



Open Science Grid

# **The Great Migration: From Pacman to RPMs**

Alain Roy

OSG Software Coordinator





# The big picture

---

1. We are stopping new Pacman-based releases of the VDT.

2. By December 2011, we will provide RPMs for most of the VDT.

You can try many of them out *now*.

Let's flesh that out a bit and explain the background.



# 1. We are stopping new Pacman-based releases of the VDT

---

- The existing Pacman caches will remain as long as needed (a year or more)
- We will do urgent updates as necessary
  - Urgent means “security” or “really important”



## 2. By December 2011, we will provide RPMs for most of the VDT

---

- Supported OSes (covers majority of sites):
  - Red Hat Enterprise Linux 5
  - CentOS 5
  - Scientific Linux 5
  - Both 32 & 64 bit
- “Most of the VDT” means:
  - Worker node/glexec
  - Client
  - VOMS
  - Compute Element/RSV
  - Bestman & Xrootd
  - HDFS (moving from existing CMS repo into new repo)
- What remains?
  - GUMS
  - Gratia Collector



# What about other OSes?

---

- Will consider other OSes in early 2012
- Possible candidates:
  - RHEL 6 & variants
  - Debian 6
  - Mac OS X client-only
- We need feedback on OS support:
  - Do you care about 32-bit anymore?
  - When do you need SL6 support?
  - I hear you might move services to SL6 much sooner than you are allowed to move your worker nodes.



# Why are we moving to RPMs?

---

- Pacman was very nice:
  - Multi-platform
  - Root & non-root installs
  - Install in any location
- But Pacman didn't mesh well with the larger community:
  - People are already using RPMs for the rest of the OS
  - We were repackaging an increasingly large set of software already available as RPMs
  - It was hard for people to contribute to the VDT



# Principle of Community Packaging

The OSG Software Team should be a good community citizen when it comes to packaging: When possible, we should use packages from existing and/or broader communities; when that is not possible, we should make our own packaging but contribute them back to the broader communities.

Therefore, we should package software only when one of the following is true:

- The software is not already packaged; or
- The software is packaged but needs significant changes to be acceptable to our users. (Different version, extra patches, etc...)

Otherwise we should use the existing packaging provided by external developers or software repositories.



# In short: leverage the community

---

## **Borrow**

The VDT should borrow packages from the larger community whenever possible.

