

WEB SEMINAR:

Intelligent Multi-Core for Intelligent Networks

Kerry Johnson, QNX Software Systems
Steve Kilnger, Cavium



QNX and Cavium – PARTNERS in MULTI-CORE



- **Pioneers and leaders in multi-core technology**
 - QNX multi-core technology solutions initially introduced 1997 – offering includes robust OS and rich multi-core toolset
 - Cavium has pioneered and successfully introduced five generations of multi-core processors, used by all top-tier networking OEMs
- **Together, Cavium and QNX are leveraging multi-core technology and to offer integrated solutions for the communications market**
- **AVAILABLE NOW: early access QNX Neutrino RTOS 6.4.1 BSP on the OCTEON 57XX EVB**

- Demand Drivers
 - 10x increase in line rate in LAN (1G), WAN and Data Center
 - Integrated Layer3 to Layer7 security (VPN, SSL, IDS/IPS/AV)
 - Layer4 to Layer7 data processing for application and content-awareness
- Architectural Trends
 - Increased total device computation performance required
 - Integration and control and data plane
 - Application specific on-chip coprocessors required for packet processing and Layer4 to Layer7 data and security processing
 - Application layer and content-based QoS enforcement
 - Low power consumption

Multi-Core Well Suited to Handle Control Plane and Data Plane




- **Multi-Core processing capabilities:**
 - Run a full operating system, and/or
 - Run tuned data-plane-like code
 - Provide $n \times$ Core Control Plane Performance for SMP ready code
 - Provide high throughput Packet processing
- **Example OCTEON programming models:**
 - All cores run data-plane code
 - All cores run multiprocessor OS
 - 1-2 cores run OS, remaining cores run data-plane code
 - Hybrid OS threads running data-plane code


Data-plane only:




Operating system only:



Mixed:



Combining Five Key Functions into SOC



Traditional Approach

↑ Cost ↑ Power
 ↑ Area ↓ Performance
 ↓ Fragmented Software
 Results in Long TTM

OCTEON Solution

↓ Cost ↓ Power
 ↓ Area ↑ Performance
 ↑ Common Software
 Enables Rapid TTM

OCTEON Plus Multi-Core Family of Processors



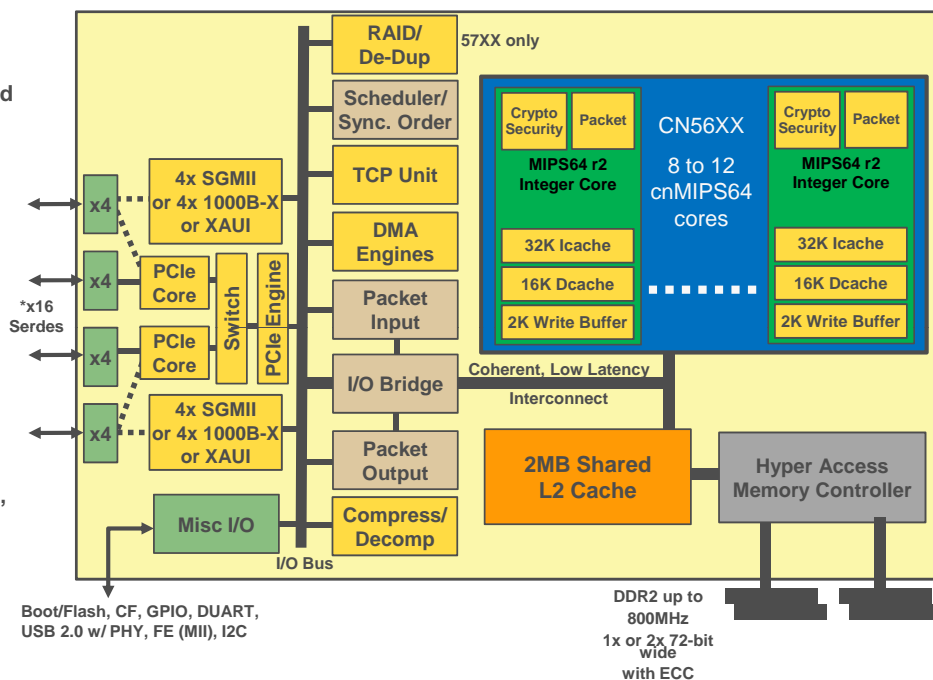
- 1 to 16 cnMIPS 64-bit cores
- 300MHz to 800MHz per core, 300MHz to 12.8GHz total CPU compute performance
- 128KB to 2MB L2 Cache, 36-bit to 144-bit DDR2
- Most advanced hardware acceleration integrated
 - L2 to L7 processing
 - Load balancing and multi-core scaling
 - Compression/Decompression
 - Security
 - Pattern Matching
 - QoS
 - RAID and De-duplication
- Integrated networking and system I/O
 - GbE, 10GbE, PCI-X, PCI Express
- Seven families, all fully production-released
 - CN50XX, CN58XX (PCI/PCI-X, NON-SERDES I/O)
 - CN56/57XX, CN54/5XX, CN52XX (PCIe, SERDES-based I/O)
- Power consumption from 3W to 40W
- 564-BGA (27mm x 27mm) to 1521-BGA (40mm x 40mm)

OCTEON™ Plus CN56/57XX Product Family



Key Features

- 8 to 12 cnMIPS cores
 - MIPS64 R2 + enhanced instructions
 - 600 to 800MHz
- Up to 9.6GHz per chip
- Integrated Serial I/O
 - Up to 2 PCIe x8
 - Root Complex or Endpoint Mode
 - Up to 2 XAUI
 - Up to 8 SGMII
- Most Advanced HW acceleration including:
 - 15Gbps security, TCP, packet processing
 - 8GB/sec RAID
 - 10Gbps compression



