Accelerator Controls and Global Networks –

State of the Art

LINAC 2004, Lubeck, Germany

Dave Gurd

August 20, 2004

Outline

- Two Interpretations of the Controls “Standard Model”
  - » EPICS and LHC
- Issues in New Machines
- Common Trends
- Collaboration and the “GAN”
- Summary and Conclusions

Everyone's Control System – the "CSM"
(as illustrated by the LHC Design)

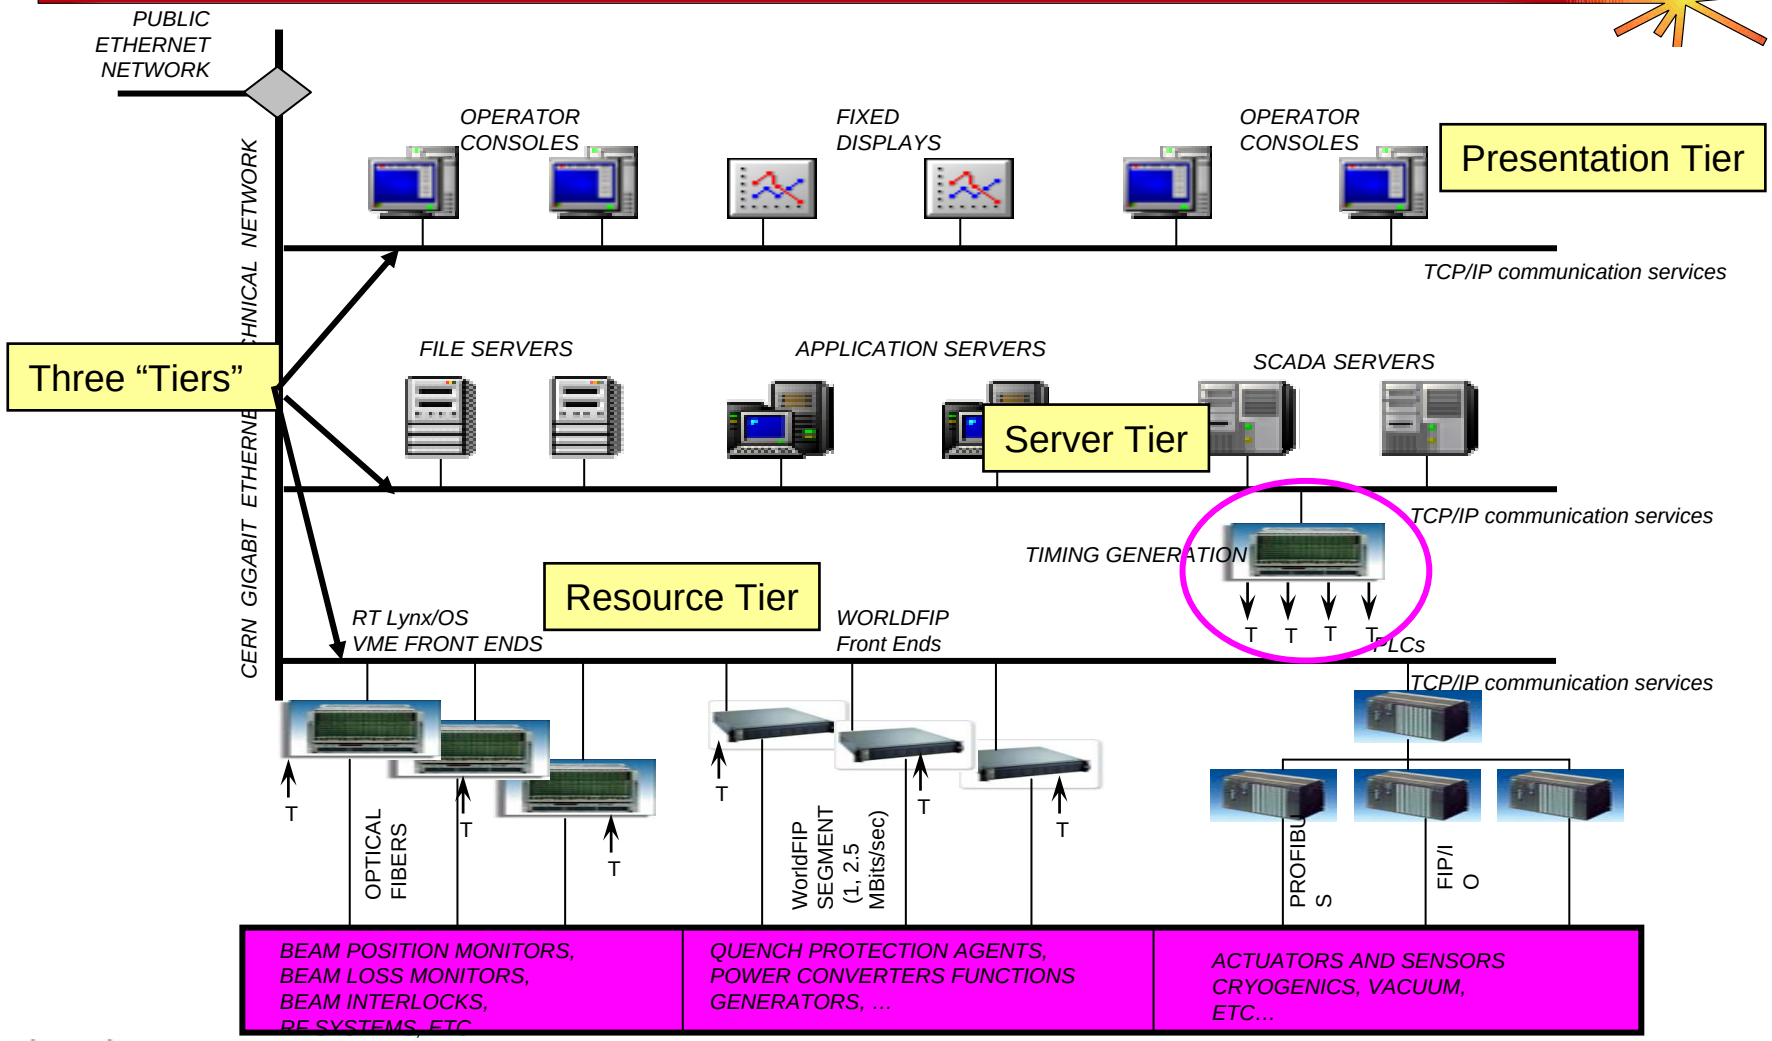

Two Hardware Implementations of the CSM
(One EPICS and one not)

