Redundant EPICS IOCs

Matthias Clausen
Gongfa Liu
Bernd Schoeneburg
(DESY)

Agenda

- High Availability
- Choose the 'right' approach
- Design
- Implementation
- Functionality
  - Redundancy Monitor Task
  - Continuous Control Executive
  - SNL Executive
- Redundancy management
  - Diagnostic
  - SNL Debugger
- Outlook

High Availability Example: File-Server

- Main Service: NFS File Server (1)
  - must be cluster service
- Adding Services: Archiver (2)
  - as cluster service?
  - as '*managed*' service!
- Adding Services: LDAP-HA
  - must be cluster service

  right choice for a file server?

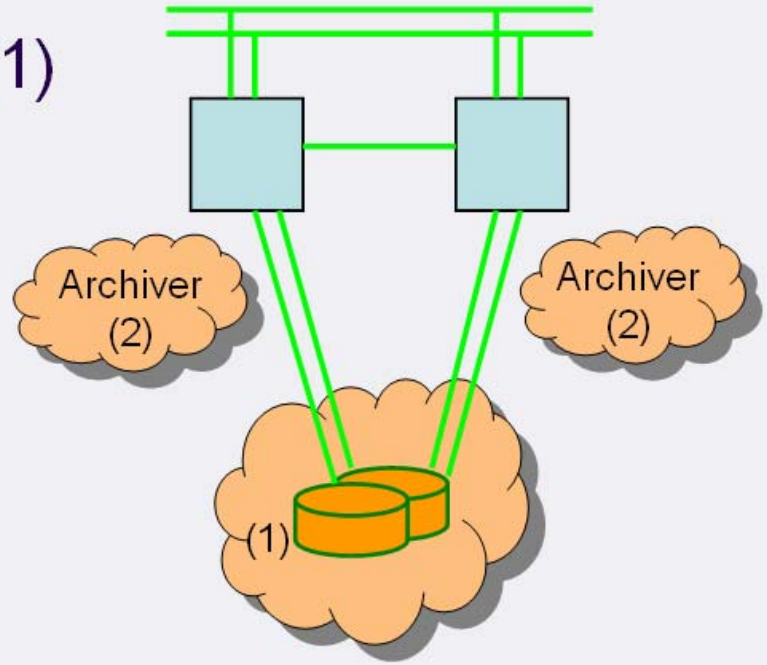

➔ Keep your eye on the main service. Do not allow other services to interfere with the main service.

High Availability: How to implement it?

- *Increase Availability by Following Mill Specs ?*
- Redundant Components !

High Availability: Why?

In our case:

- 24/7 Cryogenic operations for more then one year of operation without any interruption

Necessary for:

- FLASH Cryogenic Plant
  Will be converted from (redundant) commercial to redundant EPICS next year.(1/3 of the system)
- XFEL Cryogenic Plant
  Will be converted from (redundant) commercial to redundant EPICS in 2010.(remaining part)
- XFEL Cryogenic (and possibly Utility) Controls in the XFEL Tunnel

When using redundant IOCs?

- In applications, where high availability is needed and the failure of an IOC can cause a long plant breakdown.
- If you have to be able to maintain the system during operation. Like exchanging a power supply, or loading new software versions.
- If a risk of failure is increased; e.g. in areas, where ionizing radiation can be present.

Design goal:

The redundant IOC pair must be more reliable than a standalone system!

Project Schedule:

- Design Phase (June 2005)
- Identifying the main components:
  - Redundancy Monitor Task
  - Continuous Control Executive
  - SNL Executive
- Implementation Phase (March-September 2006)
  - Redundancy Monitor Task: Industry
  - Continuous Control Executive: Bob and his brother
  - SNL Executive: DESY with SLAC support
- Testing (2006-2007)
- Porting RMT and CCE to other OS (cooperation with KEK)
- Testing (2007)
- Porting to uTCA System (planned)
- Production: Middle of 2008

