RSA CONFERENCE
CHINA 2012
RSA信息安全大会2012

THE GREAT CIPHER

MIGHTIER THAN THE SWORD

伟大的密码胜于利剑

Maximizing Multicore Technology for Network Intelligence and Security

Fu Lizheng
Wind River China

RSA CONFERENCE CHINA 2012
RSA信息安全大会2012

Session ID:

Session Classification:

The Growing Network Security Challenge

Security Function Under Pressure

- Malware
- Intrusions
- Viruses

- More Inspection
- CPU Overload
- Bottlenecks
- Increasing Costs

New "More Intelligent" Approach Required

RSA CONFERENCE CHINA 2012

TRADITIONAL
Multiple Security Appliances within the Network

NEXT GENERATION
Consolidate into one with advanced multicore technologies

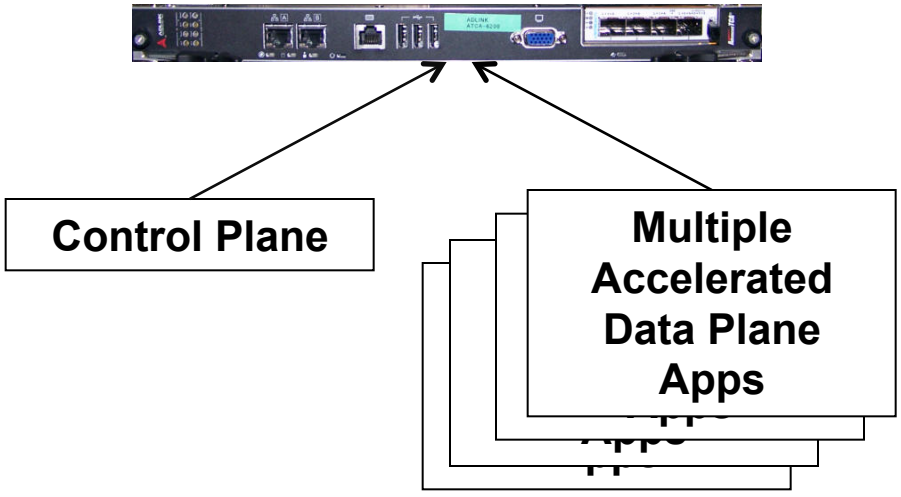

Next Generation Based on Advanced Multicore Technologies

Advanced multicore technologies enable network intelligence needed to ACCELERATE – ANALYZE – SECURE network data

