


# Multicore Enablement Challenges for Semiconductor Manufacturers

January 29, 2008



---

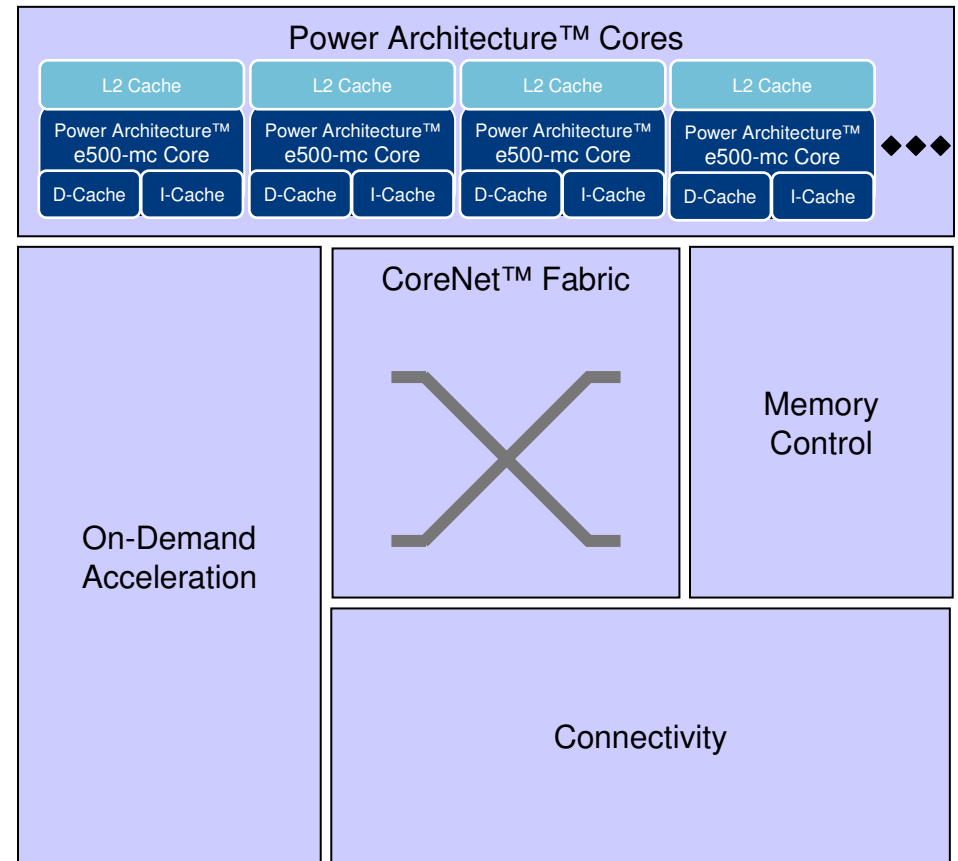
**Steve Furr**

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © Freescale Semiconductor, Inc. 2006.



# Freescale Multicore Platform

- Supporting more than 32 Power Architecture™ cores with L2 back-side cache per core
- Introducing a scalable CoreNet™ Fabric for concurrent, non-blocking, hardware-based 100% cache-coherent platform connectivity
- Expected to enable 2-3 times improved system performance



# Problem: Multicore Affects How Users Design, Write and Debug Software

## Partitioning the problem across multiple cores


- Requires new programming models
- Not well supported by today's tools – state of practice is low
- Hard, manual challenge

## Critical impact on system level performance and time to market

- Complexity of programming model makes diagnosis difficult (longer)
- Partitioning approach is difficult to validate and tune (longer)
- Adding cores increases synchronization and latency sources – driving down scalability (longer and slower)

