A Generic Multi-node State Monitoring System

James A. Hamilton

Gregory P. Dubois-Felsmann

Rainer Bartoldus

For the BaBar Computing Group

BaBar Online Structure

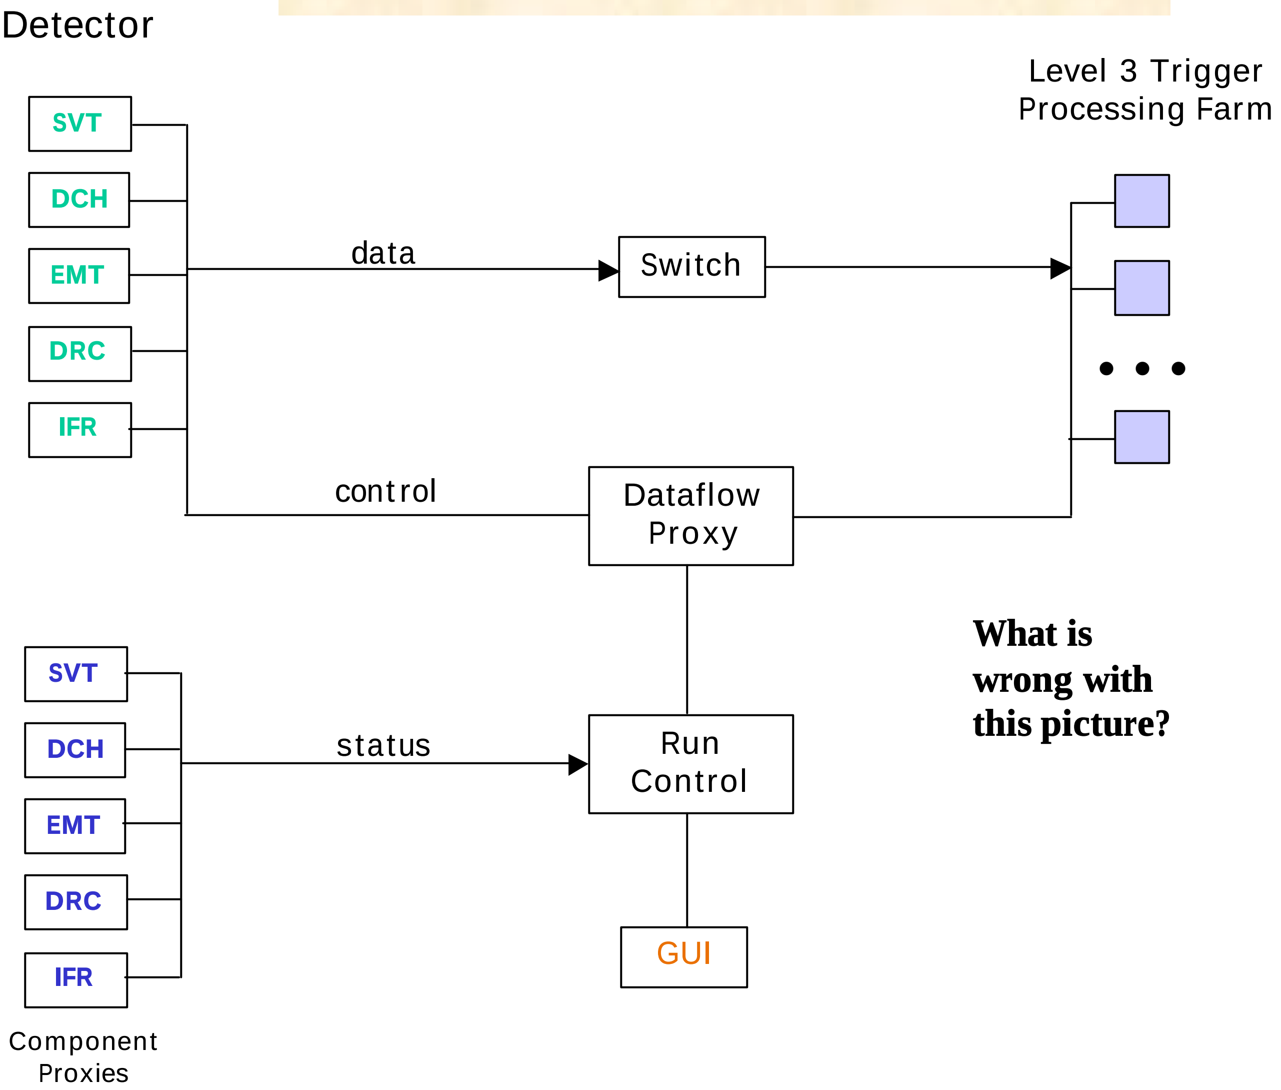

What is wrong with this picture?

