Integrating Condor into the Debian operating system

Michael Hanke

Debian Developer, Dartmouth College

CondorWeek, University of Wisconsin
Madison, Wisconsin

May 4th 2011

Background: Neuro-imaging

Background: Neuro-imaging

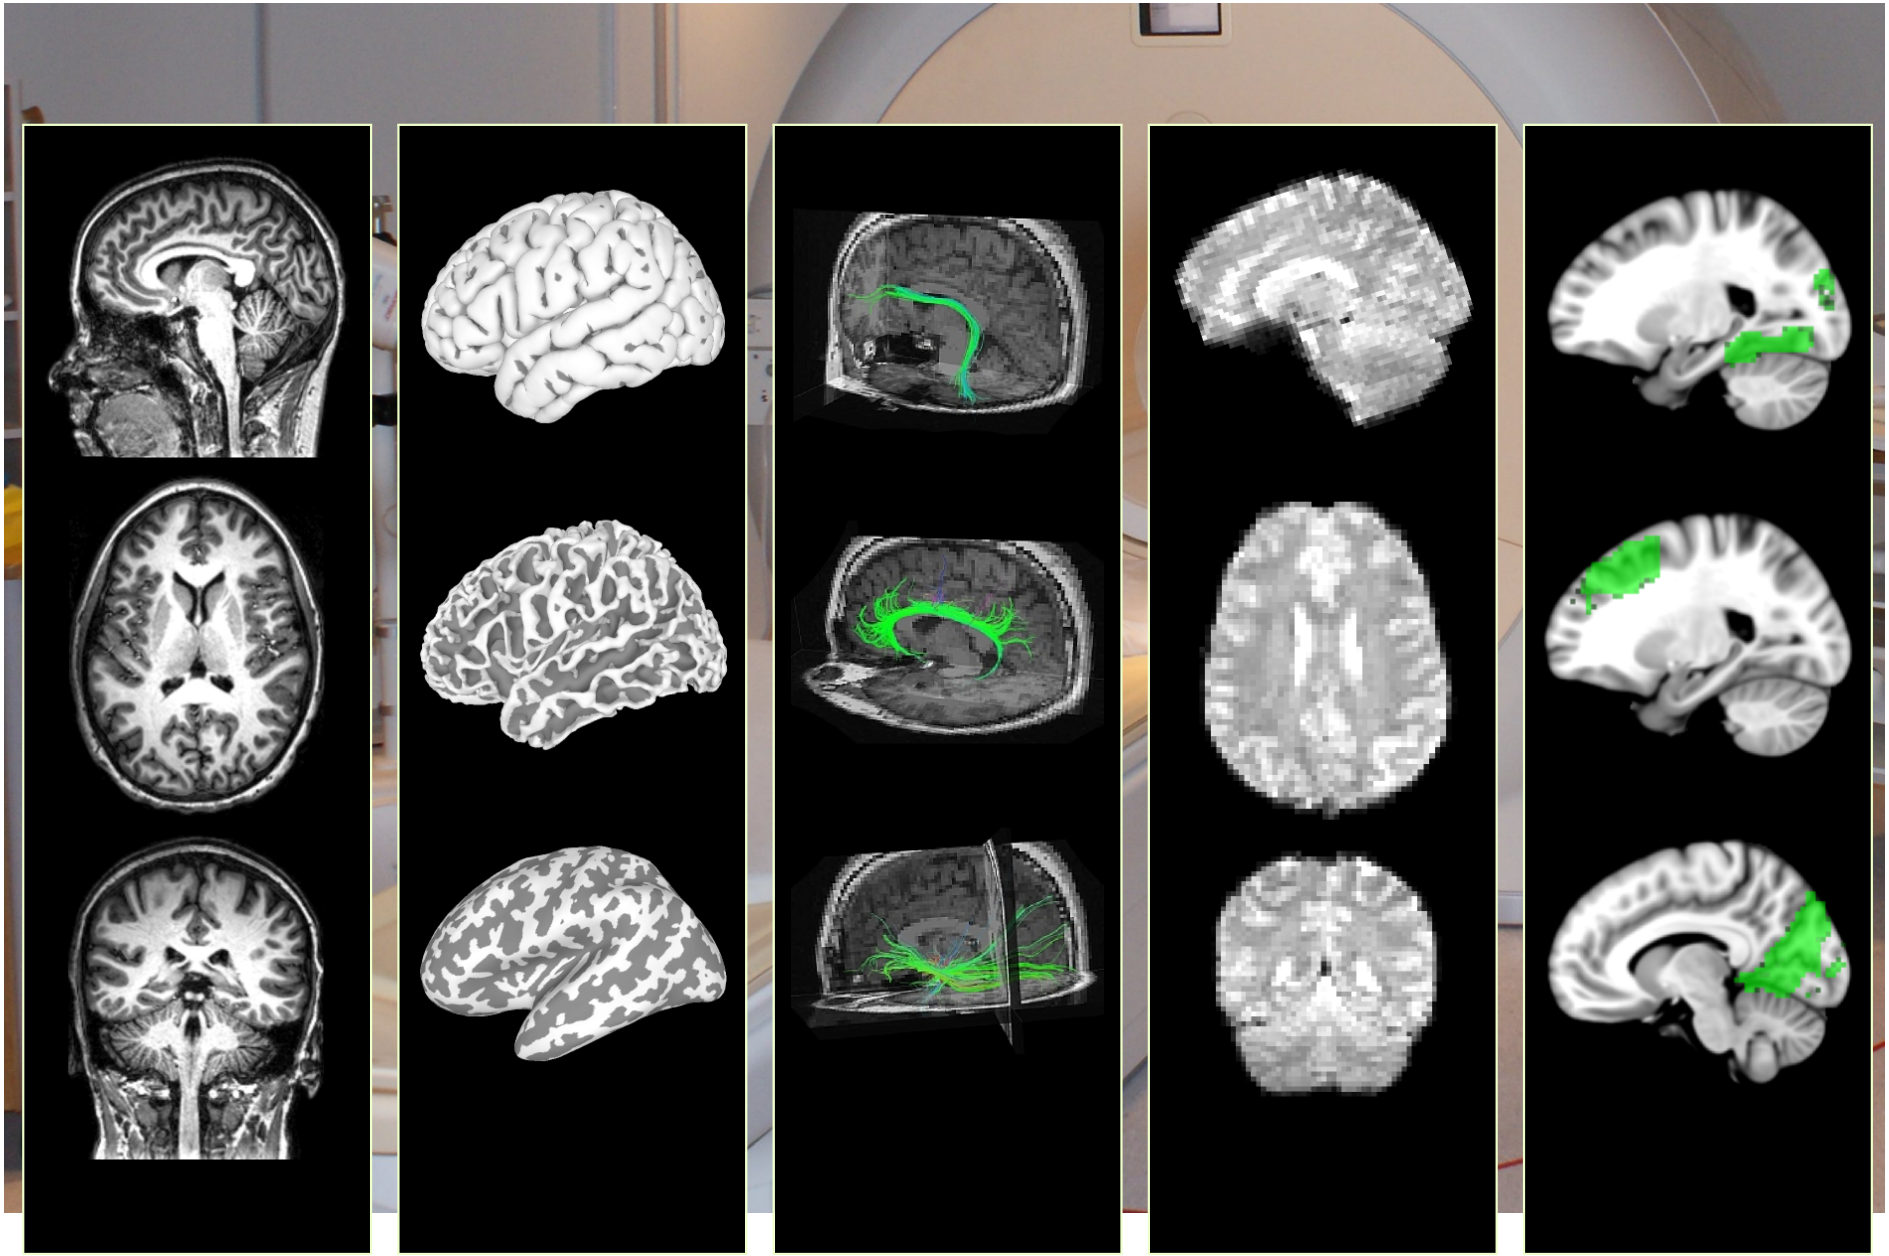

Research platform: Issues

Problem

- Complex analysis software suites
- Complicated, non-standard, or non-existing installation and update procedures
- Limited, non-uniform set of "supported platforms"
- Typical users have little technical background

`http://neuro.debian.net`

Research platform: NeuroDebian

Problem

- Complex analysis software suites
- Complicated, non-standard, or non-existing installation and update procedures
- Limited, non-uniform set of "supported platforms"
- Typical users have little technical background

Solution

- Integrate all relevant software in a common environment
- Make manual maintenance tasks trivial, or superfluous

$\longrightarrow$ Bring everything into Debian

`http://neuro.debian.net`

Why debian?

- Vast archive of maintained software (≈30000 binary packages)
- Origin of most active GNU/Linux distributions (63%; distrowatch.org)
- "Do-ocracy" instead of steering (commercial) entity – nevertheless 17-year release history
- No 2nd-class software
- Debian Science, Debian Med, ...

http://www.debian.org

Why debian?

