F. Chlebana
Aug 18 2006

Status of the CSL Upgrade

I. Bizjak, F.Chlebana, G. Guglielmo, J.Lee, T.Masubuchi,

K.McFarland, R.Sarkis, JJ.Schmidt, M.Shimojima,

R.Snider, G.Yu, D.Zhang,

DataComm, Computing Division

Why Upgrade the CSL

- Current CSL is limited to about 24 MB/s
  → Need to be able to handle higher logging rates

Run 211554 ($\mathcal{L} = 176 \times 10^{30}$)

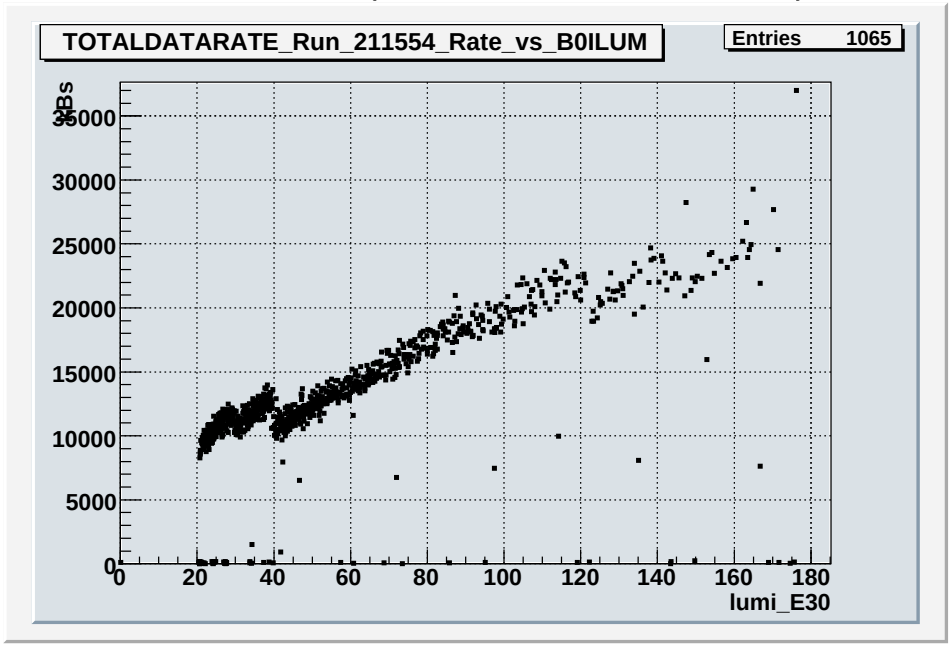

What logging rate will we need?

250 KB/Event
1000 Hz
L3 rejection of 4-5 times

→ 50 - 63 MB/s

Currently can write to buffer disks at higher rates (45 MB/s), the rate from the buffer disks to tape is limited to about 24 MB/s.

CSL upgrade design target is 80 MB/s

EVB Rate	1000 Hz	1000 Hz	1200 Hz
Rejection at L3	4-5x -> 250-200Hz	4x -> 250 Hz	4x -> 300 Hz
Event size	200 KB	250 KB	250 KB
Required Bandwidth	40-50 MB/s	63 MB/s	75 MB/s
Datahandling limit	80 MB/s		

Provide enough logging capacity so that the CSL is not the bottleneck in the system

- Uses Silicon Graphics and IRIX
  → Becoming more difficult to support
  → Data handling (SAM) software not ported to IRIX

- Using aging hardware
  → "end of life" of hardware approaching, difficult to support

Reliability Issues

The current CSL has been extremely reliable and we have not encountered any significant downtime because of problems with the CSL.

In the last ∼four years, the following error situations with b0dau32 and the RAID disks have been encountered

- Single disk failure (a few times per year)
  → Disk can be replaced and RAID rebuilt with no effect on the data taking.

- Disk controller failure (∼once for the last four years)
  → Stream is automatically redirected to another disk. Any open files under the controller are lost.

- Unexpected power outage (a few times for last four years)
  → All the open files are lost.

- One of the streams flooded due to a trigger PS problem. (once or twice for last four years during data taking, many times during cosmics)
  → Not really a csl problem, but novice can get fooled because what happens is L3 nodes turn to green. Monitoring can get smarter to tell the aces what's really wrong.

- Memory running out
  → Happened a few times but not understood. Memories should be monitored as in the b0dau32_detailed history web page, in order to recognize this problem.

- Hardware problem (bad memory)
  → It happened last summer, causing kernel panics. It was not easy to pin point the problem.

