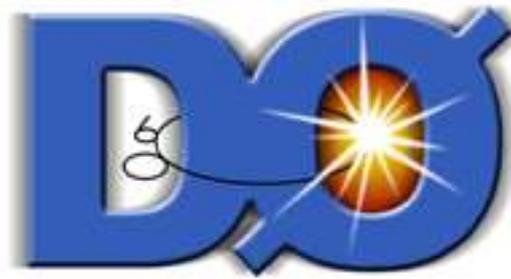


DAQ Shifter Tutorial: Level3/DAQ



Thomas Gadfort
University of Washington
DAQ Shifter Tutorial
October 11, 2005

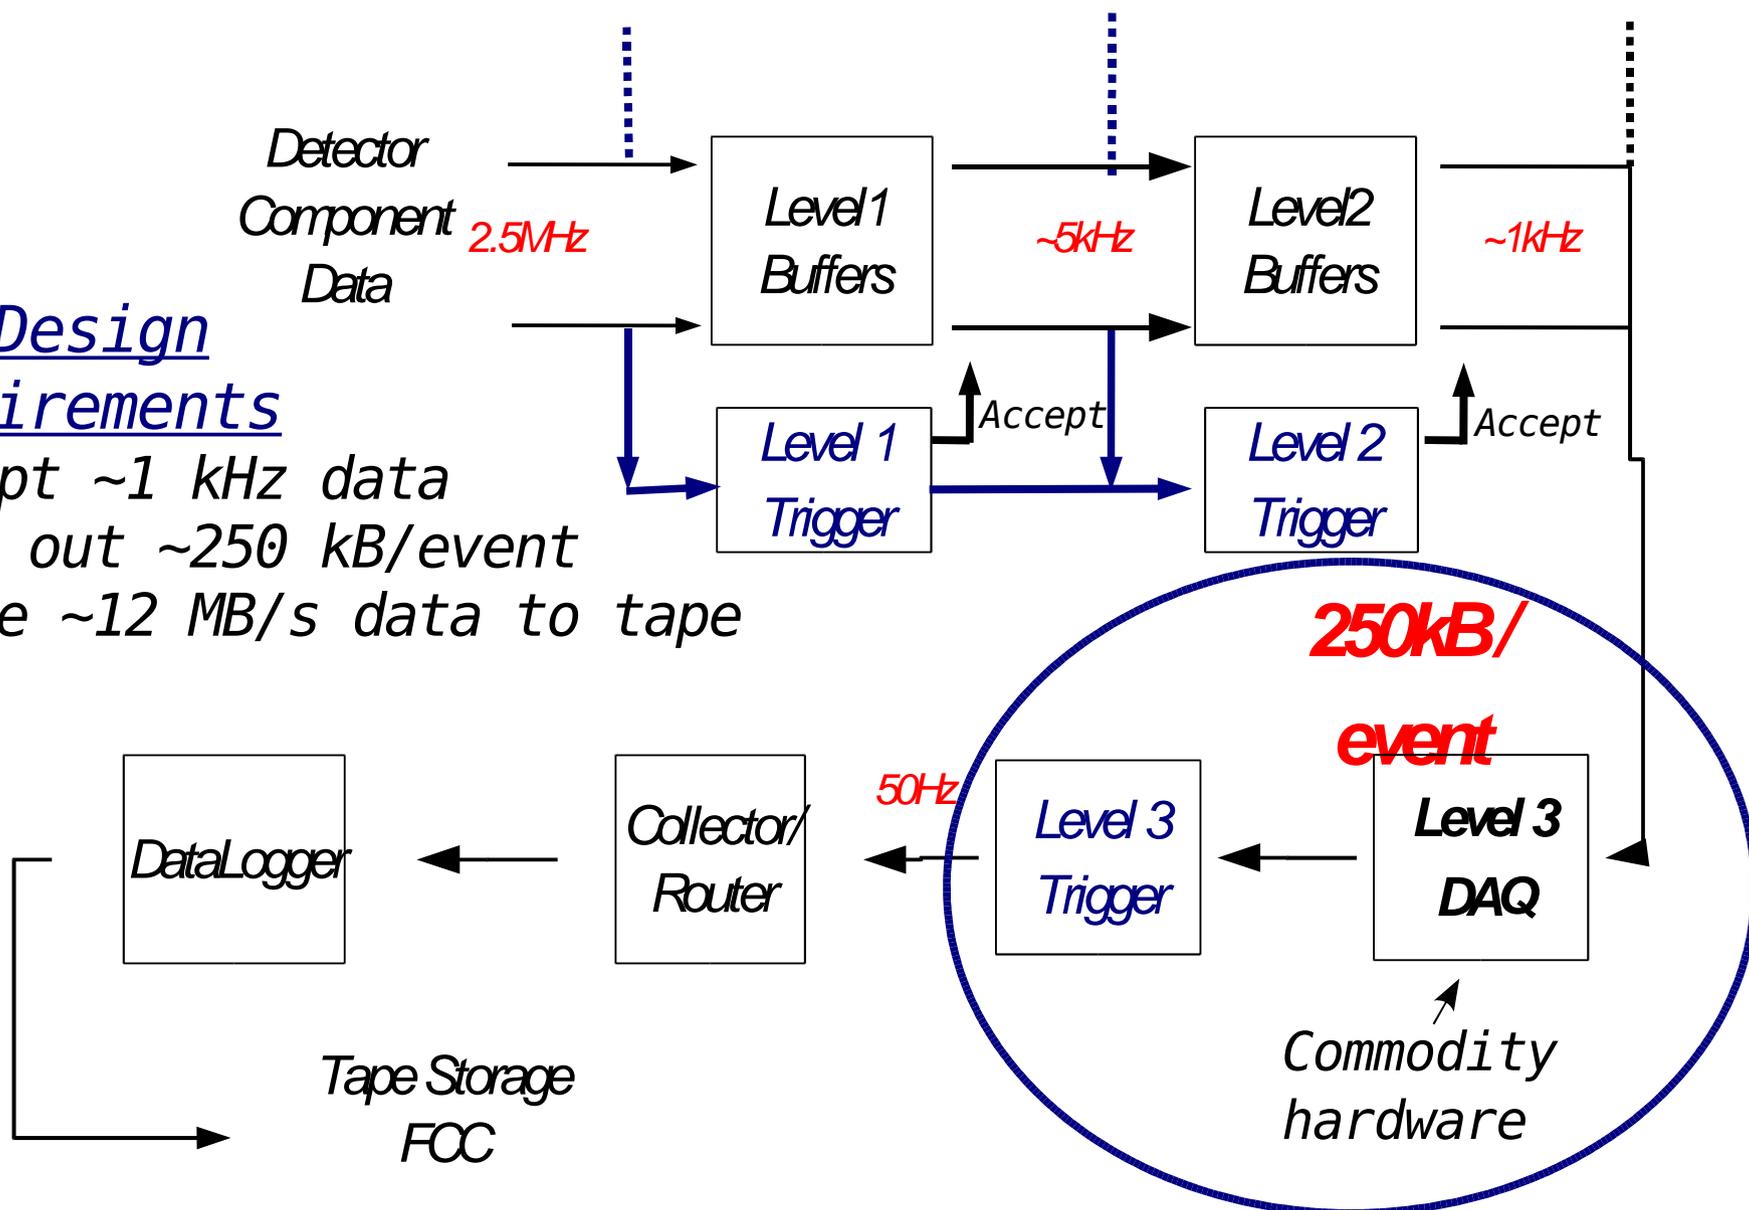
Outline

- D0 Data Acquisition System
 - L3DAQ Overview
- Level3 Components
 - SBC, Farm Nodes, Routing Master, Supervisor, SBC software, Event Builder
- Monitoring
 - Monitor Server, daqAI, üMon, füMon
- What To Do When
 - Starting/Stopping/Reseting SBCs or Farm Nodes, Contact List

