

DØ Level 3 Status

T. Bose, G. Brooijmans, D. Cutts,
T. Gadfort, A. Garcia-Bellido
A. Haas, *G. Watts*, Y. Xu
With Help From Many Others



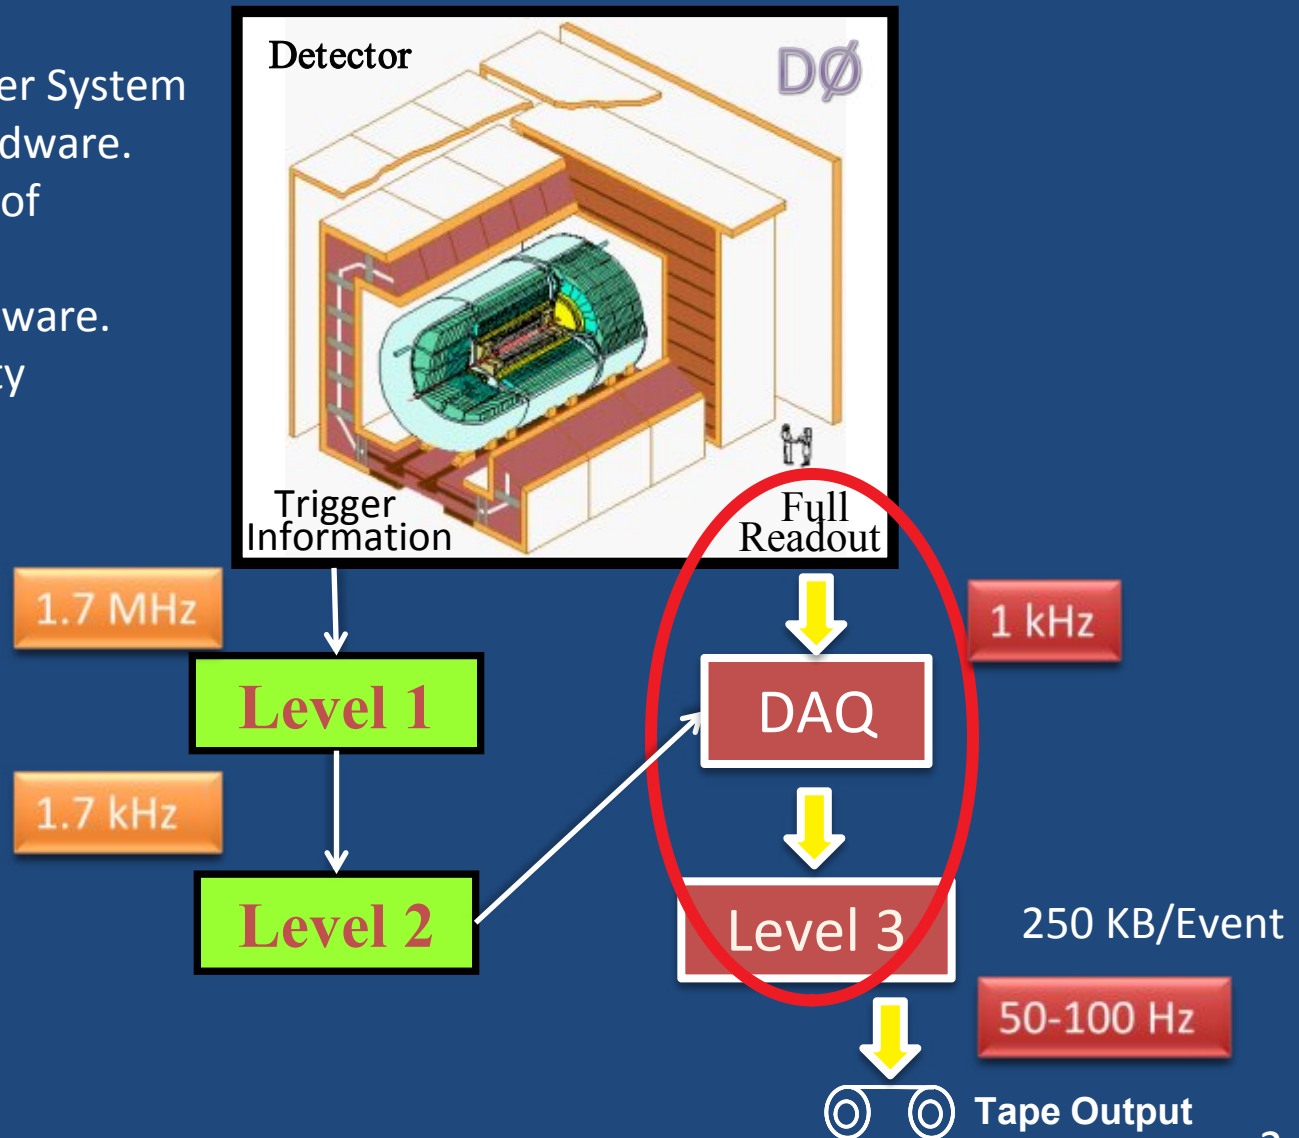
The DØ Trigger/DAQ System

Standard Tiered HEP Trigger System

- Level 1 is specialized hardware.
- Level 2 is a combination of specialized hardware and commodity hardware/software.
- Level 3 is fully commodity hardware and software.

The DAQ System

- Full detector readout occurs only after a L2 accept.
- A single Node in Level 3 farm makes the L3 trigger decision.



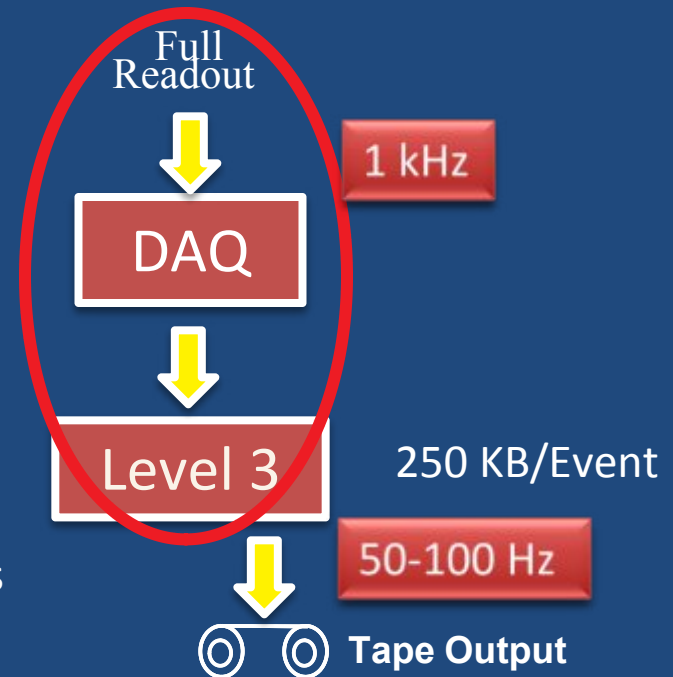
Overview

System has been fully running since March 2002

- 1 KHz Input
- 250 KB/event
- 100 Hz output

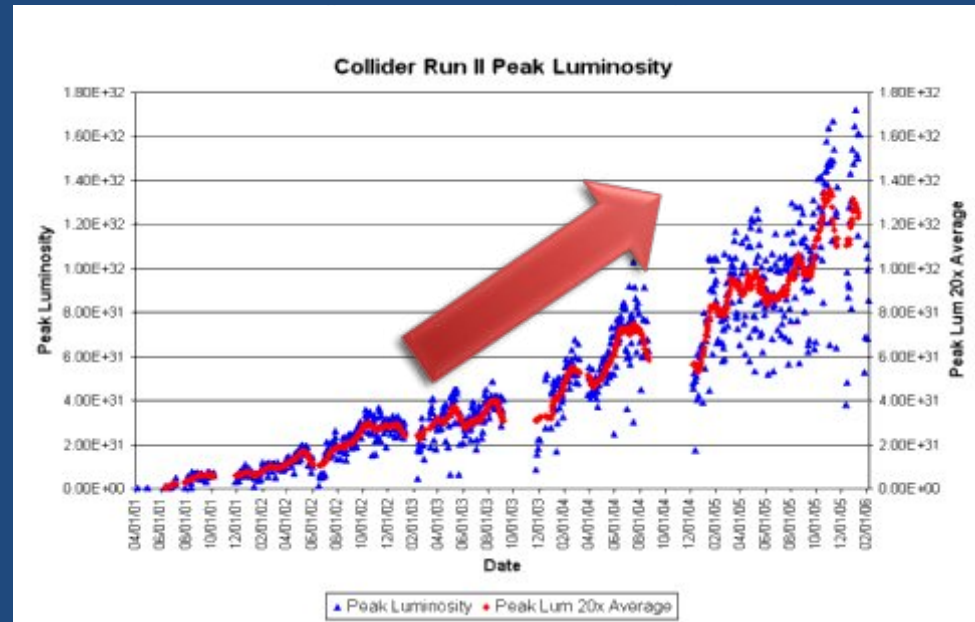
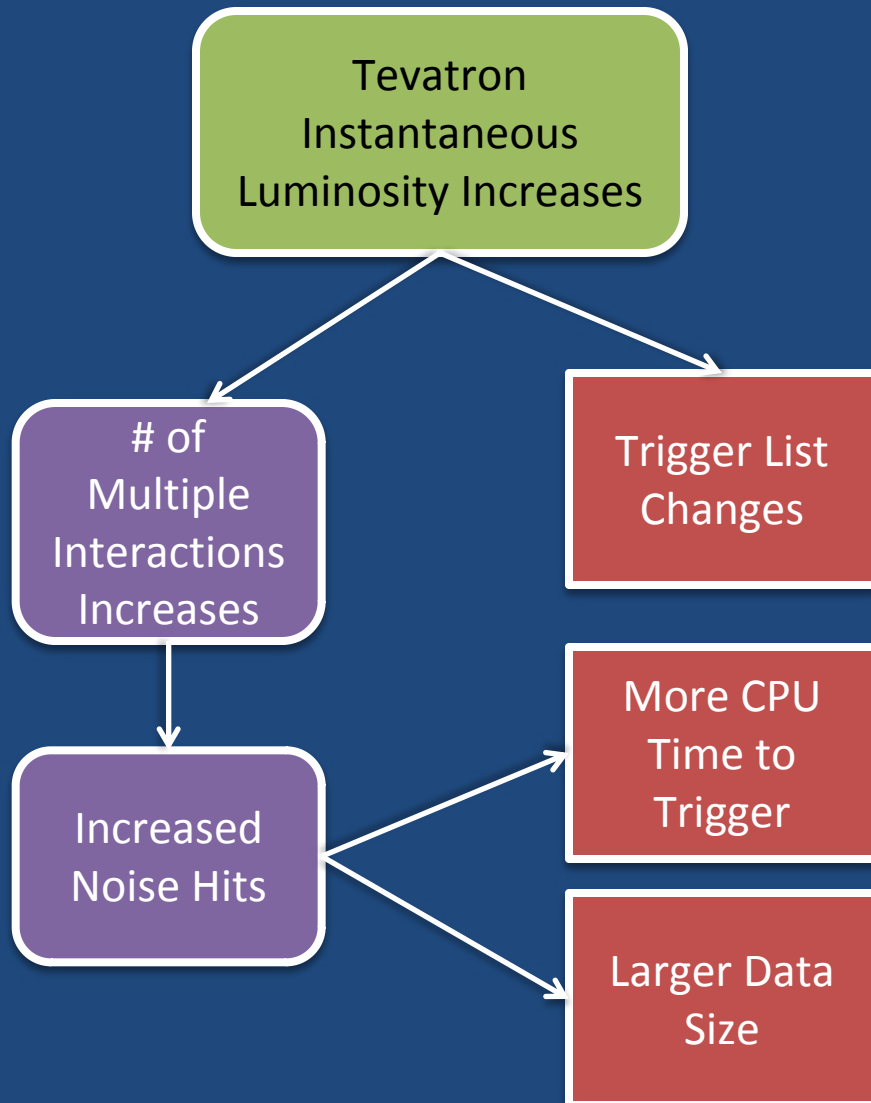
System has been fully operational since March 2002