- Vast archive of maintained software (≈30000 binary packages)
- Origin of most active GNU/Linux distributions (63%; distrowatch.org)
- "Do-ocracy" instead of steering (commercial) entity – nevertheless 17-year release history
- No 2nd-class software
- Debian Science, Debian Med, ...

Release notes Debian 6.0 (squeeze)

Debian GNU/Linux 6.0 is the first GNU/Linux distribution release ever to offer comprehensive support for magnetic resonance imaging (MRI) based neuroimaging research.

`http://www.debian.org`

What makes Condor attractive?

- Free and open-source software with a standard license
- Active research project
- Transparent development (VCS, bug tracker)
- Evidence for prosperous future

What makes Condor attractive?

- Free and open-source software with a standard license
- Active research project
- Transparent development (VCS, bug tracker)
- Evidence for prosperous future

- Ease of maintenance
- Support for whole DAG submissions
- Support for heterogeneous computing environments
  - "University-style" computing environments are a supported use case
  - Option for the cloud

How can Condor benefit from Debian integration?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈120 derivative distributions (distrowatch.org)

How can Condor benefit from Debian integration?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈120 derivative distributions (distrowatch.org)
- Mutual awareness
  - Explicitly documented dependencies
  - Synchronized transitions

How can Condor benefit from Debian integration?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈120 derivative distributions (distrowatch.org)
- Mutual awareness
  - Explicitly documented dependencies
  - Synchronized transitions
- Less maintenance work through modularity
  - 3rd-party software in dedicated packages maintained by someone else

How can Condor benefit from Debian integration?

- Extended reach
  - one stable release, two rolling "release" flavors
  - ≈120 derivative distributions (distrowatch.org)
- Mutual awareness
  - Explicitly documented dependencies
  - Synchronized transitions
- Less maintenance work through modularity
  - 3rd-party software in dedicated packages maintained by someone else
- Continuous integration testing
  - 13 hardware architectures
  - Three kernels
  - Continuous automated testing for
    - Build success
    - Clean installation/de-installation, Availability of dependencies
    - Policy compliance
    - Package conflicts

How much do you need to/should involve yourself?

How much do you need to/should involve yourself?

Who maintains the package?

- Upstream interest in Debian-integration is essential for a high-quality package
- Need to track Debian development

How much do you need to/should involve yourself?

Wh

- ity
-

Debian "stable"

Release:	when ready
Updates:	every two months (only critical fixes)
Security support:	one year after next stable (≈ 3 years)

How much do you need to/should involve yourself?

