

Automating Linux Installations at CERN

Lionel Cons, German Cancio, Philippe Defert, Mark Olive, Ignacio Reguero, Cedric Rossi

CERN, European Laboratory for Particle Physics, Geneva, Switzerland

Abstract

An automated installation service is being used to install the Linux farms in the CERN computing centre and part of the desktop systems since end of 1998. The system is based on the Kickstart software from the RedHat distribution. The configurations for 370 systems are currently registered. The system has been adapted to automatically install CERN's Standard Unix Environment (SUE). A number of Kickstart problems have been solved or circumvented. Work is currently done on a more general system suitable for desktop PCs.

Keywords: System Administration, Automatic installations, Linux, Kickstart

1 Introduction

Linux is becoming a key operating system at CERN. Some 900 PC's running Linux are currently installed on the site, and some 50 new machines are added monthly. A big percentage are part of data processing farms placed at CERN's computing centre, the rest being used as servers or desktop workstations. With the freezing of other UNIXes and the data and CPU requirements imposed by new experiments like COMPASS and specially LHC's, the amount of Linux PC's is expected to increase dramatically over the next years. Several problems will have to be overcome, one being the system installation. A typical Linux installation takes around 20 minutes and requires human intervention in order to enter configuration data and monitor the execution. This is too inefficient and error prone with hundreds or even thousands of systems to install or reinstall. An automated procedure should therefore be used.

2 Requirements

An automated installation procedure should fulfill several conditions, namely:

- *unattended run*: Installations should be run without (or minimal) human interaction.
- *concurrent installations*: Multiple systems can be installed in parallel.
- *scalable*: New clients (and OS releases) can be integrated easily.
- *flexible*: A feature-rich configuration should be offered.
- *customisable*: It should be possible to adapt the system to a specific environment and settings.

Examples of automated installation systems include *JumpStart* for Solaris, *Ignite-UX* for HP-UX and *Kickstart* for RedHat Linux [4], the latter being the Linux distribution used at CERN.

3 Kickstart

Kickstart is the automated installation method for Linux developed by RedHat. A system administrator creates a file (called a Kickstart file) which describes the configuration and the required

setup.

This file contains the required disk partition setup, the packages to be installed and other configuration items like the language, time zone, mouse and keyboard setup etc.

A Kickstart file may also contain a post installation procedure, in which a list of shell level commands may be specified which are executed once the normal installation process has finished.

The target machine is bootstrapped with a down sized kernel and a minimised root file system, the installation server is then NFS-mounted. The installation program starts and is driven by the contents of the Kickstart file.

The Kickstart file may be stored on the installation diskette or in a central repository accessed using the network via BOOTP and NFS. Often, a single Kickstart file can be used for installing multiple machines with similar configuration.

4 Architecture

Figure 1 depicts the architecture used at CERN for Kickstart. The BOOTP server provides the network information to the client and the location of its Kickstart file. The NFS servers host the Kickstart files and the operating system distribution.

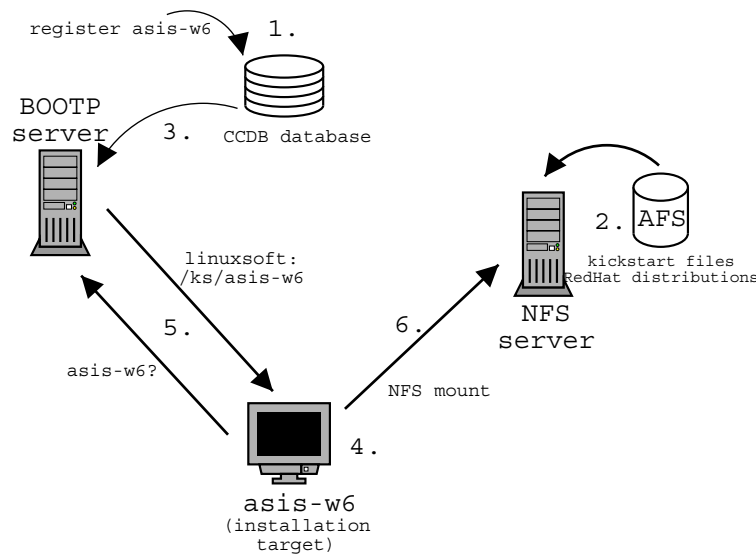


Figure 1: Kickstart architecture at CERN.