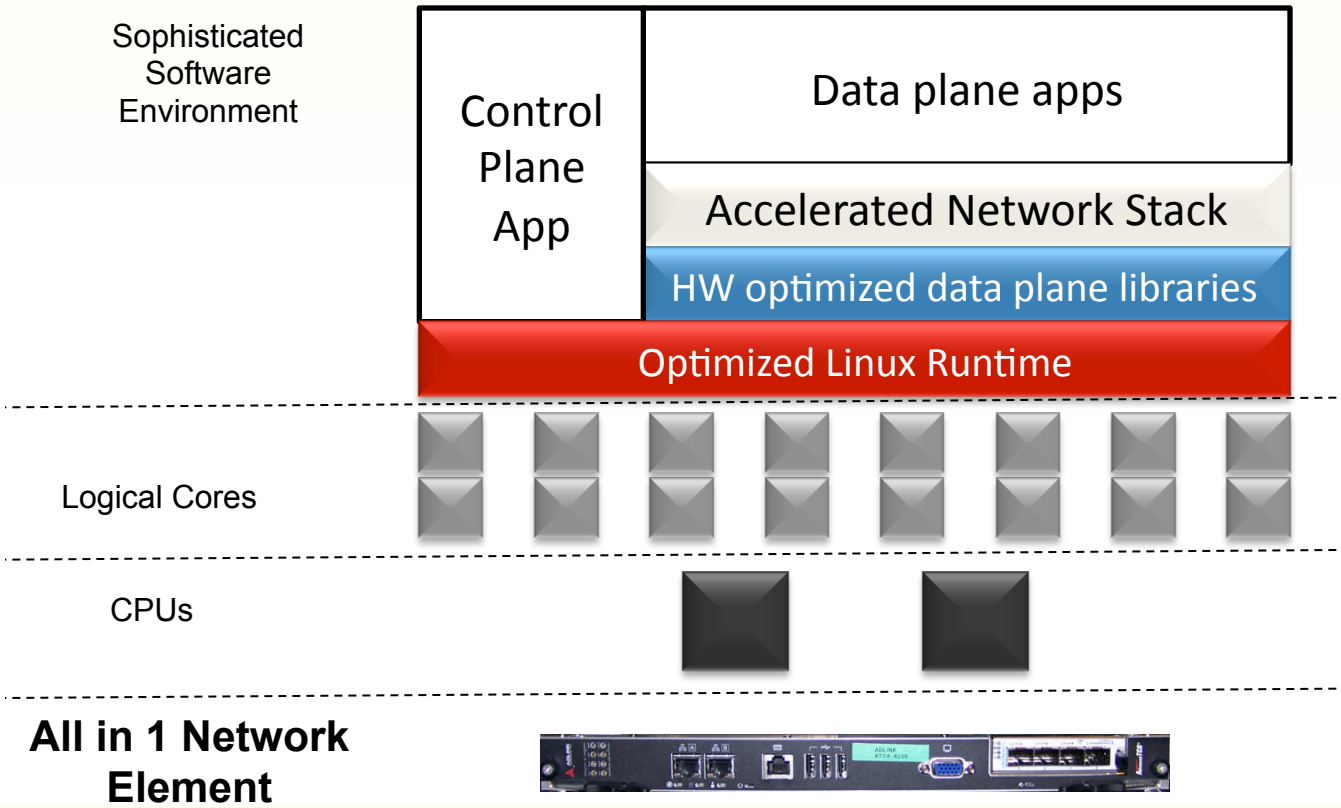

Advantages of a Consolidated System

- Higher performance
- One environment for control and data plane
  -Easier to manage and debug
- Enables multiple Deep Packet Inspection apps at line rate
  -Pattern Matching
  -Traffic Shaping
- Enables a wide-variety of functionality
  -IPS/IDS, FW, AV, UTM
- Scales across product families