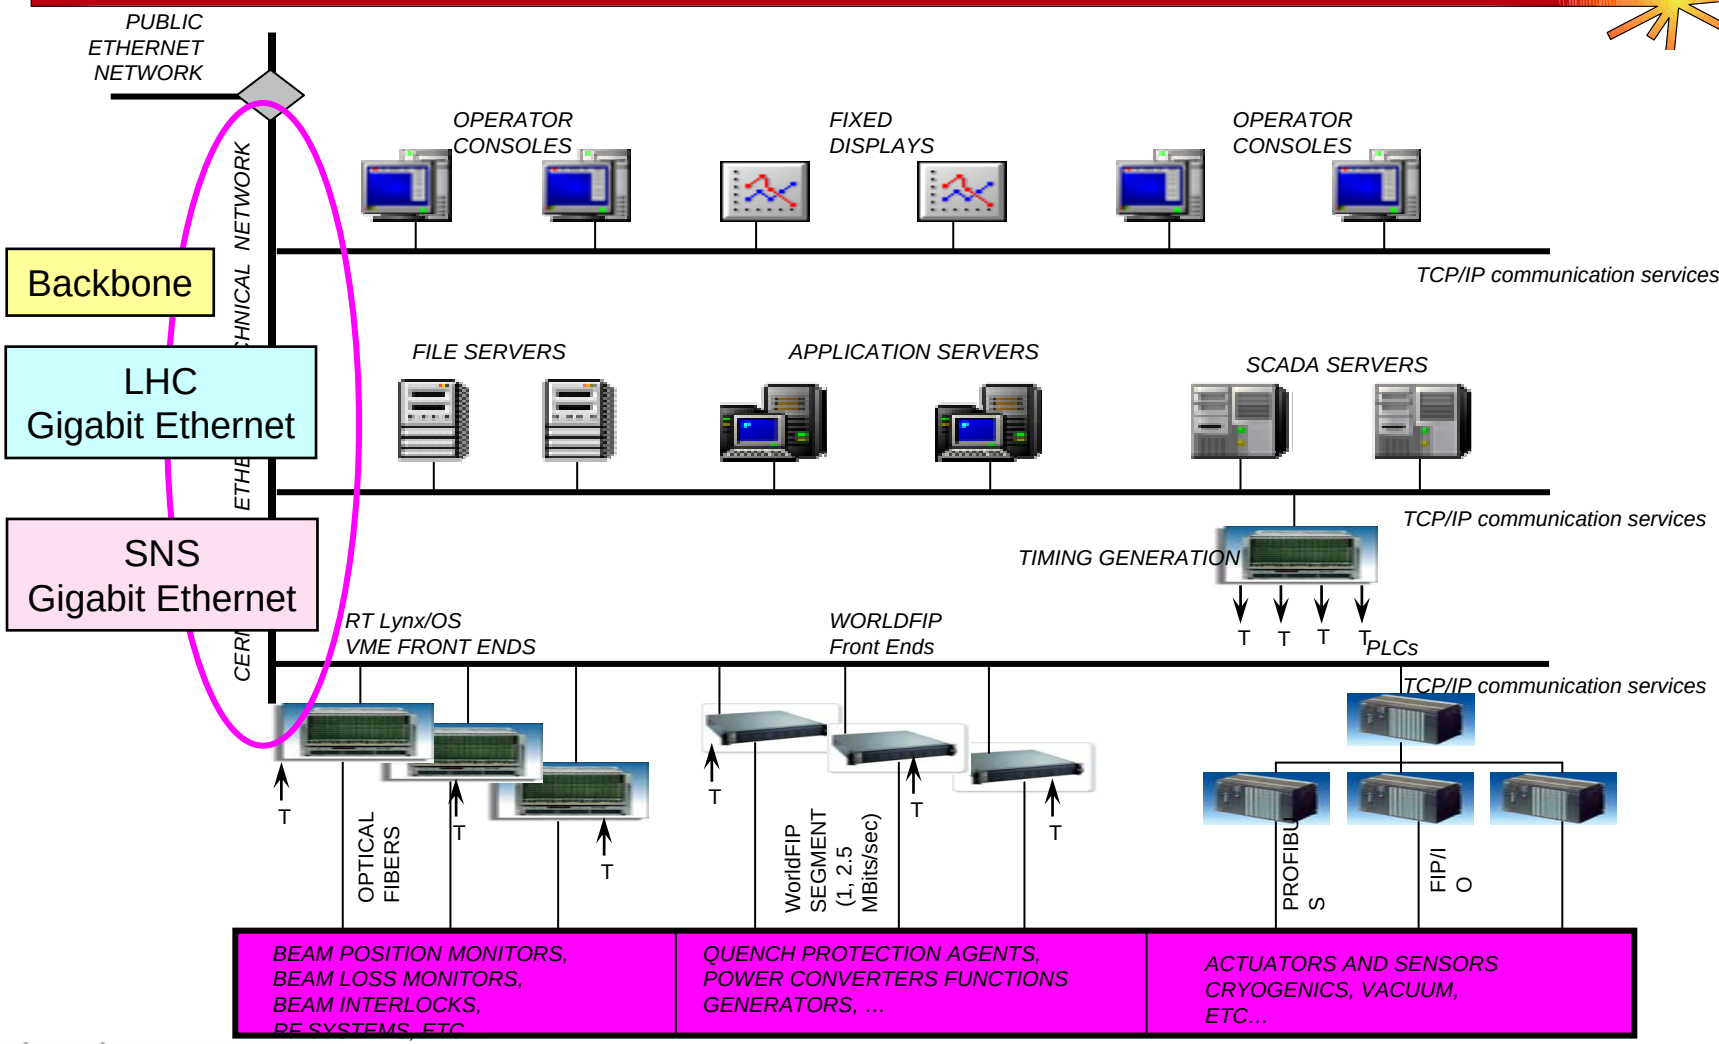

Two Hardware Implementations of the CSM
(One EPICS and one not)

