




Use of Fedora's Tools and Policies outside Fedora

FudCon 09, Berlin



Steve Traylen
EGEE Automation Team
CERN, IT-GD
steve.traylen@cern.ch



- CERN
 - LHC GRID or LCG.
 - Scale
 - Operating 300 Sites
-
- Integration and Deployment
 - Fedora's Packaging Polices
 - Deployment Tools - koji, mock.



Grid
Deployment

CERN stands for 50 years of..

- fundamental research and discoveries
- technological innovation
- training and education
- bringing the world together



1954 Rebuilding Europe
First meeting of the
CERN Council



1980 East meets West
Visit of a delegation from Beijing



**2008 Global
Collaboration**
The Large Hadron
Collider involves over 80
countries

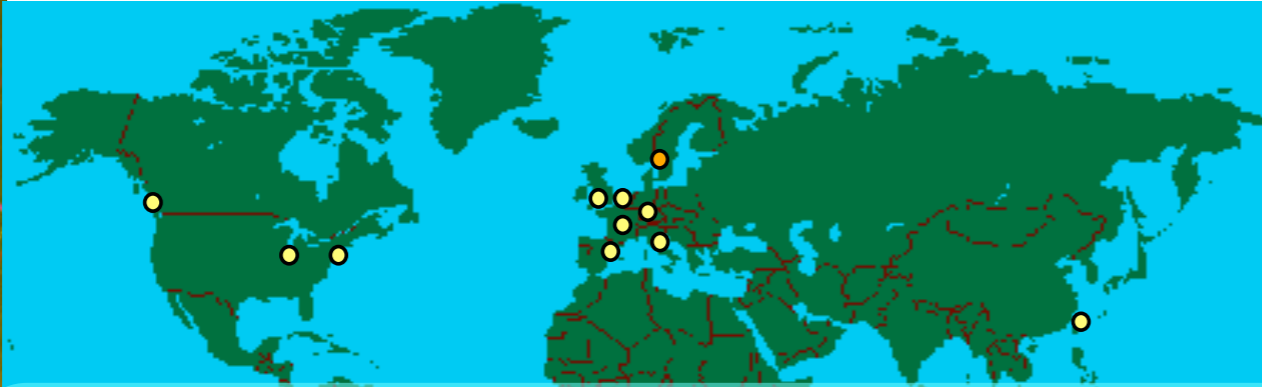
Grid
Deployment

CERN's mission in Science

- Understand the fundamental laws of nature
 - Accelerate elementary particles and make them collide
 - Compare the results with the theory.
 - Provide a world-class laboratory to researchers in Europe and beyond
-
- **2500 employees:** physicists, engineers, technicians, craftsmen, administrators, secretaries, ... (shrinking)
 - **6500 visiting scientists:** (half of the world's particle physicists representing 500 universities and over 80 nationalities - (increasing)
 - Budget: **~1 Billion Swiss Francs** per year.
 - Additional contributions by participating institutes.

Grid
Deployment

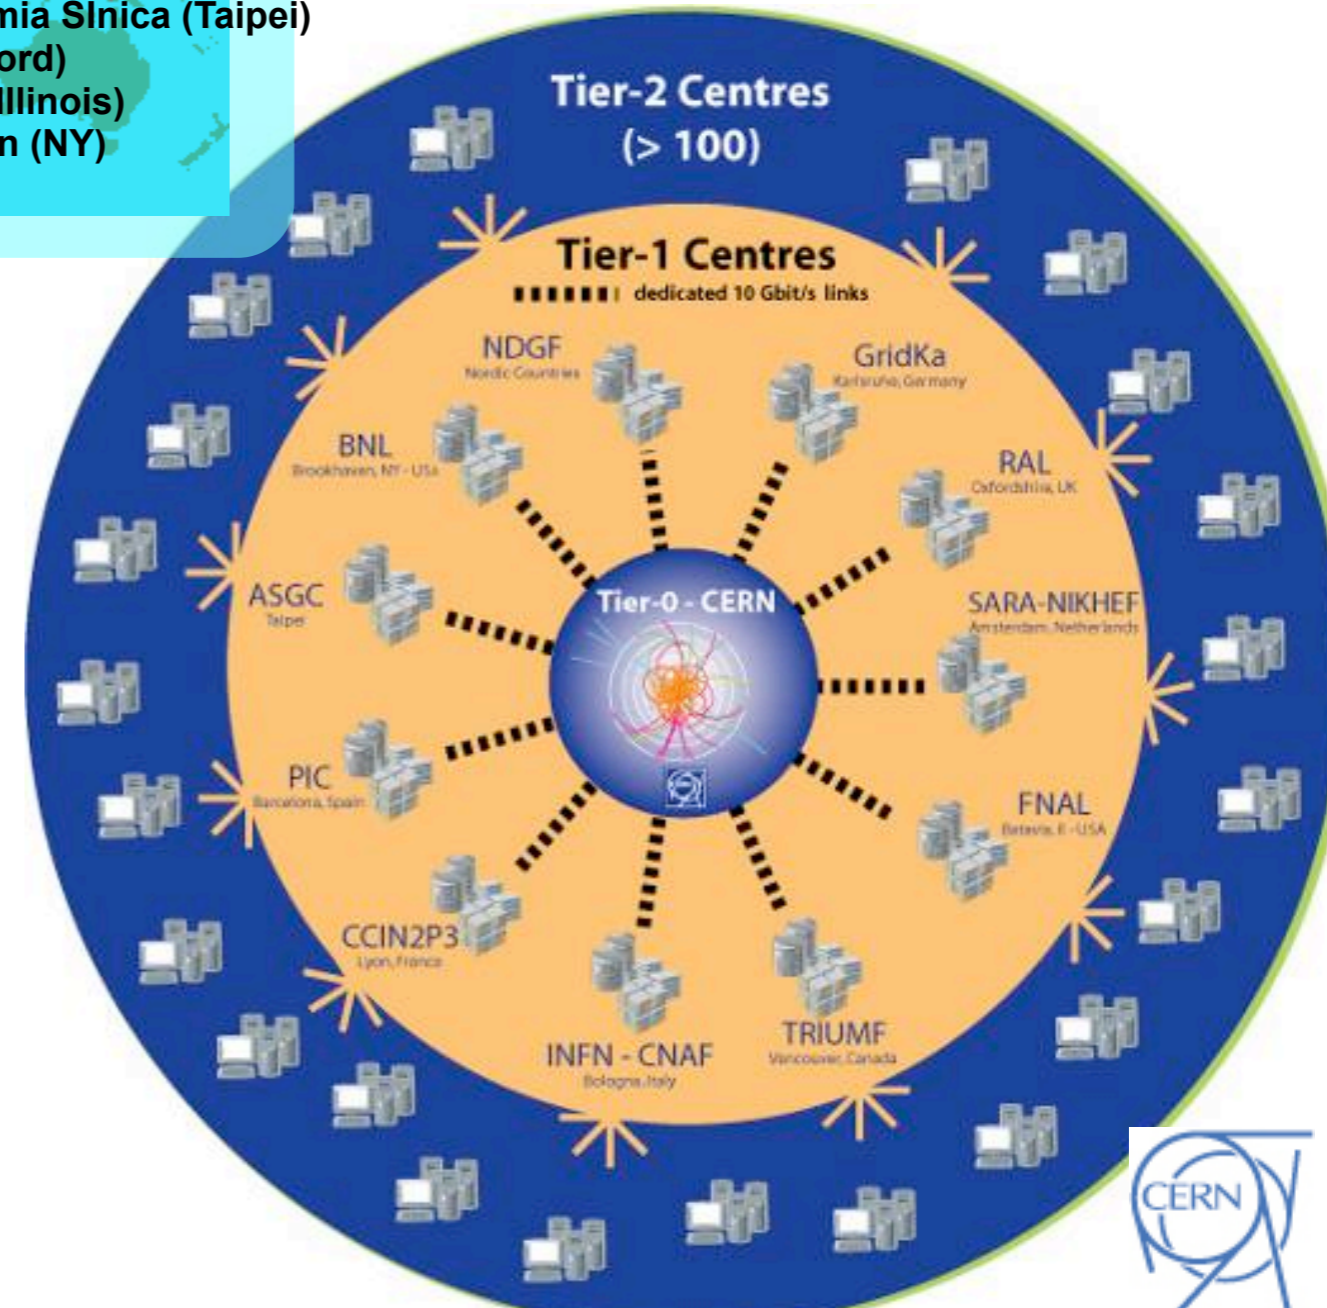
LHC Computing GRID (LCG)