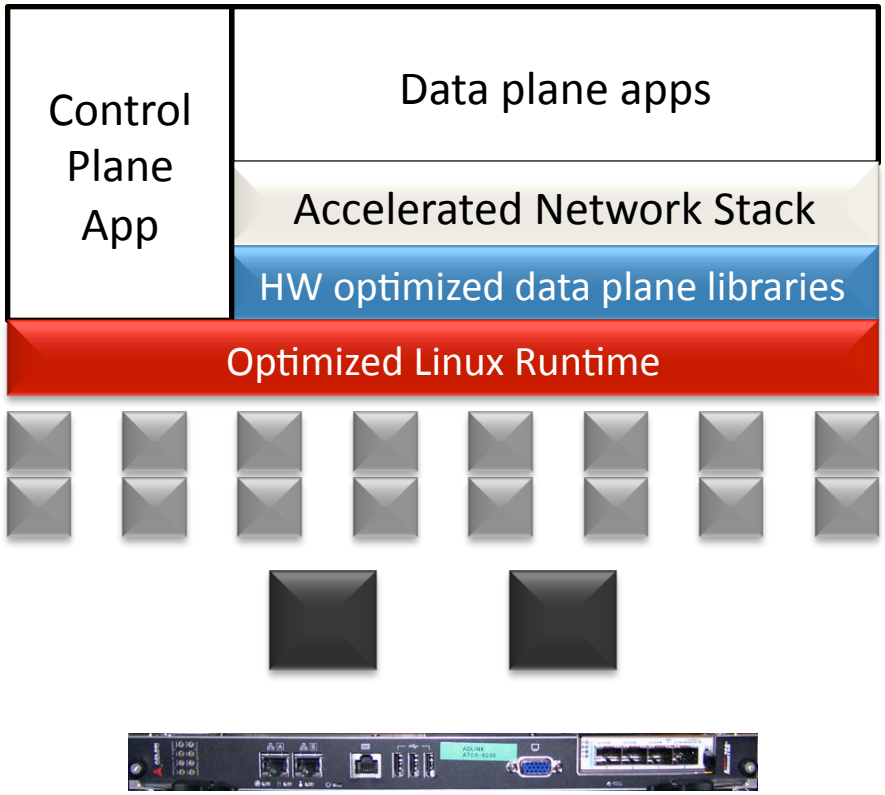

Another Look at the High Level Architecture

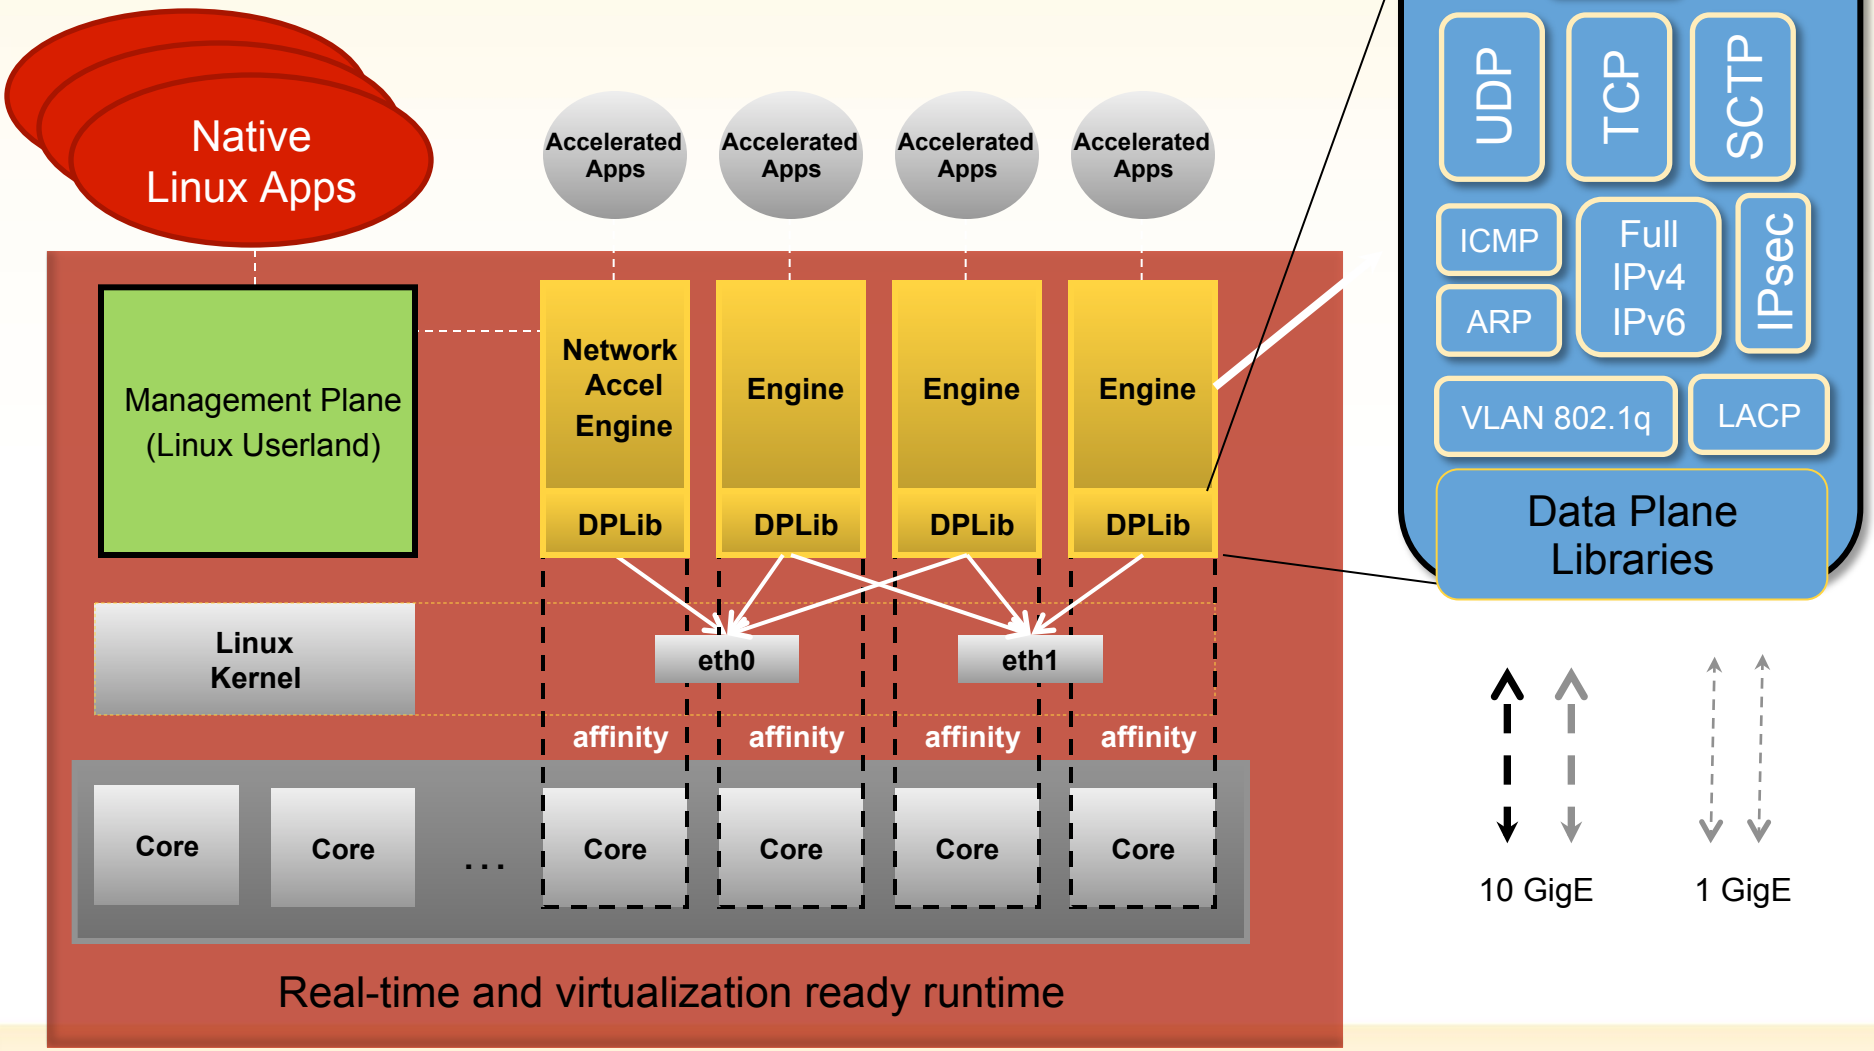

An Even Deeper Look at a Multicore Accelerated Software System

- Adds spt for multiple contexts within the same thread/process/core
- enables implementing a multi-tasking environment and the BSD socket API with blocking calls

