A Short Introduction to OSG3

Brian Bockelman, OSG Technology Area Lead

For April 2012 GDB

What is OSG3?

- OSG3 is the new version of the software stack from the OSG Consortium.
- It represents major departures from previous OSG release series in:
  - Distribution ideology.
  - Packaging technology.
  - Included software.

(Where did the 3 come from?)

- The last major release series from the OSG was 1.2.
- The OSG also produced a separate distribution, VDT 2.x, that OSG 1.2 was based on.
- OSG3 merges VDT and OSG into one repository.
  - So, we take (max(1,2)+1) = 3.

Distribution Ideology

- Previously, we distributed software in a niche packaging and repository format called "pacman".
- Pacman packages were cross-platform in the same way Java GUIs are cross-platform: mostly works almost anywhere, but feels wrong everywhere.
    - For example, each site must custom integrate them with configuration management tools.
- New ideology: use the OS's packaging tools to give as "native" experience as possible.

Distribution Ideology

- On the RHEL-clone platforms used by WLCG, this means packages in RPM format distributed via yum.
    - Additionally, we adhere to Fedora/EPEL guidelines wherever plausible.
- It also means integration with system tools instead of shipping our own.
    - Limits our choices, but it does mean we no longer bear the weight of packaging Apache, MySQL, Tomcat, Java, Python, etc.
- We ship no package that comes with RHEL.

Package Repositories

- Where possible, we reuse packages found in high-quality community repositories (for RHEL, this is EPEL).
    - However, if we are worried about their quality, stability, or have a valid technical reason (e.g., OSG-specific patches unacceptable to upstream), we will fork the package and ship it ourselves.
- Over the long term, we strive to fork as few packages as possible!

Community Packaging

- The native packaging approach is part of a larger idea of *community packaging*.
  - We want to be net contributors to the larger community.
  - We want to be in position to take advantages of innovations done by the community.
- For example, our Globus build is based on Globus's RPMs (which was, in turn, based on IGE's work). That was not possible for Pacman!
  - We submit patches upstream when possible, but we are otherwise a bit weaker in contributing back.

Packaging Technology

- Previously, VDT packages were built using a custom set of scripts on top of a niche platform at the Wisconsin Build and Test Lab (Batlab).
  - Advantage: huge number of platforms supported.
  - Disadvantage: no one outside Wisconsin knew how to use it. Cannot build RPMs.

Packaging Technology

- We now use a piece of software called Koji.
  - Koji is the build software that RedHat uses to build RHEL and Fedora uses to build Fedora.
  - Combines "mock" for building and a custom database / web interface for repository management.
  - "mash" is used to push final product to the mirror.
  - Results in a highly automated build and release process, with little room for humans to mess things up.
- As of late-April release, all our supported software packages will be available via Koji.
  - The most popular packages (OSG CE, worker node client, and others) were released by last November. Mostly just stragglers.