- First full detector readout
 - L1 and L2 use fast outs for some detectors
- Standard Scatter - Gather Architecture
 - ~70 readout crates
 - Data sent to one of 250 nodes
- Trigger Software written by large group of physicists
- Continuous Upgrades since operation started
 - Added ~10 new crates
 - Added almost 180 new nodes
- No Major Unplanned Outages



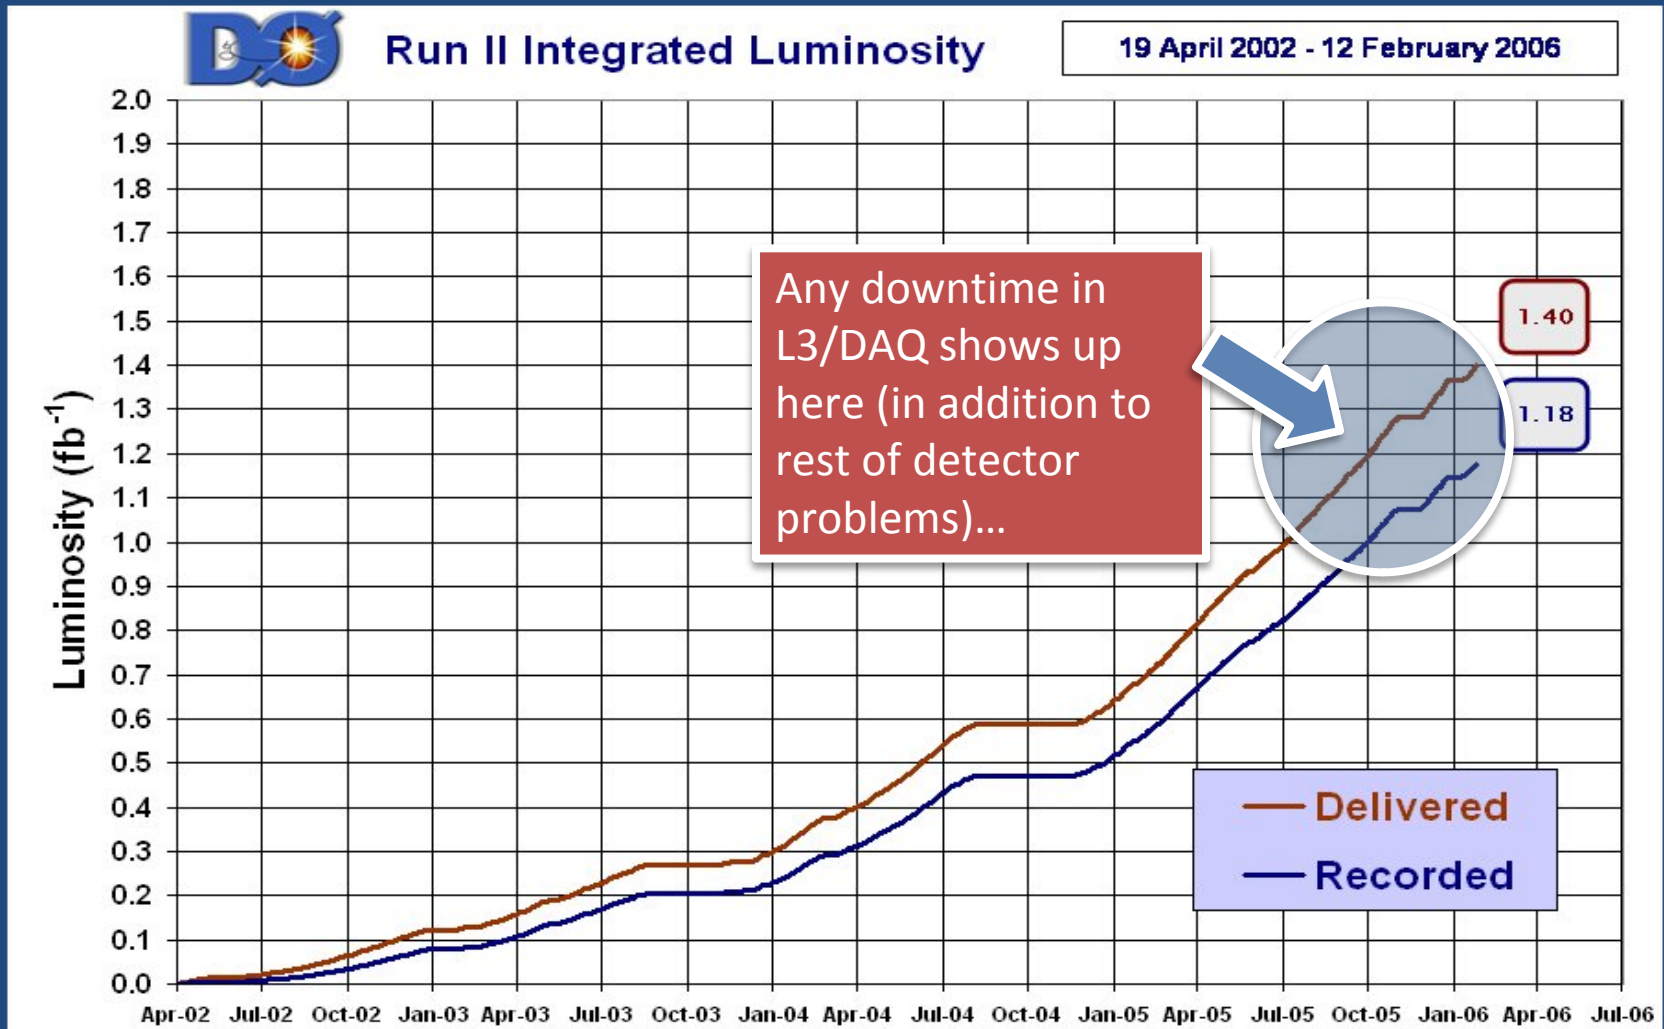
An Overwhelming Success

The Tevatron Is Doing Well...



All of these put pressure on the detectors and the DAQ and the Trigger decision system!

And The Constant Pressure...



Basic Operation

Two Types of Messages In the System

Data Flow

- Directed, unidirectional flow
- Minimize copies of data
- Buffered at origin and at destination

Control Flow

- 100% raw TCP/IP
- Group small data messages into bundles to increase network efficiency.
- Compress messages via preconfigured lookup tables

Endpoint Layout

Readout Crates receive data from detector via detector specific communications

- All Located in **Counting House**
- CH and detector are electrically isolated!
- Digitized on detector
- Analog to counting house
- ROC: where data format becomes uniform

Farm Node

Farm Node

Farmnodes send data to online system and tape.

- Located about 100 m away
- Not electrically isolated

Farm Node

ROC

ROC

ROC

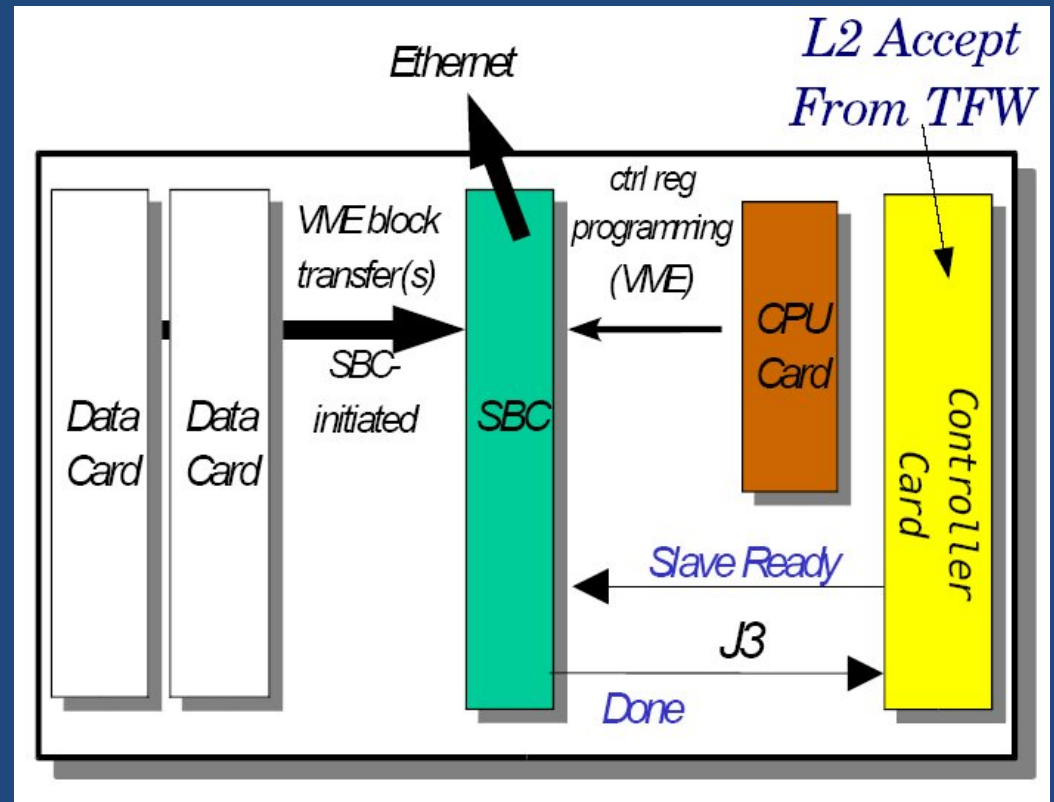
ROC

ROC

Readout Crate

The interface between detector groups and the DAQ/Trigger Group

- VME backplane (32 and 64 bit)
- 2 Special lines on the J3 backplane
 - For control only
- Crate Specific Controller card runs everything
- DAQ Single Board Computer
 - Reads out data cards
 - After Controller request
 - Data is always read out
 - VMIC PIII 933 MHz
 - 128 – 256 MB Ram