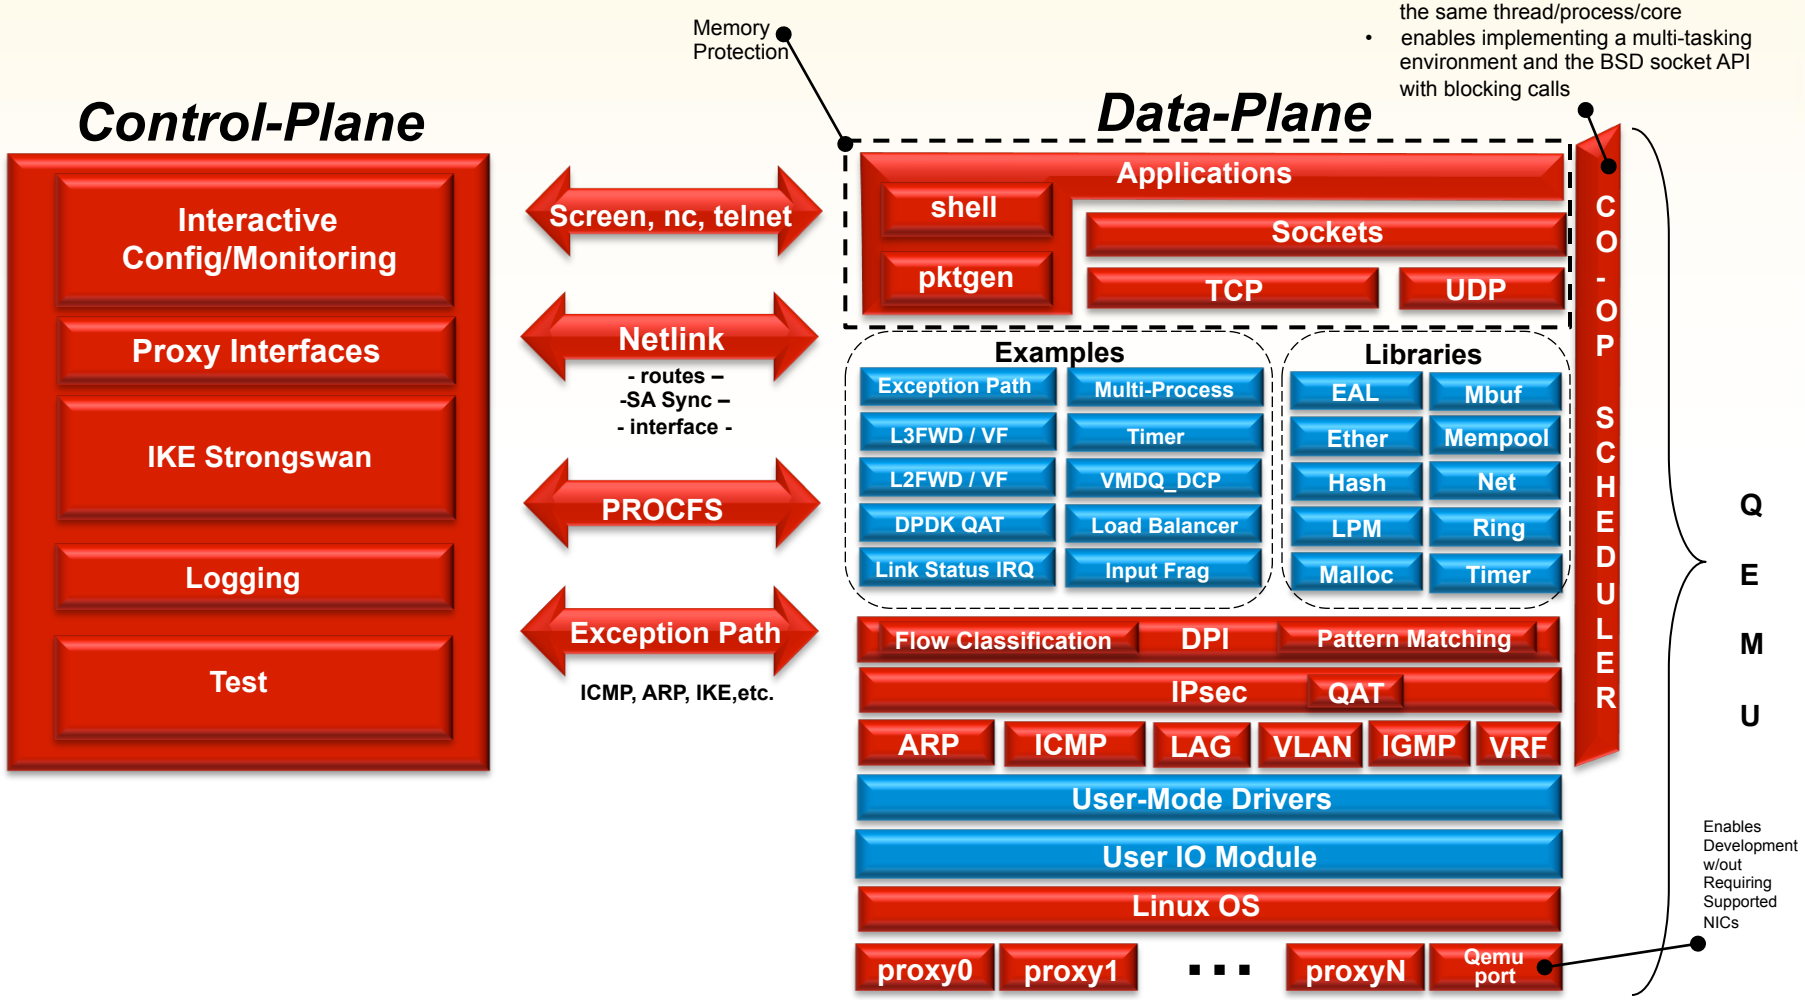

Network Protocols Accelerated

- IP checksum validation
- IPv4 and IPv6 support
- Proxy Interfaces
- Exception Path

  -which enables: ARP, ICMP, SSH, etc. on the acceleration-plane
- Layer 4 acceleration: UDP, TCP, SCTP
- IPsec/IKE

Benefits of Using an Accelerated Software Solution

- Reserves cores of a multicore CPU for accelerated packet processing
  - uses an optimized TCP/IP stack on each core of the accelerated plane
  - Scales with high-volume traffic and core count
- Goes beyond data plane libraries
  - Integrated data plane libraries for optimized silicon-level packet movement
  - supports existing Linux applications
  - Uses standard Linux constructs and tools
- More efficient system
  - Scales to higher throughput (~16Mpps/core)
  - Uses fewer accelerated; leave more for control plane apps

Performance vs. Native Linux (L3 FWD)