- Canada – Triumf (Vancouver)
- France – IN2P3 (Lyon)
- Germany – Forschungszentrum Karlsruhe
- Italy – CNAF (Bologna)
- Netherlands – NIKHEF/SARA (Amsterdam)
- Nordic countries – distributed Tier-1
- Spain – PIC (Barcelona)
- Taiwan – Academia Sinica (Taipei)
- UK – CLRC (Oxford)
- US – FermiLab (Illinois) – Brookhaven (NY)

- 1 Tier0 @ CERN
- 11 Tier1s
- > 250 Tier2s.

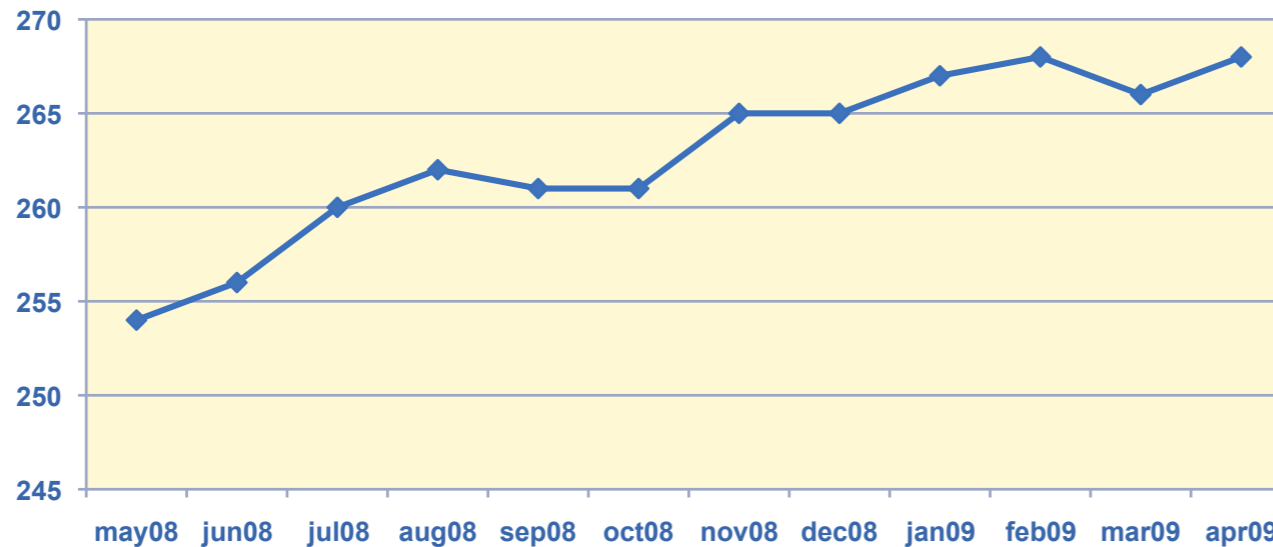
- 15 PB of data/year
- 11 * 10 Gbit links
 - Tier0->Tier1
- Worldwide analysis
 - Local funding.
 - Grid




Grid Deployment



Scale of LHC Computing Grid



- **268 Certified Sites.**



- **12 Million Jobs per Month.**

- 268 Sites
- 45-50 countries.
- 12,000,000 jobs per month
- 139,000 CPU cores
- 25 Pb disk
- 38 Pb tape

Grid
Deployment

Monitoring ~300 Sites.

