

AccessionIndex: TCD-SCSS-T.20121208.095

Accession Date: 8-Dec-2012

Accession By: Dr.Brian Coghlan

Object name: csTCDie Beowulf Cluster

Vintage: 1998

Synopsis: Departmental cluster using 100Mbps Ethernet as interconnect, the second cluster constructed in the Department of Computer Science, Trinity College Dublin.

Description:

In the 1994 Thomas Sterling and Donald Becker of NASA created their *Beowulf* cluster [1] with multiple commodity nodes exchanging messages over a commodity interconnect. Beowulf clusters run Linux or another UNIX-lookalike plus messaging libraries such as MPI or PVM. This innovation prompted the building of this cluster.

The first Irish Beowulf cluster was constructed in 1997 by the School of Cosmic Physics in the Dublin Institute for Advanced Studies. The second Irish Beowulf cluster was constructed in the middle of the same year by the Department of Computer Science in Trinity College Dublin, a highly experimental configuration of 4-nodes that used a SCSI interconnect, the first of a number of clusters constructed by the department, some very production-oriented, others more adventurous, see elsewhere in this catalog.

The next cluster constructed in the department was a more straightforward Beowulf cluster that used 100Mbps Ethernet as interconnect. It was constructed with strong support from Prof.J.G.Byrne, for use by the Dept.Computer Science, Trinity College Dublin, but also to acquire knowledge of Beowulf clustering techniques.

This 16 cluster nodes were generic PC/AT-compatible machines with (**Pentium ?**) CPUs, **??MB** of memory, **??GB** of disk, interconnected with a 100Mbps Ethernet switch, all mounted on two custom-built shelving units. Account management was handled by primary and secondary NIS servers, and there were also a local DNS cache and an NTP time server, all mounted alongside UPSes below the cluster.

A common shared NFS-compatible filesystem was provided by a Windows-NT server with a RAID-5 subsystem using Vortex RAID controllers and running an Intersys NFS-to-NTFS service. This service was easy to administer but not bug-free. The RAID array was inherited from earlier research on a *Stable Disk* within the EU-funded FASST project, see elsewhere in this catalog.

A lot was learnt from the operation of this cluster. The initial support for parallel programming focussed on PVM, but MPI soon followed. Links with like-minded academics were also begun that continued through EU projects for the next 15 years. One of the most interesting episodes occurred when hackers outside the department but inside the university intensively attacked the cluster, initially bringing it to its knees within two minutes of a reboot. This forced the rapid evolution of a Linux `proxy-arp` firewall on a DEC Alpha machine pretending to be a Windows x86 machine (our thanks to Mark Gantly of DEC for quickly donating an Alpha server).

This cluster was superseded in 1999 by a cluster with firstly a ring-structured, then a second generation switched, SCI interconnect, see elsewhere in this catalog.

Apart from the RAID array, it was thought that no remnants of this cluster remained, but then two decades after decommissioning, the primary NIS server (cagraidsvr04, a Tatung “*pizza box*” PC with Cyrix MediaGXm 266MHz CPU, 32MB of memory, ST34321A 4.3GB IDE disk, 3.5” floppy disk and 24X CD-ROM drive) has been discovered (see figures below), and surprisingly found to be still operational and running Linux 2.0.6 (released on 12-Jul-1996 [2]), a really early Linux kernel only used in Redhat 3.x releases [3][4][5] (known to have been installed from a RedHat CD [6]). Given the early kernel, the original IDE hard disk has been cloned to an SSD, which has then been connected in place of the original hard disk so that the kernel is preserved working ‘live’ for posterity and to be available for detailed investigation in the potentially distant future. There are mail logs beginning on 13-Oct-1998 through to 2004. The disk contents have been found to be fully accessible, with what appear to be backups of the cluster software; digging into the OS and these backups may reveal more about the configuration of this early cluster as a whole.

The homepage for this catalog is at: <https://www.scss.tcd.ie/SCSSTreasuresCatalog/> Click 'Accession Index' (1st column listed) for related folder, or 'About' for further guidance. Some of the items below may be more properly part of other categories of this catalog, but are listed here for convenience.