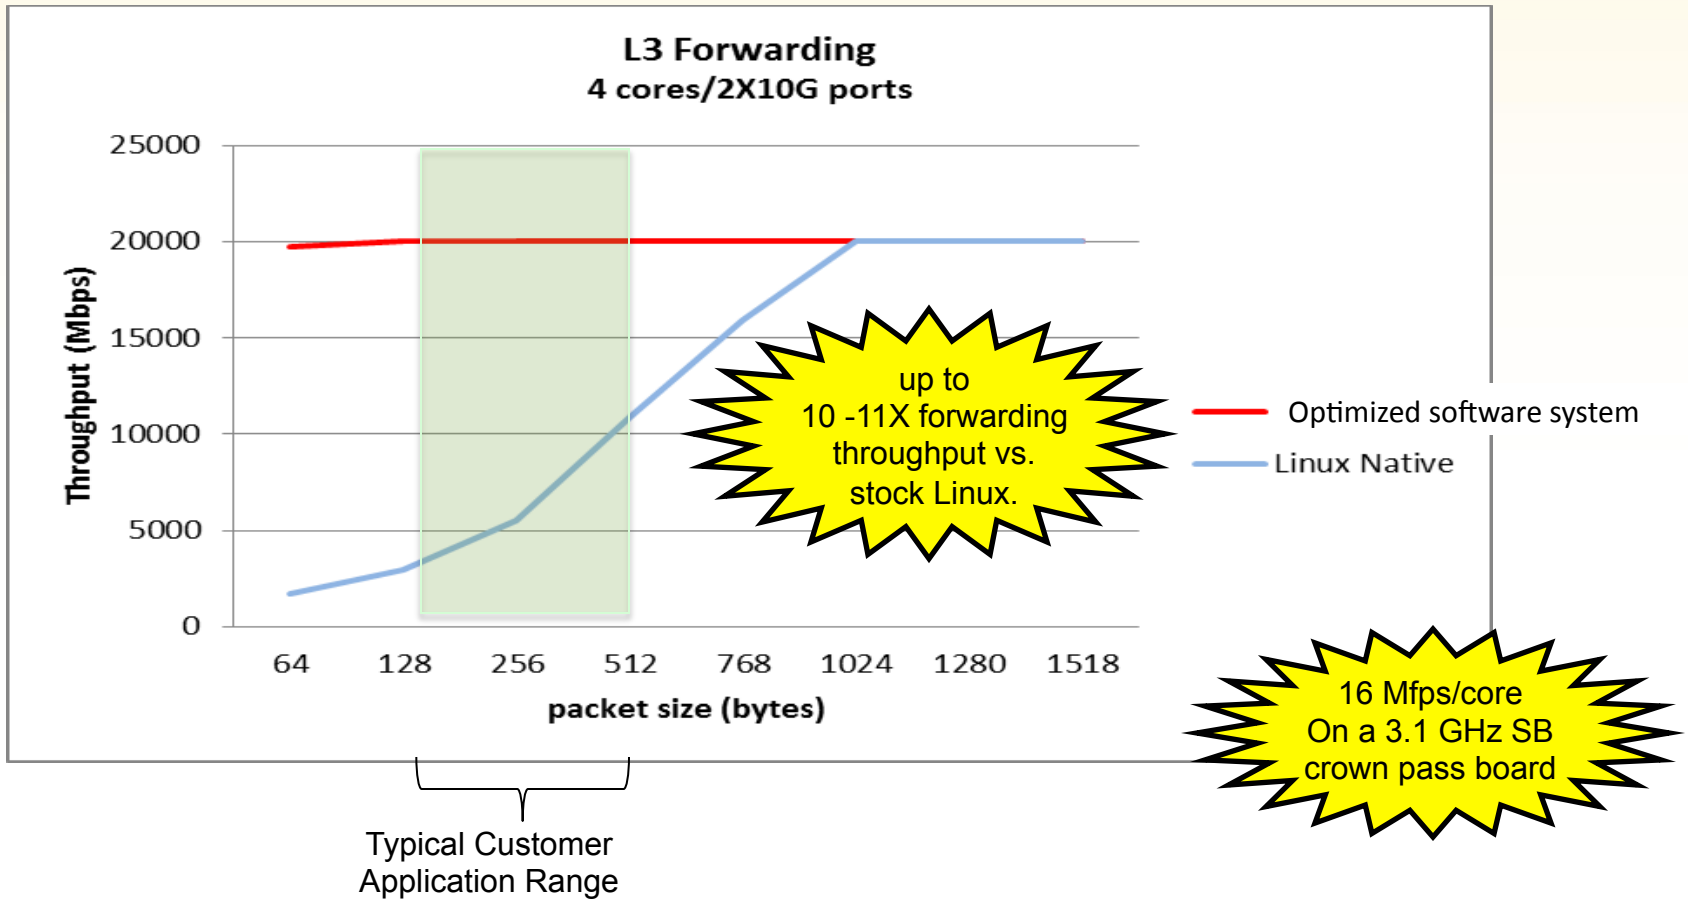

This benchmark is comparing L3 forwarding throughput performance of Linux native vs. optimized software system. The setup is using 2X10G Ethernet ports and 4 hyperthreads on 4 separate cores on a Sandy Bridge based platform.

Performance vs. Native Linux (UDP)

