Tools and Techniques for Managing Clusters for SciDAC Lattice QCD at Fermilab

D Holmgren, R Rechenmacher, A Singh, S Epsteyn

Amitoj Singh amitoj@fnal.gov

Fermi National Accelerator Laboratory, Batavia, IL

FNAL SciDAC Lattice QCD clusters

80 node Pentium III cluster

176 node Xeon cluster

Introduction

- Tools
  - for hardware management tasks.
  - for OS installation/upgrade and reloading BIOS/firmware.
  - tools that work in conjunction with the PBS batch queue system.
- Integration of tools for remote administration.

Remote node management

- serial links to each worker node.
- BIOS/console redirection.
- support both IPMI 0.9 and 1.5 versions.
- IPMI – remote power on, power off, reset.

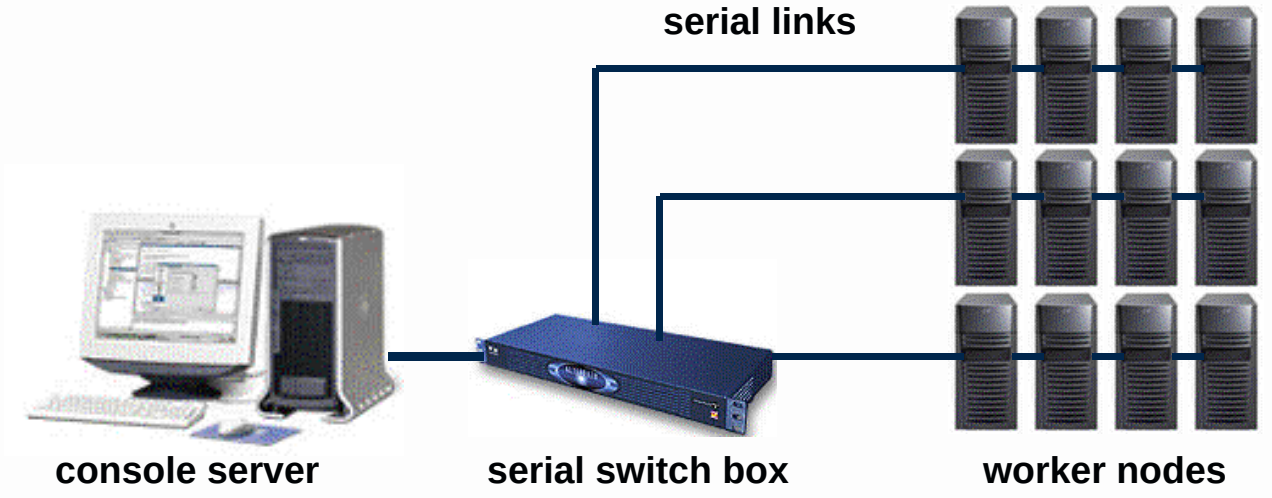

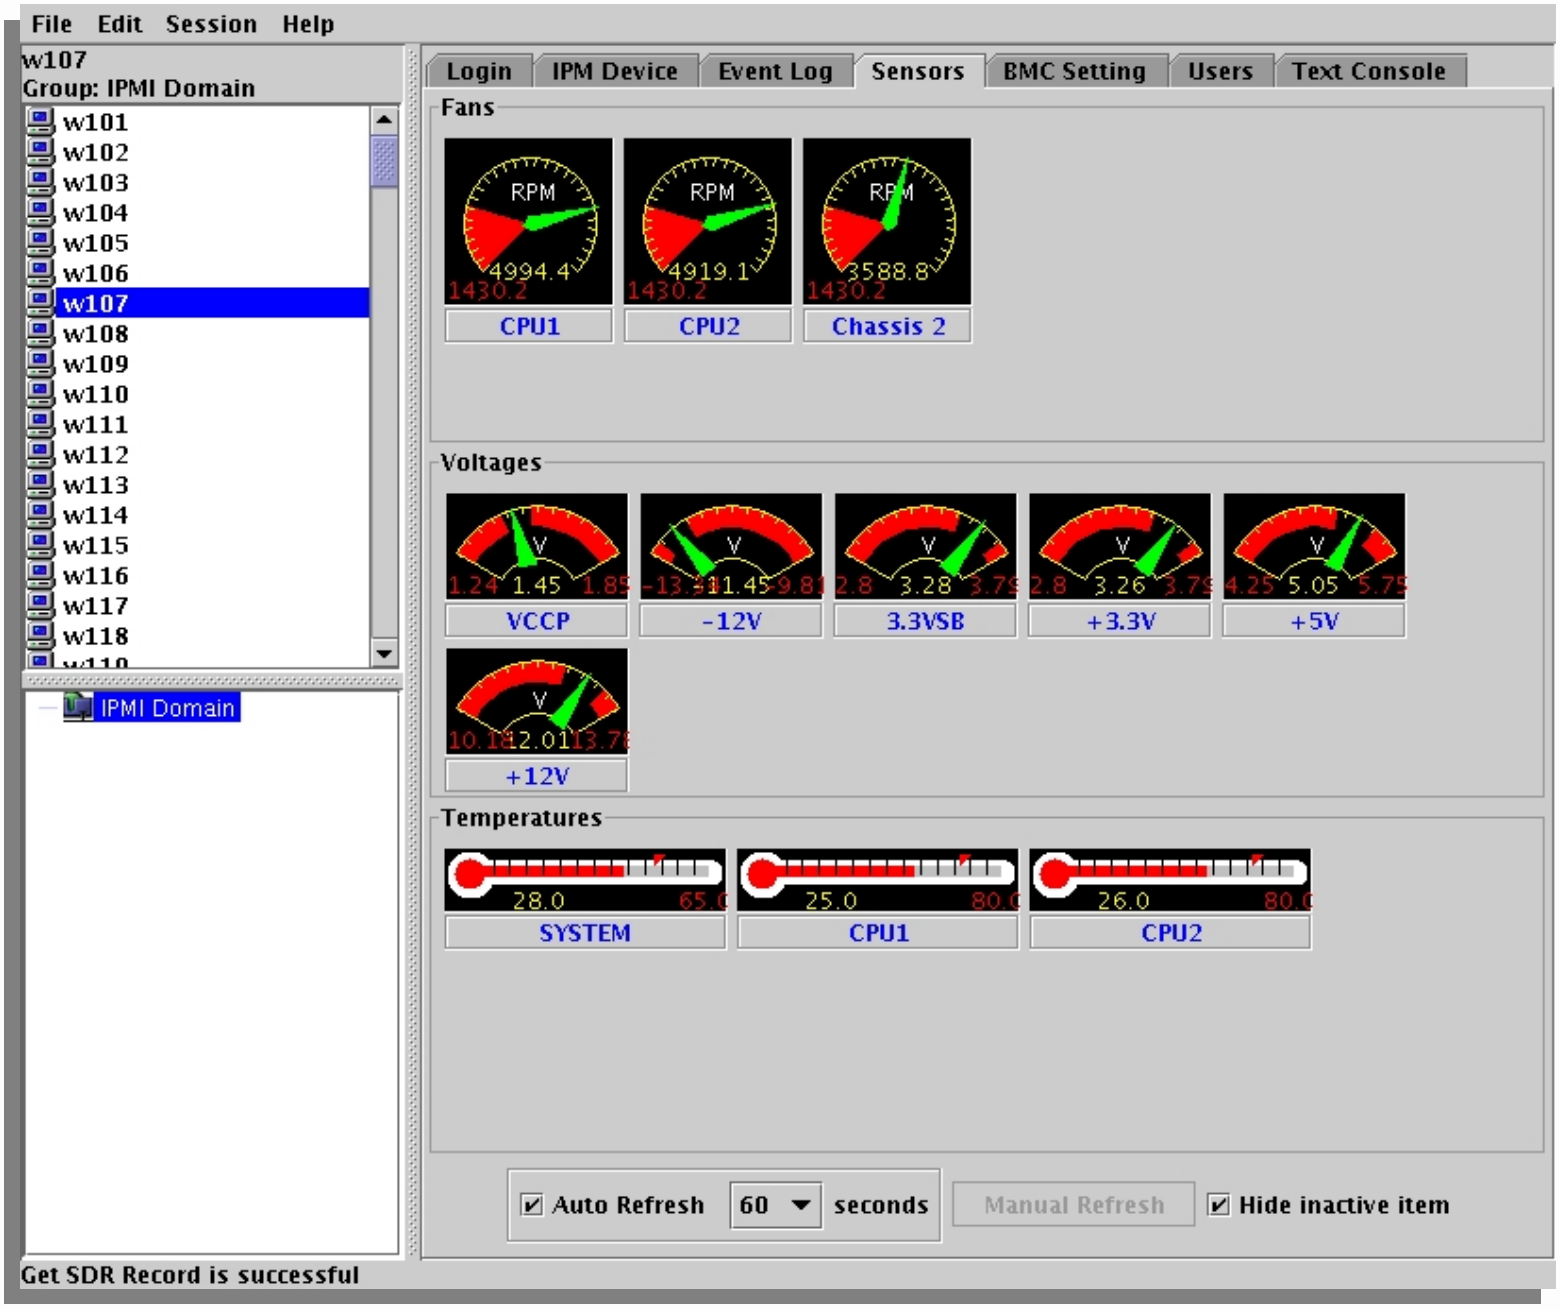

Super Micro's **IPMI View** – GUI based management Interface over LAN

Network boot

- PXE (Pre-boot Execution Environment)

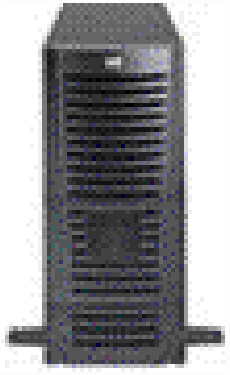

DHCP/BOOTP tftp server

PXE DHCP/BOOTP request

returns host IP, bootloader file location

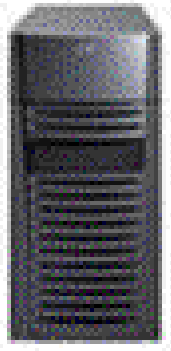

client

	OS Install	**BIOS/Firmware Install**
PXE loads	bootloader pxelinux	bootloader pxegrub
bootloader loads	linux kernel image + initrd. Launches kernel bootstrap	mknbi (from etherboot) DOS image. Launches DOS bootstrap
install	via init/rc.d, brings in install script, which -partitions disk -makes file system -rcp's tar files -explodes file system -does host IP config -rsh's to bootp server to comment itself out under bootptab. reboot.	autoexec.bat executes -does firmware/BIOS upgrade/install -using bootp discovers host IP -ftp's stdout/stderr to bootp server. -rsh's to bootp server to comment itself out under bootptab. reboot.

Fermi Tools

- **rgang** – (milc - minimum level of complexity) execute the same command on all of the nodes. Coded in python.
  - two modes :
    - command mode.
    - copy mode.
  - n-way option.
  - rgang can be used to install itself.

