



# Continuum Security Processor Micro-Architecture Overview

Srinivas Mantripragada, Architect



**NETCONTINUUM**

# Outline

- Motivation
- Obstacles and Limitations
- Platform
- Continuum Security Processor (CSP)
- Compiler Technology
- Summary

# Motivation

- ❑ Today's enterprise server farms require high performance equipment to perform fast TCP and HTTP protocol processing.
- ❑ Security features such as high performance SSL termination and intrusion prevention are fast becoming de-facto requirements.
- ❑ Equally important requirements are control-plane tasks, e.g. easy traffic management / classification capabilities combined with powerful policy management.

# Obstacles and Limitations

## ❑ Limitations on general purpose solutions:

### — **General purpose architectures:**

- Single-threaded with good instruction-level parallelism (ILP). Don't scale well for web applications with massive amounts of thread-level parallelism (TLP) present.
- Incur large context switch overheads, interrupt latencies.
- Insufficient memory and I/O bandwidth.

### — **General purpose protocol solutions:**

- Quickly suffer from poor code and data locality with increase in traffic rate.
- Software-only SSL solution approach is compute intensive with poor performance, e.g. 100x slower than hardware approaches.

## ❑ Limitations on network packet processors:

— Perform well for stateless applications. Pre-dedicated code memories way too small for stateful TCP and higher layer protocol processing.

— Multi-processor designs lack good support for inter-process synchronization.