Custom Driver and Universe II chip
do all VME interaction

ROC-SBC

PIII 933MHz
128MB RAM
128MB flash

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

Reset
Button



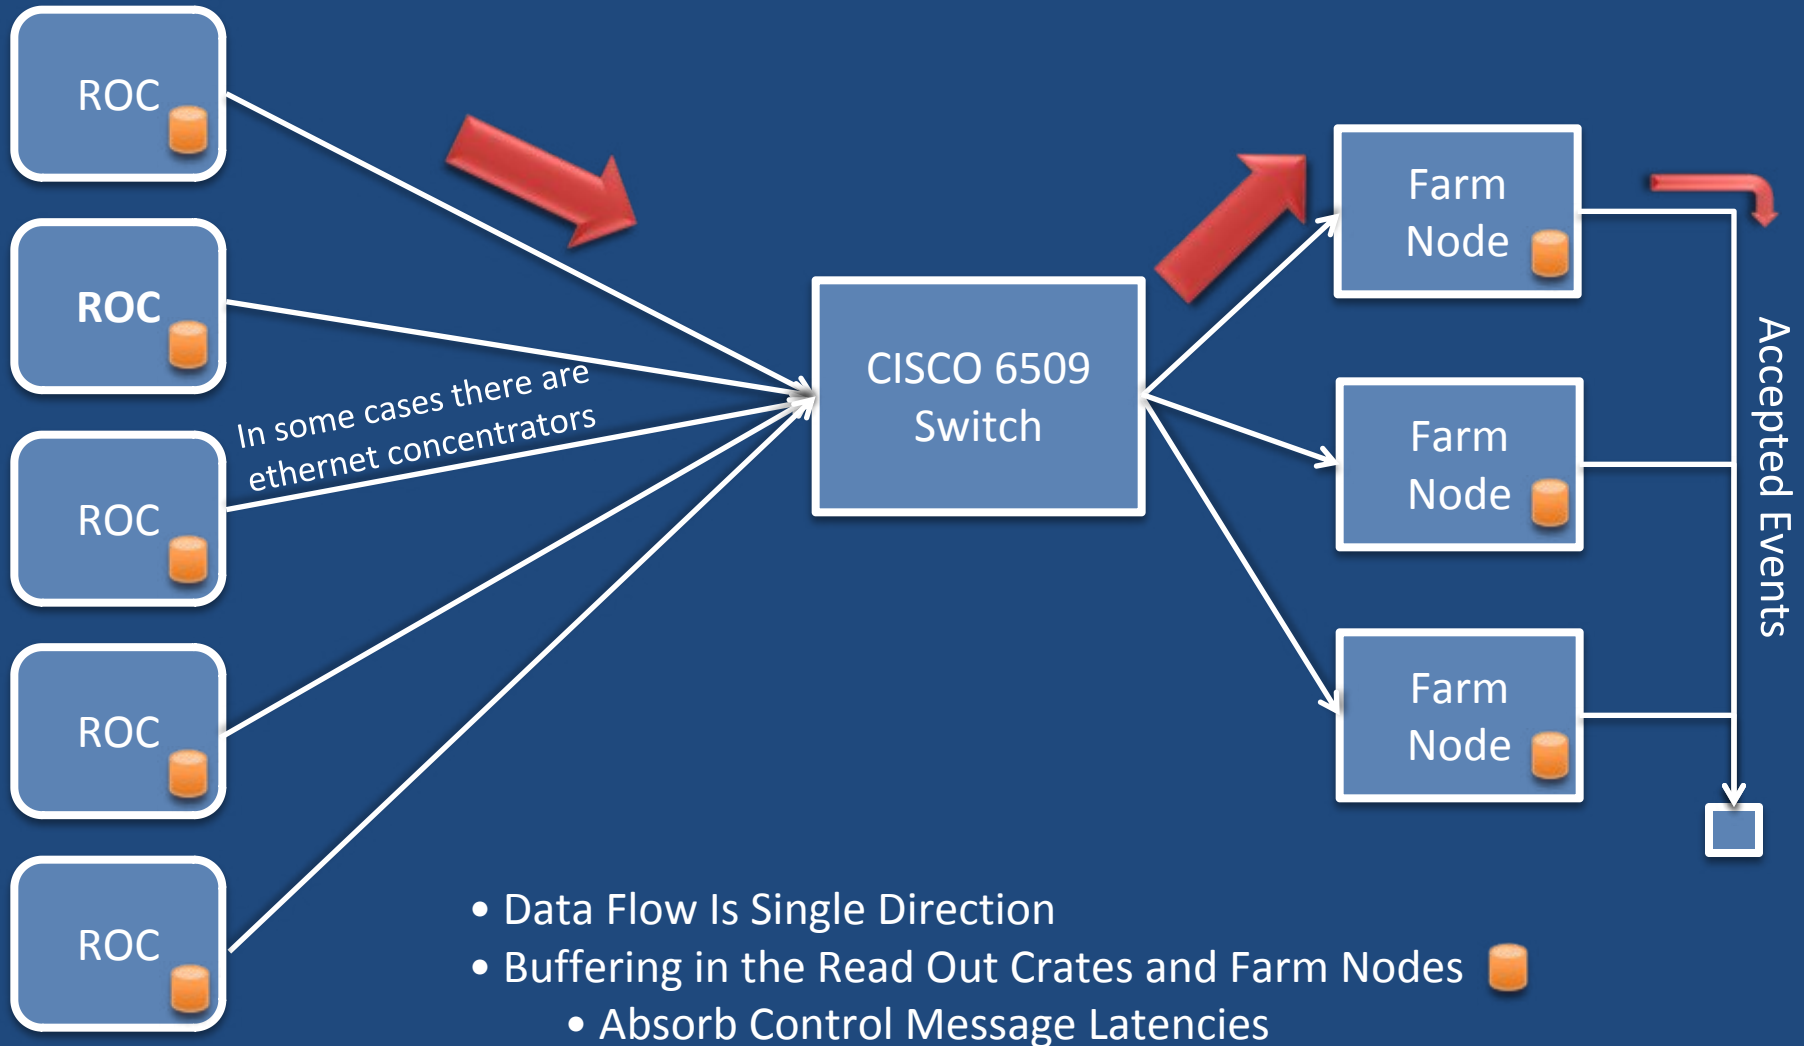
128MB Flash ROM
•Stores
configuration file

Tundra Universe II
•PCI-VME Interface
•DMA controller

PMC Digital I/O Card
•Coordinates VME
Readout

J3 Handshake
Slave Ready &
Done Lines

Data Flow



Control Flow

ROC

ROC

ROC

ROC

ROC

Supervisor

- Begin/End Run
- Per Event Messages

Attempt to keep a consistent model of ownership.

Routing Master

DØ Run Control

Farm Node

Farm Node

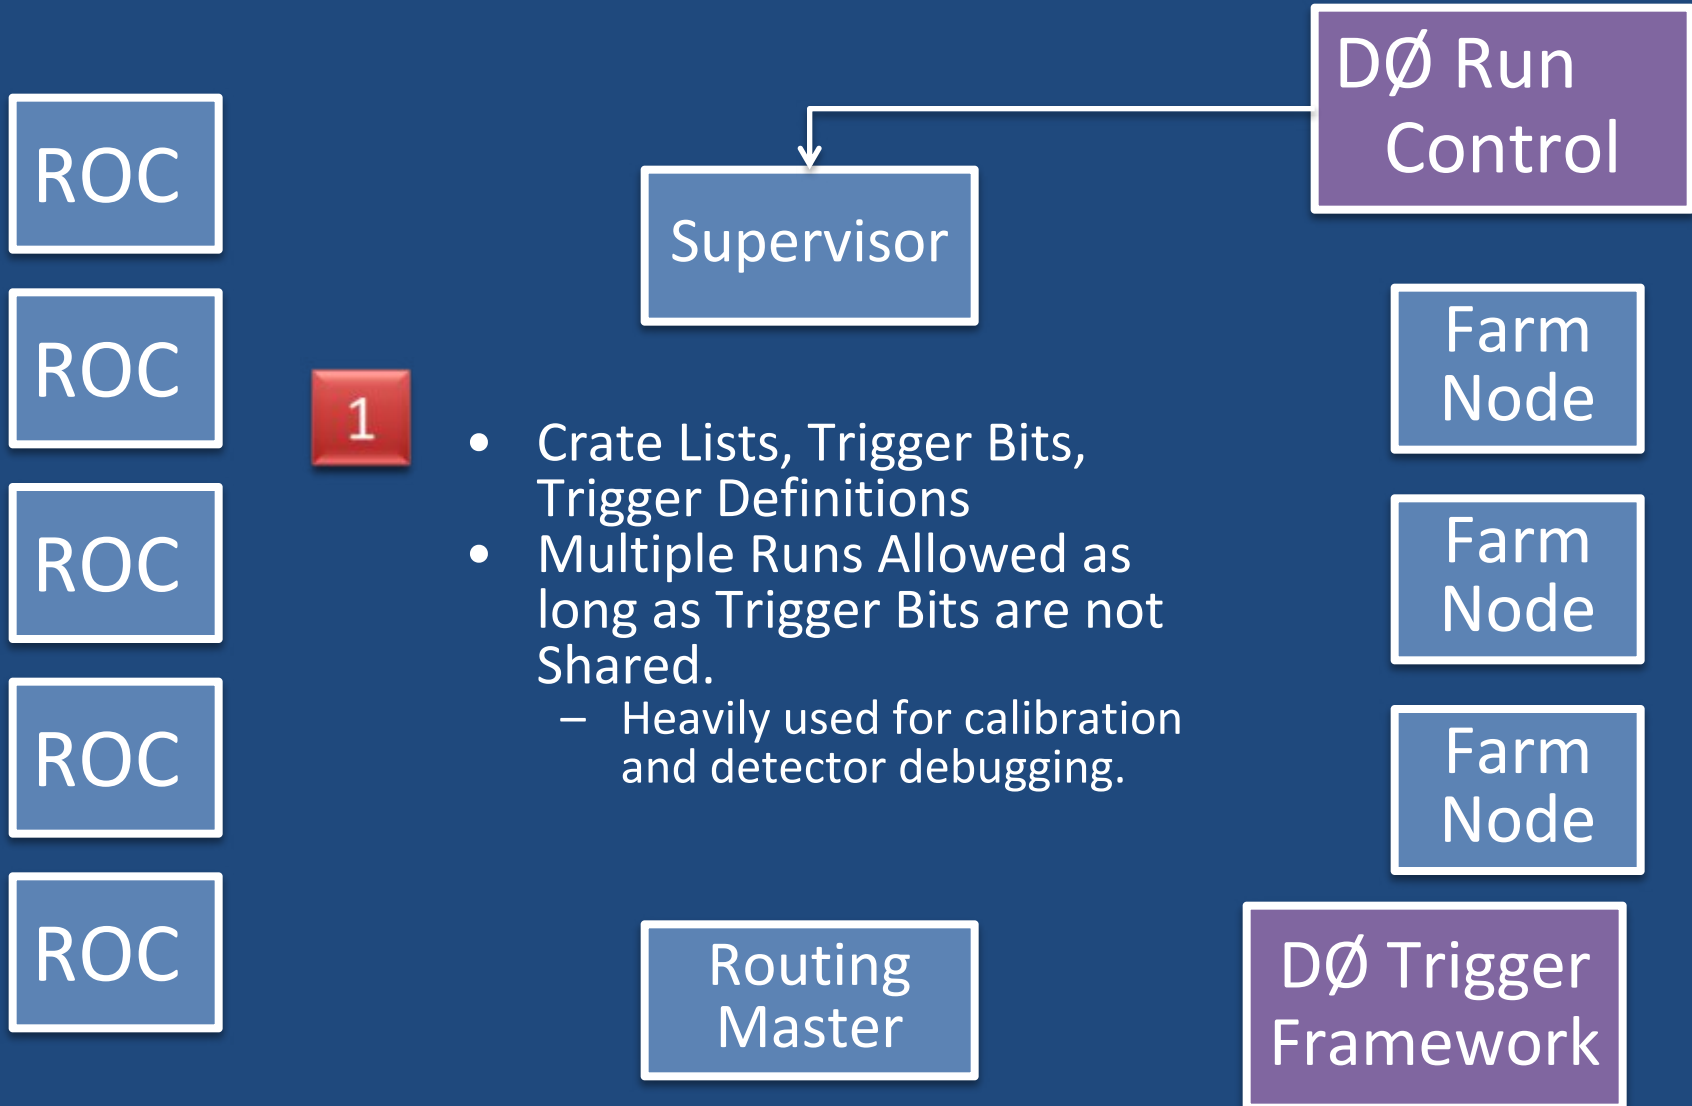
Farm Node

DØ Trigger Framework

Run Configuration

- Make sure all components are ready
 - All components of the system are touched
 - There is only success
 - Even a single node will cause the process to fail
 - Even though the run will proceed fine with a missing node.
 - Node failure during run doesn't interrupt run at all.
 - Everything occurs synchronously.
 - Typically done before beam is in the machine!
- Distribute Lookup Tables
 - Keep per-event control messages as small as possible by using indices in lookup tables