BaBar Online Structure

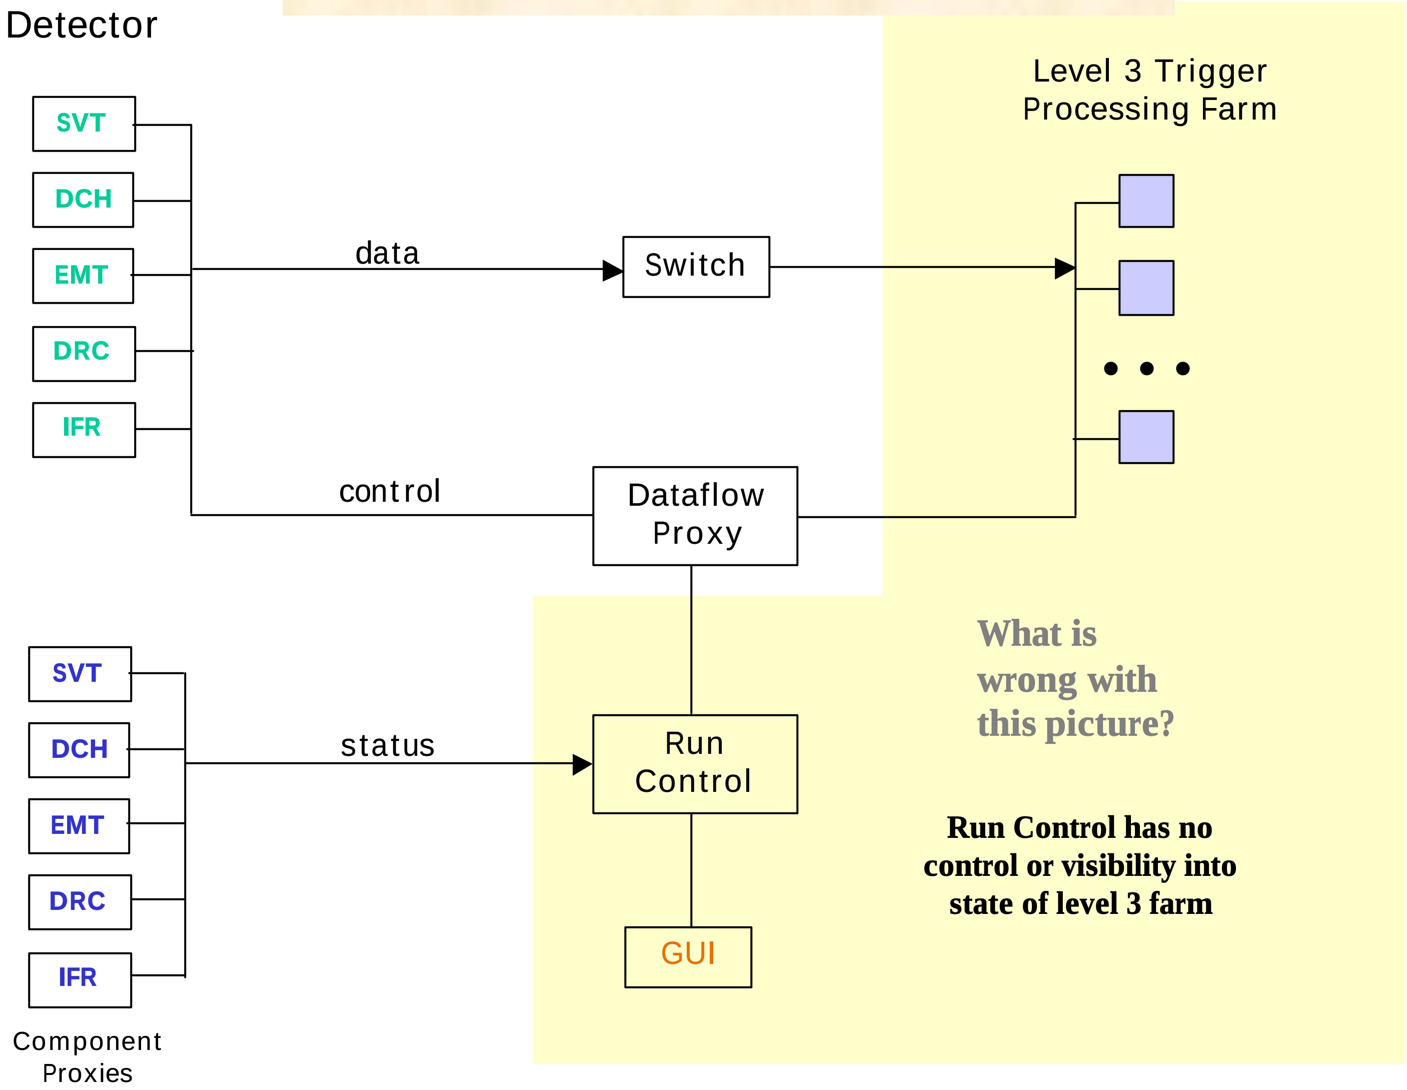

What is wrong with this picture?

