



Setting up UV2000 at NASA Ames

Herbert Yeung

Herbert.Yeung@nasa.gov

May 2013





Overview

- Endeavour configuration
- Problems encountered during installation
- Tuning recommendations





Why UV2000?

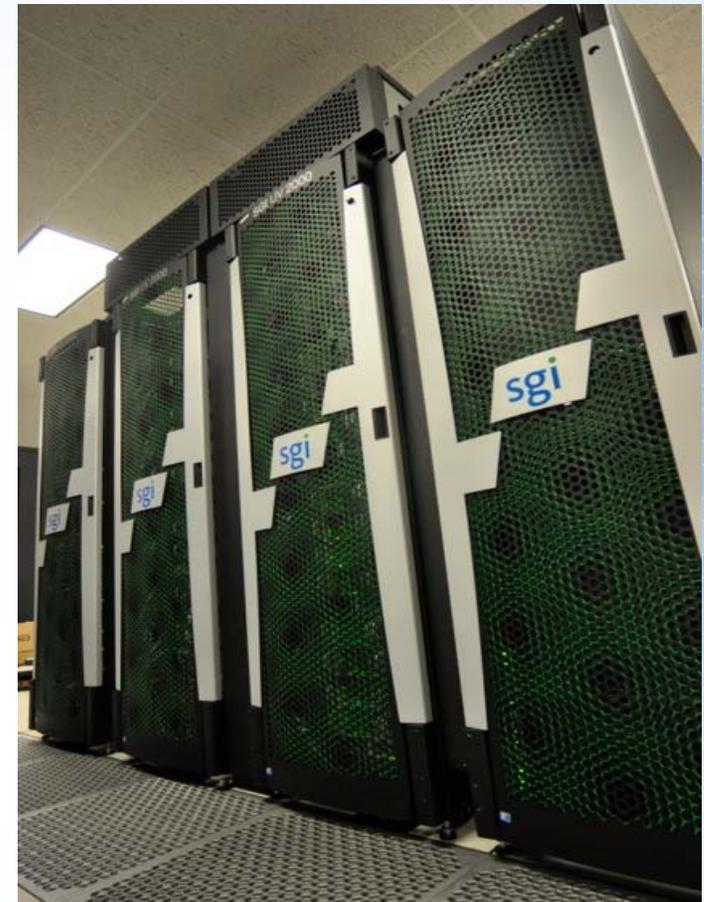
- To provide resources for applications that need access to large cache-coherent, global shared-memory capabilities in a single system image (SSI).
- Replace Columbia
 - First installed in 2004
 - Remaining systems, Columbia21-24 30TF
 - Endeavour provides 32TF





Endeavour Hardware

- <http://www.nas.nasa.gov/hecc/support/kb/entry/410>
- Endeavour1
 - 512P SSI UV2000
 - 2.6 GHZ Sandy Bridge
 - 2 TB memory
- Endeavour2
 - 1024P SSI UV2000
 - 2.6 GHZ Sandy Bridge
 - 4 TB memory
- Over 6.7 PB scratch space (lustre + NFS)
- Dual port 10GbE Intel Corporation I350
- Two/Four FDR Dual port ConnectX3 IB/Mellanox Infiniband HCAs





Endeavour Software

- SLES 11SP2
- STOUT706
- SGIMC 1.7
- 3.0.51-0.7.9.1.20130221-nasuv kernel
- Lustre 2.3.0-2.1
 - <https://github.com/jlan/lustre-nas/branches>
- PBSPro 11.3.0.120133_nas





System Management Node

- SLES 11SP2
- STOUT706
- SGIMC 1.7
- 3.0.38-0.5-default kernel
- conserver-8.1.18





Software Issues

- Automatic Boot
- policykilld
- Crashdumps
- mgrclient fails to start
- Pulling in Image





Automatic Boot

- On Columbia, we could create boot options via EFI boot menu
- Not available for UV2000 due to changed boot process

```
efibootmgr -c -l '\efi\SuSE\elilo.efi' -L elilo
```





policykilld

- This software attempts to supply the missing "cpuset policy kill" semantics for Linux.
- Written by Bron Nelson, who was an onsite SGI employee at the time.
- If interested in policykilld, contact your SGI salespeople or contact SGI Service.





Crashdump, first problem

- Ran out of memory on Endeavour2 (4TB)
- Initially reserved 512MB of memory
- Turns out we failed to even reserve 512MB
- Kdump kernel and all of its data needs to reside in “low” memory region
- SGI patched kernel to consolidate free memory in “low memory” region to create a large chunk of memory
 - Able to reserve 880MB for kdump
- SGI patched makedumpfile to move some of its operations to kernel
- `crashkernel_maximize` and `makedumpfile` patch available from SGI until community fixes issue of kdump being restricted to “low memory” region.





