



# Git Branching for Continuous Delivery

Sarah Goff-Dupont

*Automation Enthusiast*

Hello everyone! I'll be talking about how teams at Atlassian use Git branches for continuous delivery.

My name is Sarah, and I'm on the product marketing team. Before Atlassian, I spent almost 10 years in QA, including a few years as an automated test engineer. I don't do much coding now, which is ok because I still get paid to learn, think and talk about CI, CD, automation and agile development.

# Agenda

- ① A branch for every issue
- ② CI in the mix
- ③ Multiple-repo projects
- ④ From process to culture

I'll walk you through the branch-and-merge development workflow and talk about why we use it. Then go over some of the things we've learned about incorporating CI, working with projects that span multiple repos, and end with a few words on process and culture.

# A branch for every issue

We usually think of CD in terms of killer automation, but it starts further back in the dev cycle.



Every team at Atlassian is now using Git, and we are bullish on exploiting the power of branch-and-merge workflows. In fact, we've taken to creating a dev branch for each issue we work on – be it bug, user story, or technical task. And that's helping us deliver faster than ever before.

By keeping work in progress isolated on dev branches, we can keep master in a clean and releasable state. And by developing in small, modular pieces we avoid having lots of tightly coupled code that has to be released all together.

This means our SaaS teams like Marketplace and Bitbucket can release every day if they want to, and our teams that make installed products can continuously deliver to our dogfooding instances and ship to customers every 6 or 7 weeks.



If you haven't worked with Git yet, you've probably heard about it. And for the next few minutes, all you need to know about Git is that branching and merging are really easy.

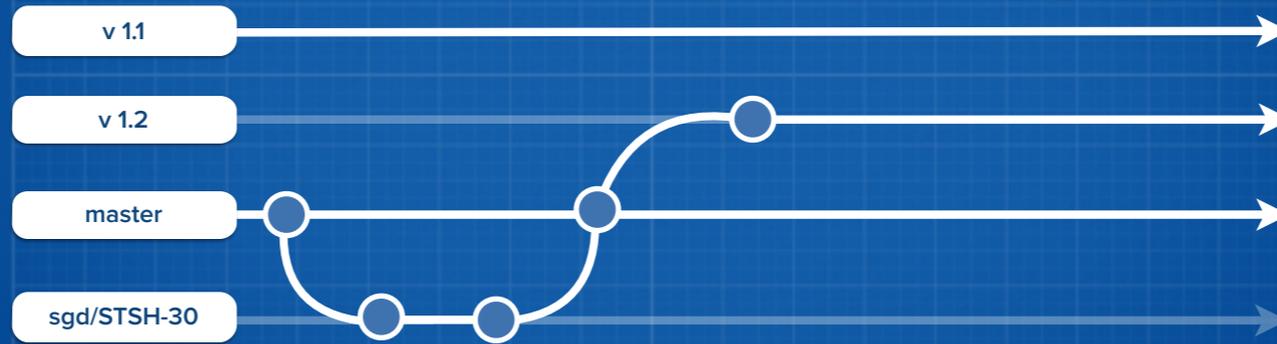
You won't find the day-long code freezes and hellish merges and heavy administrative controls that are associated with centralized VCSs like Subversion.

# ① Atlassian Marketplace Workflow



(call out naming convention)