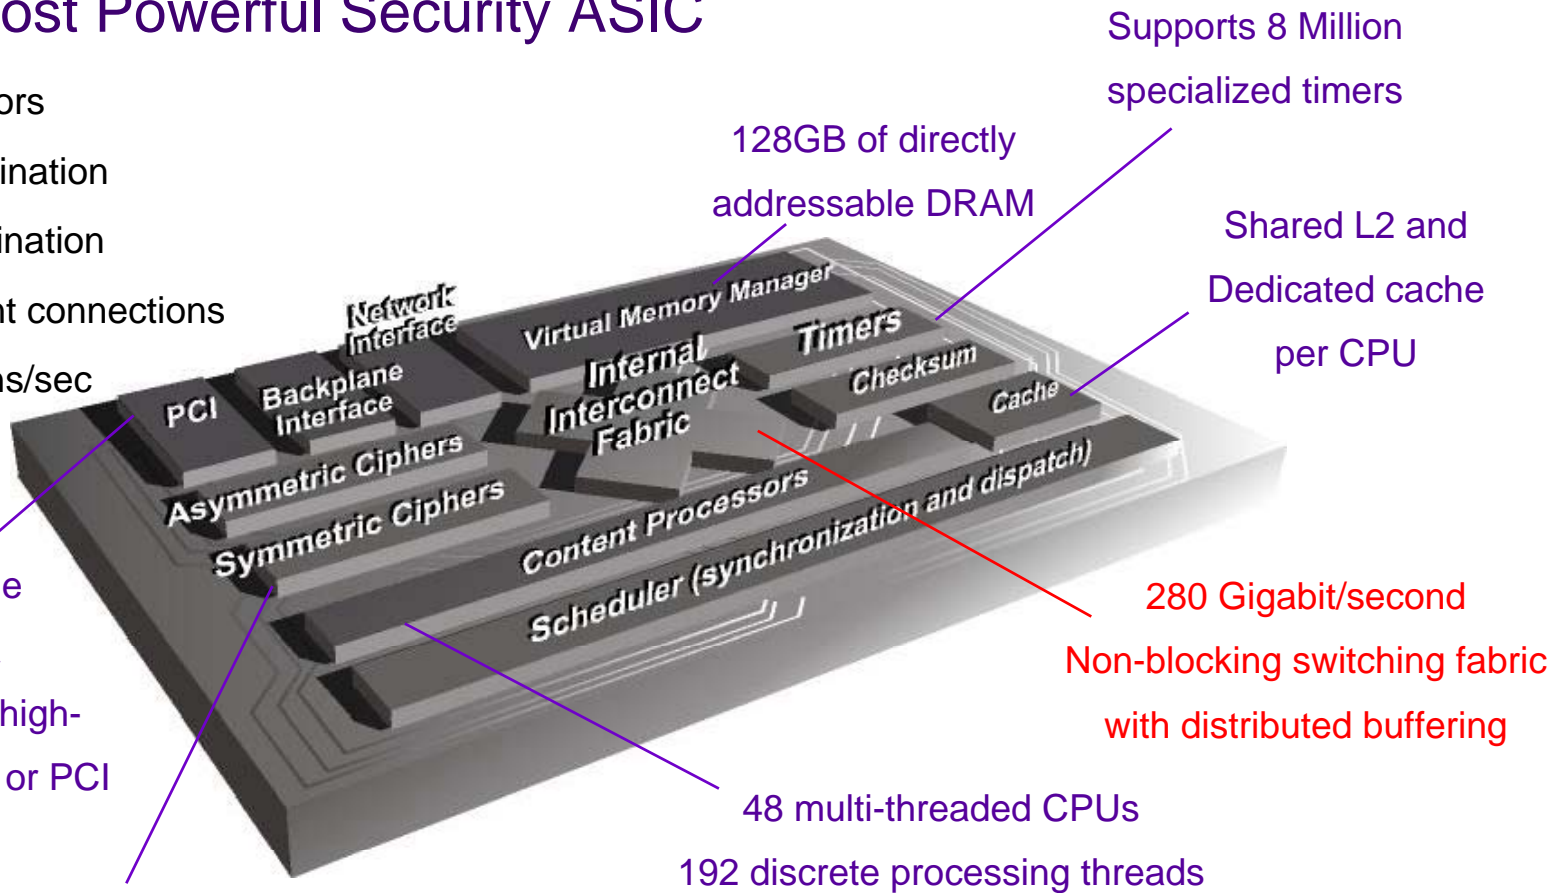
# Architectural Foundation

- ❑ Primary goal is to achieve high performance TCP, SSL and upper layer protocol connection rate.
- ❑ A new architecture which allows high extraction of TLP and ILP. Multithreading and multiprocessing single chip solutions seemed obvious choice.
  - **Deep multithreading to hide intra-thread memory latencies.**
  - **Extensive multiprocessing to facilitate large amounts of thread parallelism.**
- ❑ Efficient shared cache design to achieve large amounts of scatter-gather thread parallelism with relatively low inter-process synchronization and data-sharing delays.
- ❑ Good software (OS, Compiler) support to efficiently map software threads to hardware.
- ❑ Security and key TCP features done in hardware.