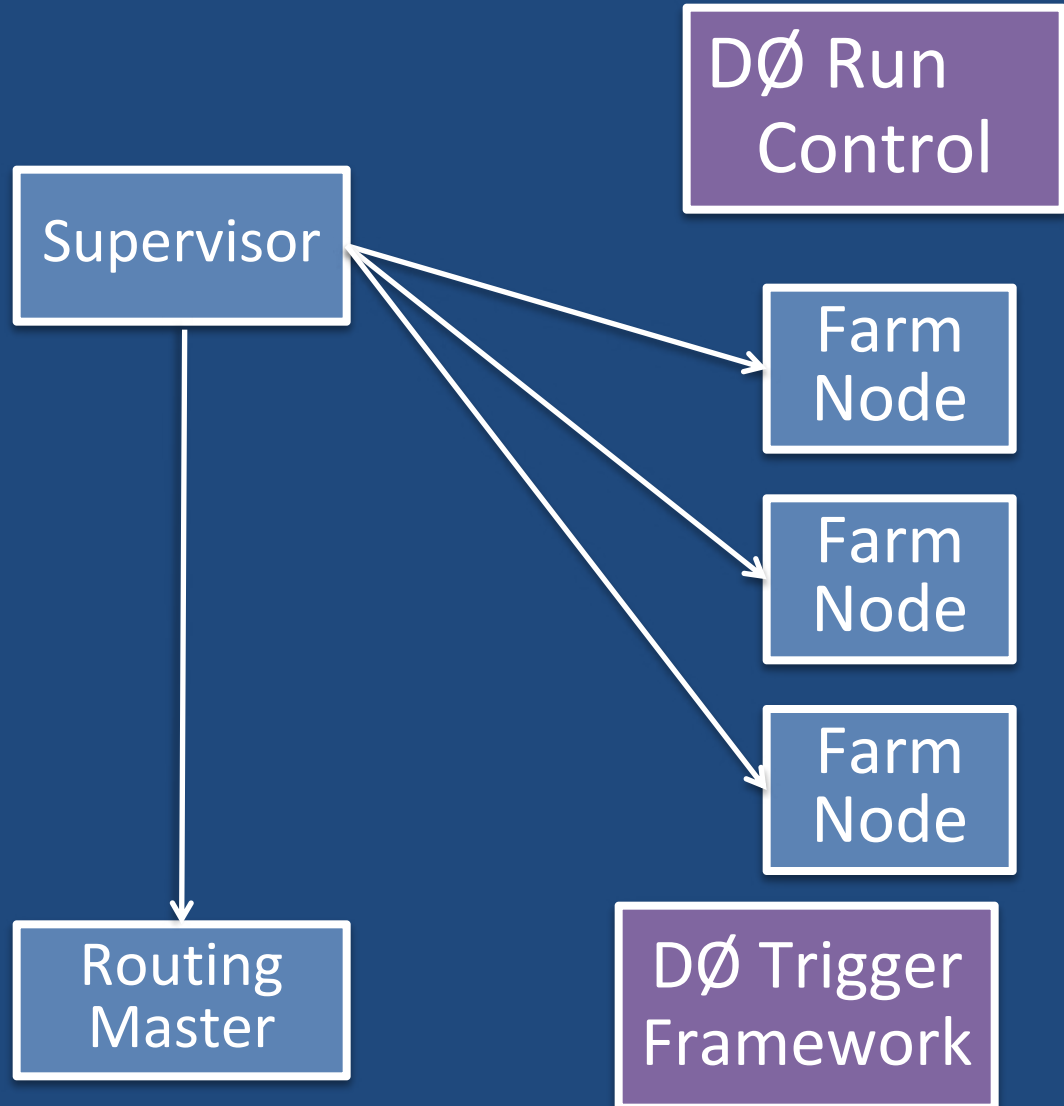
Run Configuration



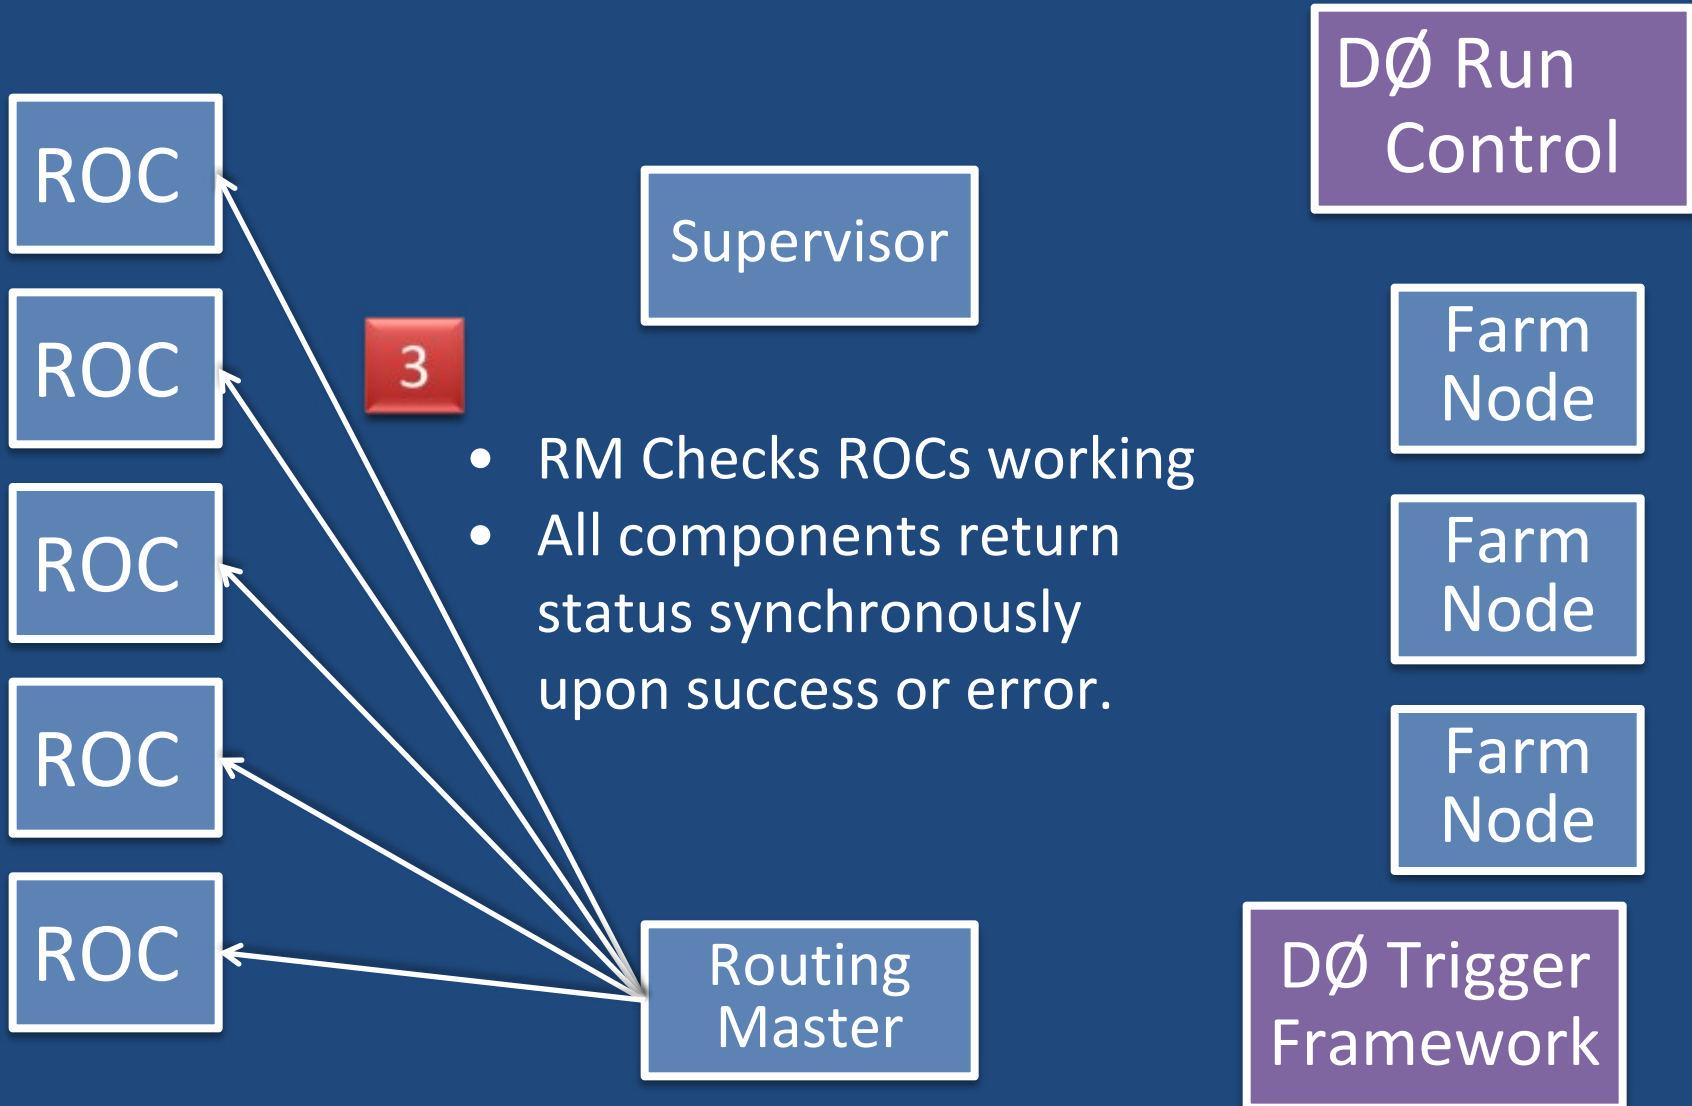
Run Configuration

2

- Farm Nodes get trigger bits and expected crate lists
- RM gets trigger bits, crate lists, and the sub-set of nodes participating in the run