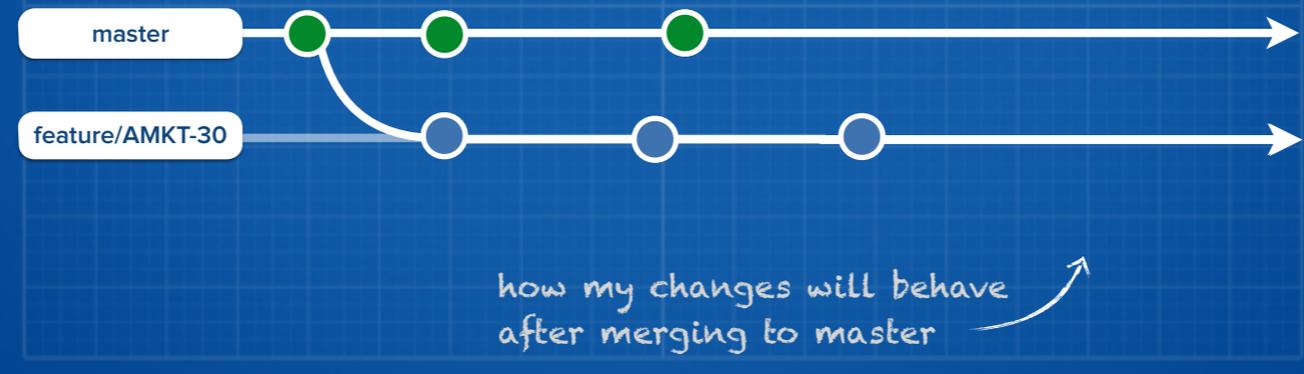
# ② Atlassian Stash Workflow



```
$ git fetch origin  
$ git rebase master
```

When each developer has their own branch, merge conflicts are bound to happen. We've found that doing frequent fetch & rebasing of the dev branch helps avoid surprises during the final merge to master. It also helps catch integration issues faster and lets us get those resolved down on the branch.

# Rebase before merging



Almost all developers do a fetch and rebase when updating their dev branches – not a fetch merge. Rebase is like resetting the commit of origin for your branch, then the new changes on the branch are re-played on top of that. (CLICK)

This provides a more accurate preview of how the new code will work with what's already on master.

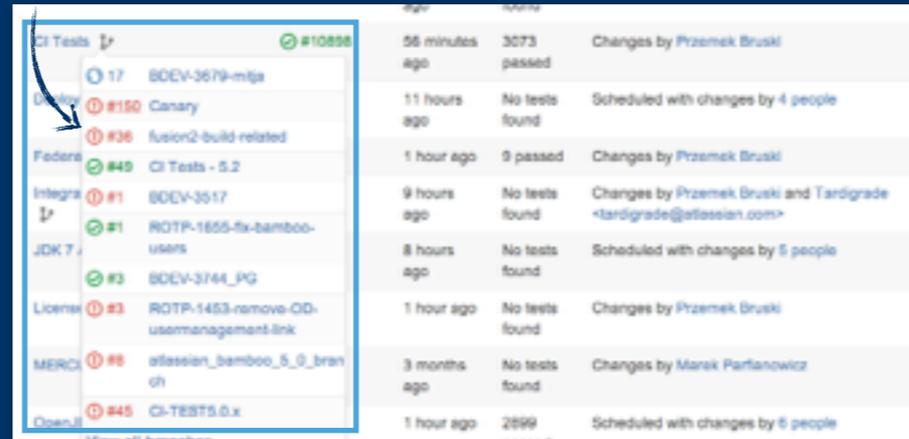
And btw, our teams prefer to use the explicit merge when dev work is merged to master or a stable release branch so the history of what was done on the branch is preserved. With fast-forward merges, you lose that.

# CI in the mix

Now that we have all these branches, how does that effect continuous integration?