CSL Upgrades

Retain original software design and port onto new hardware

→ Current software structure serves our purpose

→ Limited resources for any major rewrite

Increase data throughput

→ Use modular distributed logging architecture

Increased buffer capacity

→ Buffer increased from 3.7 TB to 24 TB

→ 80 MB/s → ∼7 TB/day

Improve maintainability and robustness

→ Replace aging equipment with new servers/disk arrays

→ Use Linux instead of IRIX

→ Redundant hardware and hot spares

Expandable

→ Can increase number of logger nodes and disk buffer

Hardware Status

Performance was first evaluated using a *prototype system* and hardware elsewhere at the lab.

Hardware has arrived and is installed and configured

10 Servers, 3.4 GHz (dual processor) with 4 GB memory
Using Scientific Linux 4.2 (2.6 kernel)

2 SATABeast disk arrays with 28 - 500 GB disks per unit (14 empty slots) $\rightarrow$ *Total of 24 TB of buffer capacity*

Data transported over Gigabit ethernet ($\sim$125 MB/s capacity) , using multiple Network Interface Cards to increase bandwidth

Servers connected to the Disk Arrays through a Fibre Channel Switch, providing a redundant path and failover protection.

Have available two fully configured hot spares, one Receiver node and one Logger node.

Architecture Overview

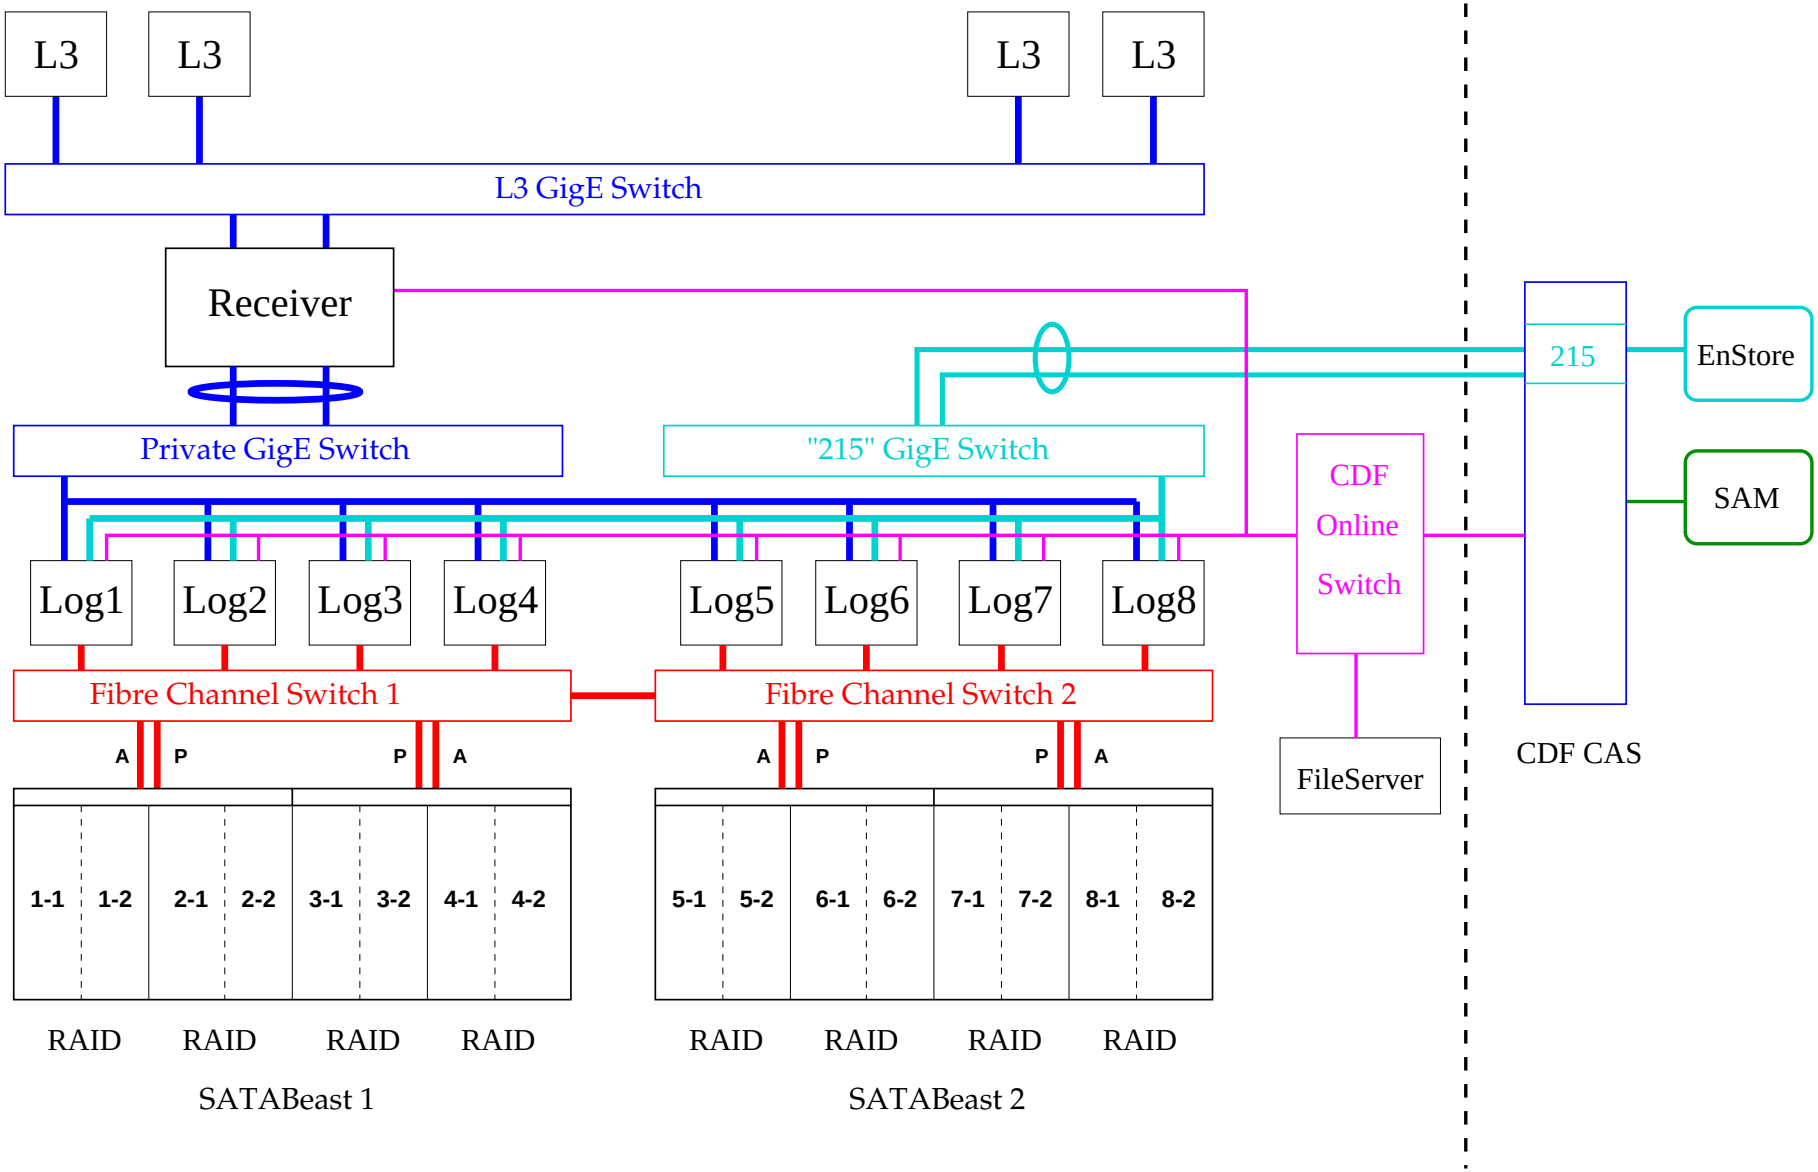

Commissioning Steps

Step 1) System burn in tests

→ *Configuration*

→ *Infant mortality*

→ *Verify hardware performance*

Step 2) L3Sender sends data from a file through the system

→ *Software functionality*

→ *Interface between CSL and logging software*

→ *Functionality*

