

# **A coLinux/Condor Computer Cluster at the University of South Alabama**

**C. M. Jenkins**

Department of Physics  
University of South Alabama

**Horst Severini**

Homer L. Dodge Department of Physics & Astronomy  
The University of Oklahoma

**Joshua Alexander and Henry Neeman**

OU Supercomputing Center for Education and Research (OSCER)  
The University of Oklahoma

**Chris Franklin**

OU Information Technology  
The University of Oklahoma

**Representing Distributed Organization for Scientific and Academic Research  
(DOSAR)**

This work funded in part by DoE grant  
DE-FG02-96ER40970

# DOSAR

- **D**istributed **O**rganization for **S**cientific and **A**cademic **R**esearch.
- **<http://www.dosar.org>**
- “A ‘grass-roots’ grid organization that focuses on community and campus grids and promotes a wide range of interdisciplinary and educational activities within the organization and member institutions.”
- Member institutions:
  - Iowa State University
  - University of Johannesburg
  - Langston University
  - Louisiana State University (CCT)
  - Louisiana Tech University
  - University of Mississippi
  - University of Oklahoma
  - Universidade Estadual Paulista (UNESP) (SPRACE GridUNESP)
  - **University of South Alabama**
  - Susquehanna University
  - University of Texas at Arlington

# Condor

- A project that supports high throughput computing
  - <http://www.cs.wisc.edu/condor>
  - Open source software: condor
- A condor cluster:
  - Condor central manager
    - Monitor the condor cluster
    - Submit jobs to the condor worker machines
    - Monitor/manage jobs running on the cluster
  - Condor worker machines (many machines)
    - Accept and execute jobs from the condor central manager
    - Return results back to the condor central manager after job completion

# coLinux

- Cooperative Linux
  - <http://www.colinux.org>
  - Open source software
  - A native port of the Linux kernel to Windows
  - Runs on top of Windows OS
  - Allows the user to run linux operating system

# coLinux/Condor Package

- Developed by OU Supercomputing Center for Education & Research (OSCER) and the University of Nebraska.
  - <http://www.oscer.ou.edu>
- The package allows the user to setup a linux operating system on a Windows OS machine that runs the condor program.
- The objective is to allow the Windows PC's to become worker nodes in a condor cluster.
- Use Windows-based PC's in student computer labs.
- **Harvest idle CPU cycles – late at night when machines are not used.**

# Local Platforms for the USA

## coLinux/Condor Cluster

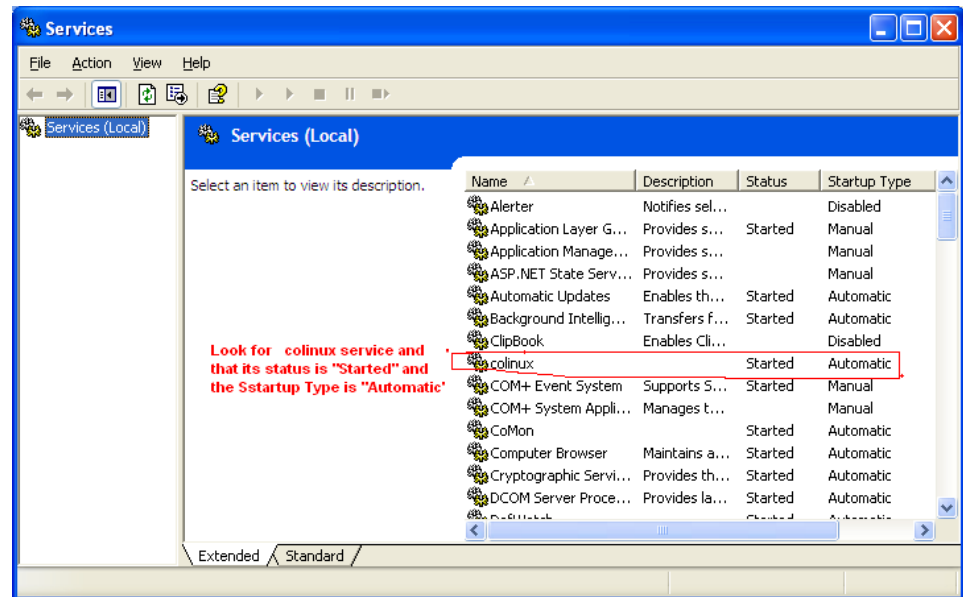
- **Condor central manager**
  - Dell Precision 350
  - Pentium 4 - 2.80 GHz
  - 512 Mb RAM
  - 120 Gb disk space
  - Dual Boot OS:
    - Red Hat Linux 8.0
    - Scientific Linux 3.0.9
- Run condor 7.0.4 with slc 3.0.9 OS
- **coLinux/Condor worker nodes**
  - Four PC's in Advanced/Modern Physics instructional laboratory
    - Received Physics Department Chair's permission to use before implementing.
    - **These are primarily instructional computers!**
  - Dell Dimension 2400
  - Pentium 4 – 2.8 GHz
  - 512 MB RAM
  - 33.6 Gb disk space
  - Window XP Operating System