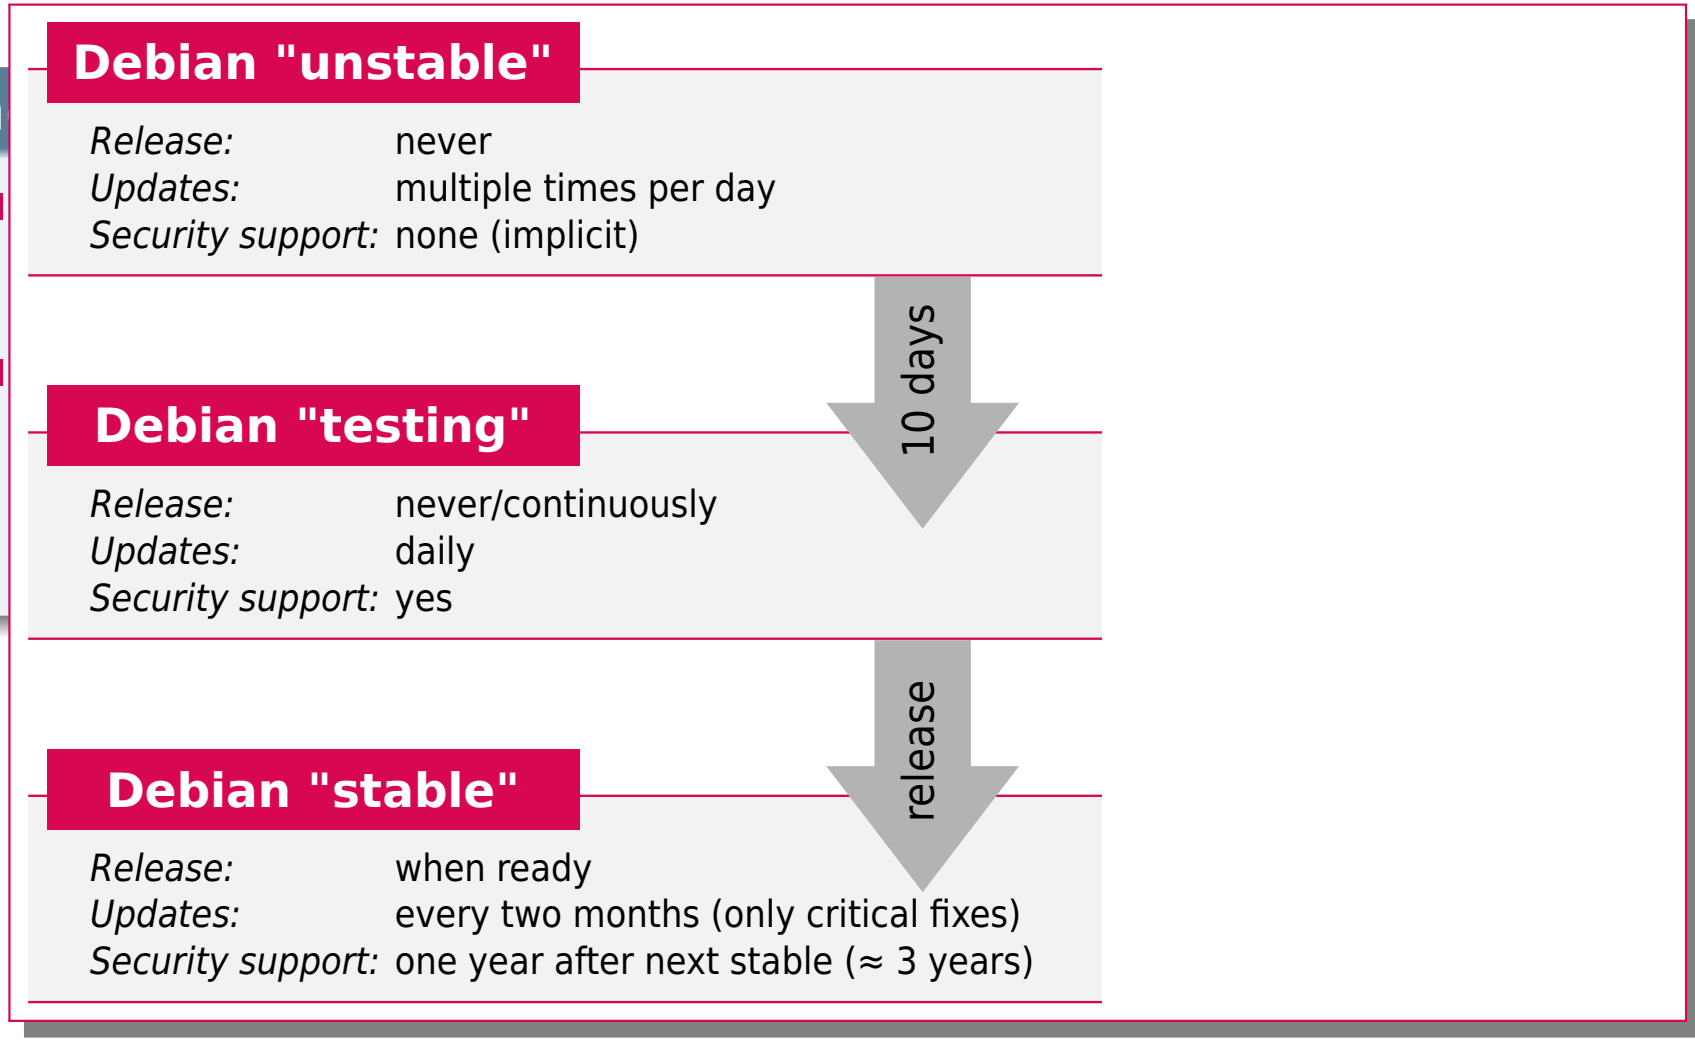

How much do you need to/should involve yourself?