Accession Index	Object with Identification
TCD-SCSS-T.20121208.095	csTCDie Beowulf Cluster, Departmental cluster using 100Mbps Ethernet as interconnect, the second cluster constructed in the Department of Computer Science, Trinity College Dublin, 1998.
TCD-SCSS-T.20121208.095.01	csTCDie Beowulf Cluster NIS server. Login name server from second cluster constructed in the department. 1998.
TCD-SCSS-T.20121208.095.02	csTCDie Beowulf Cluster NIS server. Backup of login name server IDE disk from second cluster constructed in the department. 1998.
TCD-SCSS-T.20150217.003	Prototype RAID, First RAID chassis built in the Dept.Computer Science, TCD, with five SCSI-1 disks and PC/AT power supply. PSU S/N: 7X705993, c.1990
TCD-SCSS-T.20121208.094	Experimental SCSI Cluster, 4-node prototype cluster using SCSI as interconnect, the first cluster constructed in the Department of Computer Science, Trinity College Dublin, and second cluster constructed in the Republic of Ireland, 1997.
TCD-SCSS-T.20141120.003	csTCDie Grid-Ireland SCI Cluster, 16-node cluster using 400MB/s SCI switched interconnect, the third cluster constructed in the Department of Computer Science, Trinity College Dublin, c.1999.
TCD-SCSS-T.20121208.097	VRengine, 9-node virtual reality engine using 600MB/s SCI 2-d toroidal interconnect, c.2005.
TCD-SCSS-T.20121208.098	csTCDie Grid Site Beowulf Clusters and Datastore, Complex of clusters & storage (1500 cores/600 TB) using 1Gbps Ethernet interconnect and 10Gbps backbone, participant in DataGrid, EGEE, EGI, and CERN LHC computing. From 2013 repurposed as SCSS Cloud, c.2009.
TCD-SCSS-T.20121208.106	csTCDie PS3 Cluster, Ten nodes from a 16-node Sony Playstation PS3 cluster plus build machine, using 1Gbps

TCD-SCSS-T.20121208.099	Ethernet interconnect and running Yellow Dog Linux, c.2009. csTCDie GPU Cluster, 64-core/32-GPU/16-node cluster using 1Gbps Ethernet interconnect, c.2011.
---	--

References:

1. Wikipedia, *Beowulf cluster*, see:
https://en.wikipedia.org/wiki/Beowulf_cluster
Last browsed to on 30-May-2023.
2. mirror.cs.msu.ru, *Linux Kernel Version History: 2.0 series kernels*, (table last updated on 30-Jan-2001 at 00:53:59 GMT), see:
<https://mirror.cs.msu.ru/oldlinux.org/Linux.old/docs/history/2.0.html>
Last browsed to on 30-May-2023.
3. Wikipedia, *Red Hat Linux*, see:
https://en.m.wikipedia.org/wiki/Red_Hat_Linux
Last browsed to on 30-May-2023.
4. OpenSkills, *RedHat: History of releases*, see:
<https://openskills.info/infobox.php?ID=1094>
Last browsed to on 30-May-2023.
5. Fedora Project Wiki, *History of Red Hat Linux*, see:
https://fedoraproject.org/wiki/History_of_Red_Hat_Linux
Last browsed to on 30-May-2023.
6. Brian Coghlan, *Personal recollection*, 30-May-2023.



Figure 1: csTCDie Beowulf Cluster NIS server, top three-quarter view



Figure 2: csTCDie Beowulf Cluster NIS server, front view



Figure 3: csTCDie Beowulf Cluster NIS server, rear view




Figure 4: csTCDie Beowulf Cluster NIS server, rear view closeup



Figure 5: csTCDie Beowulf Cluster NIS server, manufacturing label




Figure 7: csTCDie Beowulf Cluster NIS server, internal view, CPU and memory



Figure 8: csTCDie Beowulf Cluster NIS server, internal view, Ethernet adapter



Figure 9: csTCDie Beowulf Cluster NIS server, internal view, motherboard




Figure 10: csTCDie Beowulf Cluster NIS server, internal view, motherboard closeup




Figure 11: csTCDie Beowulf Cluster NIS server, internal view, memory




Figure 12: csTCDie Beowulf Cluster NIS server, internal view, partially obscured motherboard identification



Figure 13: csTCDie Beowulf Cluster NIS server, internal view, hard disk identification
Seagate ST34321A 4.3GB IDE hard disk, 5 heads, 17 sectors of 512 bytes, 33344 cylinders

/hda1: blocks 1-890,	37768+ blocks,	ID=83 (Linux)
/hda2: blocks 891-98725,	4158000 blocks,	ID=5 (ext)
/hda5: blocks 891-3914,	128488 blocks,	ID=82 (swap)
/hda6: blocks 3915-6938,	128488 blocks,	ID=82 (swap)
/hda7: blocks 6939-9962,	128488 blocks,	ID=82 (swap)
/hda8: blocks 9963-12986,	128488 blocks,	ID=82 (swap)
/hda9: blocks 12987-98369,	3628768 blocks,	ID=83 (Linux)