# Running CI on dev branches

all active branches are under test



| Branch                                 | Build ID | Status | Age            | Tests          | Changes by   |
|--|----------|--------|----------------|----------------|--|
| CI Tests                               | #10838   | Passed | 56 minutes ago | 3073 passed    | Changes by Przemek Bruski  |
| BDEV-3679-mija                         | 17       | Failed | 11 hours ago   | No tests found | Scheduled with changes by 4 people                                 |
| Canary                                 | #150     | Failed | 11 hours ago   | No tests found | Scheduled with changes by 4 people                                 |
| fusion2-build-related                  | #36      | Failed | 1 hour ago     | 9 passed       | Changes by Przemek Bruski  |
| CI Tests - 5.2                         | #45      | Passed | 1 hour ago     | 9 passed       | Changes by Przemek Bruski  |
| BDEV-3517                              | #1       | Failed | 9 hours ago    | No tests found | Changes by Przemek Bruski and Tardigrade <tardigrade@afession.com> |
| ROTP-1655-fix-bamboo-users             | #1       | Passed | 8 hours ago    | No tests found | Scheduled with changes by 5 people                                 |
| BDEV-3744_PG                           | #3       | Failed | 8 hours ago    | No tests found | Scheduled with changes by 5 people                                 |
| ROTP-1453-remove-OD-usamanagement-link | #3       | Failed | 1 hour ago     | No tests found | Changes by Przemek Bruski  |
| afession_bamboo_5_0_branch             | #8       | Failed | 3 months ago   | No tests found | Changes by Marek Parfanowicz                                       |
| CI-TESTS.0.x                           | #45      | Failed | 1 hour ago     | 2899           | Scheduled with changes by 6 people                                 |

Probably goes without saying, but it's really important to be running CI against dev branches. Otherwise you'll be hit with all kinds of surprises if you merge to master and do your first-ever CI run there - which defeats the whole purpose of this whole branching strategy.



It's pretty painful to clone your CI configs every time you start work on a new issue.

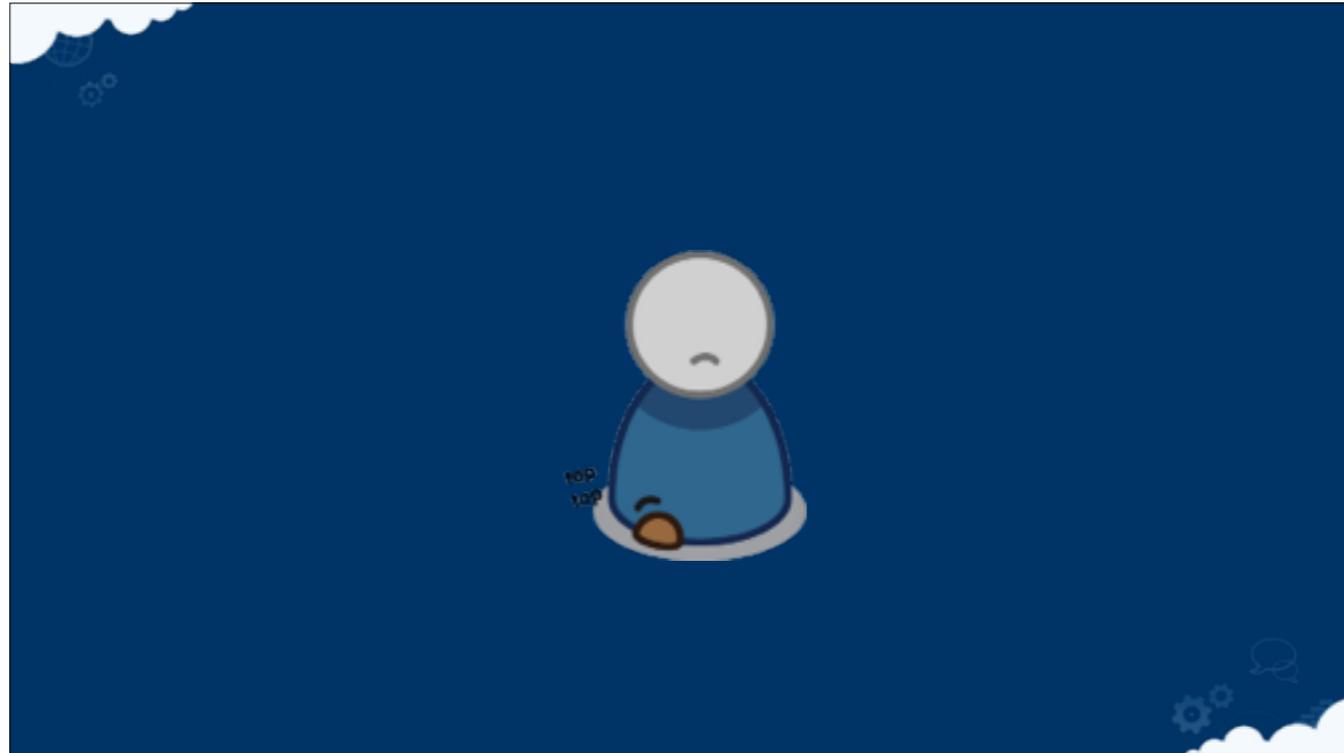
But you can automate that with a Jenkins plugin, or with a Git hook that one of our developers put up on Bitbucket.