Layout: RMT, CC-Exec, SNL-Exec

The Redundancy Monitor Task (RMT) is an implementation Independent from EPICS

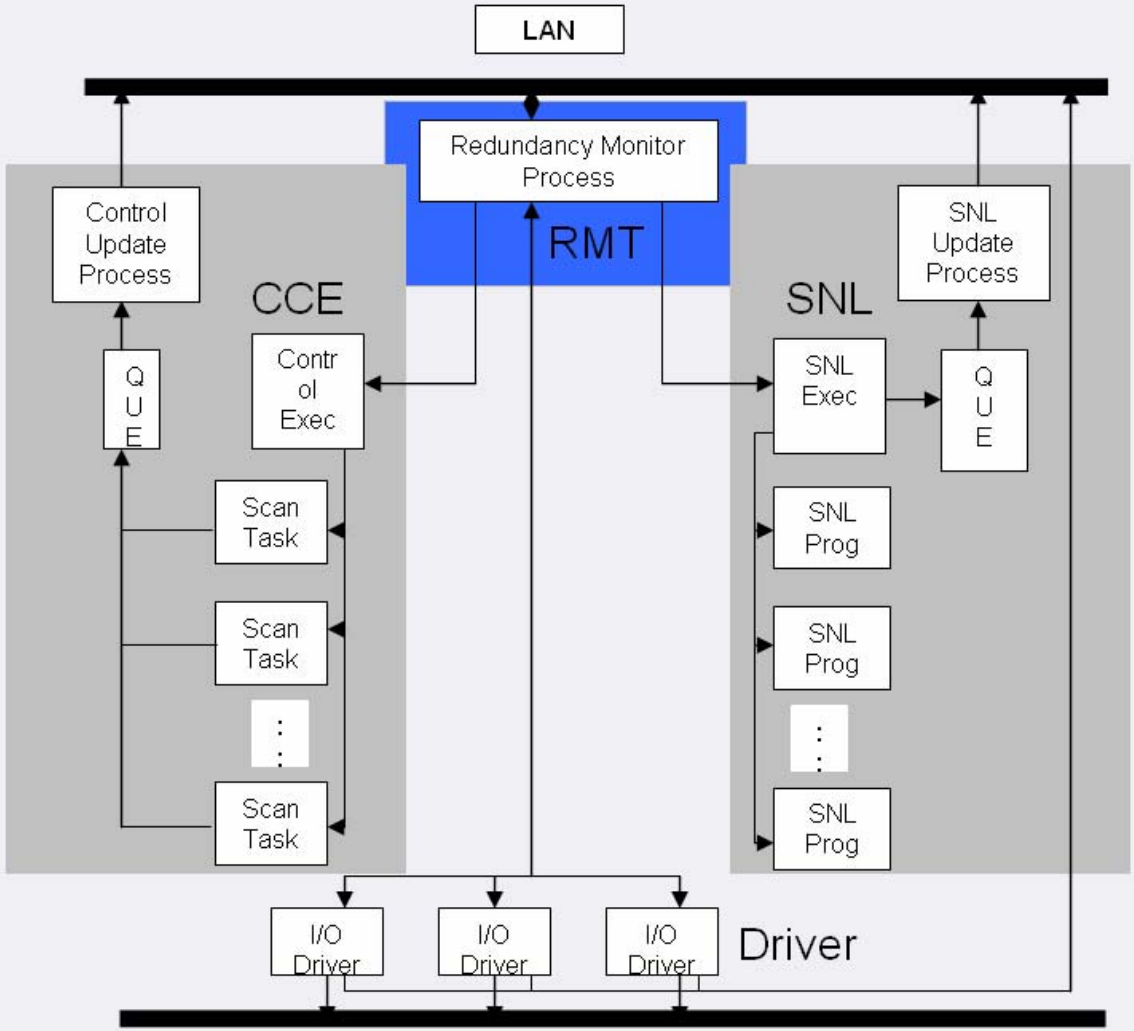

Basic Layout

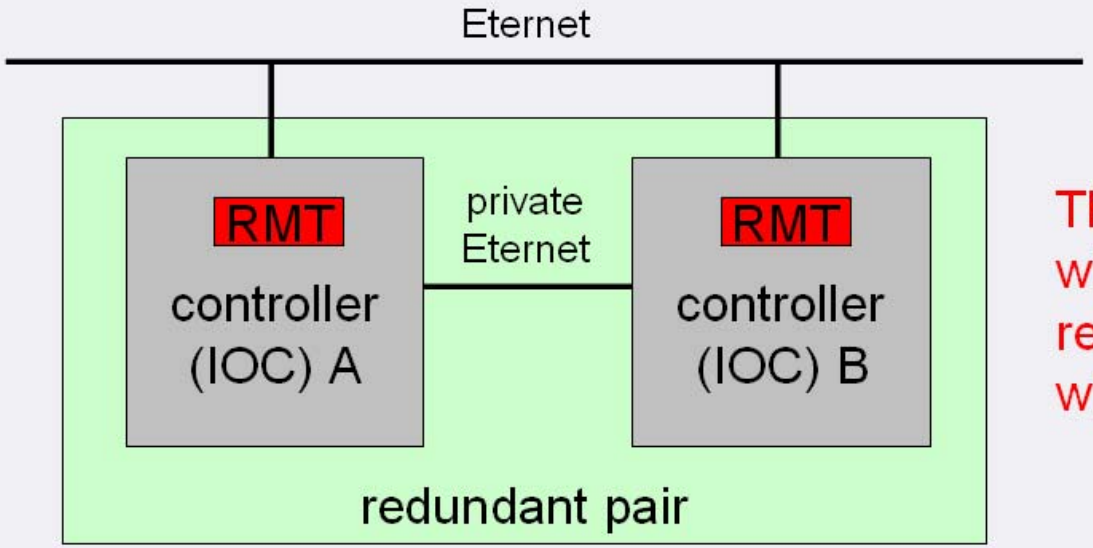

The RMT is a process which handles all redundancy issues within the IOC

- Two IOC form a redundant pair
- One controller is active (Master state)
- The other IOC keeps synchronized with the Master