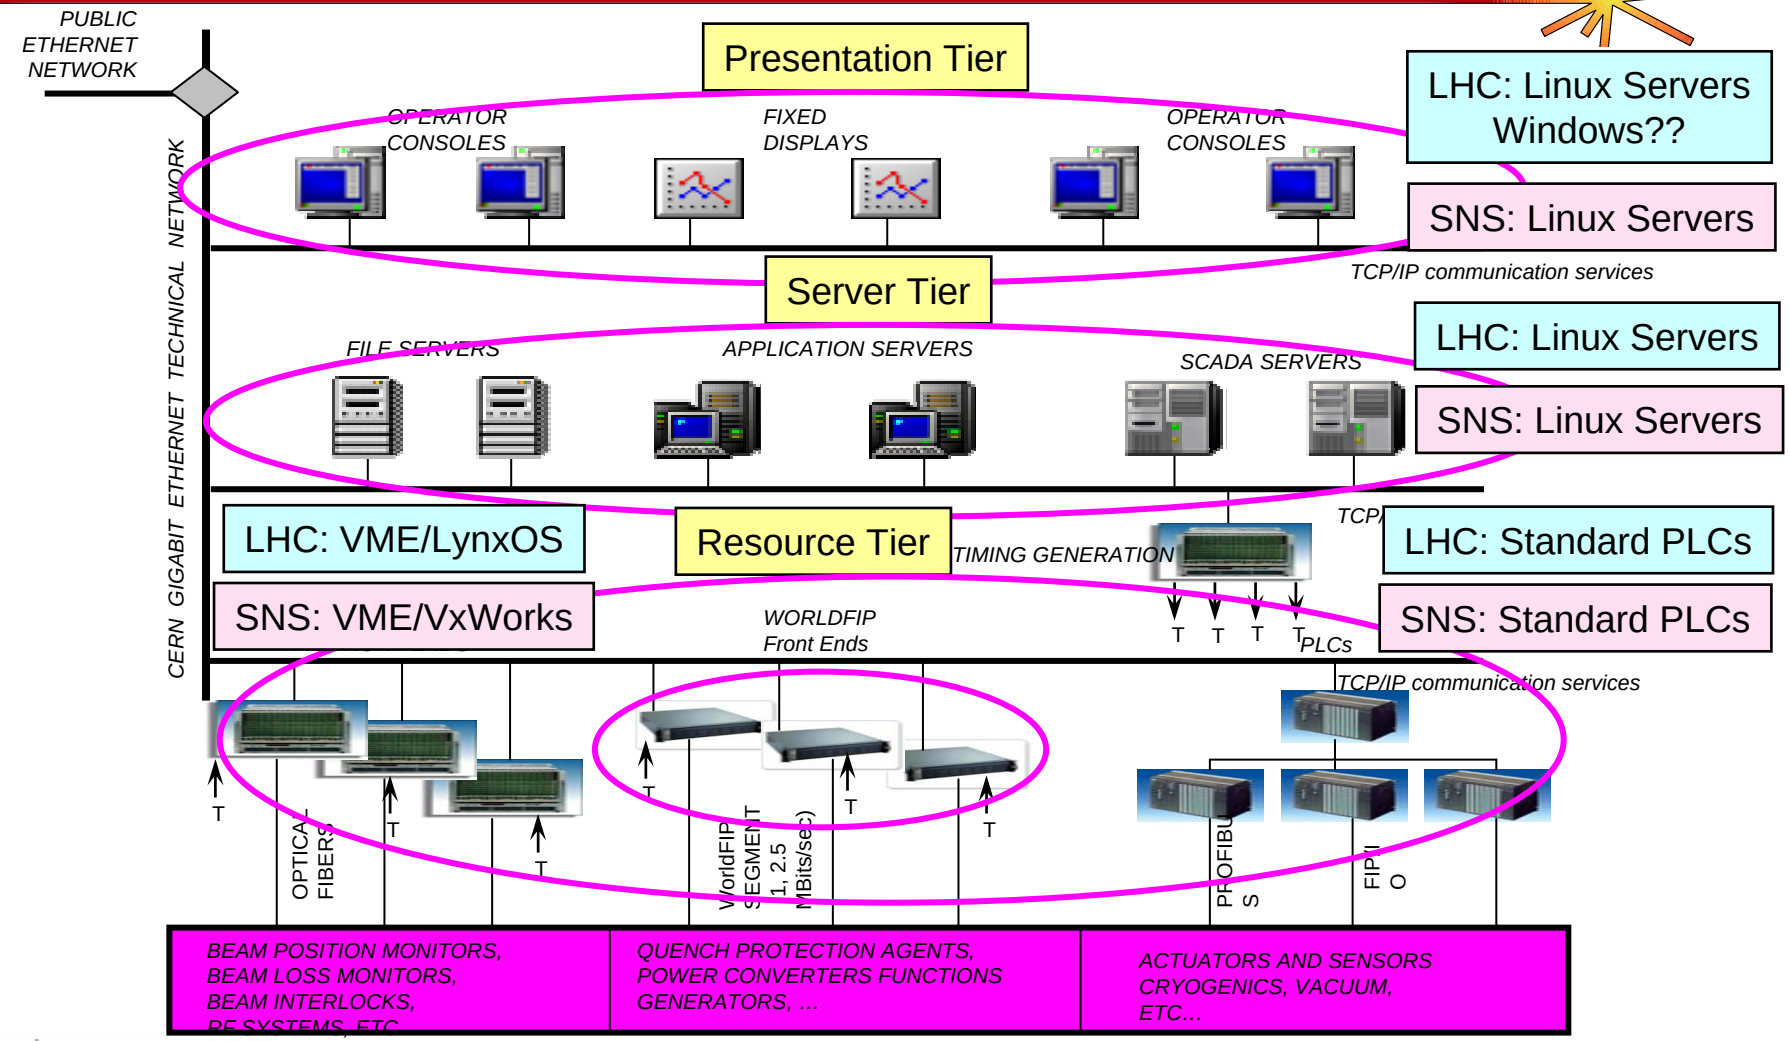

The Difference is in the Software…

- SNS uses EPICS
  - » The first and only example of a successful controls collaboration
  - » After 15 yrs, still the choice of most new facilities (>100 users)
  - » Clear interfaces have made EPICS adaptable to new hardware
  - » EPICS “core” is its distributed database and “Channel Access”
  - » After that, the choice of tools is open – Still room for creativity
  - » Channel Access is optimized for performance

BUT…

  - » EPICS has no concept of an accelerator or its “devices” …
  - » … so it has been difficult to integrate models, applications
  - » “EPICS 2010” committee is looking to the future

The Difference is in the Software… (2)

- LHC uses an “OO Controls Middleware” (CMW)
  - » The collection of protocols, APIs and frameworks that allow the layers to communicate
  - » All services provided, so programmer concentrates on Apps
  - » Concept of a device is inherent, so model is implicit and more natural
  - » Relevant standards for LHC are J2EE, CORBA
  - » There is a communications overhead – but is it “fast enough?”
  - » LHC is not alone - a whole session of the last ICALEPCS was devoted to Controls Middleware
  - » LHC has extended the concept to process control – UNICOS

  - » SNS has developed “XAL” on top of EPICS for modeling

Complexity and Scale Drive Developments

- Number of processors (IOCs or FECs) goes from hundreds to thousands
- Data volume increases linearly
- Data paths increase exponentially
  - » Does the network scale??
- Happily, the pace of network technology mitigates
  - » Switched networks
  - » Gigabit Ethernet – 10 Gbit in the wings

Complexity and Scale (2)

- Configuration Management
  - » Network configuration critical
  - » Many concurrent software versions – need CVS or equivalent
  - » Use of RDB (Oracle) for configuration
  - » “Crawlers” to assure DB is current (APS)