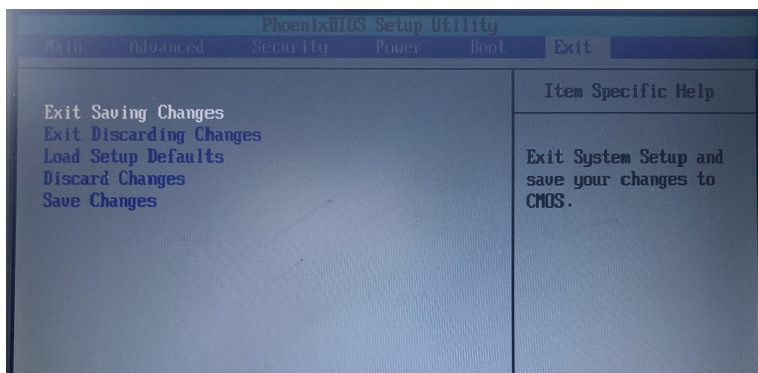
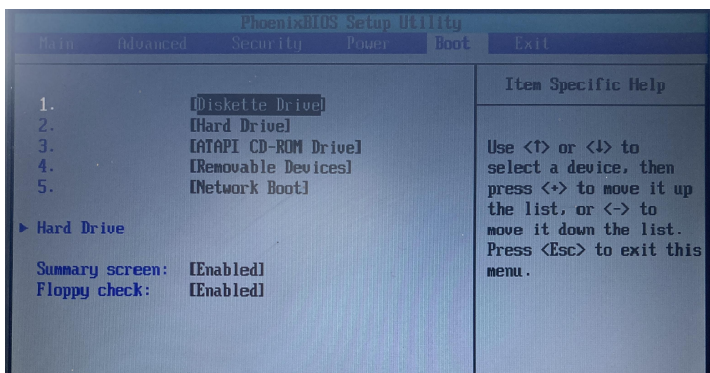
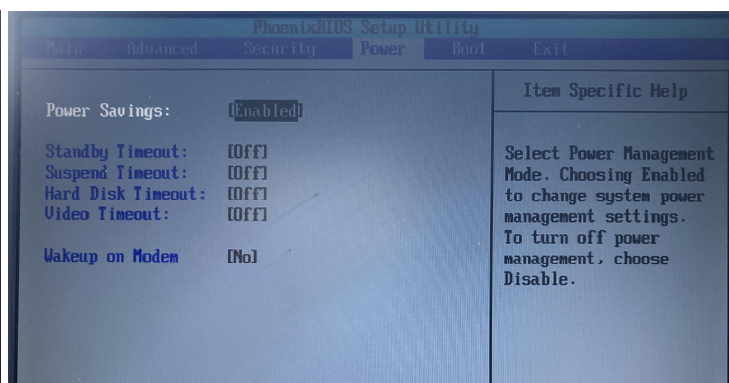
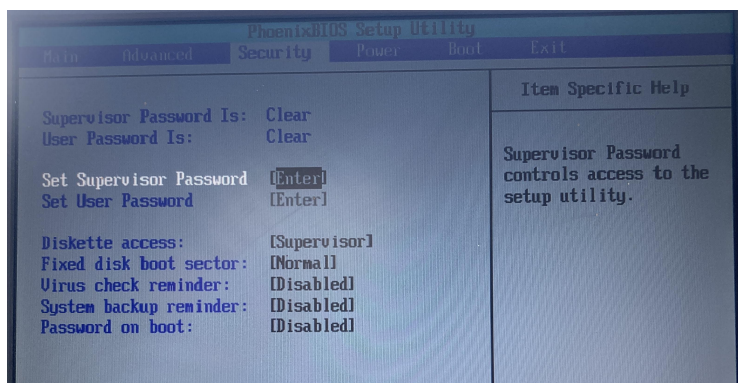
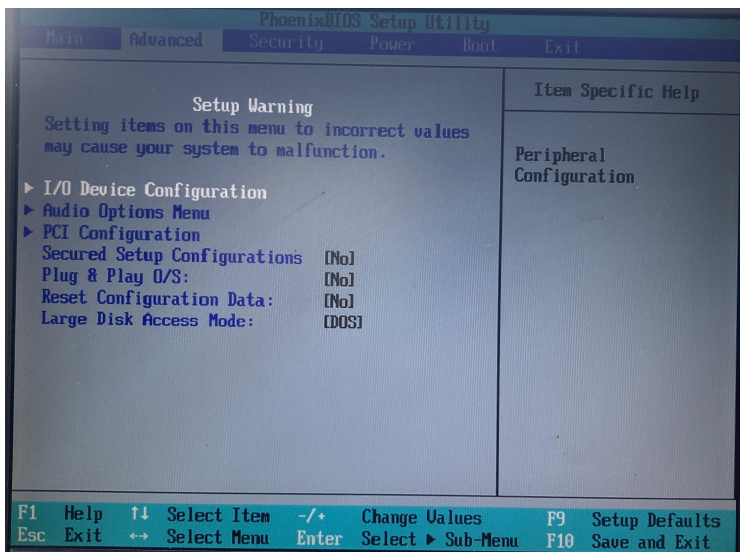
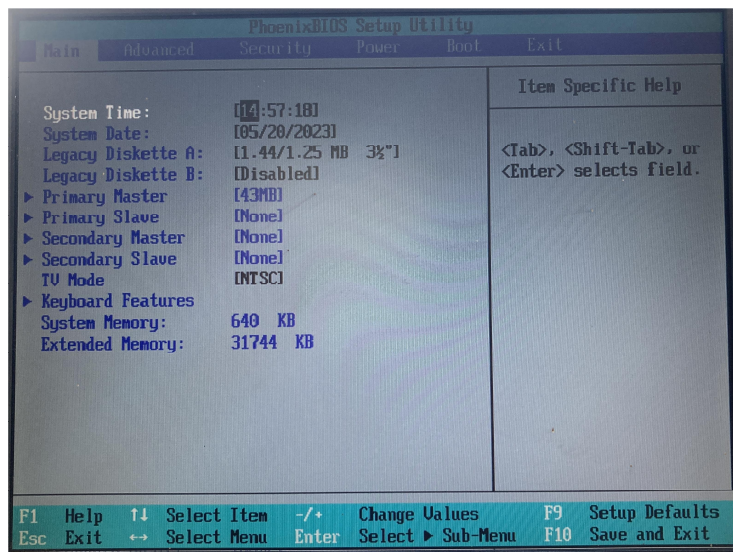


