IBM Power Systems

Mainline Functional Verification of IBM's POWER7 Processor Core

John Ludden
Senior Technical Staff Member
Hardware Verification
IBM Systems & Technology Group

E-mail: ludden@us.ibm.com

IBM Exec once said: "It's not rocket science"

- **1997 Austin American Statesman Classifieds Ad**

Overview

1. Background: POWER Processor History / Roadmap
2. Verification Methodology
3. Verification Execution
4. Verification Advances
5. Summary

20+ Years of POWER Processors

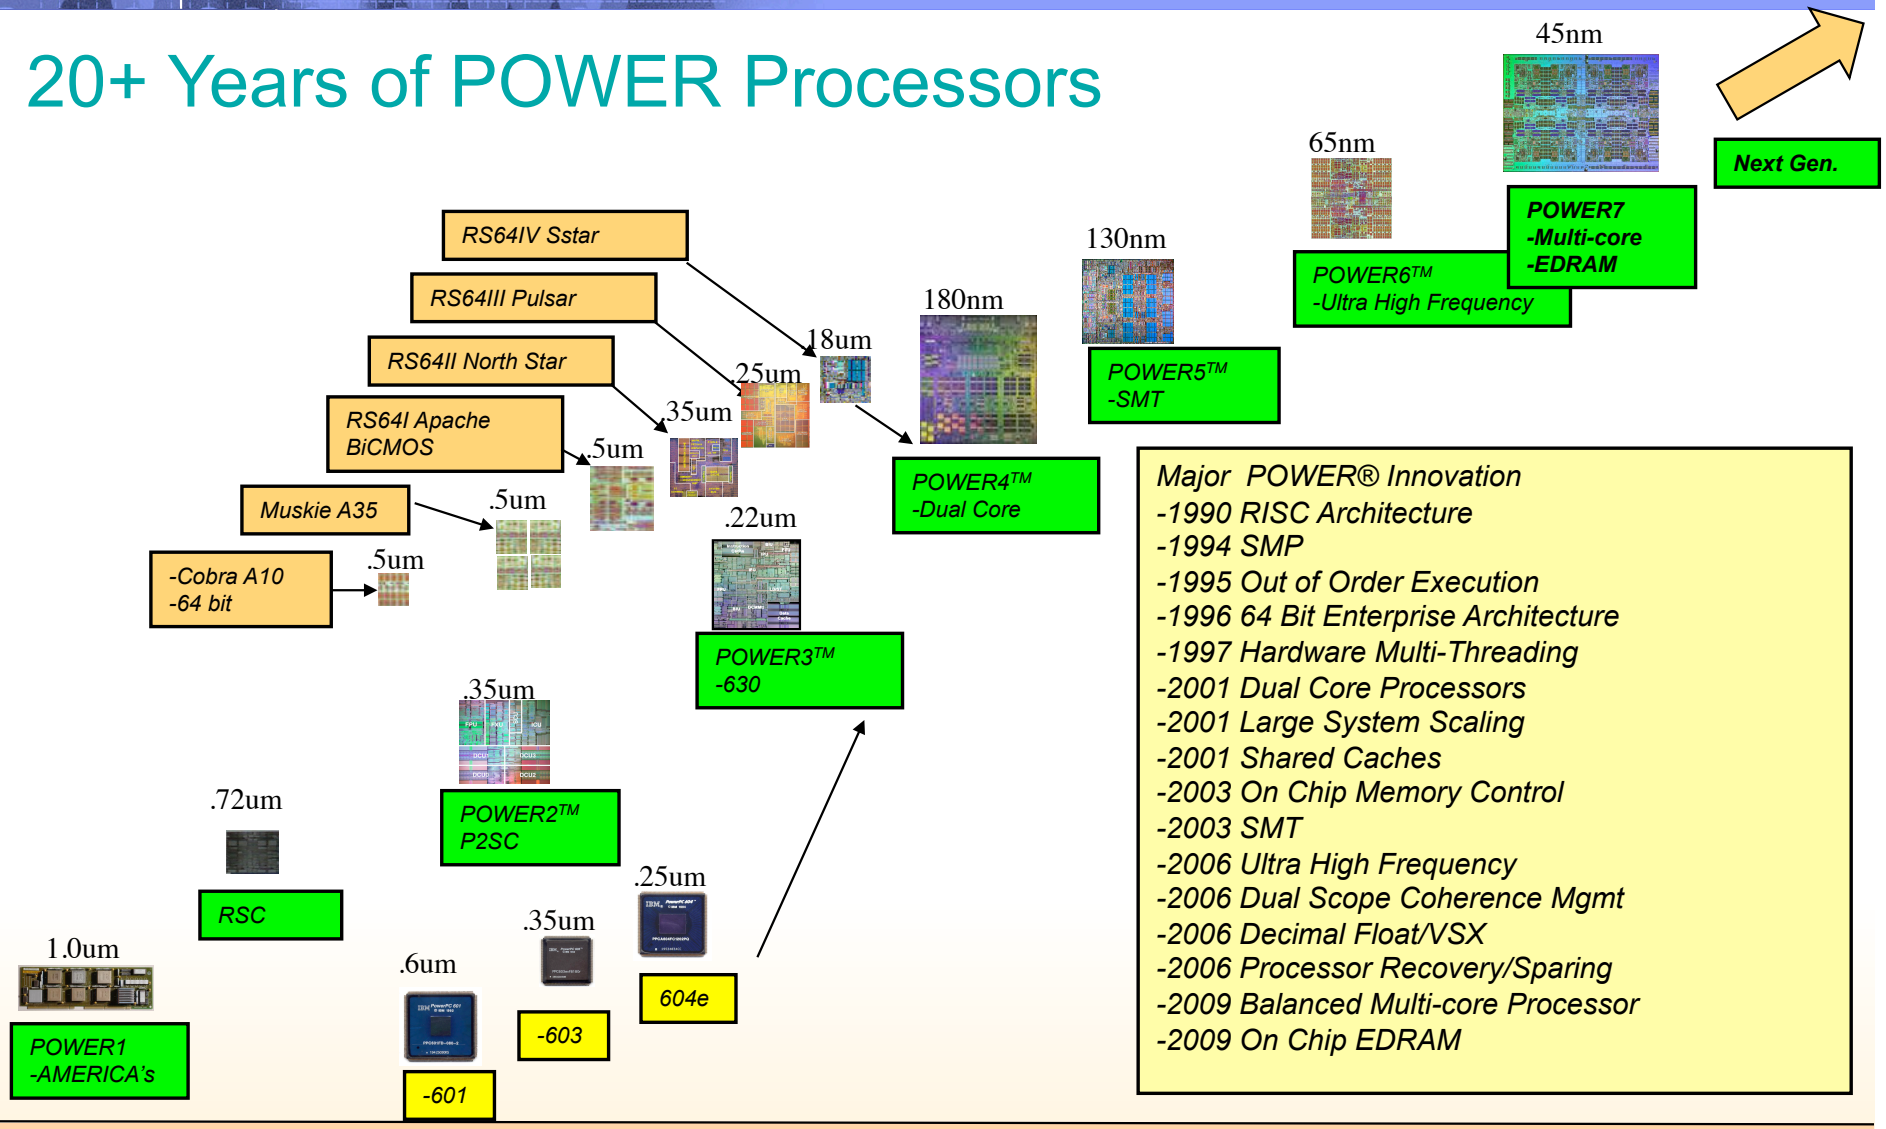

* Dates represent approximate processor power-on dates, not system availability

Processor Technology Roadmap

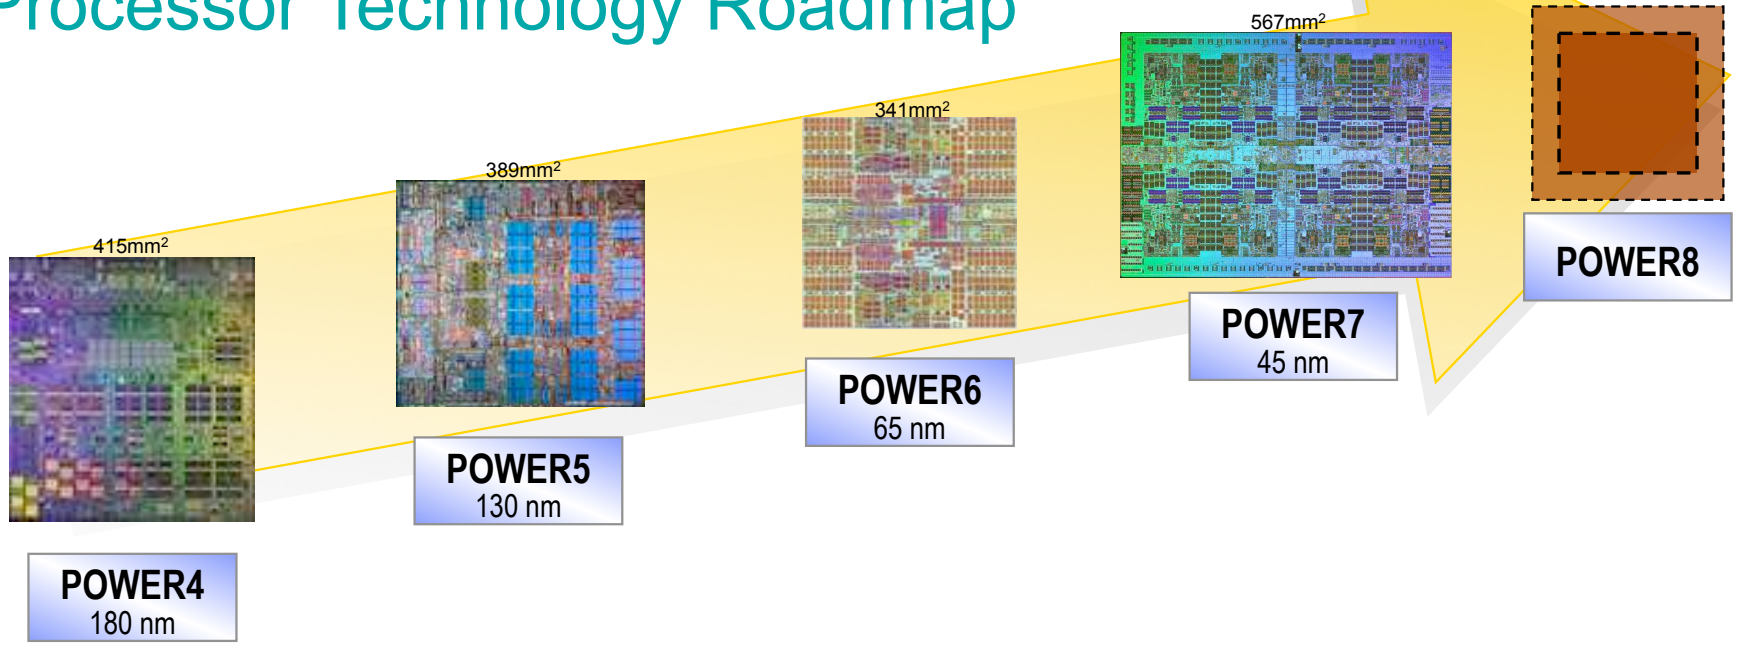

2001	2004	2007	2010	
▪ Dual Core ▪ Chip Multi Processing ▪ Distributed Switch ▪ Shared L2 ▪ Dynamic LPARs (32)	▪ Dual Core ▪ Enhanced Scaling ▪ SMT ▪ Distributed Switch + ▪ Core Parallelism + ▪ FP Performance + ▪ Memory bandwidth + ▪ Virtualization	▪ Dual Core ▪ High Frequencies ▪ Virtualization + ▪ Memory Subsystem + ▪ VMX (Altivec) ▪ Instruction Retry ▪ Dyn Energy Mgmt ▪ SMT + ▪ Protection Keys	▪ Multi Core ▪ On-Chip eDRAM ▪ Power Optimized Cores ▪ Mem Subsystem ++ ▪ SMT++ ▪ Reliability + ▪ VSX & VMX (AltiVec) ▪ Protection Keys+	▪ Concept Phase Exit ▪ On Schedule ▪ Core running in Sim ▪ Continued Leadership

Processor Designs

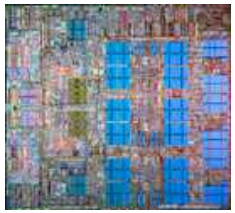
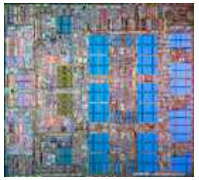
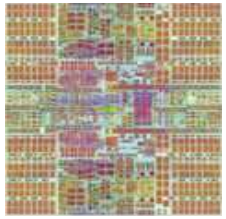
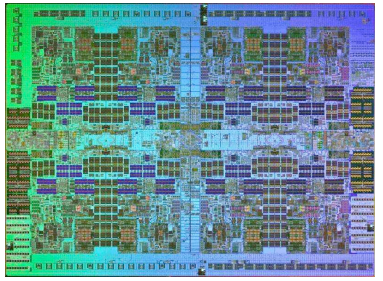

	POWER5	POWER5+	POWER6	POWER7
Technology	**130 nm**	**90 nm**	**60 nm**	**45 nm**
Size	**389 mm^2**	**245 mm^2**	**341 mm^2**	**567 mm^2**
Transistors	**276 M**	**276 M**	**790 M**	**1.2 B**
Cores	**2**	**2**	**2**	**4 / 6 / 8**
Max Threads Per Core (Chip)	**2 (4)**	**2 (4)**	**2 (4)**	**4 (32)**
Frequencies	**1.65 GHz**	**1.9 GHz**	**3-5 GHz**	**3-4 GHz**
L2 Cache	**1.9 MB Shared**	**1.9 MB Shared**	**4 MB / Core**	**256 KB / Core**
L3 Cache	**36 MB**	**36 MB**	**32 MB**	**32 MB**
Memory Cntrl	**1**	**1**	**2 / 1**	**2**