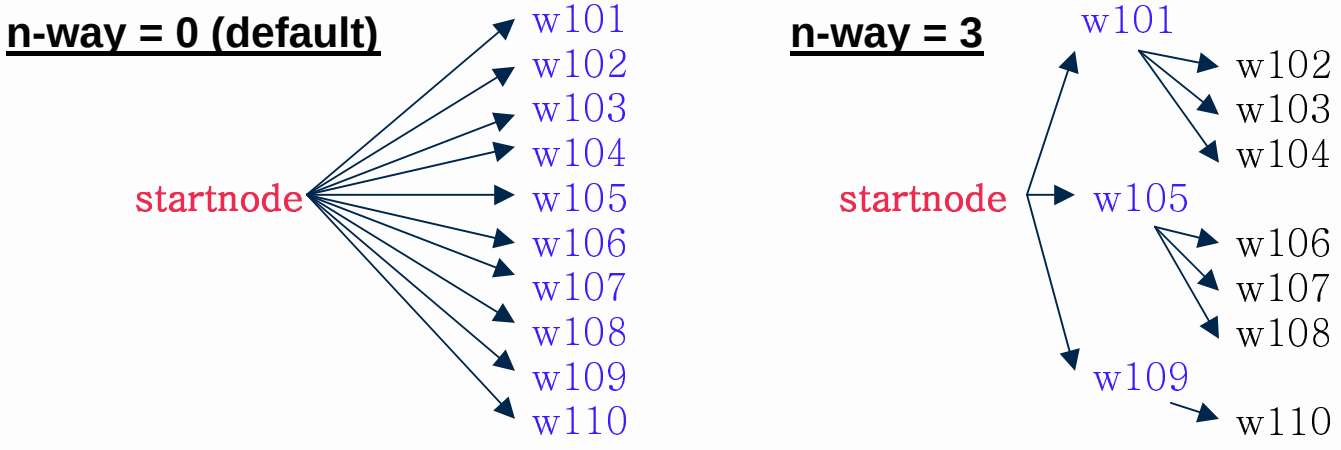

Fermi Tools

- **fermistat** - list cluster resource usage by users and running jobs. Works in conjunction with PBS batch queue system.
- **fermitrack** – poor man's project accounting. Works in conjunction with PBS batch queue accounting system.

  > qsub –A "myproject" –l nodes=n mypbsscript

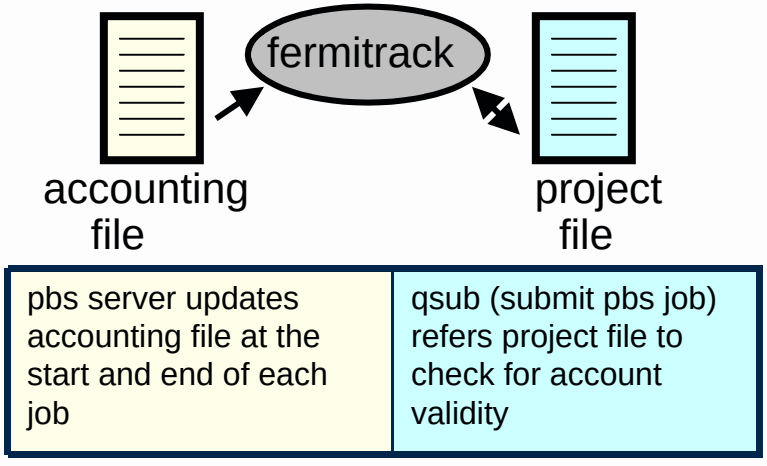

Integration-How it all comes together.

head node checks worker node health data & resource usage are within safe limits

Head node

head node **rgang's** to all worker nodes to check resource usage

health data transferred over to head node over **syslog** udp socket

worker nodes

fermitrack uses PBS accounting to check per project resource usage.

each worker node collects health data issuing IPMI commands to the BMC interface at set intervals

http://lqcd.fnal.gov

File Edit View Go Communicator Help

lqcd.fnal.gov Node Map

w101		w102		w103		w104		w105		w106		w107		w108		w109		w110		w111			
w112		w113		w114		w115		w116	1	w117	1	w118	1	w119	1	w120	1	w121	1	w122	1		
w123	1	w124	1	w201	1	w202	1	w203	1	w204	1	w205	1	w206	1	w207	1	w208	1	w209	1		
w210	1	w211	1	w212	1	w213	1	w214	1	w215	1	w216	1	w217	1	w218	1	w219	1	w220	1		
w221	1	w222	1	w223	1	w224	2	w301	2	w302	2	w303	2	w304	2	w305	2	w306	2	w307	2		
w308	2	w309	2	w310	2	w311	2	w312	2	w313	2	w314	2	w315	2	w316	2	w317	2	w318	2		
w319	2	w320	2	w321	2	w322	2	w323	2	w324	2	w401	2	w402	2	w403	2	w404	2	w405	2		
w406	2	w407	2	w408	3	w409	3	w410	3	w411	3	w412	3	w413	3	w414	3	w415	3	w416	3		
w417	3	w418	3	w419	3	w420	3	w421	3	w422	3	w423	3	w424	3	w501	3	w502	3	w503	3		
w504	3	w505	3	w506		w507	3	w508	3	w509	3	w510	3	w511	3	w512	3	w513	3	w514	3		
w515	3	w516	3	w517	4	w518	4	w519	4	w520	4	w521	4	w522	4	w523	4	w524	4	w601	4		
w602	4	w603	4	w604	4	w605	4	w606	4	w607	4	w608	4	nqcd0101		nqcd0102		nqcd0103		nqcd0104			
nqcd0105		nqcd0106		nqcd0201		nqcd0202		nqcd0203		nqcd0204		nqcd0205		nqcd0206		nqcd0301		nqcd0302		nqcd0303			
nqcd0304		nqcd0305	5	nqcd0306	5	nqcd0401	5	nqcd0402	5	nqcd0403	5	nqcd0404	5	nqcd0405	5	nqcd0406	5	nqcd0501	5	nqcd0502	5		
nqcd0503	5	nqcd0504	5	nqcd0505	5	nqcd0506	5	nqcd0601	5	nqcd0602	5	nqcd0603	5	nqcd0604	5	nqcd0605	5	nqcd0606	5	nqcd0701	5		
nqcd0702	5	nqcd0703	5	nqcd0704	5	nqcd0705	5	nqcd0706	5	nqcd0801	5	nqcd0802	5	nqcd0803	5	nqcd0804	5	nqcd0805	5	nqcd0806	5		