## Reduced visibility into activity on-chip

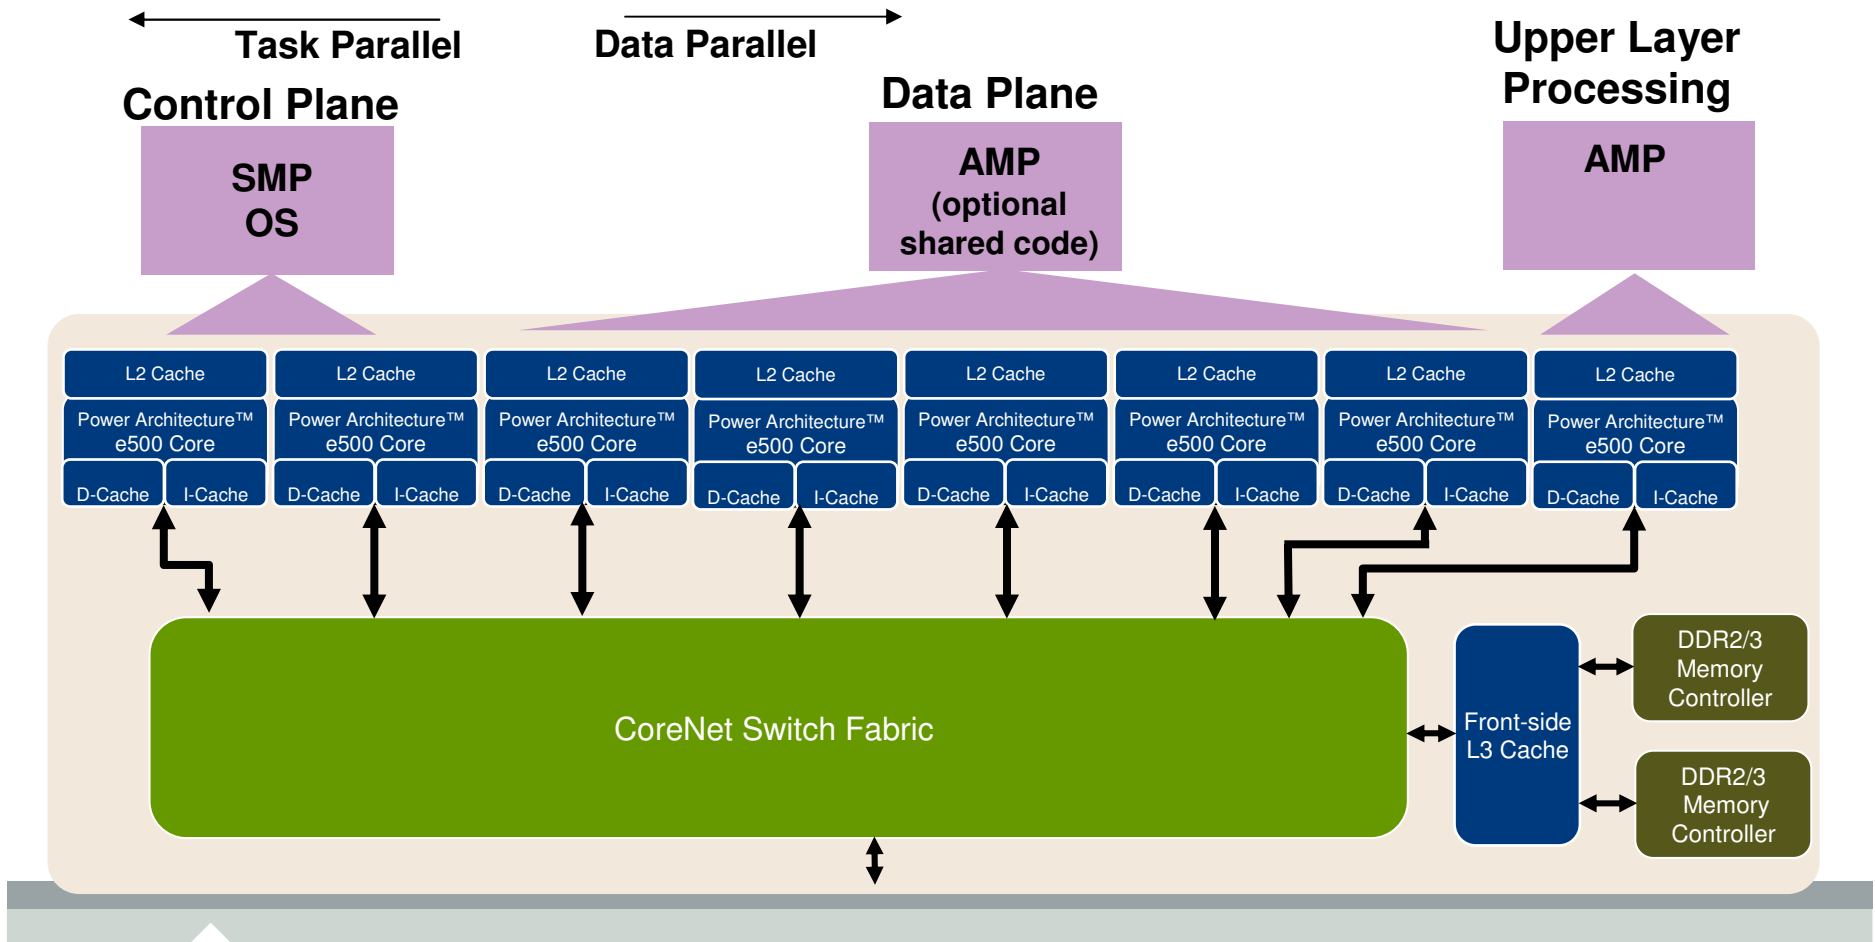
- Tracing/monitoring at the system and application level
- Finding bottlenecks in performance – correlating to application behavior
- Correlating application behavior to silicon behavior (cache, memory hierarchy, internal buses)



# Networking Core Usage Model

Flexible choice of OS model on cores

Assumed common use case (number of cores in categories may vary):



# Use cases - networking

1. Software debugging
2. **High-level application profile performance analysis with correlated counts**
3. **High-level transactional analysis with counts**
4. **IP block utilization**
5. Correlation of events from trace, high-level counts, and IP block utilization
6. Setting up advanced triggering scenarios
7. Instruction trace from core
8. **Packet-flow trace**
9. Trace traffic on CoreNet
10. **Buffer recovery (garbage collection)**
11. **Cache analysis**
12. **DDR analysis**
13. Access to entire contents of all memories
14. Multiprocessor programming  
malbehavior detection
15. Debugging when system busy or hung
16. Power-On-Debug (POD)
17. Profiling to control debugging outside of SoC scope
18. Security through conditional debug using challenge/response
19. Understanding virtualization, ultravisor, and access control partitioning
20. Migrate silicon state to simulation models
21. Allow customers to leverage their own analysis tools
22. Helper core that can reset data-plane hardware and cores
23. Helper core gathers run-time statistics
24. **Core pipeline visualization**