→ *Develop monitoring*

→ *Also sent events to a dummy consumer!*

Step 2) Connect L3 Output

→ *Change the source of the data from L3Sender to EVB Output*

→ *Interface with RC*

→ *Operational experience*

Step 3) End of store test

→ *Demonstrate full functionality*

Network Tests

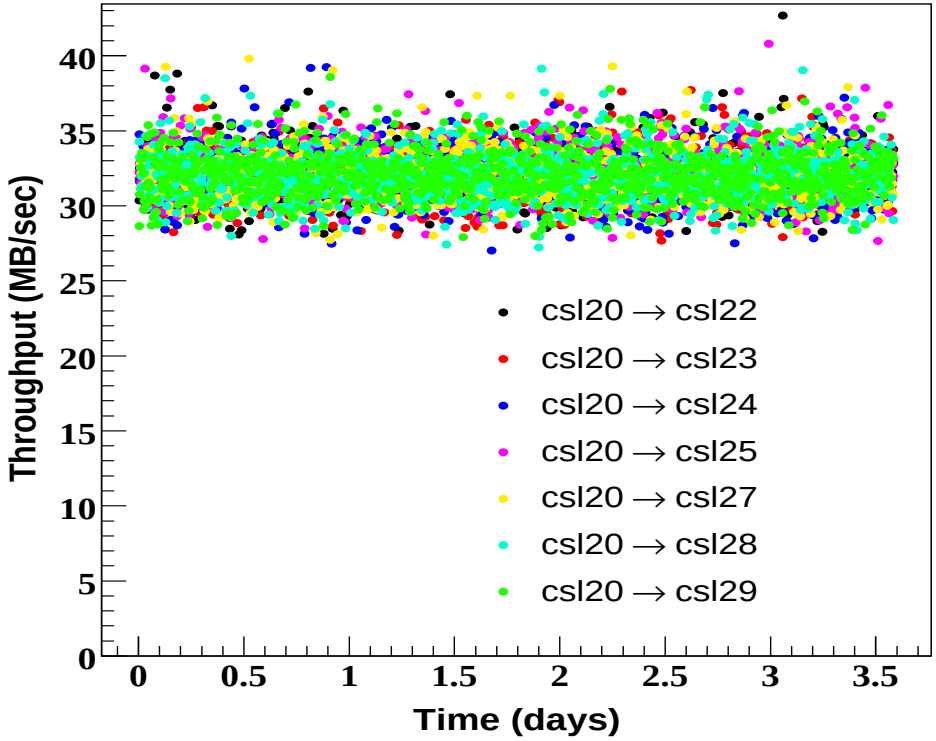

Sending data from receiver node to *7 logger nodes*

Bandwidth is equally shared

Using two GigE links *bonded* together

Total network rate: $7 \times 32 = 224$ MB/s

Disk IO Tests

Have been running long term tests in various configurations

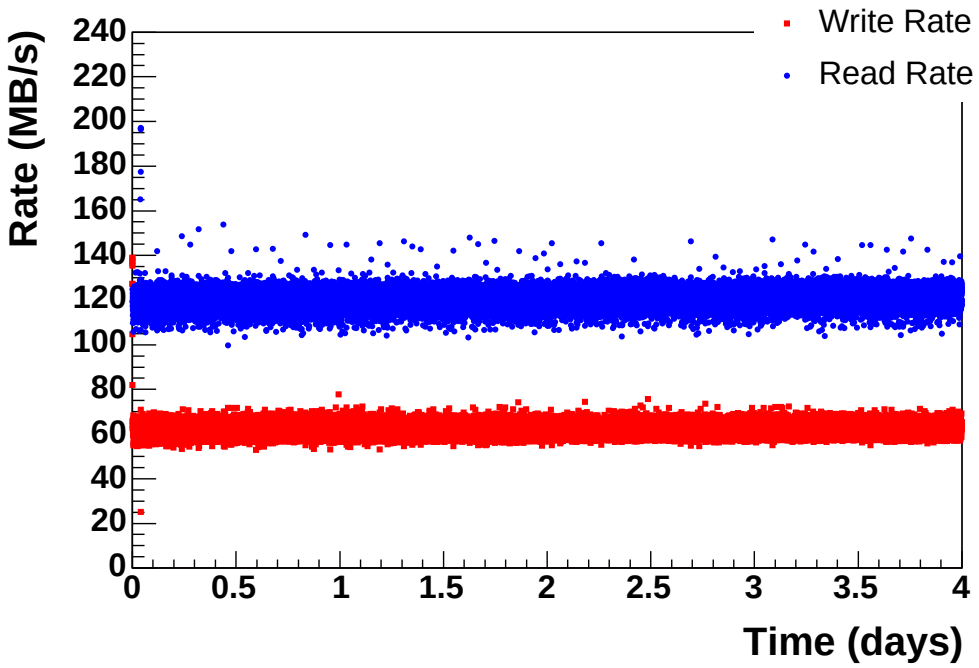

Using 4 servers to write to 2 SATABeast through 4 controllers

Servers are accessing the disk arrays through the Fibre Channel Switch

Tests running on/off for about two weeks

So far looks good!