The steps to be followed for a complete installation are:

1. The machine to be installed has to be registered into the CERN network database (CCDB) in order to obtain a IP address and a DNS name.
2. After writing the Kickstart file for the machine, it has to be placed on the NFS servers. It is also possible to use a symlink to an already existing Kickstart file.
3. The machine is registered into the BOOTP server using a wrapper which looks up in CCDB all the necessary information (DNS name, IP address, ethernet address).
4. The machine is rebooted and the installation program loaded. This may be done using a boot installation diskette or booting from a partition containing the boot system. TFTP might also be used with boot PROMS.
5. Using BOOTP, the network information and the location of the Kickstart file are retrieved.

6. The Kickstart file is copied from NFS, and the RedHat distribution is mounted via NFS from the installation server.
7. The installation is carried out, including the post installation procedure.

Currently, the kickstart files and the RedHat distribution are stored on AFS¹ [2] and made accessible via NFS on the installation server using an AFS to NFS exporter. This allows us to profit from AFS's improved backup, security, scalability and accesibility features.

5 Configuration Example

At CERN, a template file is provided for standard configurations (see figure ??).

This standard configuration includes the installation of SUE, CERN's Standard Unix Environment² [1]. The %post procedure adds a system startup file and records the SUE features to be used (like network and mail setup, access to the ASIS³ [3] repository). During the first reboot after installation, the AFS client gets installed and SUE bootstrapped.

```
# language: default en_US
lang en_US

# network information
network --bootproto static --ip 1.2.3.4 --netmask 255.255.0.0 \
--gateway 1.2.1.1 --nameserver 1.1.16.5

# the NFS install server.
nfs --server linuxsoft --dir /redhat/cern/6.1

# US keyboard.
keyboard us

# clear all linux partitions
clearpart --linux

# define partitions
part swap      --size 256
part /         --size 500 --grow
part /var      --size 40
part /afscache --size 100
part /home     --size 1000

# do full install
install

# mouse
mouse generic3ps/2

# timezone
timezone --utc Europe/Paris

# X configurator (broken!!)
xconfig --server "Mach64" --hsync "30-60" --vsync "50-120"

# encrypted root password.. Change ASAP!
rootpw --iscrypted !@#*4

# no NIS, shadow passwd..
auth

# put LILO on the MBR (may cause problems on dual boot systems)
lilo --location partition

# reboot automatically after installation.
reboot

# package group to take: CERN recommended setup.
%packages
@ CERN Recommended Setup

# post install procedure
%post

# fix /etc/hosts
# replace <MYMACH> by your machine's DNS name
```

¹AFS, the Andrew File System, is a distributed, location-independent wide area file system. It was originally developed at Carnegie-Mellon University and is now comercialised by Transarc/IBM.

²SUE is a set of software components, configuration files and utility programs which together form a ready-to-use, site-customized Unix system.