But since we happen to make a continuous delivery server, we decided to just build that in and save ourselves (and our customers) the hassle.



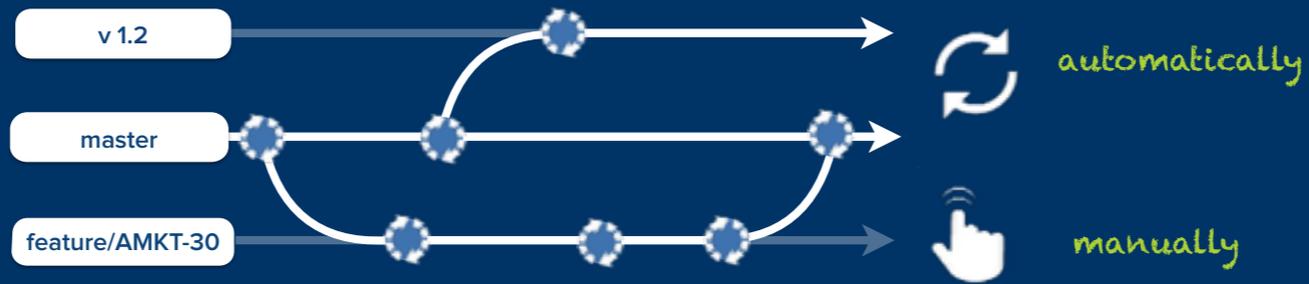
Either way, the point is to spend your time on more interesting work.



But running CI on every branch can use up a TON of resources. How do you keep your build queue from getting clogged up?

Naturally, you could increase your number of build agents (or slaves). But let's think about efficiency first.

# Using build power efficiently



We've found that a good way to balance testing rigor with resource conservation is to have push-button builds on development branches. This is where the most change activity is happening, and where the biggest savings can be had. Our devs find that it fits naturally into their workflow and they like the extra control/flexibility this gives them.

Builds to the most critical branches, however, are triggered automatically with each change. This includes master, and any stable release branches. Most, if not all, the changes pushed to these branches are going to be merges from dev branches, so it's important to get feedback on that quickly - every time.

# Multiple-repo projects



Most of us work with multiple-repo projects, so dependency management will come into play.

```
11 <groupId>com.atlassian.stash</groupId>
12 <artifactId>stash-parent</artifactId>
13 <version>3.0.0-SNAPSHOT</version>
14 <name>Atlassian Stash Parent</name>
15 <packaging>pom</packaging>
16
17 <organization>
18   <name>Atlassian</name>
19   <url>http://www.atlassian.com/</url>
20 </organization>
21
22 <modules>
23   <module>api</module>
24   <module>api</module>
25   <module>modcli</module>
26   <module>dao-api</module>
27   <module>dao-impl</module>
```

We use Maven for this.

Here's a shot of the parent pom for Stash...