I: Industrial temp versions available

Key OCTEON Processor Features



CPU Core, Cache, and Interconnect

Optimized cnMIPS 64-bit cores

- Fast, dual-issue cores with highest # cores/chip
- MIPS64 R2 + additional Cavium instructions for optimal performance

Low-Latency, High-Bandwidth Coherent Cache Subsystem

- Large, L2 with ECC and highly-associative L1 caches
- Flexible locking, partitioning and pre-fetching features

Scalable Hyper-connect

- Scalable, ultra high-bandwidth internal chip interconnect
- Minimizes latency and ensures linear performance scaling

Flexible, Standards-Based I/O

High-Bandwidth Memory Controllers

- Based on industry-standard DRAM technologies
- ECC-protected, optimized for low-cost commodity DIMMs

Advanced Network and System I/O

- Latest serial technologies for highest-performance
- Integrated GbE, 10GbE, channelized interfaces, PCI & PCI Express, USB2.0

Key Differentiated OCTEON Processor Features



Powerful Hardware Application Acceleration

Packet Input & Output Processing

- Flexible hardware packet parsing, checking, classification
- QoS-based processing and traffic scheduling

Application Acceleration Manager

- Automatically orders & load-balances work to be processed by the cores
- Eliminates need for SW 'locking' which kills performance

Comprehensive Security Acceleration

- Per-core engines with complete set of crypto algorithms
- Highest throughput for key security applications including IPsec, SSL

Pattern Matching/DPI

- Integrated HW pattern matching acceleration
- Highest performance for IDS/IPS, AV, Content Inspection

TCP Acceleration

- Hardware acceleration for key TCP functions
- Enables TCP termination at industry-highest throughput

Compression/Decompression RAID + De-dup

- Standards-based support for GZIP/PKZIP + variants
- Best compression ratios and highest throughput
- RAID 5/6 XOR and Galois Field acceleration