Large range of site sizes with varying quality of services

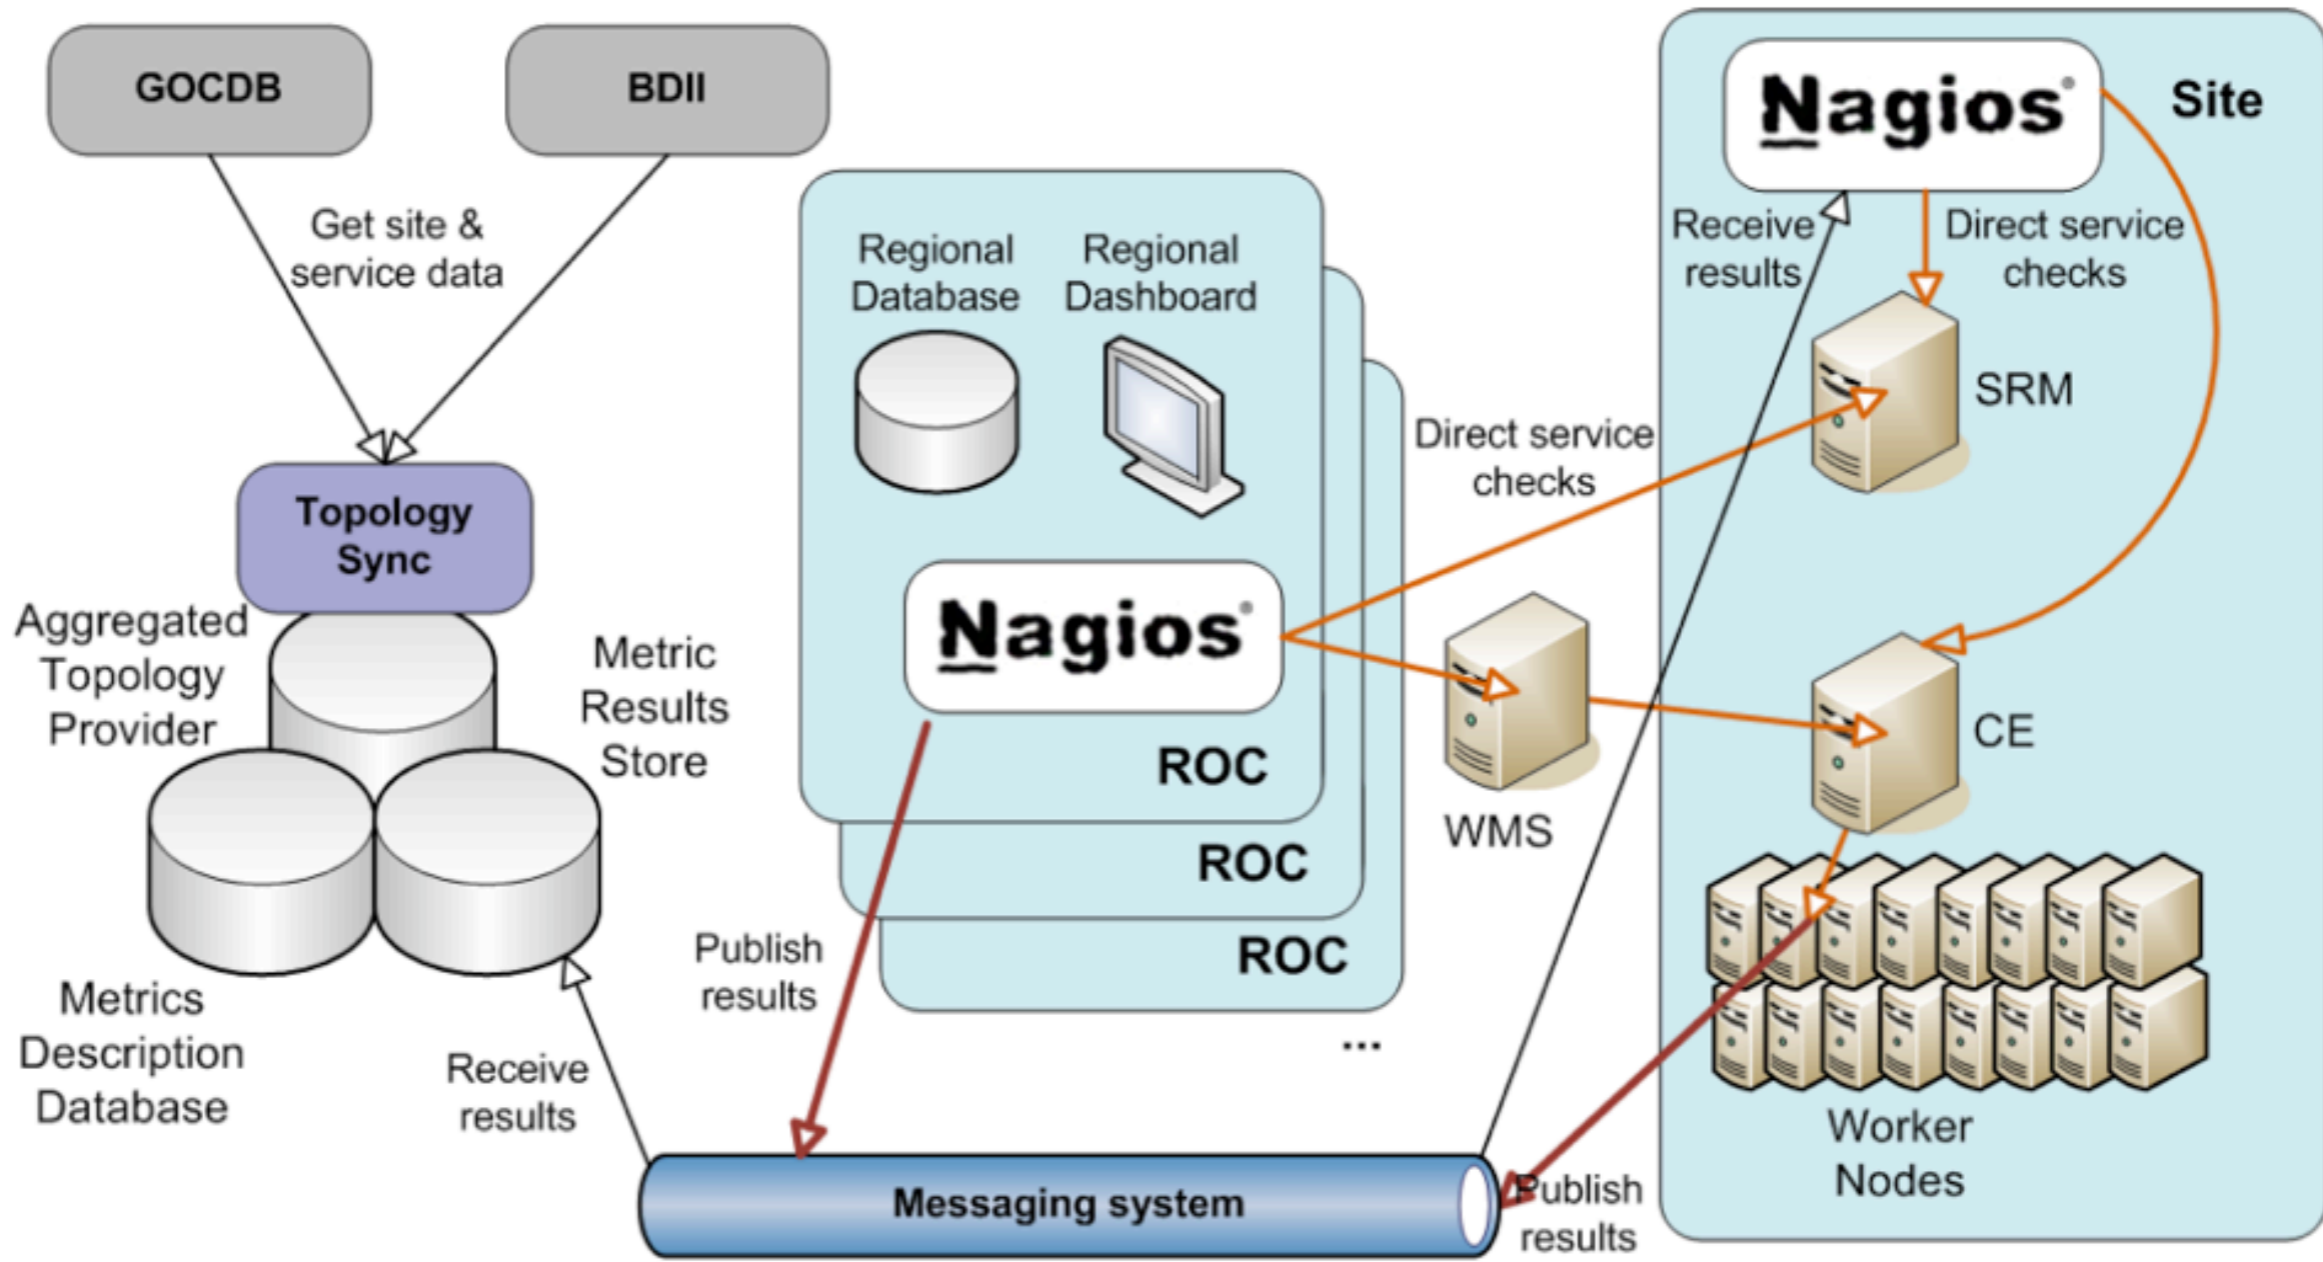
Small Sites @ University Dept'
Site operated by a MSc student.
No or little config mgt or monitoring.

Large Sites @ CERN + Tier1s
Puppet, cfengine, quattor, ...
24hour callout support

- Operation & Automation team's role.
 - Measure service availability.
 - Present results to sites.
 - Encourages sites to improve
- Improve service for users.