D0 Data Acquisition System

DAQ Design Requirements

- Accept ~ 1 kHz data
- Read out ~ 250 kB/event
- Write ~ 12 MB/s data to tape

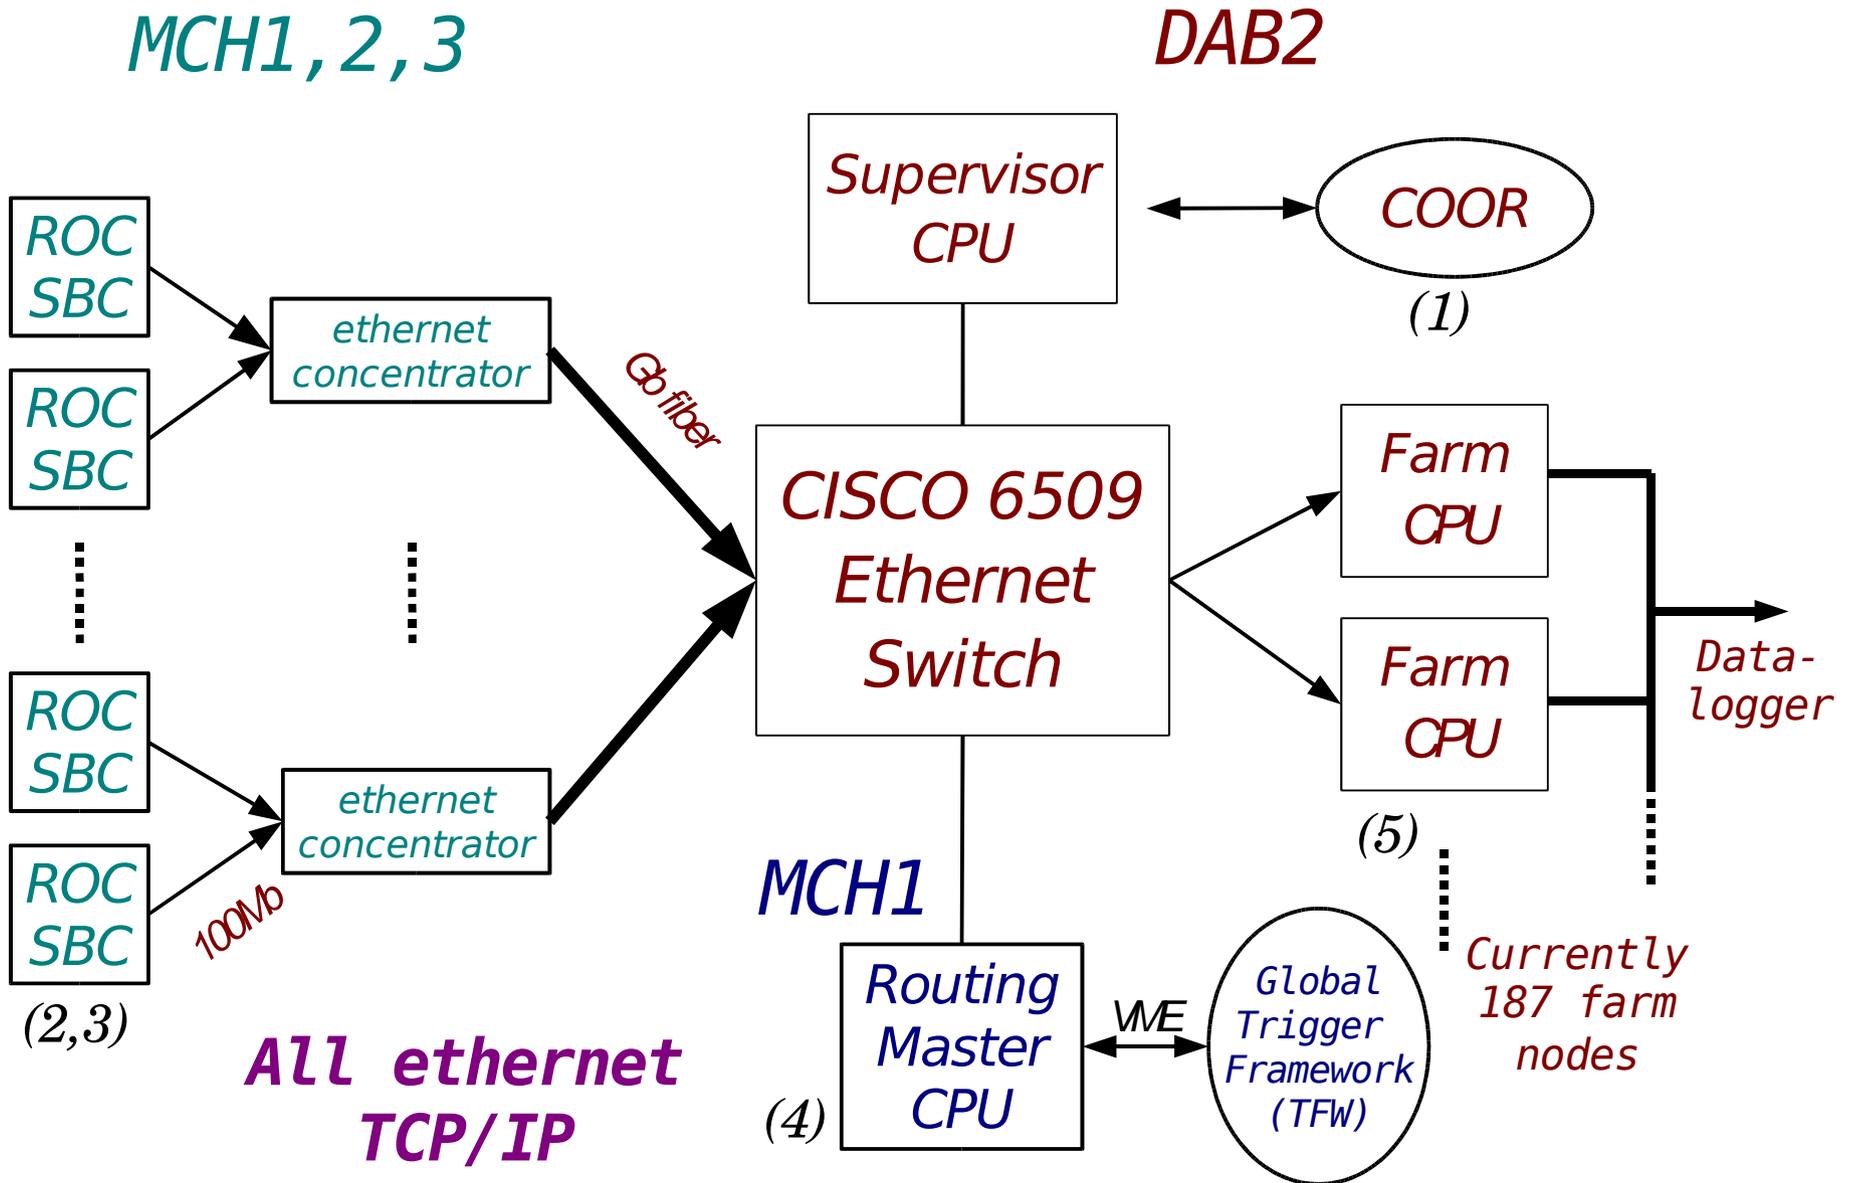


Level3 Components

Level3 has both a hardware and software side

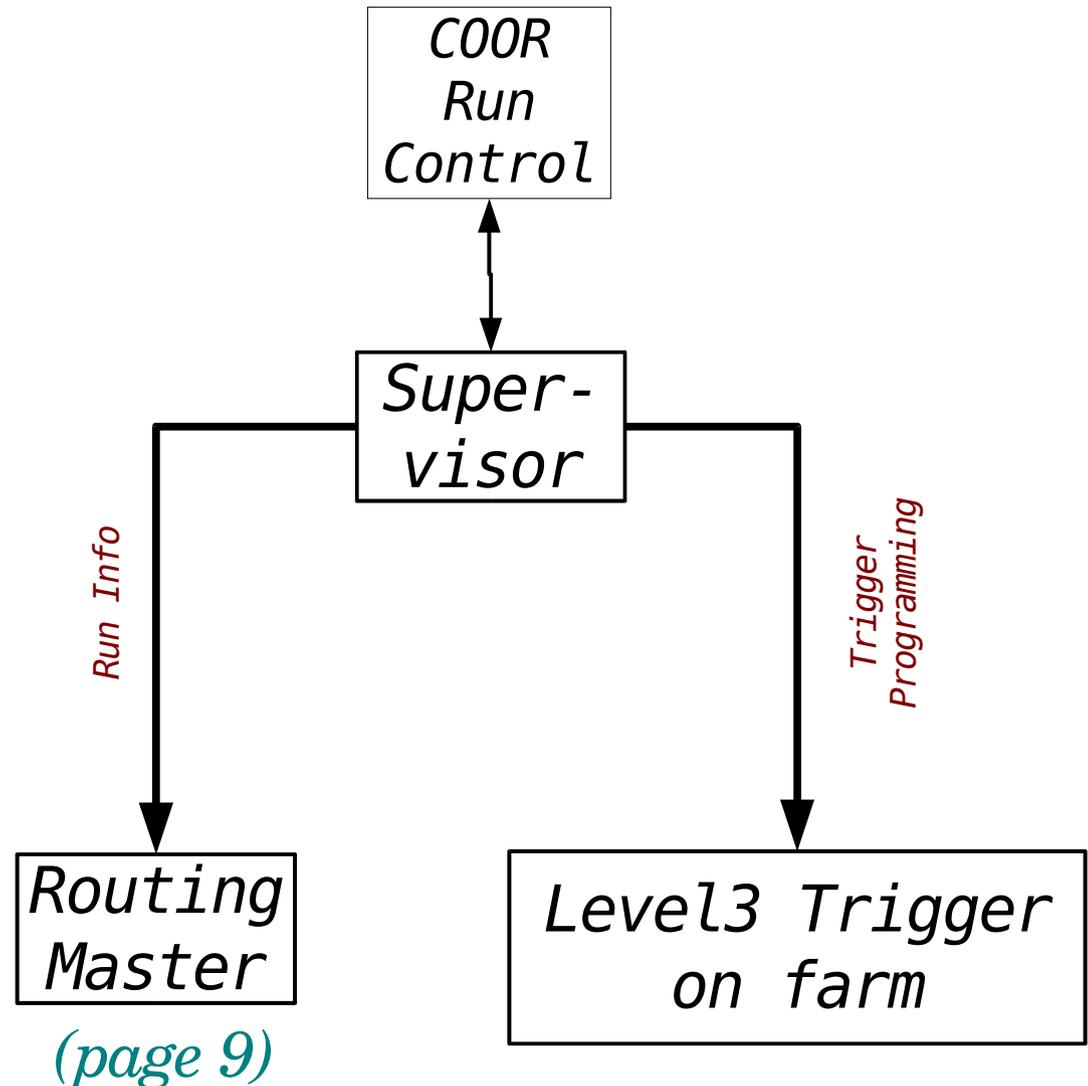
- Hardware
 - SBCs
(Single Board Computer)
 - 2948G Ethernet concentrator switches
 - 6509 CISCO Ethernet switch
 - Farm Nodes
- Software
 - Routing Master
 - Supervisor
 - Event Builder
 - SBC Event Sender
 - Level3 Trigger
 - Monitoring

Level3/DAQ System Flow Chart



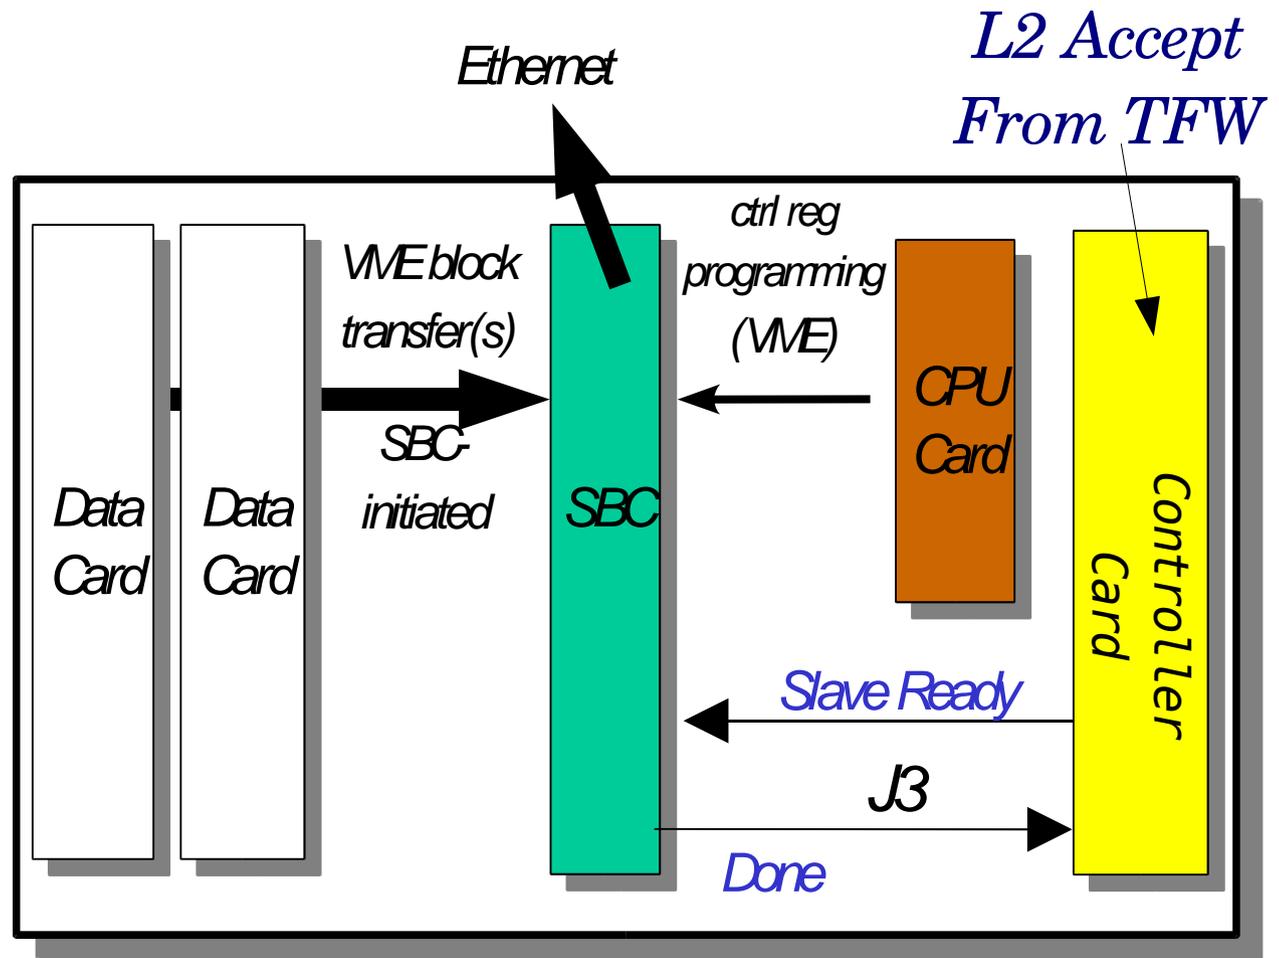
Start/Stop/Pause a Run – L3 Supervisor

- You decide you want to start a run
 - Pick crates with CRATER and download trigger with TAKER
 - COOR tells the supervisor who is in the run
 - COOR checks that there are a minimum of 9 nodes in the run
 - Supervisor sends the crate list, trigger info to the RM
 - Supervisor sends the trigger programming to the level3 trigger running on the farm



VME Readout Crate

- In the run, we now want to read out data if L2 accept
 - Controller card checks crate data size
 - Level2 accept comes from trigger framework
 - Controller card asserts a slave ready to the SBC
 - This tells the SBC to read crate data cards.
 - Data is stored in RAM memory on the SBC
 - SBC event sender software starts running (page 10)



- Data card, CPU, and Controller cards specific to crate type

Single Board Computer (SBC)