Basic Layout

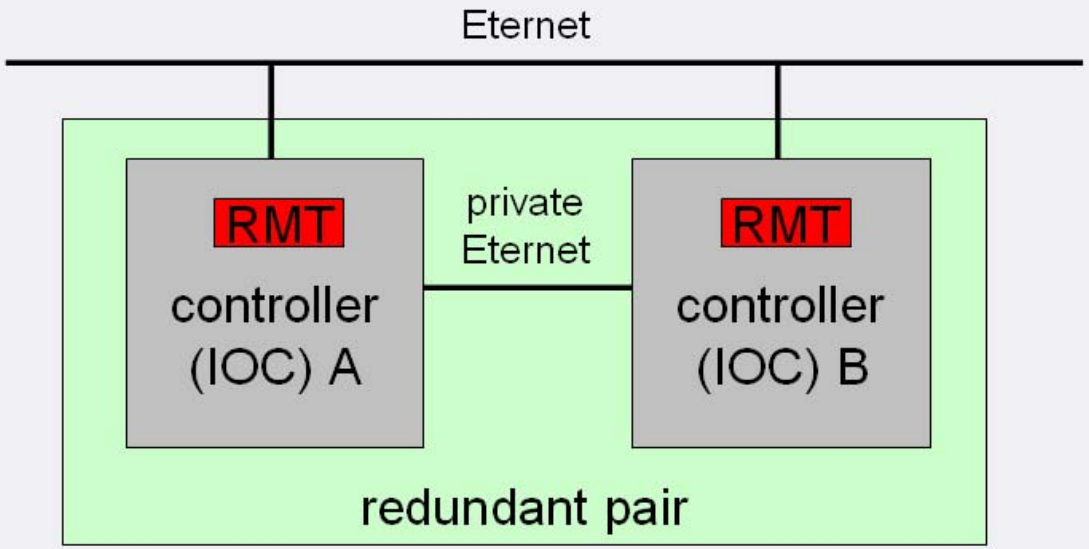

Ethernet Connectivity

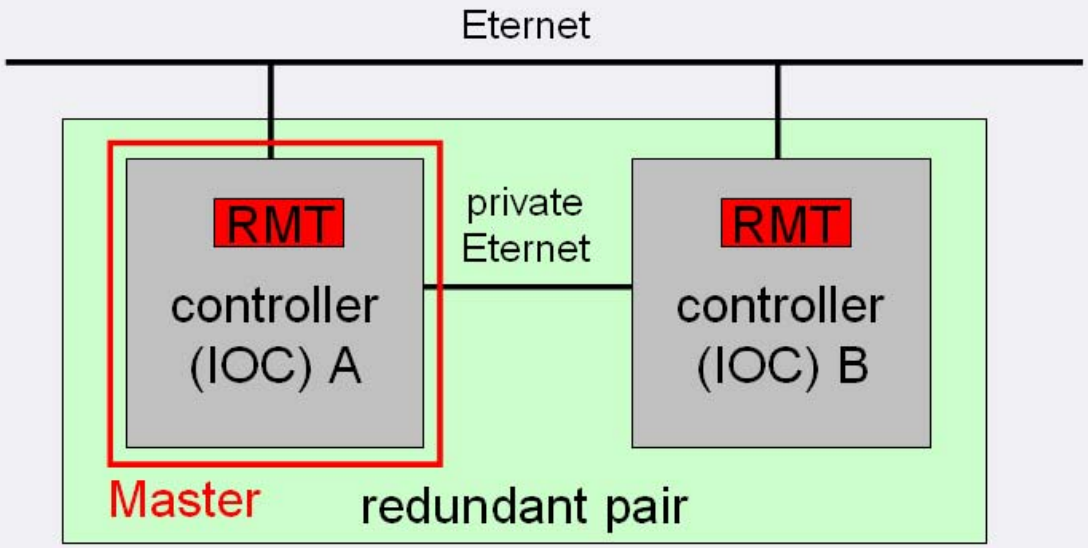

Ethernet Connectivity

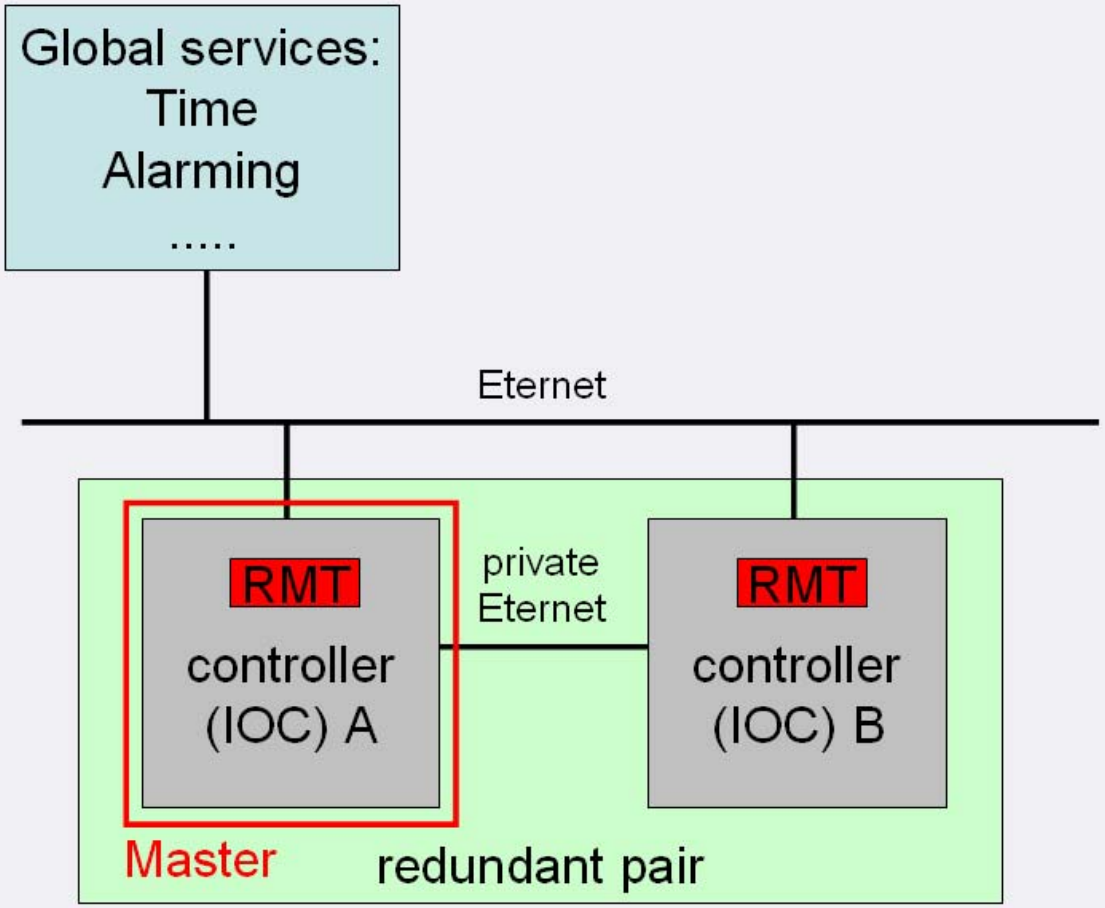

Ethernet Connectivity