Grid
Deployment

Nagios and ActiveMQ Chosen

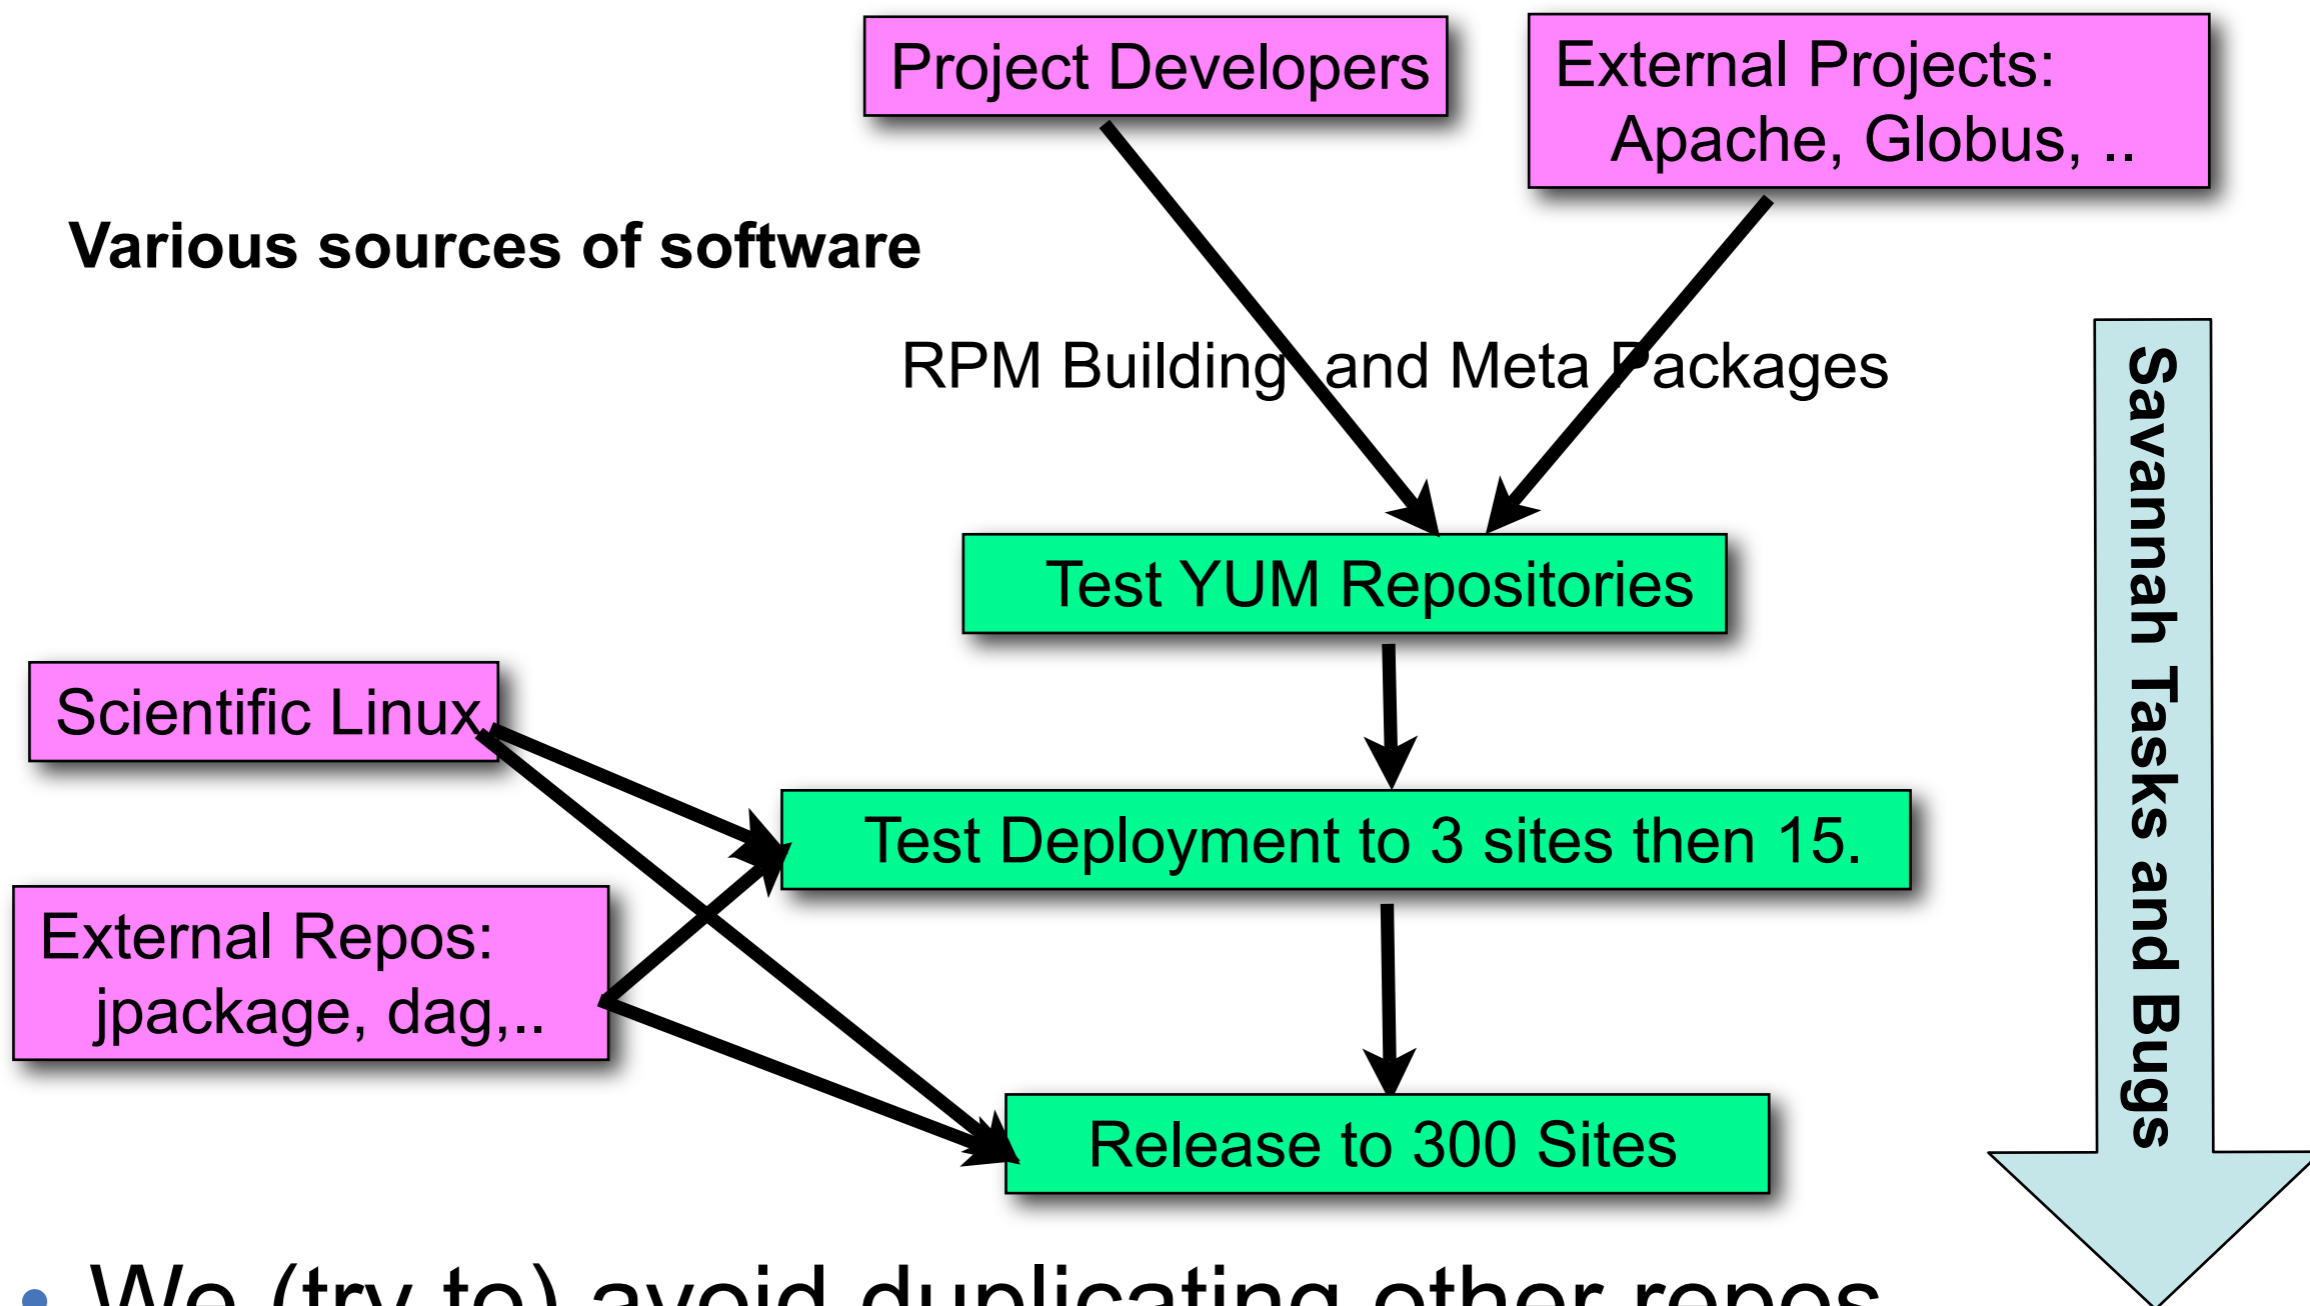


© Emir Imamagic / John Shade

- ROC - Regional Centers - Monitor Regions
- Results aggregated centrally via msg-bus.

Grid Deployment



Various sources of software



- We (try to) avoid duplicating other repos.
 - We use dag and jpackage but this has hit us.
- EPEL hopefully better, missing just a few.

RPM Guidelines set in 2003

- Admins were worried about software quality.
- Developers did not want to trash the farms.
- There were no public RPM guidelines
 - FHS of course did exist.



- Compromise “Don’t touch the OS!!!”
 - %post, %pre scripts are banned
 - Install all files in /opt/<ProjectName>
 - e.g /opt/globus/ , /opt/egee, /opt/lcg
 - Config’ mgt’ tools (quattor) were expected to fill the gap.

- These simple RPMS (almost tar balls) resulted in:
 - A lot of extra documentation
 - A lot of extra configuration management
 - **67,000!!! lines of bash** written to configure.
 - Constant headaches with:
 - PATH, LD_LIBRARY_PATH, PYTHONPATH, ...
 - 32bit to 64bit was and still is bad.
 - Configuration being difficult has been
 - Blamed for site failures.