The highest performing 4-socket system on the planet

POWER7 continues to break the rules with more performance

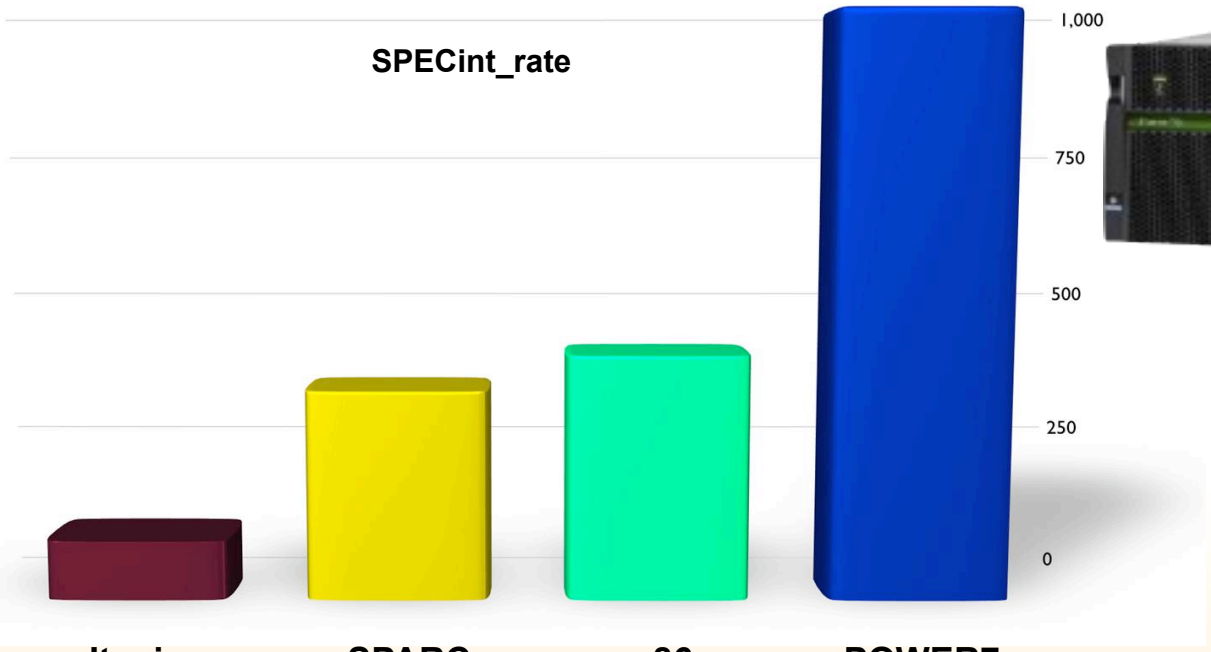

Power 750

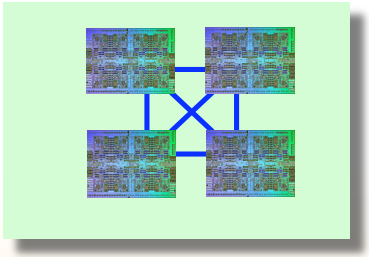

4 POWER7 Chip Interconnect

POWER7 Processor Chip Overview

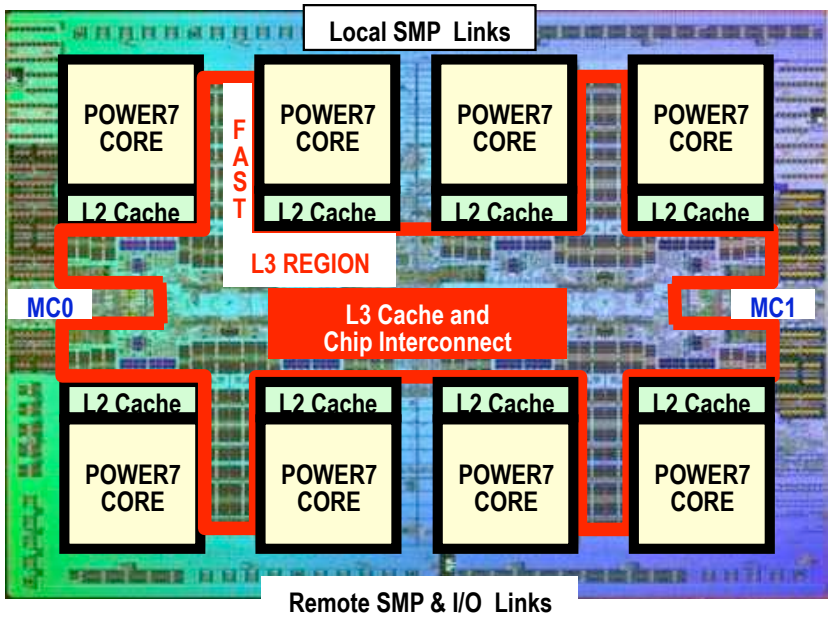

Binary Compatibility with POWER6

- **Cores : 8 (4 / 6 core options)**
- **567mm² Technology:**
  - 45nm lithography, Cu, SOI, eDRAM
- **Transistors: 1.2 B**
  - Equivalent function of 2.7B
  - eDRAM efficiency
- **Eight processor cores**
  - 12 execution units per core
  - 4 Way SMT per core – up to 4 threads per core
  - 32 Threads per chip
  - L1: 32 KB I Cache / 32 KB D Cache
  - L2: 256 KB per core
  - L3: Shared 32MB on chip eDRAM
- **Dual DDR3 Memory Controllers**
  - 100 GB/s Memory bandwidth per chip
- **Scalability up to 32 Sockets**
  - 360 GB/s SMP bandwidth/chip
  - 20,000 coherent operations in flight

POWER7 Core Details

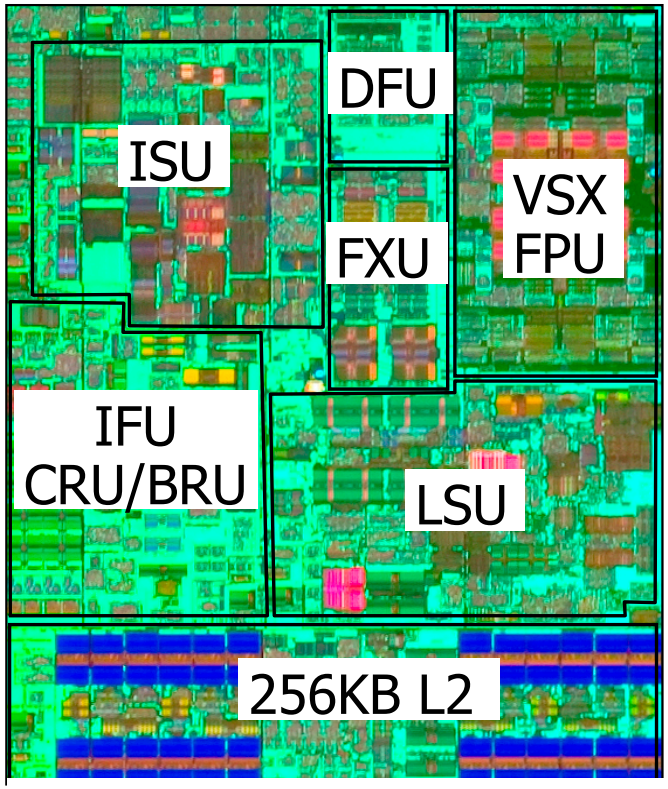

- 64-bit Power ISA Architecture v2.06
- **Out of Order Execution (POWER5-like)**
- **Up to 4 threads per core**
- 12 Execution Units
  - 2 Fixed Point Units
  - 2 Load Store Units also do Simple FX ops
  - 4 Double Precision Floating Point Pipes
  - 1 Branch
  - 1 Condition Register
  - 1 Vector Unit
  - 1 Decimal Floating Point Unit
  - 6 Wide Dispatch (2 branches per group)
  - 8 Wide issue
  - Units include distributed Recovery Function

- POWER7 continues to support VMX / Extends SIMD support with VSX
  - 2 VSX units that can each handle
    - 2 Double-Precision FP calculations
    - 4 Single-Precision FP calculations
  - 64 architected registers