# Installing the coLinux/Condor package

- Follow instructions on the OSCER web site:
  - [http://www.oscer.ou.edu/CondorInstall/condor\\_colinux\\_howto.php](http://www.oscer.ou.edu/CondorInstall/condor_colinux_howto.php)
- **Major steps (some details missing):**
  - Download the coLinux/Condor file
    - Unzip the file
  - Configure the Condor configuration files
    - More on this later
  - Modify the first.sh script file
  - Modify the sethostname file
  - Copy the script on the `condor_colinux_howto.php` page.
  - Save this file locally as `condor_script.bat` or equivalent name and modify
    - Set the size of physical RAM
    - Reserve the size of disk area for condor to run in
    - Modify for the local area connection for the ethernet connection
  - Modify the boot.ini file
  - Dell computers with recovery partition:
    - partition(1) must be changed to partition(2)
    - **This is very important – so as not to write over the boot sector!**
  - Execute the `condor_script.bat` file:
  - Bring up DOS window
  - `cd c:\condor\colinux3`
  - `Condor_script.bat >> nodeNameCondorInstall.Log.txt`

# Starting the coLinux/Condor System

- Reboot the computer
- **Is coLinux running?**
- Look for condor service
  - Start Button →
  - Control Panel (Classic) →
  - Administrative Tools →
  - Services
- Service is present and started.



- The OSCER coLinux/Condor package installs
  - Fedora Core release 6 (Zod)
  - Condor 6.8.4



# Setting up the condor central manager

- Use Dell Precision 350
  - Pentium 4 - 2.80 GHz
  - slc 3.0.9
- Follow instructions at condor web site
  - <http://www.cs.wisc.edu/condor>
- Download rpm and install rpm
  - Make condor account
  - Modify the configuration files
    - More on this later
  - Setup condor environment
  - Start up condor
    - `%source /opt/condor/condor.sh`
    - `%<release_dir>/sbin/condor_master`
  - See if condor is running:
    - `%ps -ef | egrep condor`
    - Look for condor processes

```
$ ps -ef | egrep condor
condor 4923 1 0 Jan21 ? 00:00:20 /opt/condor-7.0.4/sbin/condor_master
condor 4960 4923 0 Jan21 ? 00:02:57 condor_collector -f
condor 4991 4923 0 Jan21 ? 00:01:26 condor_negotiator -f
condor 4992 4923 0 Jan21 ? 00:00:00 condor_schedd -f
condor 4993 4923 0 Jan21 ? 00:20:10 condor_startd -f
root 4994 4992 0 Jan21 ? 00:00:00 condor_procd -A /tmp/condor-
lock.orion0.453315811305405/procd_pipe.SCHEDD-S 60 -C 502
root 30139 4993 0 Jan28 ? 00:00:00 condor_procd -A /tmp/condor-
lock.orion0.453315811305405/procd_pipe.STARTD-S 60 -C 502
root 31740 4797 0 08:05 ? 00:00:00 sshd: condor [priv]
condor 31742 31740 0 08:06 ? 00:00:00 sshd: condor@pts/0
condor 31744 31742 0 08:06 pts/0 00:00:00 -tcsh
condor 31767 31744 0 08:06 pts/0 00:00:00 ps -ef
```

