

CASE STUDY

Etsy—Self-Service Developer Deployment: An Example of Continuous Deployment (2014)

Unlike at Facebook, where deployments are managed by release engineers, at Etsy deployments are performed by anyone who wants to perform a deployment, such as Development, Operations, or Infosec. The deployment process at Etsy became so safe and routine that new engineers could perform a production deployment on their first day at work. Etsy board members and even dogs have deployed to production!²⁰

As Noah Sussman, a test architect at Etsy, wrote, “By the time 8am rolls around on a normal business day, 15 or so people and dogs are starting to queue up, all of them expecting to collectively deploy up to 25 changesets before the day is done.”²¹

Engineers who want to deploy their code first go to a chat room, where engineers add themselves to the deploy queue, see the deployment activity in progress, see who else is in the queue, broadcast their activities, and get help from other engineers when they need it. When it’s an engineer’s turn to deploy, they are notified in the chat room.²²

The goal at Etsy has been to make it easy and safe to deploy into production with the fewest number of steps and the least amount of ceremony. Likely before the developer even checks in code, they will run on their workstation all 4,500 unit tests, which takes less than one minute. All calls to external systems, such as databases, have been stubbed out.²³

After they check their changes into trunk in version control, over seven thousand automated trunk tests are instantly run on their continuous integration (CI) servers. Sussman writes, “Through trial-and-error, we’ve settled on about 11 min-

utes as the longest that the automated tests can run during a push. That leaves time to re-run the tests once during a deployment [if someone breaks something and needs to fix it], without going too far past the 20 minute time limit.”²⁴

If all the tests were run sequentially, Sussman states that “the 7,000 trunk tests would take about half an hour to execute. So we split these tests up into subsets, and distribute those onto the 10 machines in our Jenkins [CI] cluster. . . . Splitting up our test suite and running many tests in parallel, gives us the desired 11 minute runtime.”²⁵

The next tests to run are the *smoke tests*, which are system-level tests that run cURL to execute PHPUnit test cases. Following these tests, the functional tests are run, which execute end-to-end GUI-driven tests on a live server—this server is either their QA environment or staging environment (nicknamed “Princess”), which is actually a production server that has been taken out of rotation, ensuring that it exactly matches the production environment.

Once it is an engineer’s turn to deploy, Erik Kastner writes, “you go to Deployinator [an internally developed tool, see [Figure 12.4](#)] and push the button to get it on QA. From there it visits Princess. . . . Then, when it’s ready to go live, you hit the ‘Prod’ button and soon your code is live, and everyone in IRC [chat channel] knows who pushed what code, complete with a link to the diff. For anyone not on IRC, there’s the email that everyone gets with the same information.”²⁶

Deployinator

The screenshot displays the Deployinator console interface. On the left, there are three deployment buttons: 'Deploy to QA (Trunk)' with a 'Message:' input field and a 'Push to QA →' button; 'Princess is in the other castle' with a 'Save the Princess →' button; and 'Deploy to Production' with a 'PROD!!! →' button. On the right, there is a 'Log' section with an 'Add log message:' input field and a 'Log' button. Below the input field is a list of log entries, each showing a timestamp, environment, user, and deployment details.

```
Log  
Add log message:  
[Log]  
• [web] 2010-05-18 22:20:50 | PRODUCTION | sandrews | Production deploy: old 25134, new: 25145 diff  
• [web] 2010-05-18 22:18:00 | PRINCESS | sandrews | Princess Deploy: old: 25144, new: 25145 diff  
• [web] 2010-05-18 22:17:22 | QA | sandrews | kyles bug fix old: , new: 25145 diff  
• [web] 2010-05-18 22:12:03 | QA | sandrews | pushing again -- banned user cache busting old: , new: 25144 diff  
• [web] 2010-05-18 22:08:39 | PRINCESS | sandrews | Princess Deploy: old:25134, new:  
• [web] 2010-05-18 22:02:35 | QA | sandrews | //////////////// old: 25134 , new: 25144 diff  
• [web] 2010-05-18 20:56:50 | PRODUCTION | cmunns | Production deploy: old 25134, new: 25134 diff  
• [web] 2010-05-18 20:49:02 | PRODUCTION | cmunns | Production deploy: old 25134, new: 25134 diff  
• [web] 2010-05-18 20:44:43 | PRODUCTION | cmunns | Production deploy: old 25030, new: 25134 diff  
• [web] 2010-05-18 20:41:17 | PRINCESS | ahashim | Princess Deploy: old: 25030, new: 25134  
• [web] 2010-05-18 20:40:38 | QA | ahashim |Penut butter jelly time!!!!: old: 25030, new: 25134  
• [web] 2010-05-17 15:23:26 | PRODUCTION | sandrews | Production deploy: old 24951, new: 25030 diff  
• [web] 2010-05-17 15:08:05 | PRINCESS | sandrews | Production deploy: old 24951, new: 25030 diff
```

Figure 12.4: The Deployinator Console at Etsy

Source: Erik Kastner, “Quantum of Deployment,” [CodeasCraft.com](http://codeascraft.com), May 20, 2010, <https://codeascraft.com/2010/05/20/quantum-of-deployment/>.

In 2009, the deployment process at Etsy was a cause of stress and fear. By 2011, it had become a routine operation, happening twenty-five to fifty times per day, helping engineers get their code quickly into production, delivering value to their customers.

Automating the majority of the deployment process and including extensive automated testing created an easy