POWER7 Core Block Diagram

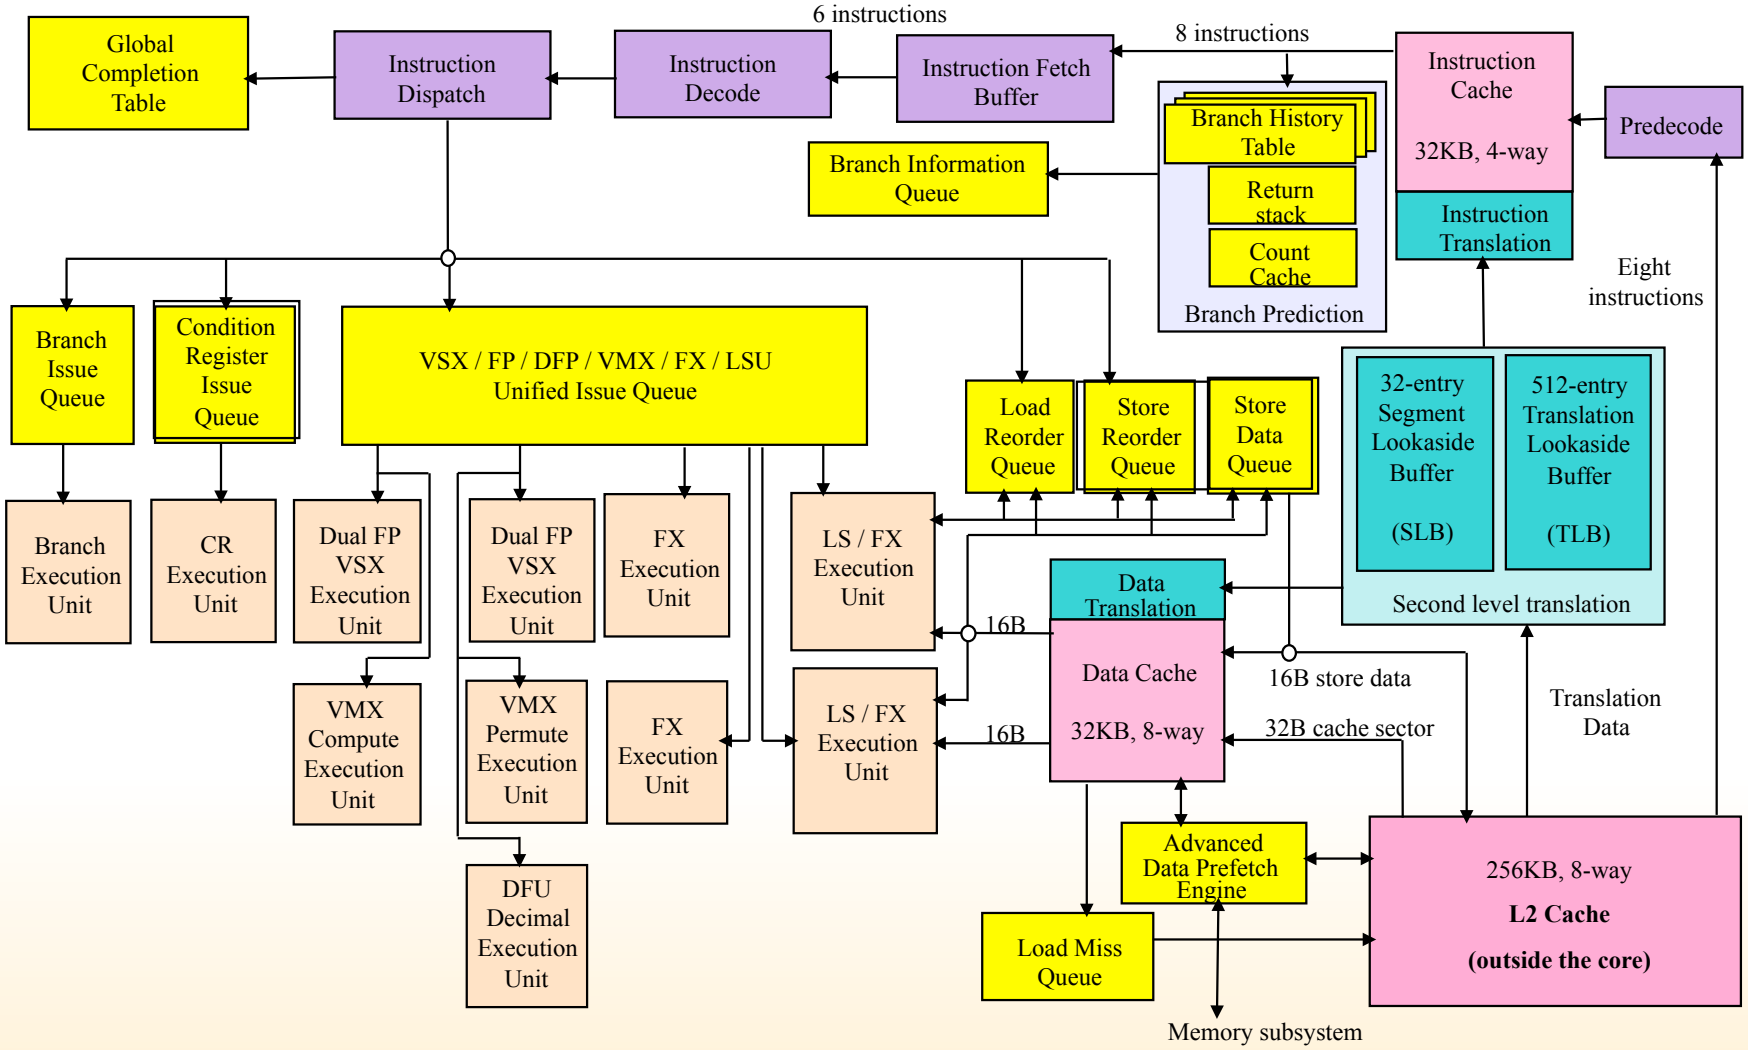

Overview

1. Background: POWER Processor History / Roadmap

2. Verification Methodology

3. Verification Execution

4. Verification Advances

5. Summary

POWER7 RTL verification technology

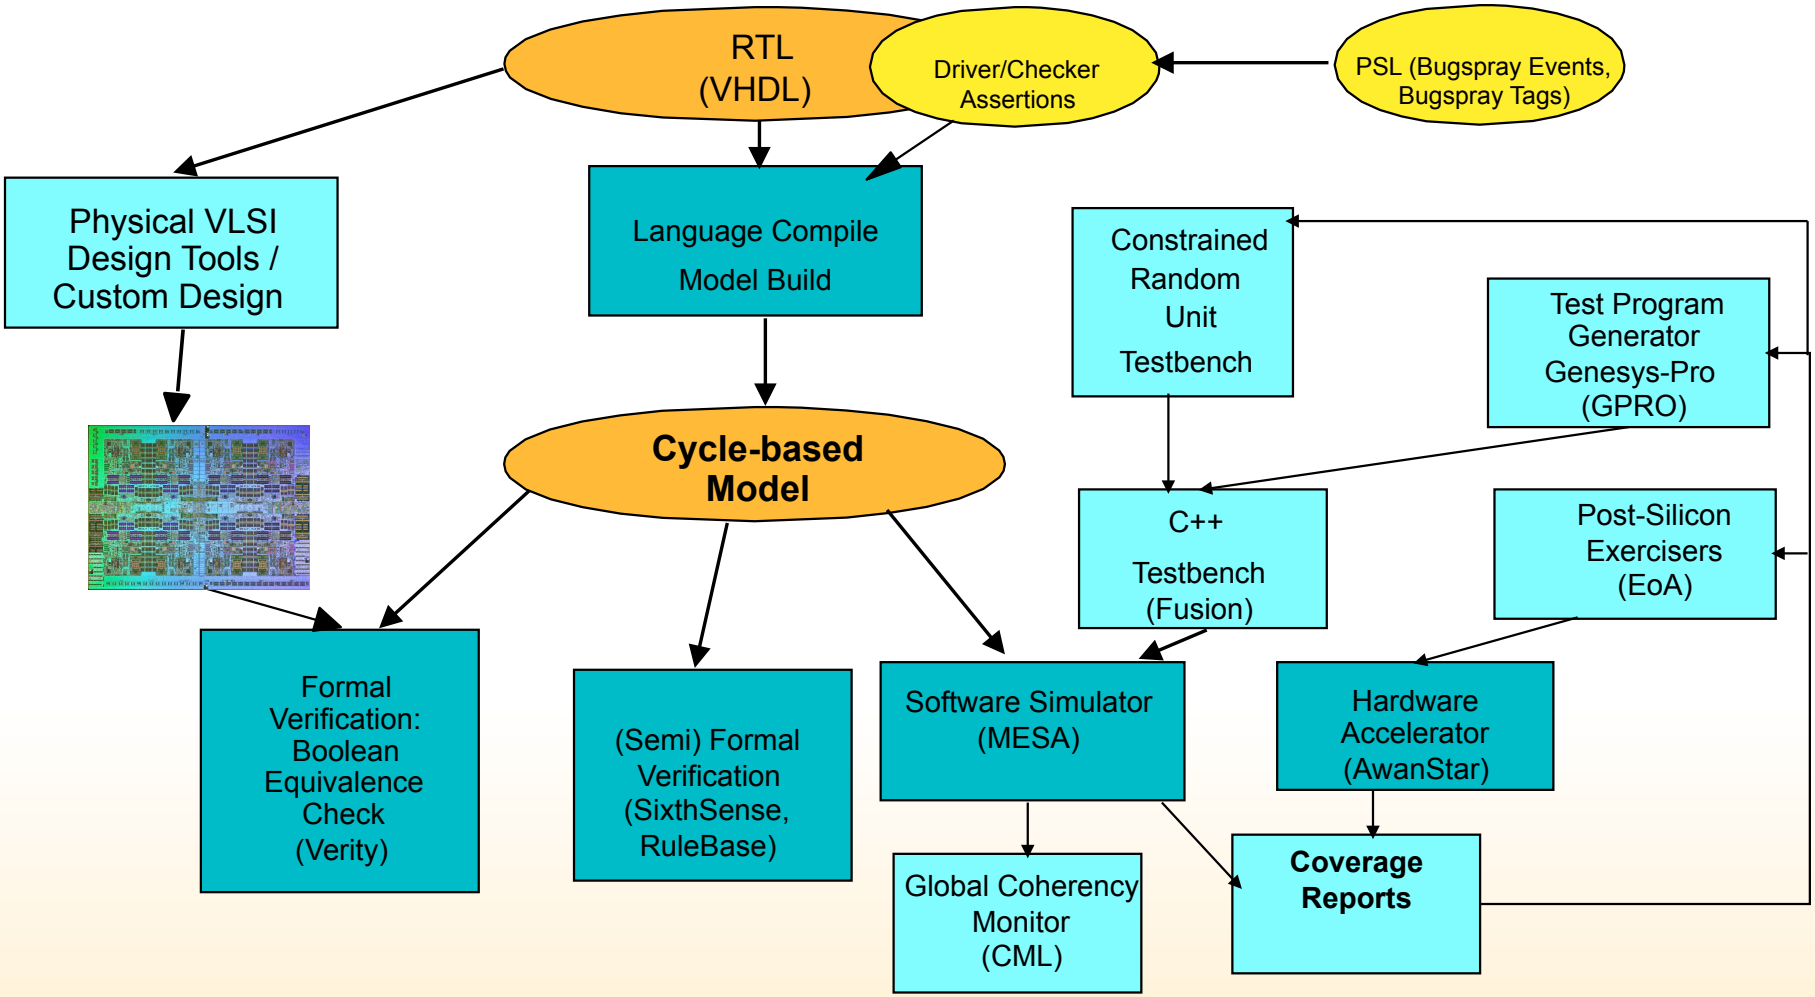

POWER7 Unit Simulation

- **Two Types:**
  - Instruction Level Test case Driven: (Inherited from POWER5)
    - Employs "Core-Common Code" to drive "unit under test"
      - Instruction level Architecture Verification Programs (AVPs) created by Genesys-Pro (GPRO) pseudo-random test generator
      - Instruction-By-Instruction (IBI) checking against AVP results
    - Low level microarchitecture checkers written in FUSION
    - Employed by: IFU, LSU, ISU, FXU, VSU(FPU)
  - Constrained Random Command Driven Environments (Modeled after POWER6)
    - C++ FUSION Drivers and Checkers
    - "On-the-fly" test generation and result checking
    - Employed by: IFU, LSU, ISU

POWER7 Full Core Software Simulation Model

Instruction Level Architecture Verification Program (AVP) Based

- Default test generation via GPRO pseudo-random generator
  - IBI checking against AVP results for up to four independent instruction threads contained within single test
  - Low level microarchitecture FUSION checkers promoted from unit sim environments
  - Global Coherency Checker (CML) looks for architectural storage rule violations
- SMT: True-sharing scenarios, lock testing and storage access ("weak") ordering checked
  - GPRO employed but….
    - IBI checking of these accesses is limited or not possible:
      - › Non-unique or non-deterministic results
      - › CML employed to detect the "right answer" as a post-simulation rule check

Hardware Acceleration Platform for Core Simulation

- **AwanStar: Hardware-accelerated simulation platform**
  - Characteristics
    - 20 times faster than previous generation accelerator
    - 1700 times faster than SW sim, but need less intrusive driving/checking to not slow down hardware box.
    - Enables testing / debug of post-silicon exercisers prior to tape-out
      - Coverage-driven focused exerciser shifts
  - Slower than an FPGA-solution, but
    - FPGA would be a very large platform
    - Full simulation model trace capability for debug (same as SW sim)
    - Ability to leverage coverage (Bugspray) the same as other levels of sim
    - Ability to leverage Bugspray "fail events" (assertions) (same as SW sim)
    - Leveraged same L2 behavioral / driver as utilized for unit and core level sim testing

Synthesizable L2 Cache Behavioral

- **POWER4 thru POWER7 have:**
  - "write-through" L1 Data Cache
    - No modified data
  - independent L1 Instruction Cache
    - No modified data
  - L2 cache:
    - "inclusive" of L1 caches
    - contains modified data
- **POWER4 and POWER5 used C++ developed behavioral**
  - Not portable to accelerator platforms
- **POWER6 and POWER7 developed synthesizable VHDL L2 behavorial**
  - Key component enabling re-use across **all levels** of P7 Core Simulation (including Accelerator platform)

Formal Verification

Formal methods are a vital complement to simulation

- Largest ever application of formal on any IBM project
- Synergistic application as a "mainstream" verification discipline
- POWER7 deep dives early to identify areas ripe for FV
- Reuse of FV assertions by cycle simulation environments
- Sequential Equivalence Checking for late design changes
- Leveraged extensively for lab bring-up bug re-creation
  - Often faster reproduction than simulation based approaches
  - Aids in root cause analysis
  - High-coverage / proof of side-effect-free fixes

Overview

1. Background: POWER Processor History / Roadmap

2. Verification Methodology

3. Verification Execution

4. Verification Advances

5. Summary

POWER7: "We need to change our way of thinking"

- **Four threads per core instead of two**
- **POWER7 Compressed Schedule**
  - Approximately 1 year less than for POWER6 to meet "time to market" objectives
  - "**Shift left**": term used to describe pulling the schedule in
- **POWER6 Lessons Learned**
  - Need to find a better way to interlock verification plan and coverage
- **Need to increase exploitation of Simulation Acceleration Platform**
  - Ability to run same exercisers in pre-silicon that will eventually run on silicon
    - EoA: **Exercisers on Accelerators**
  - Fast platform, but what about coverage?

DTP: "The Glue"

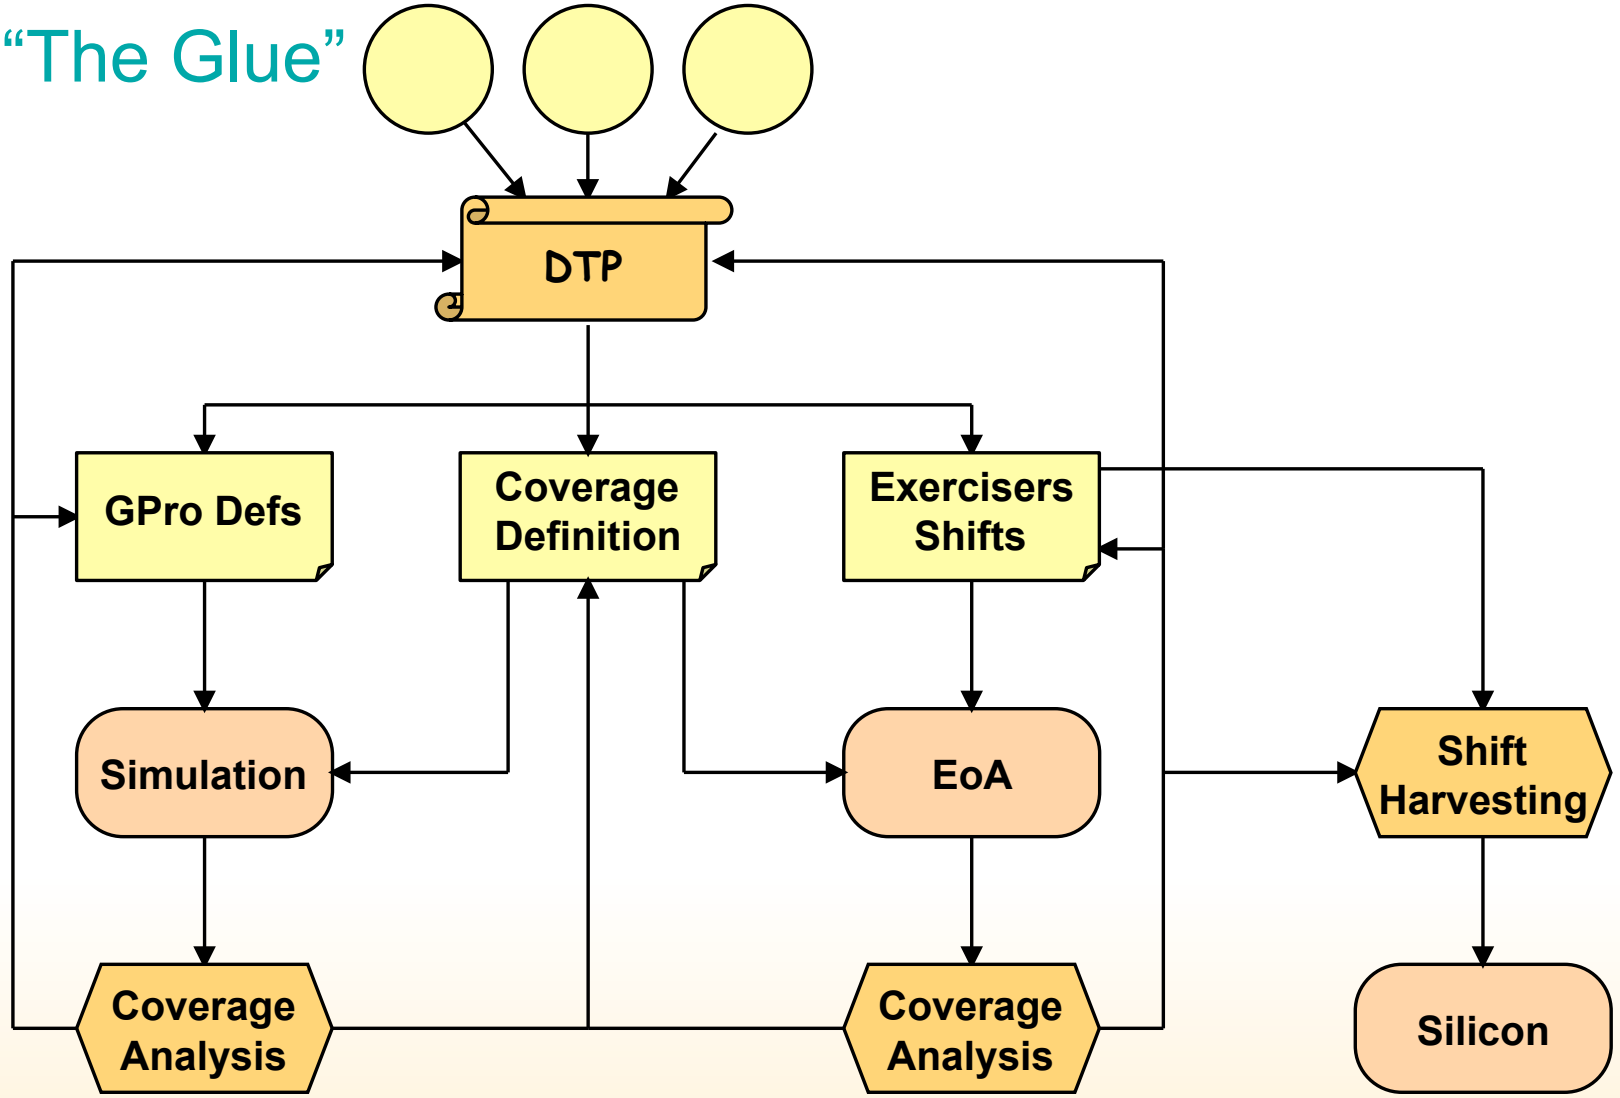

DTP: "The Glue"