Grid
Deployment

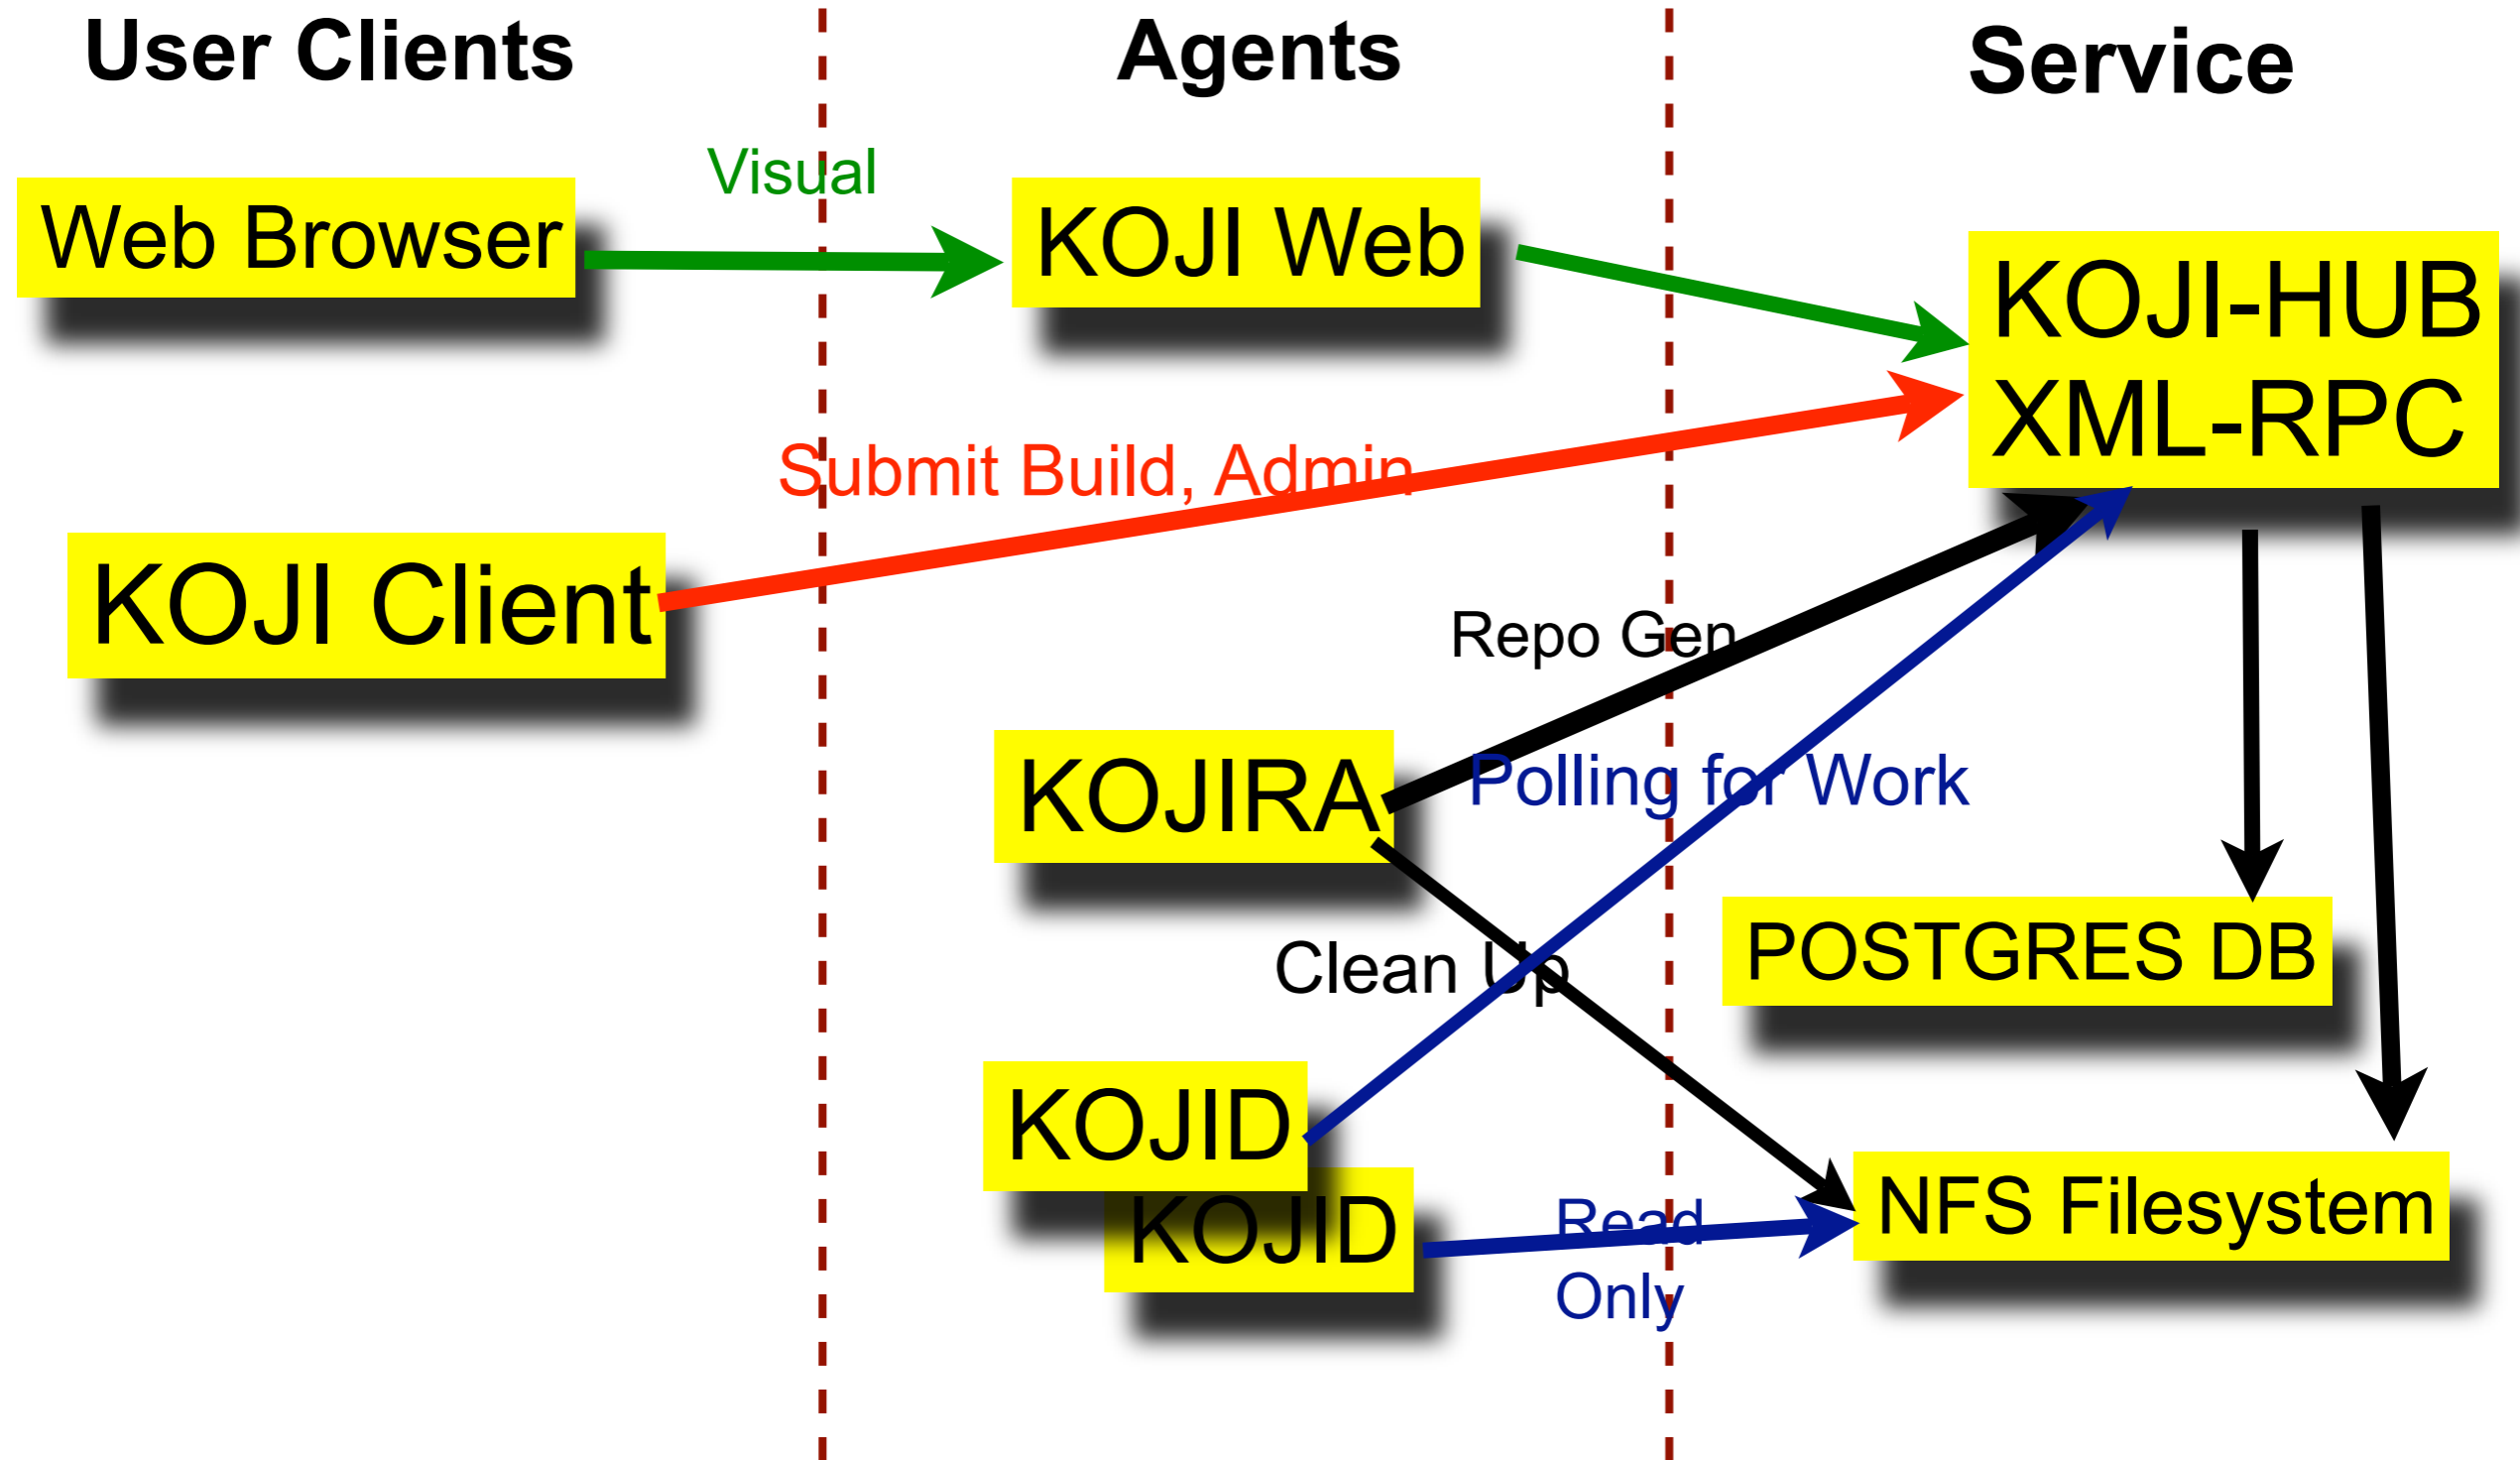
- Proposing & evaluating Fedora's pkg guidelines
 - their application to our grid deployment.
 - they are the most comprehensive today?
- Advantages:
 - Guidelines with a 3rd party.
 - We can influence.
 - We make more use of tricks and tools.
 - `/etc/cron.d` , `/etc/logrotate.d`
 - `rpmlint`
 - Easier to subsequently **contribute to Fedora.**
- Disadvantages:
 - Change - we are meant to be stabilizing for LHC (re-)startup in Autumn.

Grid
Deployment

Use Fedora Pkg Guidelines?