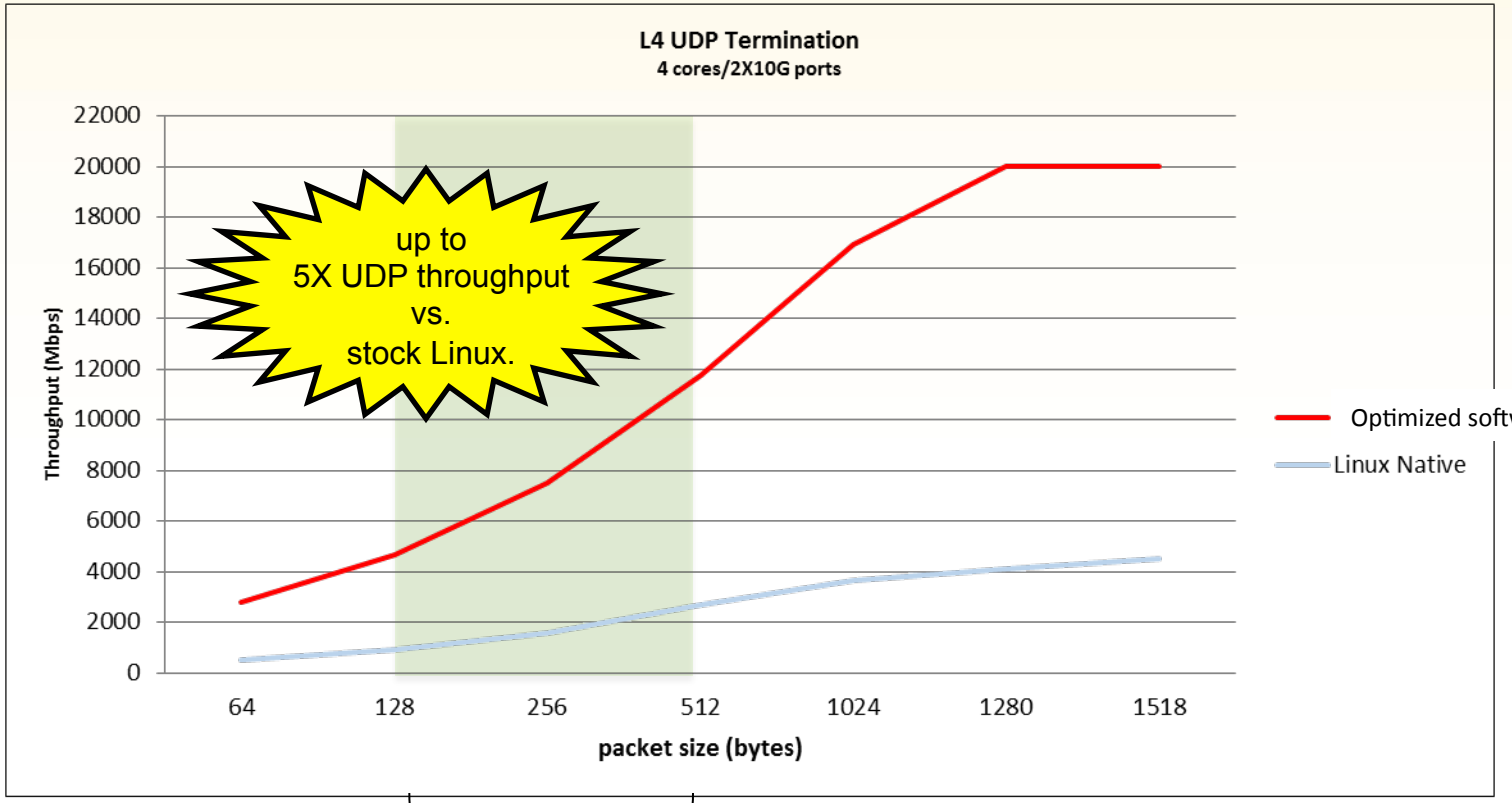

Typical Customer Application Range

This benchmark is comparing L3 forwarding throughput performance of Linux native vs. optimized software system. The setup is using 2X10G Ethernet ports and 4 hyperthreads on 4 separate cores on a Sandy Bridge based platform.

Why this approach is faster?

1. Software system requires fewer cycles than native Linux
2. Faster context switching
3. No interrupts)
4. Much fewer misses in L1 cache than Linux

Performance vs. Native Linux (TCP)

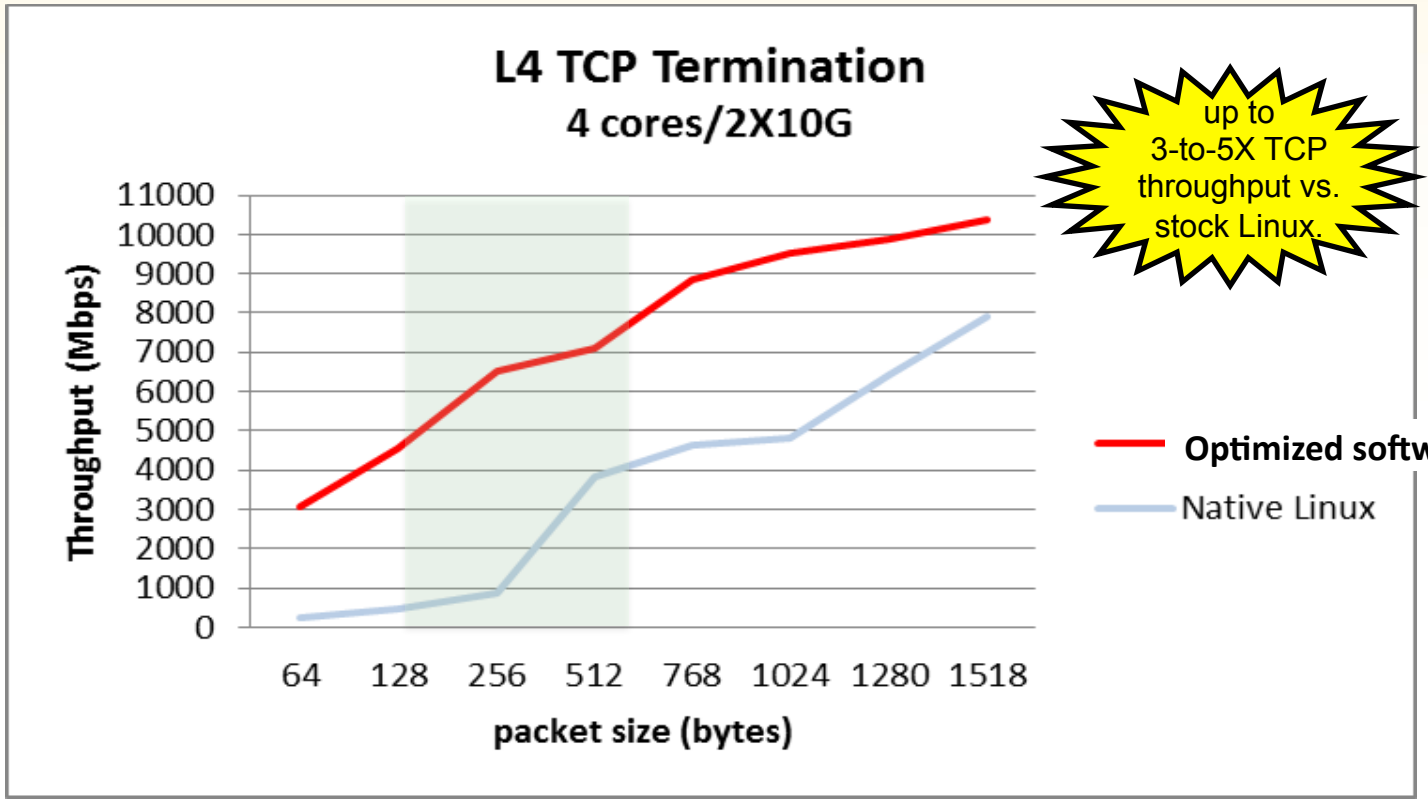

This benchmark is comparing L3 forwarding throughput performance of Linux native vs. optimized software system. The setup is using 2X10G Ethernet ports and 4 hyperthreads on 4 separate cores on a Sandy Bridge based platform.