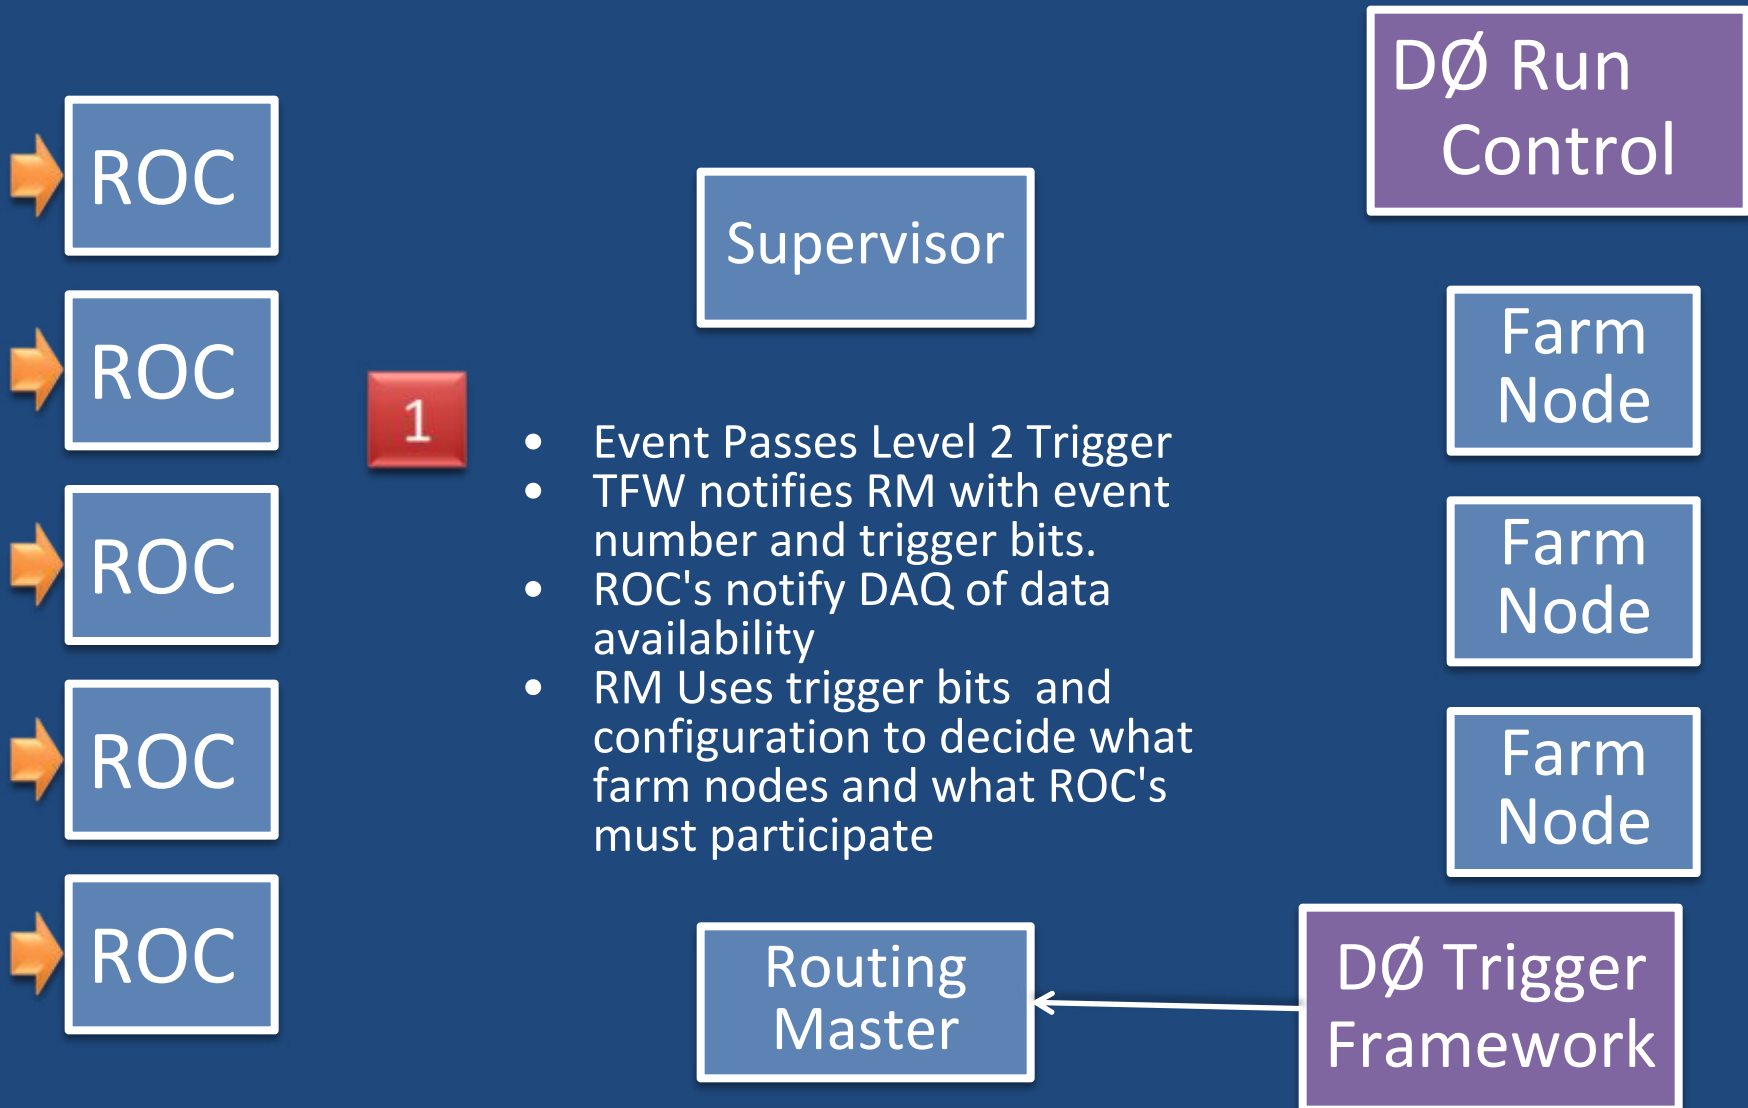
Run Configuration



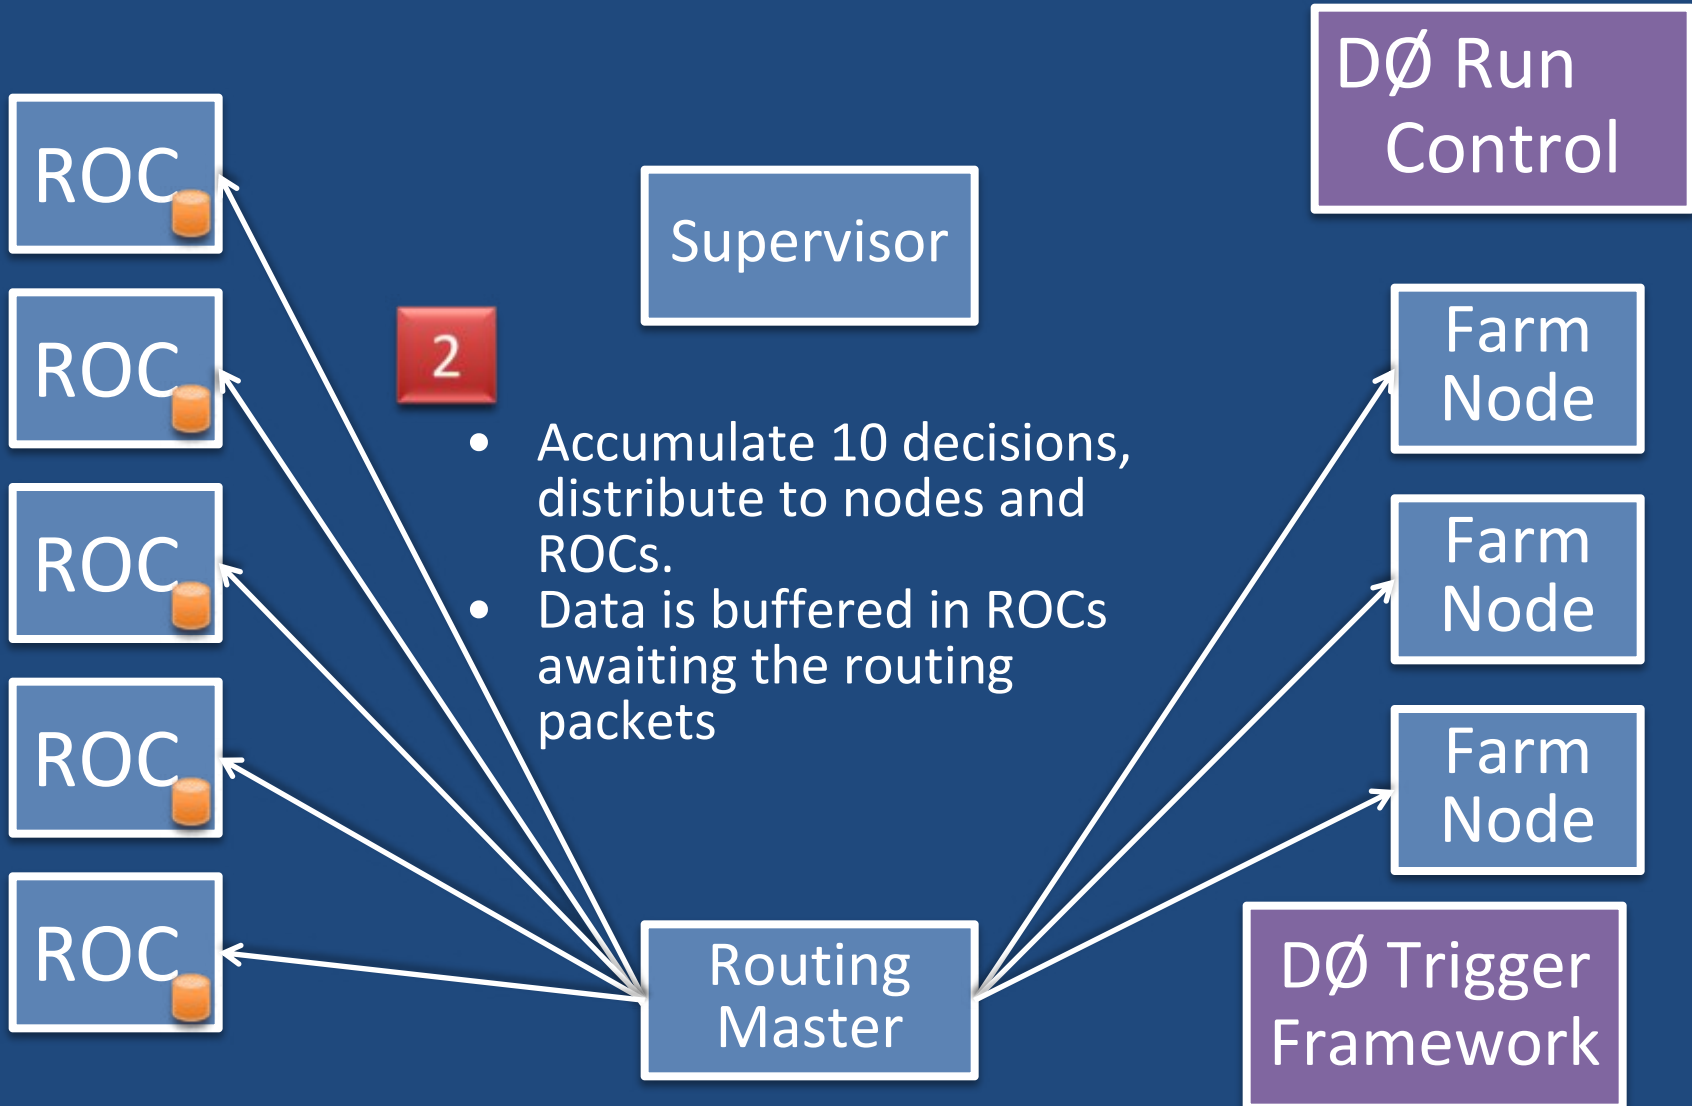
Event Flow

- The Routing Master acts as the traffic director.
 - Assigns nodes to process particular triggers
 - Performs node load balancing
 - Will stop the DAQ if too few nodes are available.
 - All intelligence in the system is located here.
- ROC and the Farm Nodes buffer data
 - To give the RM time to make its decisions
 - And to keep the network flow smooth
 - Don't back up the switch buffers!