4 logger nodes writing/reading to/from one SATABeast

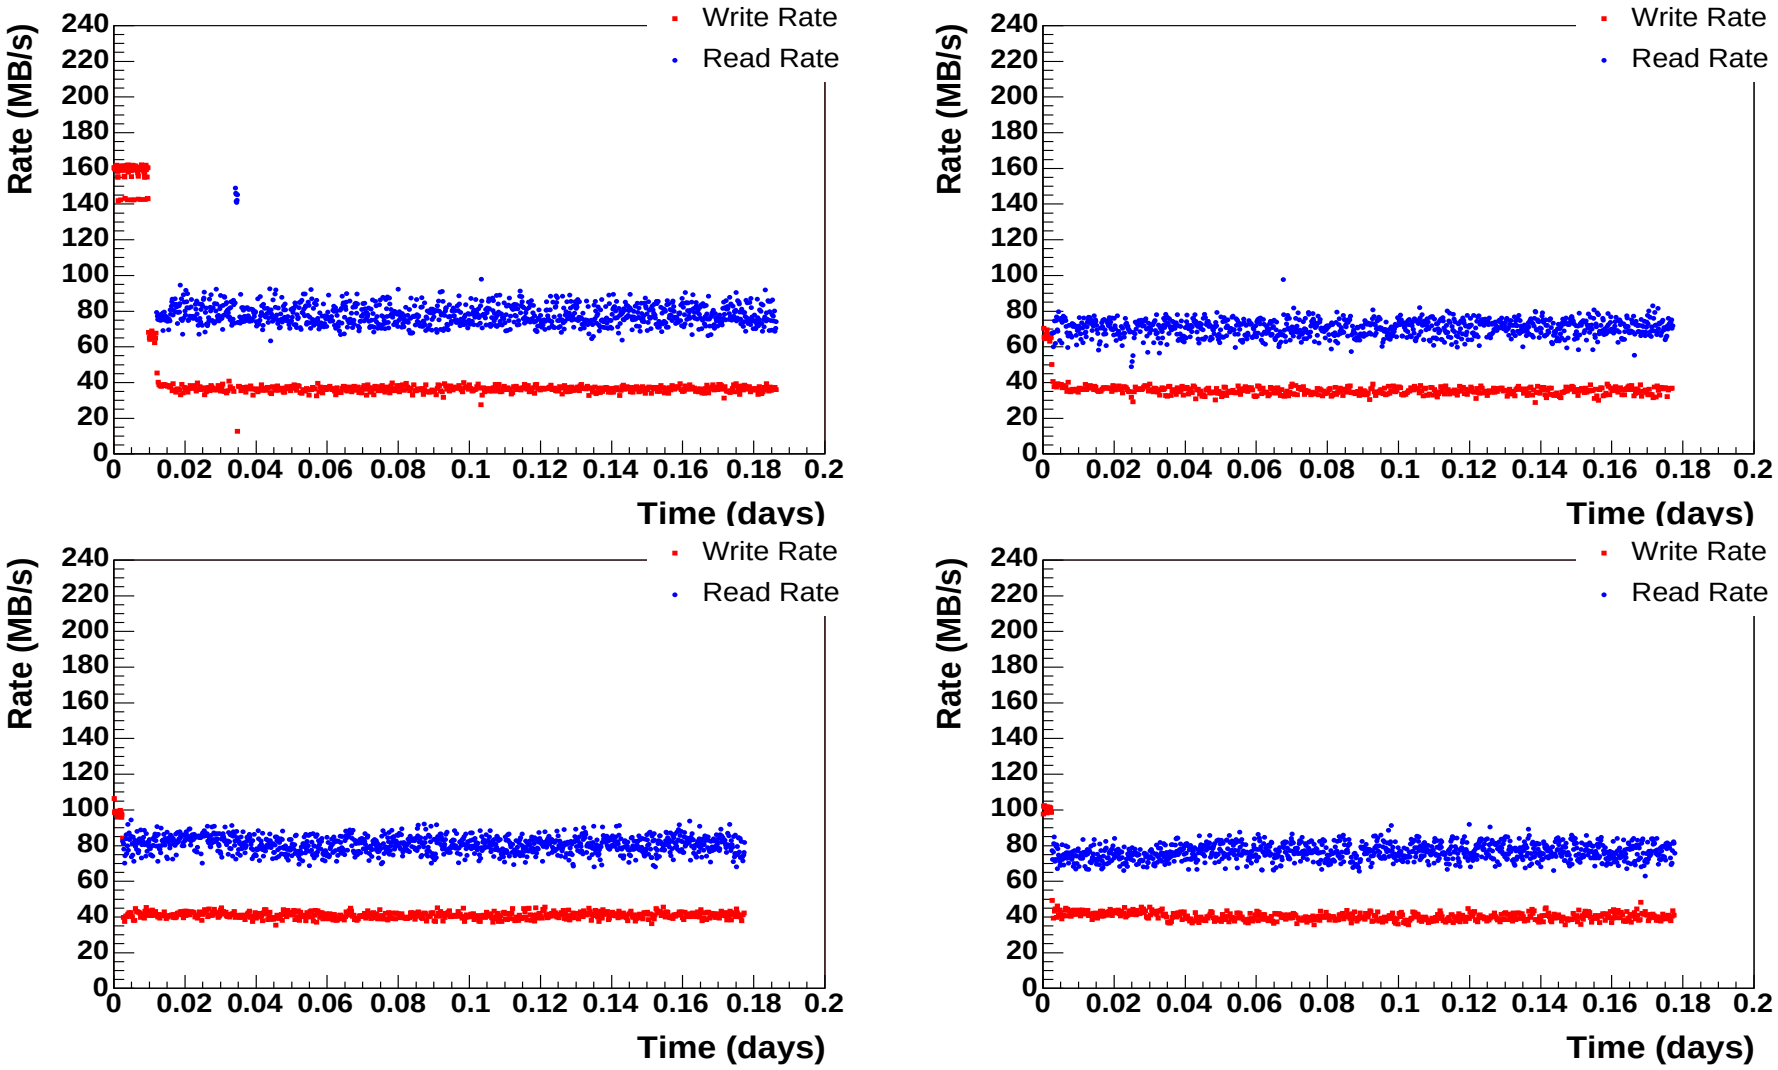

Total Disk Bandwidth using 8 loggers: > 300 MB/s

This depends on some configuration details (io scheduler) and may change... We want to emphasize reliability while ensureing we meet the specs.

Integration Tests

L3Sender → CSL Receiver → Distributed to 8 logger nodes → written to