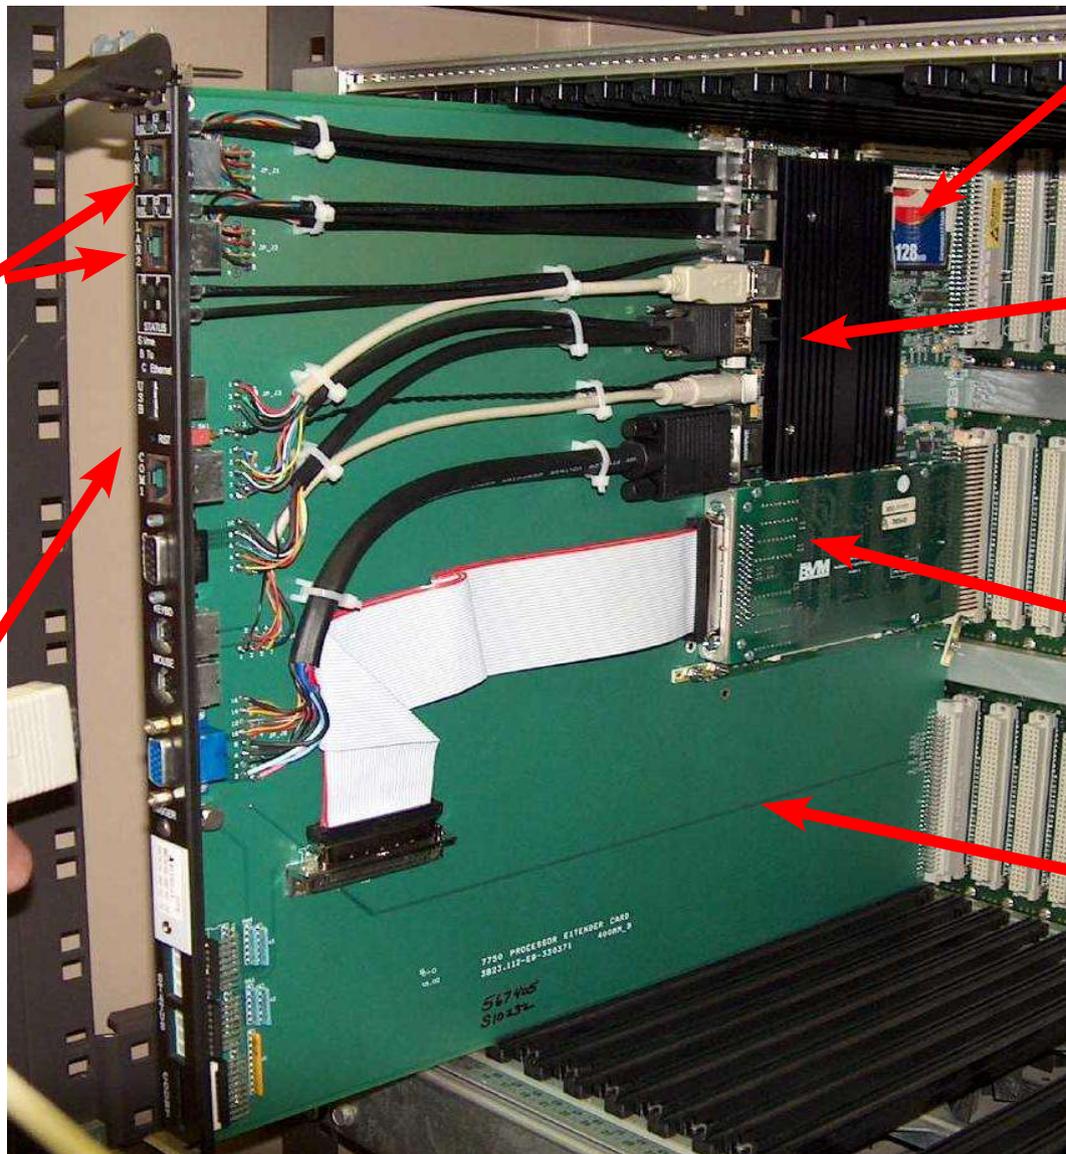
PIII 933MHz
128MB RAM
128MB flash

Dual 100Mb/s
ethernet
(24MB/s)
(Data sent to
farms over
ethernet)

(page 10)

Reset
Button

*(We'll talk
about this
later)*



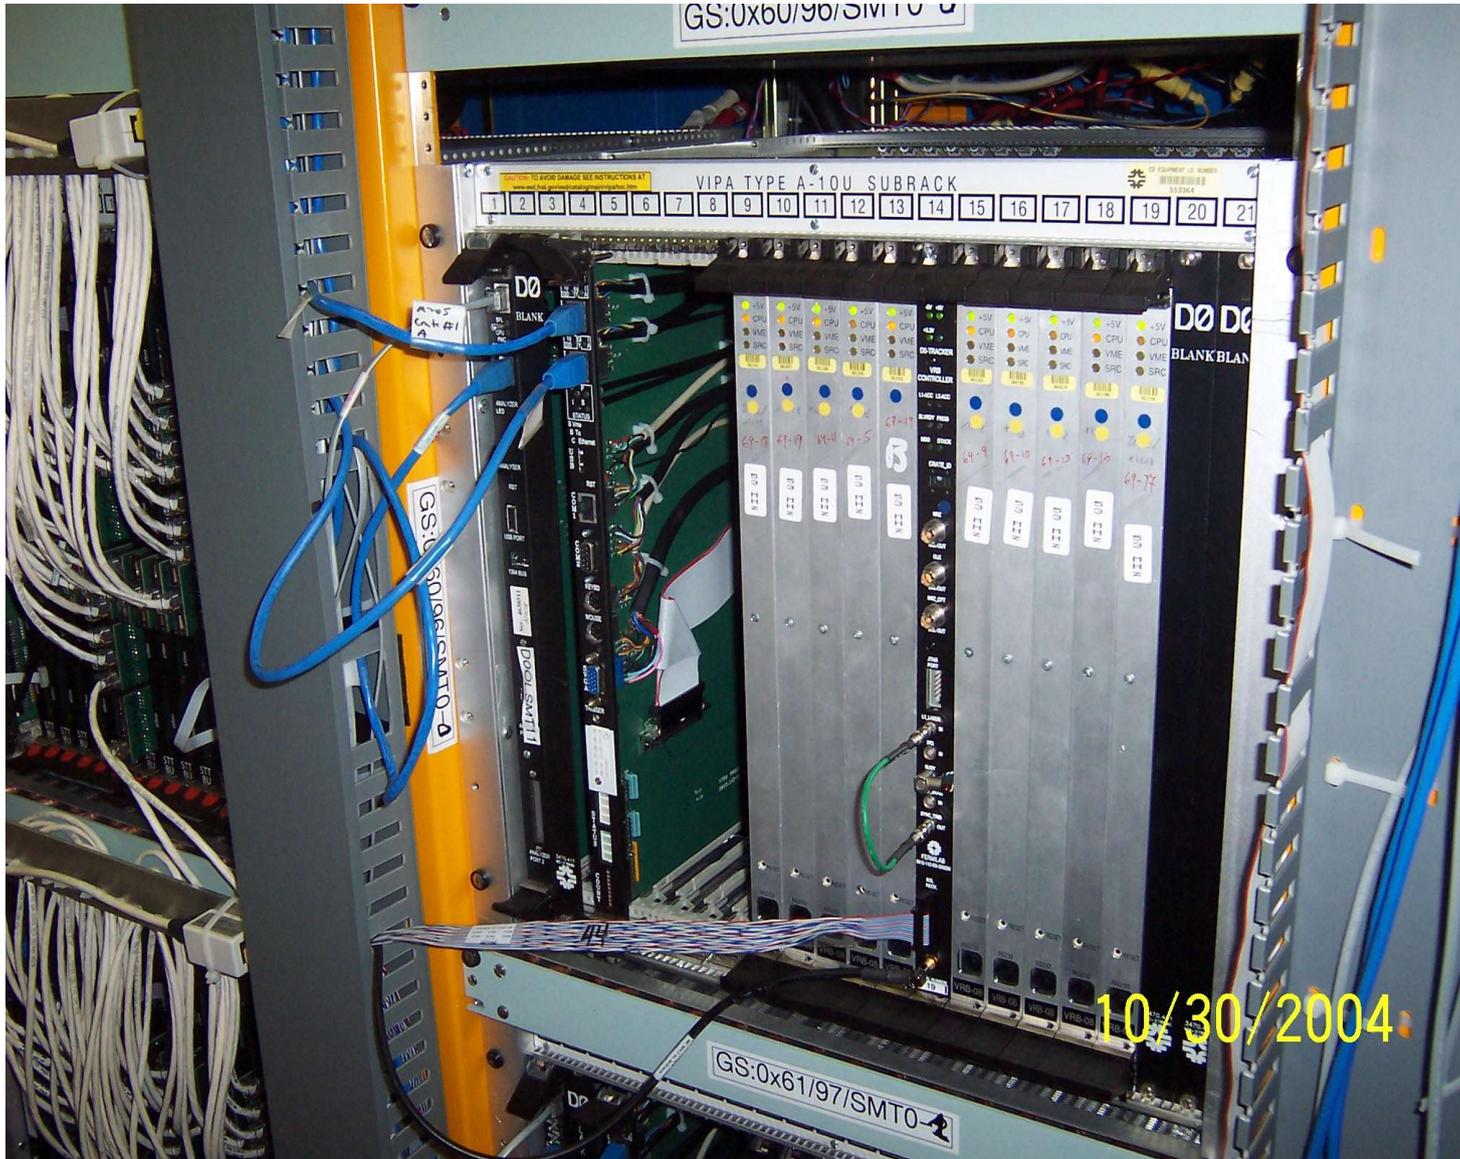
128MB Flash ROM
•Stores
configuration file

Tundra Universe II
•PCI-VME Interface
•DMA controller

PMC Digital IO Card
•Coordinates VME
Readout

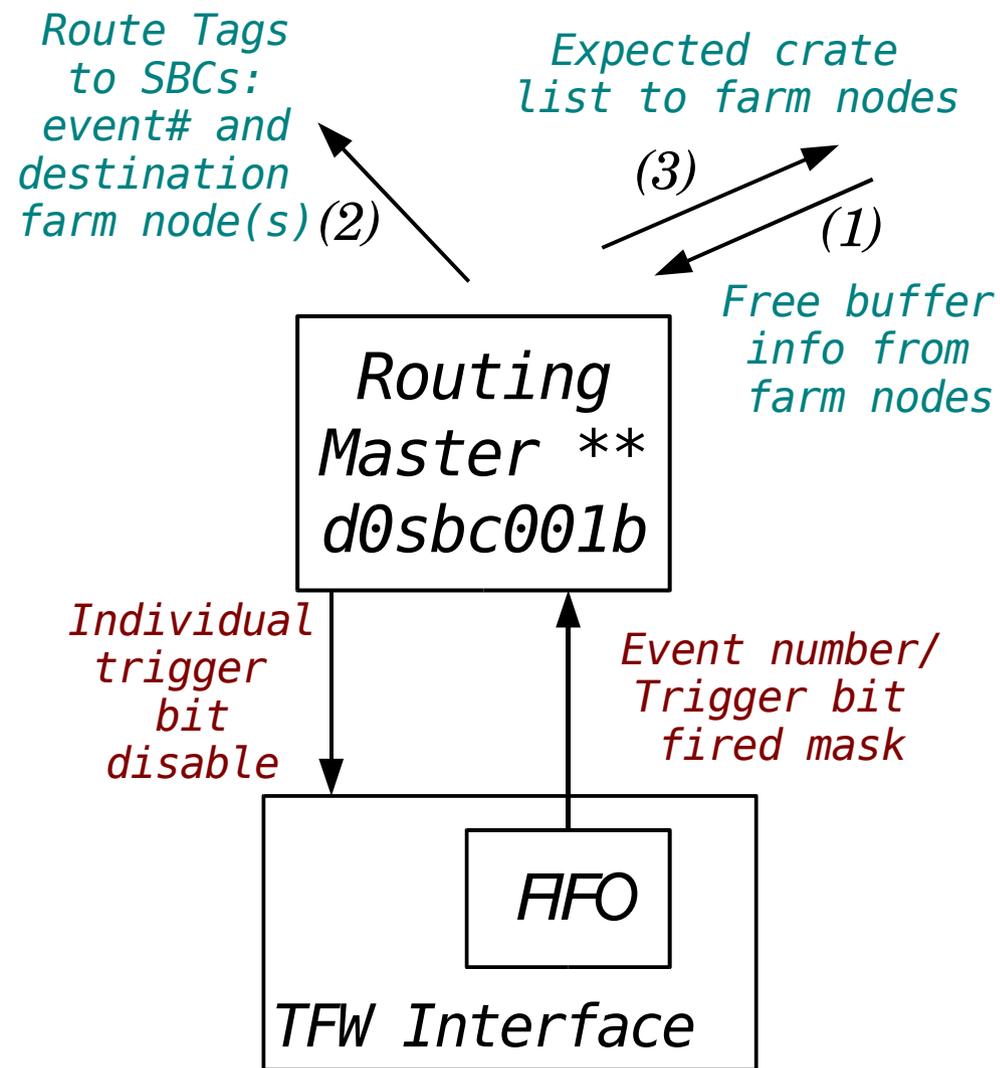
J3 Handshake
Slave Ready &
Done Lines

SBC w/ VME Cards in ROC



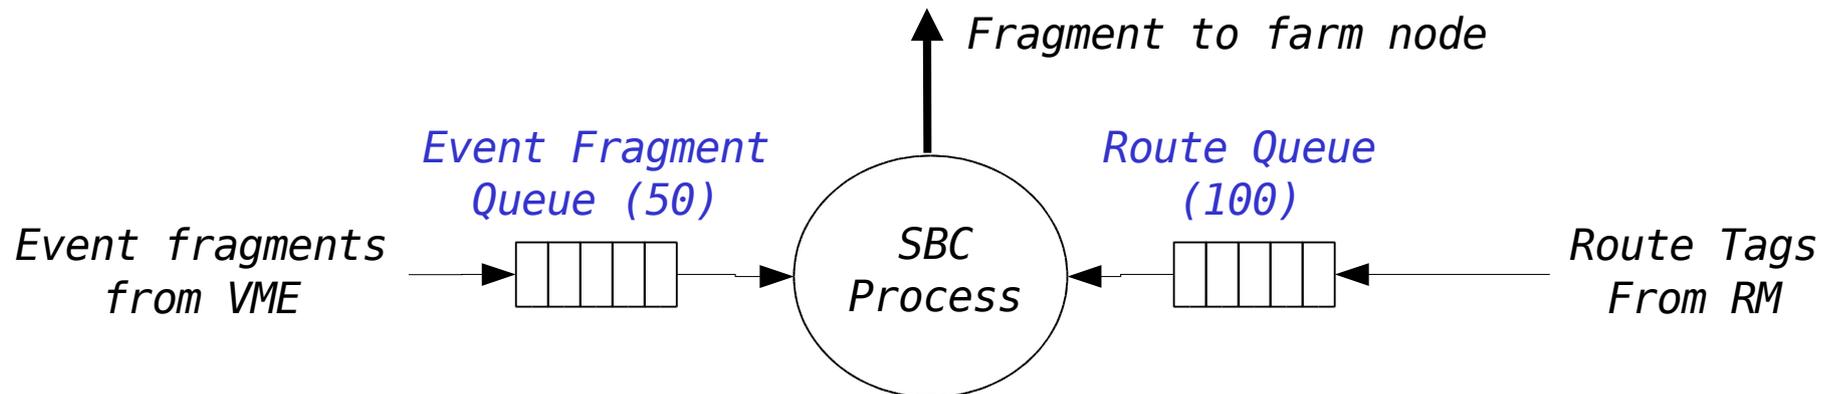
Routing Master

- Serves two purposes
 - Connect to TWF
 - Event routing
- Event Routing
 - Decides which farm node gets the event based on free buffers from farm nodes
 - Sends event number, expected crate list, and fired triggers
 - Tells SBCs which nodes to send their data
- Connection to TFW
 - TFW sends event number and fired trigger bit mask
 - Can disable all L1 triggers if number of free buffers is too low (back pressure)