Figure 14: csTCDie Beowulf Cluster NIS server, BIOS screenshots 1-6


```

hd0 at 0x1f0-0x1f7,0x3f6 on irq 14
hd1 at 0x170-0x177,0x376 on irq 15
Floppy drive(s): fd0 is 1.44M
FDC 0 is a post-1991 82077
md driver 0.36.3 MAX_MD_DEV=4, MAX_REAL=0
scsi : 0 hosts.
scsi : detected total.
Partition check:
  hda: hda1 hda2 < hda5 hda6 hda7 hda8 hda9 >
UFS: Mounted root (ext2 filesystem) readonly.
INIT: version 2.74 booting
Activating swap partitions
Adding Swap: 128484k swap-space (priority -1)
Adding Swap: 128484k swap-space (priority -2)
Adding Swap: 128484k swap-space (priority -3)
Adding Swap: 128484k swap-space (priority -4)
hostname: cagraidsvr04.cs.tcd.ie
Checking root filesystems.
Parallelizing fsck version 1.12 (9-Jul-98)
[/sbin/fsck.ext2] fsck.ext2 -a /dev/hda9
/dev/hda9 has gone too long without being checked, check forced.
/dev/hda9: 64513/907264 files (0.8% non-contiguous), 1546183/3628768 blocks
Turning on user and group quotas for root filesystem
Remounting root filesystem in read-write mode.
Finding module dependencies...

```

Screenshot 1

```

Finding module dependencies... done
Starting kernel, version 2.1.05 (pid 36)
Checking filesystems.
Parallelizing fsck version 1.12 (9-Jul-98)
Checking all file systems.
-----
[/sbin/fsck.ext2] fsck.ext2 -a /dev/hda1
/dev/hda1 has gone too long without being checked, check forced.
/dev/hda1: 19/9488 files (0.8% non-contiguous), 1799/37768 blocks
-----
Mounting local filesystems.
Turning on user and group quotas for local filesystems
Setting clock: Sat May 20 14:22:23 GMT 2023
Enabling swap space.
Initializing random number generator...
INIT: Entering runlevel: 3
Starting up APM daemon: apmd No APM support in kernel

Disabling IPv4 packet forwarding.
sysctl: ip forwarding off
Swansea University Computer Society IPX 0.34 for NET3.035
IPX Portions Copyright (c) 1995 Caldera, Inc.
Appletalk 0.17 for Linux NET3.035