- condor 7.0.4

# Condor Configure Files

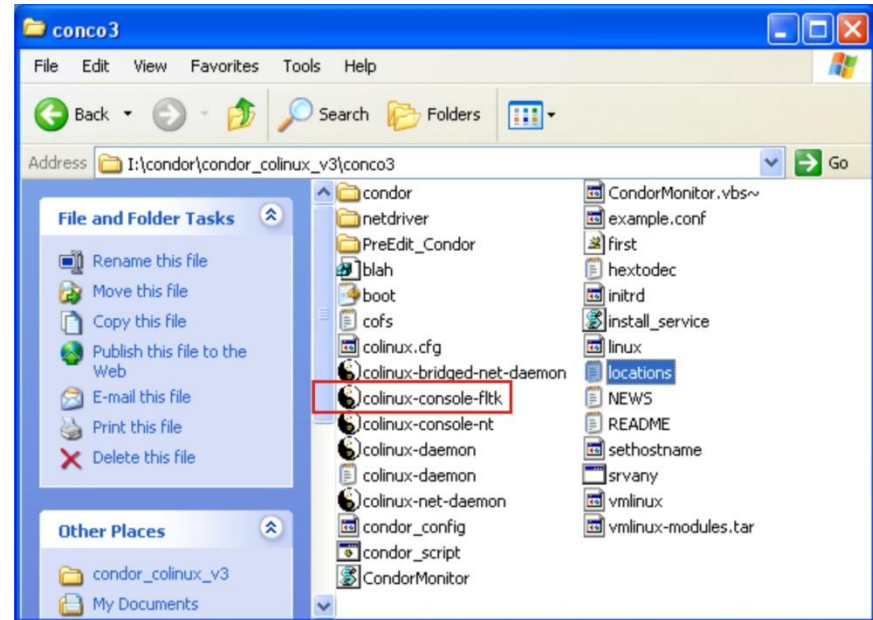
- condor config files must be setup properly for the condor cluster to work.
- <release\_dir>/etc/condor\_config
- condor\_config variables to define:
  - LOCAL\_DIR
  - FLOCK\_FROM
  - FLOCK\_TO
  - HOSTALLOW\_READ
  - HOSTALLOW\_WRITE
- condor\_config.local variables to define:
  - CONDOR\_HOST
  - CONDOR\_ADMIN
  - UID\_DOMAIN
  - COLLECTOR\_NAME

# Initial Problems with the Cluster

- The coLinux/condor worker nodes were missing.
- Set up a second slc 3.0.9 node
- Remove possible coLinux/condor problems
  - Received some help from USA Academic Computing
  - Resolving network issues
- Trouble shooting greatly helped by locating condor log files.
- /opt/condorVersion/local.node/log/collectorLog
- Helped in resolving condor issues from network issues
  - Proper setup of condor config files on both machines
    - Add node on host.allow
- Second node appeared on condor cluster
- Condor working!

# Resolving coLinux/condor Issues

- On the Windows PC
- Start up a linux console
- “click” **coLinux-console-ftp** icon
  - Login
  - Startup condor
  - Look at the condor error log files
    - Separate network problems from condor problems
    - Correct problems in the condor config files
  - Eventually found:
    - coLinux IP address different than Windows IP address
    - No hostname defined for coLinux!