Log Analysis

Profiling

# Solution: Provide greater visibility

## Core behavior

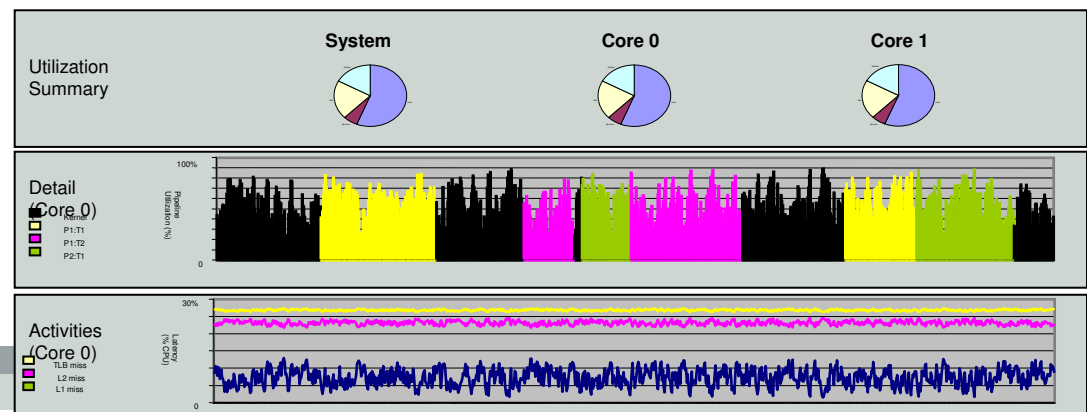

- Trace program (branch trace) or data (read/write)
- Count core event occurrences (I-cache, D-cache, L2, ...)