** Dedicated SBC (001b) Located in MCH1 Rack 100 (If problem, CALL!!!)

SBC Event Sender Process

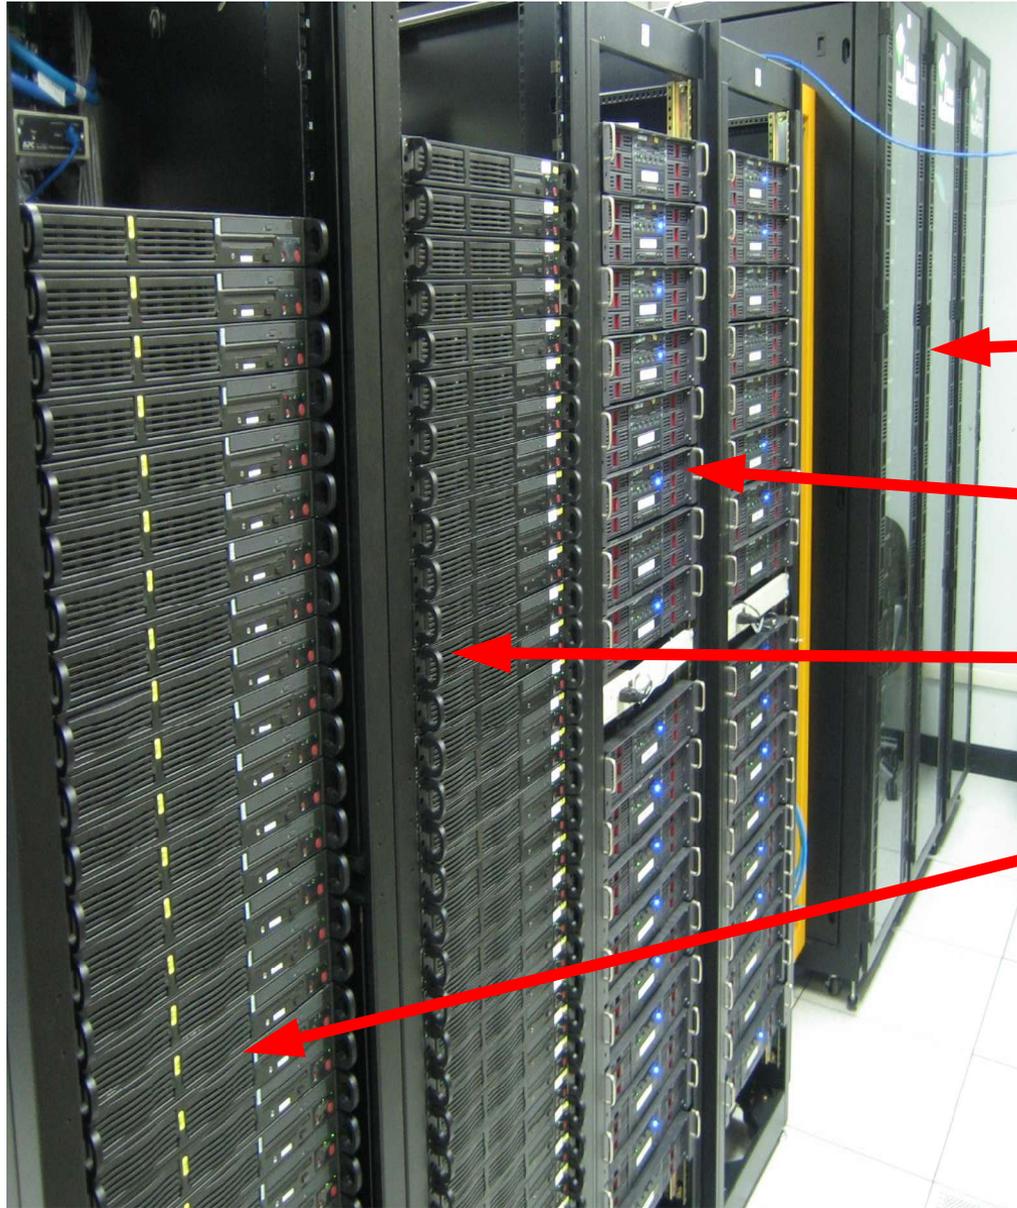


- Event fragment queue stores event fragments from VME crate with event number tag
- Route queue stores routing tags from the RM with event number tag
- SBC Event Sender Process compares the head event fragment tag with the head routing tag
- If the tags match, then it will send the fragments to a particular farm node. Otherwise it will drop the tag with the higher event number tag

- This causes the crate box in üMon to go red



Level3 Farm



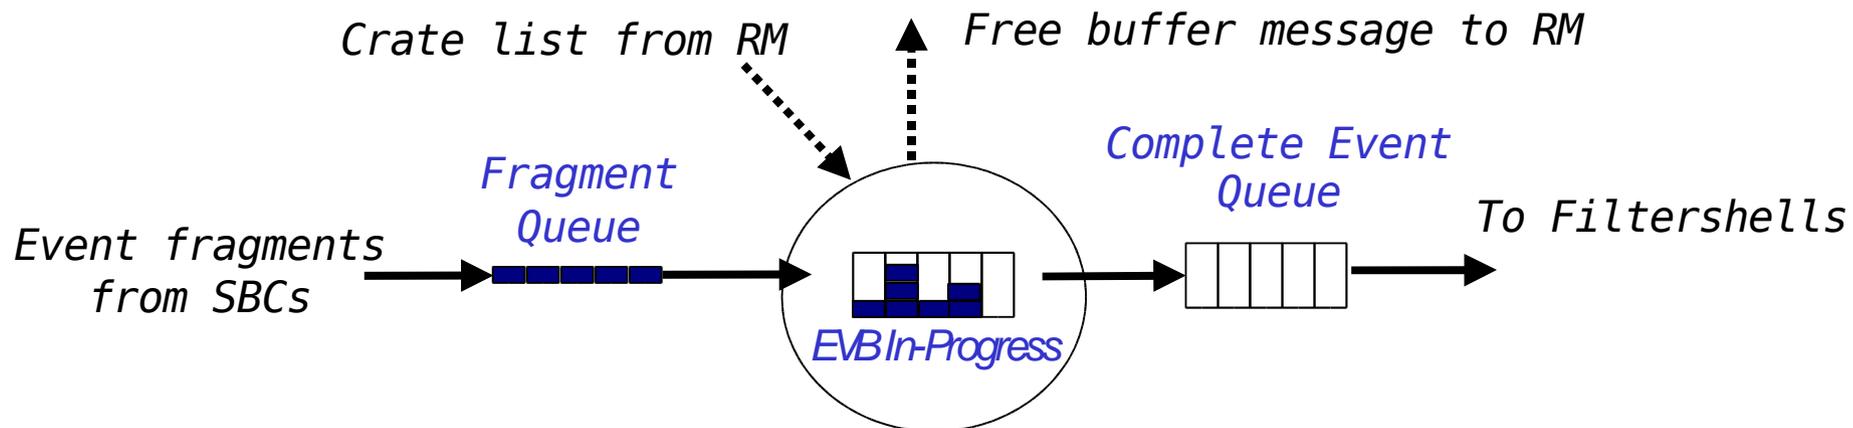
- Runs level3 software
- 114 + (32) Nodes + 128 (next page)
 - 48 Dual 1 GHz PIII Processors (nodes 1-48)
 - 34 Dual 2 GHz AMD Athlon Processors (nodes 49-82)
 - 32 Dual 2.8 GHz Xeon Processors (nodes 83-114)
 - 32 2.7 Ghz CAB Nodes (on loan)
- Dual 100Mb ethernet ports
 - One connects to L3DAQ (in)
 - One connects to online (out)