Crashdump, second problem

- Once able to drop into kdump kernel, was not able to write to disk
- Our Megaraid controller sits on PCI segment 1 while the kdump kernel could only see devices on PCI segment 0
- Patch to kexec-tools by SGI
- Put “megaraid_sas” under INITRD_MODULES in /etc/sysconfig/kernel
- kexec-tools patch pushed to SuSE and upstream. Not yet released to public





Crashdump, third problem

- Now able to start crashdumps
- Ran out of disk space – filled up root with over 190GB of data after taking over 7 hours
- makedumpfile failed to recognize hugepage pages as user data, which we do not want to save
- Another makedumpfile patch by SGI
- Always specify `KDUMP_DUMPLEVEL="31"` in `/etc/sysconfig/kdump` on a large memory system
- Creating the vmcore for Endeavour2 now takes about 13 minutes and uses up about 7GB.
- makedumpfile patch pushed to upstream Makedumpfile. Not yet released.





mgrclient fails to start

- Shows splash screen with error:

Could not communicate with the SGI Management Center server. Check that the server is running and try again.

- Errors from mgr

```
smn2 # /etc/init.d/mgr status
```

The service 'DistributionService.provisioning-00' is not responding.

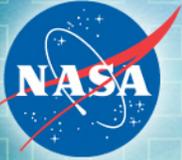
The service 'DistributionService.provisioning-01' is not responding.

The service 'FileService.UV00000044-P000' is not responding.

The service 'FileService.admin' is not responding.

- Need to include LD_LIBRARY_PATH to scripts and user's environment
- Problem caused by /etc/profile.d/mgr.sh not sourced by /etc/profile





Pulling in Image

- From mgrclient GUI:

Error: Unable to connect to the Version Service (UV00000044-P000)

- Errors from mgr:

```
smn2 # /etc/init.d/mgr status
```

The service 'InstrumentationService.UV00000044-P000' is not responding.

The service 'RemoteProcessService.UV00000044-P000' is not responding.

- From SGIMC-server.log:

```
com.xeroone.ComponentInstantiationException: VersionService.UV00000044-P000
```

```
at com.xeroone.BrokerService$BrokerServiceImpl.activate(Unknown Source)
```

```
at com.xeroone.BrokerService$BrokerServiceImpl$BrokerInterface_1Impl.lookup(Unknown Source)
```

```
at com.Inxi.payload.server.PayloadAdministrationService$PayloadAdministrationServiceInterface_1Impl$InstallationCreateThread.run(PayloadAdministrationService.java:5774)
```

```
Caused by: java.lang.IllegalArgumentException
```

```
at com.xeroone.BrokerReference.<init>(Unknown Source)
```

- Problem caused by /etc/profile.d/mgr.sh not sourced by /etc/profile





Pulling in Image, second issue

- From mgrclient GUI:

Error: Unable to connect to the Version Service (UV00000044-P000)

- Errors from mgr on client and server:

```
endeavour2 # /etc/init.d/mgr status
```

The service 'DNA.172.21.1.0' is not responding.

The service 'HostAdministrationService.UV00000044-P000' is not responding.

The service 'InstrumentationService.UV00000044-P000' is not responding.

```
smn2 # /etc/init.d/mgr status
```

The service 'DNA.10.150.63.50' is not responding.

The service 'DNA.172.21.1.0' is not responding.

The service 'FileService.endeavour2' is not responding.

The service 'HostAdministrationService.UV00000044-P000' is not responding.

The service 'HostAdministrationService.endeavour2' is not responding.

- On Endeavour2, need to use serial number instead of hostname in files
/opt/sgi/sgimc/\@genesis-profile & /opt/sgi/sgimc/Activator.profile
- On Smn2, removed /opt/sgi/sgimc/*.rna
- Restarted mgr on client and server

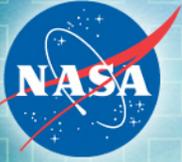




Pulling in Image, third issue

- From GUI on smn1:
Error: Unable to connect to the Version Service (UV00000043-P000)
- No errors in logs
- No errors with status check of mgr
- Problem caused by iptables file on Endeavour1
- Opened up management interface
INPUT -s 172.21.1.0/255.255.0.0 -j ACCEPT
to
INPUT -i eth1 -j ACCEPT





Hardware Issues

- Unknown Beeping
- CATERRs
- Failed reboots





Unknown Beeping

- Heard beeping coming from near base I/O blade for Endeavour1
- No warning lights visible
- Nothing should be able to beep from blade
- SGI tracked it down to disk in bad state via MegaCli

```
endeavour1 # /opt/MegaRAID/MegaCli/MegaCli64 -ldinfo -lall -aall
```

```
State          : Degraded
```