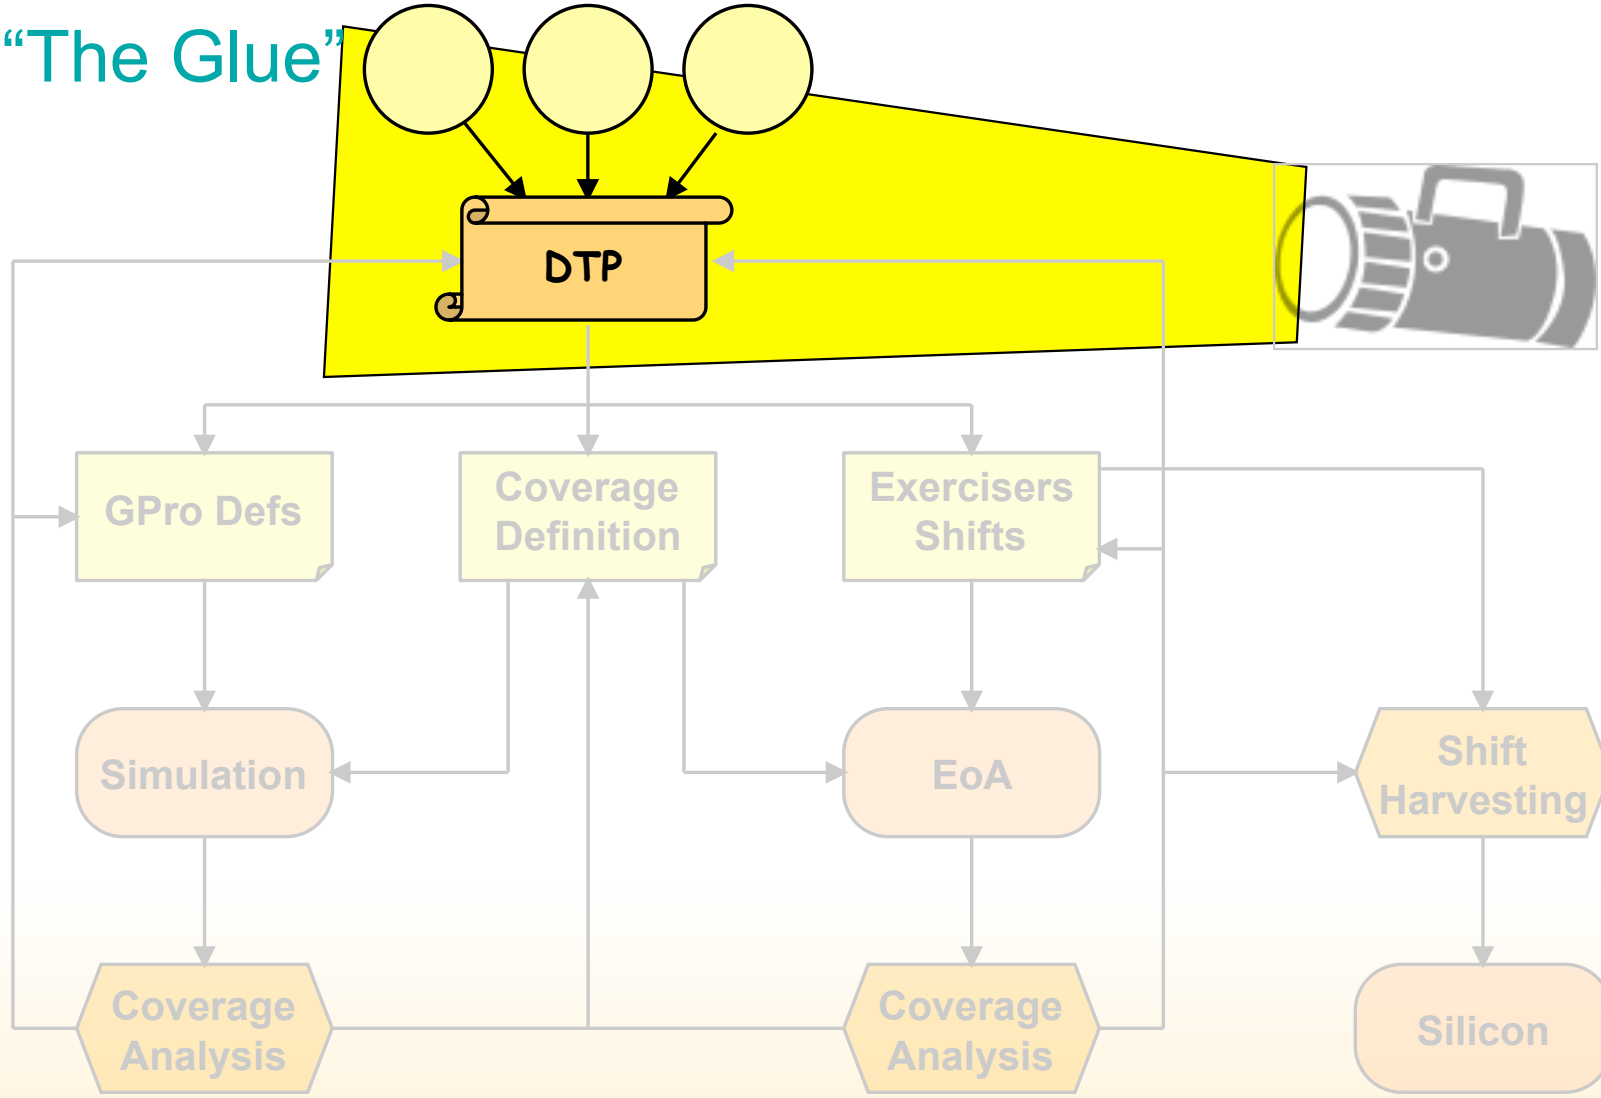

Core Directed Test Plan (DTP)

- **A new way of steering a big ship**
  - A working document (spreadsheet)
  - Not a pdf that gets written and never looked-at again
- **Essentially one all-encompassing test plan in spreadsheet format used to**
  - Hierarchically divide the test plan into major sections and their components
    - Major sections: typically unit or major function boundaries
      - Each gets its own tab
      - Each line item typically requires a unique set of targeted tests (GPRO test templates)
      - One or more line items assigned a coverage tag (Bugspray Tag)
        - Can span unit boundaries
  - Close the loop between the verification plan and the coverage plan for both Core sim and EoA (Exercisers on Accelerators)
  - Measurably gauge test plan completeness for Project Management
- **Created by Verification Lead based on:**
  - Design documentation
  - Microarchitecture walk-thru's with design team

What is a Bugspray Tag?

- **Bugspray Events:**
  - Discrete coverage points we want to see hit in the design which are coded by the designers in VHDL.
  - 3 Types:
    - **Count:** Most commonly used
    - Harvest: Usage is waning
    - Fail: Causes simulation failure if event occurs….even on Accelerator platform
- **Bugspray Tags:**
  - Groups of **Count Events** that are related to covering a particular function in the hardware.
  - These events could exist within a single unit or spans multiple units.
  - Typically corresponds to one or more line items in DTP

Example of Bugspray Tag

- Example:
  - If there is "flush" related logic in both ISU and LSU, those events should get the same tag: "flush". These tags will become line items in the core Directed Test Plan (DTP).
- What if tag applies to multiple line items in DTP?
  - Not a problem
  - This just shows the items are related
  - Counts against you multiple times when coverage is poor. But, coverage improves quickly when you get the correct stimulus on line.
- Litmus test:
  - **Q:** Do I need a tag for something?
  - **A: Rule of thumb:** If I need to write a unique set of tests (drivers/stimuli) to target it, then give it a new tag.

DTP: "The Glue"

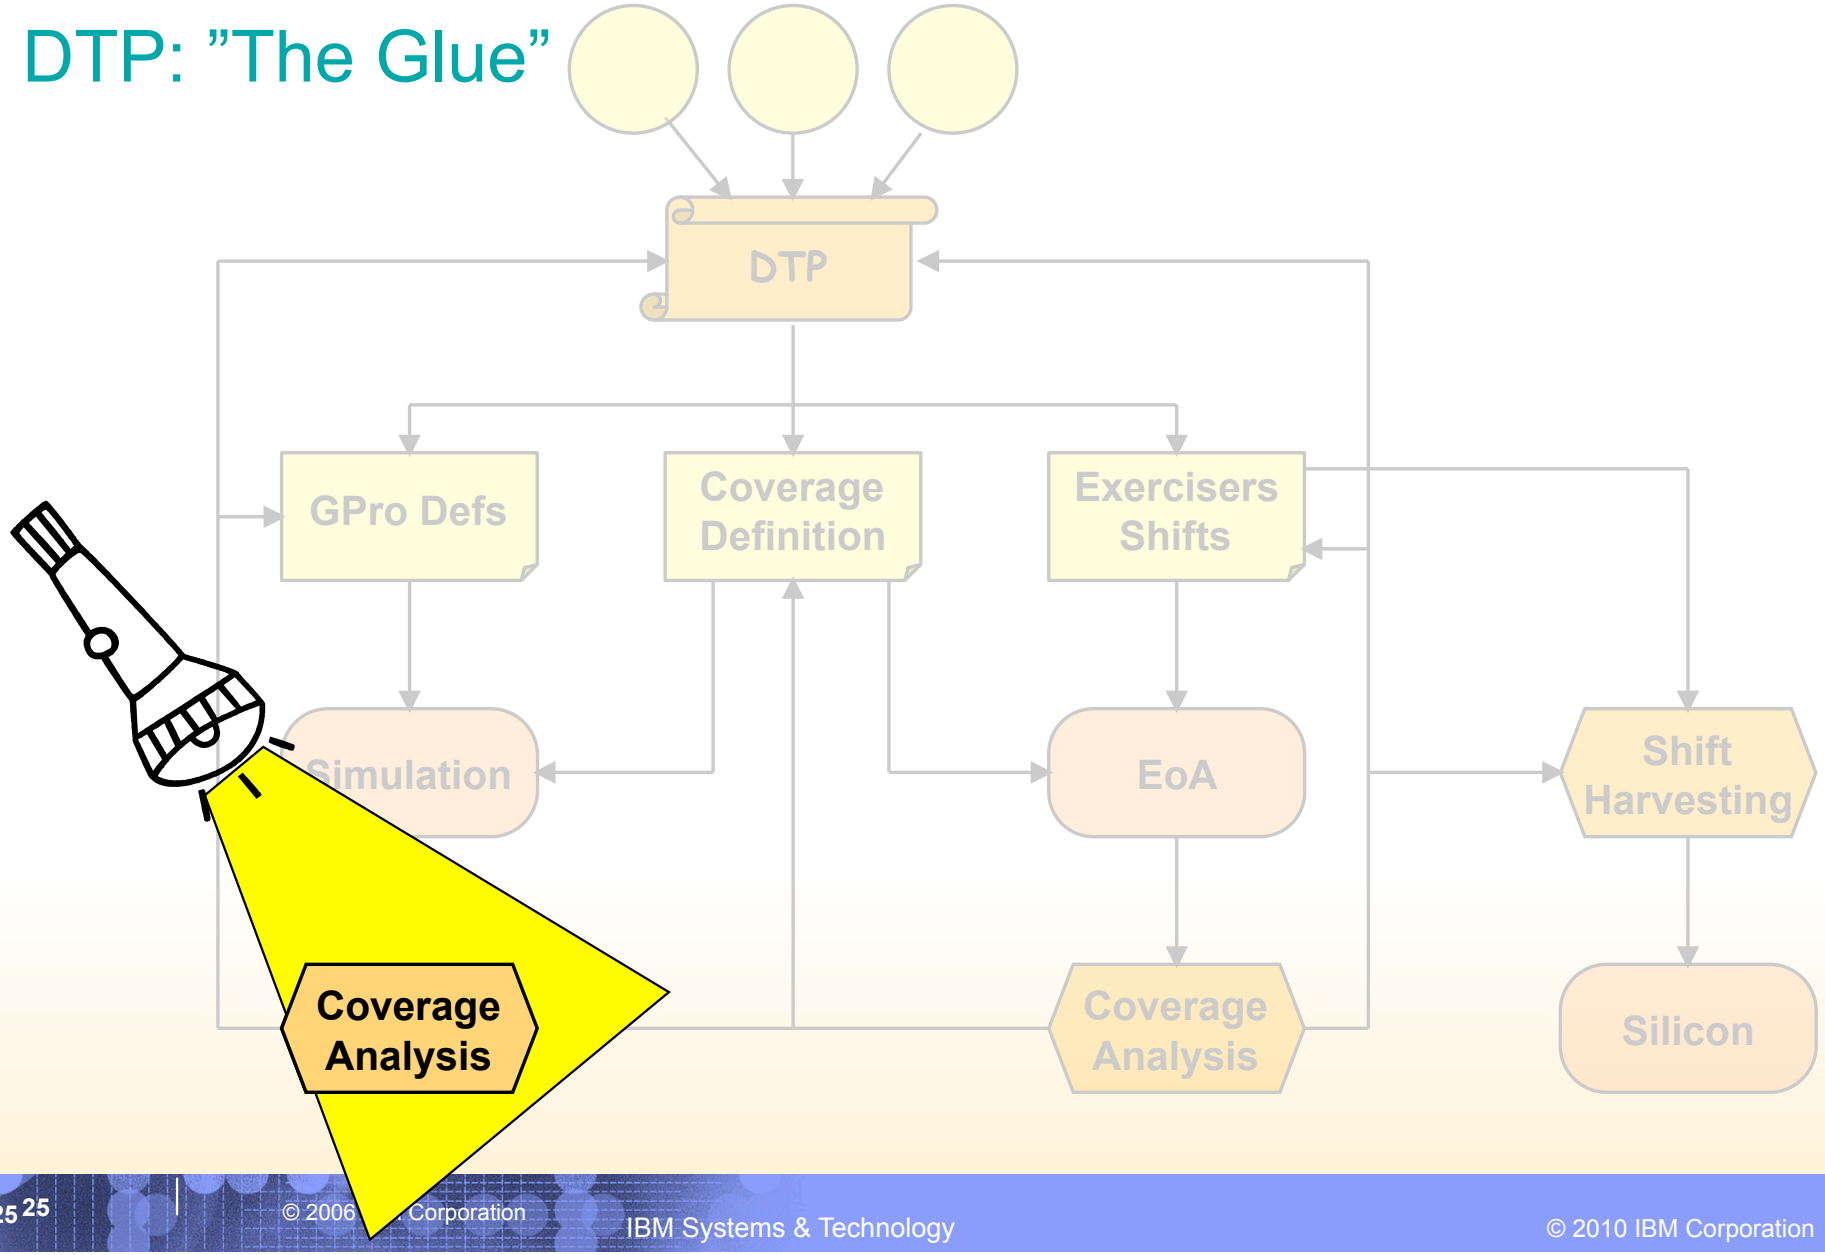

Coverage Analysis

- **The Bugspray Tags allowed a "big picture" view**
  - Can see the "forest through the trees"
- **Coverage Reports presented data in terms of**
  - Major sections of the DTP
  - Bugspray tags
  - Side-by-side comparisons of coverage by platform (unit, core, EoA, system)
    - Allowed quick analysis of unreachable events if un-hit on all platforms
      - Bugspray coding errors
      - Microarchitecture Restrictions Prevent Occurrance
    - Identified platform dependent tool / driver deficiencies
    - Enabled priority calls to be made when schedule was tight
      - Event or Tag has good coverage on multiple platforms vs. poor coverage on all platforms
        - › Risk mitigation

Category Report

LSU

	Tags	DarkRed	Red	Yellow	Green	DarkGreen	Above 90%	At 100%
Unit Sim	24	0	11	8	1	4	15	7
Core Sim	24	0	6	6	3	9	20	15
EoA	24	0	9	4	1	10	15	11
Total	24	0	3	5	5	11	22	18

