

Case study of Ned Collyer

Who are you?

I'm Ned Collyer. I specialize in web technologies - accessibility, HTML, JavaScript, Java frameworks etc. We build search products and solutions. These are mainly written in Java.

What do you use Hudson for?

We use Hudson for our CI needs. We have projects requiring a large number of configurable modules that need to run on different environments, with around 20 developers.

Previously we were using CruiseControl. We migrated because of the great failure reporting and central management that Hudson provides. Also configuring the lava lamps from a central XML feed just feels nicer, and the improving maven support is something to be excited about (order the module build order correctly - then it will be perfect 🍌)

Build system

We use maven 2. The dependency management is fantastic. If you know how to use maven effectively it saves days of effort setting up build environments, tracking down dependencies (and/or their updates) and versioning internal modules.

Distributed builds

We have slaves for PostgreSQL, MySQL, Oracle running against Linux, and also a VM running Windows 2003 with MSSQL. These are running on a rack of beastly dual quadcore Xeons with a 4GB RAM each. We currently have 3 dedicated machines for this. A full build cycle takes around 40 minutes depending on the settings. The execution of the master node is significantly faster than the slaves so we use it for running Selenium integration tests (through selenese/maven/cargo). This is configured via MAVEN_OPTS on one of the sub modules.

These all poll SVN every 5 minutes - I could tweak it a bit to improve performance, but at the moment I know everything is getting tested when it happens.

Managing slaves

The slaves are autonomous. After I set them up there is no problem keeping them going. If the slave is turned on, Hudson can connect to it. I connect to the slaves via ssh, after copying the public keys from master to slave, there has been no issues.

Unfortunately, many of our modules test against the same database so we cannot take advantage multiple executors, and need to tie each build against a specific slave. The one thing that bugs me is probably not Hudson related. The Windows filesystem and how paths/classpath relate to Cygwin shell is horrible. I can't stand Windows!

Using virtual machines as slaves

Because our servers all run Linux, we use a VM to run our Windows test environment. We also use a VM to run Oracle - it just makes things easier 🍌. Each VM has its own IP so its just like any other physical machine on the network. If you have enough hardware running VMs is great! (eg: If you need to reinstall software, like Oracle, you can have it back up and running within minutes.)