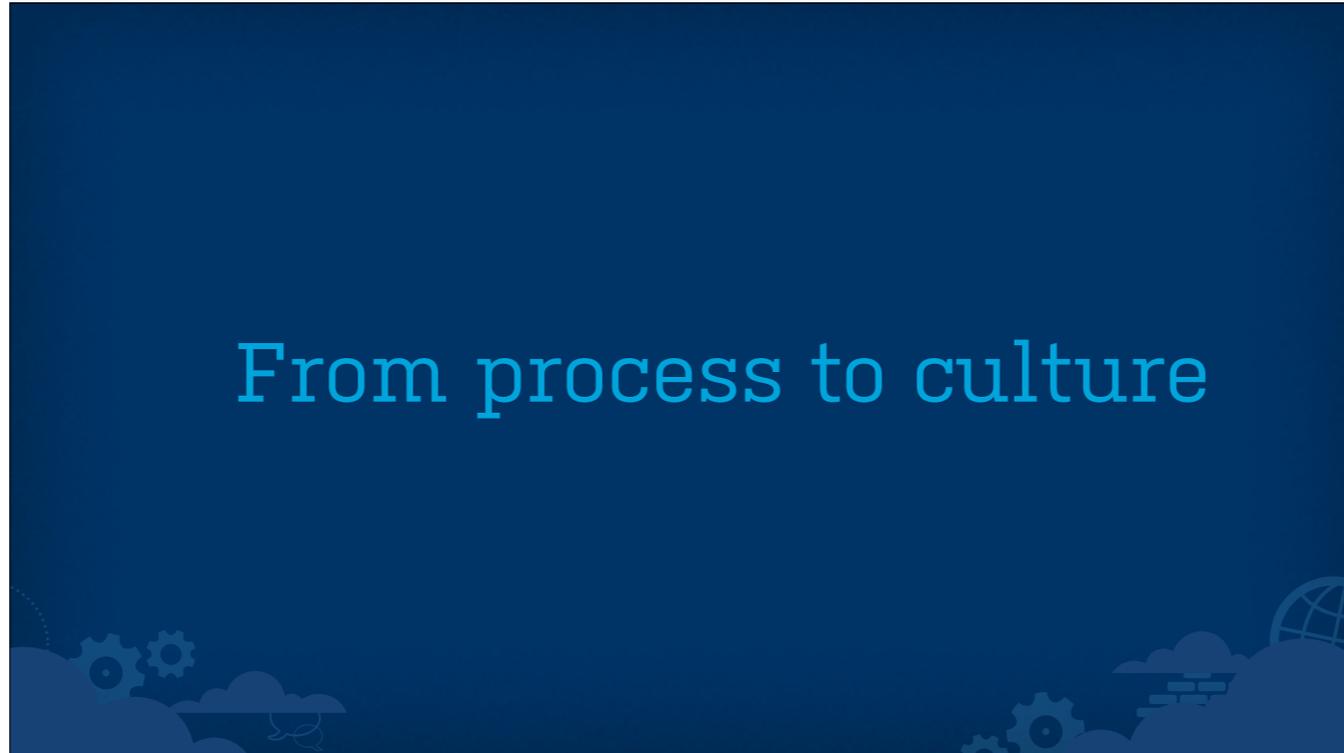
```
742 </dependencyManagement>
743 <dependencies>
744 <!-- List Stash modules first. These artifacts will be used for
745 <dependency>
746 <groupId>com.atlassian.stash</groupId>
747 <artifactId>stash-ee-common</artifactId>
748 <version>1.0.0-SNAPSHOT</version>
749 </dependency>
750 <dependency>
751 <groupId>com.atlassian.stash</groupId>
752 <artifactId>stash-ee</artifactId>
753 <version>1.0.0-SNAPSHOT</version>
754 </dependency>
755 <dependency>
756 <groupId>com.atlassian.stash</groupId>
757 <artifactId>stash-ee-crowd-ee</artifactId>
758 <version>1.0.0-SNAPSHOT</version>
```

And then all the dependent modules.