# Platform

## World's Most Powerful Security ASIC

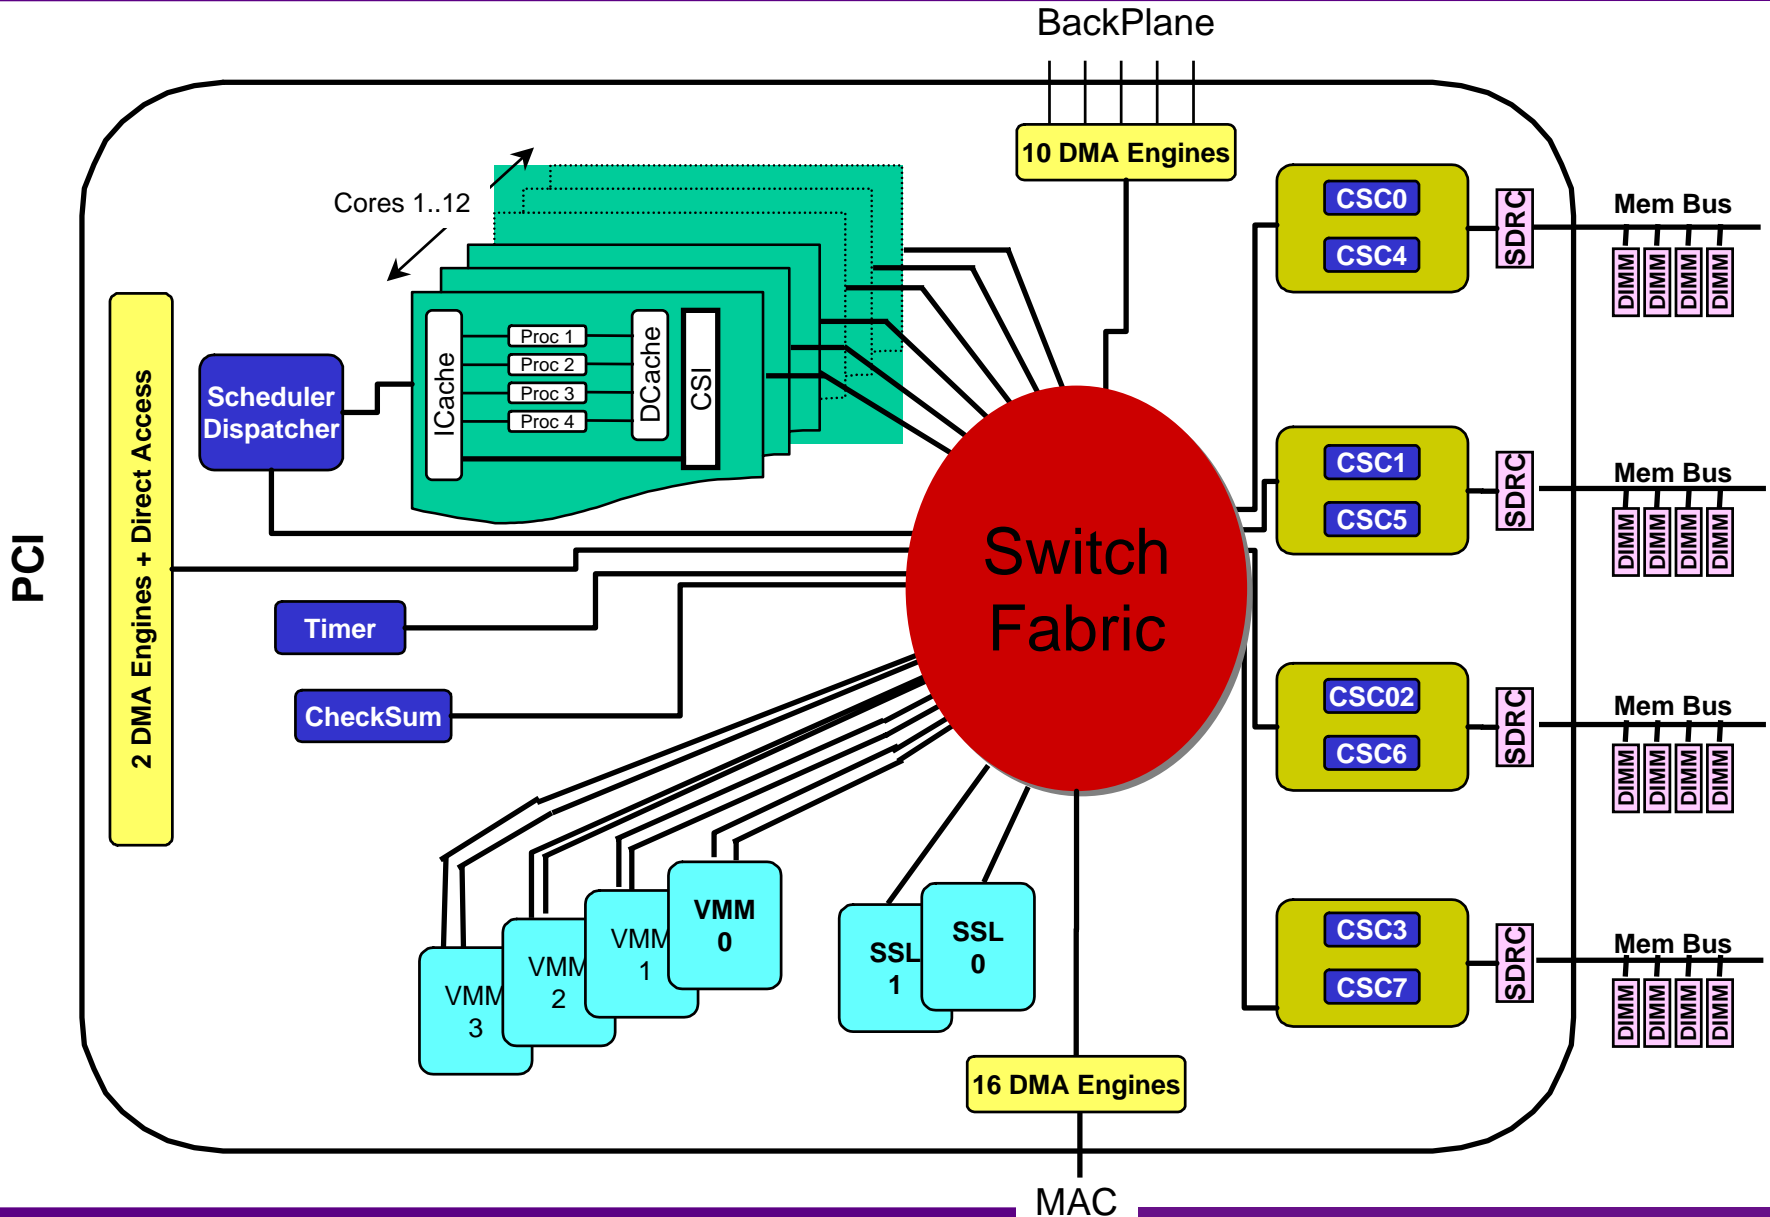
- 60M+ Transistors
- Full TCP Termination
- Full SSL Termination
- 1M+ concurrent connections
- 6,000 SSL trans/sec



