

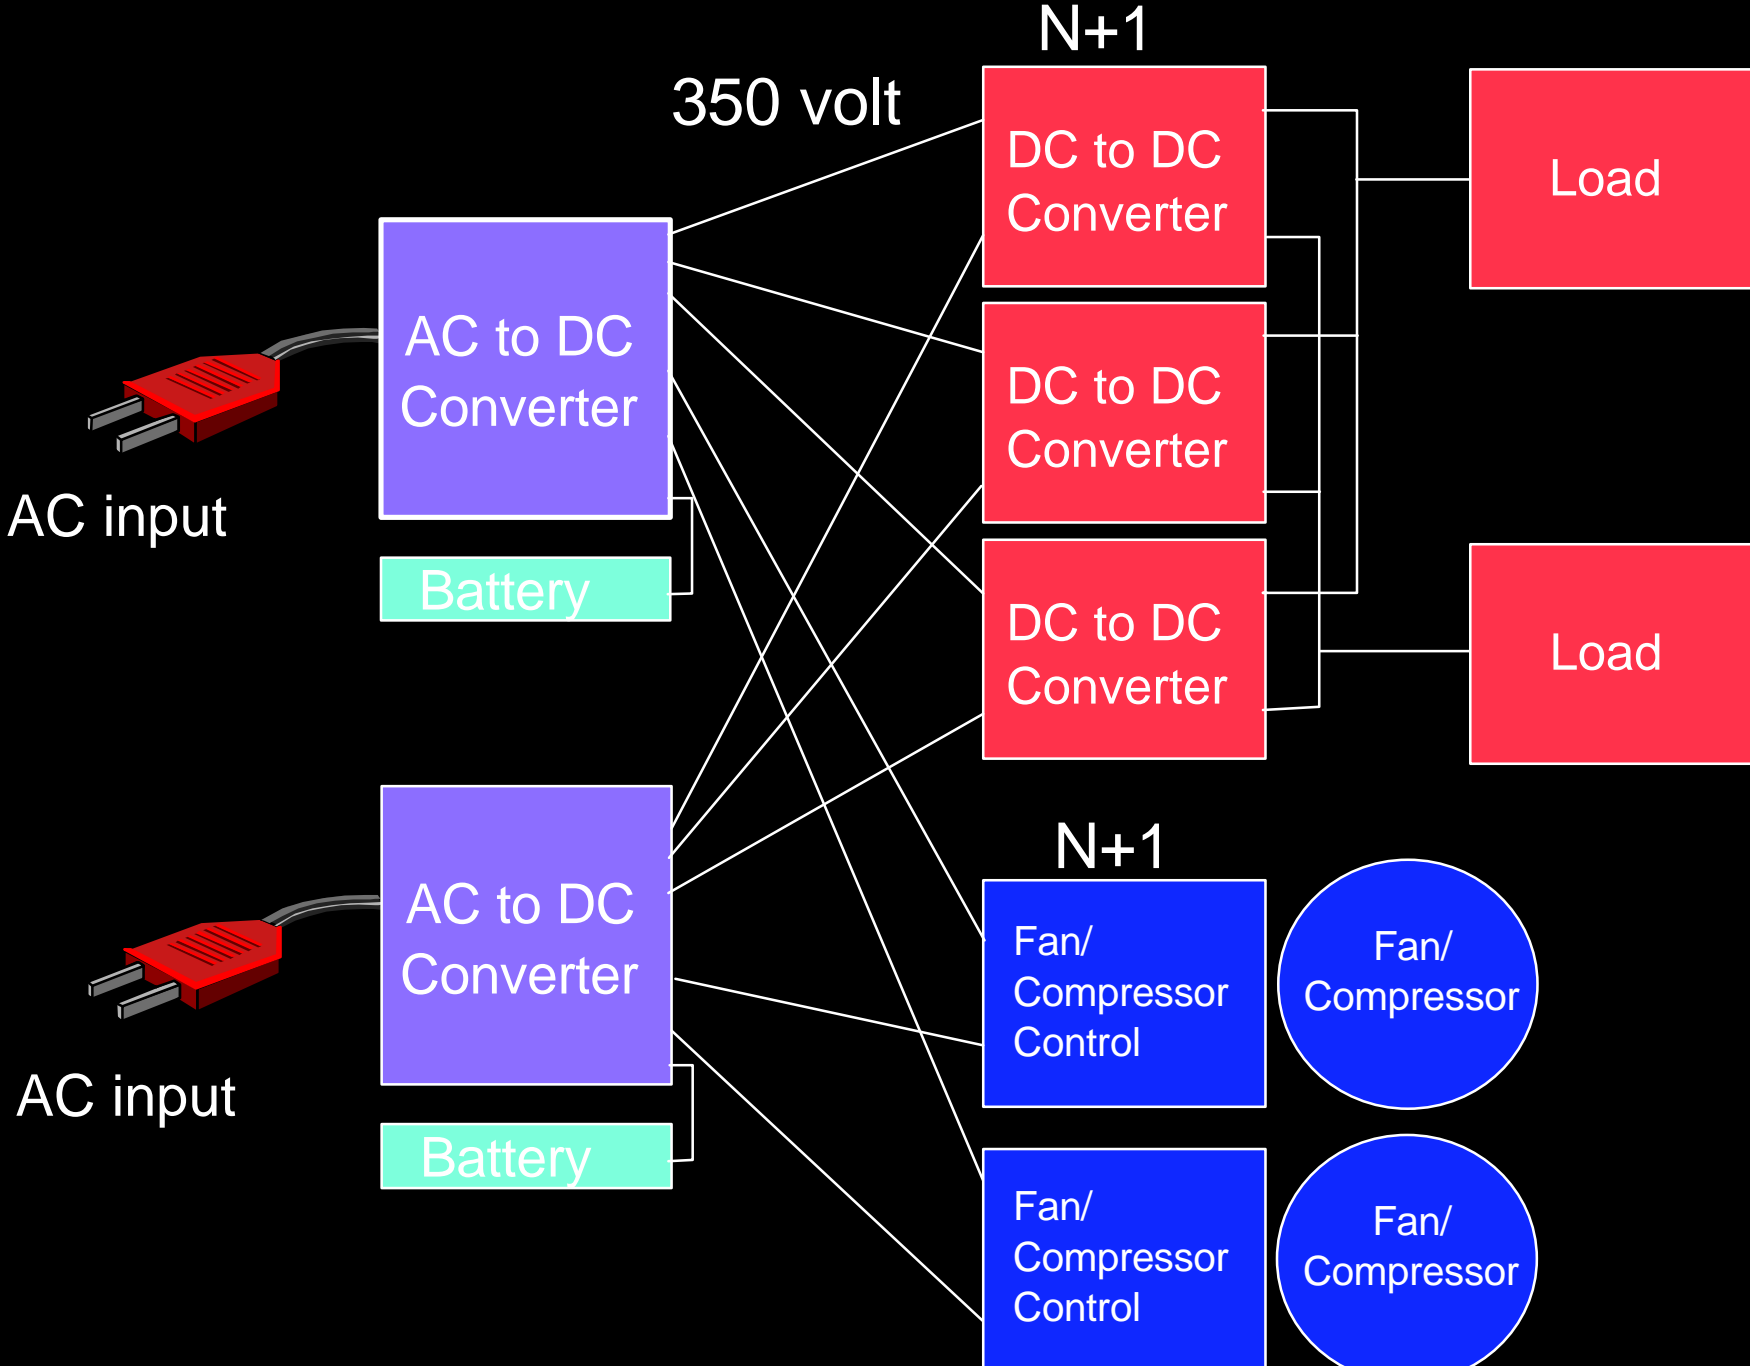
IBM zSeries Fault Tolerant Design

Lisa Spainhower

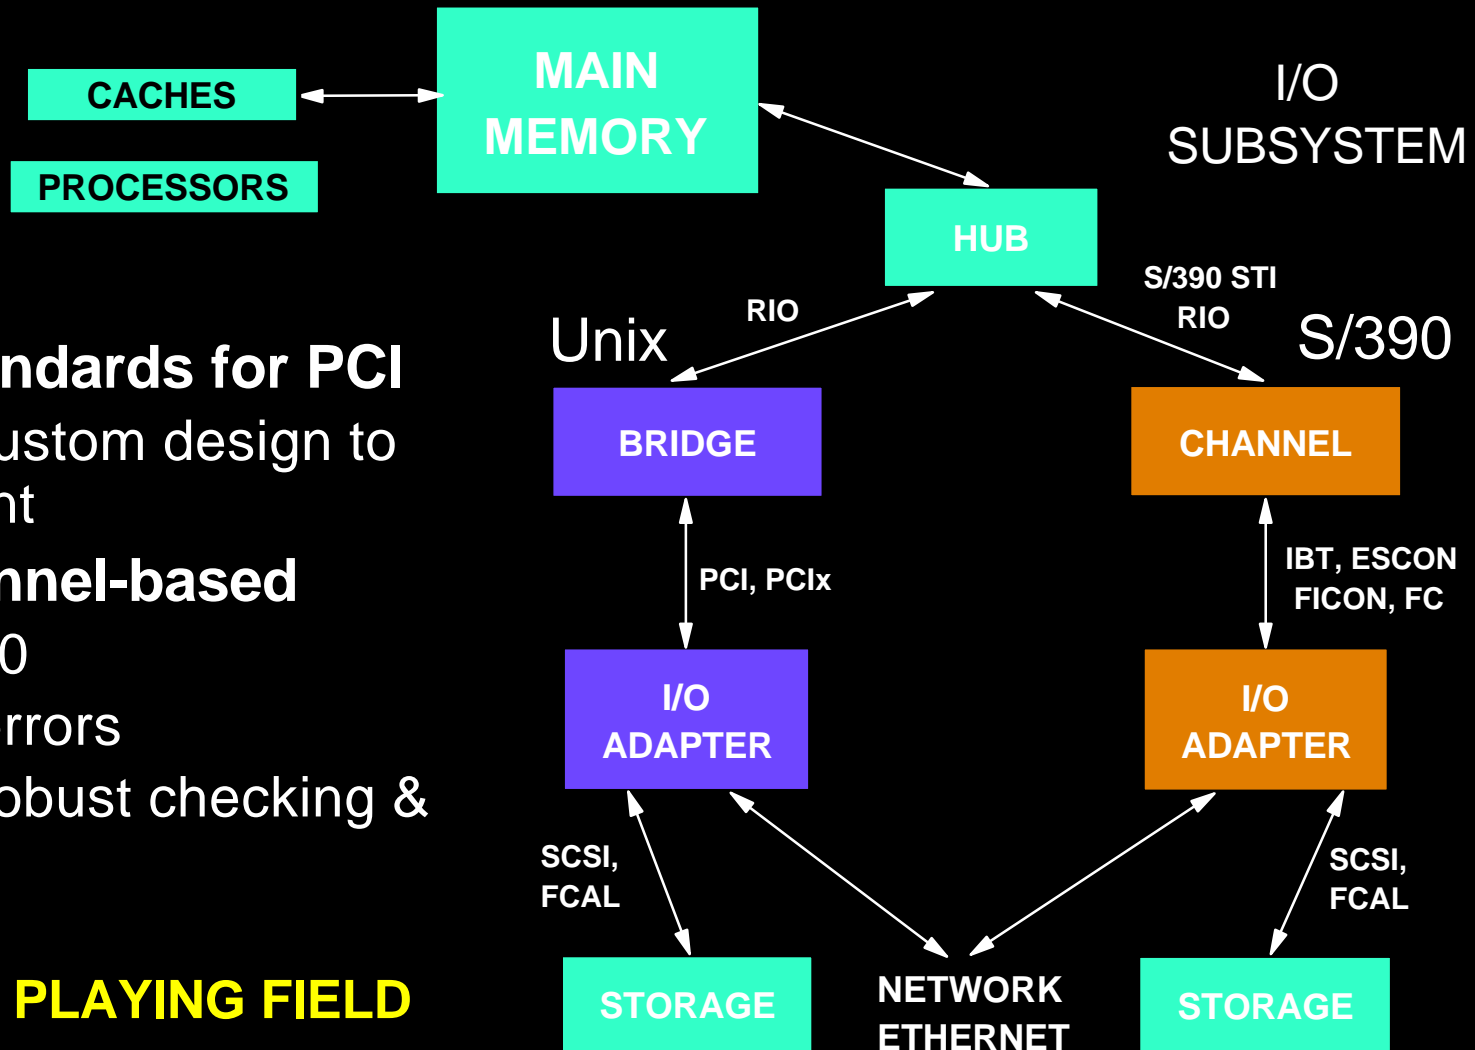
m ^ Technology