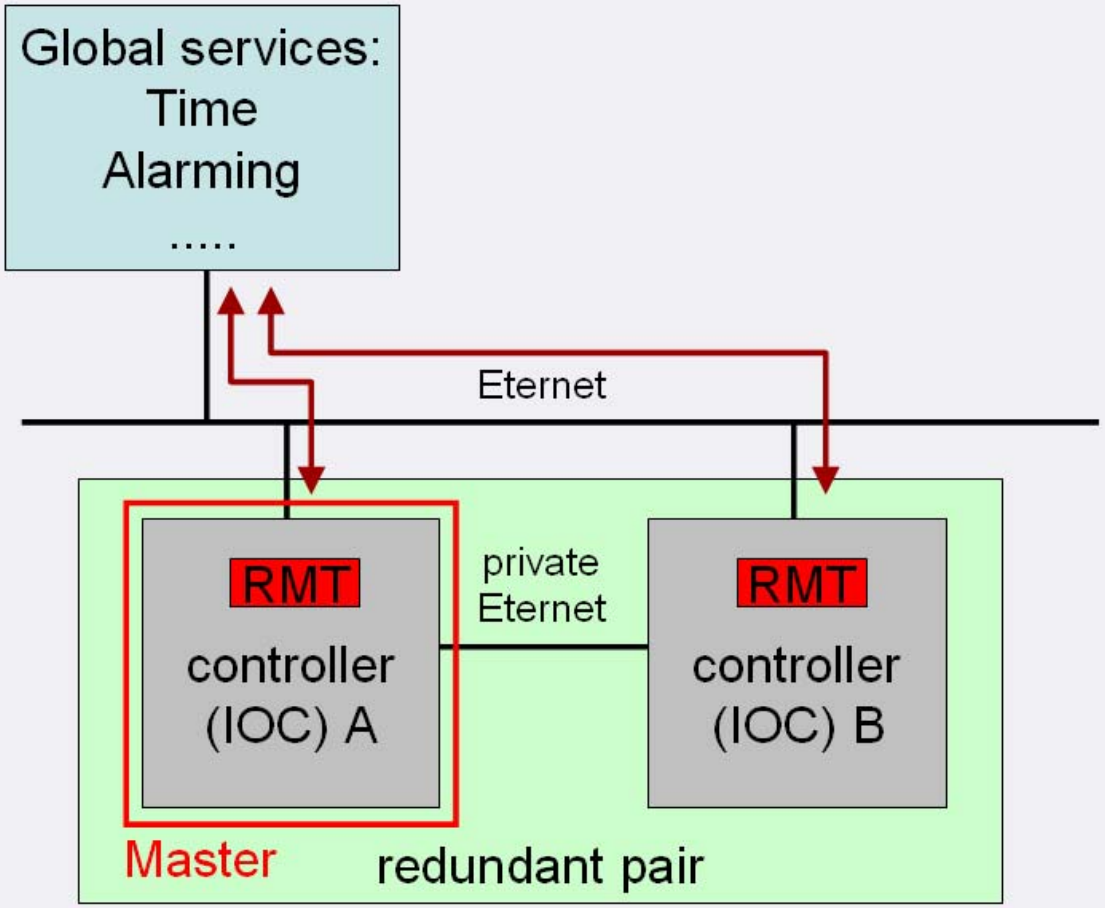

- check global Ethernet

Ethernet Connectivity

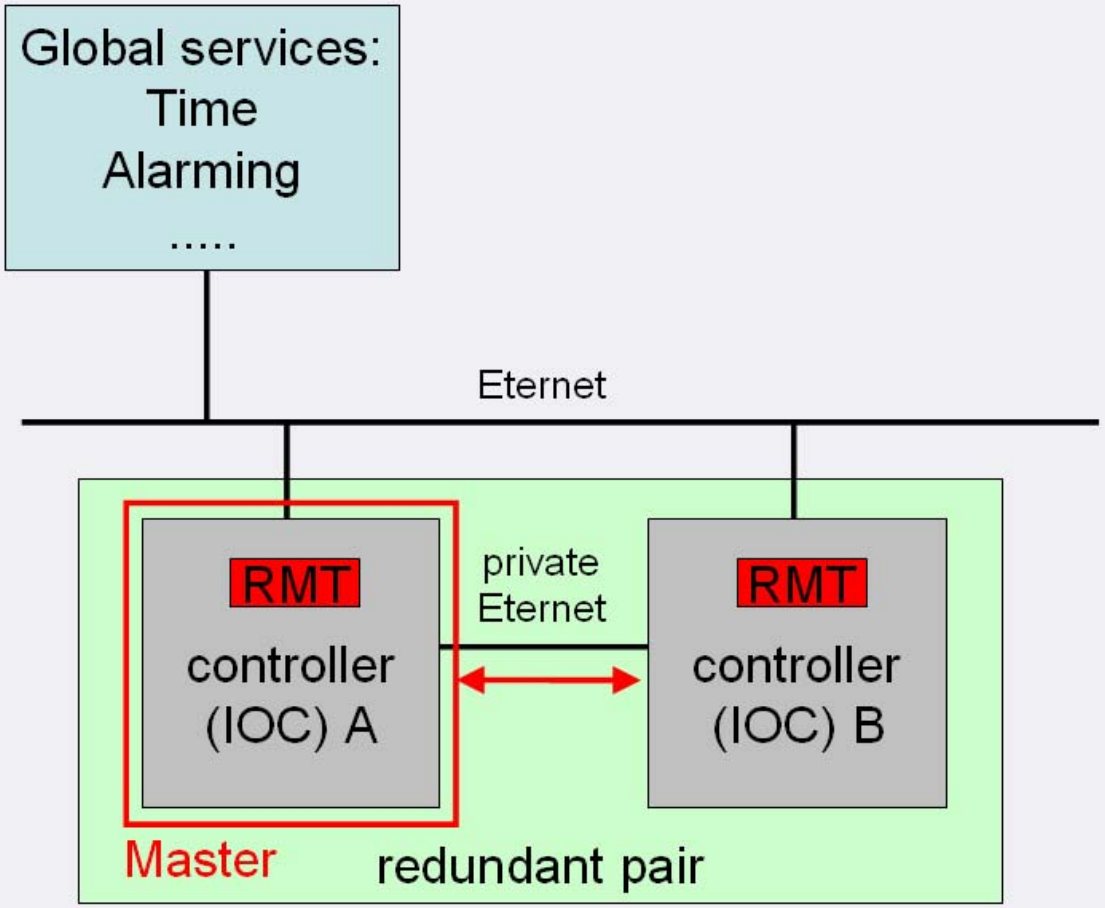

- check global Ethernet
- check public Ethernet