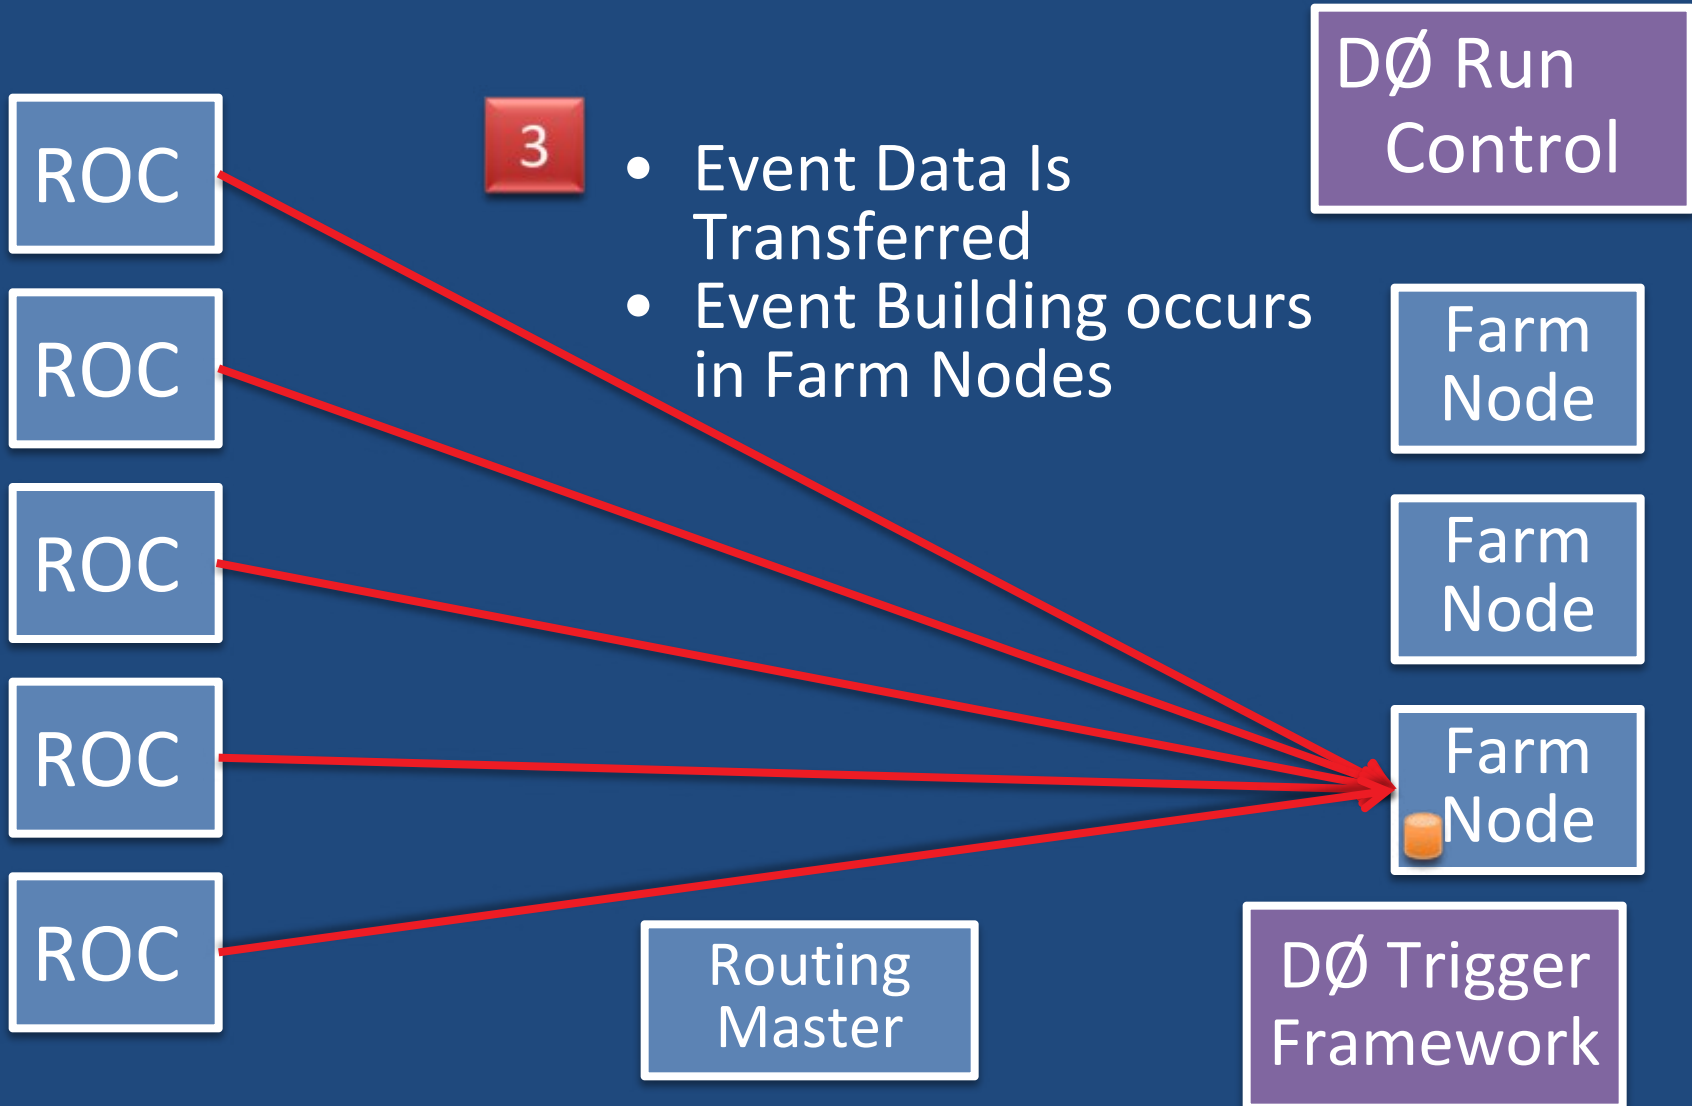
Event Flow



Event Flow



Event Flow



Event Flow

ROC

ROC

ROC

ROC

ROC

4

- Node notifies RM when all bits of event have arrived or the event has timed out
- Done only if node has free space for a new event!

DØ Run Control

Farm Node

Farm Node

Farm Node

Routing Master

DØ Trigger Framework

Event Flow

ROC

ROC

ROC

ROC

ROC

5

- Trigger software processes the event
- Sent to online if it passes trigger criteria