New Farm Nodes

- 128 2.7 Ghz Xeon Machines



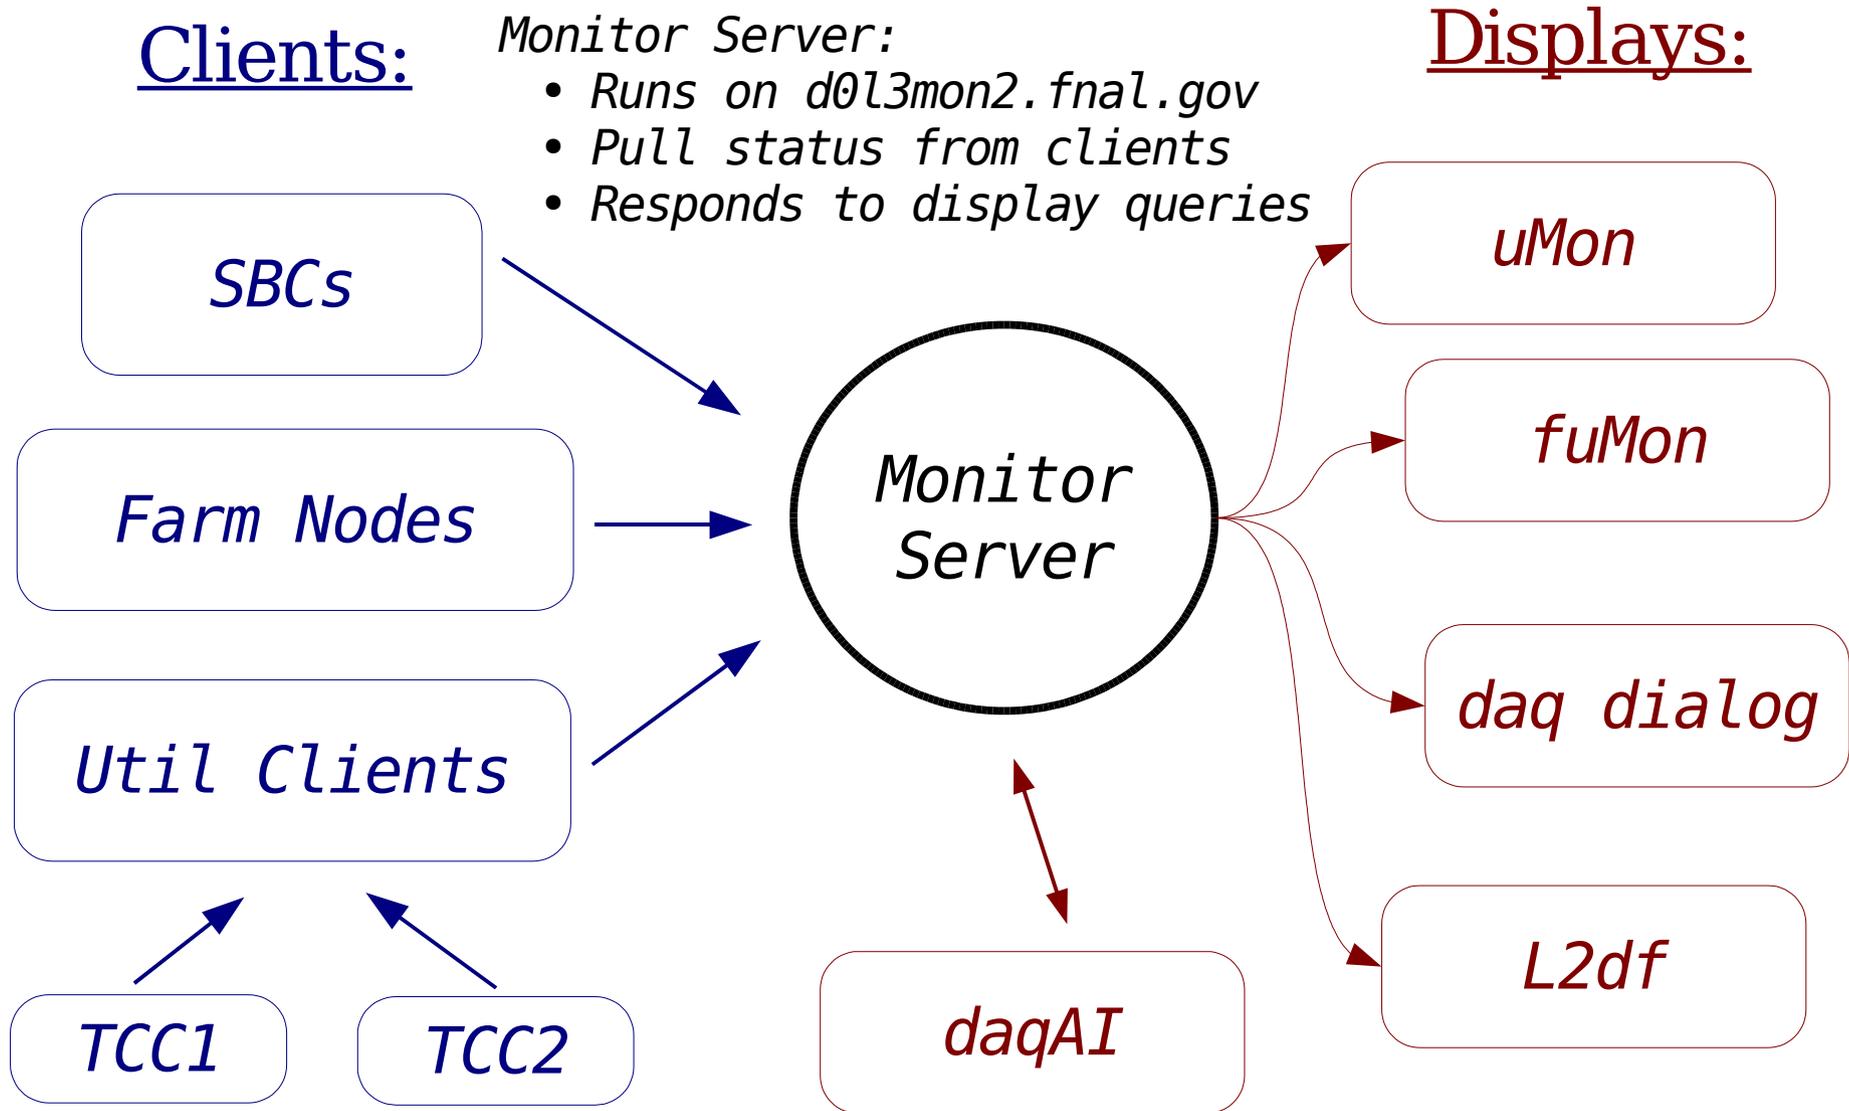
Event Builder



- Recap: The RM told each SBC where to send their data and it told the particular farm node which SBCs are sending data.
- EventBuilder strings together all the SBC data fragments into an “event”.
- Now at level3, all detector components are available
- Event will wait for 1 second for the SBCs to send their data.
 - If it does not receive what is expected, the event is dropped
 - This will show up as a red node box in füMon.



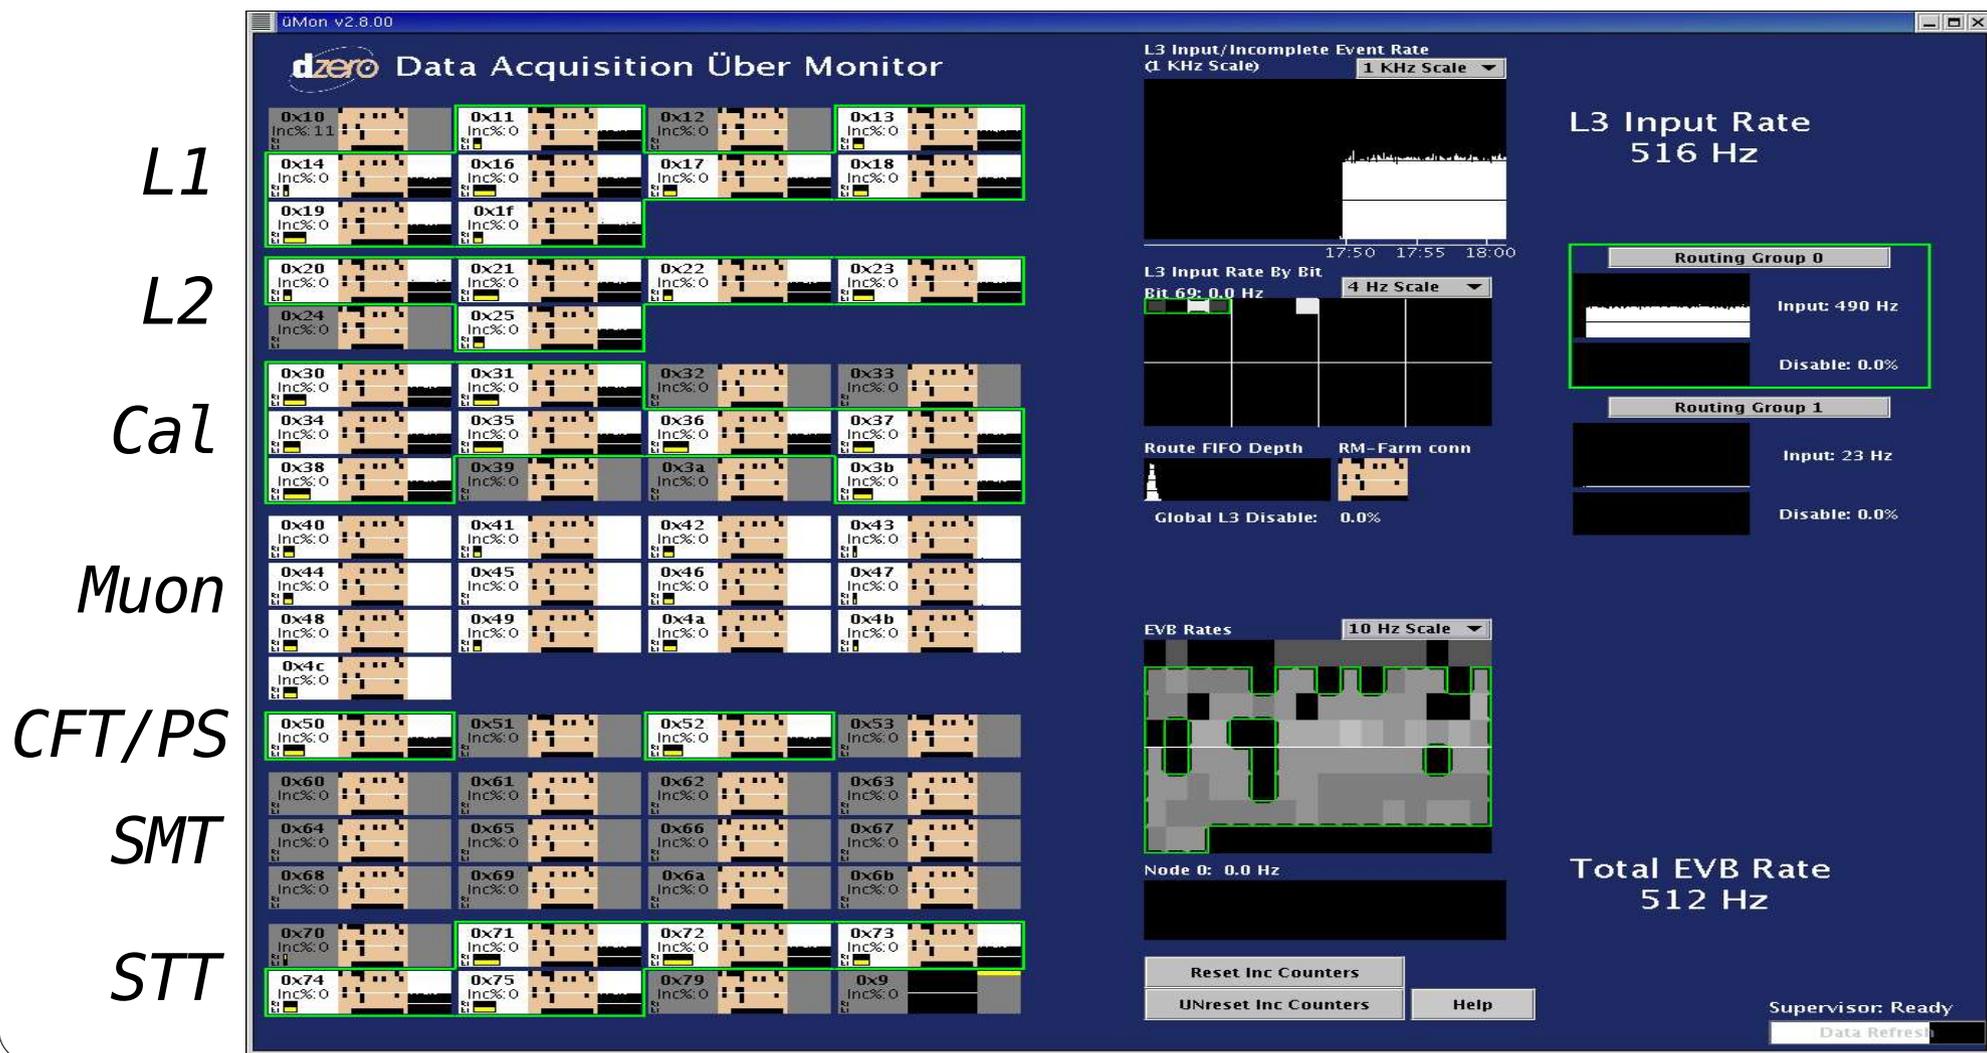
L3DAQ Monitoring – 99% of your time



Monitoring Tools: ümon – SBC monitoring

Sits at top left black monitor at daq console

- Outlined green crates are “in the run”*



Monitoring Tools: ümon – Crate Info



Farm node Connections:

- *Beige means live connections*
- *Black means no connection*
- *Lower right is RM connection*

*Crate
Number*

*Event
Size*

*Route queue
depth*

Event queue depth



*Rate at
which SBC
reads VME*

Monitoring Tools: ümon – Various States



*Normal crate
in the run* →



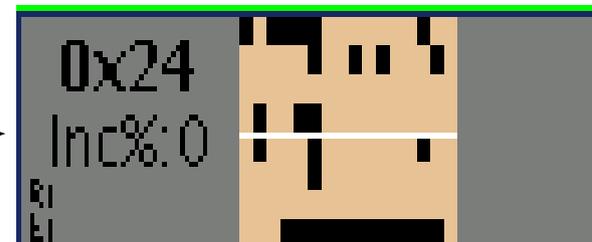
*Mis-matched
RQ & EQ: Drop* →



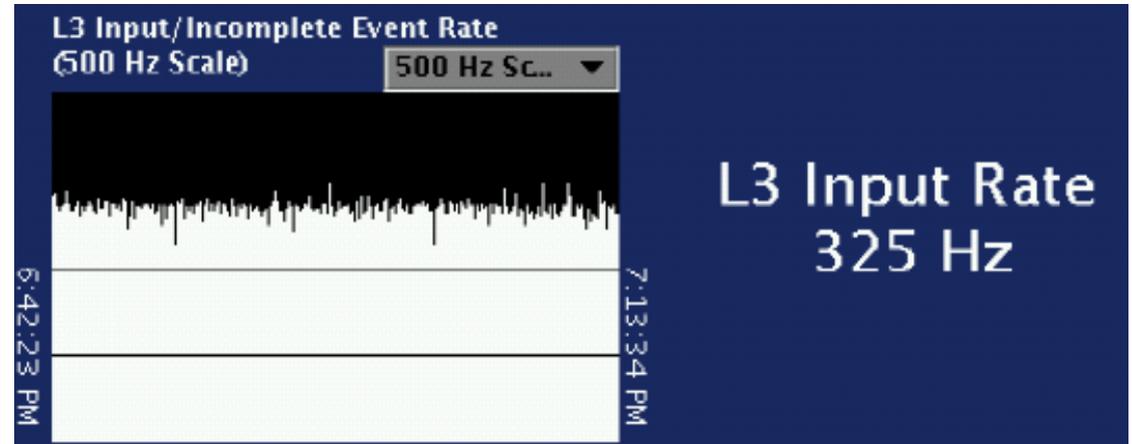
*Crate is not in the run &
there is no monitoring info* →