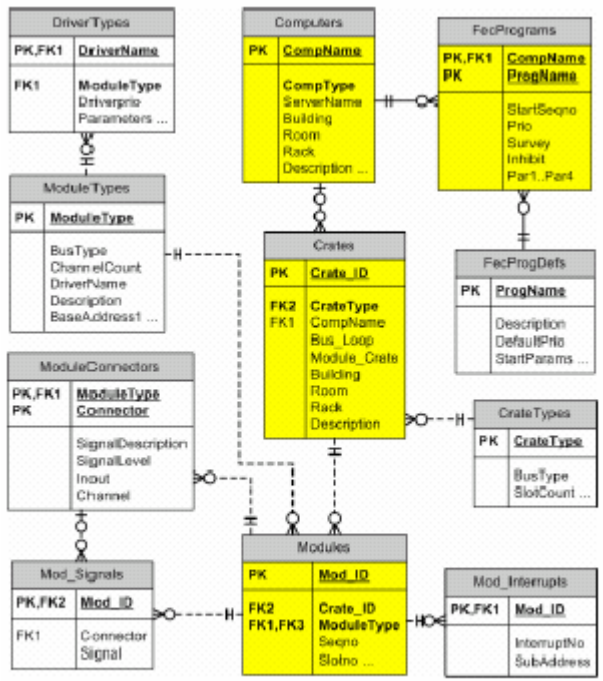

Figure 6: Configuration of FECs, hardware crates and modules.

Complexity and Scale (3)

- Data Management
  - » RHIC 2004 Au/Au run logged ~9 Gbytes/day (SNS same)
  - » LHC anticipates 10^5 – 10^6 variables “routinely” logged plus “several Gigabytes” on events such as quenches
  - » One NLC estimate was 30 *Petabytes (!!)* / year
  - » Data is “bursty,” so need high speed disc I/O
    - – “Entry-level” fiber channel storage array gets 320 Mbytes/sec data transfer rates and 6 Terabyte capacity
  - » Issue is strategies for long term storage, decimation, etc
  - » Good news: Accelerator data volume and rates still << HEP

Complexity and Scale (4) - Reliability

- Controls is not expected to contribute (>99.8%)
- It's the numbers (scale again) that get you
- Demanding requirements (XADS <5 beam trips / yr!!!)
- Industrial hardware
- Redundant hardware
- Hot swap (VME64x)
- SW Testing and QA (NIF)
- No reboots!!

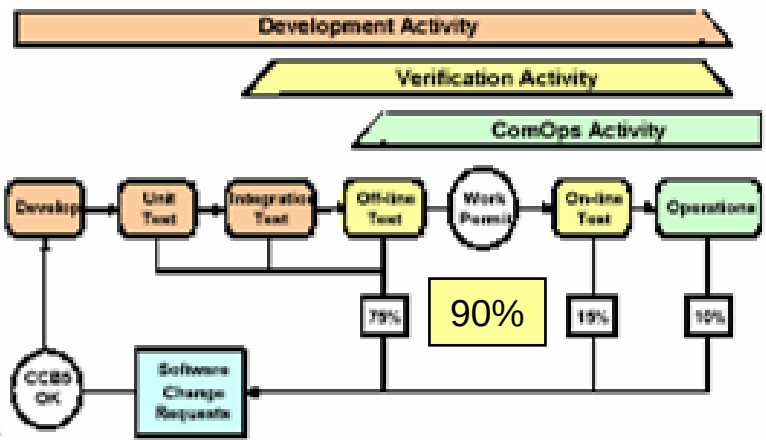

Complexity and Scale (5) – Beam Control

- Diagnostics and Feedback
  - » New machines very demanding
  - » 10X improvement in position?
  - » Single pulse feedback
  - » More embedded (DSP, FPGA)
  - » Dedicated communications

Consider papers by:
Krejcik MO201
Peters MO203
Ross TU302

- Timing and RF Stability
  - » $10e^{-4}$ amplitude; 0.01deg phase
  - » Subpicosecond timing (LCLS, XFEL)
  - » Timing jitter in femtoseconds

Consider papers by:
Ross TU302
Simrock WE103
Bocchetta WE203

Some Trends, Some Fads

- Open (non-proprietary) Systems
  - » Linux – commodity PCs
  - » EPICS itself is open
  - » Open kernels - EPICS dependency on VxWorks removed (CLS/RTEMS)
  - » J2EE is vendor-independent
  - » ORACLE© seems to be a significant exception