- check private Ethernet

Ethernet Connectivity

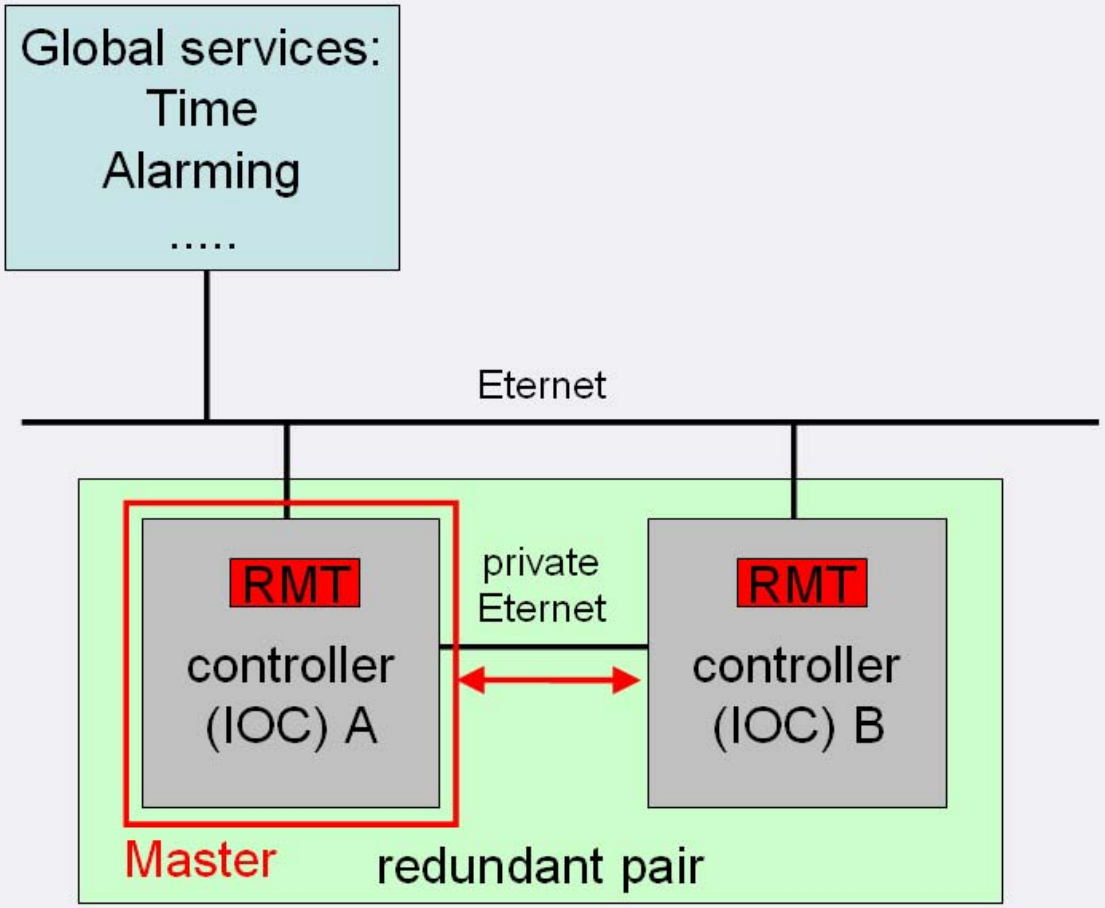

- check global Ethernet
- check public Ethernet
- check private Ethernet
- RMT communication
- synchronization of data