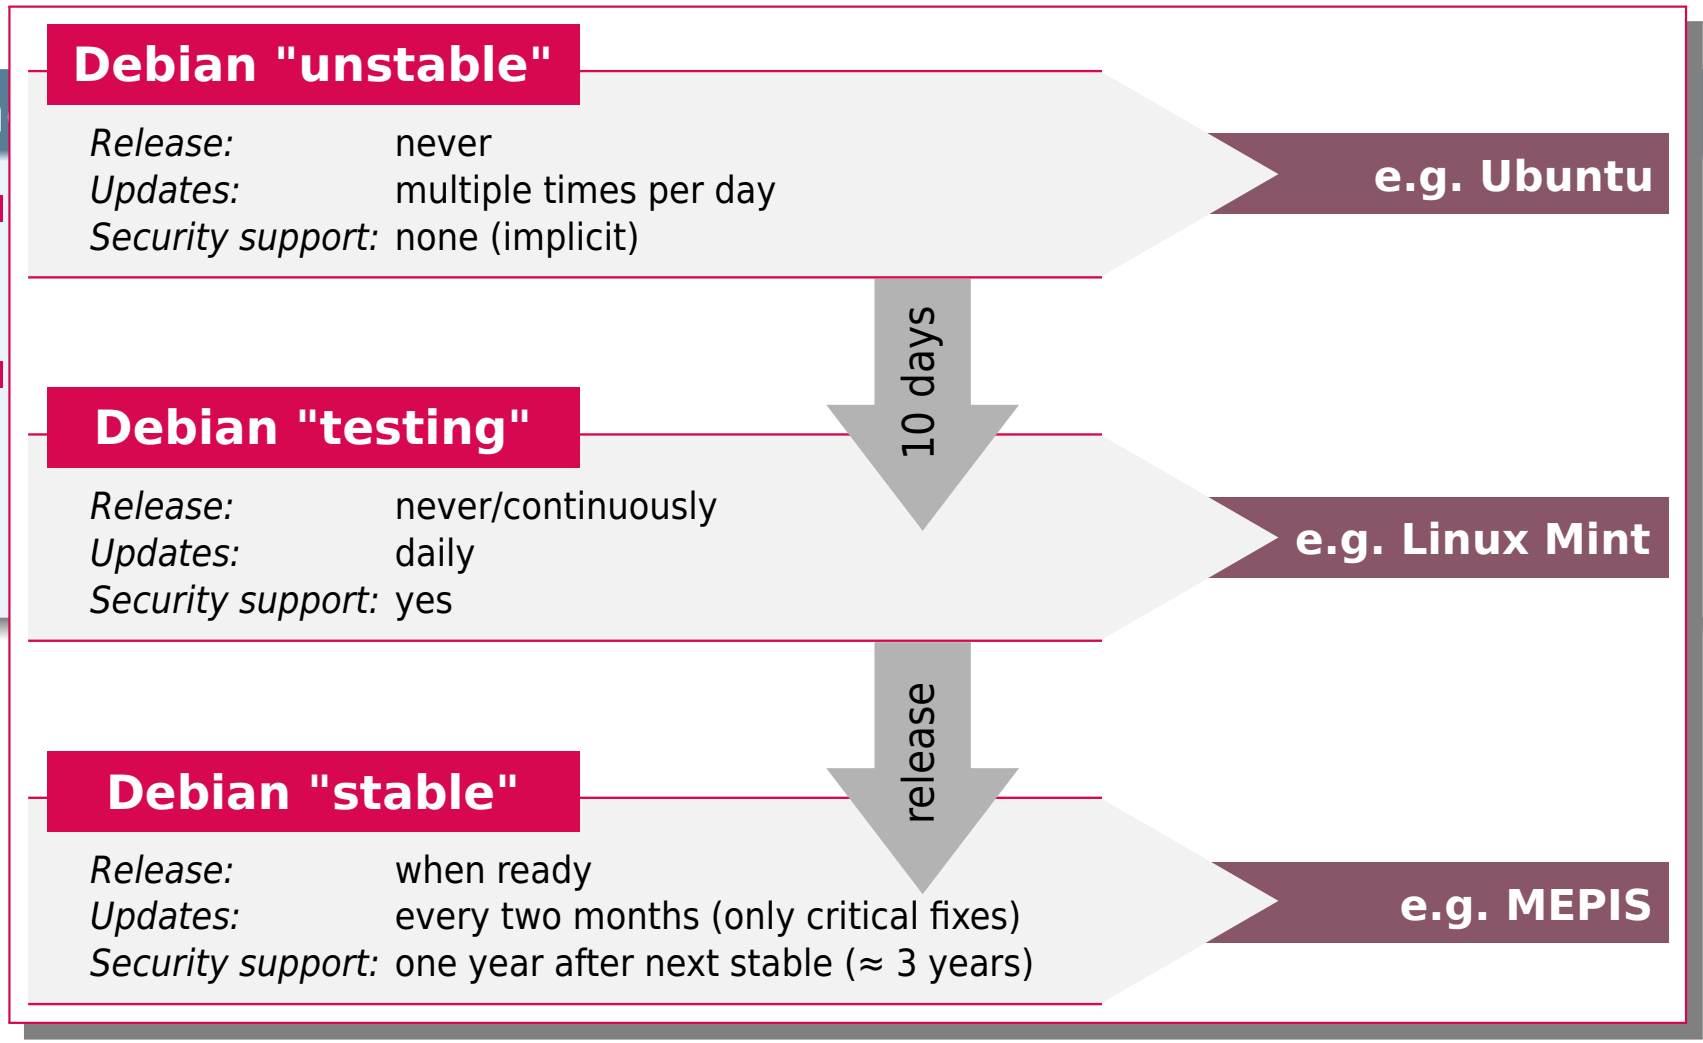

How much do you need to/should involve yourself?

Who maintains the package?

- Upstream interest in Debian-integration is essential for a high-quality package
- Need to track Debian development
- Need upload privileges, but sponsor could be enough

$\longrightarrow$ Ideally: team maintenance with public VCS

How much do you need to/should involve yourself?

Who maintains the package?

- Upstream interest in Debian-integration is essential for a high-quality package
- Need to track Debian development
- Need upload privileges, but sponsor could be enough

$\longrightarrow$ Ideally: team maintenance with public VCS

Just ship source code or also binary packages?

- Why not? If necessary.
- But from a common source package!
- backports.debian.org, neuro.debian.net

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
   - `classad`, `globus`, `krb5`, `libvirt`, ...

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
   - Condor 7.6.0 (clipped)
   - Based on previous packaging attempts
   - Debconf-based setup
   - `condor`, `condor-doc`, `condor-dbg` (, `condor-tests`)