Crate is not in the run →

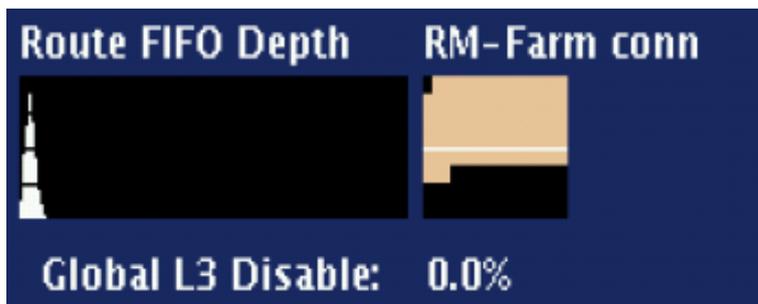


Monitoring Tools: ümon – Other



Level 3 Input Rate

*This is the RM send/receive rate
Typically ~700 Hz for physics runs*



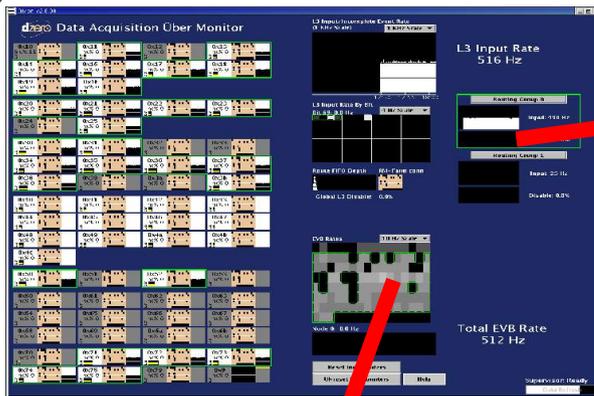
RM-Farm Connection:

*Beige means RM is connected
to that farm node*

FIFO Depth:

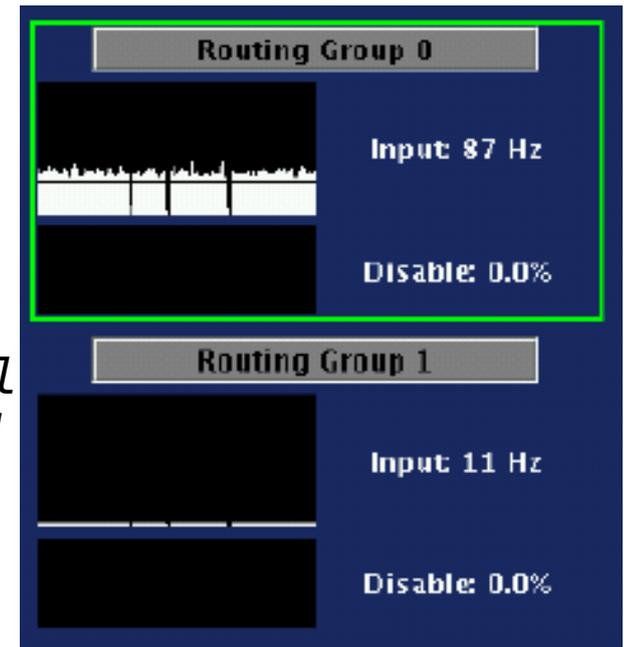
*Histogram of the number of events in the
TFW FIFO as polled by the Routing Master*

Monitoring Tools: ümon – Other



Routing Groups:

- Input rate and disable Fraction for each group
- Clicking on a group will Highlight the associated crates



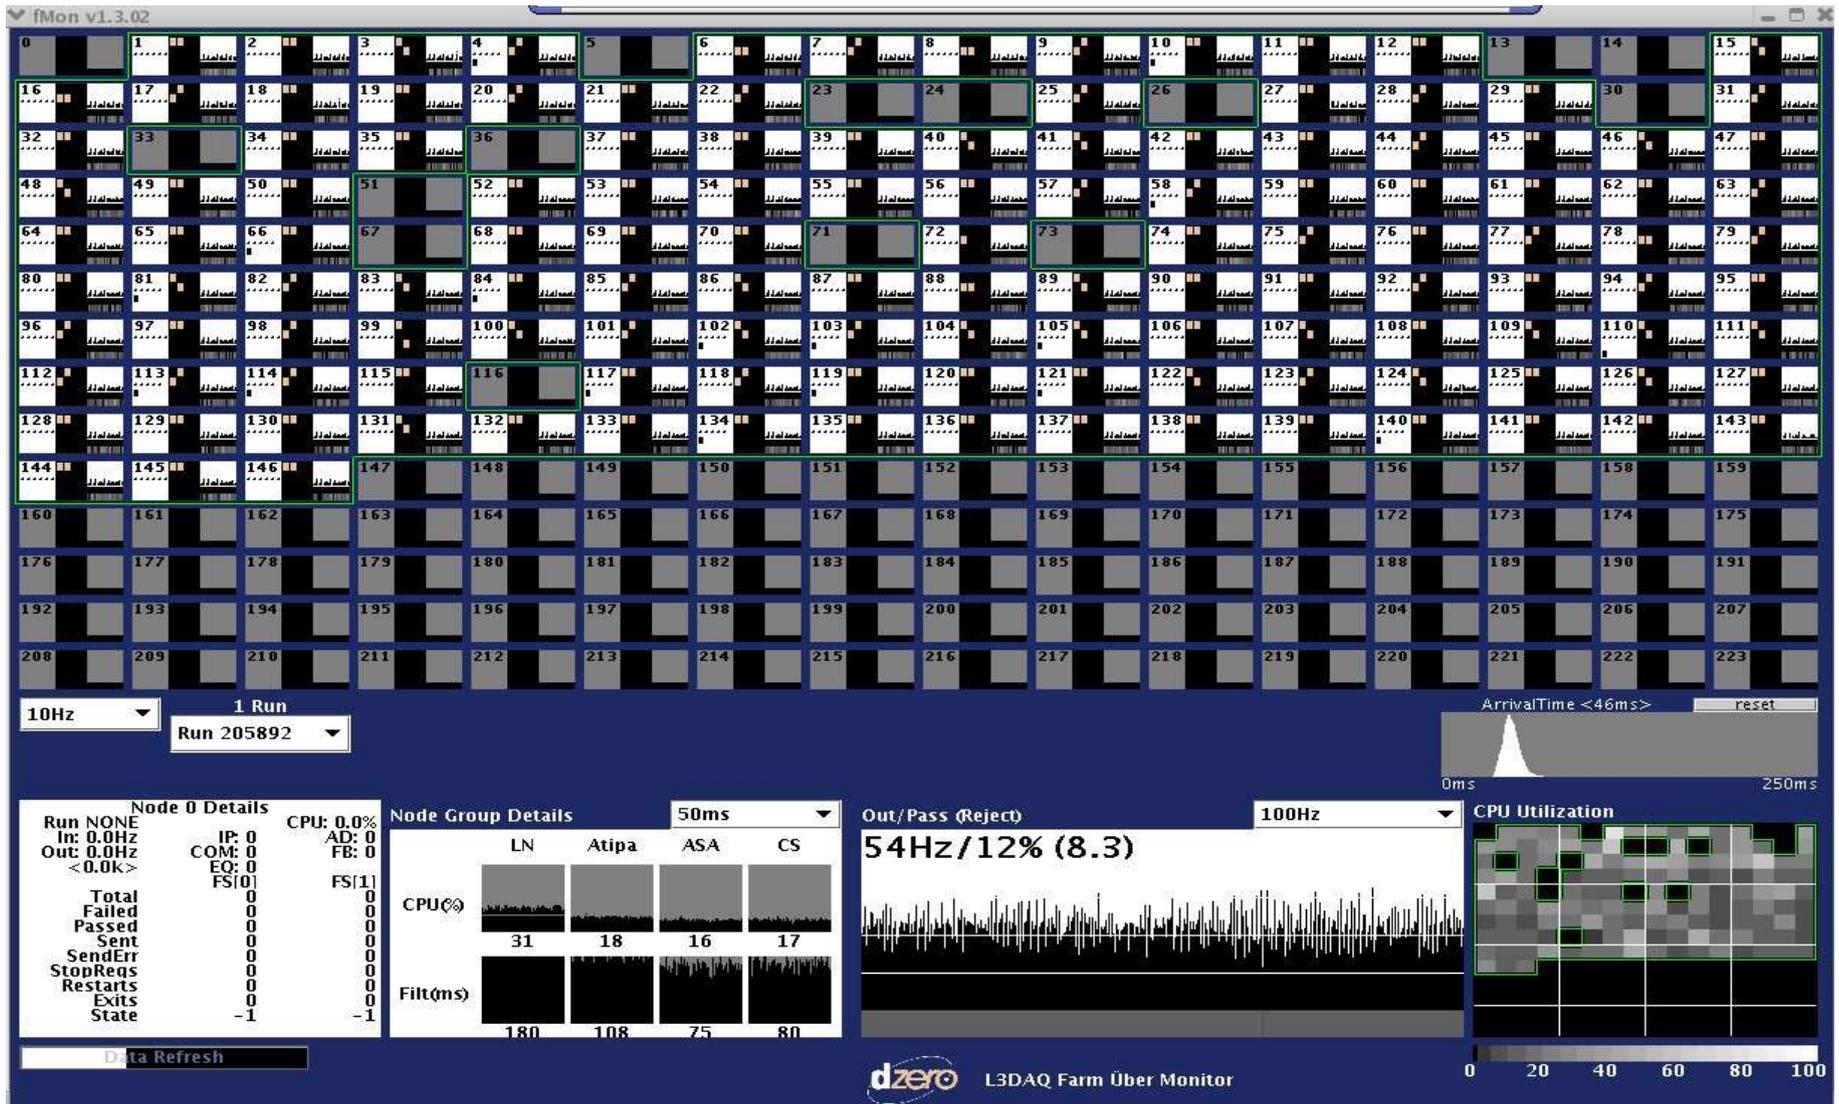
EventBuilder Rates:

- This is the EventBuilder input rate for each farm node. Scale is chosen by the pull-down menu
- The strip chart shows the total EVB rate over the last 20 minutes for a particular node
- Also shown is the total EVB rate.

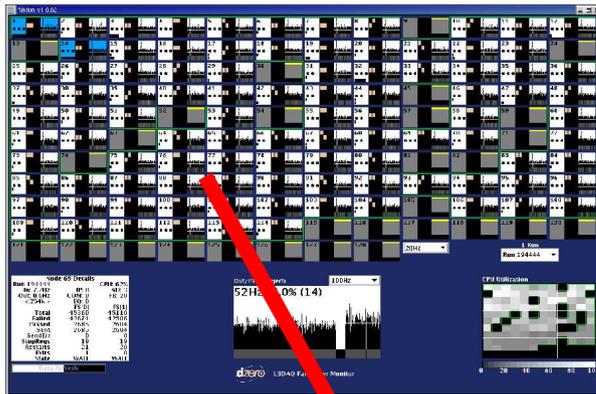
Monitoring Tools: fūmon

Sits at top right black monitor at daq console

- *Outlined green nodes are “in the run”*



Monitoring Tools: fūmon – Node Info



Two beige boxes represent the two Level3 filtershell states:

- 1st Waiting
- 2nd Filtering
- 3rd Flattening
- 4th Sending

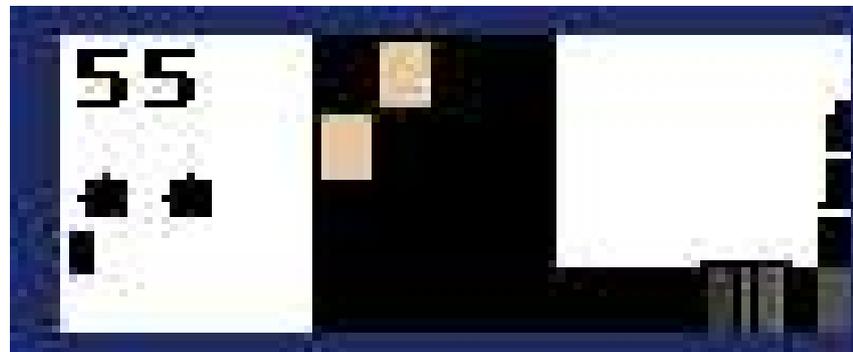
Node Number

Event rate into filtershell

Number of advertised buffers

EVB buffer status

Filtershell pass fraction



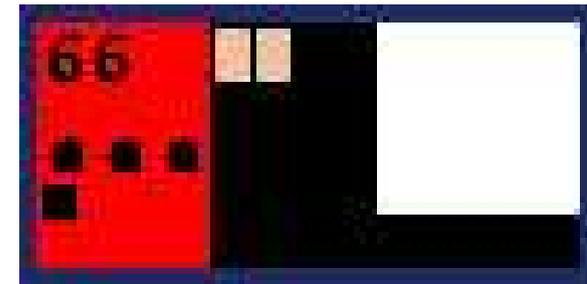
Monitoring Tools: fūmon – Various States



