

Status Ubuntu 8.04 LTS mit FAI

FAI-Setup for MPP-Computers, currently focused on ILC/Belle

Martin Ritter

2009-02-06

FAI

Fully Automatic Install is a Framework to install and maintain debian-based distributions. It's written and maintained by Thomas Lange: <http://www.informatik.uni-koeln.de/fai/>

Advantages

- ▶ Fully unattended Installation
- ▶ Powerful, class-based System allows to manage different configurations easily
- ▶ Update-facility to update running systems whenever convenient
- ▶ Already used in large environments

Installation using FAI

Once you have set up the Server, installation consists of

1. Setup DHCP-Configuration for client using MAC-address
2. Setup PXE-boot-configuration: `"fai-chboot -IF hostname"` on Server
3. Boot Client using PXE
4. Wait until installation is finished
5. Tell client to boot from hard-disk: `"fai-chboot -d hostname"` on Server
6. Reboot Client

Steps 4-6 can be automated using `faimond`

Updates using FAI

Once you have installed the client, updates consist of

1. Run `“fai softupdate”` on client

Step one can be automated using `cron`

Why Ubuntu?

Ubuntu is a community-developed, debian-based distribution under the lead of Canonical Ltd.

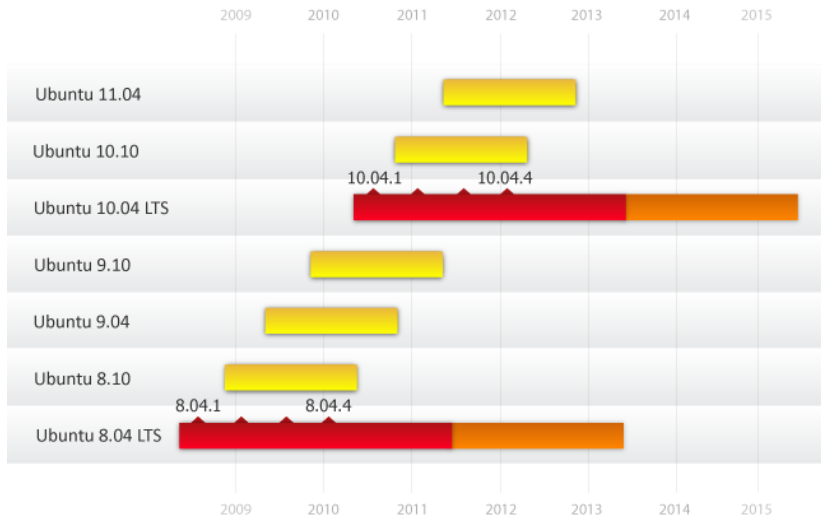
Advantages




- ▶ Supported until April 2011 for Desktops and 2013 for Server
- ▶ Next LTS-Release in April 2010
- ▶ Commercial support available
- ▶ Huge number of packages ready to install
- ▶ Very easy backporting for single applications
- ▶ Smooth upgrade to newer Versions

Disadvantages

- ▶ End-user-centered Distribution
- ▶ No guaranteed SL-compatibility (except LSB-conforming software)

Ubuntu Lifecycle



 LTS Desktop and Server  Standard Release  Point Release  LTS Server

Incompatibility between RHEL5 and older distributions

- ▶ The default hash-tables produced by the linker have changed to speed up relocation.
- ▶ Default is “gnu” on RHEL5 with standard compiler, incompatible with older systems.
- ▶ On Ubuntu, default is “both”
- ▶ SLES10 does not know about gnu hash maps

Simple example compiled on SL5.2

```
#include <stdio.h>
int main() { printf("Hello World\n"); }
```

compiled with

```
gcc -o hello hello.c
gcc -Wl,--hash-style=both -o hello.both hello.c
gcc -Wl,--hash-style=sysv -o hello.sysv hello.c
```

Incompatibility between RHEL5 and older distributions

Output on SL 5.2 and Ubuntu 8.04/8.10

```
[pcilc9:~]$ ./hello
Hello World
[pcilc9:~]$ ./hello.sysv
Hello World
[pcilc9:~]$ ./hello.both
Hello World
```

Output on SLES10 (Garching)

```
[at01:~]$ ./hello
Floating point exception
[at01:~]$ ./hello.sysv
Hello World
[at01:~]$ ./hello.both
Hello World
```


Current Status

- ▶ Server running on **fai.mpp.mpg.de**
- ▶ 3 Clients installed and managed: **pclh1-2**, **pcilc9** and **pcilc12**
- ▶ Fully integrated into MPP-Services
- ▶ Belle-Library and ILC-Software is running
- ▶ ATLAS installed but still problems validating. Various successful user-tests

Server

Install- and update-server is currently running in a virtual machine on **pcilc12** and is accessible as **fai.mpg.mpp.de**