Run Control has no control or visibility into state of level 3 farm

BaBar Online Structure

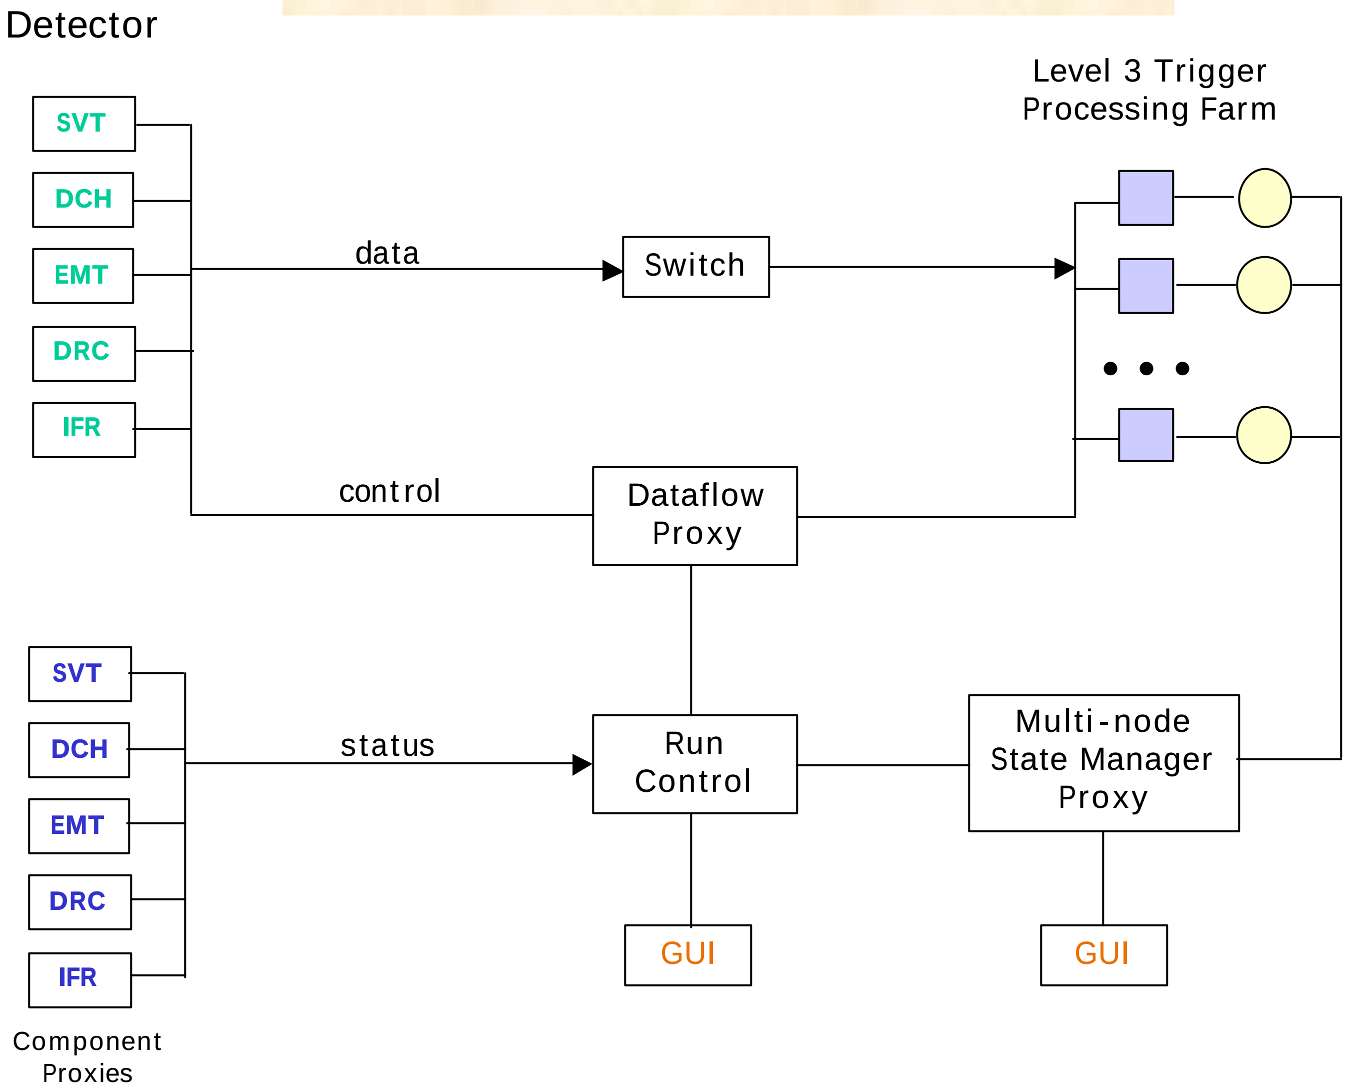

Alternative Implementations

- Fold into dataflow proxy
  - *Destabilize existing complex subsystem*
  - *No GUI*
- Implement directly in Run Control
  - *Uses SMI (State Management Interface – Boda Franek)*
  - *Difficult to support desired variability*
  - *Need for timeouts in state monitoring*