	Unit Sim					Core Sim					EoA					Total				
Tag	**Status**	**Total**	**No Hit**	**Light Hit**	**Hit**	**Status**	**Total**	**No Hit**	**Light Hit**	**Hit**	**Status**	**Total**	**No Hit**	**Light Hit**	**Hit**	**Status**	**Total**	**No Hit**	**Light Hit**	**Hit**
autogen	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)
cacheops	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)
ciloads	Red	12	2 (16.67%)	3 (25.00%)	7 (58.33%)	DarkGreen	12	0 (0.00%)	0 (0.00%)	12 (100.00%)	DarkGreen	12	0 (0.00%)	0 (0.00%)	12 (100.00%)	DarkGreen	12	0 (0.00%)	0 (0.00%)	12 (100.00%)
dcache	Red	110	22 (20.00%)	21 (19.09%)	67 (60.91%)	Red	110	12 (10.91%)	14 (12.73%)	84 (76.36%)	Red	110	18 (16.36%)	2 (1.82%)	90 (81.82%)	Yellow	110	10 (9.09%)	4 (3.64%)	96 (87.27%)
derat	Yellow	41	3 (7.32%)	4 (9.76%)	34 (82.93%)	Green	41	0 (0.00%)	1 (2.44%)	40 (97.56%)	Red	41	23 (56.10%)	0 (0.00%)	18 (43.90%)	DarkGreen	41	0 (0.00%)	0 (0.00%)	41 (100.00%)
flush	Red	335	44 (13.13%)	69 (20.60%)	222 (66.27%)	Red	335	34 (10.15%)	42 (12.54%)	259 (77.31%)	Red	335	46 (13.73%)	13 (3.88%)	276 (82.39%)	Yellow	335	20 (5.97%)	21 (6.27%)	294 (87.76%)
fxulsu	DarkGreen	24	0 (0.00%)	0 (0.00%)	24 (100.00%)	DarkGreen	24	0 (0.00%)	0 (0.00%)	24 (100.00%)	DarkGreen	24	0 (0.00%)	0 (0.00%)	24 (100.00%)	DarkGreen	24	0 (0.00%)	0 (0.00%)	24 (100.00%)
icswx	Red	13	1 (7.69%)	12 (92.31%)	0 (0.00%)	Yellow	13	0 (0.00%)	5 (38.46%)	8 (61.54%)	Red	13	13 (100.00%)	0 (0.00%)	0 (0.00%)	Yellow	13	0 (0.00%)	5 (38.46%)	8 (61.54%)
larxstcx	Green	14	0 (0.00%)	2 (14.29%)	12 (85.71%)	DarkGreen	14	0 (0.00%)	0 (0.00%)	14 (100.00%)	DarkGreen	14	0 (0.00%)	0 (0.00%)	14 (100.00%)	DarkGreen	14	0 (0.00%)	0 (0.00%)	14 (100.00%)
lhs	Yellow	13	0 (0.00%)	4 (30.77%)	9 (69.23%)	DarkGreen	13	0 (0.00%)	0 (0.00%)	13 (100.00%)	DarkGreen	13	0 (0.00%)	0 (0.00%)	13 (100.00%)	DarkGreen	13	0 (0.00%)	0 (0.00%)	13 (100.00%)
lmq	Red	121	25 (20.66%)	21 (17.36%)	75 (61.98%)	Red	121	21 (17.36%)	6 (4.96%)	94 (77.69%)	Red	121	22 (18.18%)	0 (0.00%)	99 (81.82%)	Red	121	16 (13.22%)	4 (3.31%)	101 (83.47%)
lqstq	Red	7	2 (28.57%)	1 (14.29%)	4 (57.14%)	Red	7	0 (0.00%)	4 (57.14%)	3 (42.86%)	DarkGreen	7	0 (0.00%)	0 (0.00%)	7 (100.00%)	DarkGreen	7	0 (0.00%)	0 (0.00%)	7 (100.00%)

Bigger Category Report

LSU

	Tags	DarkRed	Red	Yellow	Green	DarkGreen	Above 90%	At 100%
Unit Sim	24	0	11	8	1	4	15	7
Core Sim	24	0	5	7	3	9	21	15
EoA	24	0	9	4	1	10	15	11
Total	24	0	3	5	5	11	22	18

	Unit Sim					Core Sim				
Tag	Status	Total	No Hit	Light Hit	Hit	Status	Total	No Hit	Light Hit	Hit
autogen	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)	DarkGreen	2	0 (0.00%)	0 (0.00%)	2 (100.00%)
cacheops	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)	DarkGreen	4	0 (0.00%)	0 (0.00%)	4 (100.00%)
ciloads	Red	12	2 (16.67%)	3 (25.00%)	7 (58.33%)	DarkGreen	12	0 (0.00%)	0 (0.00%)	12 (100.00%)
dcache	Red	110	22 (20.00%)	21 (19.09%)	67 (60.91%)	Red	110	12 (10.91%)	6 (5.45%)	92 (83.64%)

Tag Report

dcache

	Total	No hit	Light hit	Hit
Unit Sim	110	22	21	67
Core Sim	110	12	14	84
EoA	110	18	2	90
Total	110	10	4	96

Tags

- No events defined
- > 10% no hit events
- Some no hit events
- All events Hit once
- All events Hit 100 times

Events

- Aged out
- Not hit
- Hit less than 100 times
- Hit at least 100 times

Unit	Entity	Class	Var	LSU Sim	Core Sim	EoA	Total
lsu	ls	ls_load	e1	64850	1373	0	66223
lsu	ls	ls_load	e2	0	0	0	0
lsu	ls	ls_load	e3	4127	18991	45207320	45230438
lsu	ls	ls_load	e4	0	0	0	0
lsu	ls	ls_load	e5	35	124	0	159
lsu	ls	ls_load	e6	15	46	0	61
lsu	ls	ls_load	e7	2587	Waived	22906016	22928167
lsu	ls	ls_load	e8	4	141	N/A	1111
lsu	ls	ls_load	e9	11485	973	0	12458

DTP: "The Glue"

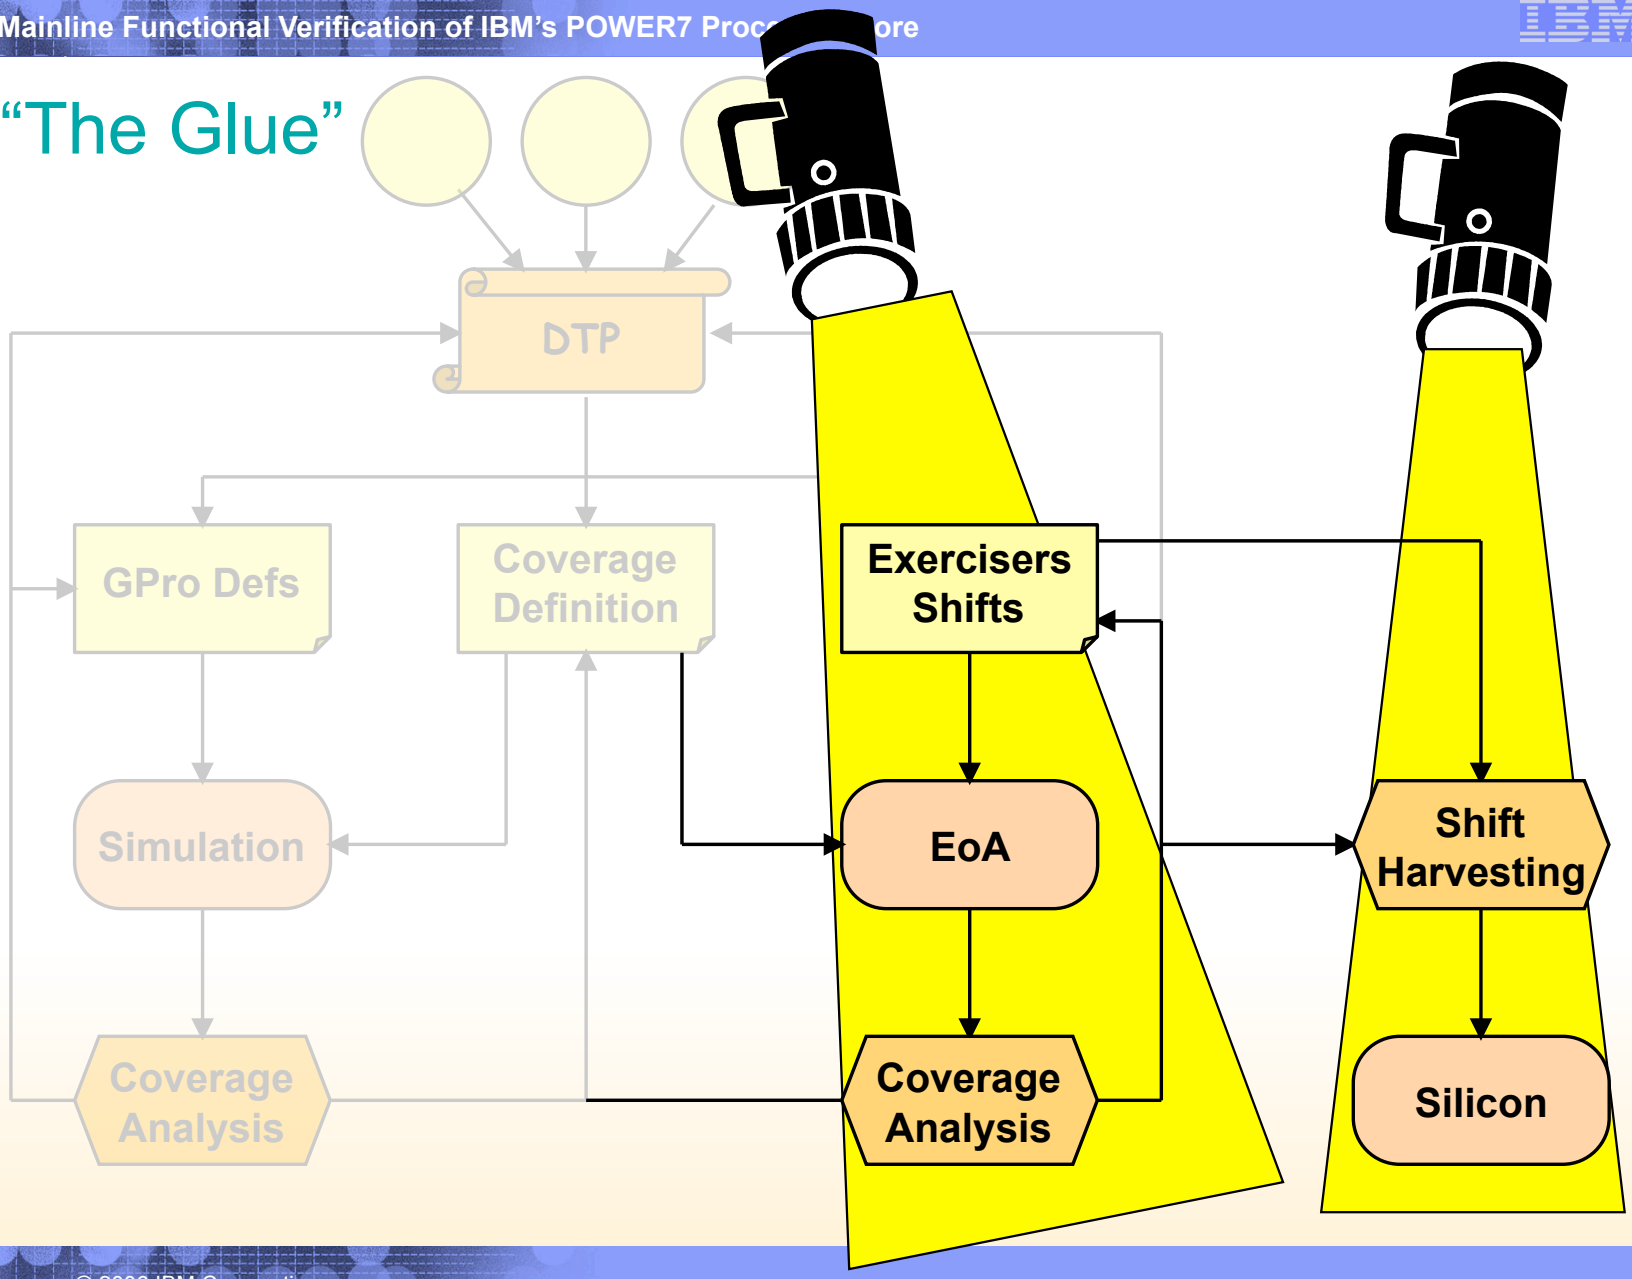

Exerciser Shift Development Workflow

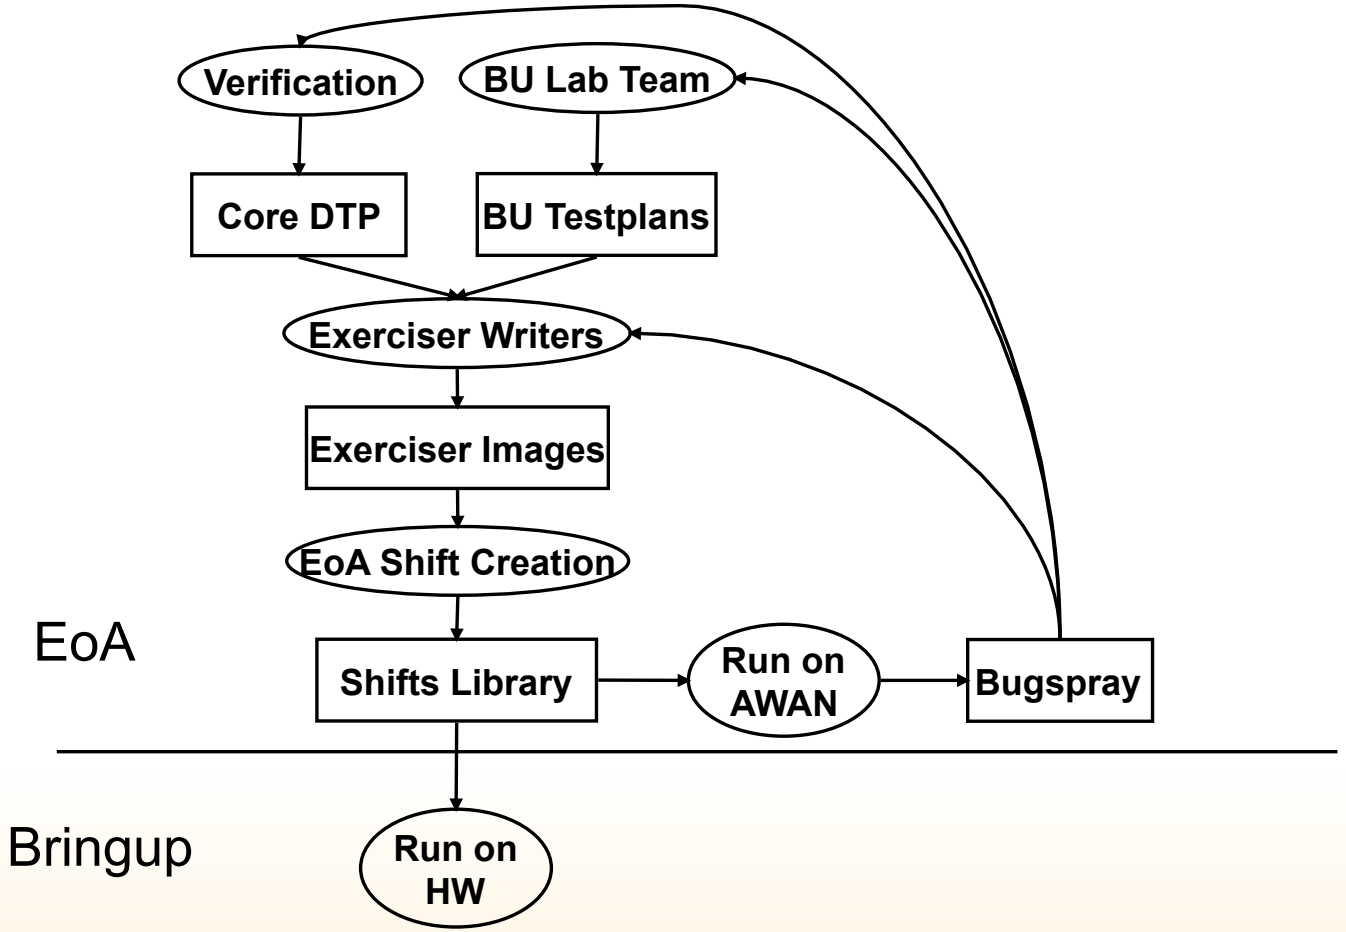

How did EoA use coverage reports?