RMT: the Redundancy Supervisor

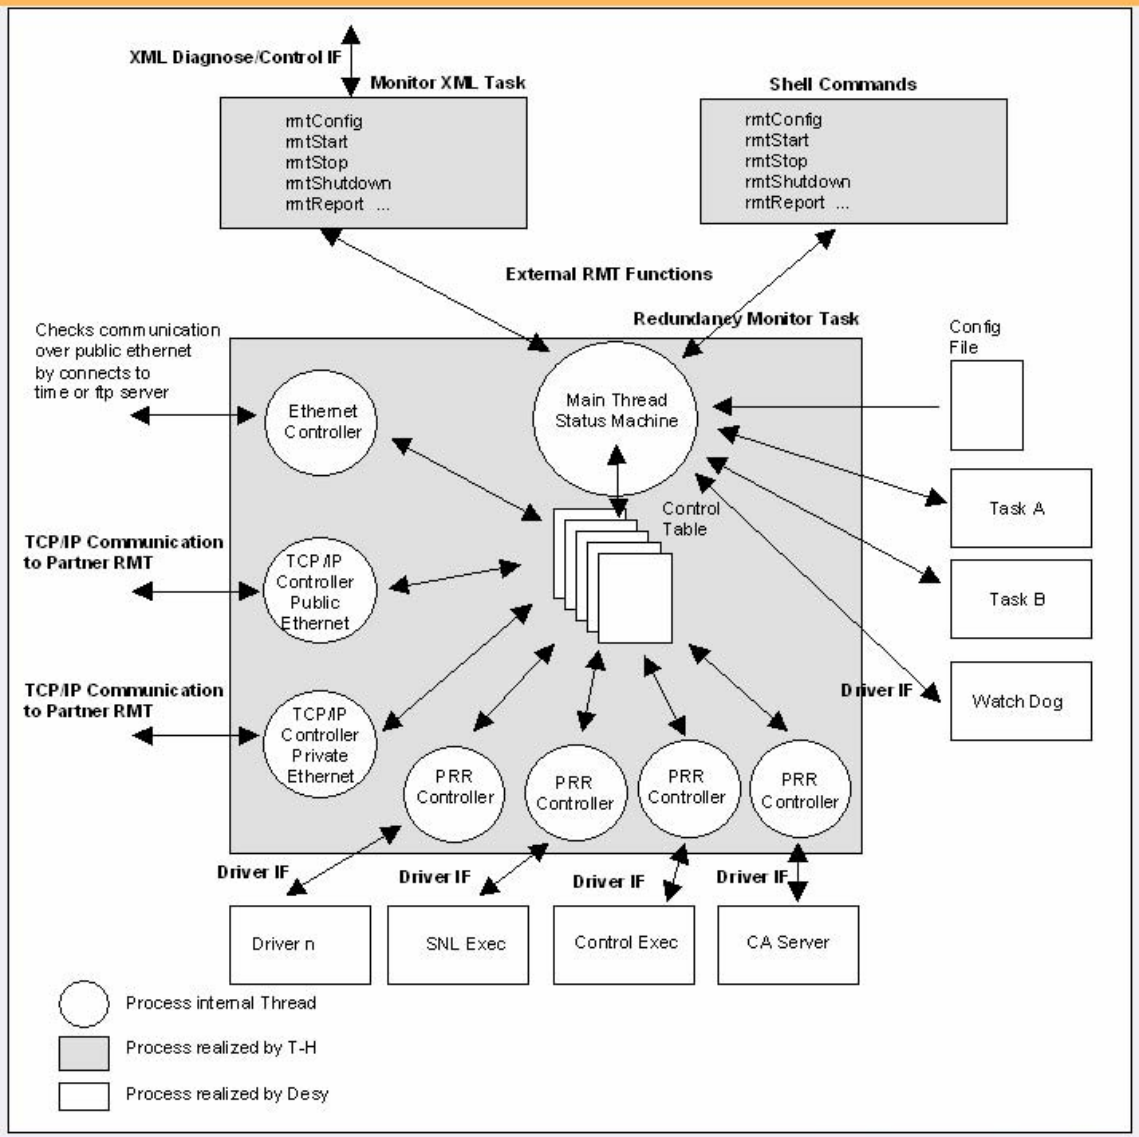

How the RMT works

- Control the processes of interest for redundancy
  - Processes register themselves to be controlled
- Communicate with the RMT in the other IOC
- Set the IOC in the Master- or Slave-state (manage switch-over)
- Monitor network connections (slide before)

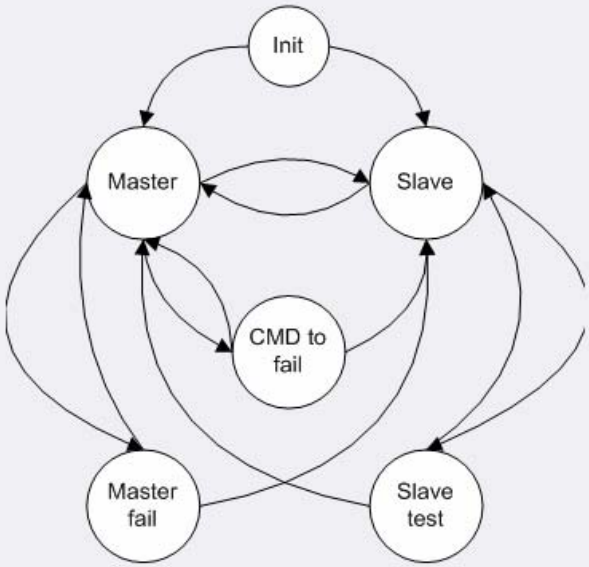

RMT State Machine
(sinplified)

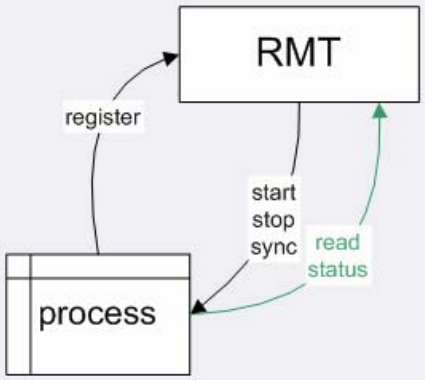

In a redundant IOC PRR-processes register by calling a RMT function and wait for a start command.

Some processes need to be synchronized with their partner process in the other IOC. Synchronization over the private Ethernet is controlled by the RMT.

When to fail over?