OCTEON/OCTEON Plus Multi-Core Families



CN38XX

400-600MHz
15 to 40W
512KB-1MB L2

CN58XX

600MHz-800MHz
15 to 40W
2MB L2

CN56/57XX

600MHz-800MHz
12 to 30W
2MB L2

CN54/55XX

500MHz-700MHz
10 to 20 W
1MB L2

CN52XX

500MHz-750MHz
6 to 13 W
512KB L2

Up to 8x RGMI, 2x SPI-4.2, PCI-X 64b/133MHz
1521-pin FCBGA

16
Core



12
Core



8
Core



4
Core



NSP, EXP, SCP
options

Up to 8x SGMII, 2xXAUI,
Up to 2xPCIe x8
1217-pin FCBGA

12
Core



10
Core



8
Core



NSP, CP (54/56XX)
SSP, SP (55/57XX)
options

Up to 4x SGMII, 1x XAUI
Up to 2x PCIe x8
1217-pin FCBGA

6
Core



4
Core



Up to 4x SGMII, 1x XAUI
Up to 2x PCIe x2
729-pin BGA

4
Core

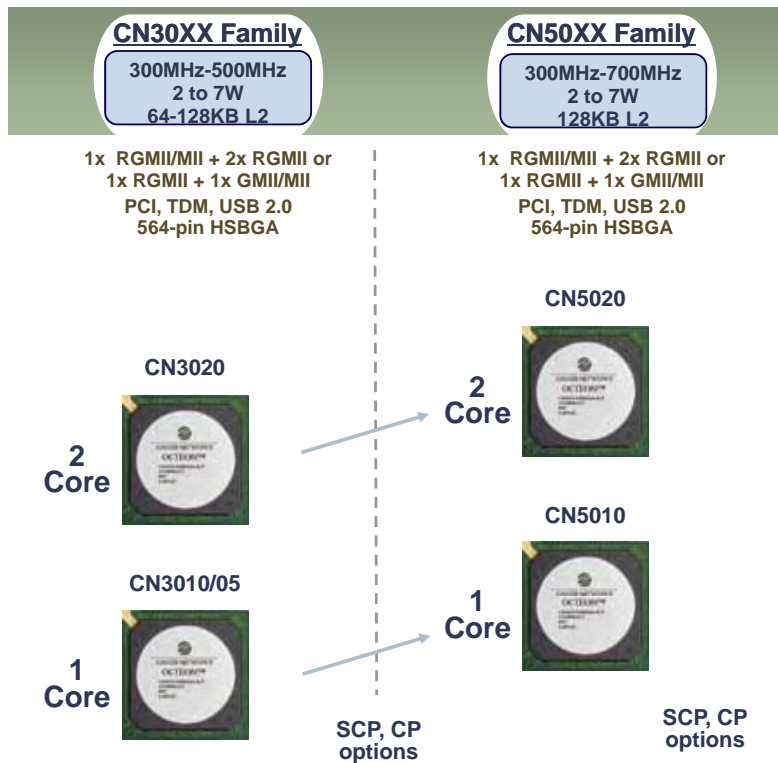


2
Core

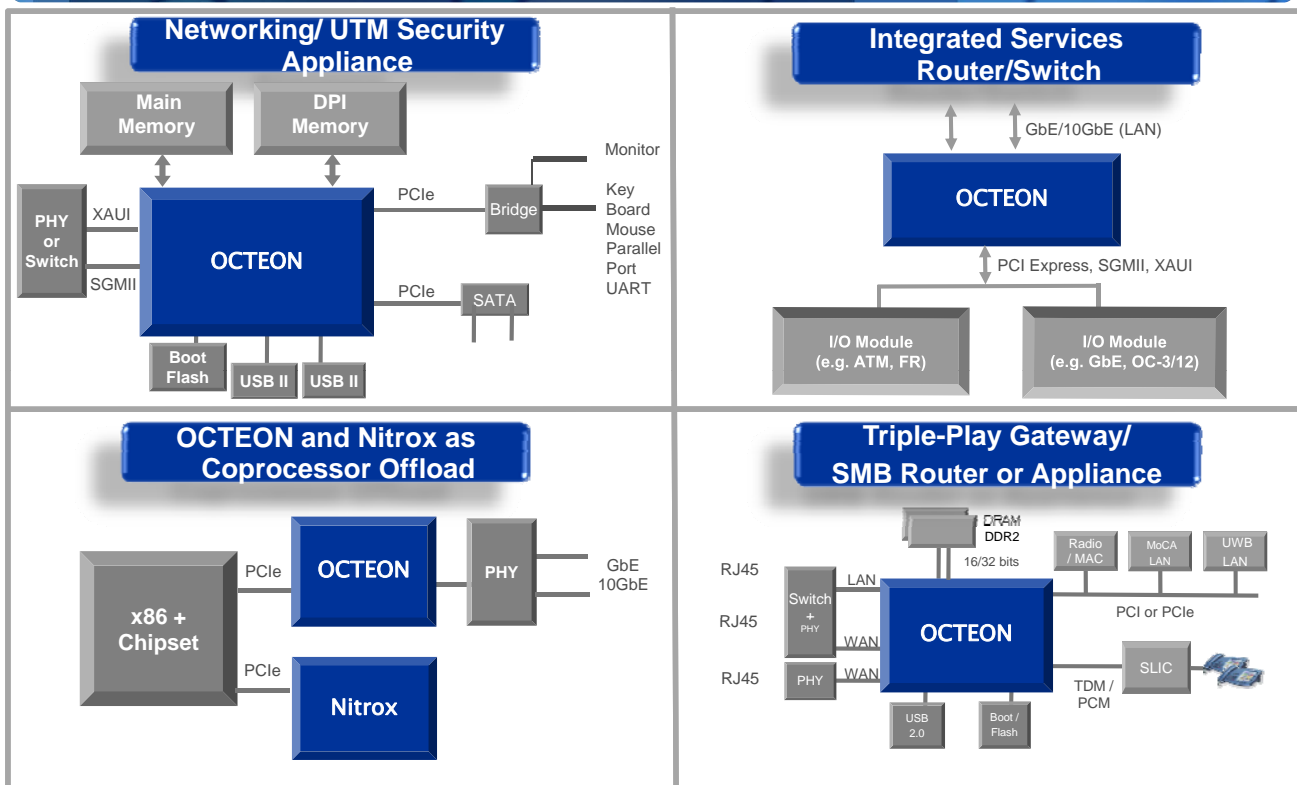


SCP, CP
options

OCTEON/OCTEON Plus Single and Dual Core Families



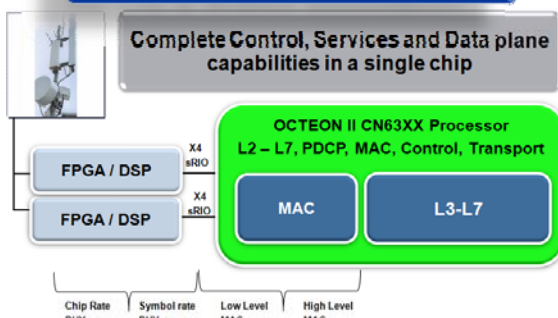
OCTEON Networking/Security Solution Examples