Highly extensible interconnect via patent-pending high-speed interface or PCI

SSL, TLS, RSA  
3DES, RC4, SHA-1, MD5  
Random Number Generation

# Continuum Security Processor

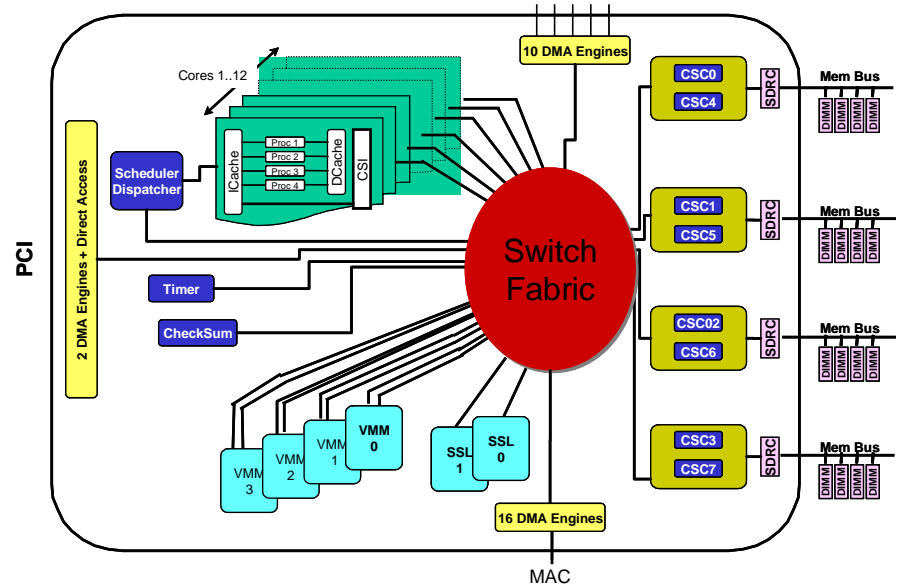


# ISA

- ❑ Variable length instructions, 1,2 or 3-byte. Average: ~1.8 bytes.
- ❑ Accumulator-based architecture.
- ❑ Both direct or segmented-based addressing schemes allowed:
  - **Allows indexing up to 63-bits of virtual address space.**
  - **Segment registers also defines the address space attributes, e.g. cacheability, write-thru, inexact etc.**