## Platform behavior

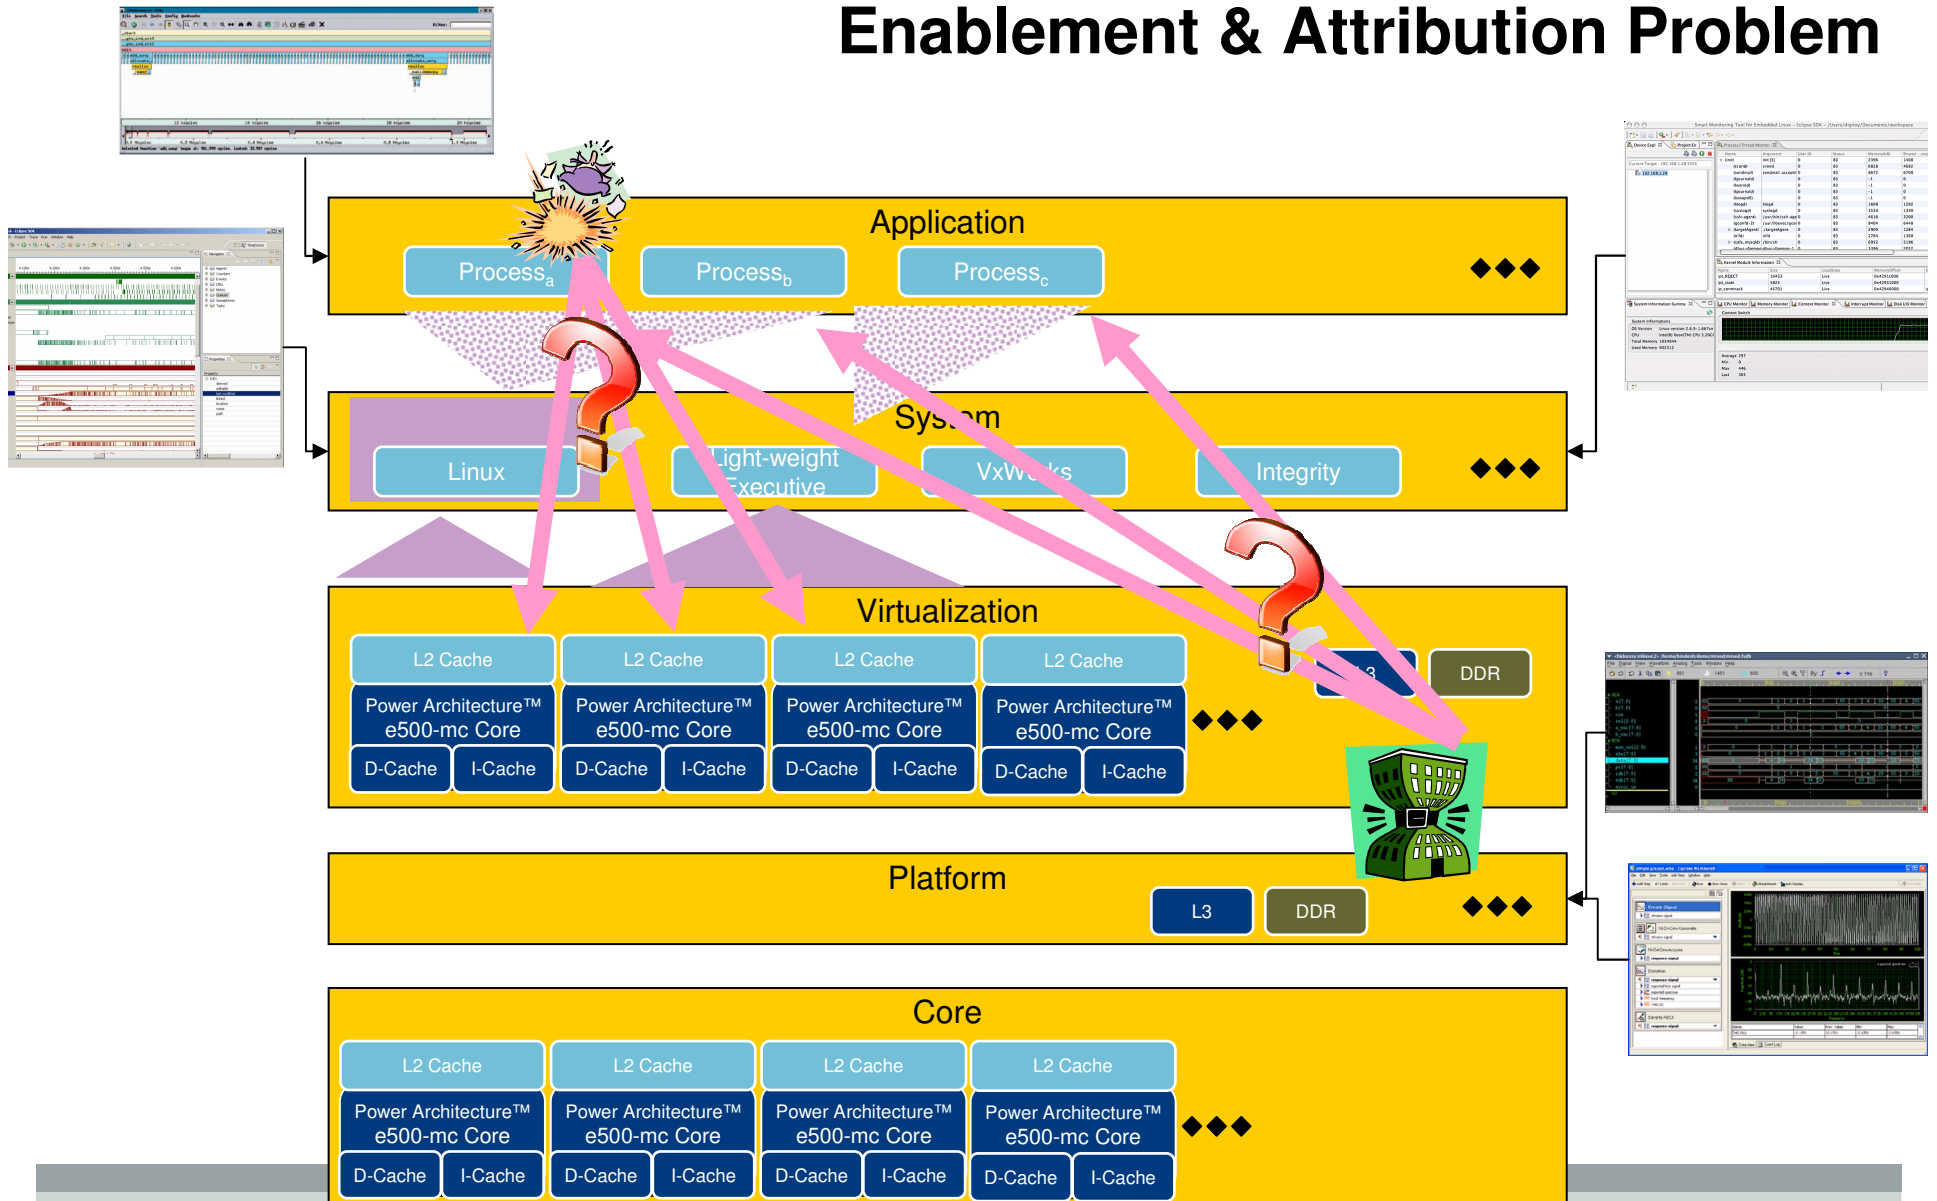
- Trace CoreNet™ transactions
- Trace external bus transactions (Serial RapidIO, DDR, etc.)
- Count platform event occurrences



# SO NOW WHAT?

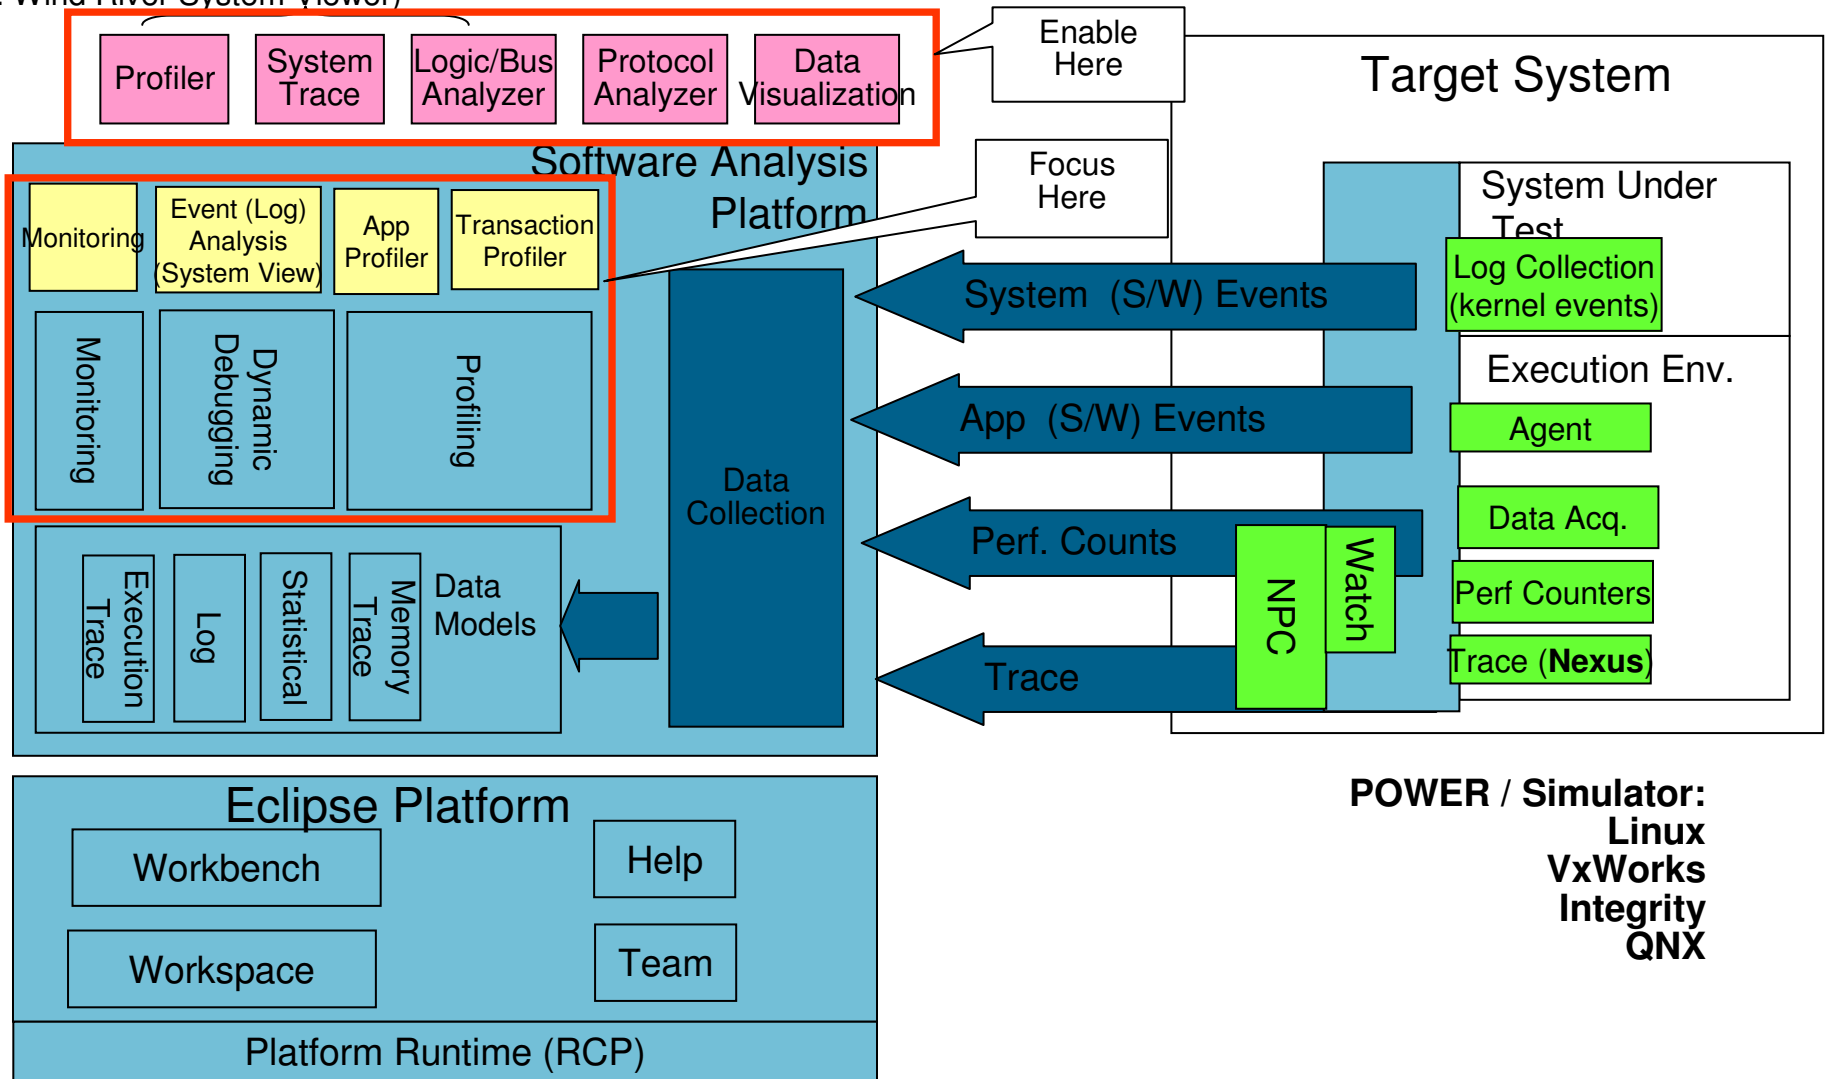


# Enablement & Attribution Problem



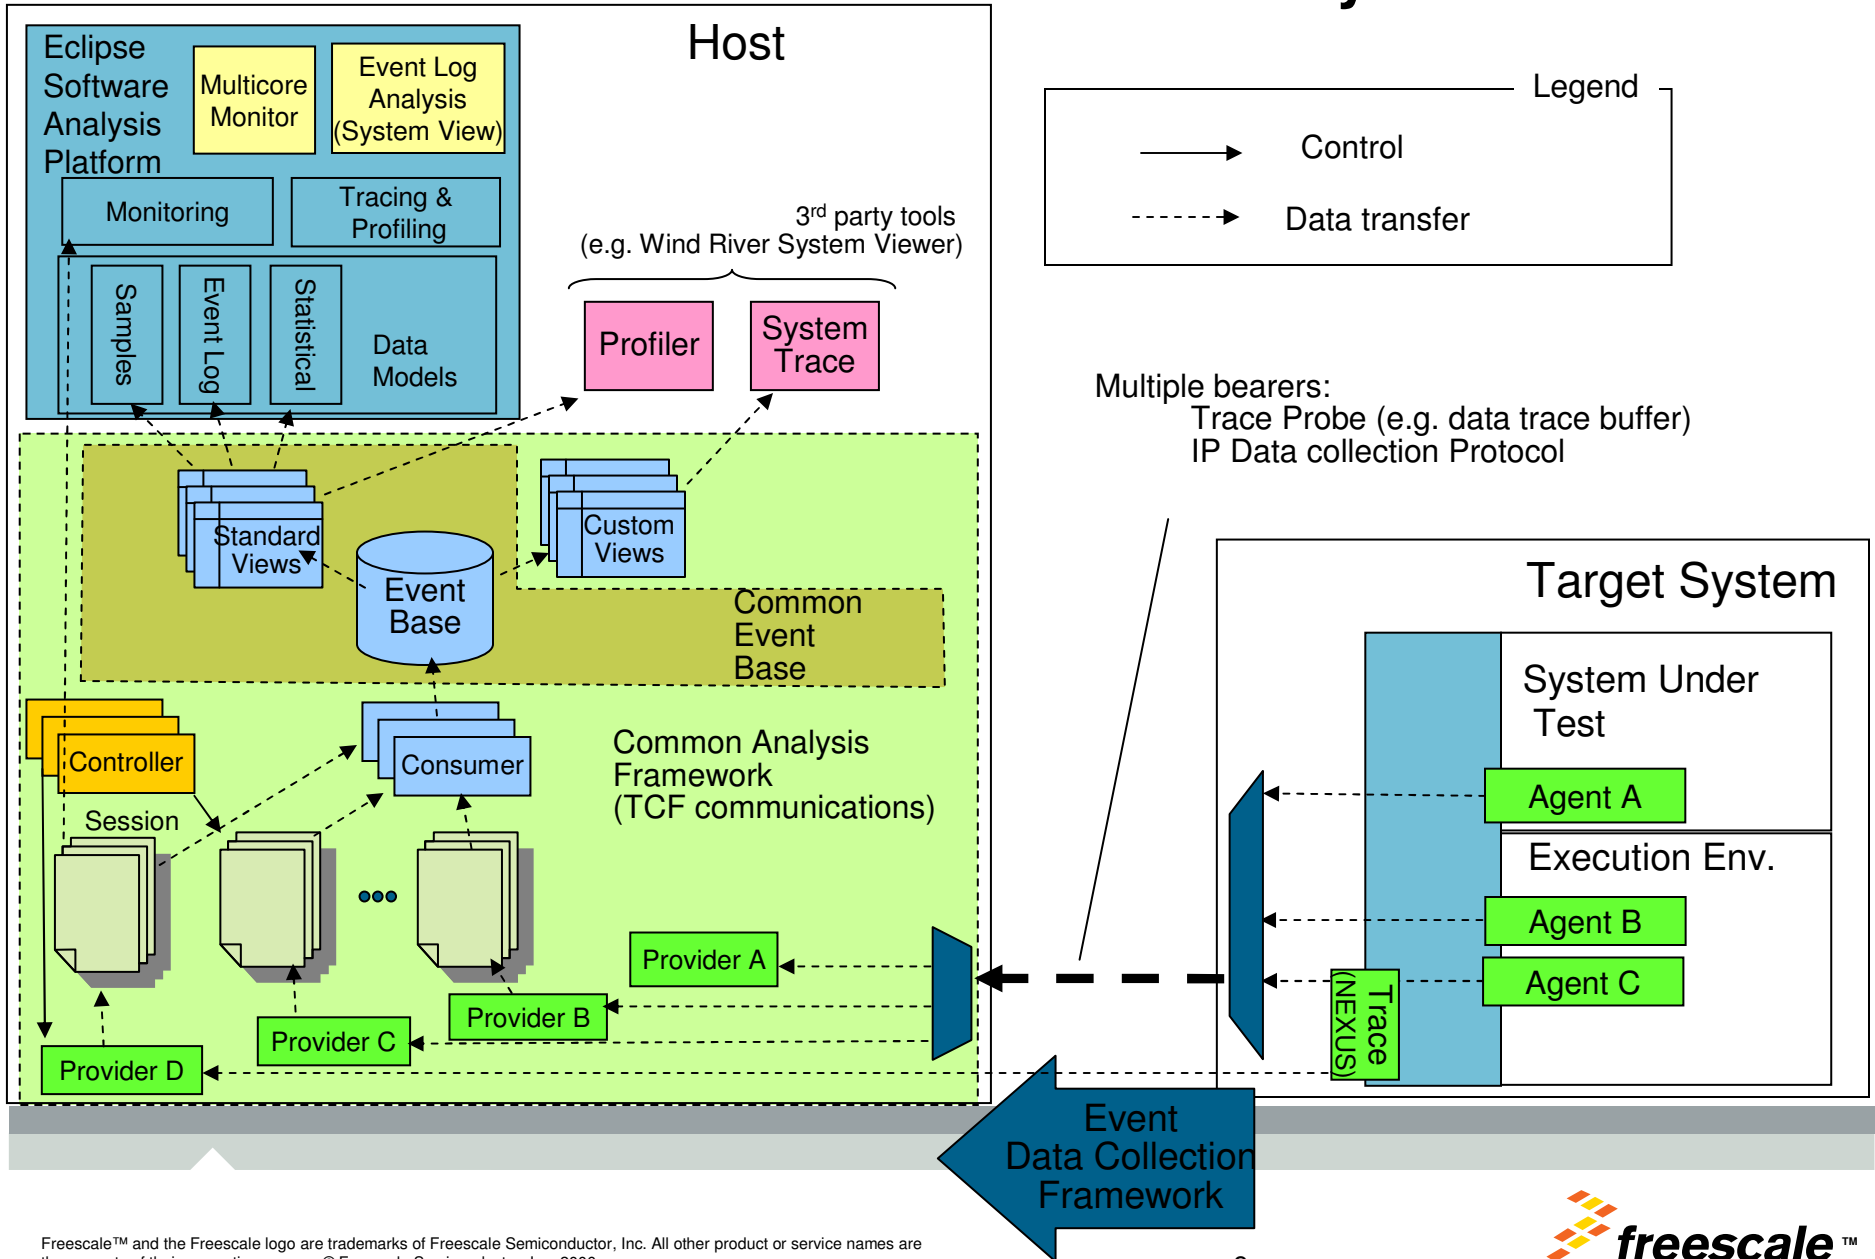
# Common Analysis Framework

3<sup>rd</sup> party tools  
(e.g. Wind River System Viewer)



**POWER / Simulator:**  
Linux  
VxWorks  
Integrity  
QNX

# Common Analysis Architecture



# Challenges

Establish a global ordering of events

- Synchronization of different time bases
- Correlation of different events based on global time stamp
- Provide high resolution time base for runtime events

Manage events from many different sources

- Platform and language neutral data collection methodology
- Scalability to different target environments (probe, embedded target)

Consolidate data feeds for different forms of tooling

- Standardized data schema for specific application domains
- Extensible formats that allow tools to ignore data they cannot process

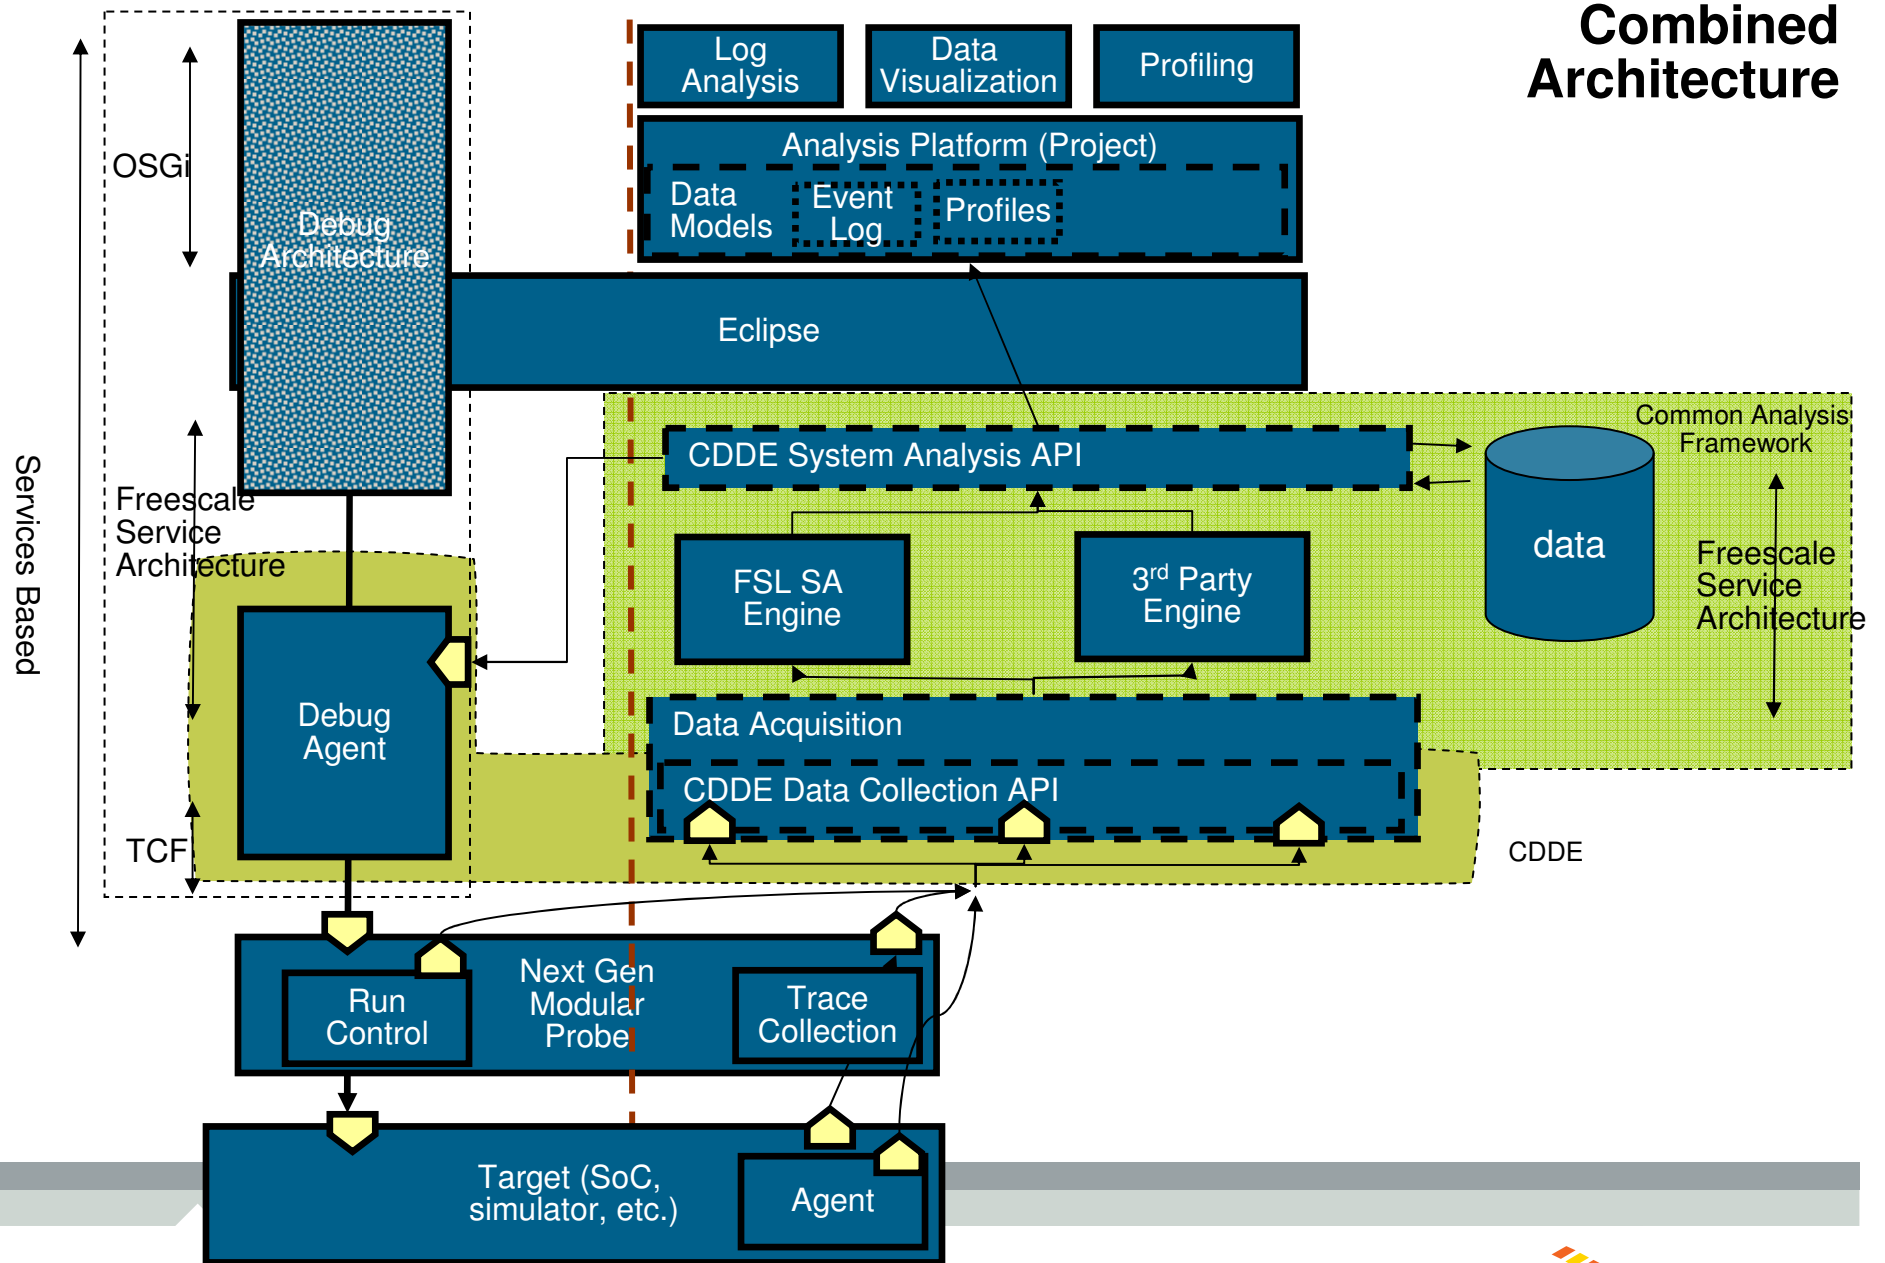
Scalability

- Handle large volumes of data
  - Some trace data sets (e.g. program/data trace) can produce two orders of magnitude more data than kernel event trace
- Continuous or stored data sets (e.g. traces)

Configuration

- Enable static and dynamic configuration of data feeds
- Target-specific information about data feeds

# Combined Architecture



Questions?