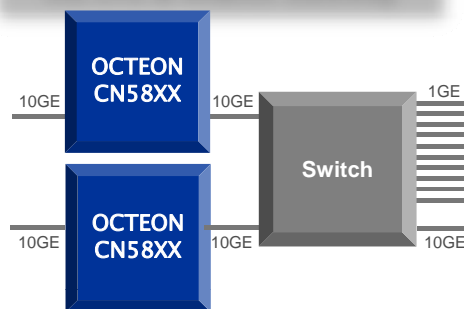
OCTEON Wireless Solution Examples




3G, LTE & WiMAX Basestation




3G, LTE & WiMAX Gateway



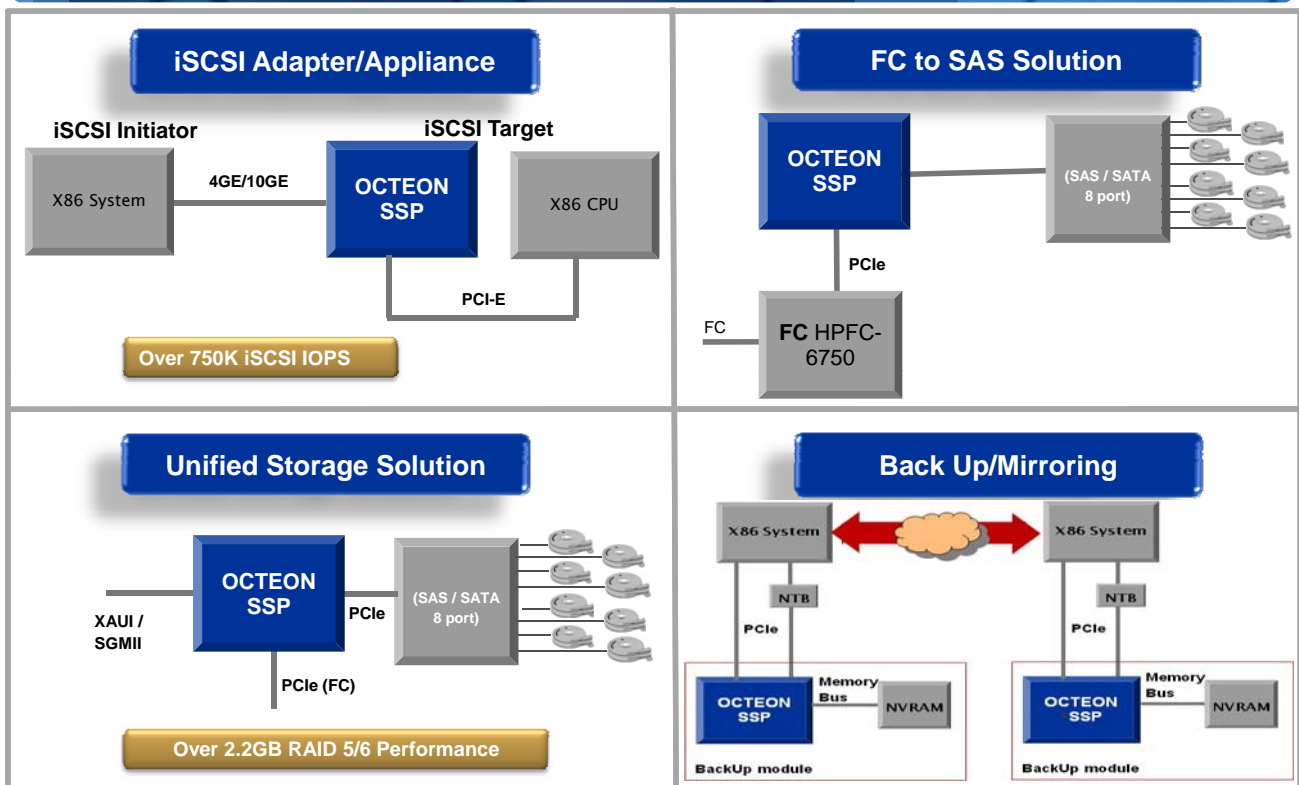
802.11n Multi-radio AP



WLAN Controller



OCTEON Storage Solution Examples

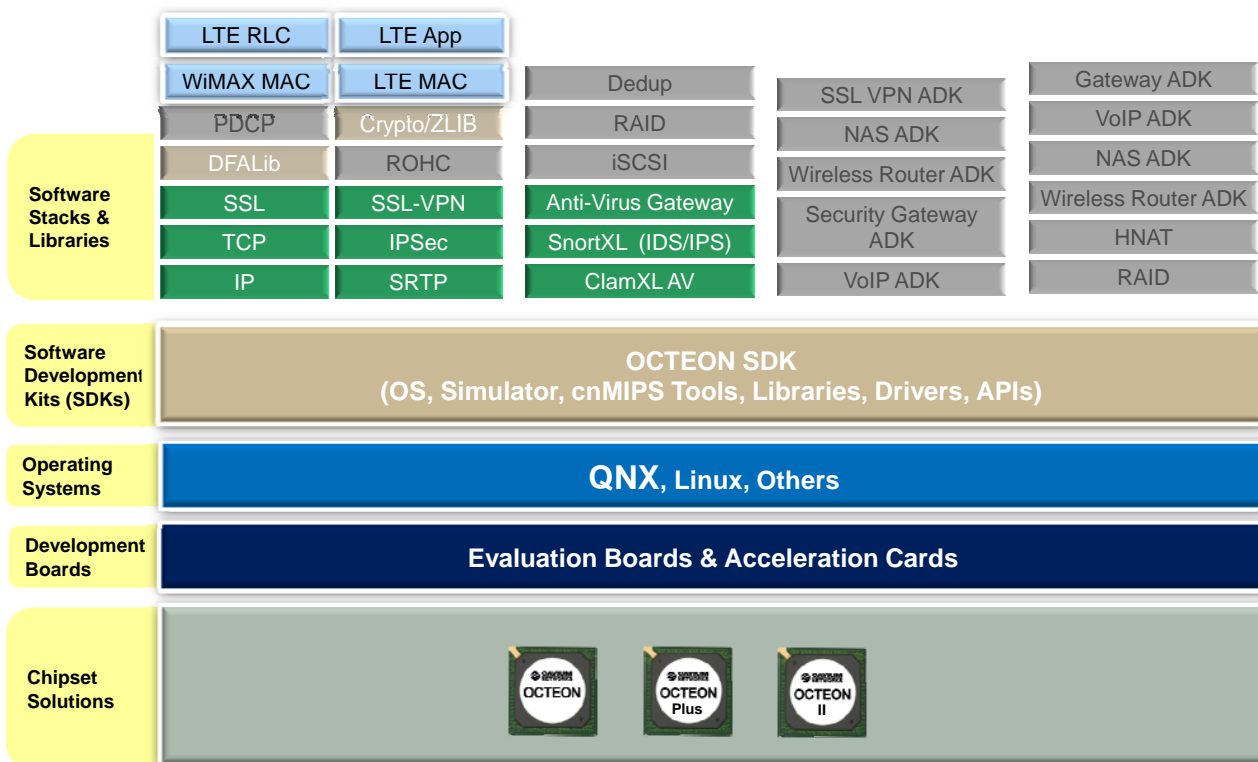


Software Standards Supported



- Industry standard MIPS32/MIPS64 instruction set architecture
 - ANSI C, C++
- Industry standard tools supported
 - GNU, DDD, EJTAG
- Support for all Key Operating Systems
- Support for BSD socket APIs, OpenSSL, IPv4 & IPv6
 - Defacto networking standards