- **Measurement & Feedback**
  - Indication of what is not being hit
  - Coverage visibility to the exerciser writers
  - Drives exerciser improvement
- **Proof of the exercisers to verification & design**
  - Previously, only make statements about what the exercisers were testing
  - Now, well conceived coverage events provide proof
- **Tagging and Organization of Coverage Events Essential**
  - Gives exerciser writers direction
  - Eliminates first pass interaction with designers as to which events are important.
- ***Allowed Development of High Quality Exerciser Shifts for Post-Silicon Validation Effort***
  - Coverage proven directed exerciser shifts
  - Ready to quickly exploit high throughput post-silicon offers

Overview

1. Background: POWER Processor History / Roadmap

2. Verification Methodology

3. Verification Execution

4. Verification Advances

5. Summary

Traditional Approach to SMT Verification

The Traditional Approach:

- "symmetric" instruction streams on each thread
  - A particular resource targeted (e.g., GPR rename registers)
    - 100 random load/store instructions on each thread

- Hindsight: This methodology had weaknesses….
  - No guarantee instructions streams from all threads run same number of cycles
  - Not the best way to achieve desired cross thread interaction
  - Not effective at targeting live-locks, hangs and "worst case scenarios" such as thread starvation

Traditional approach to SMT verification

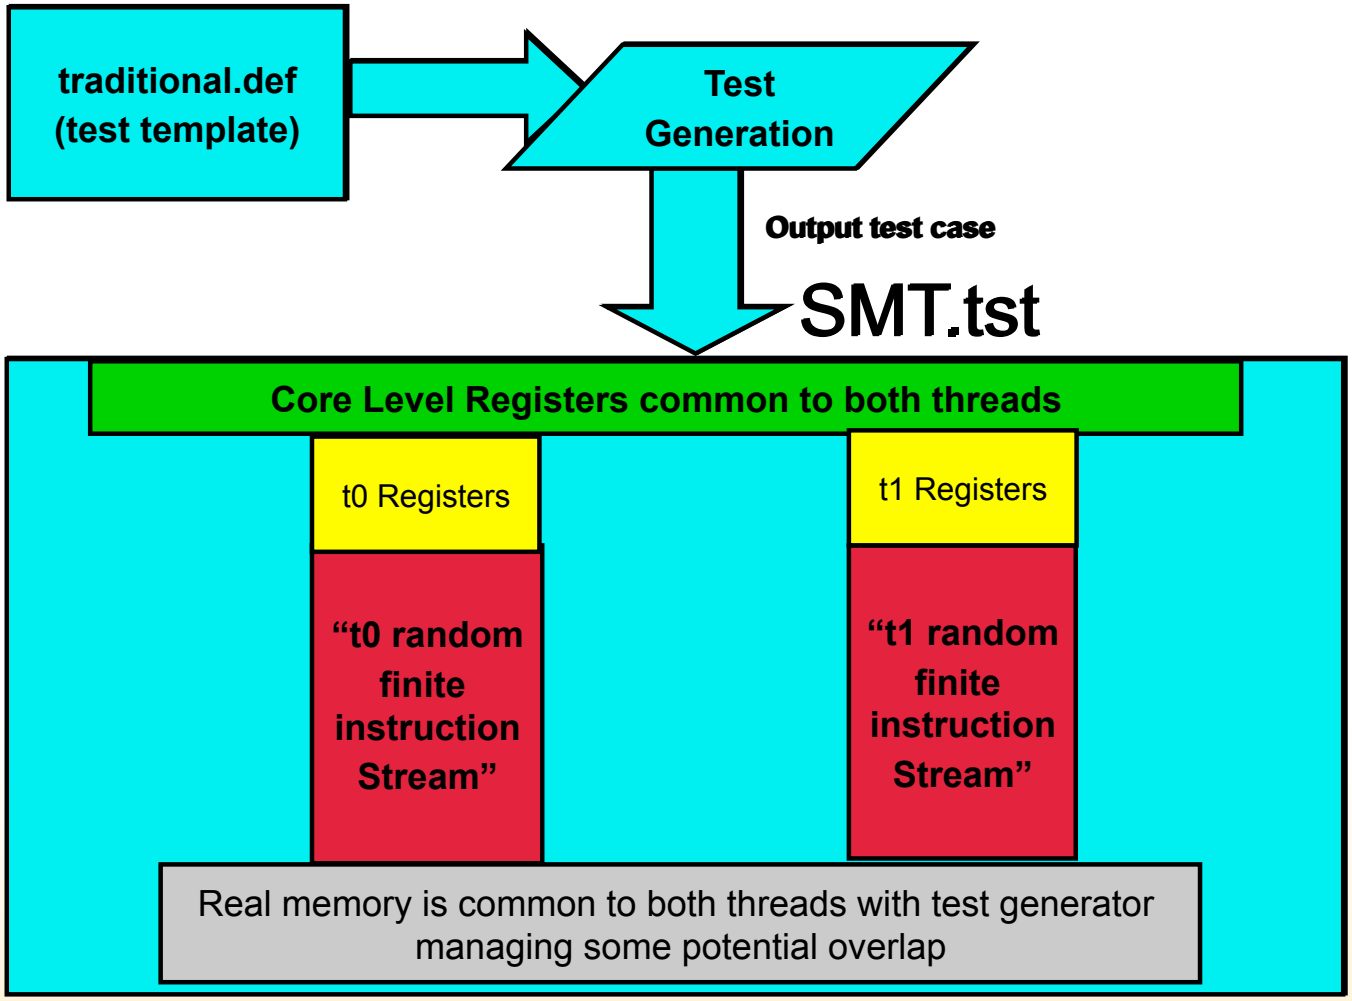

Irritator thread example for two-threaded core (e.g. POWER6)

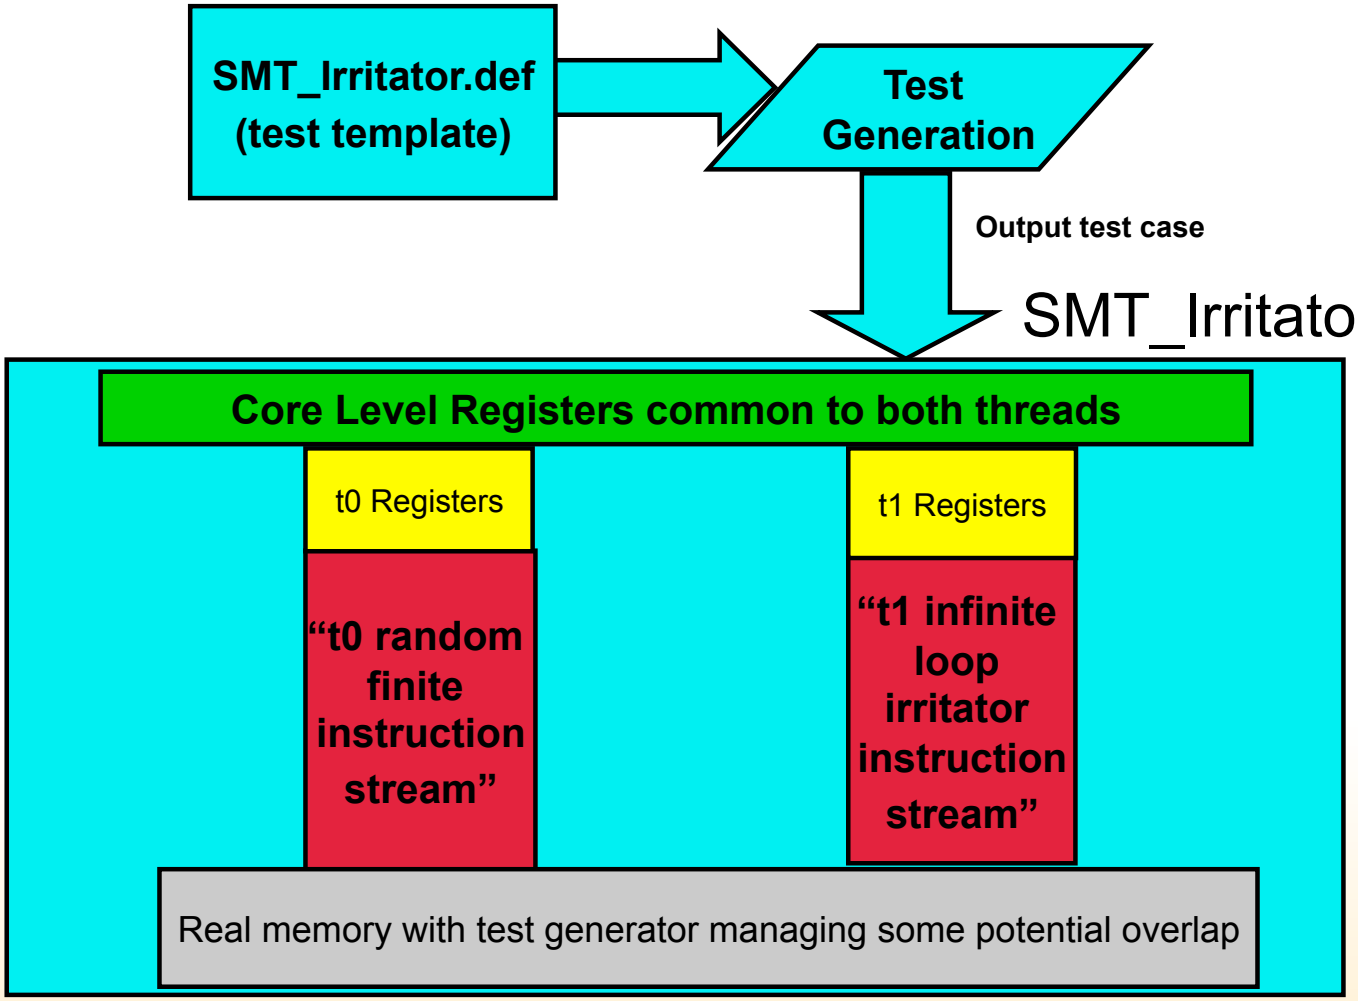

Irritator thread restrictions

- Cannot cause unexpected exceptions
- Cannot modify memory read by random thread
- Cannot modify registers shared with other threads
- Architected results may be undefined

Guaranteed Worst-Case Cross-Thread Resource Contention and SMT Balanced Test Length

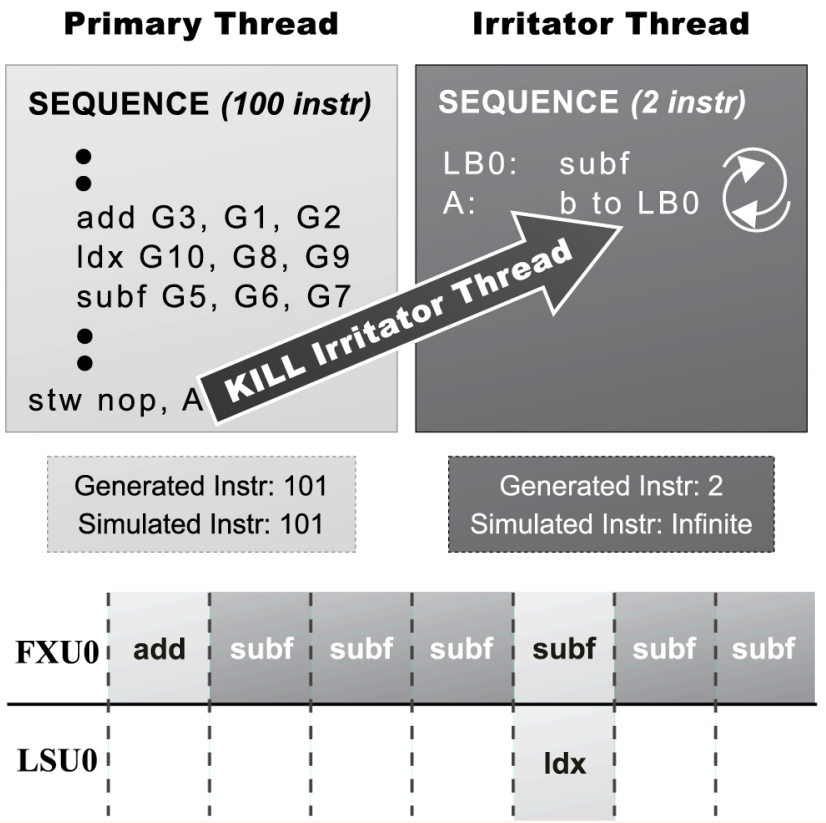

- **Irritator Thread Test Generation**
  - Description:
    - Random generation on Primary Thread
    - Tight infinite loop on one or more "irritator" threads
    - Primary Thread kills Irritator Thread(s) by terminating infinite loop
  - Balanced Test Length Among All Threads
  - Targets worst case (tight infinite loop) cross thread interaction
    - Impossible to miss:
      - live-locks
      - hangs
      - thread starvation
  - Efficient Test Generation:
    - Only generate 1 pass of irr thread loop
    - Net effect is test generation time governed by Primary Thread
      - 2x perf improvement for 2 threads
      - 4x perf improvement for 4 threads