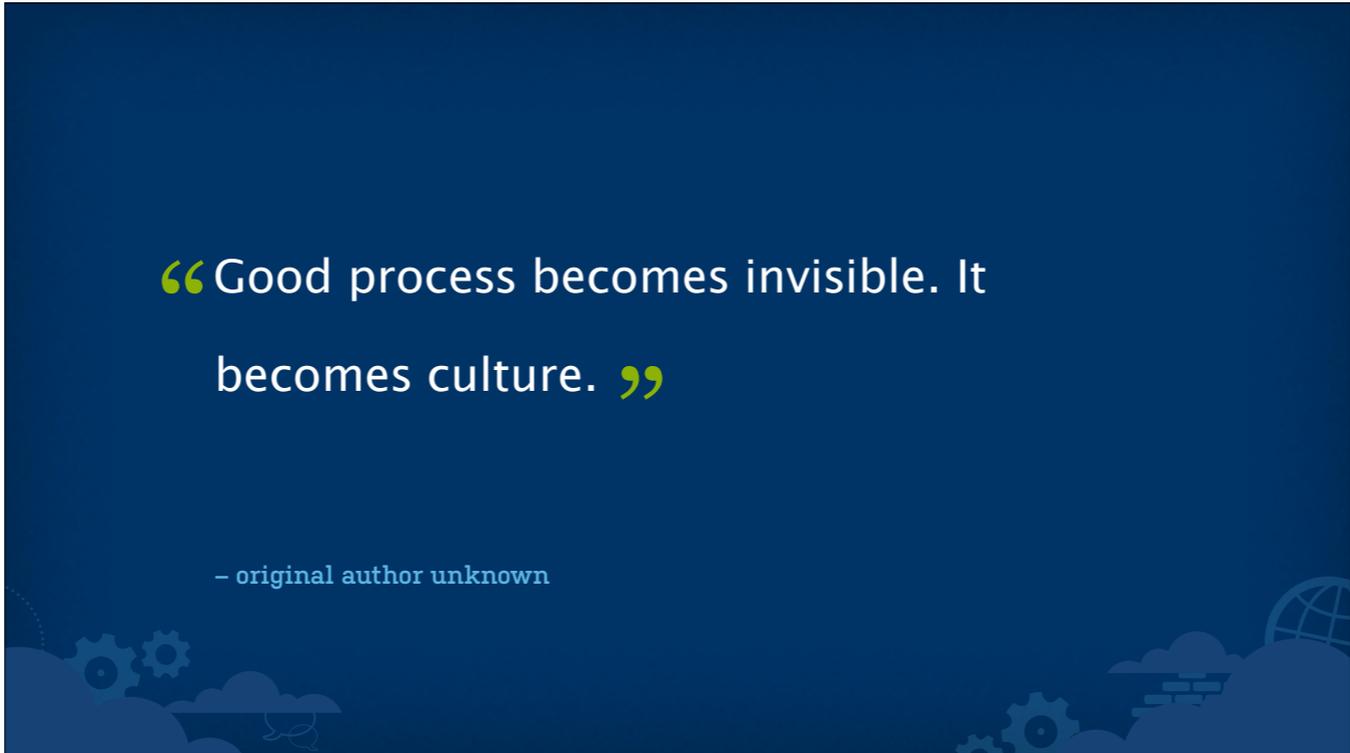
So this has all the pros n' cons of maven.

This isn't a company standard, but as far as I know, none of the teams use Git submodules for dependency management because that would force everything into one build, and that build would take longer than we'd like. Maybe submodules are better suited for dynamic or interpreted languages.

# From process to culture



I want to end with a few words about the branch-per-issue model at a high level.



