

Radish Knowledge Base

A collection of longer articles, howtos and notes relating to Atlassian products and system administration.

- [Confluence](#)
 - [Confluence attachment filesystem paths](#) — Find the filesystem path for Confluence attachments.
 - [Confluence User Macros](#) — Some useful Confluence User Macros.
 - [Allow macro content inside any other macro](#) — This page describes a way of allowing arbitrary Confluence content, including macros, in the bodies of macros that don't normally allow this.
 - [Atlassian Product Releases Macro](#) — A Confluence user macro that displays the latest Atlassian product release versions (with release notes), and release dates.
 - [HTTP Parameter Macros](#)
 - [Editing the Confluence dashboard](#) — How to edit the Confluence dashboard.
 - [What to do when your Jira or Confluence is hacked](#) — Strange running processes? 100% CPU use? Security emails? Odds are your server has been hacked. This guide is for system administrators who are not security experts, but who nevertheless need to recover from a hacked Jira/Confluence installation.
- [General](#)
 - [Dropbox to Amazon S3 Syncer](#) — Over on Github https://github.com/redradishtech/dropbox_to_s3_syncer I've published a Docker-based system for syncing a Dropbox folder with a S3 bucket. Drop a file in a designated Dropbox subdirectory, and it is automatically uploaded to S3.
 - [Dropbox Installation error: Trouble connecting to Dropbox servers](#)
 - [-Duser.timezone considered harmful, and why your Atlassian server shouldn't be on UTC](#)
 - [Parsing JIRA JSON dates](#) — strptime code for parsing dates as found in JIRA REST API results.
 - [PostgreSQL](#)
 - [Encoding HTML in PostgreSQL](#) — this function is the best way to encode HTML in PostgreSQL
 - [Temporarily Switching Users in Jira/Confluence](#) — A quick JSP allowing administrators to temporarily take on the identity of a user. Versions for JIRA and Confluence.
 - [Using database diffs to see what JIRA is doing](#) — This is a small but very, very useful tip. Often there's situations where you need to know what happened in the database, for a particular JIRA operation (or any database-using webapp). This is how you do it.
 - [Versioning your JIRA and Confluence app directories](#) — I recently published https://hg.sr.ht/~redradishtech/versioned_directories, a script to manage the seemingly trivial task of keeping a `/opt/atlassian/jira/current` symlink pointing to the correct `/opt/atlassian/jira/$version`; or more generally, managing a directory containing N versions of something.
- [Goodbye Server](#)
 - [Atlassian Future Pricing Calculations worked example: 50 user Jira and Confluence](#) — This is the essence of a report I wrote up for one of my customers, with recommendations on what to do about Atlassian's move to Cloud.
 - [Atlassian Future Pricing Calculations worked example: 2000 user Jira and Confluence](#) — This is the essence of a report I wrote up for one of my customers, with recommendations on what to do about Atlassian's move to Cloud.
- [JIRA](#)
 - [Creating interactive Jira reports in Confluence using free tools](#) — This is a tutorial on how to use Confluence as a query / reporting engine, querying SQL data sources like the Jira database. For our example we query JIRA's database to build a Monthly Worklogs Report, showing hours worked per day for every user in a given month. We use the free Play SQL Base <https://marketplace.atlassian.com/apps/1211703/play-sql-base?hosting=server&tab=overview> plugin.
 - [Crowdsourcing triage: A useful pattern for Jira issue ownership](#) — When a new issue is created, who owns it? The traditional answer is 'nobody, until it is claimed'. Here we suggest a better model, where ownership is collective until claimed.
 - [Fixing JIRA graph gadgets broken by JRA-59364](#) — JIRA 7 bug causes most 'Created vs. Resolved Chart' gadgets on dashboards to break with the error:
 - [Groovy Scripts](#) — The ScriptRunner for JIRA <https://marketplace.atlassian.com/plugins/com.onresolve.jira.groovy.groovyrunner/server/overview> plugin is indispensable for any non-trivial JIRA use (and an excellent reason on its own to self-host your own JIRA). This page collects some of the more useful Groovy snippets we've written over the years.
 - [A Jira bot for Google Chat](#) — Send a message to a Google Chat (<https://chat.google.com>) room when Jira is updated.
 - [Automatically deactivating inactive Jira users](#) — To save on licensing costs, it is sometimes useful to automatically deactivate Jira users who haven't logged in within a certain period, say 6 months. Here we review the options, and provide a ScriptRunner script (source in github <https://github.com/redradishtech/jira-user-deactivator-groovy>) that does the job.
 - [Groovy Script - simplified views of fields](#) — How to create a read-only "summary" version of a field, for display in space-constrained places like gadgets or search results.

Recent space activity



Jeff Turner

[Atlassian Future Pricing Calculations worked example: 2000 user Jira and Confluence](#) updated Nov 19, 2020 • [view change](#)

[Atlassian Future Pricing Calculations worked example: 50 user Jira and Confluence](#) updated Nov 19, 2020 • [view change](#)

[Goodbye Server](#) updated Nov 09, 2020 • [view change](#)

[A guide to setting up Postgres-backed OpenLDAP with the back-sql backend](#) updated Oct 23, 2020 • [view change](#)

[JIRA Checklist Plugin Comparison](#) updated Oct 16, 2020 • [view change](#)