- Industrial Systems, Commercial Systems
  - » PLCs for slow, asynchronous processes
  - » LHC extends this to complete SCADA systems
  - » Integration is still required
  - » Turnkey systems including controls – eg SLS linac from ACCEL

Some Trends, Some Fads (2)

- Ethernet as a Fieldbus
  - » Commercial controllers
  - » Timing on Ethernet (PSI)
  - » Achtung!!
    - – Extra traffic
    - – May not handle traffic

J-Parc

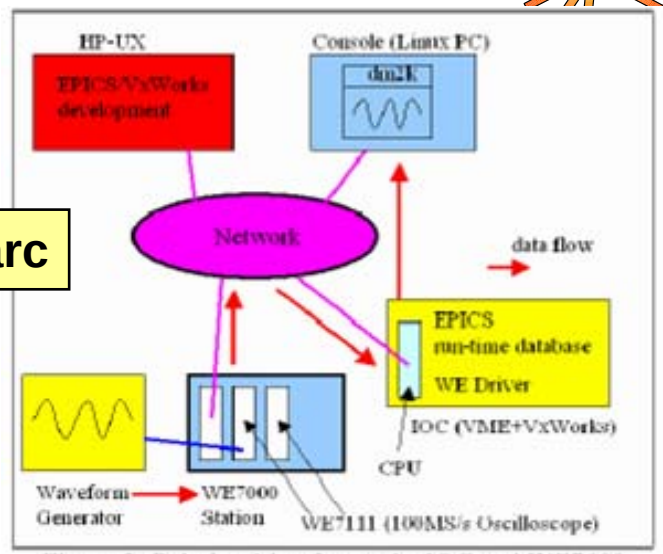

Figure 3: Relationship of console, IOC and WE7000

- PC-Based Controllers
  - » Commercial controllers
  - » SNS “NADS”
  - » Labview© integration
  - » Achtung!
    - – Security

SNS

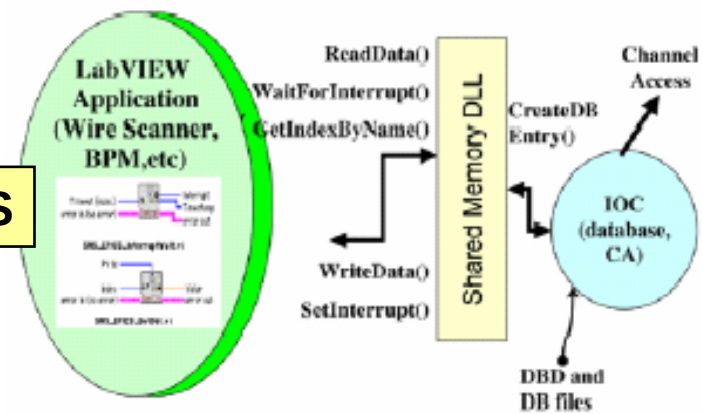

Some Trends, Some Fads (3) – "Middleware"