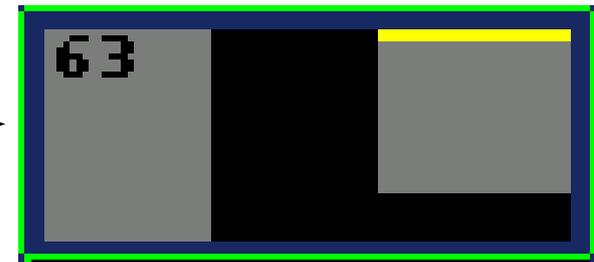
*Normal node
in the run*



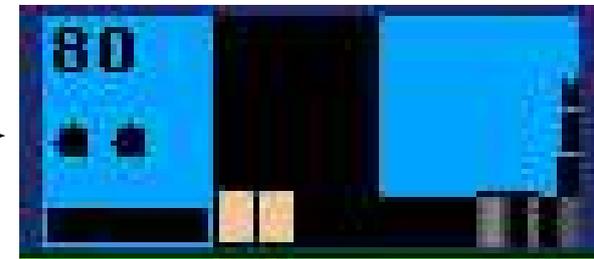
*Missed event
from SBC: Drop*



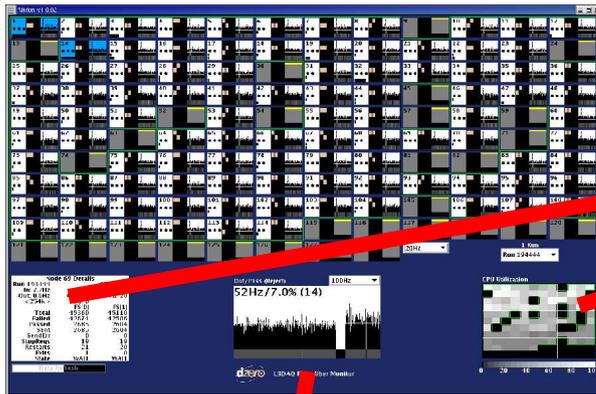
*Node is not in the run &
there is no monitoring info*



*10 or more EVB buffers
are filled up*

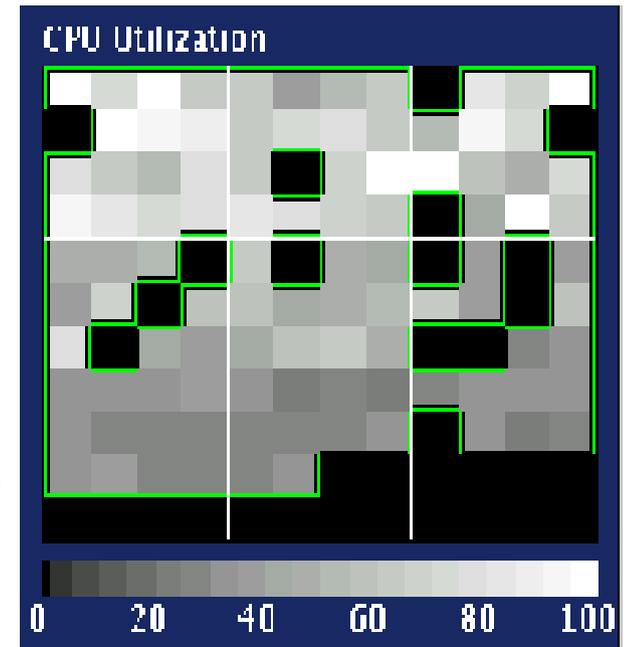


Monitoring Tools: fūmon – Output/CPU



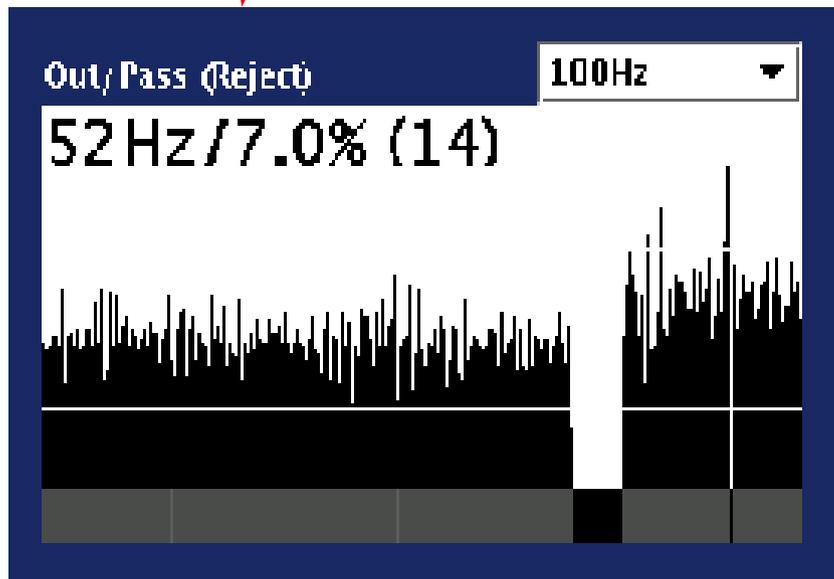
CPU Utilization:

- This is the % use of the CPUs on the nodes

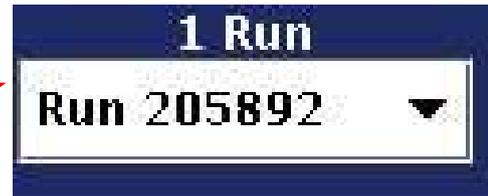
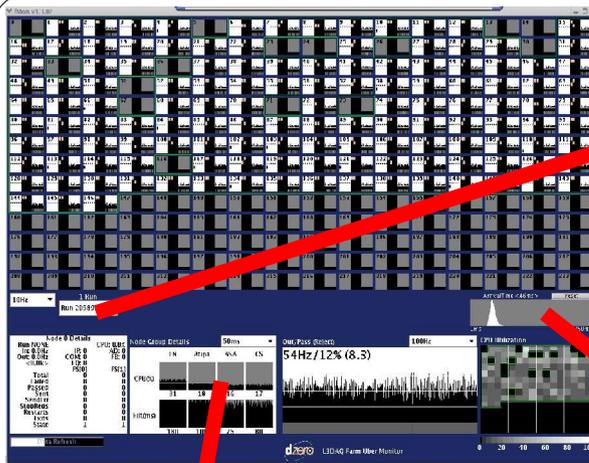


Level3 output rate:

- The 52 Hz is the rate at which events are passing the l3 filter
- The 7.0% is the pass fraction
- The (14) means 1/14 events pass the level3 filter



Monitoring Tools: fūmon – Filter Times

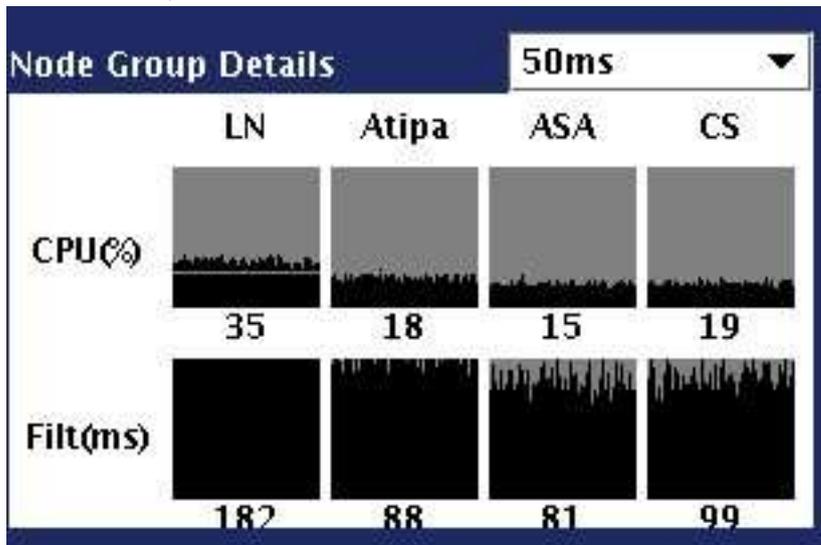


RUN Info:

- How many runs
- Run number



Arrival Time Info: Last crate arrival time (ms)

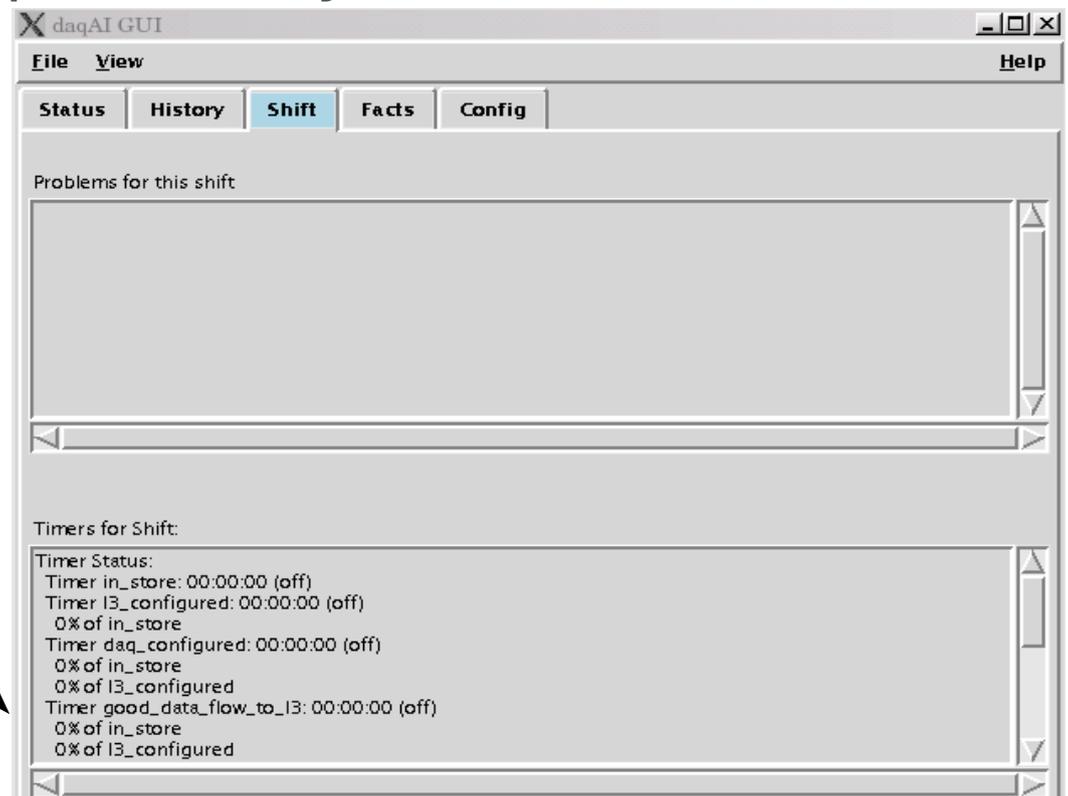


Node Group Info:

- The average CPU usage (%) for each node group
- The average Filtering time (ms) for each node group

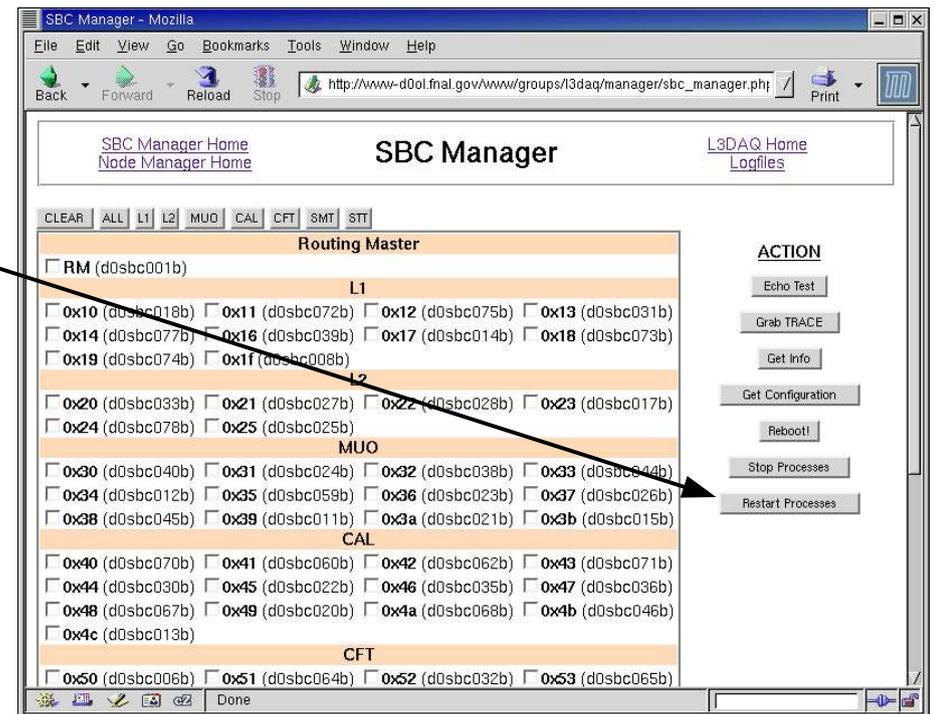
daqAI

- Runs on predefined rules in the background
 - Checks for crates loosing sync, sbcs that crash, missing slic inputs, etc...
- The Stephen Hawking style voice you will hear from time to time.
- Takes care of most of the problems you'll encounter on shift.
 - SCLimits (x24 Lost sync)
- Check your efficiency.
 - Gets you on Arnd's RunII best webpage
- Only on during physics runs
 - > start_daq daqai
 - > start_daq daqai_gui