Configuration

- ▶ Plain Ubuntu 8.04 with fai-packages from <http://www.informatik.uni-koeln.de/fai/> (3.2.16)
- ▶ Runs apt-proxy to act as sole package-server for all clients
- ▶ Serves NFSROOT for Installation and Fai-config for install/updates via nfs
- ▶ tftpd-daemon serves PXE-boot image for network-boot
- ▶ Server not yet integrated into NIS/nfs
- ▶ Module e1000e was added to NFSROOT-initrd (needed by DELL Optiplex 755)

Clients

The Clients are fully integrated into the network-infrastructure

Software includes

- ▶ Standard Ubuntu Desktop with gnome (default), kde and xfce4
- ▶ Linux-Standard-Base 3.1
- ▶ Openoffice 3 (to read OOXML), Acrobat-Reader 8
- ▶ alpine 2.0, Thunderbird 2 with Lightning, Sunbird, Firefox 3
- ▶ GCC Versions 3.3.6, 3.4.6, 4.1.3 and 4.2.4
- ▶ G77 Version 3.4.6, GFortran Versions 4.1.3 and 4.2.4
- ▶ TeXLive 2007
- ▶ VirtualBox 2.1

Clients

Some simple Modifications were needed to adapt Ubuntu to our needs.

Local modifications

- ▶ Added openafs-modules-source as dkms-module
- ▶ Manually backported some packages from newer Ubuntu: root-system (5.18), nvidia-glx (180.22), flashplugin-nonfree (10.0.15)
- ▶ Added Tivoli Storage-Manager (using alien)
- ▶ Custom debian-packages for ILC/BELLE-Software (managed by Kolja Prothmann)
- ▶ Disabled fast-user-switch-applet (Bug #203217)
- ▶ Dynamically add desktop-groups on login
- ▶ System wide alpine-configuration using IMAP (courtesy Stefan Kluth)



MAX-PLANCK-GESELLSCHAFT

Werner-Heisenberg-Institut für Physik

Username:

Options

Session

Language



Configuration

- ▶ Possible to serve different architectures with one config-space (multiple NFSROOTS)
- ▶ The current configuration is managed via git and can be checked out using

```
git-clone /home/pcilc9/ritter/faiconfig
```
- ▶ Upon *Installation*, FAI runs all scripts in directory `class`, all output will be used as class-names
- ▶ To reassign classes later, run “`fai -N softupdate`”
- ▶ After bootstrapping System, fai will partition, install and configure according to assigned classes.

Each class can consist of...

Disk-config in `disk_config/$CLASSNAME`

Variables in `classes/CLASSNAME.var`

Packages in `package_config/$CLASSNAME` (can be installed or removed)

Debconf-values in `debconf/$CLASSNAME`

Scripts in `scripts/$CLASSNAME`

Files in `files/$FILENAME/$CLASSNAME`

Hooks in `hooks/$HOOK.$CLASSNAME` to run specific scripts before certain install-tasks

- ▶ In case of conflict: Last assigned class wins.

Configuration classes

- FAIBASE** Base-class which contains tasks/packages common to every client installed. Almost unmodified.
- DHCP** Indicates that the network will be configured using dhcp. Not touched at all.
- MPPNET** Services common to all machines in the MPP-Network: NIS, nfs, AFS, mail, cups, backup and ssh.
- LAPTOP** Install things commonly found on Laptops, for example Bluetooth. (Not tested yet)
- WORKSTATION** Packages common to all Desktop-machines, including X11, gnome, kde, root etc.
- ILCBELLE** Contains specific software for the ILCBELLE-group: Geant4, clhep, cernlib, Belle-Library, ilcinstall.
- LAST** Generic class to do some clean-up.

Hardware-specific classes

- BIGMEM** For machines with ≥ 4 GiB RAM, the “server”-kernel is installed.
- NVIDIA** For machines with NVidia-graphics-card, installs binary driver and configures X11 to use it.
- ATI** Same as NVIDIA for ATI-cards.
- I386** Install normal x86-Architecture, not touched.
- AMD64** Install 64bit-System, not touched, needs additional work.

Additional, each client is assigned his hostname as a class to allow host-based configurations (currently `xorg.conf` for Dualhead).

Todo/wishlist

Todo

Things, which are not finished or tested include

- ▶ Manage faiserver using FAI
- ▶ Use UUIDs in `/etc/fstab`
- ▶ Laptop-class to install laptops/virtual machines without `nis/autofs/tivoli`
- ▶ Nvidia-Cards older then GeForce 6000
- ▶ Policy concerning virtual machines

Todo/wishlist

Wishlist

These items would be great to have, but are not necessary

- ▶ Central Wiki-Server
- ▶ Firefox/Konqueror home page with first steps for new users
- ▶ LDAP-Server to serve addresses for mail-client
- ▶ Central CUPS printing-server
- ▶ Thunderbird autoconfig
- ▶ Maybe internal instant-messaging (depends on LDAP)

Experiences

- ▶ Overall very straightforward
- ▶ Some subtleties writing scripts (`$target`, `$ROOTCMD`)
- ▶ e1000e module gave quite some headache
- ▶ Almost “out of the box” setup