Cavium Networks' Production Ready Suite of Development Tools & Software Stacks



3rd Party SW


International Company

16

Complexity → Capacity → Multi-core



**Remote
Network Management**
SNMP
TL1




Fault Management

Diagnostics
Health Monitoring
Alarm Detection &
Reporting



Control Loops



**Network Topology
and Restoration**

G-MPLS
CR-LDP
OSPF

Performance Monitoring

Packet Statistics
SONET/SDH



**Local User
Interface**

CLI
GUI

Multi-core history



- 12 years of symmetric multiprocessing in realtime embedded systems
 - QNX introduced SMP on dual socket systems in 1997
- QNX definition of multi-core
 - Symmetric multiprocessing on homogenous processors
 - Does not, *technically speaking*, require a multi-core processor at all - don't forget about multi-socket SMP systems
 - Combination of general purpose processor plus acceleration engines
- It's been done, you use it today likely without knowing it
- QNX SMP has been field deployed in thousands of systems already

Medical



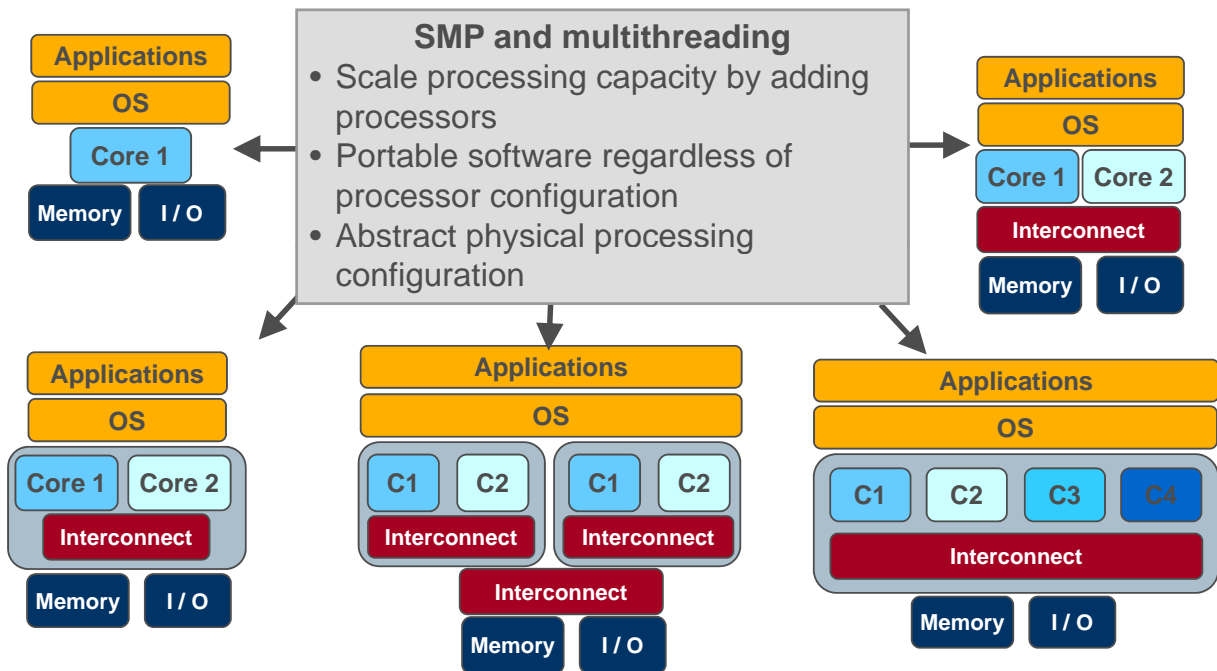
Networking



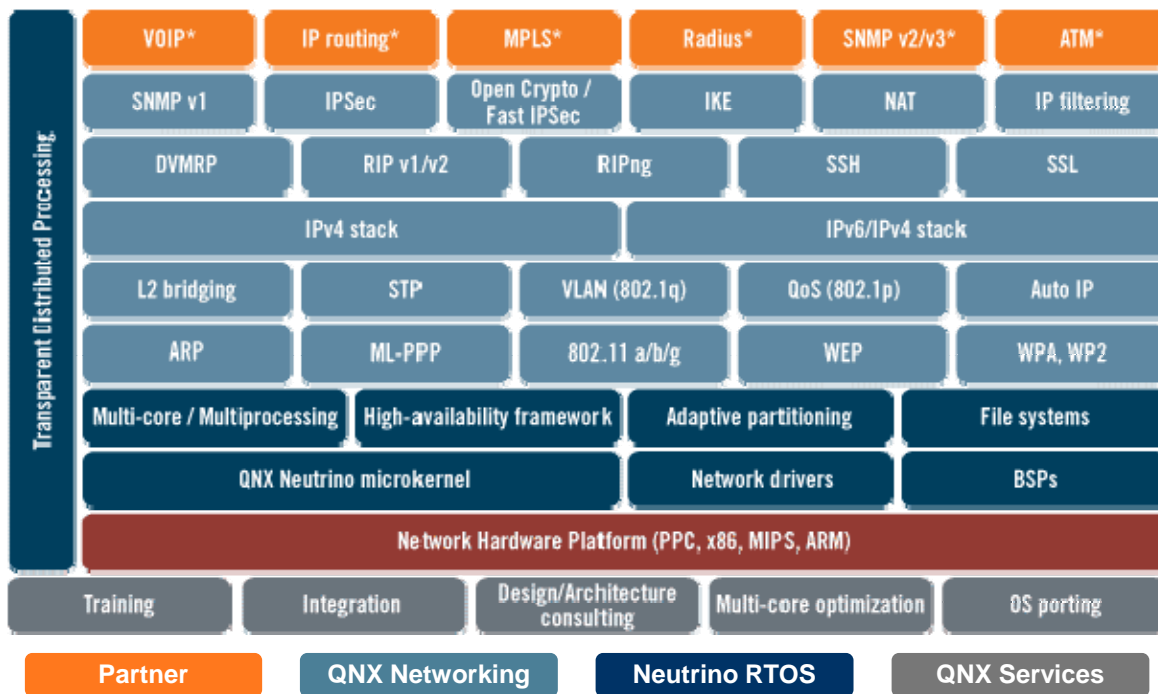
Industrial



- How many “one off”, unique programming models can you support?