```

Screenshot 2

```

Turning on user and group quotas for local filesystems
Setting clock: Sat May 20 14:22:23 GMT 2023
Enabling swap space.
Initializing random number generator...
INIT: Entering runlevel: 3
Starting up APM daemon: apmd No APM support in kernel

Disabling IPv4 packet forwarding.
sysctl: ip forwarding off
Swansea University Computer Society IPX 0.34 for NET3.035
IPX Portions Copyright (c) 1995 Caldera, Inc.
Appletalk 0.17 for Linux NET3.035
3c59x.c:v0.99E 5/12/98 Donald Becker http://cesdis.gsfc.nasa.gov/linux/drivers/vortex.html
No 3Com Vortex/Boomerang cards found.
insmod: /lib/modules/preferred/net/3c59x.o: init_module: Device or resource busy
Delaying eth0 initialization.
Starting portmapper: portmap
Mounting remote filesystems.
Starting system loggers: syslogd klogd
Starting YP passwd service: rpc.yppasswdd
Starting YP server services: ypsserv
Starting at daemon: atd
Starting cron daemon: crond
Starting INET services: inetd
Starting lpd: lpd
Starting NFS services: rpc.mountd rpc.nfsd
Loading keymap: Loading /usr/lib/kbd/keymaps/i386/qwerty/uk.map.gz
Starting gpm mouse services: gpm gpm: /dev/mouse: Operation not supported by device

```

Screenshot 3

```

Starting sound configuration: sound
Binding to the NIS domain... ypbind
Starting xntpd: xntpd
Linuxconf final setup
Shutting down xntpd: xntpd

```

Screenshot 4

```

cagraidsvr04 root@tcd.ie:~# more
Memory: sized by init3 800h
Console: 16 point font, 400 scans
Console: colour VGA+ 88x25, 1 virtual console (max 63)
pcibios_init : BIOS32 Service Directory structure at 0x800f7780
pcibios_init : BIOS32 Service Directory entry at 0xf47b0
pcibios_init : PCI BIOS revision 2.10 entry at 0xf49b3
Probing PCI hardware.
Calibrating delay loop.. ok - 88.88 BogoMIPS
Memory: 38796k/32768k available (748k kernel code, 384k reserved, 840k data)
This processor honours the WP bit even in supervisor mode. Good.
Swansea University Computer Society NET3.035 for Linux 2.0
NET3: Unix domain sockets 0.13 for Linux NET3.035.
Swansea University Computer Society TCP/IP for NET3.034
IP Protocols: IGMP, ICMP, UDP, TCP
Linux IP multicast router 0.07.
UFS: Diskquotas version dquot 5.6.0 initialized
Checking 386/387 coupling... Ok, fpu using exception 16 error reporting.
Checking 'hlt' instruction... Ok.
Linux version 2.0.36 (root@porky.redhat.com) (gcc version 2.7.2.3) #1 Tue Oct 13 22:17:11 EDT 1998
Starting kswapd v 1.4.2.2
Serial driver version 4.13 with no serial options enabled
tty00 at 0x03f8 (irq = 4) is a 16550A
--More--

```

Screenshot 5

```

hd0 at 0x3f6 (irq = 4) is a 16550A
Real Time Clock Driver v1.09
Ramdisk Driver initialized : 16 ramdisks of 4096K size
hda: ST34321A, 4103MB w/128kB Cache, CHS=33344/5/17
hdc: TATUNG CD-1624E, ATAPI CDROM drive
ide0 at 0x1f0-0x1f7,0x3f6 on irq 14
ide1 at 0x170-0x177,0x376 on irq 15
Floppy drive(s): fd0 is 1.44M
FDC 0 is a post-1991 82077
md driver 0.36.3 MAX_MD_DEV=4, MAX_REAL=0
scsi : 0 hosts.
scsi : detected total.
Partition check:
  hda: hda1 hda2 < hda5 hda6 hda7 hda8 hda9 >
UFS: Mounted root (ext2 filesystem) readonly.
Adding Swap: 128484k swap-space (priority -1)
Adding Swap: 128484k swap-space (priority -2)
Adding Swap: 128484k swap-space (priority -3)
Adding Swap: 128484k swap-space (priority -4)
sysctl: ip forwarding off
Swansea University Computer Society IPX 0.34 for NET3.035
IPX Portions Copyright (c) 1995 Caldera, Inc.
Appletalk 0.17 for Linux NET3.035
3c59x.c:v0.99E 5/12/98 Donald Becker http://cesdis.gsfc.nasa.gov/linux/drivers/vortex.html
--More--


```

Screenshot 6

Figure 15: csTCDie Beowulf Cluster NIS server, Linux boot screenshots 1-6



Screenshot 7



Screenshot 8

Figure 16: csTCDie Beowulf Cluster NIS server, Linux console screenshots 7-9



Figure 17: csTCDie Beowulf Cluster NIS server, RAID array
[TCD-SCSS-T.20150217.003](#)