

# ApacheGump

## Apache Gump

[Gump](#) is the regular clean build of everything.

Gump tries to rebuild many open source projects from scratch, checking the files out from CVS and SVN repositories, bootstrapping Ant, then building them using the bootstrapped Ant, NAnt, Maven or make in order of their dependencies. It also runs unit tests. If a build or test fails, the project gets email. Web pages and RSS feeds also track status of projects.

Gump ensures that the HEAD versions of all these projects are sync, by "recreating early, recreating often". Any change that breaks a downstream project is immediately caught and resolved, somehow.

Any OSS project with a public CVS/SVN repository can join the Gump build. Apache committers have access to the Gump repository; other developers can submit their own gump descriptors via bug reports.

In-house projects can host their own version of Gump, integrating private code with public repositories. There are some interesting security implications there -best to run Gump in the DMZ, or an isolated Xen/VMWare machine image which only has net access to the public net and the private SCM servers.

Gump serves as the nightly integration test of Ant. Build files that work on Gump should work on Ant product releases. Trouble that gump does not catch is mostly limited to

- where changes to Ant's classloading break things, because ant-on-gump loads classes differently.
- the batch files to start ant
- builds in directories with spaces, unicode chars.
- installations of Ant with bad classpaths, version conflicts, etc.
- Windows/OSX/Netware-specific problems. While Gump is supposed to be cross-platform the ASF currently only uses instances running on Linux and Solaris.

Also, the quality of build files used in Gump is often higher than many personal projects. OSS projects have to address portability early on, and tend to be less brittle and build files written by and for one person, or for a team where all hardware/OS images are nearly identical.