Integration roadmap

1. H…e]
2. In…

Integration roadmap

1. H…e]
2. In…

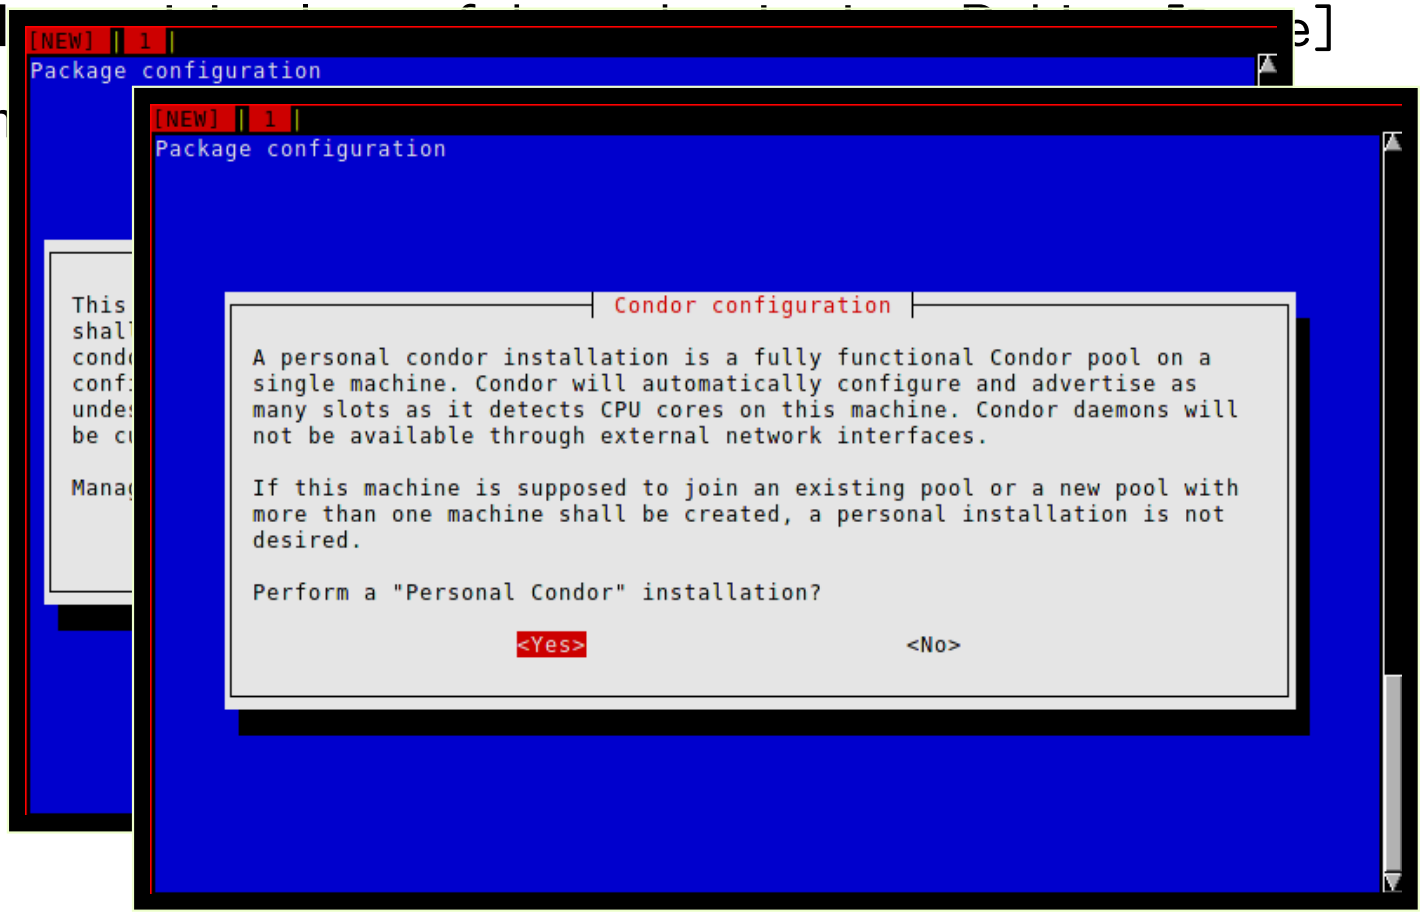

Integration roadmap

1. H…e]
2. In…

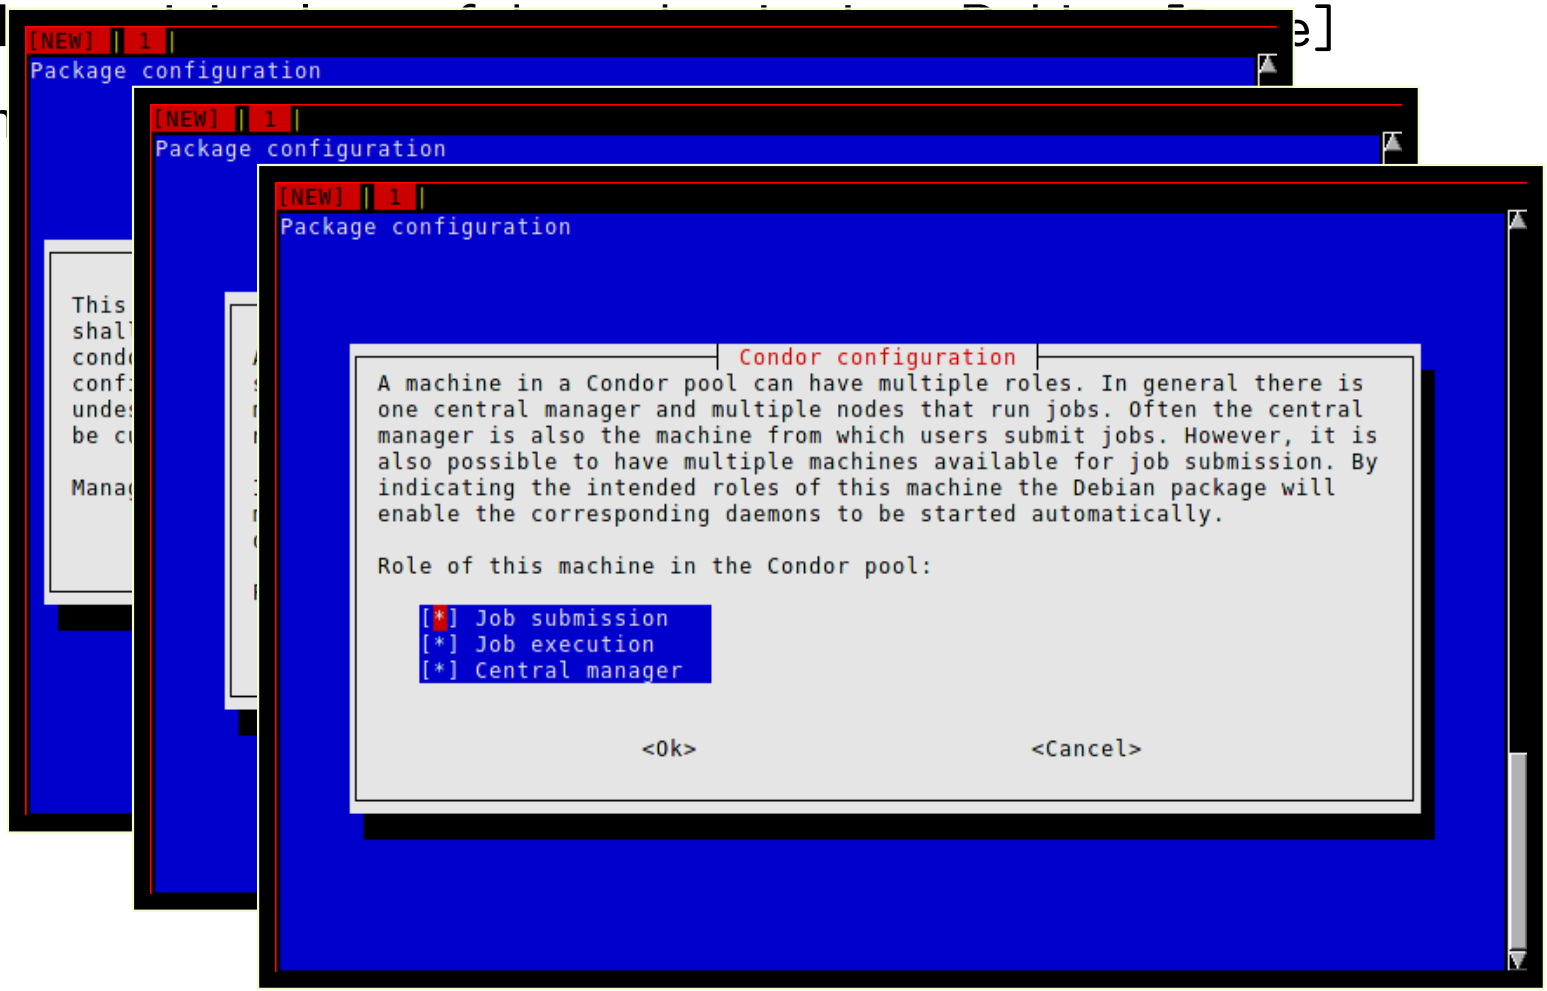

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
   - Package in use since 4 months
   - Needs more testers

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
   - dmtcp [Accepted in Debian]
   - gsoap [Uploaded]
   - cctools [Upload pending]
   - nmi

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
6. Improve "just works" experience [WiP]
   - Engine management for IPython
   - DAGMan support for NiPyPE
   - Makeflow support for FSL
   - VM universe with VirtualBox

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