# Content Processors

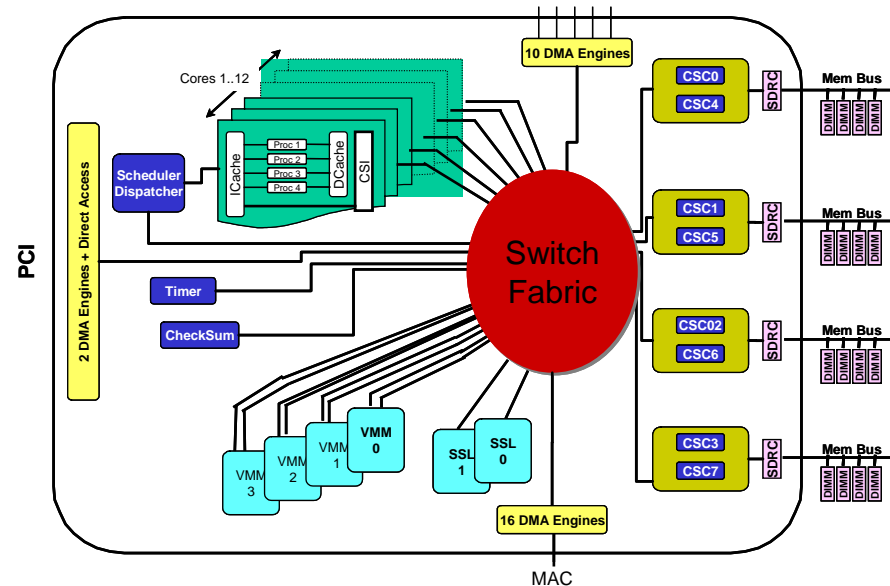
- ❑ 48 multithreaded processors, total 192 threads.
- ❑ Each thread has dedicated GP registers (8), segment registers (4) and private stack area.
- ❑ Execution resources shared by multiple threads, thread-swapping rules include:
  - **Taken branches and call instructions cause a thread-swap.**
  - **Load requests which miss in DCache causes a thread-swap.**
- ❑ All hazards, e.g. register/load/thread hazards resolved by compiler at static time.



# Instruction and Data Cache

## ❑ Instruction Cache

- Fully associative, 64-bytes/ line.
  - Shared by 16 threads among 4 processors.
  - Thread-arbitration engine round-robins the Ifetch requests coming from 4 processors.
  - Each processor does its own inter-thread arbitration.
- Supply 8 instruction bytes to decoder per cycle, for current active thread.
- A separate 8-byte prefetch buffer.

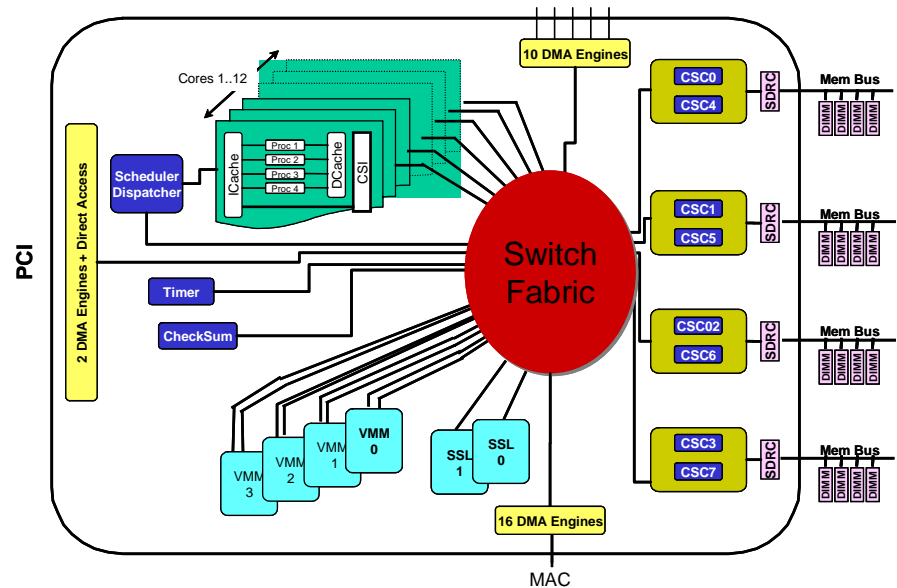


## ❑ Data Cache

- 4-way set associative, 16-bytes/set.
  - Shared by 16 threads among 4 processors.
  - Thread-arbitration engine round-robins the Dfetch requests coming from 4 processors.
  - Each processor does its own inter-thread arbitration.
- Multiple Dfetch miss requests to same address are coalesced.