The Generic Solution

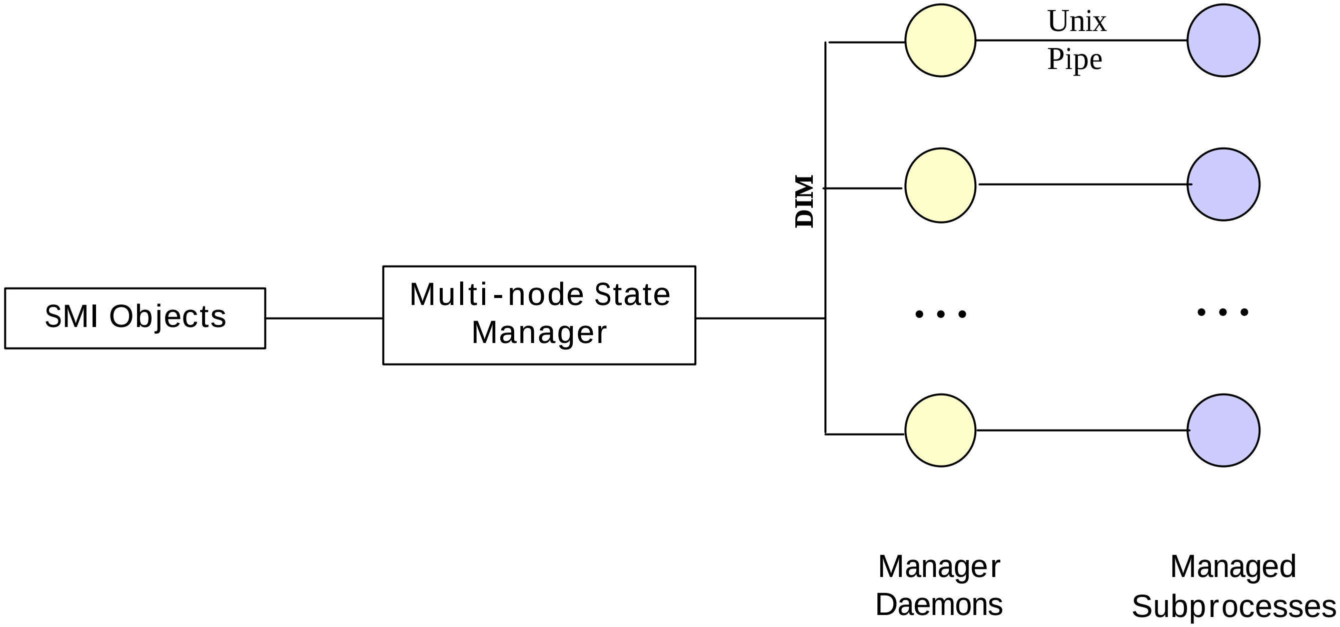

Communication Subsystem

- Used for sending commands and reporting state
- Wanted TCP level of connection
- DIM (Distributed Information Management System) from CERN
  - *Thin layer over TCP*
  - *Provides useful naming service*
  - *Robust disconnect/reconnect support*
  - *Would prefer TCP-like read/write API*

Manager Operation

- NOT another state machine
- Adopts state machine of daemons
  - *Subset of daemon states become manager states*
- Understands READY and ERROR states
- Understands START and RESET commands
  - *Other commands passed through to daemons*

Daemon Operation

- Yet another state machine
- Controlled by SMF (State Machine Framework – Alex Samuel) state diagram
- Starts and monitors subprocess to do actual work.

State Reporting

- Major states
  - *States that are reported back to SMI objects*
    - Examples: ALLOCATED, CONFIGURED, RUNNING
- Minor states
  - *States that display only on manager GUI*
    - Examples: CONNECTING, MAPPED
- Micro states
  - *Internal states of the daemon*
    - Examples: starting
- Error states
  - *Cause the manager to enter ERROR state*
    - Examples: ERROR_EXIT, SIGNALED

Daemon State Transitions

- Commands from Manager
  - *START, RESET, DISCONNECT, typical*
- Subprocess termination
- Subprocess event reports (e.g. dataflow transition events)

Error Conditions

- Multiple daemons report either of the following:
  - *Any error state*
  - *Disappearance (network disconnect)*
- Timeout after START
- Daemons enter conflicting states
- Timeout after some, but not all, daemons change state.
- No commands sent to daemons until RESET

Detection of Hung Daemons

- If a daemon fails to respond to RESET by going to READY, it is marked unavailable, and will not be used in future operations.
- It can be restarted from the GUI.