Guaranteed Cross-Thread Interaction

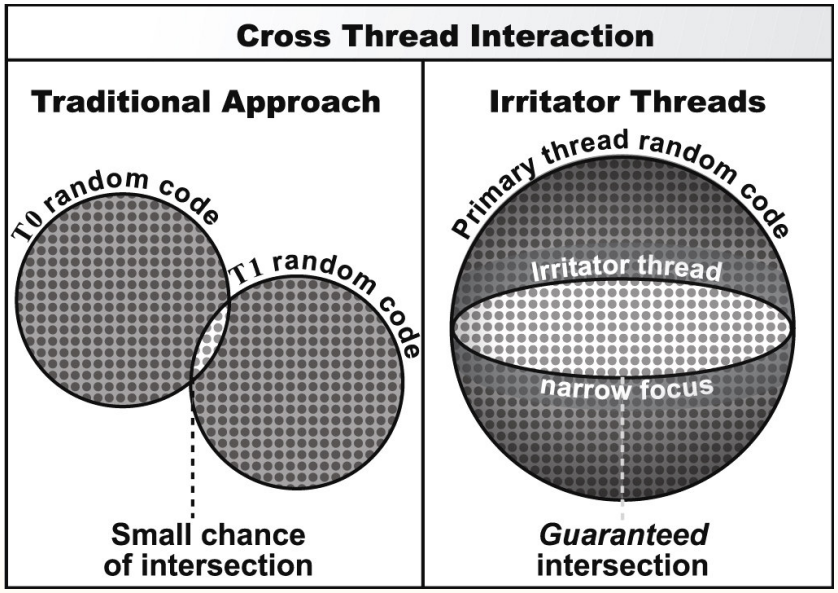

- **Extremely effective cross-thread interaction when compared to traditional approach**

P7 SMT4 "Irritator Thread"

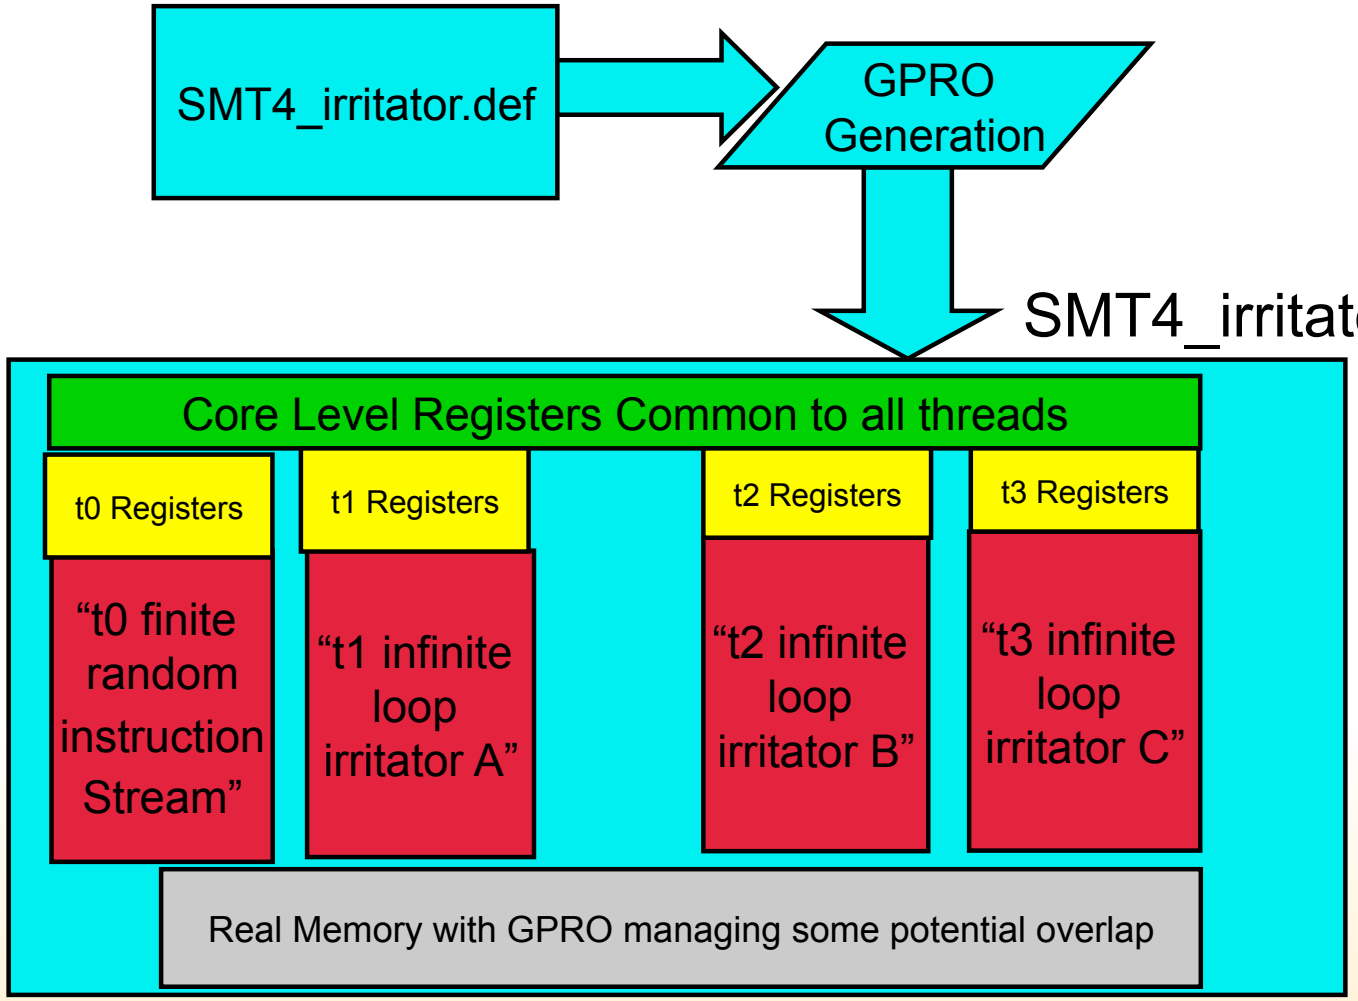

Irritator Thread Restrictions

- Cannot cause unexpected exceptions
- Cannot modify memory read by Random threads
- Cannot modify registers shared with other threads
- Architected results may be undefined

P7 SMT4 "Irritator Thread" Example

Long Random Thread

```
SEQUENCE
    REPEAT 100
        SELECT
            Group_All
    stw nop, A
    stw nop, B
    stw nop, C
```

Generated Instr: 103
Simulated Instr: 103

Irritator Thread A

```
LB0: SEQUENCE
        fdiv
        add
        ld
A:      b to LB0
```

Generated Instr: 4
Simulated Instr: Infinite

Irritator Thread B

```
LB1: SEQUENCE
        slbie
B:      b to LB1
```

Generated Instr: 2
Simulated Instr: Infinite

Irritator Thread C

```
SEQUENCE
C:    b2self
```

Generated Instr: 1
Simulated Instr: Infinite

POWER7 Irritator Thread Bug Summary

- With irritator threads, removed 23 "high quality" bugs traditionally seen only in post-silicon:
  - Flush Related: 6
  - Hang: 5
  - Thread Starvation: 5
  - Live-lock: 4
  - Cache Write Transition: 2
  - Branch to wrong target EA: 1
- Patent pending

Additional Advancements over POWER6 Verification

- **Software simulation on core model**
  - 2x number of simulation cycles
  - Very methodical effort driven by Core Directed Test Plan
  - 2.5x number of test cases simulated
  - Most effective use of coverage ever by POWER processor team
    - 45K discreet (Bugspray) coverage events
- **Hardware-accelerated simulation**
  - 7x number of accelerator cycles
  - 10x number of targeted exerciser shifts based on Core Directed Test Plan
    - Coverage-driven focused exerciser shifts vs. running random
- **Formal Verification**
  - Broad application of the technology across all areas of the chip
  - Extensive leverage from assertion-based verification <-> designer-level verification
  - Deep dive FV reviews conducted early in the project to identify key areas

Additional Improvements: Design for Verification

- **"POWER Architecture for Verification"**
  - Tight interlock between POWER Architecture Specification Team and Verification Team
  - Verification team actually maintained the official POWER Architecture Change list
- **Extensive "chicken switches" for back-off modes**
  - Any area considered risky from verification team
  - Degraded performance, but does not gate forward progress
- **"Hardware Irritators"**
  - Special modes built into silicon to enable stressing of otherwise rarely occurring microarchitecture events
  - "Opposite of a chicken switch"
  - Verification-only usage
    - Used for either pre-silicon (including EoA) or post-silicon testing
  - Patent pending

Overview

1. Background: POWER Processor History / Roadmap

2. Verification Methodology

3. Verification Execution

4. Verification Advances

5. Summary

Summary

1. **POWER7 verification effort is the new "gold standard" in IBM**
   - Release 1.0 hardware shipped to select customers (e.g. Rice University)
2. **Effective Utilization of Coverage Across All Simulation Levels**
   - Direct comparison of coverage from one platform to the next
3. **Irritator Threads for SMT Verification**
   - Found 23 "tough" bugs on POWER7
4. **Exercisers on Accelerator (EoA) Exploitation**
   - Allowed Development of High-Quality Exerciser Shifts later utilized for Post-Silicon Validation Effort
5. **Constant planning for Post-Silicon Validation**
   - Hardware Irritators and Back-off Mode Testing

Global Design and Verification Team

- **POWER7 team was distributed across 3 continents and multiple locations, because**
  - –Using teams of various competence centers
  - –Higher integration density with more logic combined on single chip leads to multiple teams joining single chip effort
  - –Large team not being available at one site at required point in time

- **Consequence of distributed team:**
  - –Additional communication needs
  - –Strict project management with tracking, reviews and clearly defined milestone
  - –Centralized data collection &
    Web based data analysis for
    - Defects, Coverage, Test results, ...
  - –Focus on standardization and common processes across teams

Thank You's to IBM team members

- **Borrowed charts / content input:**
  - Ron Kalla
  - Avi Ziv
  - Michal Rimon
  - Bryan Hickerson
  - John Schumann
  - Wolfgang Roesner
  - Viresh Paruthi
  - Mike Behm
  - Klaus-Dieter Schubert
  - Charles Meissner

- **Reviewers**
  - Wolfgang Roesner
  - Mendy Furmanek
  - Klaus-Dieter Schubert
  - Keith Sharp

REFERENCE MATERIAL

The IBM Power™ 750 Express is the highest performing 4-socket system on the planet. In addition it outperforms all other non-IBM 8 and 16-socket systems

System Name	Cores	Chips	Cores/ Chip	Threads/ Core	Peak*
IBM Power 750	32	4	8	4	1060
HP ProLiant DL585 G6 (2.8 GHz AMD Opteron 8439 SE)	24	4	6	1	416
HP Integrity rx6600 (1.6 GHz/24MB Dual-core Intel Itanium 2)	8	4	2	1	102
HP ProLiant DL580 G5 (2.66 GHz, Intel Xeon X7460)	24	4	6	1	291
Sun SPARC Enterprise T5440	32	4	8	8	360
Sun SPARC Enterprise M4000	16	4	4	2	152
HP ProLiant DL 785 G6 (2.8 GHz AMD Opteron 8439 SE)	48	8	6	1	800
Unisys ES7000 Model 7600R, Intel Xeon X7460, 2.66 GHz	48	8	6	1	527
Sun SPARC Enterprise M5000	32	8	4	2	296
HP Integrity rx8640 (1.6 GHz/24MB Dual-core Intel Itanium 2)	16	8	2	1	209
Unisys ES7000 Model 7600R, Intel Xeon X7460, 2.66 GHz	96	16	6	1	1049
Sun SPARC Enterprise M8000	64	16	4	2	753

*Peak – SPECint_rate2006 (Peak)

Substantiation:

- Competitive benchmark results reflect results published as of February 3, 2010. The results are the best results for four-socket single (non-clustered) systems using POWER™, Intel® x86, AMD Opteron™ x86, SPARC and Intel Itanium® processors. IBM Power 750 result submitted on February 8, 2010.
- SPEC® and the benchmark names SPECrate®, SPECint®, and SPECjbb® are registered trademarks of the Standard Performance Evaluation Corporation. For the latest SPEC benchmark results, visit http://www.spec.org

The IBM Power 750 Express is the most energy efficient 4-socket system on the planet.

System Name	Cores	Chips	Cores/ Chip	Threads/Core	Peak *	WATTs	Peak / WATT
IBM Power 750	32	4	8	4	1060	1950	0.54
HP ProLiant DL585 G6 (2.8 GHz AMD Opteron 8439 SE)	24	4	6	1	416	1548	0.26
HP Integrity rx6600 (1.6 GHz/24MB Dual-core Intel Itanium 2)	8	4	2	1	102	1600	0.06
HP ProLiant DL580 G5 (2.66 GHz, Intel Xeon X7460)	24	4	6	1	291	1412	0.20
Sun SPARC Enterprise T5440	32	4	8	8	360	2700	0.13
Sun SPARC Enterprise M4000	16	4	4	2	152	2016	0.07

* Peak = SPECint_rate2006 (Peak)

- Substantiation
- Competitive benchmark results reflect results published as of February 3, 2010. The results are the best results for four-socket single (non-clustered) systems using POWER, Intel x86, Opteron x86, SPARC and Itanium processors. IBM Power 750 result submitted on February 8, 2010.
- SPEC and the benchmark names SPECrate, SPECint, and SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation. For the latest SPEC benchmark results, visit http://www.spec.org
- Performance/WATT is calculated by dividing the performance from the tables below by the recommended maximum power usage for site planning. This defines the requirement for the power infrastructure. Actual power used by the systems will be less than this value for all of the systems. For Power Systems™ servers, this information is available in the site planning guides available through www.ibm.com. For HP systems, this information is contained in the QuickSpecs for each system available through www.hp.com.
- For Sun systems, this information is available through the Site Planning Guides available through www.sun.com.

The IBM Power 750 Express has more SAP performance than any 8-socket system in the industry – and is even comparable to a 128-core, 32-socket Sun M9000.

Substantiation:

All results are 2-tier, SAP EHP 4 for SAP ERP 6.0 (Unicode). IBM results valid as of 2/8/2010. Competitive results valid as of 2/3/2010.