September 20, 2001

Power/Cooling Fault Tolerance



I/O ED and Recovery



NO ED standards for PCI

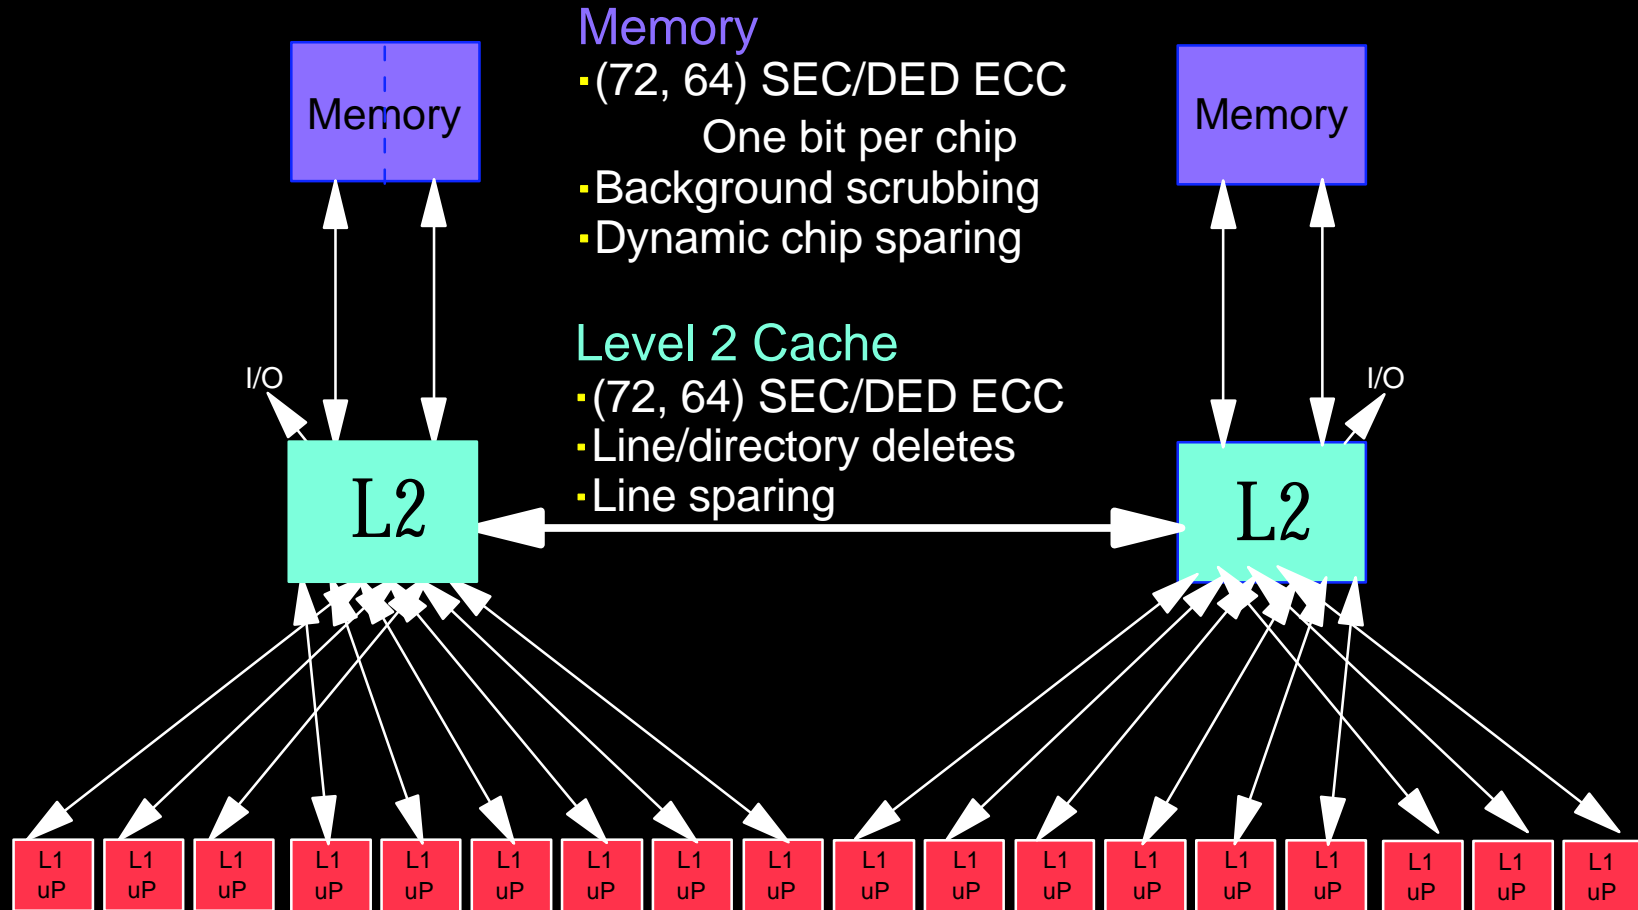
- ▶ RS/AIX custom design to circumvent

IBT is channel-based

- ▶ Like S/390
- ▶ Defined errors
- ▶ Defined robust checking & isolation

LEVEL THE PLAYING FIELD

Memory Hierarchy Fault Tolerance



Memory

- (72, 64) SEC/DED ECC
One bit per chip
- Background scrubbing
- Dynamic chip sparing

Level 2 Cache

- (72, 64) SEC/DED ECC
- Line/directory deletes
- Line sparing

Level 1 Cache

- Parity Protected
- Store-through to L2
- ECC'd Store Buffer on uP
- Line delete/sparing