disk → Three separate Stager processes transfered data to Feynmann

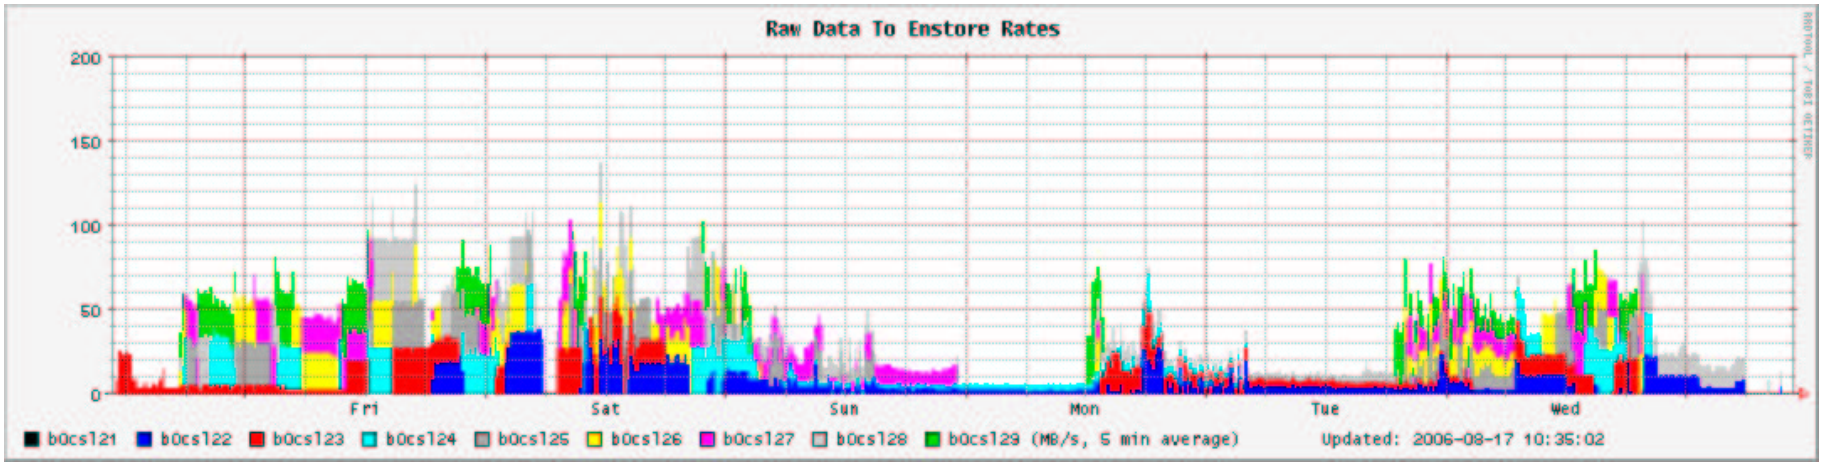

Able to write to Feynmann at 50 MB/s (only using 3 stagers)

→ In the final system we could have up to eight...

Tracking problems at:
http://ncdf76.fnal.gov/~chlebana/daq/cslUpgrade/integration/

Downstream logging

For software testing we are using null mover nodes.