- IBM Power 750 Express **certification number not available at press time and can be found at** www.sap.com/benchmarks.
  IBM Power 750 Express: **4p / 32–c / 128 – t, POWER7™, 3.55 GHz, 256 GB memory, 15,600 SD users, dialog resp.: 0.98s, line items/hour: 1,704,330, Dialog steps/hour: 5,113,000, SAPS: 85,220, DB time (dialog/ update):0.015s / 0.028s, CPU utilization: 99%, AIX® 6.1, DB2® 9.7**
- Sun SPARC Enterprise T5540: **4p / 32-c / 256 –t, UltraSPARC T2 plus OC, 1.6 GHz, 256 GB memory, 4720 SD users, dialog resp: 0.97s, line items/hour: 516,670, dialog steps/hour: 1,550,000, SAPS: 25,830, Solaris 10, Oracle 10g , cert# 2009026-1**
- HP DL585 G6: **4p / 24-c / 24-t, AMD Opteron 8439 SE, 2.8 GHz, 64 GB memory, 4665 SD users, dialog resp: 0.96s, line items/hour: 510,670, dialog steps/hour: 1,532,000, SAPS: 25,530, Windows Server 2008 EE, , SQL Server 2008, cert#: 2009025**
- HP DL785 G6: **8p / 48-c / 48-t, AMD Opteron 8439 SE, 2.8 GHz, 128 GB memory, 8280 SD users, dialog resp: 0.96s, line items/hour: 907,000, dialog steps/hour: 2,721,000, SAPS: 45,350, Windows Server 2008 EE, , SQL Server 2008, cert#: 2009035**
- Sun Fire x4640: 8**p / 48-c / 48–t, Six-core AMD Opteron 8435, 2.6 GHz, 256 GB memory, 10,000 SD users, dialog resp: 0.9s, line items/hour: 1,101,330, dialog steps/hour: 3,304,000, SAPS: 55,070, Solaris 10, Oracle 10g, cert# 2009049**
- HP DL380 G6: **2p / 8-c / 16-t, Intel Xeon® X5570, 2.93 GHz, 48 GB memory, 3171 SD users, dialog resp: 0.94s, line items/hour: 347,670, dialog steps/hour: 1,043,000, SAPS: 17,380, SUSE Linux® Enterprise Server 10, MaxDB 7.8, cert#: 2009006**
- Sun SPARC Enterprise M9000: 32**p / 128-c / 256–t, Six-core AMD Opteron 8435, 2.6 GHz, 1 TB memory, 17,430 SD users, dialog resp: 0.95s, line items/hour: 1,909,670, dialog steps/hour: 5,729,000, SAPS: 95,480, Solaris 10, Oracle 10g, cert# 2009038**

Consolidation onto POWER7 can deliver significant savings. Ninety-two Sun SPARC Enterprise T2000 servers can be consolidated into a single IBM Power 750 Express system, saving 95% of the cores for software licensing, 97% on floorspace, and 95% on energy.

Calculation Summary: the Power 750 has 30.93 better SPECjbb2005 performance than the Sun T2000. Assuming a 3x virtualization factor for greater consolidation - then 92 Sun File T2000 servers could be consolidated onto one Power 750 Express server (30.93 * 3 = 92.8 servers rounded to 92 T2000 servers)

HW System Name	JVM Instances	Cores	Processor chips	HW Threading	bops	bops/JVM
IBM Power 750 Express	32	32	4	Yes	2,478,929	77,467
Sun File T2000	4	8	1	Yes	74,365	18,591

System Name	SPECjbb2005	Max Watts	Rack space	Cores	Systems	Total Perf	Total Cores	Total Watts	Total Rack Space
IBM Power 750 Express	2,478,929	1950	4	32	1	1,380,000	32	1950	4
Sun File T2000	74,365	450	2	8	92	1,368,150	736	41,400	184
Savings with Power 750 Express							95.6%	95.2%	97.8%

Substantiation

1. SPEC and the benchmark names SPECrate, SPECint, and SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation. Competitive benchmark results stated reflect results published on www.spec.org as of February 08, 2010. The comparison presented below is based on a consolidation of a legacy 8-core Sun SPARC Enterprise T2000 UltraSPARC T1 servers into a 32 core IBM Power 750. For the latest SPEC benchmark results, visit http://www.spec.org.
2. SPECjbb2005 results are:
POWER7: IBM Power 750 Express with 4 chips, and 32 cores and four threads per core with a result of 2,300,000 bops and 71,875 bops/jvm submitted to SPEC on February 8, 2010.
SPARC: Sun Microsystems Sun SPARC Enterprise T2000 with 1 chip, 8 cores and 4 threads per core with a result of 74,356 bops and 18,591 bops/jvm
*The virtualized system count and energy savings were derived from several factors:
- A **performance ratio factor** of 30.93X was applied to the virtualization scenario. The performance factor is the SPECjbb2005 result of the Power 750 Express divided by the result of the competitive Sun SPARC Enterprise T2000 server.
- A **virtualization factor** of 3X was applied to the virtualization scenario using utilization assumptions derived from an Alinean white paper on server consolidation. The tool assumes 19% utilization of existing servers and 60% utilization of new servers. Source - www.ibm.com/services/us/cio/optimize/opt_wp_ibm_systemp.pdf.
Space calculation: The Sun T2000 is 2U in height and 21 can fit into a 42U rack. The 750 is 4U in height.
Power consumption figures of 1950W for the IBM Power 750 and 450W for the Sun T2000 were based on the maximum rates published by IBM and Sun Microsystems, respectively. This information for the Power 750 is in "Model 8233-E8B server specifications" available at http://www-01.ibm.com/common/ssi/index.wss - search for Power 750. Sun T2000 Maximum AC power consumption of 450 WATTs was sourced from Sun SPAC Enterprise T2000 Servers site planning guide at http://docs.sun.com/app/docs/doc/819-2545-11 as of 2/9/2010.

52

345 million kilowatt-hours are used yearly by the 91,920* Sun SPARC Enterprise T2000 servers shipped since 2005 above what would be used yearly if consolidated into 1,000 IBM Power 750 Express servers at the rate of 92 to 1.
That's enough electricity to supply 34,500 homes for a year.**

Substantiation

- Maximum power for 1 IBM Power 750 Express server = 1950 watts
- Maximum power for 92 Sun Fire T2000 servers = 92 x 450 = 41,400 watts
- Excess power per consolidation instance (92 Sun T2000's into one Power 750) = 39,450 watts
- Number of consolidations required = 1,000 (91,320 / 92)
- Total excess kilowatt-hours per year =
  39,450 watts x 24 hrs/day x 365 days/yr x 1,000 = 345 million kilowatt-hours per year

* Source: 3Q09 IDC Server Tracker

** Source: Wikipedia estimate of average annual household energy use of 10,000 kilowatt-hours

The IBM Power 750 Express has 28% more performance than a 64-core HP Integrity Superdome and requires only 83% as much power to run – at a fraction of the price.

Substantiation:
Notes:
1. SPEC and the benchmark names SPECrate, SPECint, and SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation. HP Integrity Superdome benchmark results stated reflect results published on www.spec.org as of February 08, 2010. For the latest SPEC benchmark results, visit http://www.spec.org.
2. SPECint_rate2006 Peak results are:
POWER7: IBM Power 750 Express with 4 chips, and 32 cores and four threads per core with a result of 1060 submitted to SPEC on February 8, 2010.
Itanium: Hewlett-Packard Integrity Superdome with 32 chips, 64 cores, and one thread per core with a result of 824.
3. The HP Integrity Superdome is a rack cabinet. The 750 is 4U in height.

Power consumption **is derived from the recommended maximum power for site planning. Actual power used by the systems will be less than this value for all of the systems.**

This information for the Power 750 Express is available at http://www-01.ibm.com/common/ssi/index.wss - search for Power 750. The maximum power requirement for the Power 750 is 1,950 Watts.

The information for the Integrity Superdome is in "QuickSpecs HP Integrity rx6600 Server" available at http://h18000.www1.hp.com/products/quickspecs/11717_div/11717_div.HTML, which shows the maximum power requirement for the Integrity Superdome of 12,196 VA. Using the Power Factor of 0.95 shown at http://www.spectra.com/pdfs/superdome.pdf, the maximum input power is 11,586 Watts.

Price comparison based on IBM analysis:
HP Superdome price estimated at $2,117,000 for the configuration described in the SPECint_rate2006 benchmark
IBM Power 750 Express U.S. list price = $275,420

The IBM Power 780 delivers leadership performance and consolidation capability vs. HP and Sun high-end servers. For example, eight HP Integrity Superdome 64-core systems utilized at 30% can be consolidated into a single IBM Power 780 server utilized at 80%, thus saving 87% of the cores for software licensing, reducing floorspace from 80 square feet to 7.6 square feet, and reducing energy costs by 92%.

Substantiation:

SPECint_rate2006 Results										
System Name	Cores	Chips	Cores / Chip	Threads / Core	Peak	Published		Wattage	Performance per watt	Performance per core
IBM Power 780	64	8	8	4	2530	February 2010		6,400	395.31	39.53
HP Integrity Superdome	64	32	2	1	824	October 2006		12,196	67.56	12.88
HP Integrity Superdome	128	64	2	1	1648	September 2006		24,392	67.56	12.88
Sun SPARC Enterprise M9000	256	64	4	4	2586	October 2009		44,800	57.72	10.10

Performance per watt is calculated by dividing the performance in the table above by the recommended maximum power for site planning. Actual power used by the systems will be less than this value for all of the systems. The maximum power requirement for the Power 780 is 6,400 Watts and is available at http://www-01.ibm.com/common/ssi/index.wss - search for Power 780.

Power consumption figures of 6400 W for the IBM Power 780, 12,196 W / 24,392 W for the HP Superdome and 44,800 W for the Sun SPARC Enterprise M9000 were based on the maximum rates published by IBM, HP and Sun Microsystems, respectively. The information for the HP Integrity Superdome is in "QuickSpecs HP Integrity Superdome Servers 16- processor, 32-processor, and 64- processor Systems" available at www.hp.com. The information for the Sun SPARC Enterprise M9000 is in the "Sun SPARC Enterprise M9000 Servers Site Planning Guide" available at www.sun.com

The virtualized system count and energy savings were derived from several factors:

A performance ratio factor was applied to the virtualization scenario based on SPECint_rate2006. The performance factor is simply the SPECint_rate2006 result per core of the Power 780 divided by the per core result of the HP or Sun system.

Power 780 (64-core, 8 chips, 8 cores per chip, 3.8 GHz) SPECint_rate2006 2,530 peak as of 2/8/2010. HP Superdome (64-core, 32 chips, 2 cores per chip) 1.6 GHz, SPECint_rate2006 824 peak published October 2006. Data valid as of 2/3/2010.

Sun SPARC Enterprise M9000 (256-core, 64 chips, 4 cores per chip) 2.88 GHz, SPECint_rate2006 2,586 peak published October 2009. Data valid as of 2/3/2010. SPEC® results available at: www.spec.org.

A virtualization factor of 3.157X was applied to the virtualization scenario using utilization assumptions derived from an Alinean white paper on server consolidation. The tool assumes 19% utilization of existing servers and 60% utilization of new servers. Source - www.ibm.com/services/us/cio/optimize/opt_wp_ibm_systemp.pdf.

Air conditioning power requirement estimated at 50% of system power requirement.

Energy cost of $.1031 per kWh is based on 2009 YTD US Average Retail price to commercial customers per US DOE at http://www.eia.doe.gov/cneaf/electricity/epm/table5_6_b.html as of 1/27/2010.

The reduction in floor space, power, cooling and software costs depends on the specific customer, environment, application requirements, and the consolidation potential. Actual numbers of virtualized systems supported will depend on workload levels for each replaced system.

System data for HP from the HP Superdome Datasheet and HP Integrity Superdome Server — specifications both available at www.hp.com. System data for Sun from the Sun SPARC Enterprise M9000 Tech Specs available at www.sun.com. Data is current as of January 27, 2010.

The modular enterprise class POWER systems have continued to deliver significant improvements year over year. With the Power 770 server, clients can consolidate four POWER5™ processor-based Power 570 systems onto one Power 770. In fact, it only takes two nodes and moving to the Power 770 still has an effective capacity increase of 50%.

Substantiation:

System Name	Cores	Nodes	rPerf	Utilization	Effective Performance	WATTs	Maintenance
IBM Power 770	24	2	261.19	60%	156.7	3200	
IBM System p® 570 (x4)	64	16	309.8	30%	92.9	20,800	
Advantage / Savings		87% Less Space			> 50% Capacity	84% Less Energy	

56

POWER7 systems deliver up to three or four times the energy efficiency of POWER6™ based systems.

Substantiation:

- SPEC and the benchmark names SPECrate, SPECint, and SPECjbb are registered trademarks of the Standard Performance Evaluation Corporation. Benchmark results stated reflect results published on www.spec.org as of February 8, 2010. The comparison used in the claim is based on a consolidation of the best high-end POWER6 result (Power 595) to the Power 780, the best mid-range POWER6 result (Power 570) to the Power 770, and the best four-socket and above POWER6 Express results with the Power 750 Express. For the latest SPEC benchmark results, visit http://www.spec.org.
- Performance/WATT is calculated by dividing the performance from the tables below by the recommended maximum power usage for site planning. This defines the requirement for the power infrastructure. Actual power used by the systems will be less than this value for all of the systems. This information is available in the site planning guides available through www.ibm.com.

	SPECint_rate2006 results as of January 7, 2010						
System Name	Cores	Chips	Cores/ chip	Threads/ Core	Peak	WATTs	Peak / WATT
IBM Power 780	64	8	8	4	2530	6400	0.39
IBM Power 595	64	32	2	2	2160	28300	0.07
IBM Power 770	64	8	8	4	2013	6400	0.31
IBM Power 570	16	8	2	2	542	5600	0.09
IBM Power 750 Express	32	4	8	4	1060	1950	0.54
IBM Power 560 Express	16	8	2	2	363	2400	0.15
IBM Power 550 Express	8	4	2	2	263	1400	0.18

POWER7 systems deliver up to three or four times the performance with less energy than POWER6 based systems.

Substantiation:

rPerf (Relative Performance) is an IBM estimate of commercial processing performance relative to other IBM UNIX® systems. The comparison used in the claim is based on these comparisons:

- 4-node Power 570 (POWER6+™) to a 1-node Power 780 (POWER7)
- 4-node Power 570 (POWER6+) to a 3-node Power 780 (POWER7)
- 4-node Power 570 (POWER6+) to a 3-node Power 770 (POWER7)
- 2-node Power 560 Express (POWER6+) to Power 750 Express (POWER7)

Performance/WATT is calculated by dividing the performance (rPerf) from the tables below by the recommended maximum power usage for site planning. This defines the requirement for the power infrastructure. Actual power used by the systems will be less than this value for all of the systems. This information is available in the site planning guides available through www.ibm.com

System Name	Nodes	Processor Technology	Proc. Freq.	Energy (Watts)	rPerf	Factor* (P7 over P6)
IBM Power 780	1	POWER7	3.8 GHz	1600	195	1.38
IBM Power 570	4	POWER6+	5.0 GHz	5600	141	-
IBM Power 780	3	POWER7	3.8 GHz	4800	523	3.7
IBM Power 570	4	POWER6+	5.0 GHz	5600	141	-
IBM Power 770	3	POWER7	3.1 GHz	4800	443	3.1
IBM Power 570	4	POWER6+	5.0 GHz	5600	141	-
IBM Power 750 Express	1	POWER7	3.55 GHz	1950	331	3.3
IBM Power 560 Express	2	POWER6+	3.6 Ghz	2400	100	-

* Factor = Performance increase factor of POWER7 system over POWER6 system for less energy

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (cont.)

Revised February 9, 2010

60

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3, AIX 5L or AIX 6 were used. All other systems used previous versions of AIX. The SPEC CPU2006, SPEC2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised March 12, 2009

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html .

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX 5L were used. All other systems used previous versions of AIX. The SPEC CPU2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seq=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	**http://hmmer.janelia.org/** http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Revised March 12, 2009

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- **rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.**

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

==========

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007