QNX networking solution



Foundry27

The community portal for QNX software developers

A Harman International Company


20

Achieving High Availability



$$\text{Availability} = \frac{\text{Mean Time Between Failure}}{\text{Mean Time Between Failure} + \text{Mean Time To Repair}}$$

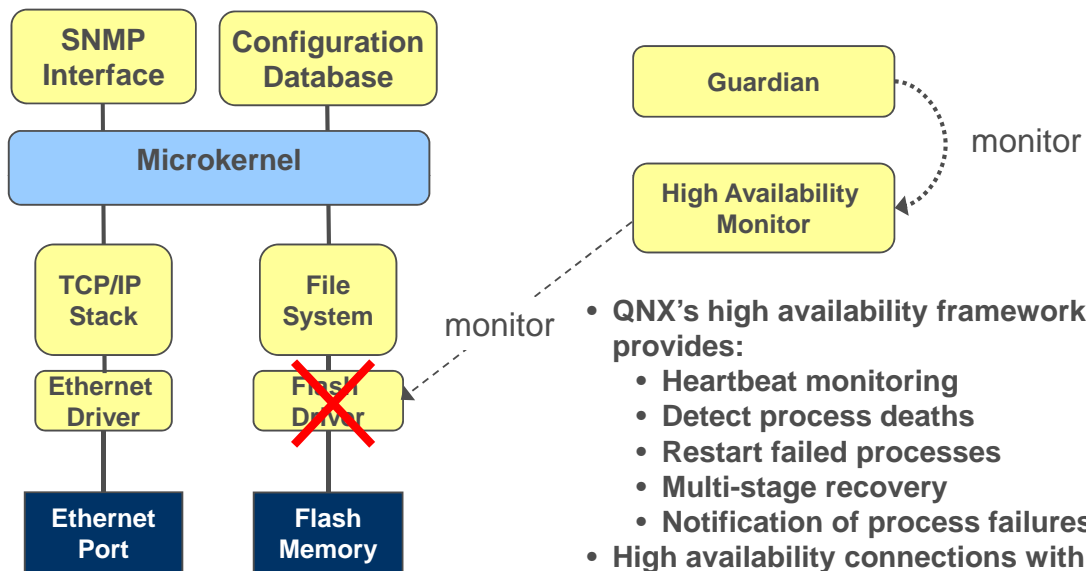
$$\text{Availability} = \frac{\text{Mean Time Between Failure}}{\text{Mean Time Between Failure} + \text{Mean Time To Repair}}$$



Achieving High Availability – microkernel + monitor



$$\text{Availability} = \frac{\text{Mean Time Between Failure}}{\text{Mean Time Between Failure} + \text{Mean Time To Repair}}$$



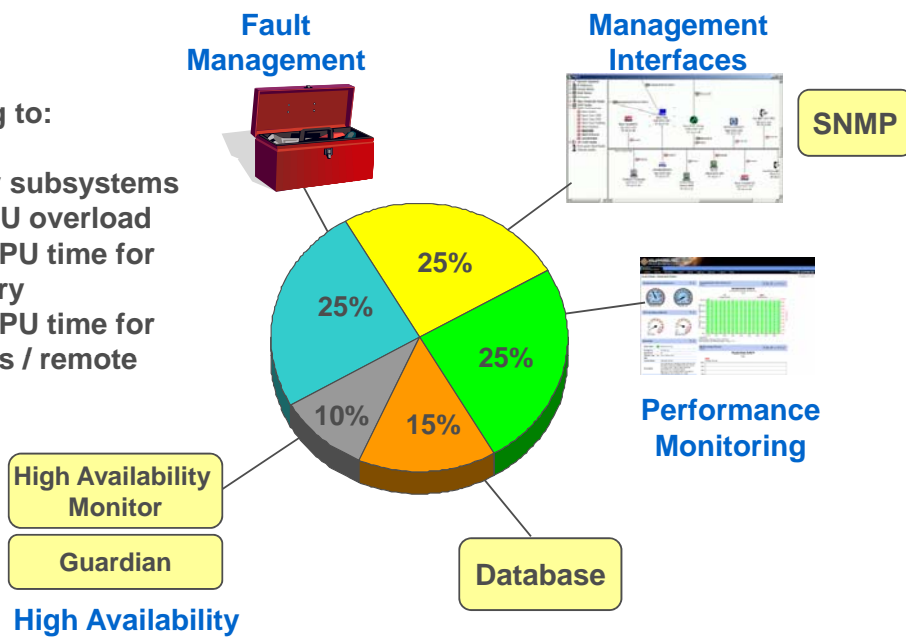
- QNX's high availability framework provides:
 - Heartbeat monitoring
 - Detect process deaths
 - Restart failed processes
 - Multi-stage recovery
 - Notification of process failures
- High availability connections with connection recovery procedures

Partitioning to increase availability



Add partitioning to:

- Contain faulty subsystems that cause CPU overload
- Guaranteed CPU time for failure recovery
- Guaranteed CPU time for user interfaces / remote monitoring