Software same as used at $D\emptyset$

We have not yet written data to tapes.

New tape drives in Feynman being commissioned
→ *Could write to old tape drives but prefer not to...*

Switching back and forth between the old and new CSL should be straight forward.
→ *Switch destination in L3 Output node*
→ *File naming convention is data driven*
→ *File location is stored in a database*

Monitoring

Content of messages will mostly remain the same, sources of information will change...

→ *Logger processes send "LogMess" to Receiver*

→ *Receiver broadcasts "CSL Status" message as before*

→ *CSLMon listens for message, interprets and display statistics*

→ *Logger can provide additional information as text files*

→ *Content formatted for display on the web*

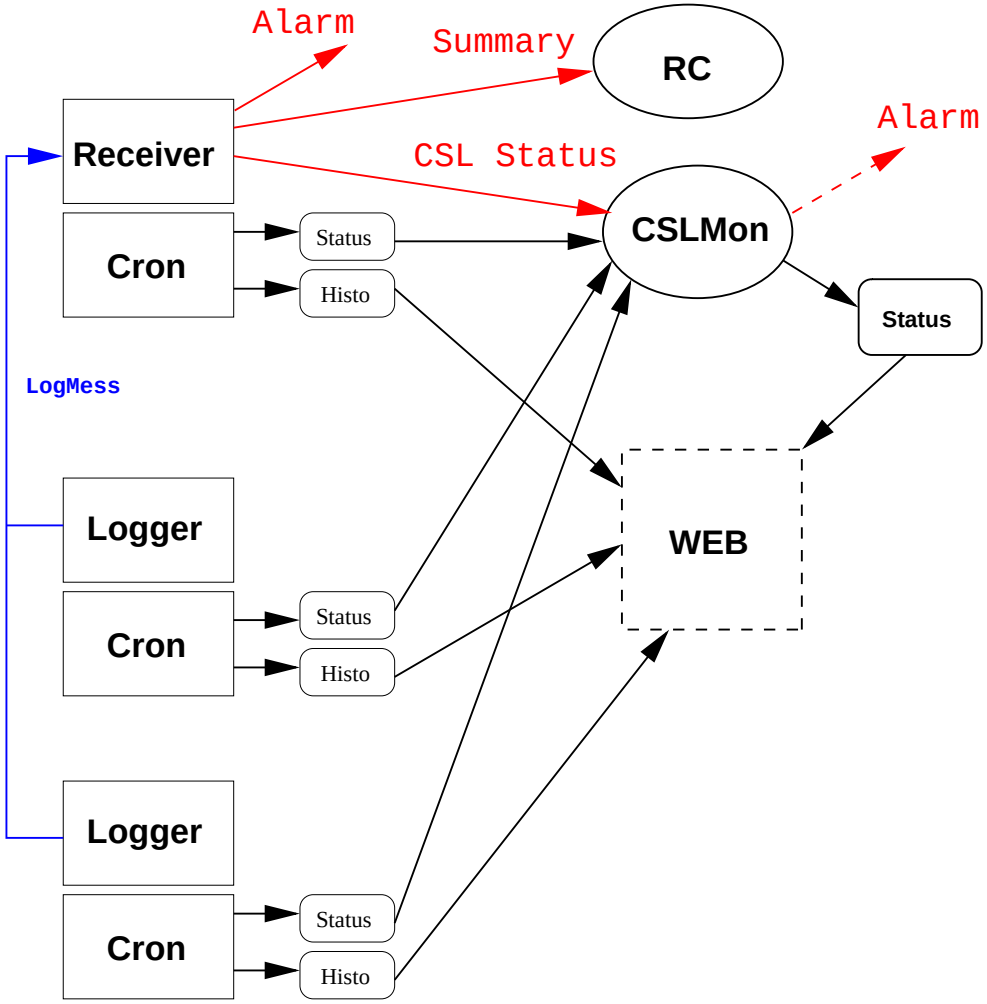

Updating CSLMon

New look same great taste...

→ *Reorganize content to make it easier for the shift crew to identify sources of problems*