Configurable Parameters

- Minimum number of nodes needed to START
- Maximum number of nodes to use
- Maximum number of errors to allow
- Length of state convergence timeout

Separate GUI

- One button for each node
  - *Allows kill, restart, and view logfile*
  - *Displays node's current state*
  - *States have colors, defined and reported by daemon*
- Set configuration parameters
- Displays manager state and description of most recent action

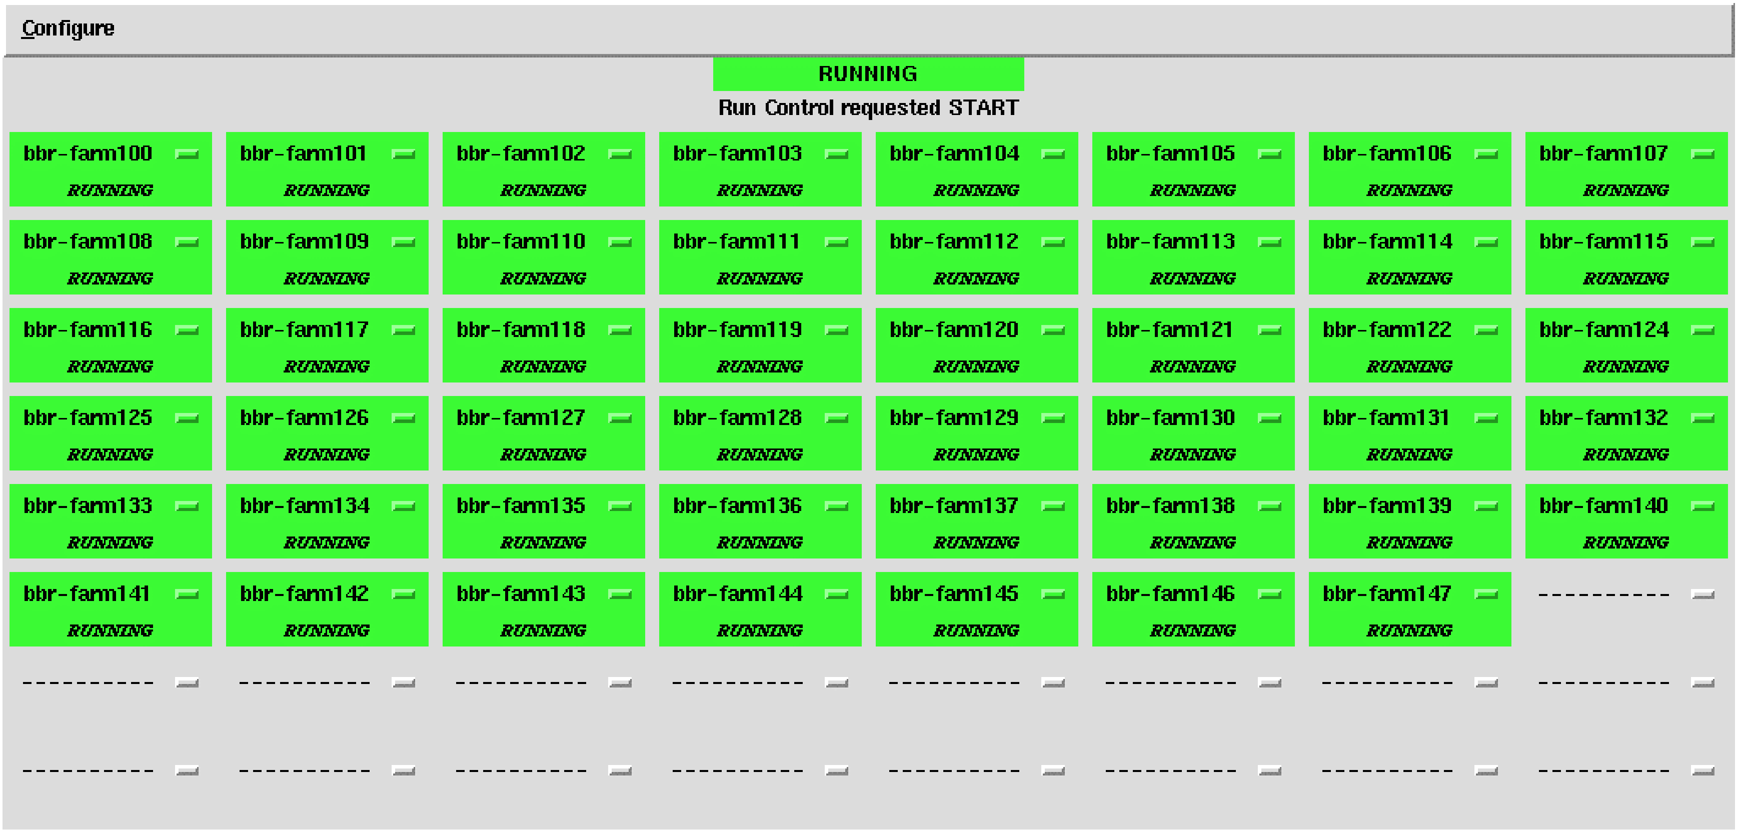

RUNNING

Run Control requested START

bbr-farm100 RUNNING	bbr-farm101 RUNNING	bbr-farm102 RUNNING	bbr-farm103 RUNNING	bbr-farm104 RUNNING	bbr-farm105 RUNNING	bbr-farm106 RUNNING	bbr-farm107 RUNNING
bbr-farm108 RUNNING	bbr-farm109 RUNNING	bbr-farm110 RUNNING	bbr-farm111 RUNNING	bbr-farm112 RUNNING	bbr-farm113 RUNNING	bbr-farm114 RUNNING	bbr-farm115 RUNNING
bbr-farm116 RUNNING	bbr-farm117 RUNNING	bbr-farm118 RUNNING	bbr-farm119 RUNNING	bbr-farm120 RUNNING	bbr-farm121 RUNNING	bbr-farm122 RUNNING	bbr-farm124 RUNNING
bbr-farm125 RUNNING	bbr-farm126 RUNNING	bbr-farm127 RUNNING	bbr-farm128 RUNNING	bbr-farm129 RUNNING	bbr-farm130 RUNNING	bbr-farm131 RUNNING	bbr-farm132 RUNNING
bbr-farm133 RUNNING	bbr-farm134 RUNNING	bbr-farm135 RUNNING	bbr-farm136 RUNNING	bbr-farm137 RUNNING	bbr-farm138 RUNNING	bbr-farm139 RUNNING	bbr-farm140 RUNNING
bbr-farm141 RUNNING	bbr-farm142 RUNNING	bbr-farm143 RUNNING	bbr-farm144 RUNNING	bbr-farm145 RUNNING	bbr-farm146 RUNNING	bbr-farm147 RUNNING	----------
----------	----------	----------	----------	----------	----------	----------	----------
----------	----------	----------	----------	----------	----------	----------	----------