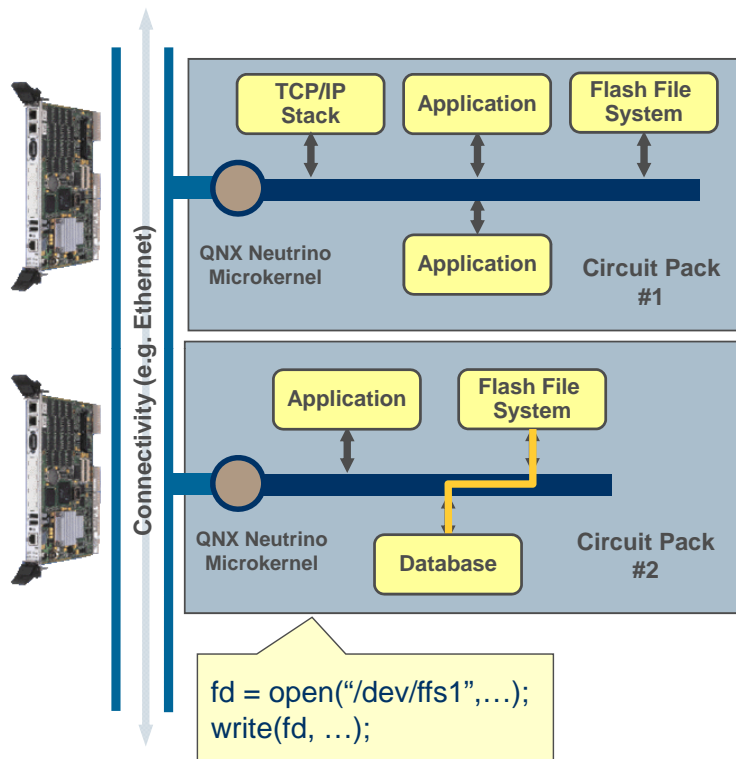


QNX's adaptive partitioning allows you to guarantee CPU time and still accommodate bursty processing needs

Achieving redundancy & scale with distributed processing

CAVIUM
NETWORKS

QNX
QNX SOFTWARE SYSTEMS

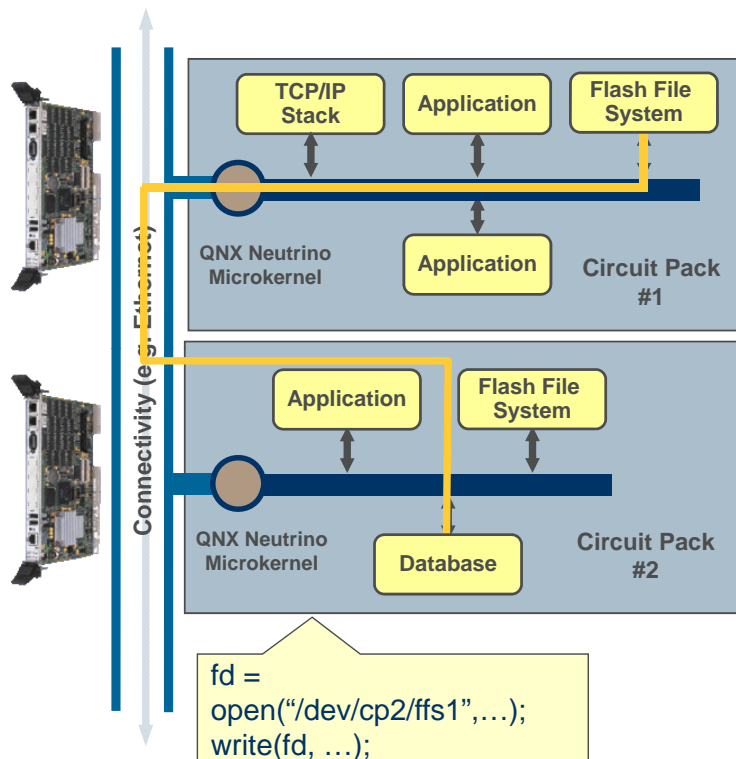


- QNX distributed processing (QNET) extends message passing over a transport layer
- Auto discovery of all nodes and published resources over Ethernet or other inter-connect
- Uses a global pathname space for networked resources
- Seamless sharing of I/O resources between nodes (e.g. use flash memory located on another node)
 - Networking stack
 - File systems
 - Hardware ports
- Applications / services can be built in a fully distributed manner without special code

Achieving redundancy & scale with distributed processing

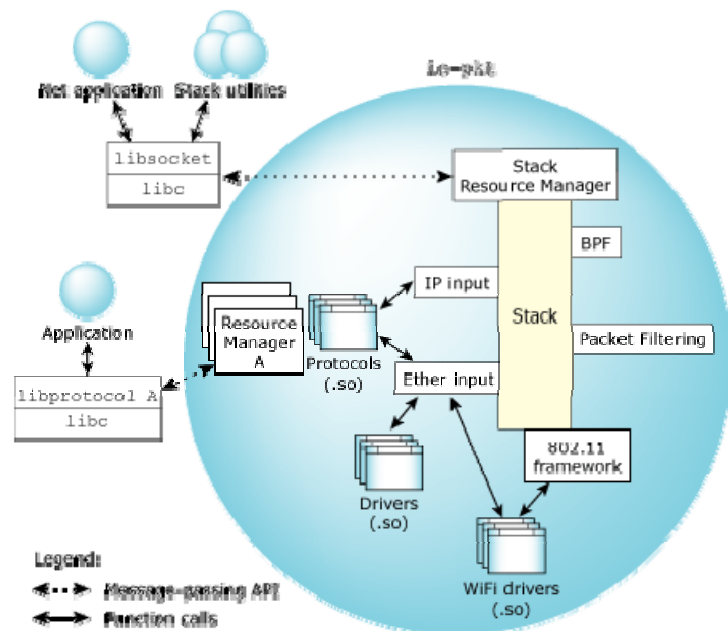
**CAVIUM
NETWORKS**

QNX
QNX SOFTWARE SYSTEMS



- QNX distributed processing (QNET) extends message passing over a transport layer
- Auto discovery of all nodes and published resources over Ethernet or other inter-connect
- Uses a global pathname space for networked resources
- Seamless sharing of I/O resources between nodes (e.g. use flash memory located on another node)
 - Networking stack
 - File systems
 - Hardware ports
- Applications / services can be built in a fully distributed manner without special code