## **Contribute**

The VDT should contribute packages to the larger community whenever possible

## **Accept Donations**

The VDT should make it easy for people to donate packages to the VDT





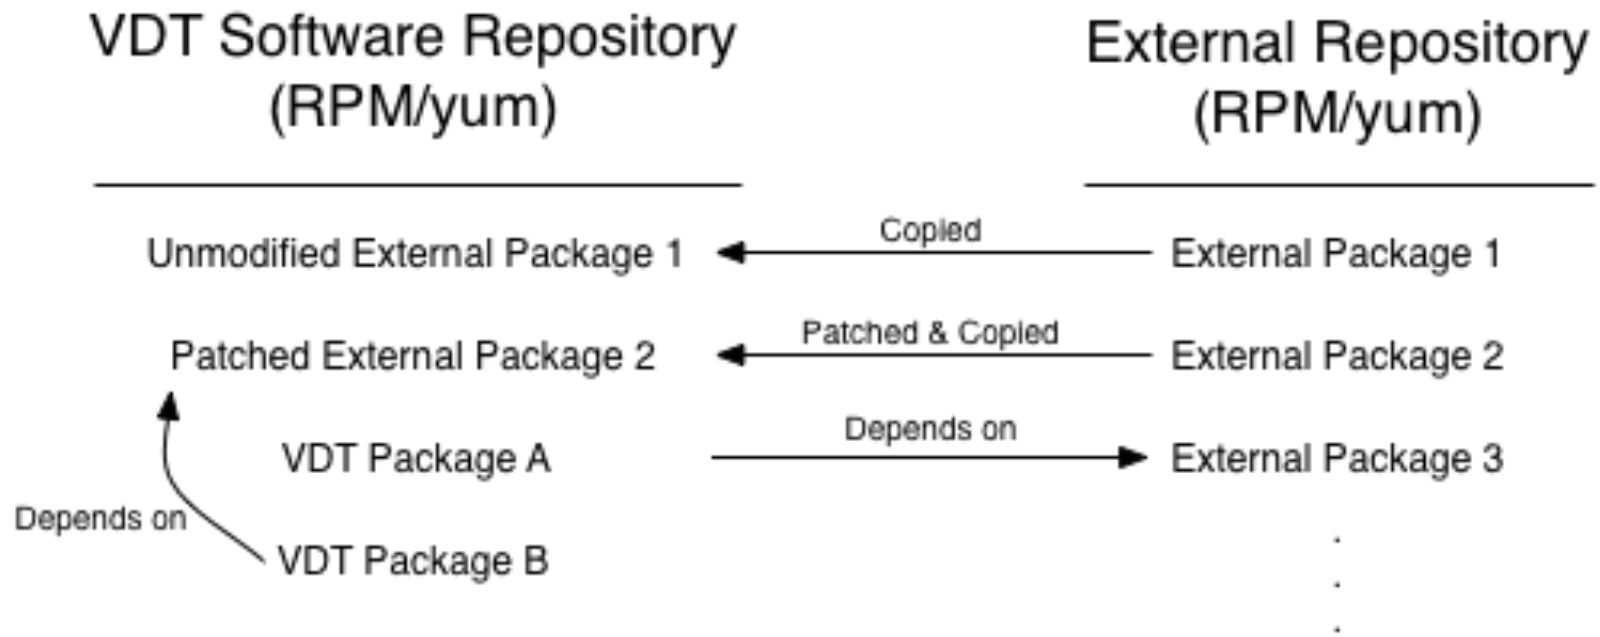
# Some technical details: RPM/yum

---

- RPMs:
  - Installed with yum
  - Not relocatable
  - Use your standard RPM tools
- We require use of:
  - OS yum repository
  - EPEL yum repository
  - OSG yum repository
  - Yum priorities
- We may not be compatible with:
  - DAG
  - rpmforge
  - jpackage

# The VDT is an amalgamation of repos

- Refer to packages in OS/EPEL
- Copy/patch packages as needed
- Create packages as needed





# Configuration

---

- For now: `osg-configure`
  - Renamed from `configure-osg`
  - Maintain continuity with existing practices
  - Unsure if it's the right long-term choice
- But:
  - `osg-configure` will be less complete than in the past (e.g. you need to use `chkconfig` instead of `vdt-control`)
  - May be some documentation to fill in gaps
  - `config.ini` is now `/etc/osg/config.d/* .ini`



# Some other changes

---

- Upgrading to Globus 5.2
  - Big new release, not yet released
  - Can't (yet) use Globus from EPEL
    - It's 5.0 and broken
    - We're trying to get 5.2 into EPEL
  - Should perform as well or better than Globus 4.0.8, but better packaged, updated
- There is no `$VDT_LOCATION`, no *setup.sh*
- RPM deployment with Condor requires 7.6.0 or later
  - Condor 7.4.x RPMs don't appear to "provide" the software
  - For glxexec, Condor 7.4.x *procd* isn't sufficient.



## Common questions

---

- Can I install the VDT in a different location (for a shared filesystem for the worker nodes)?
  - Not yet. If this is a serious problem for you, let us know. We are likely to implement this in early 2012.
  - We assume that many of you are already using cluster management systems (cfengine, puppet...) and can deploy the OSG software similarly. Is that true?
- Can I install the VDT as non-root?
  - Not yet.



# Is this just vaporware?

- *Last week* we released a subset:
  - VOMS (server & client)
  - CA Certificates
- *Today* you can try out stable test-ready versions:
  - Worker Node Client
  - Glxec
  - Client
  - Xrootd/Bestman
  - CE (mostly, GIP fix in a day or three)
  - RSV
- *Very soon*:
  - CMS HDFS repo moving into OSG repo

Brief how-to documentation:

<https://twiki.grid.iu.edu/bin/view/Documentation/Release3/>



# We need your help!

---

- Participate in our external beta testing
  - Try out the worker node? The CE?
- Participate in the integration test bed
- Ask us questions
- Give us feedback



Open Science Grid

# Questions?

---

Questions? Now or later:

[roy@cs.wisc.edu](mailto:roy@cs.wisc.edu)