# L2 Cache

- ❑ 8-way set-associative, 64 bytes/line.
  - 8 cache blocks total, 2 cache blocks per DRAM controller.
  - Shared by all clients, e.g. Processors, MAC, Scheduler etc.
  - Shared by both ICache and DCache.
- ❑ Non-blocking cache.
- ❑ Aggregate cache peak bandwidth, up to 38.4 Gbps.
- ❑ LRU replacement policy.
- ❑ Consecutive 64-byte addresses color to different banks to minimize unnecessary conflicts.
- ❑ Application-specific primitives supported, e.g. cache-line push/clear, atomic increment/decrement on 2-byte aligned addresses, etc.



# Security Features

- ❑ High performance public key processor
  - **6000 1024-bit (RSA, Diffie-Hellman, DSA) transactions per second.**
  - **High speed multiplier block controlled by dedicated firmware code.**
- ❑ Integrated symmetric key processor
  - **DES, 3DES, MD5, RC4, SHA-1**
- ❑ Concurrent public key and symmetric key processing.
- ❑ Built-in random number generator.

# Specialty Hardware

- ❑ Affinity-based hardware scheduler for efficient mapping of software threads to processors.
- ❑ Checksum block for fast TCP/IP header and payload verification (up to 4 Gbps).
- ❑ Timer block supports up to 8 million independent timer counters.
- ❑ High speed non-blocking switch fabric ties together all CSP components with aggregate bandwidth, up to 280 Gbps.
- ❑ MAC, Backplane and PCI support.

# Compiler Technology

- ❑ Industry-leading set of performance optimizations.
- ❑ GCC 3.0 front-end, CSP 2.0 backend.
- ❑ Major components:
  - **Global optimizer**
  - **CSP- specific code generator**
  - **Feedback design and hot-cold optimizations**
  - **Inter-procedural analysis and code layout**
- ❑ Single-Static Assignment (SSA) is used as intermediate representation.
- ❑ Global optimizations include:
  - **Global register allocation using SSA conflict graph.**
  - **Aggressive global scheduling before (and after) register allocation.**
  - **Tree-height reduction optimizations.**
  - **Control flow optimizations**
    - Block merging, tail-duplication, branch collapsing etc.

# Compiler Technology (contd.)

- ❑ Code generation improvements:
  - **Constant folding, dead-code removal.**
  - **Local and global common sub-expression elimination (CSE).**
  - **Iterative re-associative optimizations.**
  - **Automatic data-prefetching for fast linked-lists traversals.**
  - **Aggressive aliasing support for easy disambiguation of memory references.**
  - **CSP-specific ASM support to embed hand-tuned assembly in C.**

# Compiler Technology (contd.)

- ❑ **Feedback design and hot-cold optimizations:**
  - **Multi-thread aware insertion of instrumentation/profiling code.**
  - **Framework supports both path-profiling and call-profiling mechanisms.**
  - **Hot and cold paths are identified, CFG restructured to favor hot paths, cold paths from each procedure are then grouped into a separate cold section.**

# Compiler Technology (contd.)

- **Inter-procedural analysis and code layout:**
  - **Inter-procedural directed call-graph (IP-DCG) constructed using path and call profile data.**
  - **Initial set of hot call-nodes and call-chains identified using IP-DCG.**
  - **Code layout across procedure boundaries then performed selectively at cache-line boundary. Appropriate IP-DCG adjustments are made.**
  - **This reduces callee call overhead penalty by prefetching callee instructions before the actual call is made.**

# Implementation and Performance

## □ Implementation

- **64 M transistors**
- **.16uM CMOS**
- **150 MHz clock (shipping)**
  - Over-clocking to 200 MHz
- **1.5 V core, 3.3 V I/O**
- **10 Watts**
- **17x17 mm die size**
- **Q1/2002 production**

## □ Performance

- **1 Mil. Simultaneous TCP connections**
- **15,000 HTTP transactions**
- **6000 SSL sessions/sec**
- **Up to 4 Gbps bulk encryption rate**
- **Up to 5 Gbps of backplane interconnect traffic**

# Summary

- ❑ Continuum Security processor generates industry leading SSL performance, provides foundation for future NetContinuum family of security products.
  - **Integrated multi-processor and threaded system.**
  - **Fast security and cryptographic processing in hardware.**
  - **System architecture inherently scalable to support large amounts of ILP and TLP, going forward.**
  - **Fast back-plane channels to support multi-CSP systems.**
- ❑ Tightly coupled software and hardware
  - **Highly optimized OS, Compiler.**

Questions ?