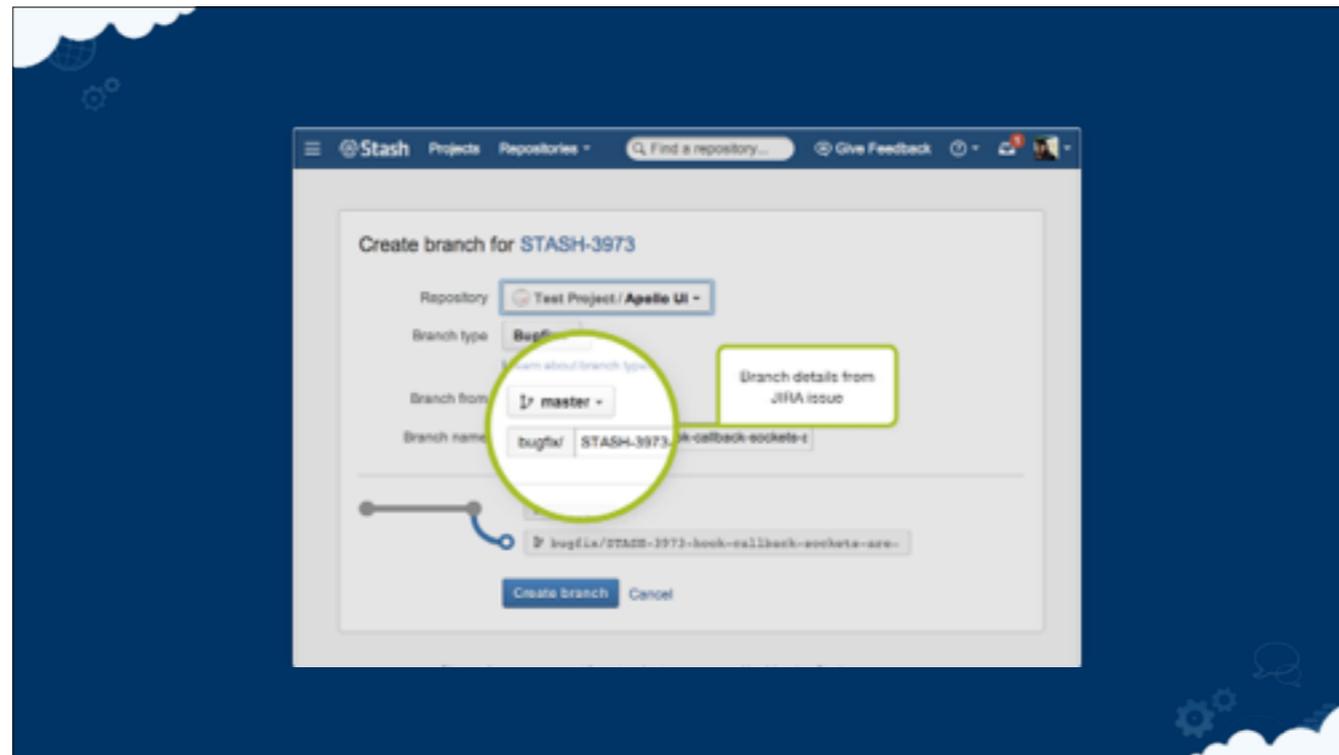
“Good process becomes invisible. It becomes culture.”

– original author unknown

My new favorite “thing someone said about software” is that good process becomes invisible – it becomes culture.

And that perfectly describes what’s happened for us with the branch-per-issue model. For us, doing the dev work on a feature branch has become part of the definition of done.

Including the JIRA issue key in our naming convention and making use of pull requests before merging to master helped us get into the swing of it.



Then after we'd been working like this for a while, we started building features into our own tools like Stash and Bamboo and Bitbucket that make it less tedious for developers to do this.

But we've never enforced it programmatically. You could, in theory. But I don't recommend that because it's inflexible and will end up getting in your way at some point.

If you're feeling like enforcement is necessary, that's a bad smell. It probably means that the process you're about to enforce isn't a good process for your team. And no amount of tooling is going to fix that. The role of tooling is merely to take the busy-work away.

More info at...

① <http://atlassian.com/git>

②  @Atlassian

we're hiring!