Koji Web Interface

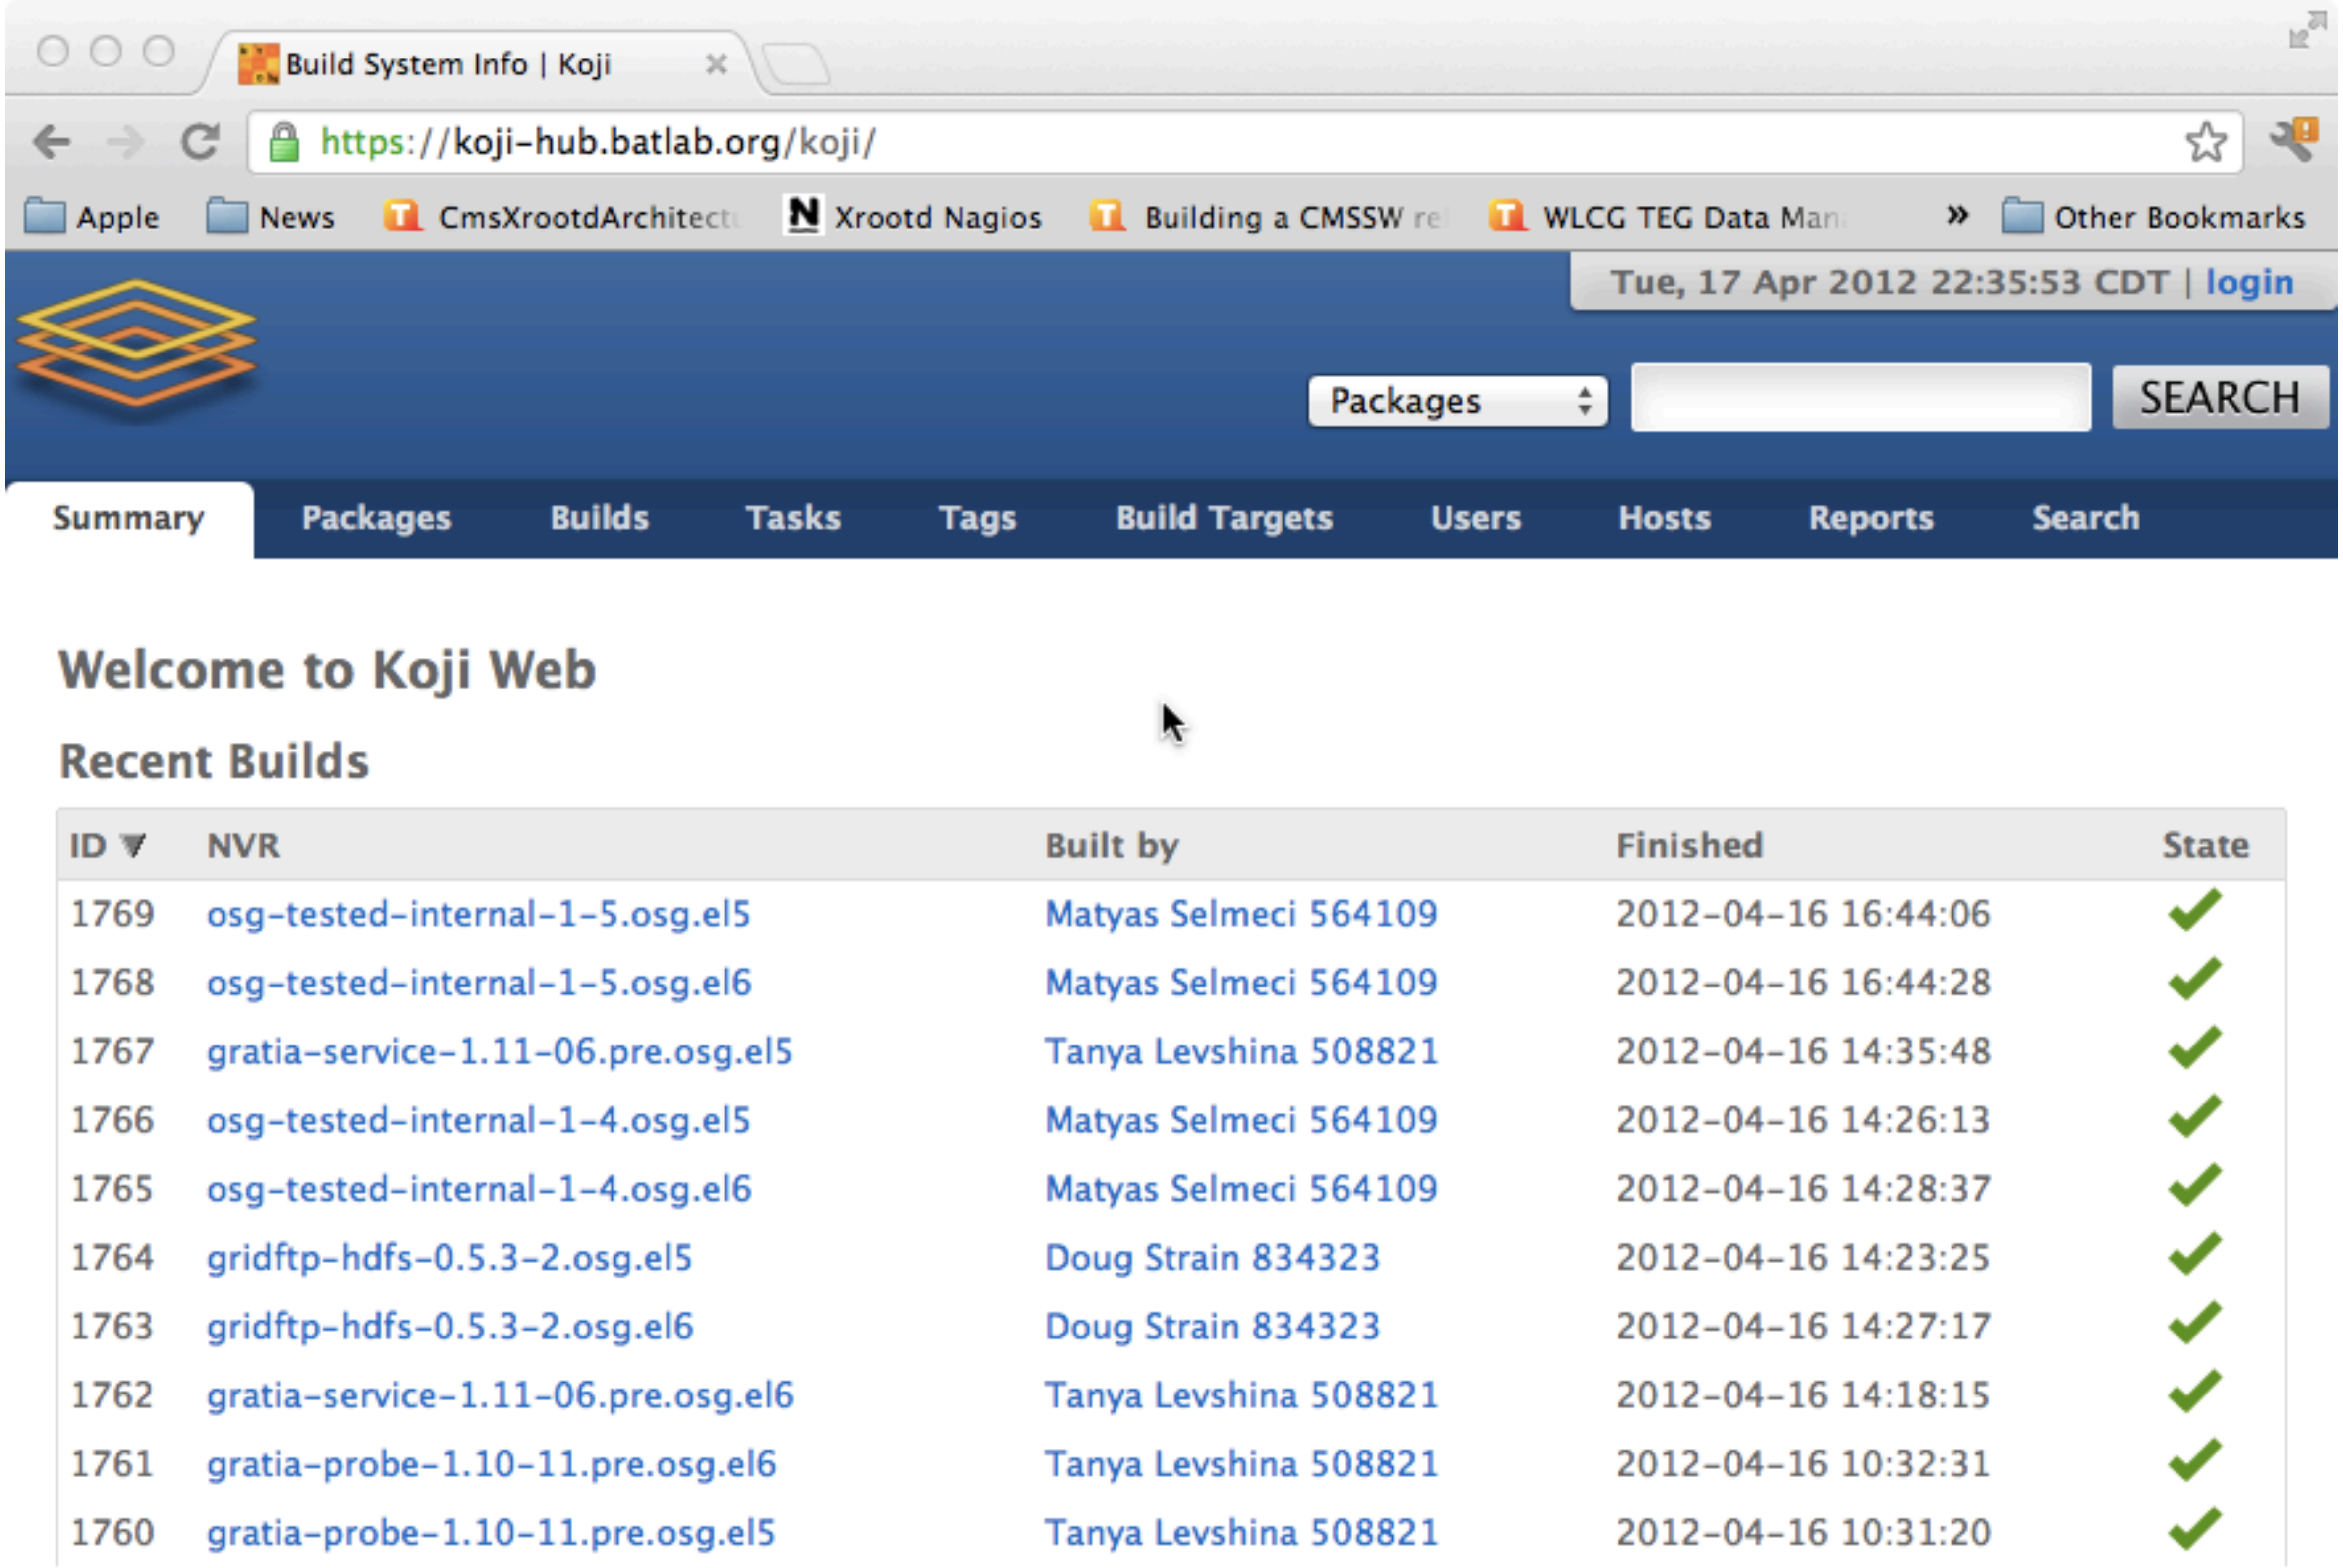

Welcome to Koji Web

Recent Builds

ID ▼	NVR	Built by	Finished	State
1769	osg-tested-internal-1-5.osg.el5	Matyas Selmeci 564109	2012-04-16 16:44:06	✓
1768	osg-tested-internal-1-5.osg.el6	Matyas Selmeci 564109	2012-04-16 16:44:28	✓
1767	gratia-service-1.11-06.pre.osg.el5	Tanya Levshina 508821	2012-04-16 14:35:48	✓
1766	osg-tested-internal-1-4.osg.el5	Matyas Selmeci 564109	2012-04-16 14:26:13	✓
1765	osg-tested-internal-1-4.osg.el6	Matyas Selmeci 564109	2012-04-16 14:28:37	✓
1764	gridftp-hdfs-0.5.3-2.osg.el5	Doug Strain 834323	2012-04-16 14:23:25	✓
1763	gridftp-hdfs-0.5.3-2.osg.el6	Doug Strain 834323	2012-04-16 14:27:17	✓
1762	gratia-service-1.11-06.pre.osg.el6	Tanya Levshina 508821	2012-04-16 14:18:15	✓
1761	gratia-probe-1.10-11.pre.osg.el6	Tanya Levshina 508821	2012-04-16 10:32:31	✓
1760	gratia-probe-1.10-11.pre.osg.el5	Tanya Levshina 508821	2012-04-16 10:31:20	✓

Platform Support

- Koji is, hands down, the best way to package and manage a software distribution for RHEL-like distros.