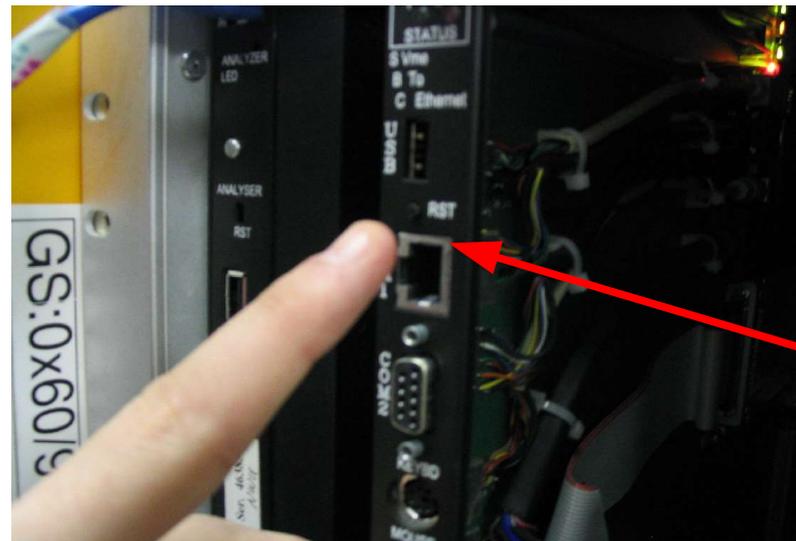
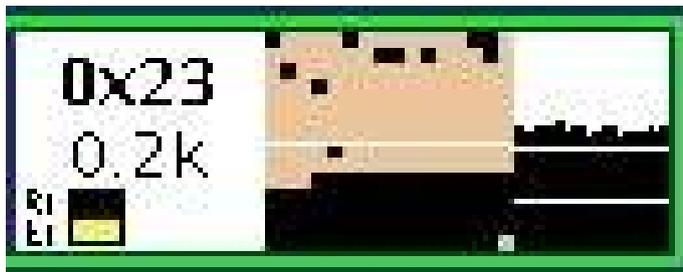
Common Problems – SBC Reset

- DaqAI tells you to reset an SBC. How do you do that?
 - From the command line:
 - > l3xdaq_reset XX <- crate number
 - From the SBC manager webpage
 - Mark the sbc
 - Click restart
 - Then click ok on the next page
 - Easier than remembering the first command (I think)



Common Problems – SBC Reboot

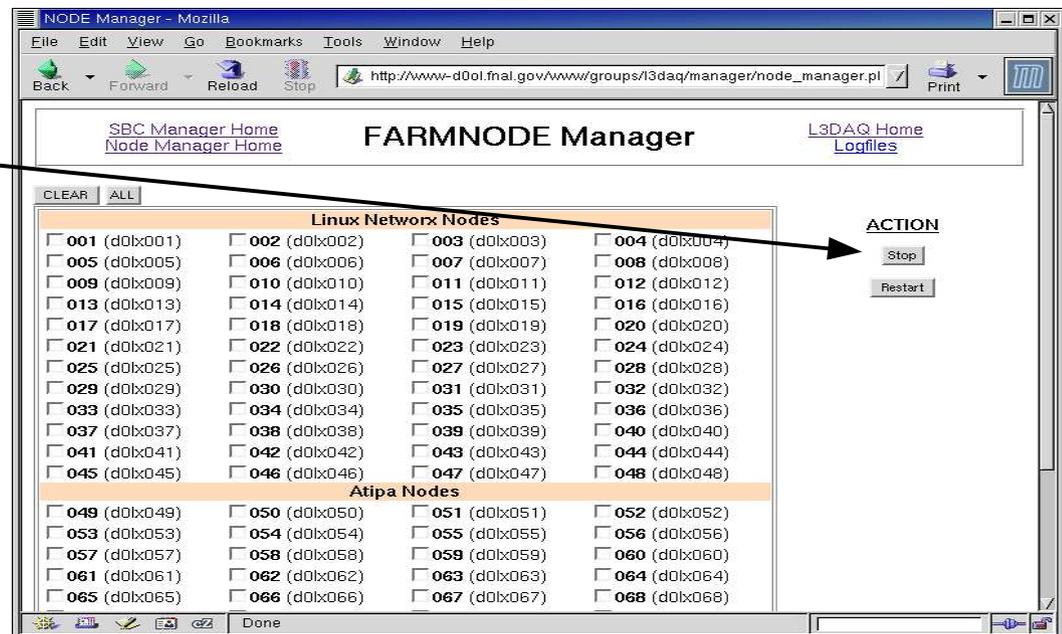
- You try to restart an SBC, but the webpage tells you the SBC is not pingable
 - This means the SBC has crashed and needs a reboot
 - Steps:
 - Find the rack with the SBC (You'll have to leave the CR)
 - Find the reset button and hold it for 5 seconds.
 - It'll take ~1 minute to come up. You can see when it does on üMon.
 - > ping d0sbc0xx



Reset
button

Common Problems – Node Stop

- Level3 expert can't get to a computer and needs you to stop a farm node from running.
 - From the command line:
 - > setup d0online
 - > l3xstop 094 <- node number
 - From the node manager webpage
 - Mark the node number
 - Click reset
 - Then click ok on the next page

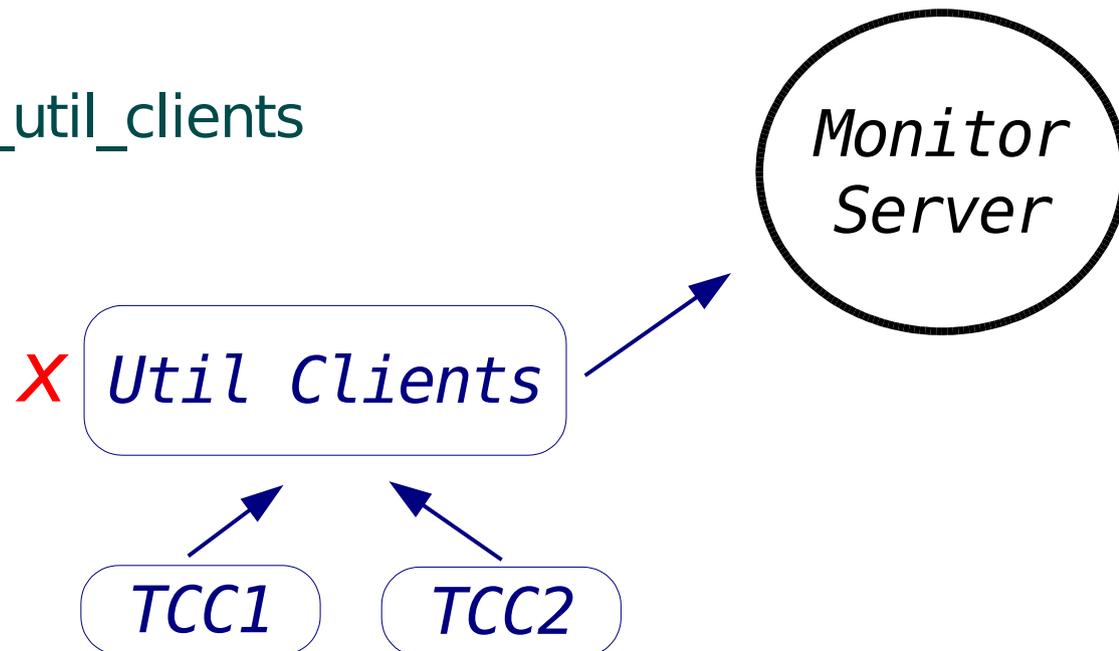


Reseting Farm Nodes

- DON'T DO IT!!!
- If you feel there aren't enough farm nodes in the run, page a level3 expert.
- If a node drops out of the run, that's fine. It happens all the time unfortunately.
 - Stop the node if you see it falling out of the run
 - Send an email to d0daq@fnal.gov

Common Problems – MS_Util_Clients

- Daq dialog goes yellow and the level2 data flow gui is gray
 - This happens all the time.
 - The culprit is the l3ms_util_clients which acts as a go between for the TCC clients and the monitor server.
 - Solution:
 - start_daq l3ms_util_clients



Common Problems – Monitor Server

- üMon and füMon go gray and say something like No Monitoring Information. DaqAI will also not respond and say something about the monitoring
- Solution:
 - Page Aran Garcia-Bellido immediately no matter what time it is.
 - Ok, I'm just kidding.
 - You will just need to restart the monitor server. There is no need stop the run. It's just the monitoring that's down. We are still taking data
 - > start_daq l3ms

Common Problems – FūMon States

- At high luminosities (beginning of the store), you may see a lot of the nodes will be blue
- This basically means the EVB complete buffers are full even though filtershell is running full steam



Common Problems – FūMon States

- You may see a lot of the nodes turn red at some point
- This probably means a readout error so the EventBuilder dropped all the events that were using that SBC.



Common Problems – Download Failed

- Another common problem is the trigger download at the beginning of the run will fail due to a timeout from the L3 Supervisor
 - The supervisor has to connect to ~200 nodes in a given amount of time.
 - Sometimes it takes too long due to network traffic, bad-hair day, whatever....
- Solution: Try again. If it fails again, then it is probably one/several nodes causing trouble and it will say so in the taker. Stop those nodes with an `l3xreset` command.

Common Problems – ScriptRunner Crash

- You may see some alarms in the alarm watcher display saying that some farm node has XX srCrashes.
- This means that ScriptRunner failed to process an event. This happens all the time now at the high luminosities.
- Solution: Make a note of it and send an email to d0daq@fnal.gov. If this happens often, we will take a look. Our newer faster machines will hopefully ease the burden of the L3 trigger at these new high luminosities.

Hi,

d0lx047 has a major alarm for srCrashes (16 according to BB) so I am e-mailing you as per guidance.

Mark Owen

EVE DAQ Shifter

markowen@fnal.gov

Problems.txt!!!

- This is an extremely useful guide for all daq shifters.

Problem: Most or all of the linux farm nodes (from 1 to 48) in the fuMon are blue. The luminosity is high and the L1, L2, and L3 rates are normal.

Solution: Look in the small white box labeled Node Group Details at the bottom of the fuMon. If the CPU(%) below LN is close to 100 there is no problem, the old Linux Nodes are just not fast enough for the high luminosity. The number of blue nodes will slowly decrease with time. If LN CPU(%) is not near 100, there is another problem and you should let an expert know! (There is a similar problem where all the nodes are blue and the solution is to restart the datalogger, do not restart the datalogger if you do not have to.)

Entry: August 08, 2004 Joe Haley

Problem: Coor lost connection with logbook. start_daq auto-elog didn't work.

Solution: l3fserver was taking all the resources of d0o103 then the logbook database couldn't run. l3fserver was killed then everything started to work.

Entry: 6/30/04 Jovan, Mario

Note: kill l3fmonitor in addition to stop_daq l3fserver. (Ammon 10/13/04)

Problem: Daq AI is not speaking up and showing "No data from Monitor Server!?" in the shift (and other) tabs. All other gui's (umon, fumon) appear to be working fine.

Solution: setup d0online
start_daq daqai

Entry: 18 April 2004 - Ike Hall

L3DAQ Webpage

- www-d0online.fnal.gov/www/groups/l3daq
- Linked off the daq shifters homepage

L3DAQ Homepage

- [Shifter Page:](#) All shifters should go here.
- [Public Talks:](#) Public talks and publications
- [Documentation:](#) Talks, Universell, SBC software
- [Monitoring:](#) The latest and greatest monitoring tools.
- [Expert Tools:](#) SBC inventory DB access
-

[Doug Chapin](#)

Last modified: Tue Dec 2 16:08:46 CST 2003

Contact List

- Linked off the DAQ shifter webpage

L3 DAQ On-Call List

[<- prev](#)

11 OCTOBER through 5 DECEMBER

	Primary	Secondary
11 OCTOBER - 24 OCTOBER	Aran	Gordon
25 OCTOBER - 7 NOVEMBER	Thomas	Aran
8 NOVEMBER - 21 NOVEMBER	Thomas	Doug
22 NOVEMBER - 5 DECEMBER	Aran	Thomas

- You'll find phone numbers at the bottom
 - If during the day, try paging or call office number
 - If after hours just call the cell phone number listed
- People: Thomas Gadfort (UW), Aran Garcia-Bellido (UW), Gordon Watts (UW), Yunhe Xie (Brown)