- Could not find problems in logs, so changed disk state to good
- Started rebuild process
- After rebuild finished, beeping stopped





CATERRs

- System hung. On console, found:

```
***** [20130116.224433] BMC r001i33b00h0: CATERR detected!
```

```
***** [20130116.224433] BMC r001i01b04h0: CATERR detected!
```

```
***** [20130116.224433] BMC r001i11b00h0: CATERR detected!
```

```
***** [20130116.224433] BMC r001i11b00h1: CATERR detected!
```

```
***** [20130116.224433] BMC r001i23b01h0: CATERR detected!
```

```
***** [20130116.224433] BMC r001i23b06h1: CATERR detected!
```

```
***** [20130116.224433] BMC r001i33b03h0: CATERR detected!
```

- Can be caused by problems with numalink fabric
- Run `uv2dump` to gather information
- Fixed with power cycle





Failed Reboots

- During reboots, see console messages:

```
***** [19700116.223009] BMC r002i23b07h0: Cold Reset via BMC
```

```
***** [20000104.203403] BMC r001i01b00: Reset *ERROR*
```

```
***** [20000104.203404] BMC r001i01b00h1: Cold Reset via BMC
```

- After reset:

```
CMC:r001i01c> power reset r001i01b00
```

```
==== r001i01b00 ====
```

```
r001i01b00:
```

```
WARNING: Patsburg sleep state detected: waking via PBG P1V5_AUX reset
```

```
Probing NL6 cables
```

```
ERROR: Patsburg in fatal sleep state: power cycle the BMC to recover
```

```
ERROR: power command failed
```

- Power cycling would not clear problem
- Original workaround was to reseal the I/O blade





Failed Reboots, continued

- Asked SGI for method that doesn't require physical access to the system
- Watch the r1i1b0 console window and the CMC command window. Should see blade output about the same time as the CMC window returns the prompt. If the windows are way out of sync, it means you have Link errors.

```
cmc> power off
```

```
cmc> power cycle bmc all
```

```
cmc> power on
```





Failed Reboots, continued

- If the 2 windows are in sync, then in the CMC window:

```
cmc>bmc nl6cmd check
```

- If this is clean and has no link errors then it may boot. If it gets an error, which will show up in the uvcon system window then issue:

```
cmc> Power sreset (this is a soft reset, it resets the sockets and leaves the nl6 fabric alone)
```

- This should make the system boot.
- Note: You may have to repeat the whole process several times.
- Issue alleviated by BMC firmware 0.8.0.





Performance tuning

- NAS Parallel Benchmarks
 - <http://www.nas.nasa.gov/publications/npb.html>
- SBU benchmarks
 - ENZO, FUN3D, GEOS-5, OVERFLOW, UMS3D, and WRF
- Performed by Application Performance and Productivity team





Hyperthreading

- We have hyperthreading disabled
- Setting changed via BIOS
- Some applications do benefit
- Keep job submittal simple
 - If enabled, users that want N physical cores need to request $2*N$ CPUs
 - PBS would need to keep track of which CPUs are on same core





Turbo boost

- To enable:
 - Load acpi-cpufreq kernel module
 - In /etc/modprobe.d/acpi-cpufreq.conf, comment out:

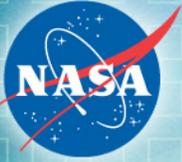
```
install acpi-cpufreq /bin/true
```

- Setup init script:

```
maxcpu=`grep processor /proc/cpuinfo | awk '{print $3}' | tail -1`  
for cpu in `seq 0 $maxcpu`  
do  
    cpufreq-set -c $cpu -g performance  
done
```

- Gives a performance boost to most codes





Transparent Huge Pages

- Currently disabled
- Still under investigation
- To disable:

```
echo never > /sys/kernel/mm/transparent_hugepage/enabled
```





Zone reclaim

- Currently enabled
- To enable in `sysctl.conf`:
`vm.zone_reclaim_mode = 1`
- No observed reduction in run-time variability, but no observed issues





Remaining Work

- Transparent huge pages
- Lustre performance
- Performance issues





Questions?

This work was in support of NASA contract #107885.2.730.B





MPT huge pages

- Currently enabled
- Still under investigation

- To enable:

```
mpt_hugepage_config -p 90 -u
```

- Allow user access via `/etc/sysctl.conf`:

```
vm.hugetlb_shm_group = 1989
```

- Bad permissions in `/etc/mpt/hugepage_mpt`

```
hugetlbfs      /etc/mpt/hugepage_mpt hugetlbfs mode=1777  0 0
```

- Users need:

```
MPI_HUGEPAGE_HEAP_SPACE = 1
```