From the Preamble of the design Document:

"Any redundant implementation must make the system more reliable than the non redundant one. Precaution must be taken especially for the detection of errors which shall initiate the failover. *This operation should only be activated if there is no doubt that keeping the actual mastership definitely causes more damage to the controlled system than an automatic failover."*

EPICS Specific Parts

RMT is an implementation Independent from EPICS

CCE: The Continuous Control Executive permanently collects changes on the master to update the client

SNL Executive: Permanently collects states and values from SNL programs on the master. Sending changes to the slave.

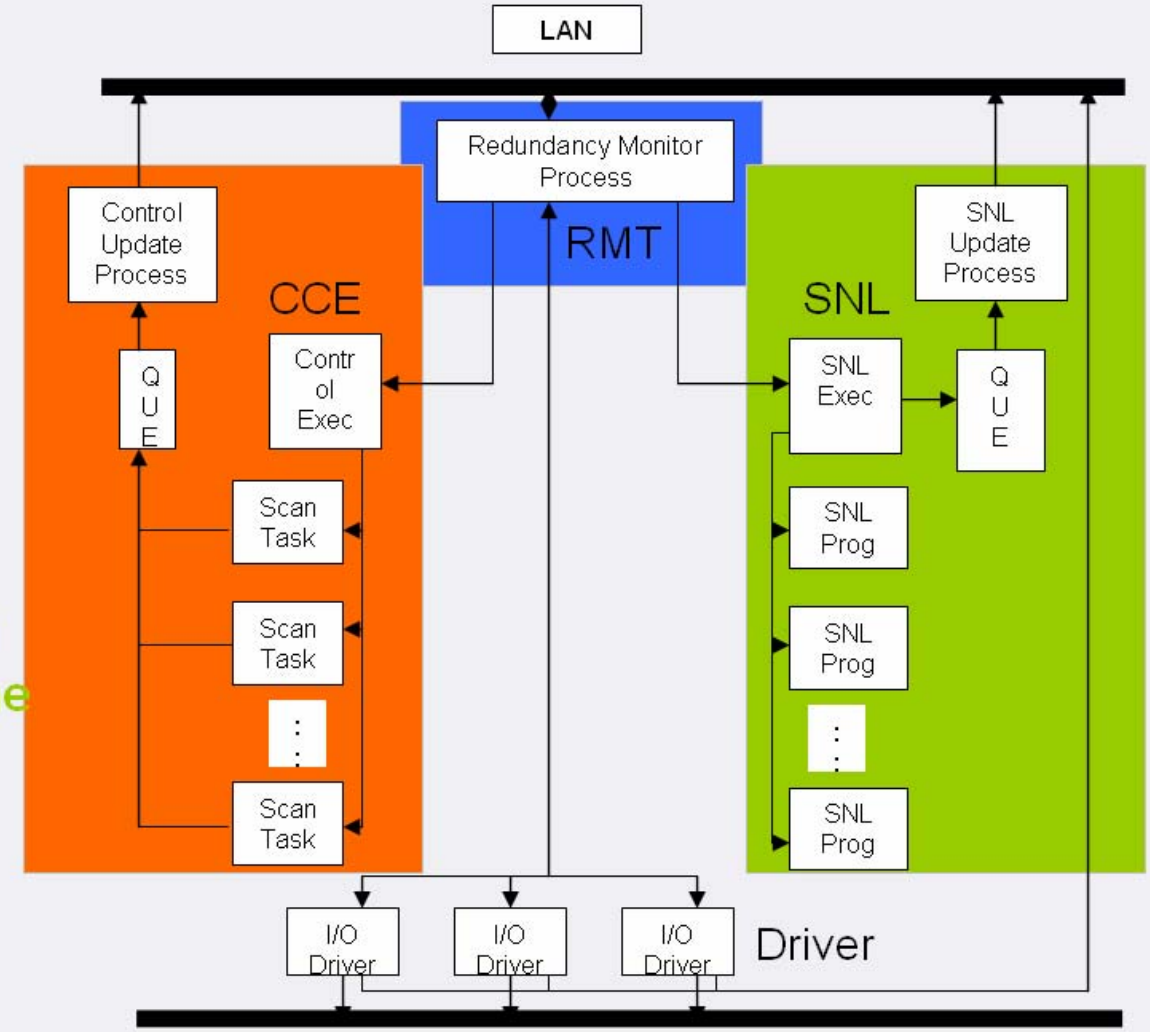

Remote Diagnostic

XML Request files can be passed though the RMT to any registered underlying process.
The final destination of the Message will generate an answer.

SNL Executive:
In this special case the remote diagnostic protocol is used to debug the SLS programs actually Executive running on the remote IOC.