free : 31 down : 0 offline : 1 reserve : 0 job-exclusive : 144 job-sharing : 0 **Usage : 82%**

Job List

Status **BUSY**

Ref Id	Job Id	Job Name	User	Time Use	S	Queue	Nodes
5	846799	falpha4.sh	trottier	13:35	R	workq	32
4	846807	par_4sc.sh	cdavies	01:02	R	workq	16
1	846801	falpha3.sh	trottier	13:32	R	workq	32
3	846800	falpha1.sh	trottier	13:34	R	workq	32
2	846797	falpha2.sh	trottier	13:37	R	workq	32

100%

http://lqcd.fnal.gov

File Edit View Go Communicator Help

all Health Statistics

Hostname	Board Tmp1	Board Tmp2	Proc Tmp1	Proc Tmp2	Fan Speed1	Fan Speed2	Chasis Fan2	Warn Code
qcd0101	42	42	36	38	4819	4794	9999	0
qcd0102	41	41	37	38	4794	4794	9999	0
qcd0103	45	45	38	39	4794	4819	9999	0
qcd0104	45	45	39	40	4819	4844	9999	0
qcd0105	45	45	38	40	4869	4844	9999	0
qcd0106	44	44	38	38	4869	4844	9999	0
qcd0201	37	37	32	33	4819	4819	9999	0
qcd0202	39	39	34	35	4794	4844	9999	0
qcd0203	39	39	34	35	4794	4819	9999	0
qcd0204	40	40	35	35	4794	4794	9999	0
qcd0205	40	40	35	36	4794	4844	9999	0
qcd0206	39	39	34	35	4819	4794	9999	0
qcd0301	35	35	29	29	4718	4794	9999	0
qcd0302	35	35	30	31	4794	4768	9999	0
qcd0303	35	35	31	31	4743	4794	9999	0
qcd0304	37	37	31	32	4794	4794	9999	0
qcd0305	45	45	49	46	4819	4819	9999	0
qcd0306	44	44	48	47	4794	4794	9999	0
qcd0401	64	64	63	64	4869	4894	9999	0
qcd0402	67	67	68	67	4919	4844	9999	0
qcd0403	67	67	67	67	4869	4869	9999	0
qcd0404	64	64	66	66	4919	4844	9999	0
qcd0405	66	66	68	67	4869	4869	9999	0
qcd0406	64	64	64	65	4869	4894	9999	0
qcd0501	59	59	58	59	4844	4844	9999	0
qcd0502	61	61	60	61	4919	4819	9999	0
qcd0503	0	0	0	0	0	0	9999	
qcd0504	59	59	61	61	4844	4844	9999	0
qcd0505	62	62	63	63	4844	4894	9999	0
qcd0506	59	59	58	59	4844	4844	9999	0
qcd0601	54	54	54	53	4844	4819	9999	0

