface and setting `structure.delegatingItemTracker.enableReindexMonitor` to `false`.

5.2. Offloading Change Stream Writing to a Background Thread

DATA CENTER

EXPERIMENTAL

In the Data Center environment, Structure running on one node needs to let Structure running on other nodes know when an item (an issue or some other object) changes. This "change stream" is communicated to other nodes via the database, asynchronous caches and occasional use of a global, one-per-cluster lock.

Normally, each change is written into the database immediately when it happens – in the "event listener". That code runs in the same execution thread as the change itself, typically as a response to a user's action. We recently worked on an incident where the Jira global locking subsystem failed and there were certain issues with the database. That made writing of the change stream "hang", which, in turn, made user request threads hang, which led to Jira being unresponsive.

In Structure 5.1, we're introducing an alternative implementation of this subsystem, which would never block a user request thread. All global lock operations and writing to the database will happen in a separate, dedicated thread of execution.

This feature is currently **not** enabled by default. For now we recommend to turn it on if a similar problem has happened in the past, or if there's a problem with Jira hanging and you suspect it might be this case.

To enable this feature, set the `structure.delegatingItemTracker.enableAsyncEventsQueue` property to `true` in the dark features interface. For more information, see [Advanced Configuration](#).

To test the change, use the same method as described in the previous section.

5.3. Lucene and Jira Search Integration

In preparations for the upcoming Jira 8 release, we have adjusted how Structure works with Lucene index and Jira search. This change should not affect performance or the users in any way. To test that everything is okay, one could run a JQL query from Structure | Query board.

5.4. Bulk Loading Rows from the Database

One of the database tables that Structure most frequently uses is `AO_8BAD1B_ROW`. Structures contain rows, and this table stores the mapping between unique numeric "row IDs" and arbitrary "item IDs" of the objects that structures contain. This table is used very often when a structure is displayed.

With Structure 5.1, we implemented a bulk operation that lets Structure get multiple row information with one database request. That greatly speeds up loading structures that have lots of manually added issues and reduces the database load.

To test this improvement, you can try creating a temporary structure, adding several thousands of issues there (manually, copying them from a search result), and then opening that structure after an upgrade or after Structure is disabled and enabled.

5.5. Testing on Staging Environment

Apart from the specific suggestions given above, you can run the generic load and stress testing on a staging environment as advised in the previous release notes.



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