  - RHEL5 is fully supported; only the worker node is supported on RHEL6. Late-April release will fill in the remaining
- However, it doesn't help much for Debian variants.
  - We are working with our stakeholders to understand what's needed versus what EMI/IGE already has gotten into Debian.
  - Until then, the pacman-based software is available.

Included Software

- Major upgrades to many pieces of software.
  - Globus 4.0.8 => 5.2.0.
    - Includes a fairly large rewrite of GRAM, our gatekeeper software.
  - glexec 0.6.x => 0.8.x
  - VOMS 1.8.8 => 2.0.5
  - PRIMA for authz to LCMAPS
  - VOMS-Admin 2.0 => 2.6

Observations

- Most admins and users really like the new approach.
    - However, some sites really like relocatable installs (i.e., to distribute worker node client via NFS).
        - We are trying to work out what would be in a relocatable tarball for this case. Nothing concrete yet.

Observations

- There are some other teething issues:
  - After training folks for 5 years to "think like pacman" instead of "think like RHEL", folks are forced through a paradigm shift in installation and configuration. We think it's better, but it's still different!
  - On the OSG side, we are throwing out 5 years of institutional knowledge in an obscure packaging format, and replacing it with 9 months experience in RPMs. Hence, some packages don't have all the "polish" of the previous stack.
    - Still seeing significant progress each month on the "rough spots". All releases have functioned, but newer ones qualitatively are better integrated.

Observations

- Our combination of major software and packaging changes is unfortunate but was unavoidable to use community packages.
    - For example, some of the "bugs" are just differences in how GRAM 5.2 works.
    - In other cases, we've uncovered some severe bugs. Because we stayed with the old GRAM version for so long, it's been awhile since we've been surprised by such things.

Conclusions

- OSG3 represents OSG Software's attempt to interact more closely with the larger grid community.
  - Changes so large it is almost like a separate project.
  - We hope this approach will pay dividends for both OSG and WLCG as a whole.