http://lqcd.fnal.gov

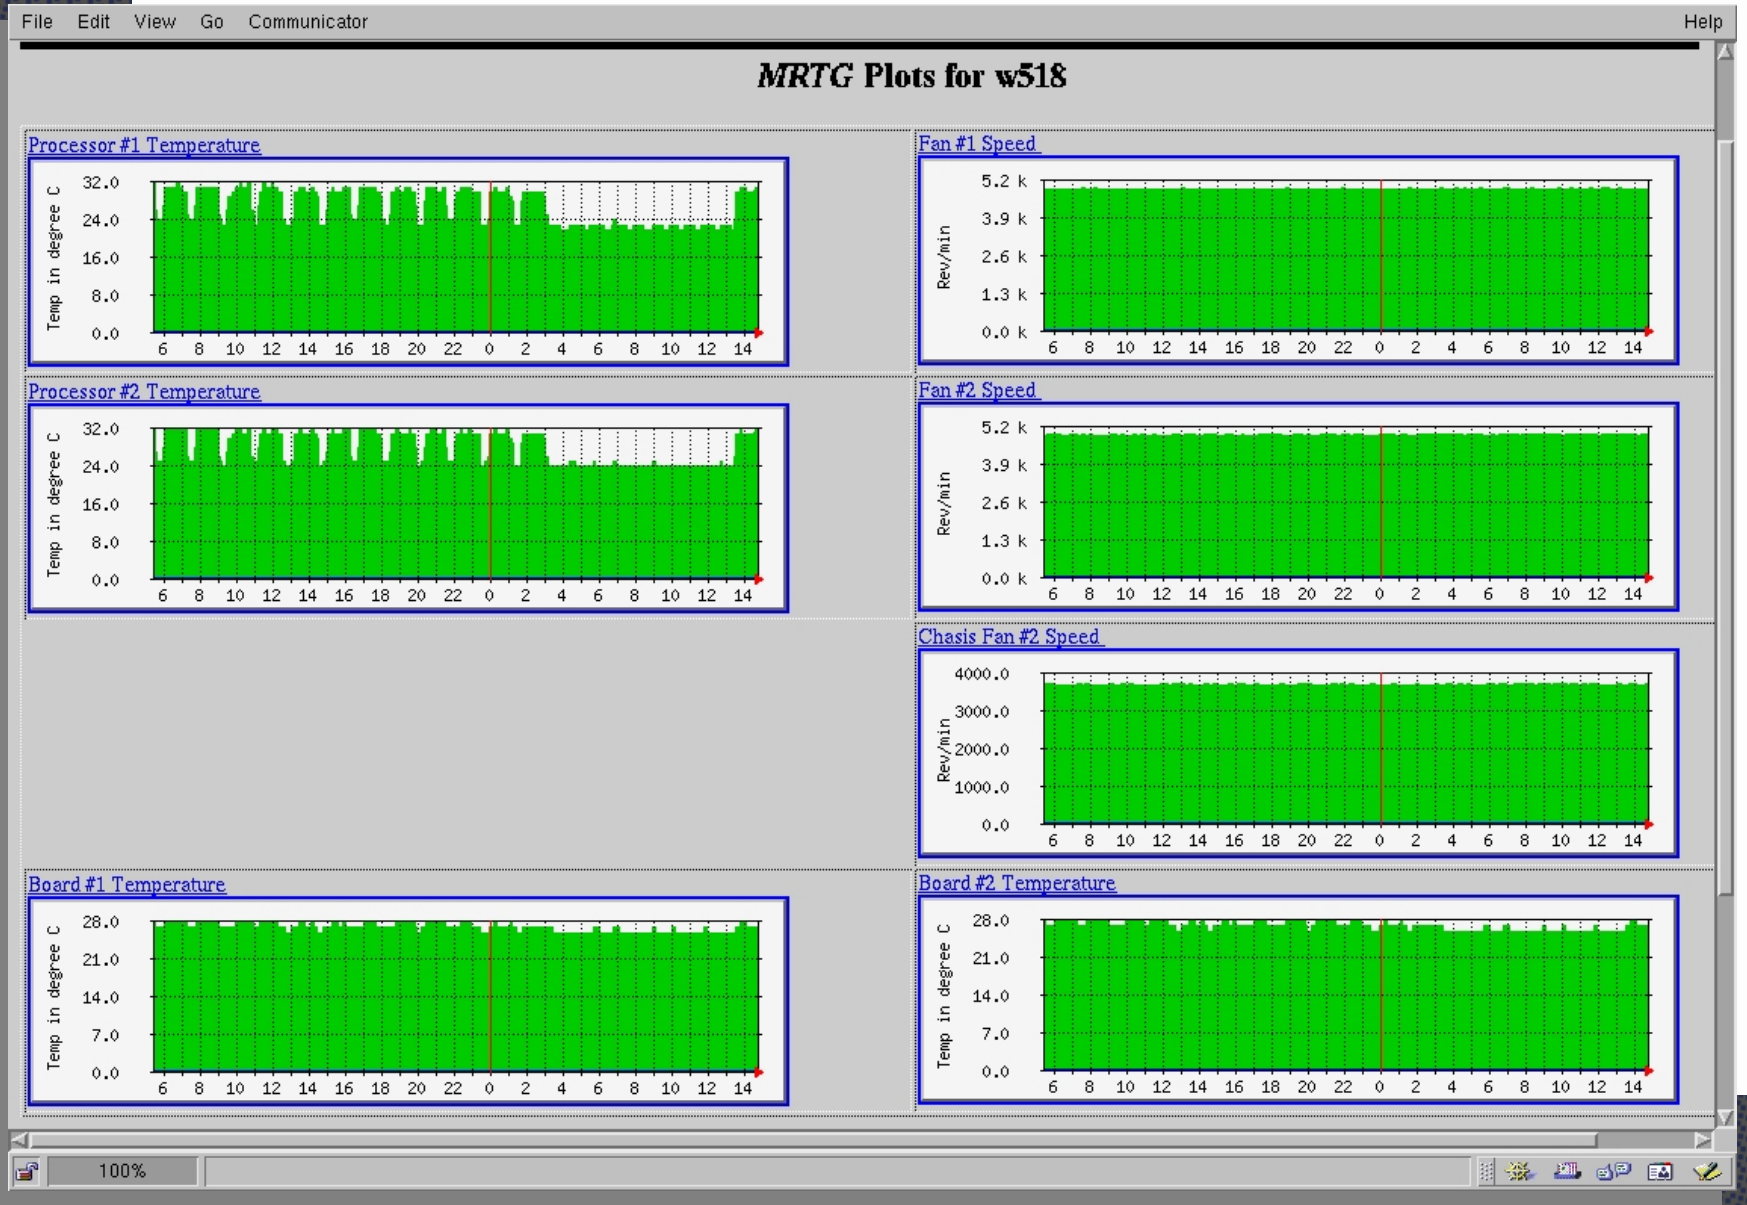

Example

- **fermistat –l <jobid> | rgang – <command>**
  - fermistat –l 8345.job | rgang – uptime

```
- - - - - - - - - - - - - - - w605 - - - - - - - - - - - - - - -
rsh w605 'uptime'
2:21pm  up 9 days,  2:02,  0 users,  load average: 1.99, 1.97, 1.91
- - - - - - - - - - - - - - - w606 - - - - - - - - - - - - - - -
rsh w606 'uptime'
2:21pm  up 9 days,  2:03,  0 users,  load average: 1.99, 1.97, 1.91
- - - - - - - - - - - - - - - w607 - - - - - - - - - - - - - - -
rsh w607 'uptime'
2:21pm  up 9 days,  2:02,  0 users,  load average: 1.99, 1.97, 1.91
- - - - - - - - - - - - - - - w608 - - - - - - - - - - - - - - -
rsh w608 'uptime'
2:21pm  up 9 days,  2:02,  0 users,  load average: 2.00, 1.97, 1.91
```

Example

- **fermistat –c <rgang style list-of-nodes>**
  - fermistat -c w1{01-04}
    - pbsnodes -c w101
    - pbsnodes -c w102
    - pbsnodes -c w103
    - pbsnodes -c w104

- **fermistat –l <jobid> | fermistat -o –**
  - fermistat -l 8345.job | fermistat -o -
    - pbsnodes -o w605
    - pbsnodes -o w606
    - pbsnodes -o w607
    - pbsnodes -o w608

Conclusion

It works . . . .

http://qcdhome.fnal.gov - 80 node Pentium III cluster
http://lqcd.fnal.gov - 176 node Xeon cluster
http://fermitools.fnal.gov - rgang and other tools

Don Holmgren
djholm@fnal.gov

Ron Rechenmacher
ron@fnal.gov

Amitoj Singh
amitoj@fnal.gov

Simon Epsteyn
seva@fnal.gov