- USA DHCP only assigned IP address
- U of Oklahoma also assigned node name
- **Solution: configure each PC:**
  - **Permanent IP address for coLinux PC**
  - **Configure with hostname**

# The coLinux/condor Cluster

- Nodes appear with `condor_status` command
- Cluster is stable

```
$ condor_status
```

Name	OpSys	Arch	State	Activity	LoadAv	Mem	ActvtyTime
ilb00500.condor.us	LINUX		INTEL	Unclaimed Idle	0.000	250	0+01:50:20
ilb00501.condor.us	LINUX		INTEL	Unclaimed Idle	0.000	250	0+02:29:02
ilb00502.condor.us	LINUX		INTEL	Unclaimed Idle	0.000	250	0+00:00:54
orion.physics.usou	LINUX		INTEL	Unclaimed Idle	0.020	499	0+00:50:04

```
Total Owner Claimed Unclaimed Matched Preempting Backfill
```

INTEL/LINUX	4	0	0	4	0	0	0
Total	4	0	0	4	0	0	

# Testing the Cluster

- Wrote a small C++ benchmark program
  - currentHost.cc
  - Loop accessing the computer clock
  - Prints out time
- Use “standard universe”
- Must build the program to link to the condor libraries.
  - %condor\_compile CC -o currentHost currentHost.cc
  - The job is controlled on the condor cluster with a condor command file
    - currentHost.cmd
    - This includes the number of jobs to submit
- To submit the condor jobs
- %condor\_submit currentHost.cmd

# Test Job Output

- The test job is called currentHost
- Output from the test job run on two nodes.

```
Max = 15000 | Modulo = 1500
Date = 2010Jan28_22_13_58
Current Host: orion
Error getting MYHOST
Current Directory: /orion2/condor/CurrentHost2010Jan28A
Error getting CONDOR_HOST
Error getting COLLECTOR_HOST
Error getting FULL_HOST_NAME
CONDOR_SCRATCH_DIR: /opt/condor-6.8.4/local.ilb00500/execute/dir_18805
last = 3
length of templn = 50
tempName = ilb00
Error getting _CONDOR_SLOT
m = 0 Time = 0.0000e+00, rtime = 2.0000e-02
m = 1500 Time = 0.0000e+00, rtime = 7.0000e-02
m = 3000 Time = 0.0000e+00, rtime = 1.0000e-01
m = 4500 Time = 0.0000e+00, rtime = 1.3000e-01
m = 6000 Time = 0.0000e+00, rtime = 2.2000e-01
m = 7500 Time = 0.0000e+00, rtime = 2.7000e-01
m = 9000 Time = 0.0000e+00, rtime = 2.7000e-01
m = 10500 Time = 0.0000e+00, rtime = 3.8000e-01
m = 12000 Time = 0.0000e+00, rtime = 4.1000e-01
m = 13500 Time = 0.0000e+00, rtime = 4.7000e-01
```

```
Max = 15000 | Modulo = 1500
Date = 2010Jan28_22_05_17
Current Host: orion
Error getting MYHOST
Current Directory: /orion2/condor/CurrentHost2010Jan28A
Error getting CONDOR_HOST
Error getting COLLECTOR_HOST
Error getting FULL_HOST_NAME
CONDOR_SCRATCH_DIR: /opt/condor-6.8.4/local.ilb00501/execute/dir_7418
last = 3
length of templn = 49
tempName = ilb00
Error getting _CONDOR_SLOT
m = 0 Time = 0.0000e+00, rtime = 3.0000e-02
m = 1500 Time = 0.0000e+00, rtime = 5.0000e-02
m = 3000 Time = 0.0000e+00, rtime = 1.2000e-01
m = 4500 Time = 0.0000e+00, rtime = 1.7000e-01
m = 6000 Time = 0.0000e+00, rtime = 2.2000e-01
m = 7500 Time = 0.0000e+00, rtime = 2.8000e-01
m = 9000 Time = 0.0000e+00, rtime = 3.6000e-01
m = 10500 Time = 1.0000e+00, rtime = 4.1000e-01
m = 12000 Time = 1.0000e+00, rtime = 4.1000e-01
m = 13500 Time = 1.0000e+00, rtime = 5.1000e-01
```

# Future Plans

- Need to build programs that use the CERN libraries.
  - Root
- Need shared disk space to run in the “vanilla universe”
  - Programs do not have to link to the condor libraries.
- Need to upgrade the linux kernel to SL 5 release
  - Will do this after it is released by the University of Oklahoma
- Need to upgrade the condor version