- Our Developers, Integrators have accepted.
- Site representative views being processed.
 - 300 sites => **a lot of views** range from..
 - Yes about bl**y time ...
 - No don't do that ever...
- The request to Fedora.
 - Please keep generic and avoid process.
 - **It is 95% good**, some counter examples:
 - A few mentions to consulting FESCO.
 - *An architecture failure, should have a bug with bug blocking on a FE-ExcludeArch-x86 bug.*
 - Not a problem in reality - common sense works.

Grid
Deployment



- KOJI-HUB - PYTHON xml-rpc
 - Interacts with database , the brain.
- KOJI Client - very simple client.
 - Adding users, submitting builds.
- KOJI Web - a client to koji-hub.
 - Visualization, Some admin, Some build resubmission.
- KOJID - Ask koji-hub for work.
 - e.g Building an RPM or rebuilding a yum repo
- KOJIRA - Clean up tasks.
 - Destroy redundant buildroots and repositories.
 - Ask koji-hub to rebuild new repositories as needed.

Our Koji Stats.

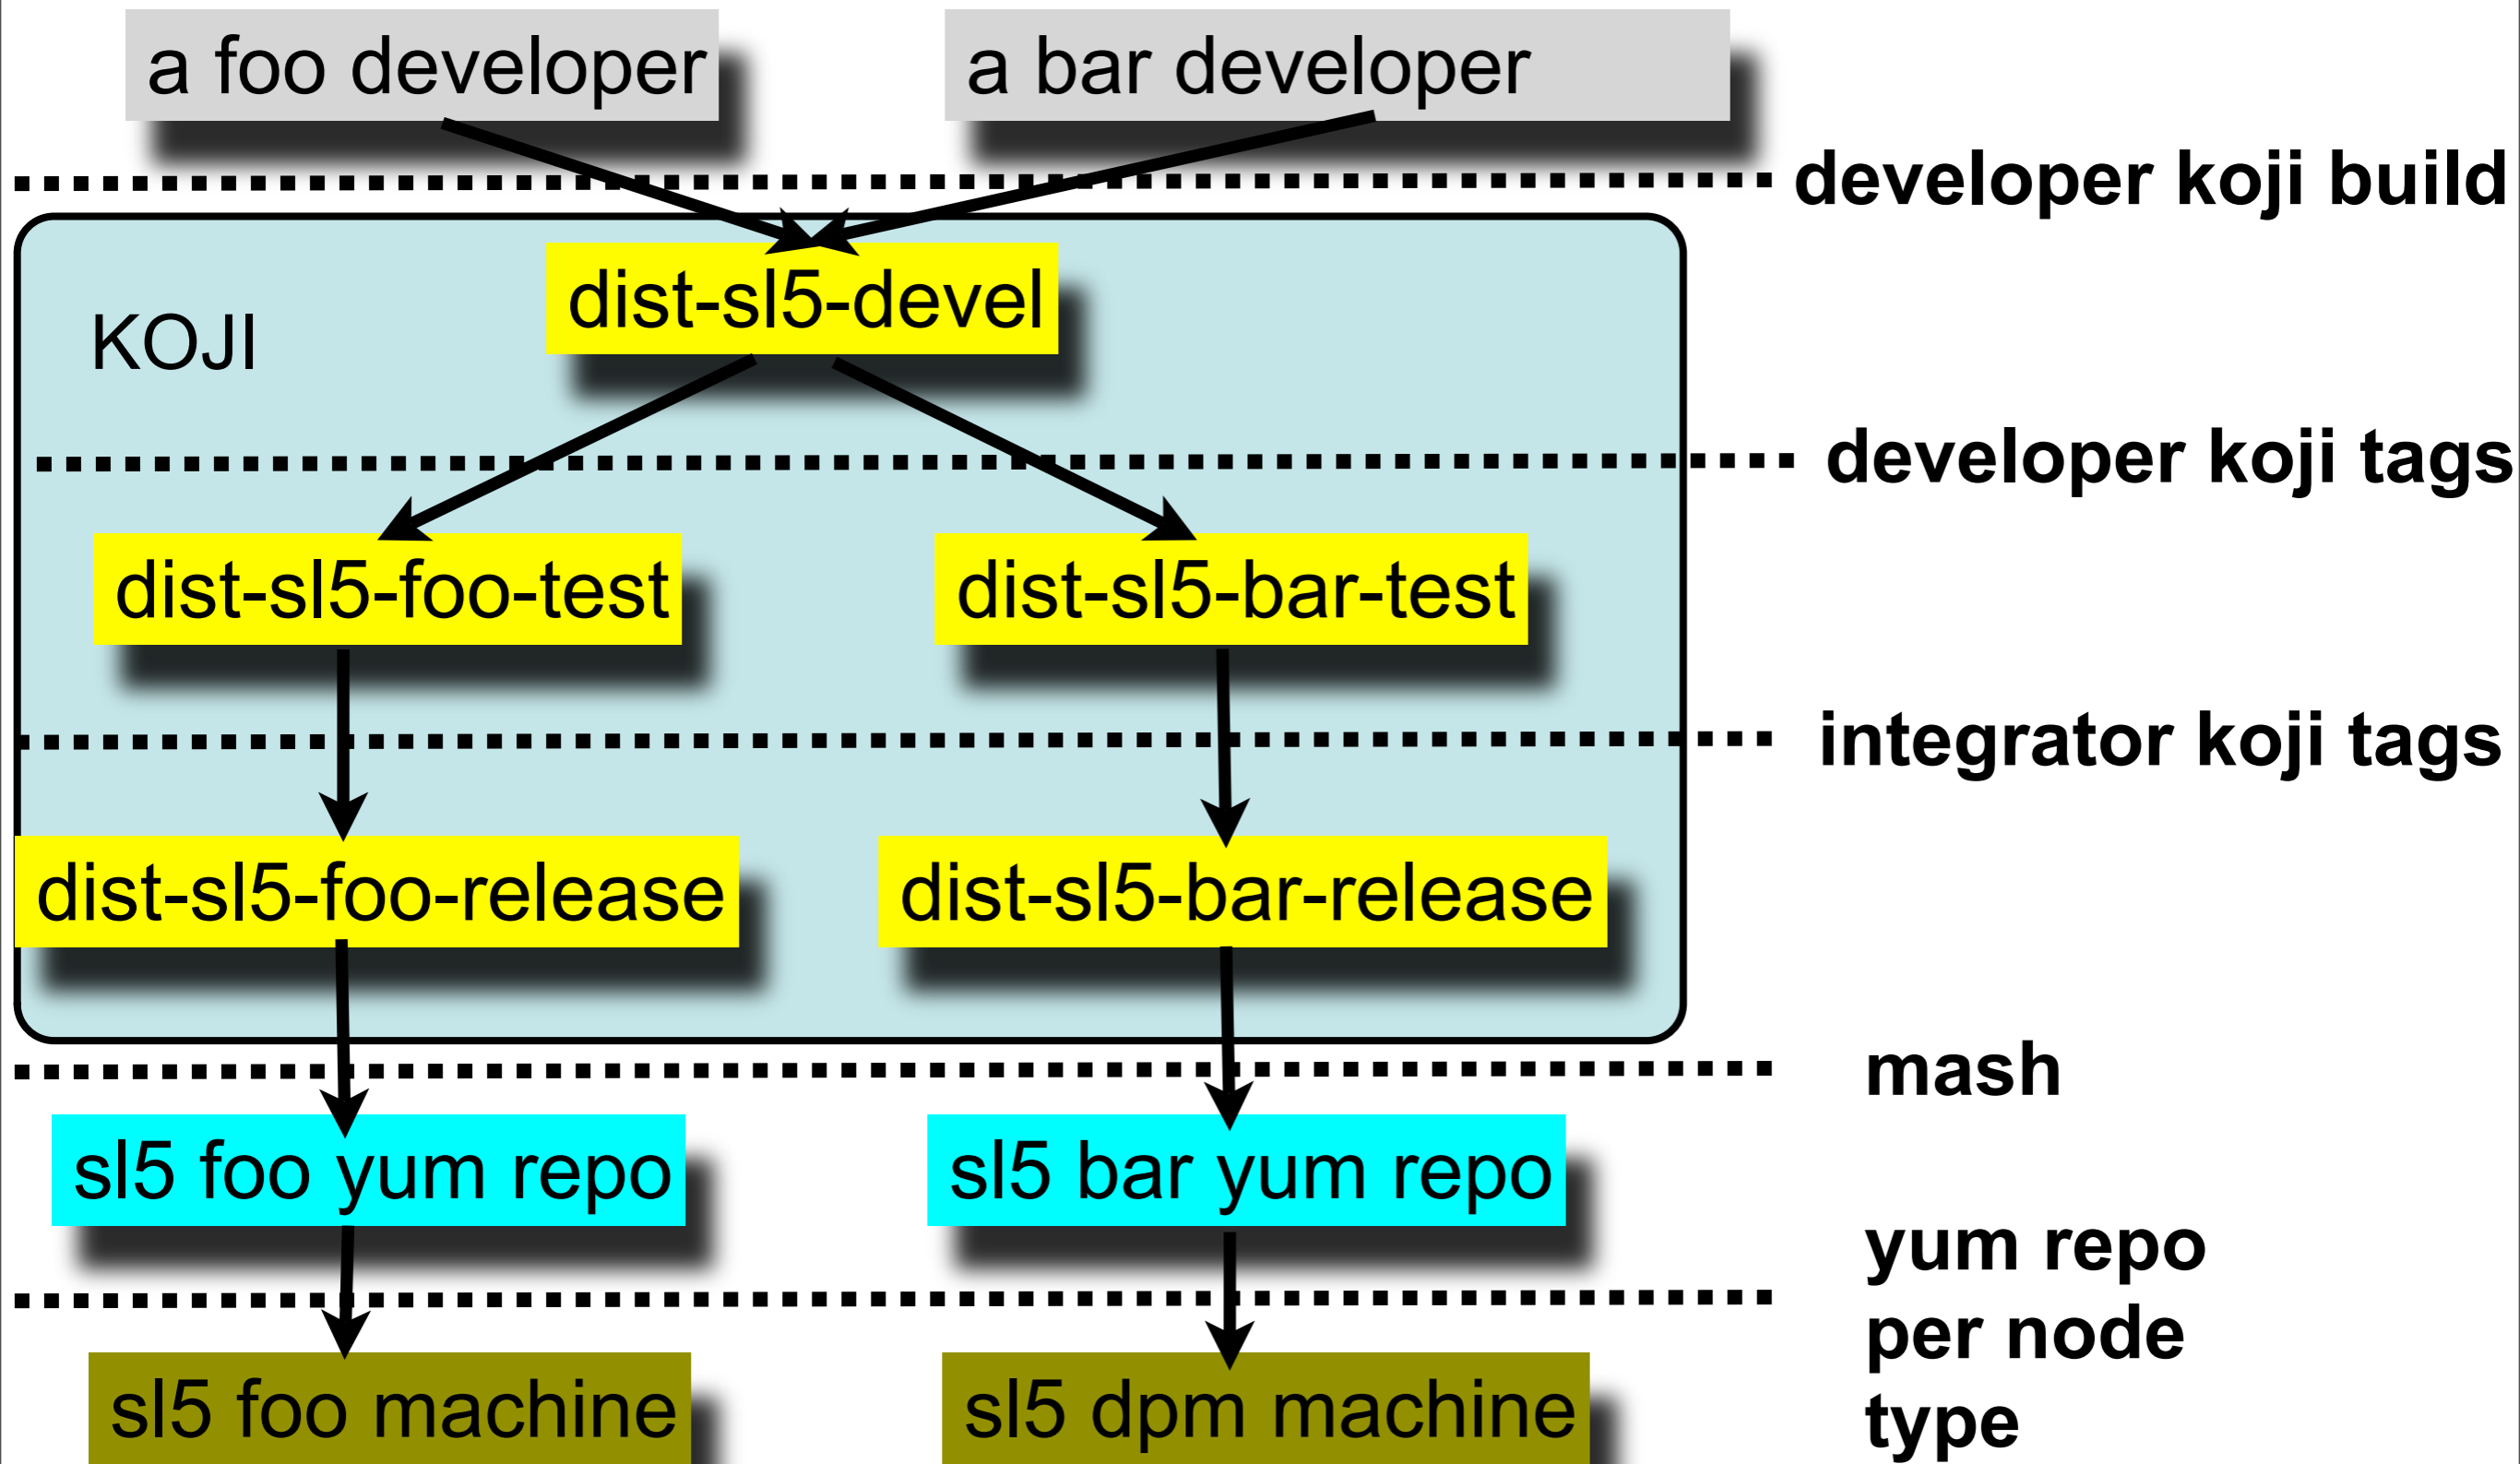
Metric	Value
Number of Users	9 (5 active)
Number of Hosts	1 for everything
Architectures	i386, x86_64
SRPM Names "N"	110, e.g gocdb-downtime
Unique Builds. "NVR"	356, e.g gocdb-downtime-0.1.el4
Build Success / Failed / Cancelled*	262 / 18 / 4
Targets	SL4, 5 and FC10.