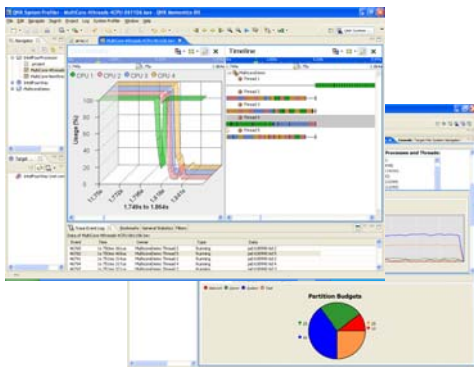
- Multi-threaded at layer 2 for SMP operation
- BSD based for driver, tooling and utility portability
- User space implementation, multiple stack instances (e.g. one per interface)
- Hardware crypto acceleration
 - BSD OpenSSL framework
- Berkley Packet Filter
 - Rules based filtering in user space
- PF Interface
 - Inspect/modify packets
 - Operates as loadable module in context of stack



The QNX software development platform



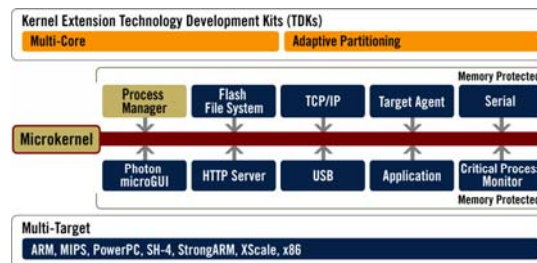
QNX® Momentics® development suite



Debug
Analyze
Optimize

Develop
Deploy


QNX® Neutrino® RTOS



Visual system analysis



QNX Momentics
on development host

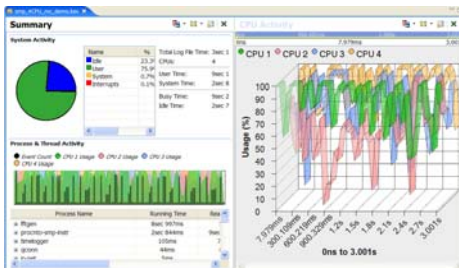


Capture Kernel Event Trace
Upload Trace File

Target system running
QNX Neutrino RTOS
instrumented kernel



Display
Trace
Information



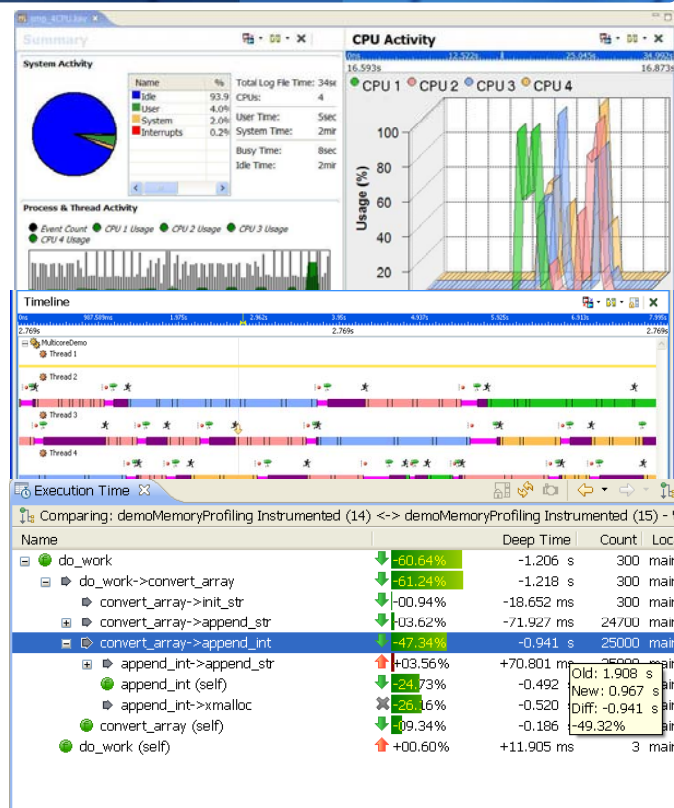
System Profiler

- Single step to capture, upload and view system trace
- Quickly visualize system interaction and behavior
- View interrupts, thread states, event timing, CPU usage, partitions, IPC, and much more....

Multi-core visualization tools



- Multiprocessor systems require more finesse than “printf” debugging
- CPU utilization across all cores to assess load balancing
- System analysis showing all processes, all cores
 - What is running where and when
 - Blocking analysis to help with shared resources
- Find slow spots using function timing – can help pinpoint opportunities to make parallel



Summary



- QNX and Cavium are providing a high performance solution for intelligent networks
- Cavium OCTEON and QNX Neutrino RTOS for multi-core control plane and data plane applications
- QNX Momentics tool suite for multi-core optimization
- **CONTACT INFO**
Kerry Johnson
Product Manager, QNX Software Systems
kjohnson@qnx.com