6. Improve "just works" experience [WiP]
7. Make Condor's test suite run and pass [ToDo]
   - Ship in dedicated package

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
6. Improve "just works" experience [WiP]
7. Make Condor's test suite run and pass [ToDo]
8. Offer "standard universe" in official Debian package [ToDo]

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
6. Improve "just works" experience [WiP]
7. Make Condor's test suite run and pass [ToDo]
8. Offer "standard universe" in official Debian package [ToDo]
9. Enable every possible Condor feature in the Debian package [ToDo]
   - Java, VM, Grid, Cloud, Rocket science
   - Missing: up-to-date Hadoop

Integration roadmap

1. Have minimal set of dependencies into Debian [Done]
2. Initial working Condor package draft [Done]
3. Test and upload to Debian [Pending]
4. Fix QA bugs
5. Enrich Condor environment [WiP]
6. Improve "just works" experience [WiP]
7. Make Condor's test suite run and pass [ToDo]
8. Offer "standard universe" in official Debian package [ToDo]
9. Enable every possible Condor feature in the Debian package [ToDo]
10. Use Condor to improve Debian [ToDo]
    - E.g. backfill jobs with package QA tests

Acknowledgements

Yarik Halchenko
Jim Haxby
Swaroop Guntupalli
Andy Connolly

Condor Team
Debian Community

Mattias Ellert
Peter Tröger
Dustin Kirkland
Ian Alderman

Thanks!

Michael Hanke
mih@debian.org
http://mih.voxindeserto.de