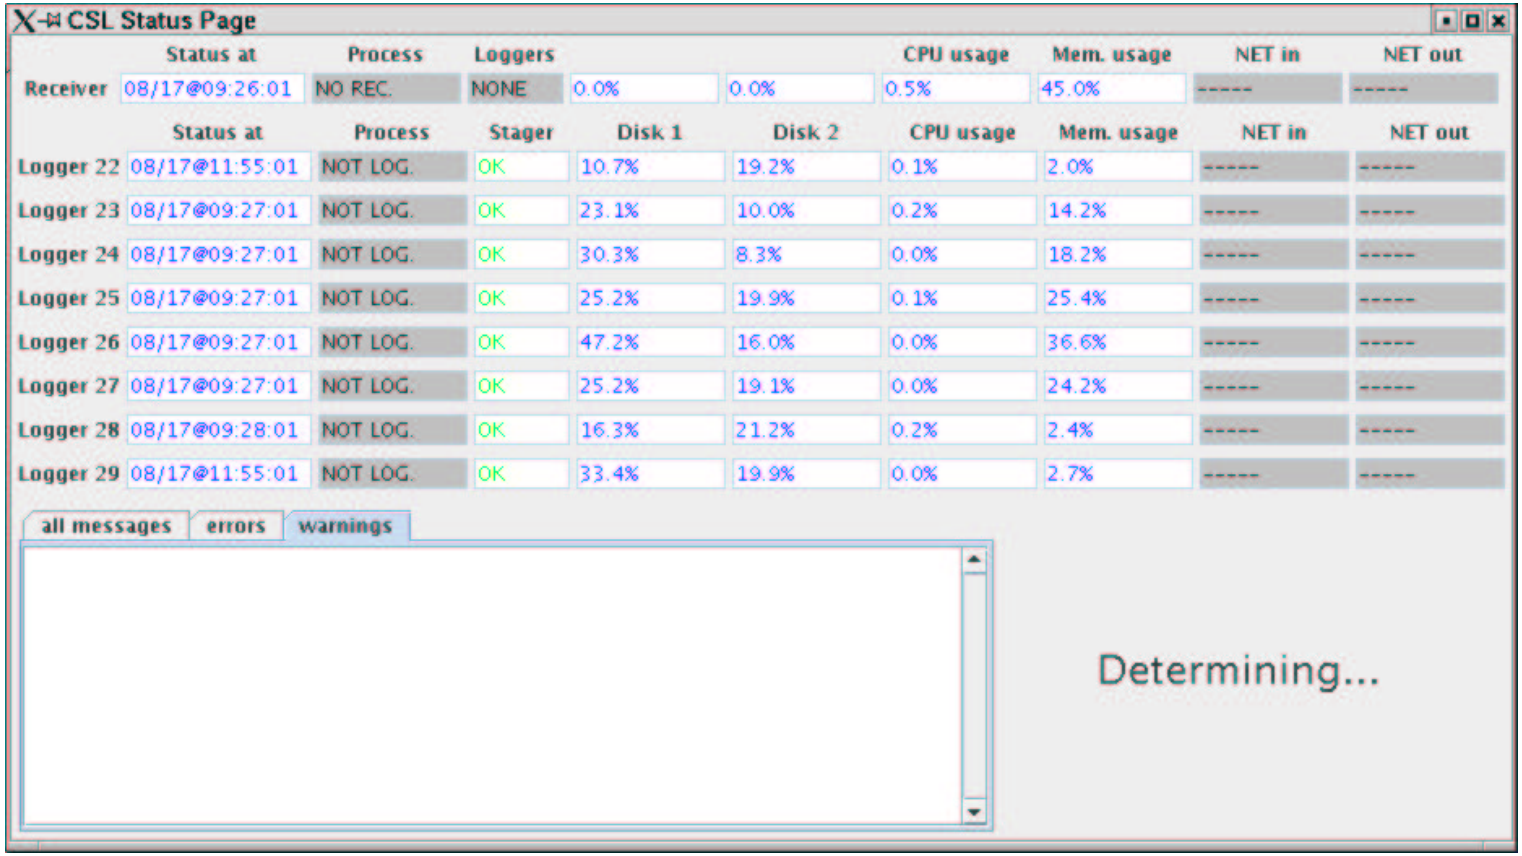

History monitoring plots available on the web

→ *Using rrdtools to generate plots*

→ *Using ganglia for system monitoring*

https://www-cdfonline.fnal.gov/internal/daq/cslmon/

Calibration CSL

We are investigating the possibility of combining the Calibration CSL and Data CSL.

Not sure we have the resources to do this...

Currently the Calibration CSL runs on a Linux server and does not use the same hardware as the main CSL.

Our main goals are to improve the logging rate, maintainability and reliability of the data CSL.

We can continue to use the calibration CSL

Summary

All equipment purchased and (nearly) installed.

Integration testing well on the way
L3Sender → Receiver → 8 logger nodes

Expect to work through problems using this configuration before interfacing with the full system.

So far hardware performing as expected and *well above* the designed target of 80 MB/s

Disk buffering:	$>$ 300 MB/s
Network:	$\sim$ 220 MB/s

Next major goals:

→ *Write data to tape*

→ *Get real data from L3*

Target delivery date: Sept 30th

Could use more time to ensure robustness