Routing Master

DØ Run Control

Farm Node

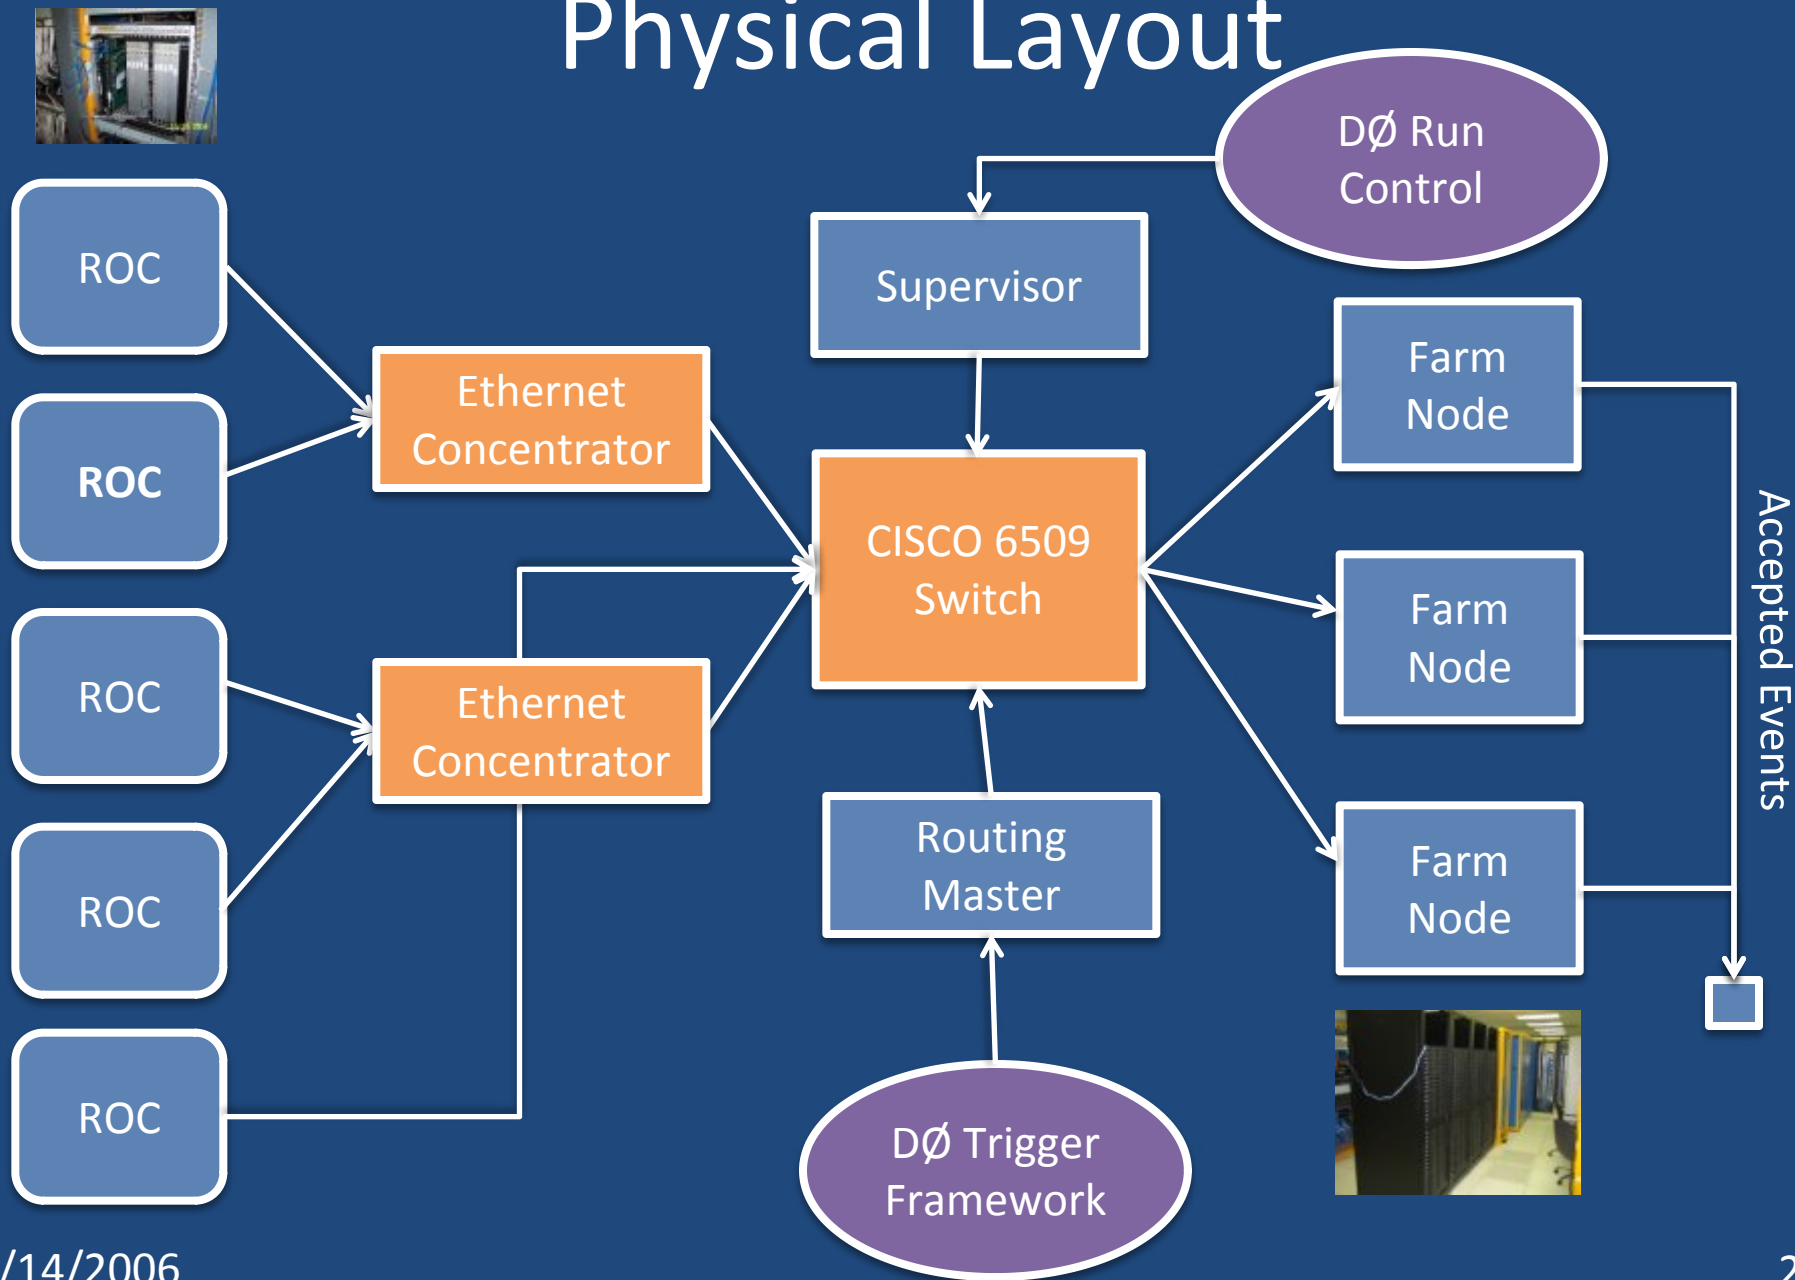
Farm Node

Farm Node



DØ Trigger Framework

Physical Layout



Performance

Exceeds DØ's current requirements

- > 1000 Hz Operation
 - Hard to test without ROC's feeding us at that rate!
- Many ROC's exceeded designed data sizes.
 - System was designed for < 5 KB
 - Some Tracking crates boarder on 24 KB.
- Latencies and Timeouts
 - We don't loose events because of the DAQ
 - Except for exceptional conditions like node crashes
 - Recover without hiccup when crate drops event

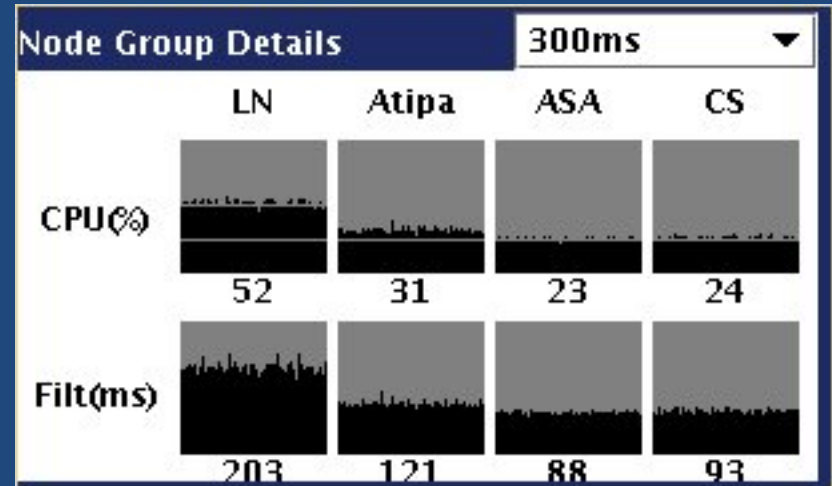
Heterogeneous Farm

Load Balancing of the Trigger Farm is Crucial

Especially with 4 different types of nodes purchased over 4 years.

Nodes Advertise Number of *Free Buffers*

Free Buffer is the ability to accept all data from a single event.



Number Free Buffers = Smallest (3, Space Left In Memory)


Limits Number of In Flight Events

Fills up more quickly on slow nodes

Event Latency

ROC's send events as soon as they have both data and routing information

Some detectors don't digitize until after L2 accept arrives!



When RM makes routing decision

Multiple events can be in flight to same node at once; throughput unaffected
Up to 3

Data Bottlenecks

VME Backplane

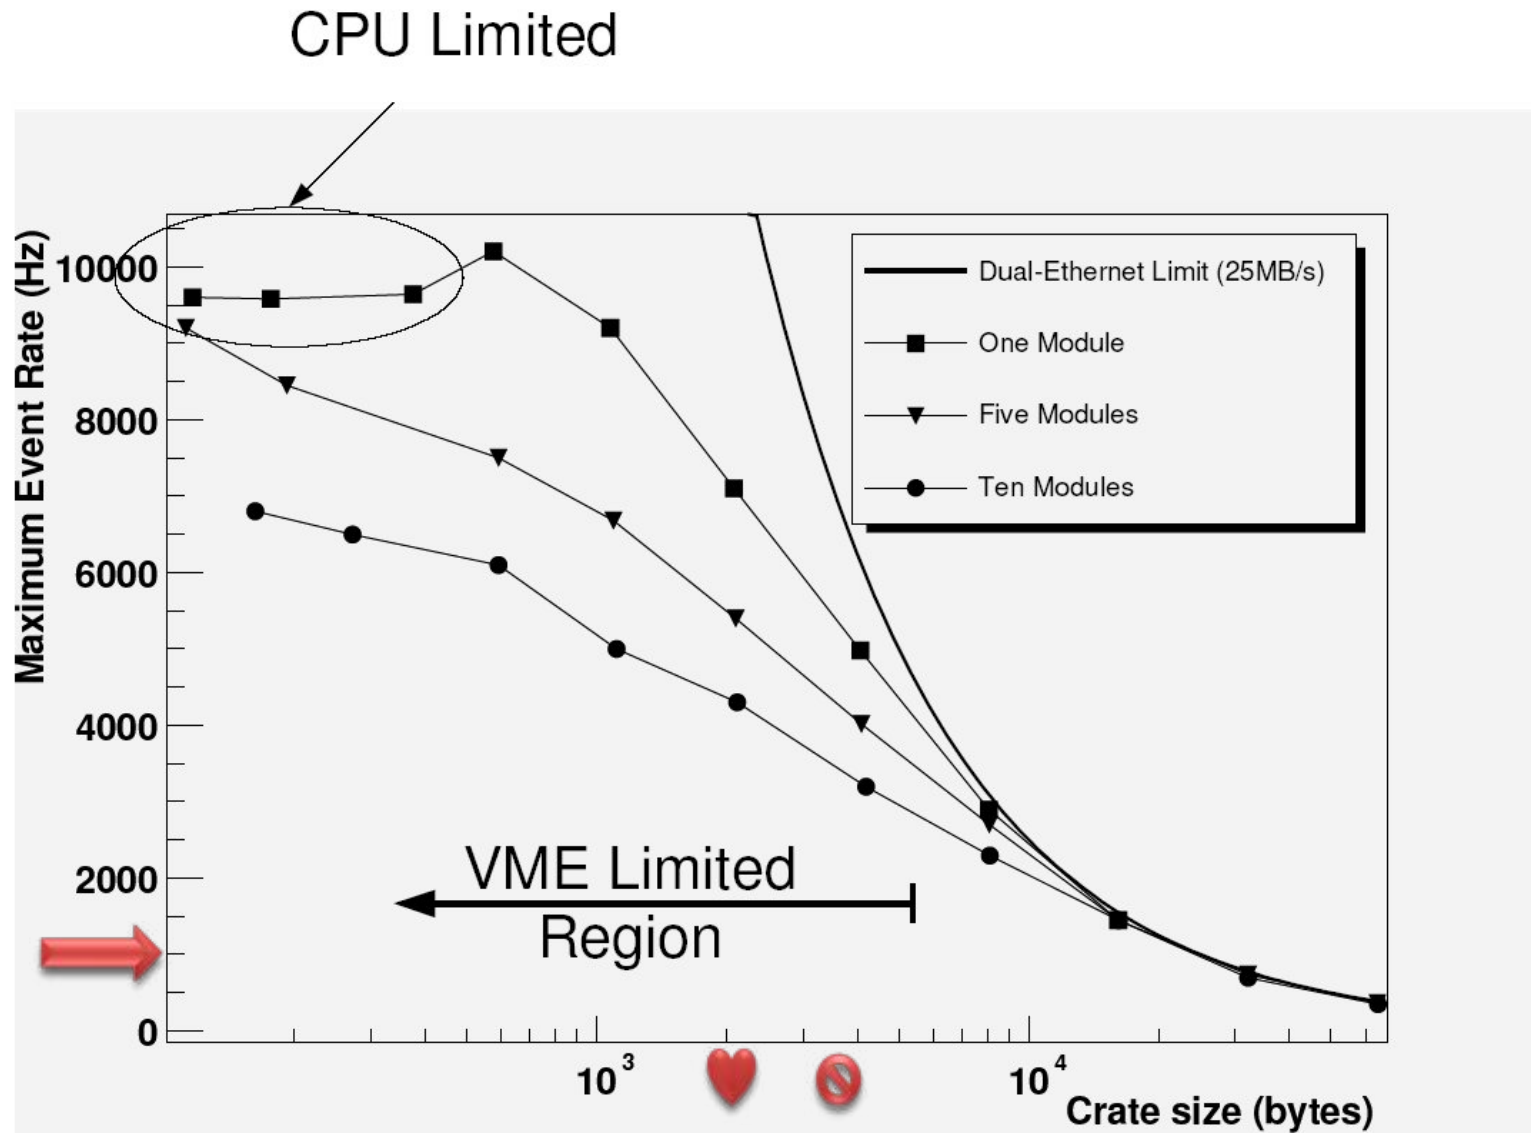
- VME Transaction Overhead
 - Per Individual Data Card
- Basic Backplane Transfer Rate Limitations

All Of Our Bottlenecks
are in the ROC's

Network Interface

- Dual 100 MB interfaces
- 23 MB/sec is approximate limit
- Installed gigabit on some tracking crates

Rate Tests



Problems

Farm Maintenance

- We have 242 Nodes (x10 less than ATLAS!)
 - Range in age from 4 years to 6 months
 - Number of failures and variety of failure modes caught us by surprise
- Nodes *die* about once every other day
 - Expert is required to recover
 - Is it just a spurious timeout?
 - Is it a real HW failure? Just a disk corruption?
 - People intensive
 - Dropped packets, extra time to setup, etc. are at the root of most of these failures
 - Expert can quickly reset these.
- Serious Failure Rate is about one per month
 - Requires retiring node or returning for maintenance
 - 3 year contracts.
 - Notice definite quality differences between purchases!
 - CPU fans one case, disks another!
- Automatic Node Recovery
 - With x10 more nodes would need automated tools to recover the nodes.

There are patterns!

Quote

[You are being] too gloomy. In 2 shifts over the weekend the L3DAQ performed flawlessly. A node lives typically for hundreds of days!