Another Application

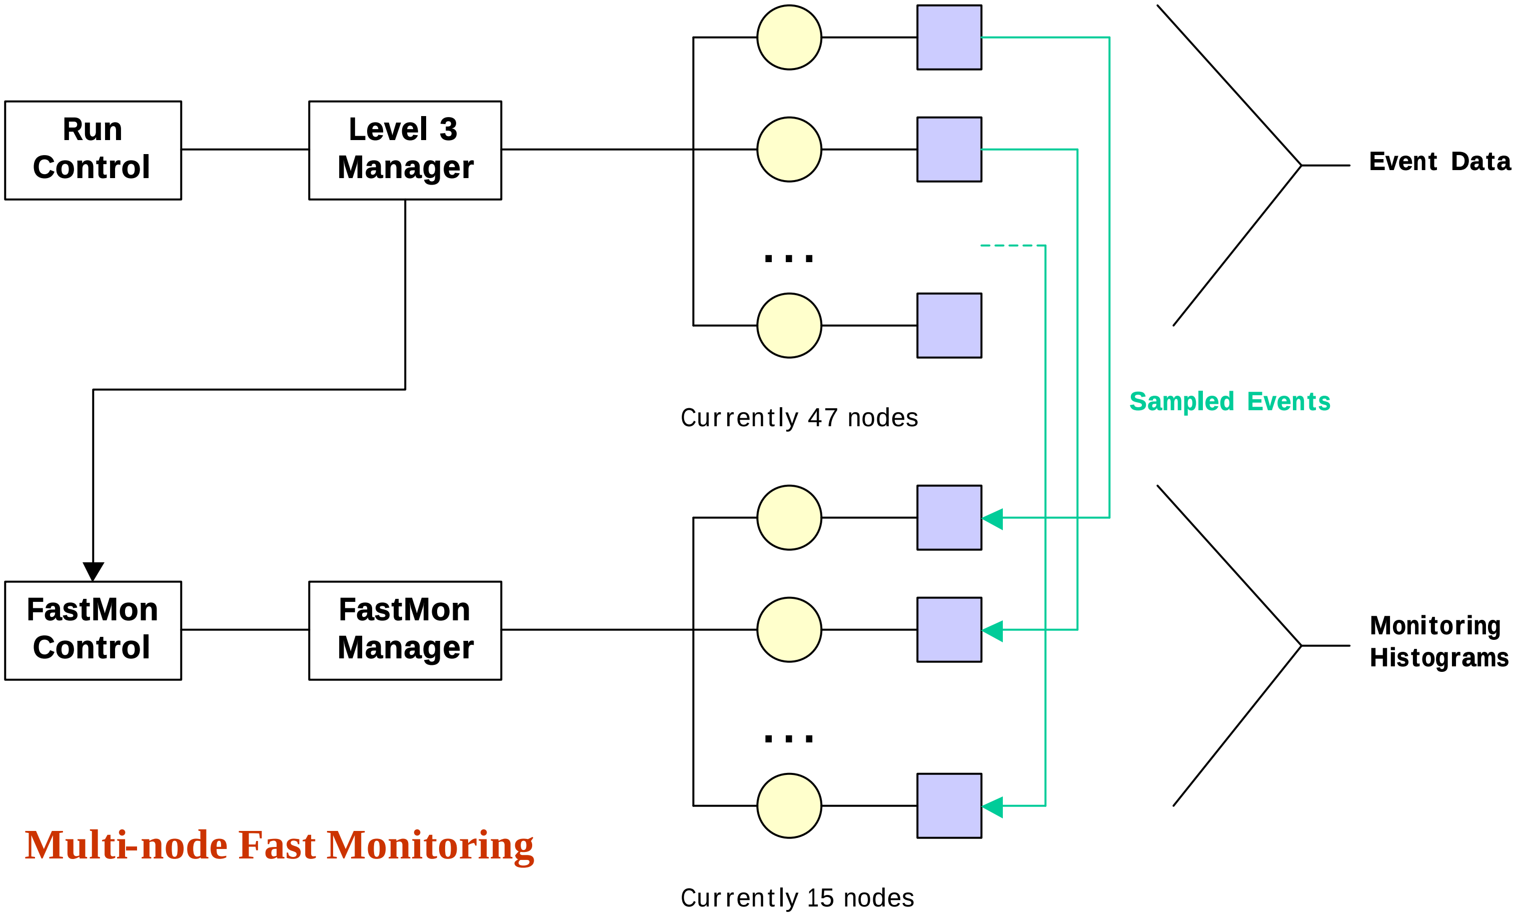

Multi-node Fast Monitoring

Conclusions

- Improved performance
  - *Cold start time reduced to process invocation*
- Improved reliability
  - *Eliminated race conditions*
  - *Robust error behavior*
- Operational since November, 2002