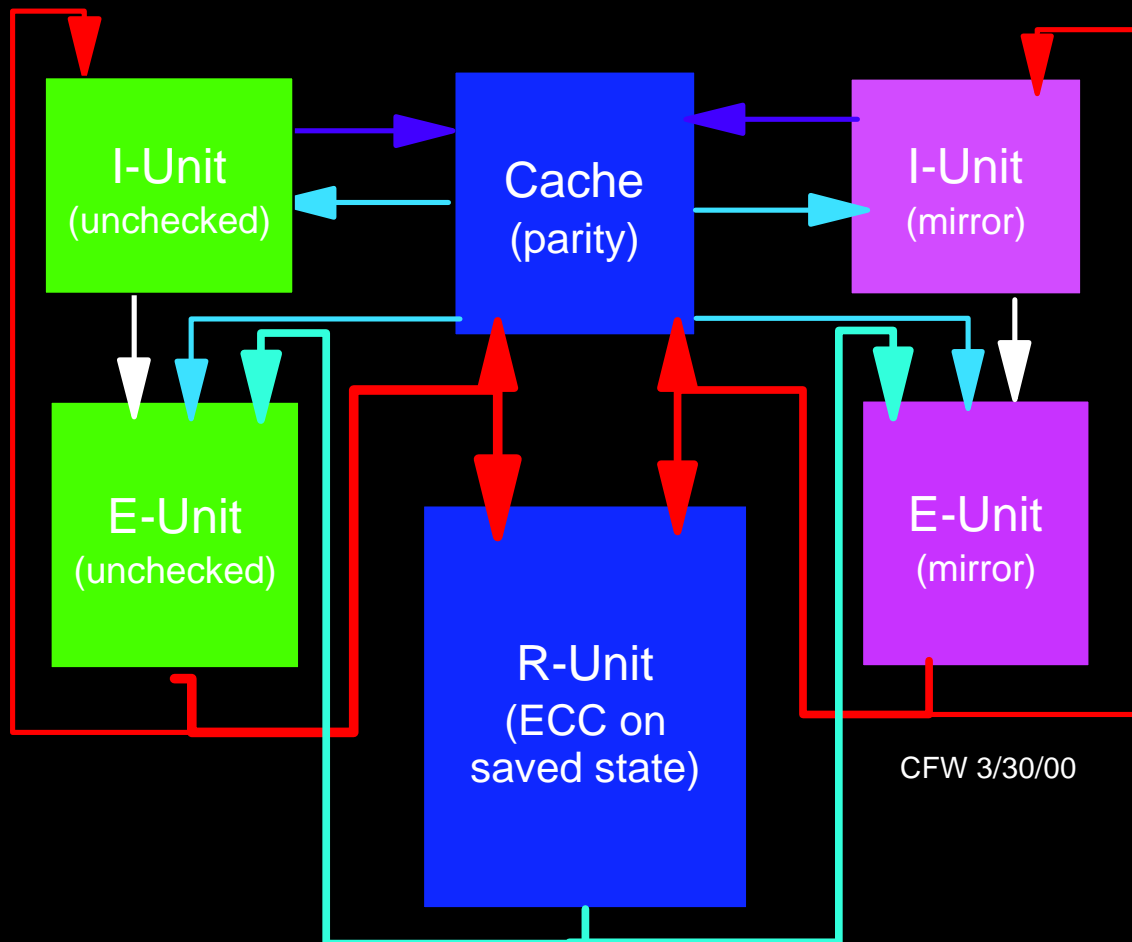
CP Error Detection & Recovery

Duplicated:


- ▶ Complex controls
- ▶ Arithmetic dataflow

Shared:

- ▶ Cache controls
- ▶ Cache data/address flow
- ▶ R-Unit



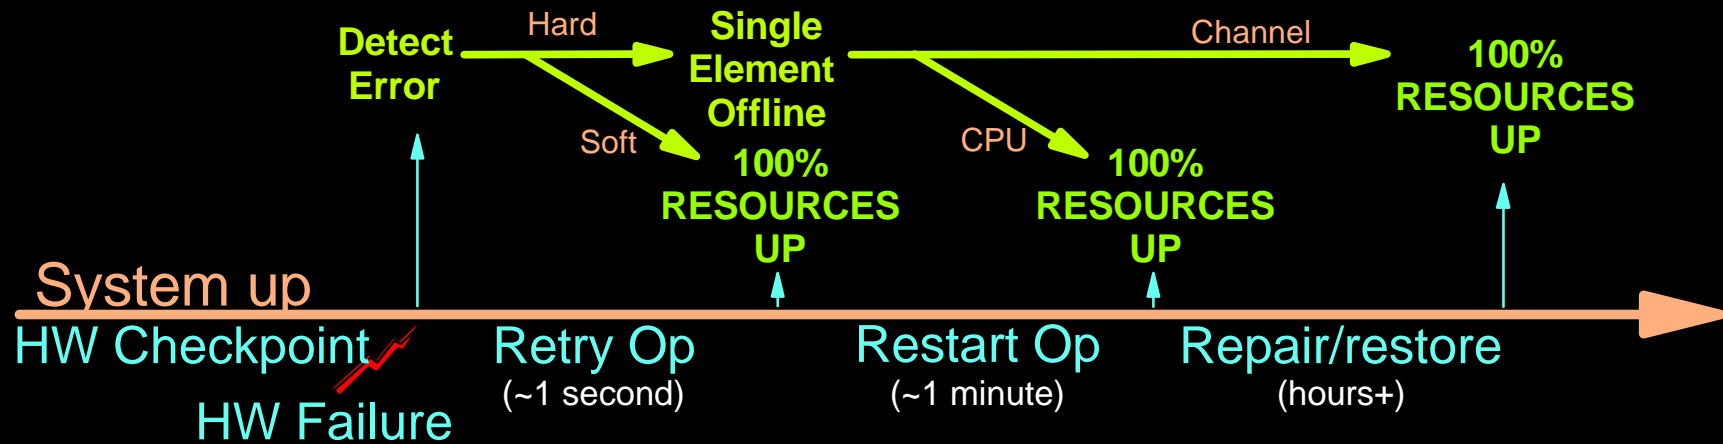
- ▶ Check all state updates
- ▶ Preserve known good state
- ▶ If error
 1. Stop state updates
 2. Refresh from saved state
 3. Restart CPU
- ▶ If error persists
 1. Extract saved state (SE)
 2. Load into spare CPU
 3. Start spare CPU