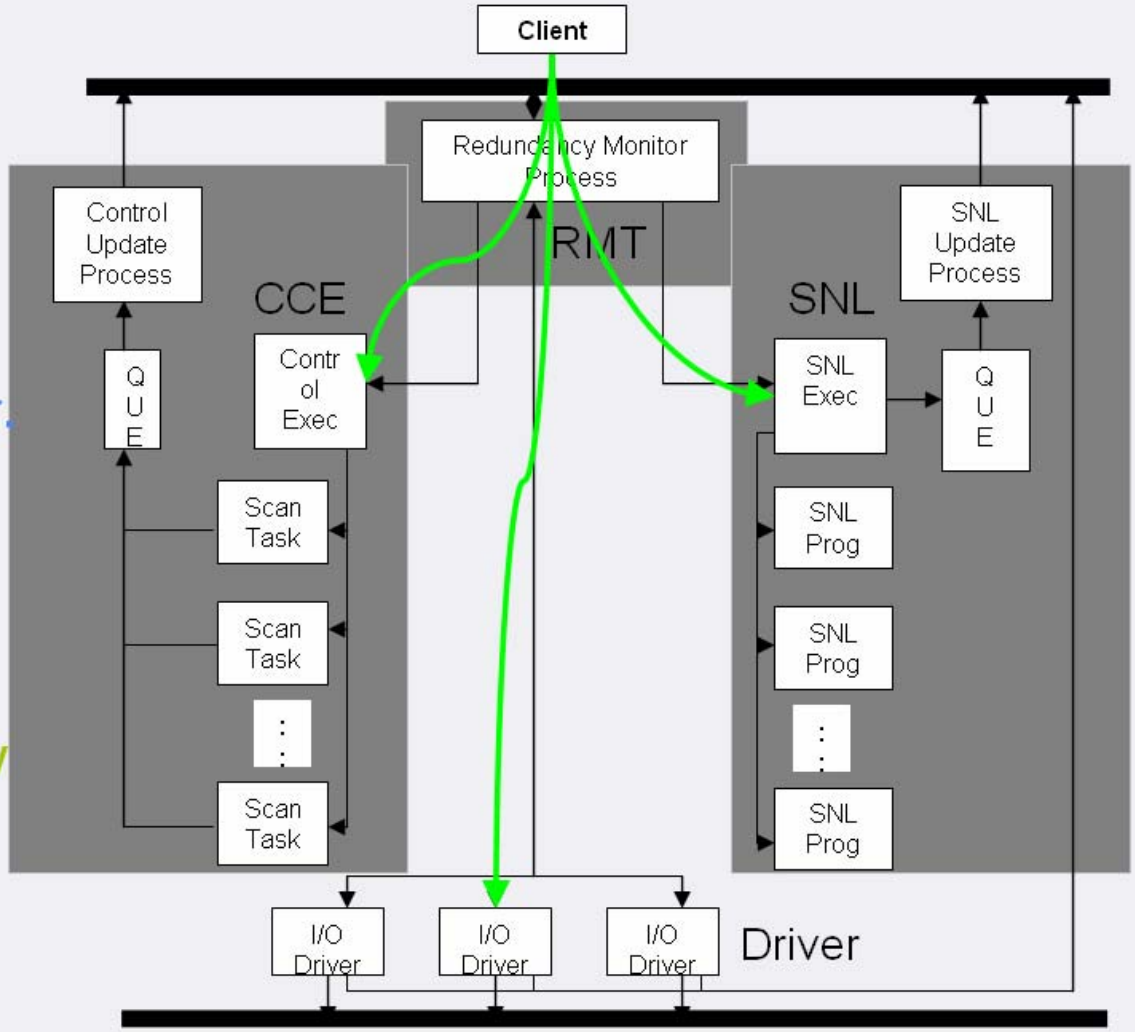

Status

- Implementing support for redundant systems is quite a challenging task.
- The current implementation is improving its maturity by continuous testing. (Gongfa Liu – (Hefei-China) DESY)
- The RMT and CCE code has bee ported to Linux by Artem Kazakov (KEK) to run redundant (soft)-IOCs on Linux. (see TPPA31)
- The ported RMT has been used to implement redundant CA-Gateways. (Artem Kazakov)
  (An option for load balancing is also available)

Outlook

- RMT can be used independent from EPICS to implement redundant applications.
- The SNL debugging features will be improved. (Joint development of SLAC and DESY)

- First production of a redundant IOC is foreseen for middle of 2008
- An implementation based/ running on uTCA is desired.

Test System

The test system consists of two Compact PCI CPUs in separated crates with redundant power supplies each.

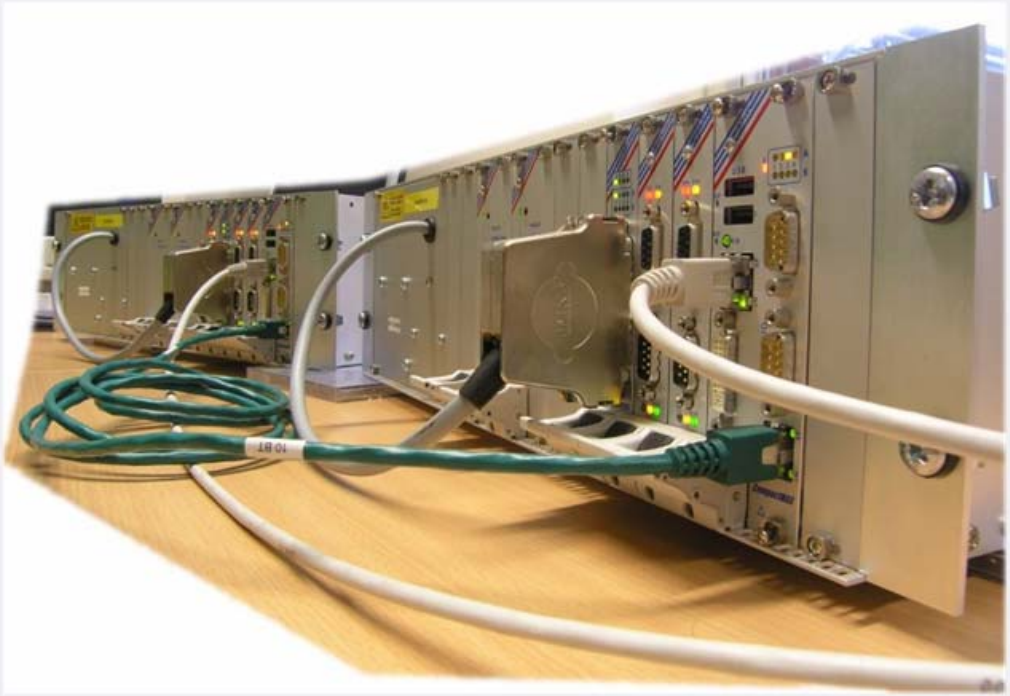

Thank you!