Author is correct: failure rates and timeouts are strongly dependent on luminosity and rate!

Node Configuration

- Each Farm Node contains trigger executable
 - Matched to trigger list
- Deployment is triggered by hand
 - But what if a node is turned off at the time?
- Doesn't happen often
 - But cost when it does is high
- Proper Solution
 - Trigger List contains Trigger Executable stamp
 - Check failure causes crash
 - Taking data with wrong trigger version is expensive
 - Automatic Configuration of Node
 - If not proper when configuration request arrives, configure it.

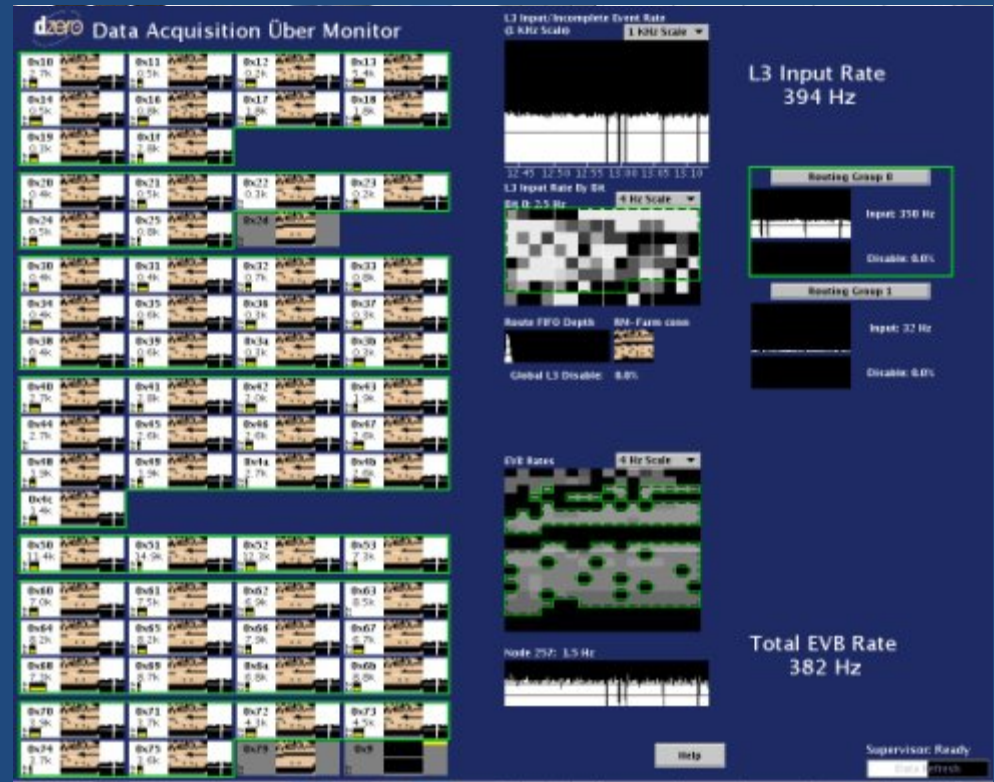
Fine as long as no one touches the system!

People

- DØ Level 3 is run by Physicists
 - Professors, Post-docs, and graduate students
- The Tevatron Has a lot of Data
 - Primary interest is analyzing the data.
- Off Loading
 - Fermilab is expert at running large farms; can we take advantage of that
 - To remove one of the most time-consuming efforts
- Project Prioritization
 - We tend to ignore small problems that are difficult to solve
 - Use people to solve them in the short term
- Constant fear that changes will cause instabilities!


Sociology

- Extensive monitor system and large number of displays in Control Room
 - See problems with all of DAQ first on a DAQ monitor
- We get called first!
 - Muon crates go bad, etc.
 - Always know when new set of shifters.



Automatic diagnostics for common problems crucial

Computer Security



- No holes through firewall unless explicitly requested and approved
 - Firewall mandated by US Department of Energy
- Two reconfigurations of firewall
 - What works one week, may not the next
- We have explicit holes for:
 - XML monitor data to external relay machine
 - NFS readonly mount of some logfile disks to an external web server.



Large And Small Runs

DAQ system supports multiple runs

- Calibration, Commissioning/testing
- Full Physics Running
- Simultaneously in some cases
 - With minimal dead time!

DAQ System Designed for 10 nodes

- Some Special Calibration Runs require 2 nodes!



There are lots of large farm assumptions
built into the DAQ routing software!

Not possible to foresee every configuration request...

➔ Pressure can be intense when quality of data is on the line

Log File Access

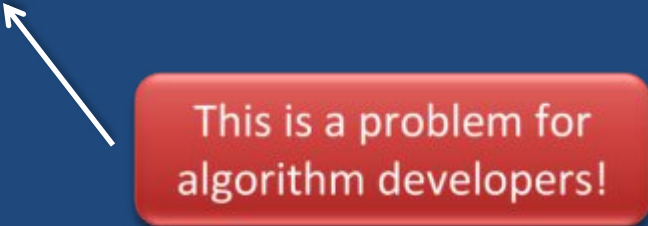
Everything produces log files

Big Effort in the DAQ group to eliminate all log messages

- About 1 kB of messages at start up.
- Nothing during normal running
- One line for error messages
- Backdoor to turn on detailed logging without restarting!

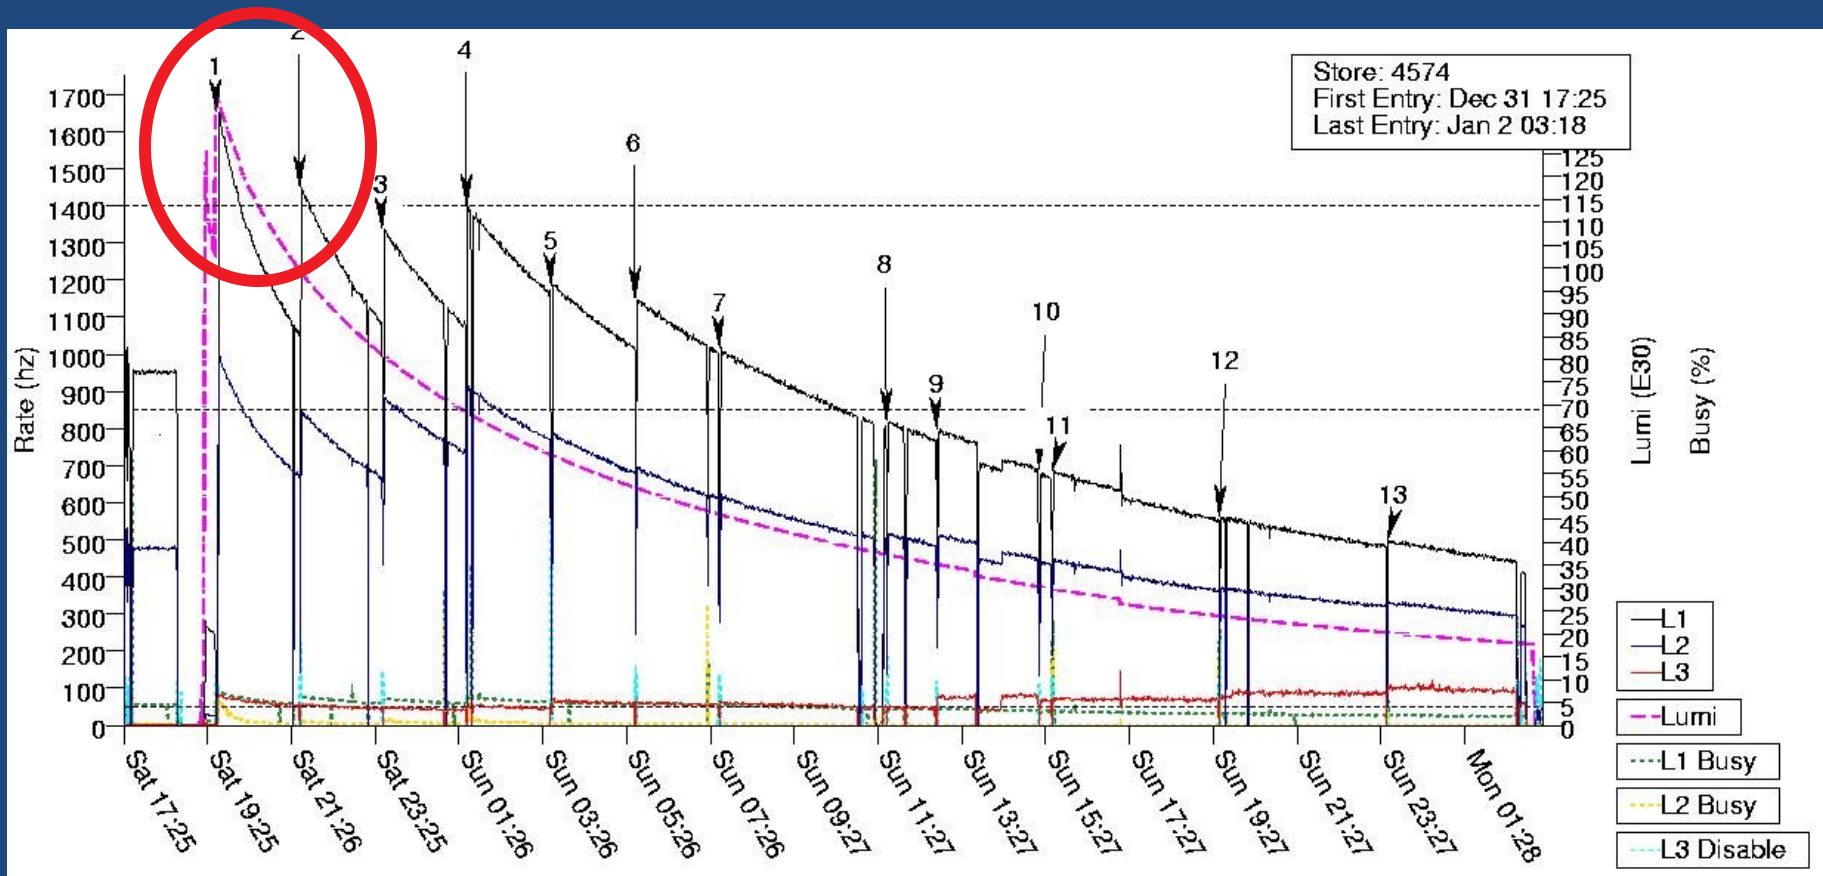
Trigger Executable Not So Fortunate

- 100's of messages per minute
- Written to local disk (NFS too failure prone).
- No good way to examine log files



This is a problem for
algorithm developers!

Trigger Leveling



- Luminosity At Start of Store Decays Very Quickly (1 hour)
- Trigger list and prescales Must be designed for Peak Luminosity
- Potential loss in physics due to quick drop
 - Would like to design for average luminosity

DAQ Load Leveling

→ Real-time Prescale Adjusting

Adiabatically adjust prescales as luminosity falls to balance trigger mix and rate.



→ Trigger Leveling

Cache Extra data at start of run, and catch up as Trigger farm utilization decreases

All of our machines have disk:

10 gig on each node is 2.25 hours of full rate events

Luminosity tracking is the Hardest Problem To Solve!

Conclusions

- The DØ DAQ System has been running well since installation.
- Design has accommodated upgrades we've needed to do.
 - CDF copied us!
- Experiment Taking Data
 - Requires numerous small changes to DAQ to accommodate special types of runs.
- A number of unexpected issues
 - Detailed Monitoring Infrastructure for debugging
 - **People!**