- CORBA
- JAVA
- RDA
- RPC (Tango)
- etc

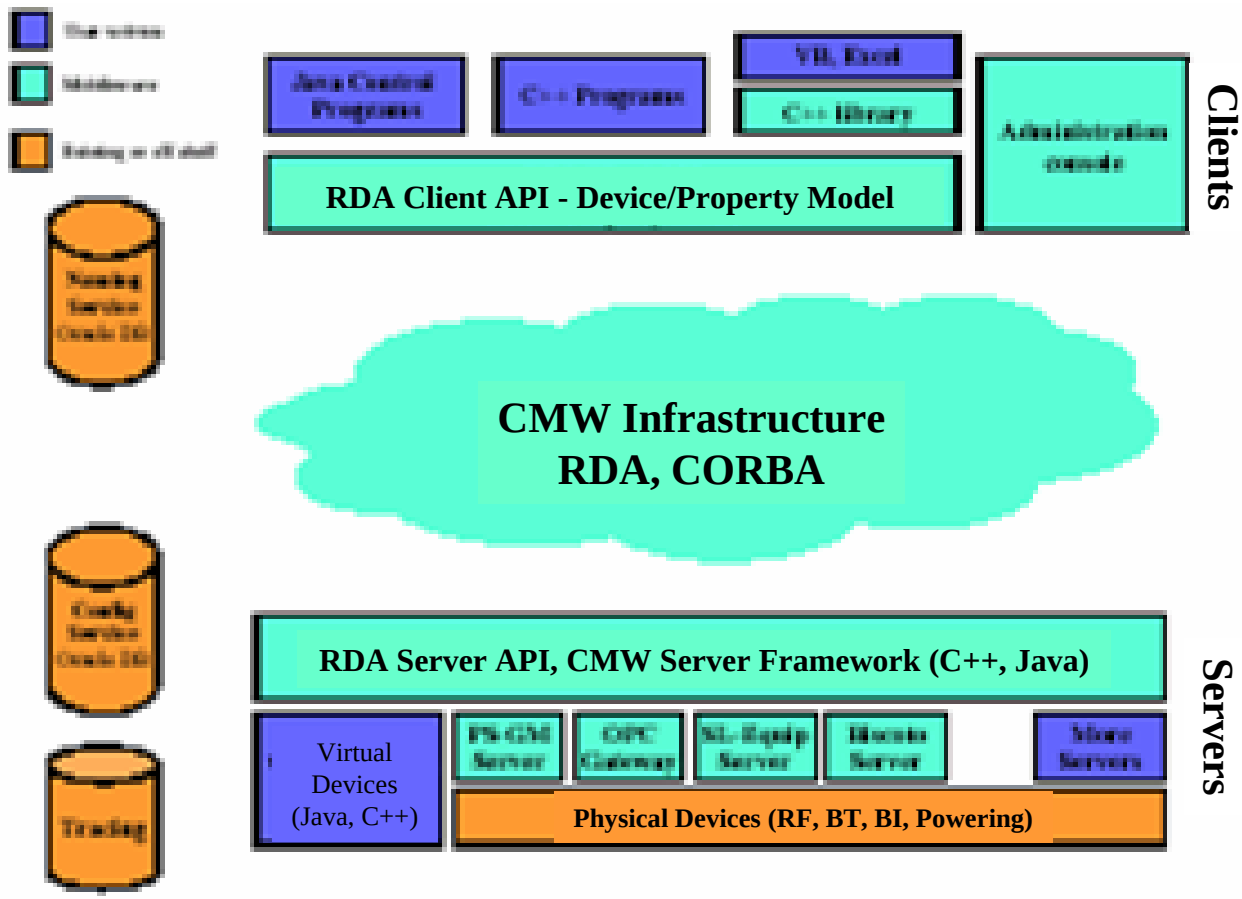

CMW Architecture and Components – LHC Design Study

Collaboration and "GAN"

- New huge machines will be international collaborations
- Can/should control system development be distributed?
- SNS was an experiment – is it a model??
  - » International will be much harder than national
- The "Global Accelerator Network" (GAN)
  - » Network of collaborating institutions – tools for collaboration
  - » Communication network for remote access and control
- A series of workshops: Cornell, Shelter Island, Trieste
- A "rotating" Main Control Room: Pie-in-the-Sky…
  - » … but remote commissioning assistance was useful at SNS
- Biggest issues are sociological and security

Summary

- The Controls “Standard Model” still prevails
  - … but there are various implementations
- EPICS still chosen for new but modest machines
  - … but it needs evolution for new challenges
- Big Issues for Big Machines
  - Precision timing, rf and diagnostics
  - Network Management
  - Data Management
  - Configuration Management
  - Collaboration Management

A Conservative Designer's Credo

"State-of-the-Art" *can* mean "cutting edge,"
but it can also mean "what actually works now."

The conservative designer should minimize risks
and use proven technologies where they will suffice.
Push the limits only where requirements dictate.

You can afford to take only one risk per project.

"Better is the enemy of good enough."

(But more fun.)

Acknowledgements

- I have shamelessly stolen illustrations from ICALEPCS Proceedings and elsewhere

- I have benefited from discussions with many colleagues – and in particular with Hamid Shoaee (LANL), Larry Hoff (BNL) and Bertrand Frammery (CERN)