Grid
Deployment

Why KOJI

- KOJI supports X509 certificates.
 - We have a PKI infrastructure already.
- Previous process problems:
 - One build system from cvs to release.
 - Code modifications went back to the developer.
The trigger in the system.
 - Even a rebuild against new lib.
- Use KOJI only for integration step.
- KOJI recently added support for external repos.
 - Done for the purposes of EPEL.
 - We do exactly the same as EPEL.
- KOJI is used by an open project - fedora.

Grid
Deployment



Koji Questions,

- Package signing.
 - Koji supports signatures but not signing.
 - Something else does it but what?
- Debian .debs.
 - We need to make them. Options:
 - Use pbuilder, debbootstrap, buildd,...
 - Add to koji, a good job for a summer student.
 - Is it feasible, interest?
 - Use something else to koji.
- bhodi
 - Should/could we use it?

Grid
Deployment

- Software is landing in fedora from many directions.
 - Globus recently added by Mattias Ellert.
 - Redhat's MRG adding Condor.
 - Sun's grid-engine now present.
- Its time the established grid software was added.
 - CERN ran webserver for 5 years before one was added to Redhat. (I think)
 - 1989 -> 1994 (Redhat 4) ?
 - (Of course CERN had the first one so not fair)
 - It's time grid software was just part of the OS.
 - Contribution the name of game of course.

Grid
Deployment

- Thank you, and Questions?
- CERN - <http://cern.ch/>
- LHC Computing Grid <http://cern.ch/lcg>
- KOJI : <http://skoji.cern.ch/koji>
- Multi-level-monitoring
<https://twiki.cern.ch/twiki/bin/view/EGEE/MultiLevelMonitoringOverview>
- Mock/Koji/Mash presentation to CERN.
<http://indico.cern.ch/event/55091>

Grid
Deployment