- [Reindex issues from specific projects](#) — A ScriptRunner 'canned script' which reindexes only issues in particular projects.
 - [View who last re-ranked the Board](#)
 - [Visualizing a Quarter Commit field](#) — Given a field like 'Q2 2016' (i.e. "second quarter of 2016"), provides a graphical 'timeline' view of upcoming year, with the relevant quarter highlighted.
- [How to write a Python script authenticating with Jira via OAuth](#) — If you are writing a script that interacts with Jira through a REST API, you should authenticate using an OAuth token, rather than an embedded username/password. Here we describe one way to do the 'oauth dance' to generate a trusted token using Python 3 - specifically the `jrashell` utility from the `jira` Python package. `jrashell` then forms a useful basis for a Python script. Our example script uses OAuth to call an undocumented REST API for querying license data.
- [Intermittent JIRA startup failures? Check your old plugins for com.atlassian.upm.plugin-license-storage-plugin](#)
- [JIRA Add-ons](#)
 - [Derby database tinkering with the Structure plugin database](#) — Notes on how to interrogate the JIRA Structure plugin <https://marketplace.atlassian.com/plugins/com.almworks.jira.structure>'s Derby database from the command-line.
 - [Disabled Jira Portfolio? Get your Initiative 'Child issues' structures back](#) — How to restore links from Initiatives to Epics, after disabling Jira Portfolio.
- [Jira administrators beware: new projects are visible by default](#)
- [JIRA-backed LDAP](#) — Notes on my attempts to have OpenLDAP use JIRA itself as its backing database, through the `back-sql` <https://linux.die.net/man/5/slapd-sql> OpenLDAP backend. This allows LDAP-enabled applications to reuse JIRA user credentials.
 - [A guide to setting up Postgres-backed OpenLDAP with the back-sql backend](#) — OpenLDAP is typically backed by a LMDB database (`mdb`). However it's backing datastore is actually pluggable. One of the most intriguing backends is `back-sql` <https://www.openldap.org/doc/admin24/backends.html#SQL>, which allows LDAP data to be fetched from an ODBC (SQL) datasources. This lets us add a LDAP facade on top of database-backed applications, like JIRA and Confluence.
 - [Migrating JIRA users to LDAP, preserving passwords](#) — JIRA gives you the choice of storing user records internally, or delegating to an external 'User Directory' like Active Directory, LDAP or Atlassian Crowd.
- [JIRA Number fields, without the commas](#) — Remove unwanted commas from a JIRA Number field
- [JIRA REST-based Reporting Scripts](#) — Scripts (mostly written in Ruby) generating reports from a remote JIRA, via JIRA's REST API.
 - [JIRA Report: overdue/on-time issue counts per priority](#) — Given a JIRA project name, a start date and an end date, find total counts of issues completed before, on or after the due date, per priority.
 - [JIRA Report: time worked per user, per sprint](#) — Given a sprint, print the total worklog hours logged by each user on the sprint's issues.
- [JIRA Sprint Capacity Planner](#) — Planning a sprint in JIRA Agile involves assigning issues. You want to spread the load fairly, not overburdening or underburdening workers. To complicate matters, some issues take longer than others (with more Story Points or greater Time Estimate), and the capacity of each worker varies per week, as they may have time off or other responsibilities.
- [Jira tinkerers beware: Tomcat does not serve raw files correctly, causing browser hangs](#) — This page is for anyone who has been tempted to add files (e.g. `js`, `css`) into JIRA's `atlassian-jira/` directory. It may come as a surprise, but JIRA suffers from the most basic of bugs: it doesn't serve static files correctly. More specifically, if Tomcat serves the file with gzip compression, the `Content-Length`: header is incorrectly set.
- [Preventing customer information leaks from Jira filter/dashboard sharing](#) — Jira allows sharing out-the-box, but with one flaw: when you're trying to completely partition customers, Jira's sharing tends to leak customer names and filter/dashboard names.
- [Recovering from an LDAP expired certificate](#) — Inspired by a real incident, this page suggests workarounds for an expired Active Directory LDAP certificate. The `socket` proxy technique at the end is, in the author's opinion, particularly neat.
- [Removing the 'None' option from JIRA select lists, or why radio buttons may be better](#) — Sometimes 'None' is a misleading word in select lists.
- [Reordering User Directories in JIRA and Confluence](#) — When LDAP breaks, sometimes you need to reorder the User Directories to be able to log in at all.
- [Replace your Jira/Confluence mail queue with Postfix -- part 1](#) — In this article we discuss the deficiencies of Jira's outgoing email handling, and how you can avoid problems by configure Jira to delegate mail sending to a local Postfix server (part 2). While we use Jira for concreteness, everything said here applies equally to Confluence.
- [Replace your Jira/Confluence mail queue with Postfix -- part 2](#) — Following on from part 1, this article discusses how to configure Postfix to act as Jira's mail transfer agent. We discuss choice of mail relay, provide relayed and non-relayed Postfix sample configurations, and discuss how Postfix can be configured on reroute email on `sandbox/staging` servers.
- [Resetting a user password in the database](#) — This page has instructions for doing a database-level user password reset in JIRA and Confluence (and other Crowd-library-using applications).
- [Testing LDAP connectivity with `ldapsearch`](#) — JIRA and Confluence can do authentication against LDAP (e.g. Active Directory), using the standard Java JNDI

library. When you're having LDAP connectivity problems, the `ldapsearch` command can sometimes be useful as a means of verifying your LDAP parameters.

- [Notes on Technical Writing](#)
- [Plugin Reviews](#)
 - [JIRA Checklist Plugin Comparison](#) — A comparative review of various todo list / checklist add-ons available for JIRA.
- [Why self-host JIRA and Confluence?](#) — When buying Atlassian products, you have a choice: host JIRA and Confluence on your own hardware, or let Atlassian do the hosting: