Niagara(T1)
A CMT PROCESSOR

Rao Shoaib
Solaris Core Technology group
rao.shoaib@sun.com

Agenda:

- Why CMT Processors
- Highlights of Sun Niagara Processor
- Performance characteristics of T1
- Need for Virtualization
- CMT & Virtualization
- Sun Virtualization Solutions
- HW and Software Network Virtualization.

Case For CMT Processors

Tradional processor behavior

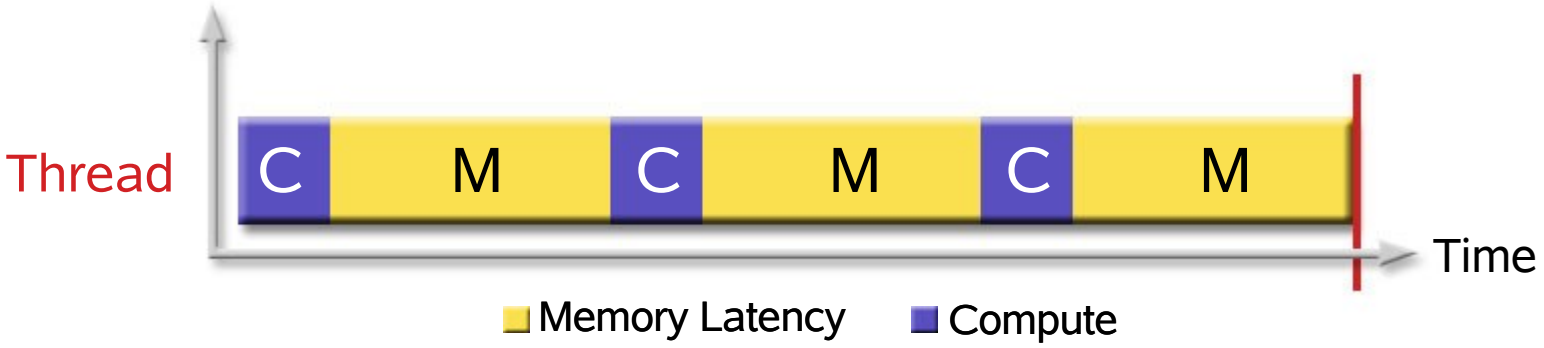

Single scalar processor

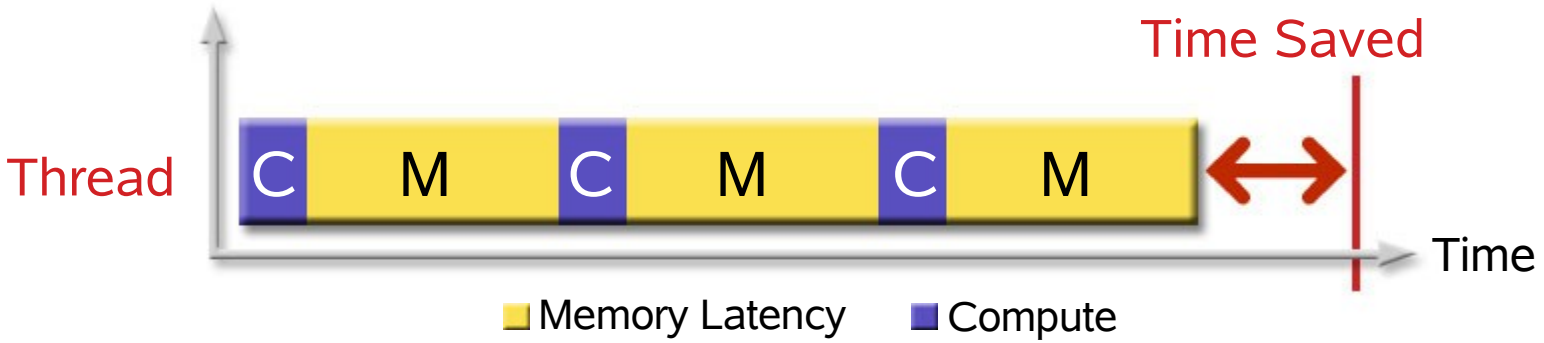

Processor optimized for ILP

Characteristics of Commercial Work Load

- **High degree of thread level parallelism (TLP)**
- **Large working sets result in poor locality of reference leading to high cache miss rates**
- **There is significant data sharing among threads resulting in coherence misses**
- **There is low instruction level parallelism (ILP) due to high cache miss rates, difficult to predict branches etc...**
- **Performance is bottle necked by stalls on memory access**

Sun Solution

NIAGARA

Chip Multi Threaded Processor

Niagara(T1)

- Uses CPU threads to exploit TLP
  - Memory and Pipeline stall times are hidden due to multiple threads
  - Shared L2 cache allows efficient data sharing between threads
- Memory system is designed for high throughput
  - High bandwidth interface to L2 cache for L1 misses
  - Highly associative L2 cache
  - High bandwidth interface to DRAM

Designed for Performance and Efficiency

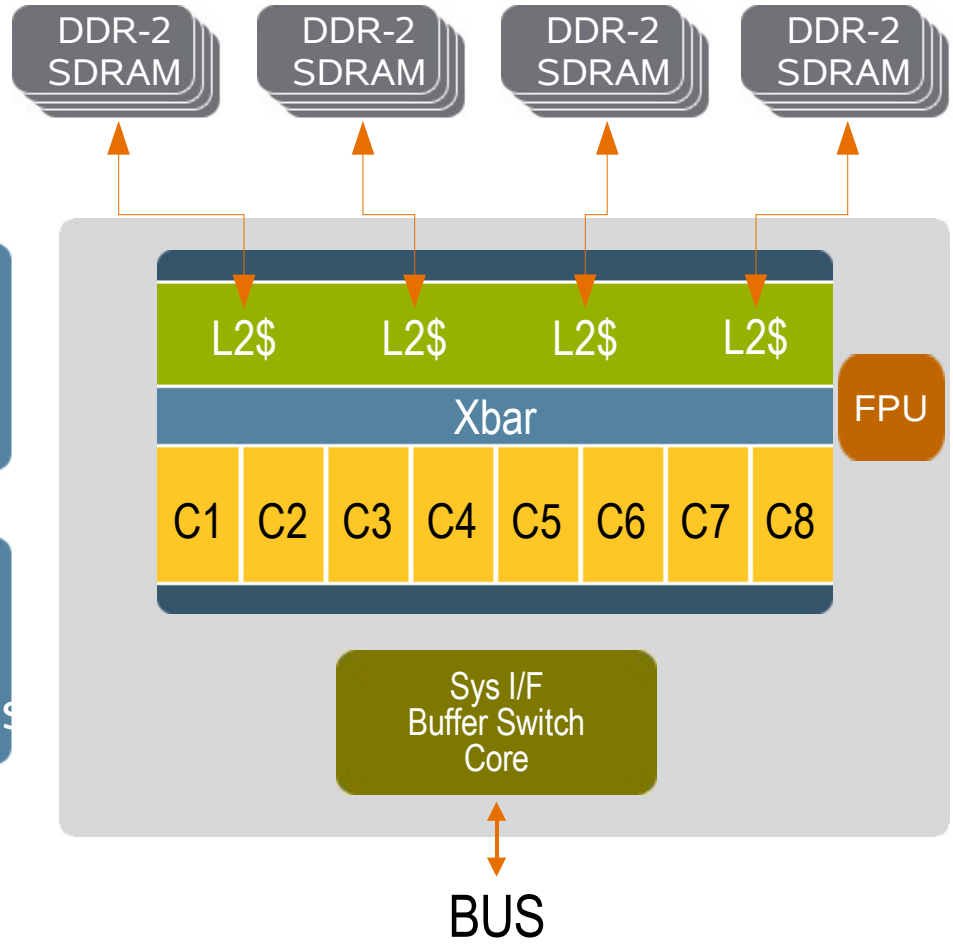

Dedicated Integrated Memory Controllers

Integrated Internal Communications

On-Chip Simplicity Means No Wait Latency

Clean Sheet Design Delivers Highest Performance, Efficiency

Niagara Specs

- Up to 32 threads, 8 cores
- Unique L1$ 16KB-I, 8KB-D per core
- Shared L2$ 3MB, 134GB/s, 12 way associative
- Radically changed cache coherency processing
- 4XDDR2 Mem on CHIP Controllers 23GB/sec
- Upto 128 GB memory
- SSL support - 7X the RSA throughput of Xeon
- Requires about 70 Watts
- Each thread requires just about 2.0 watts
- No Recompilation required

Thread Selection Policy

- CPU switches between available threads every cycle giving priority to least recently executed thread
- Threads become unavailable due to:
  - Long latency ops: loads, branch, mul, div
  - Pipeline stalls such as cache misses, traps, and resource conflicts
- Loads are speculated as cache hits, and the thread is switched in with lower priority.

Multithreaded Process on Niagara

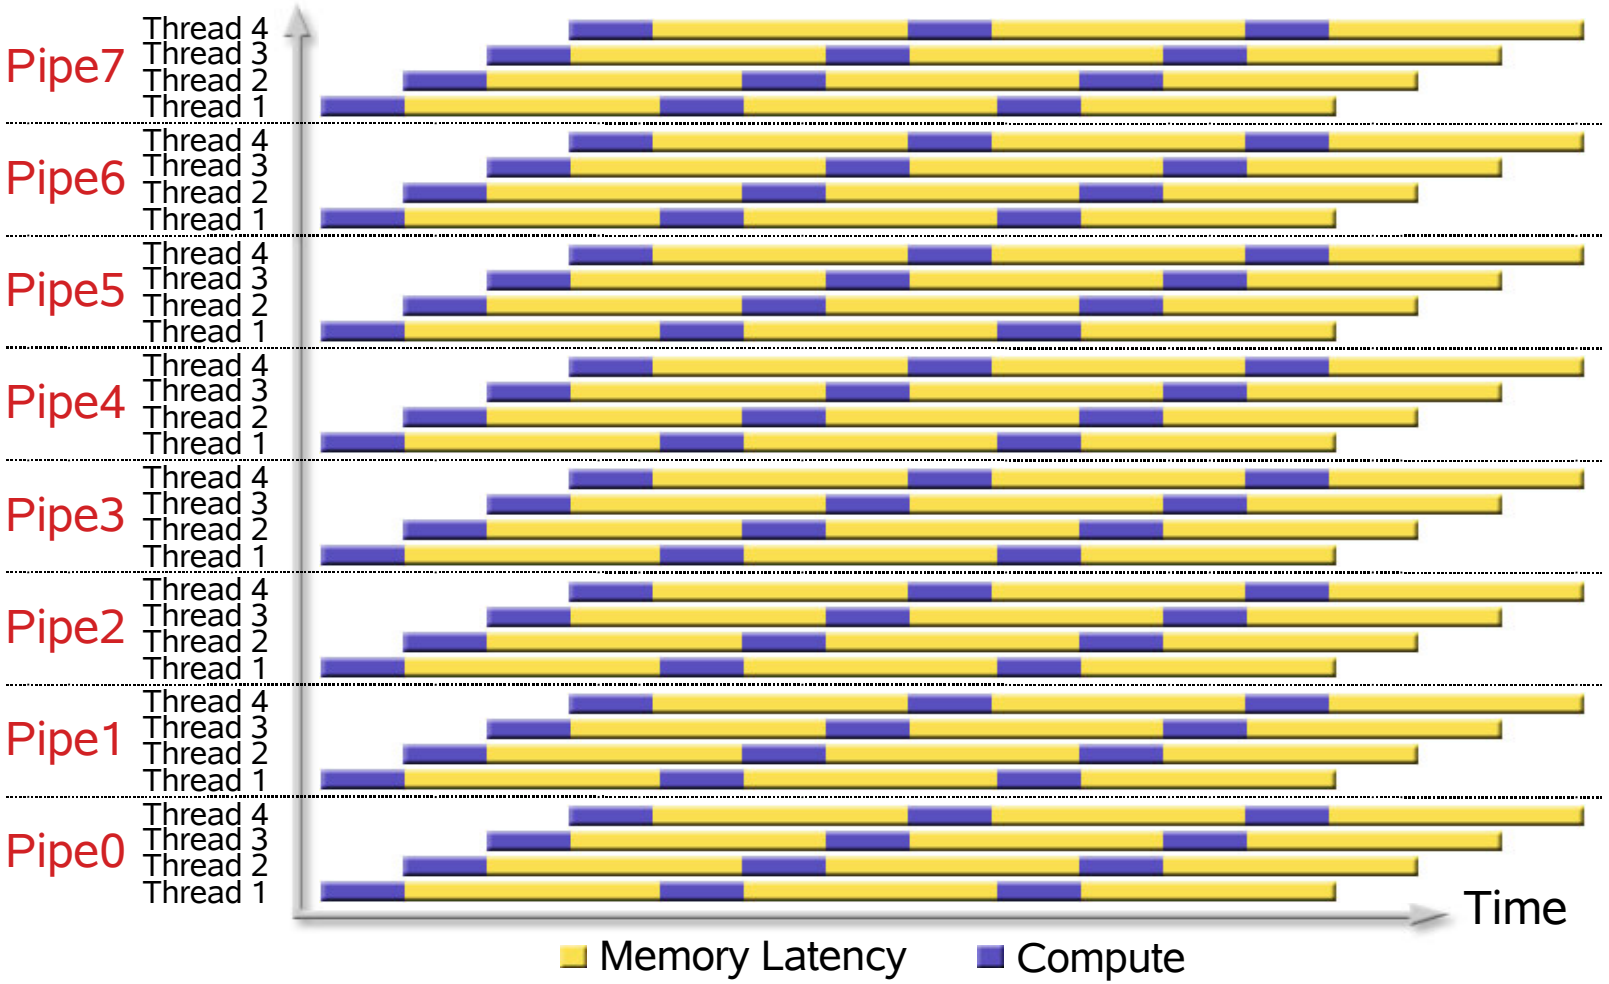

Larger number of Memory References outstanding from overlapping h/w threads leads to higher throughput

SWaP (Space, Watts and Perf)

Sun FireT2000
SWaP Rating = 30.4

Performance: 19,000 Users(1)

Space: 2RU x Watts: 312 = SWaP: 30.4

Performance/(Space*Watts) = SWaP Rating

1. LotusR6iNotes

Sun Fire T1000 Crushes Xeon and p5+

Sun Fire T1000 vs.

Dell SC1425
SPECjbb2005

IBM p5+ 520
SPECjbb2005

Performance	2.1X	1.6X
Power Usage	1/2	1/2
Space	Same	1/4
SWaP	4.4X	14X

Niagara-2 (T2): True System on a Chip

- Better performance than Niagara-1
- Up to 8 Cores
- Up to 64 threads per CPU
- Same power envelope as T1
- On chip NIC's
- And much more that I can not state

Performance Characteristics of T1

Positive Characteristics

- If a strand is stalled, its cycles can be utilized by other threads
- Multiple threads running the same application benefit by sharing text and data in L2 cache
- These characteristics make CMT ideal for throughput computing.

Not so Positive Characteristics

- If one thread is thrashing the L1 instruction cache, data cache, or TLB's on a core, it can adversely affect other threads on that core.
- If all threads run on the same core they are only getting one-quarter of the CPU time.
- So CMT is not ideal for real time applications.

Scaling issues to be aware of

- Hot locks are the most common reason applications fail to scale on CMT processors
- Tuning Critical Sections
- Apply more threads as CMT is a thread rich environment.

Server Virtualization

Benefits of Virtualization

- Virtualization is masking and sharing of server resources
- Results in
  - Server Consolidation
  - Higher server utilization
  - Increased operational efficiency
  - Improved manageability

CMT and Virtualization

- CMT provides hooks for server virtualization
- Each Strand can be a Virtual CPU
- Niagara-2 also provides support for Network Virtualization

Solaris Virtualization Solutions

- Containers (BSD Jails)
- Logical Domains (Individual OS Instance per domain)
- Xen

Logical Domains + Zones

- Partitioning capability
  - Create virtual machines each with sub-set of resources
  - Protection & Isolation using HW+firmware combination

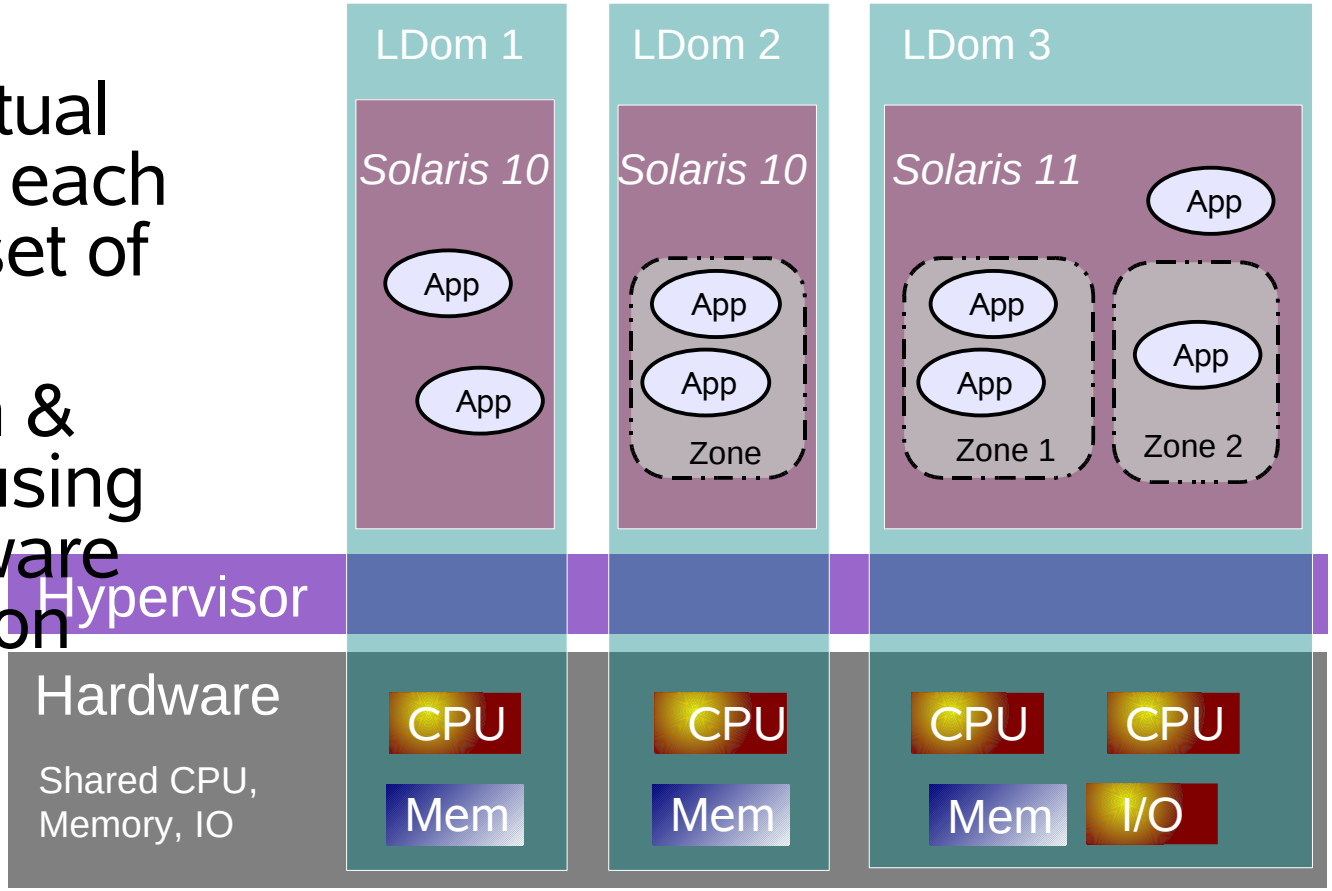

Network Virtualization

HW Based Network Virtualizarion

- Niagara-2 (T2) has on chip network interfaces
- Supports network virtualization/partitioning
  - Multiple Partitions can co-exist within a port
  - Only cable, MAC and RX FIFO's are shared.
- Virualization/Partitioning can be Based on
  - VLANS – upto 4K per port
  - MAC address – upto 16 per port
  - Service addresses (IP addresses, TCP/UDP ports) - upto 256 per device
- Interrupts for flow are sent to a particular CPU
- Full register sets are provided to control RX Rings

NIU RX Classification Model

Incoming flows are classified at layer 2, 3, or 4 and put into RX DMA channel according to classification rules that matched the flow.

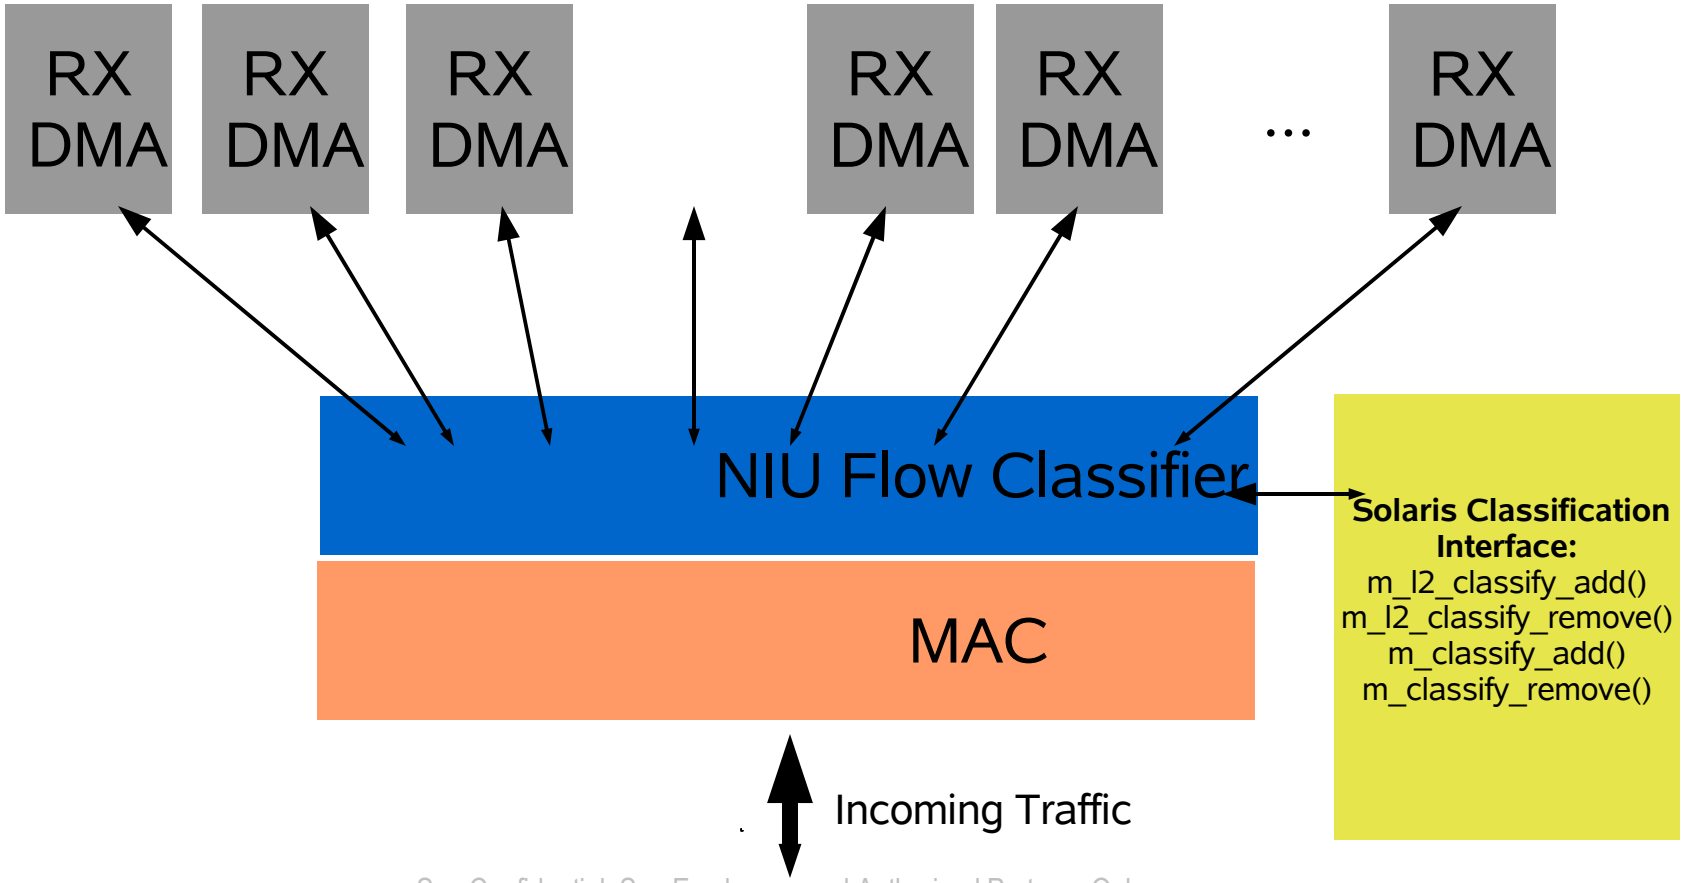

Software Based Network Virtualization

- Not All NIC's have HW support for Virtualization
- Software creates virtual stacks over 1Gb and 10Gb NIC's
- Virtual stacks are isolated from each other (for both resources and security purposes)
- Each Virtual stack can be tuned separately

Virtualized Networking

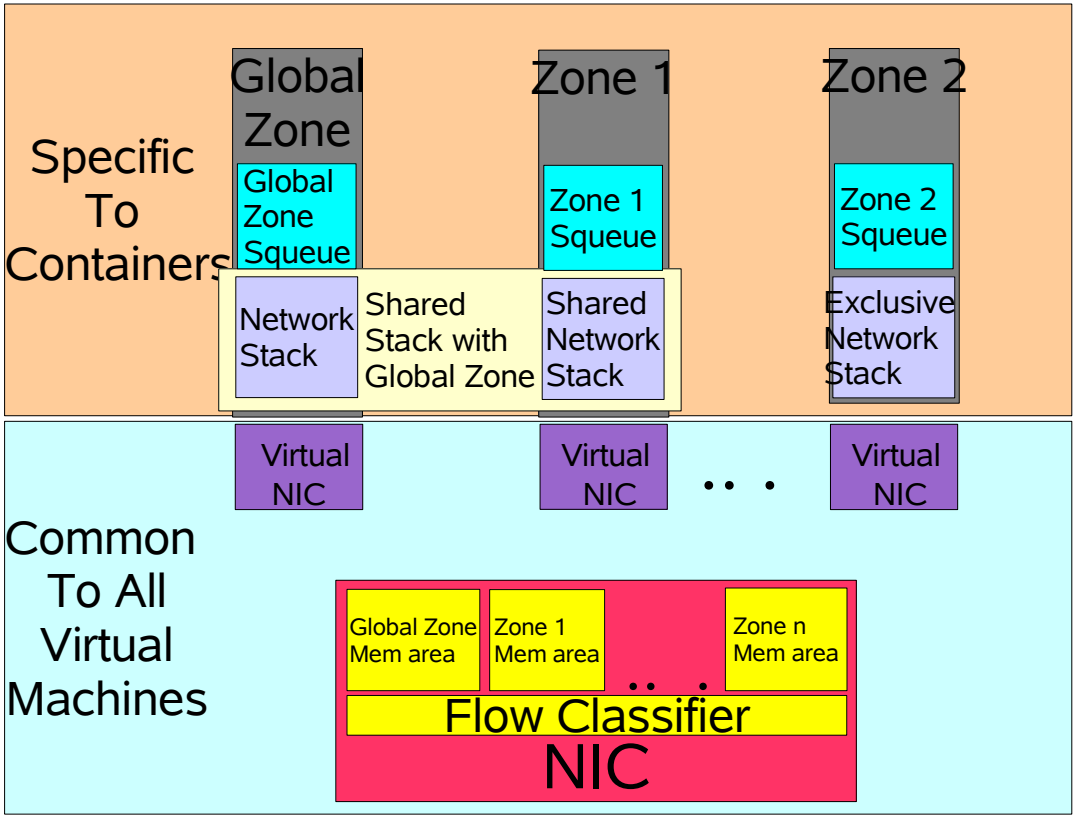

Virtual Network with XEN

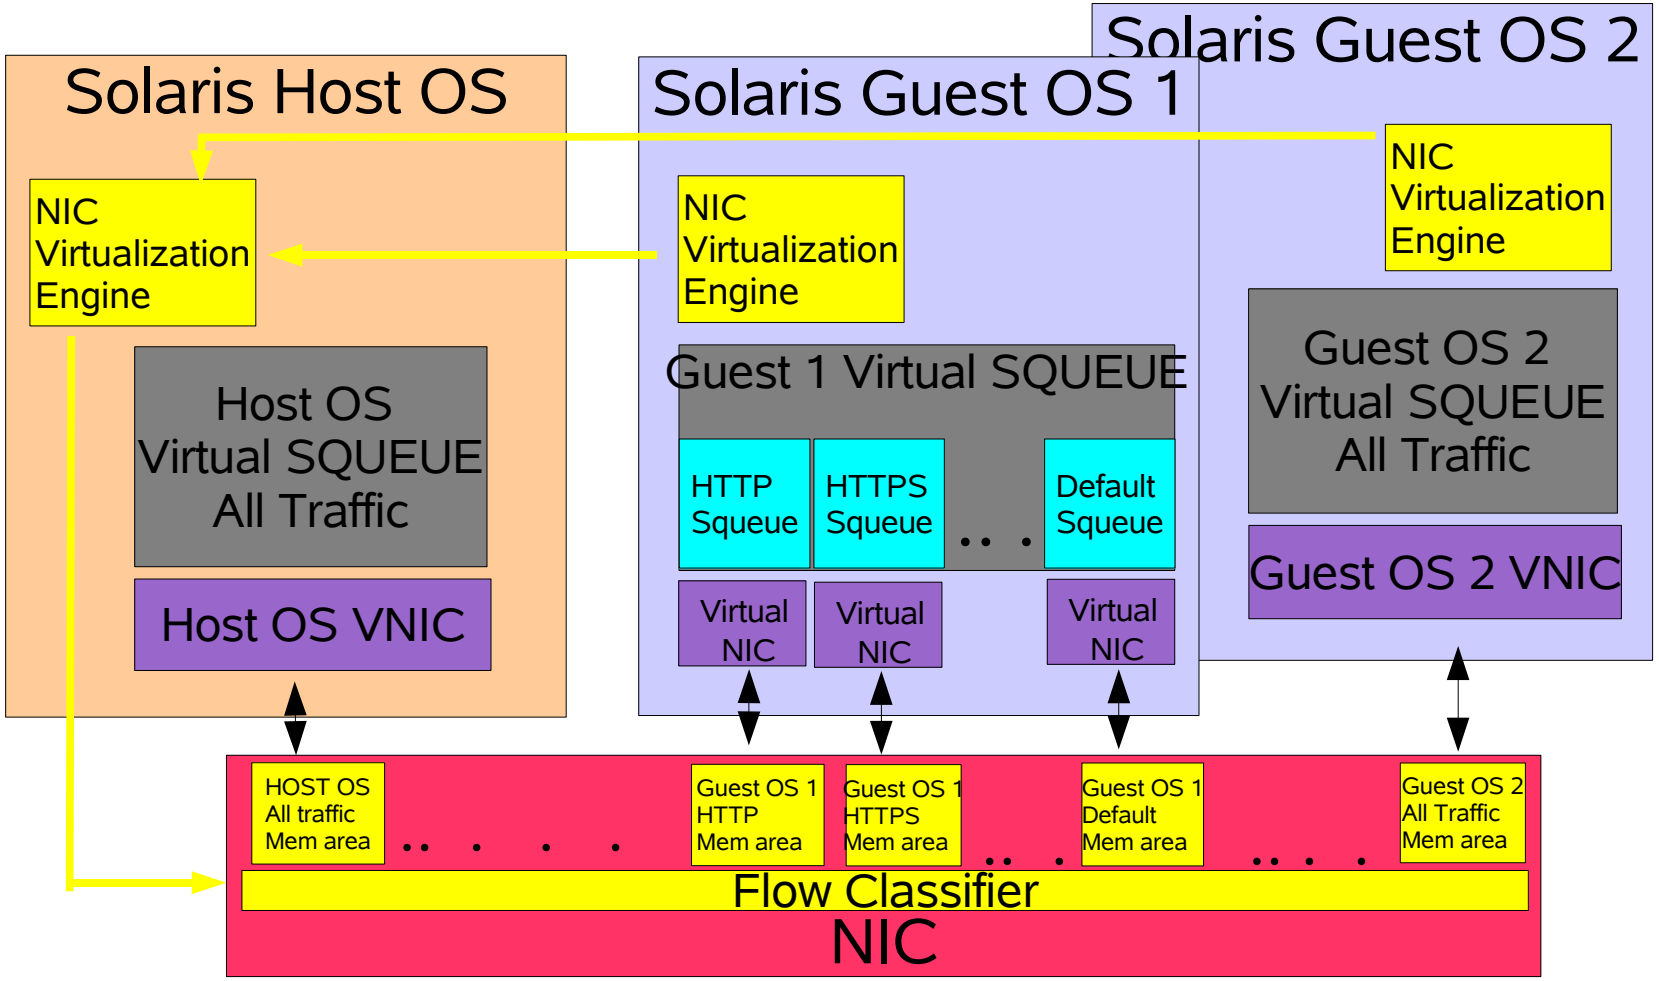

Future Work

- More work is needed to characterize different workloads on CMT processors and define best practices
- Open Interfaces are needed to implement Virtualization
- Network Bandwidth/Resource control support is needed in HW

References

- Various Sun internal and external documents and publications on Niagara

Niagara(T1)
A CMT PROCESSOR

Rao Shoaib
Solaris Core Technology group
rao.shoaib@sun.com