2Q01 zSeries Full Field Data

- MTTHardware Repair = 8 months
- 81-83% of repairs are concurrent

TYPICAL REPAIR SCENARIO



- 13-15% of repairs are deferrable
- 2-6% of repairs are app loss: MTTAL = 24 years

zSeries Error Reporting

~2 week interval "call home" recovery
data

Suppose CP hard logic (not array) fails caused app loss:

MTTAL from 24 yrs to 11 yrs

Suppose array (L1, L2, BHT) fails also caused app loss:

MTTAL from 11 yrs to 5 yrs

S/390 Evolution

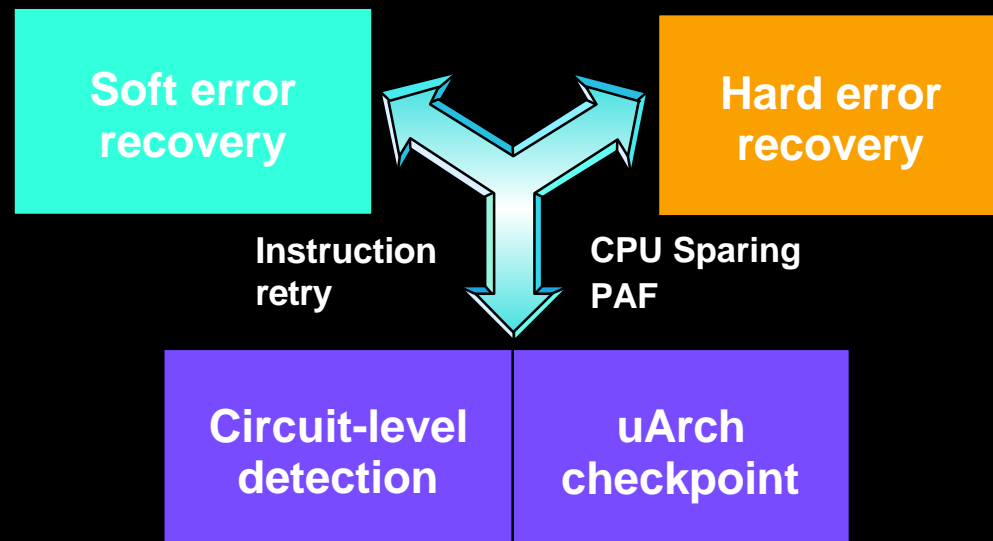
S/390 uses same technology building blocks for soft and hard error recovery

- ▶ Enhanced over past 35 years

IT'S NOT THE ONLY OPTION

- ▶ Beginning afresh, might land elsewhere
- ▶ Need to be driven by current conditions
 - ✓ Technology
 - ✓ Workload

IT'S EFFICIENT & EFFECTIVE FOR S/390



Challenges for the 00s

- Increased importance of firmware
- Circuit failure mechanisms
- State encapsulation
- On-the-fly change
- Dynamic resource allocation
- Configuration validation