

# Report on Condor Installation

John Kewley

July 6, 2004

## Summary

This document provides some notes on the multitude of Condor releases, release formats and installation possibilities for the Linux and Windows platforms.

## 1 Condor release strategy

Condor has 2 release streams: Stable and Development. The Stable releases have an even 2nd digit in their release number (6.4.x, 6.6.y), while that of the Development releases are odd (6.5.x, 6.7.y).

Stable releases are fully tested and are updated only for bug fixes and to support new platforms. Development releases, on the other hand, contain work under test. These should only be used if you need to test latest functionality, or functionality that is not supported on your platform in the stable release. It is recommended that different development releases are not mixed in the same pool; likewise the mixing of development and stable releases should be avoided.

## 2 Linux releases

On the Condor download page, there are a variety of Linux releases of Condor in the current Stable release series. Other Linuxes are also likely to work, but are not officially supported. For a given Linux, you must match the architecture with the kernel and glibc version numbers of your Linux flavour in the list below to find out the correct version to install.

**Intel x86 :**

- **Redhat 7.1, 7.2, 7.3 (Linux 2.4.x, glibc 2.2)**
- **Redhat 8 (Linux 2.4.x, glibc 2.2)**, note that rpm 4.2 or later is required to use the rpm release
- **Redhat 9 (Linux 2.4.x, glibc 2.3)**

**Alpha :**

- **RH 7.1, 7.2, 7.3 (Linux 2.4.x, glibc 2.2)** — vanilla only

## Intel Itanium :

- **RH 7.1, 7.2, 7.3 (Linux 2.4.x, glibc 2.2)** — vanilla only
- **SuSE Linux Enterprise Server 8.1 (Linux 2.4.x, glibc 2.3)**

Note that the above list is correct for version 6.6.5 at the time of writing this document, also that where releases are listed as “vanilla only”, then that means that only the vanilla universe is supported.

## 2.1 Static vs Dynamic Linking

For most supported UNIX platforms, both a static and a dynamically-linked version of Condor are available.

Guidelines for the 6.6.x series state that “On Linux, you should use the dynamically linked version”, whereas for the 6.4.x series the advice is that “You should always use the statically linked version if it is available”. This seems to be a significant change of tune/heart, and it would be interesting to find out why this recommendation has changed.

## 2.2 The 3 Installation Methods

As of the latest release series (6.6.x), the condor releases are also available in rpm format for Redhat Linux. As rpm is not supported on all Linuxes, the .tar.gz “tarball” format is also provided.

The installation documentation provides, in addition to the straightforward rpm installation, 2 methods of installing from the tarballs, making 3 installation possibilities.

rpm is the simplest installation method, but this does not lend itself to customisation as it puts the files in predefined places. It is also not available on all Linux. Latest documentation suggests running `condor_config` afterwards to move directories to where you want them, but it is not clear whether this would disrupt one of the significant advantages of the rpm install which is the possibility to completely remove the product trivially from the machine.

`Condor_install` is the “traditional” way to install Condor and asks a series of questions as to how you wish to install Condor. Some of the answers are non-trivial to answer. For instance, “Do all of the machines in your pool from your domain share a common filesystem?”, what do you answer if it has some common filestore, but not all? Ramifications of your answer are fairly significant, if filestore is shared, then installation assumes that ALL files are shared for all machines; if not, then it assumes that executables and I/O files need to be copied for each job to each machine on which the job will run. Each machine in the cluster also then needs to be installed separately.

The newer `condor_config` script acts as a halfway-house between the other two options by taking a variety of command line options (in the style of the configure scripts produced by `autoconf`), but these options do not seem to give sufficient flexibility to allow you to install Condor exactly where you want.

A presentation on a “view from the trenches” on these installation possibilities will be presented at UK Condor Week.

## 3 Windows

Until quite recently, for Windows machines, Condor was only advertised as supporting NT4. Now however it is supported (albeit limited functionality) on Windows NT4.0 Workstation and Server, Windows 2000 Professional and Server, Windows 2003 server and Windows XP Professional.

## 4 Installation activities

Condor has been installed on the following platforms at DL:

Windows 2000 Professional (ajr45vig + rjavig2)

Windows XP Professional (rjavig6)

Whitebox Linux (condor + rikster)

SuSE (tardis)

Redhat 7.2 (wk-ibm2 + its 8 nodes)

Redhat 9 (rjavig6)

The following universes have been tested: vanilla, standard and mpi. TBA - the vagaries of each will be separately listed. For example, both vanilla and MPI have changed in the latest release with respect to how file permissions are handled for file transfers (possibly as a result of my bug report) and so this behaviour needs retesting as results at first glance seem more tedious than before.

The following restrictions have been noted for MPI under Condor:

- MPICH versions 1.2.2, 1.2.3, and 1.2.4 using the `ch_p4` device, Condor does not support MPICH version 1.2.5, nor other MPICH devices
- No support for LAM
- Had to use dynamically linked version of 7.3 at one point, I believe this particular restriction has now been lifted.

### 4.1 Releases

The following condor releases have been tested on the wk-ibm2 cluster:

**6.4.7** : Initial learning about Condor

**6.5.x** : Enhanced support for MPI

**6.6.0** : Stable release containing 6.5 functionality

**6.6.1** : Full Condor support for RedHat 9

**6.6.2** : Many bugfixes

**6.6.3** :

**6.6.4** : Not installed

**6.6.5** :

**6.7.0** : Briefly tested on wk-ibm2, a problem I raised has been fixed in this version

## 5 Problems encountered

- Availability of machines: developing a condor pool is impossible without any hardware resources. I had various promises of resource which would be made available when the blade servers arrived. This eventually boiled down to 1 desktop and Rik's spare machine. I am still "on a promise" for more resources "when the new system admin guy comes".
- Access to root: To be fully effective, Condor really needs to be installed as root. As access to root passwords here at DL is limited, and the current custodians had other work to do, there were significant delays awaiting a free slot for assistance in system administration. This should be alleviated when our dedicated sysadmin arrives.
- Firewalls: I spent a week testing, analysing and checking configuration parameters to discover that one of the desktops had a firewall on it preventing the other machine getting access, likewise it took some time to work out that the wk-ibm2 machine was OUTSIDE our firewall rather than inside as I had expected.
- Dependencies for MPI: as mentioned above, there are some quite specific information as to what MPI support is given to condor. In fact one response I got from Condor about their support was that "no one here on the condor-staff has been given time to maintain the MPI support, so i fear it might not work very well at all right now". This came as a surprise after I had spent some time working on it!