³ASIS is a multiplatform and configurable software distribution system developed at CERN.

```

cat >/etc/hosts <<EOFhosts
127.0.0.1          localhost.cern.ch localhost
1.2.3.4           nuke.cern.ch nuke
EOFhosts

#
# Linux SUE installation on reboot
#
cat >/etc/rc.d/init.d/ks-post-install <<EOFinstall
#!/bin/sh -x

# execution log on tty 7
exec >/tmp/ks-post-install.log 2>&1
tail -f /tmp/ks-post-install.log >/dev/tty7 &

# which SUE profiles to install
SUEPROF="cern ssh sshd sudo"

# who should receive root's email
RESPONSIBLE='Nuke.Dukem@cern.ch'
echo "\${RESPONSIBLE}" > /root/.forward

# AFS installation
# use AFS 3.6 client
export AFSWS=afs36.t
#
echo '\$AFS_MASTER_CELLSERVDB = "","' > /usr/sue/etc/config/afs
echo "Installing AFS..."
/usr/sue/etc/sue.install afs

# run SUE
echo "running SUE bootstrap..."
/afs/cern.ch/project/sue/new/i386_redhat61/dist/etc/sue.bootstrap
echo "Installing \${SUEPROF} SUE profile"
echo "This may take a while"
/usr/sue/etc/sue.install -uselogs \${SUEPROF} </dev/null

# add AFS users <USER1>,<USER2>,...
/afs/usr/local/etc/ccdbuser nuke
/afs/usr/local/etc/ccdbuser foo
/afs/usr/local/etc/ccdbuser bar

# cleanup
rm -f /etc/rc.d/init.d/ks-post-install
rm -f /etc/rc.d/rc3.d/S*ks-post-install

# reboot once more
reboot
exit 0
EOFinstall

chmod +x /etc/rc.d/init.d/ks-post-install
ln -s /etc/rc.d/init.d/ks-post-install /etc/rc.d/rc3.d/S86ks-post-install

```

6 Problems and Limitations

Kickstart has been released recently and is still under development. Therefore, features are still missing, like:

- *Partition configuration:* There's no way to define the location of the partitions, but only their mount points and size. Specially when using multiple disks systems, this obliges to partition them by calling `fdisk` in the post install section.
- *Command syntax:* The command syntax changes from one RedHat release to another, which forces to edit and adapt the kickstart file on each operating system upgrade.

In the current release (for RedHat 6.1), there are bugs which haven't been fixed by RedHat, the most important being:

- The root password encryption doesn't work (fixed for the CERN distribution)
- The `/etc/hosts` file is generated incorrectly, which has to be fixed in the post install section (see the example file).
- The X configuration tool generates sometimes incorrect setups.

7 Future plans

Currently at CERN, more than 370 systems have been installed using Kickstart, most of them in the CERN computing centre. RedHat releases 5.1, 6.0 and 6.1 can currently be installed using Kickstart.

As more and more new systems are getting installed using this mechanism, some scalability issues have to be addressed.

When doing massive concurrent installs, the usage of an AFS to NFS exporter leads to high bandwidth consumptions as files have to be transferred first from the AFS server to the installation server and then to the installation target. A cluster of NFS installation servers is being set up and tested experimentally. The RedHat distributions are stored on local disks on those servers and are synchronised periodically with the master distribution stored in AFS; each installation target will access its nearest install server.

Another problem is Kickstart's complexity. As the configuration of the installation parameters and scripts is still non-trivial, it is done by few qualified persons. Specially for end users with a desktop system, an user-friendly front end is required. This is also useful for the installation of large PC farms or clusters, where multiple configurations have to be managed.

We are currently developing a client/server system, from which end users will be able to configure and generate a Kickstart configuration file adapted for their machine or machine group and the OS release they want to install. This system will offer users a configuration template repository, both for standard desktops and clusters, which can then be customised using a Java-based GUI. For specific setups, expert modes will be available on which extended control will be enabled. A database will be mediating between the user interface and the kickstart system, thus easing modifications and updates of the configuration of the target systems.

References

- 1 CERN IT division, "SUE Definition Document", CERN, European Organisation for Nuclear Research, <http://wwwinfo.cern.ch/dis/sue>, 1995.
- 2 Transarc Corporation, "The AFS File System in Distributed Computing Environments", <http://www.transarc.com>, 1996.
- 3 Ph. Defert et al., "ASIS, manage and distribute software in the HEP community", CHEP'97, European Organisation for Nuclear Research, <http://wwwinfo.cern.ch/dis/asis>, 1997.
- 4 RedHat Corporation, "RedHat Linux reference guide", <http://www.redhat.com>, 1999.