Sophistication of Network Acceleration Engines

- Each engine runs as a Linux process and thread
- Polling options
  - Full Mesh (default) – each NAE polls every port
  - Partial Mesh – number of NAE pollers per port is restricted **-x**
  - Manual Mesh – NAEs are manually assigned using the **–a** option
- Polls multiple receive queues in bursts
- Transmit queue for each interface it polls

Example: Initialization of Manual Mesh

- Initialize the network acceleration engine :
  - Edit /etc/unap/unap.conf – set UNAP_INSTANCES=0
  - 1 engine per port example:
    - unap-nae –c 0x1f -- -p 0x3 –f 0x5 –a 2/3
  - 2 engines per port example:
    - engine –c 0x1f -p 0x3 –f 0x5 –a 1,2/3,4

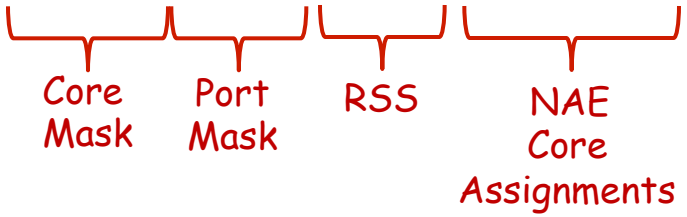

- a 2/3 means "assign the core 2 to NAE 1 and core 3 to NAE 2".
- a 1,2/3,4 means "assign core 1 and 2 to NAE1, core 3 and 4 to NAE 2".

- Assign IP addresses to the proxy interfaces:
  - ifconfig proxy0 –inet 10.1.1.1 up
  - ifconfig proxy1 –inet 10.2.1.1 up

Network Acceleration Enables Deep Packet Inspection

High Performance

- Up to 100Gbps throughput

Highly Scalable

- Linearly Scales from low-end processors to high-end
- Introduce high-speed DPI onto an entire product line with one integration cycle

Low Latency and Low Overhead

- Data is processed directly on the CPU: low latency, low overhead, Low compile time

Easy to Manage

- Simple integration process
- Fast time to market of new features/functions through software upgrades

Accelerated Pattern Matching

- **Software DPI/Pattern Matching library**
  - Higher performance, more comprehensive than most HW solutions
  - Regular Expressions library
  - Portable
  - Highly scalable (low-end to high-end)
- **Massively Parallel Matching**
  - Supports hundreds of thousands of simultaneous patterns and matches concurrently
- **Multi-gigabit pattern matching throughput**
  - L7 DPI; linear scaling for most CPU architectures
- **Low Latency and overhead**
  - Particularly compared to hardware pattern match
- **Wide Applicability**
  - Use across a wide range of architectures/platforms (Appliances, Routers, Switches, Servers)
  - Wide range of applications (IPS/IDS, FW, AV, UTM)

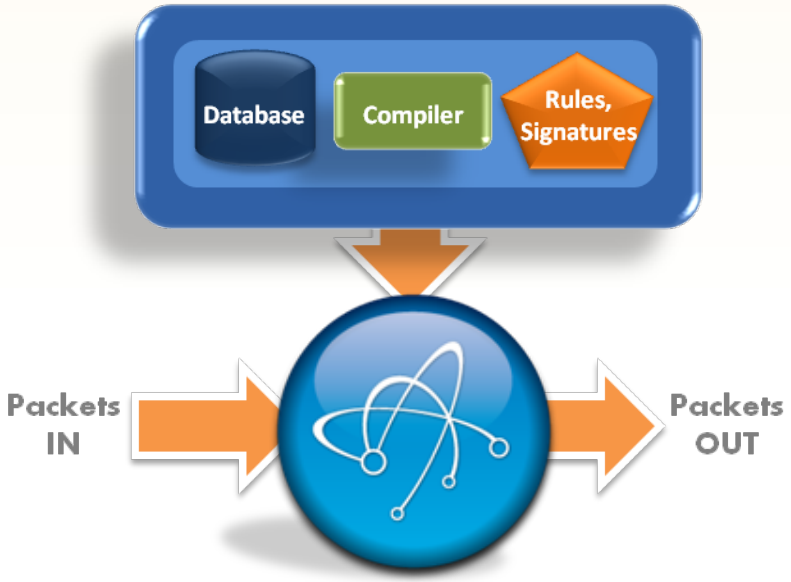

Accelerated Pattern Matching

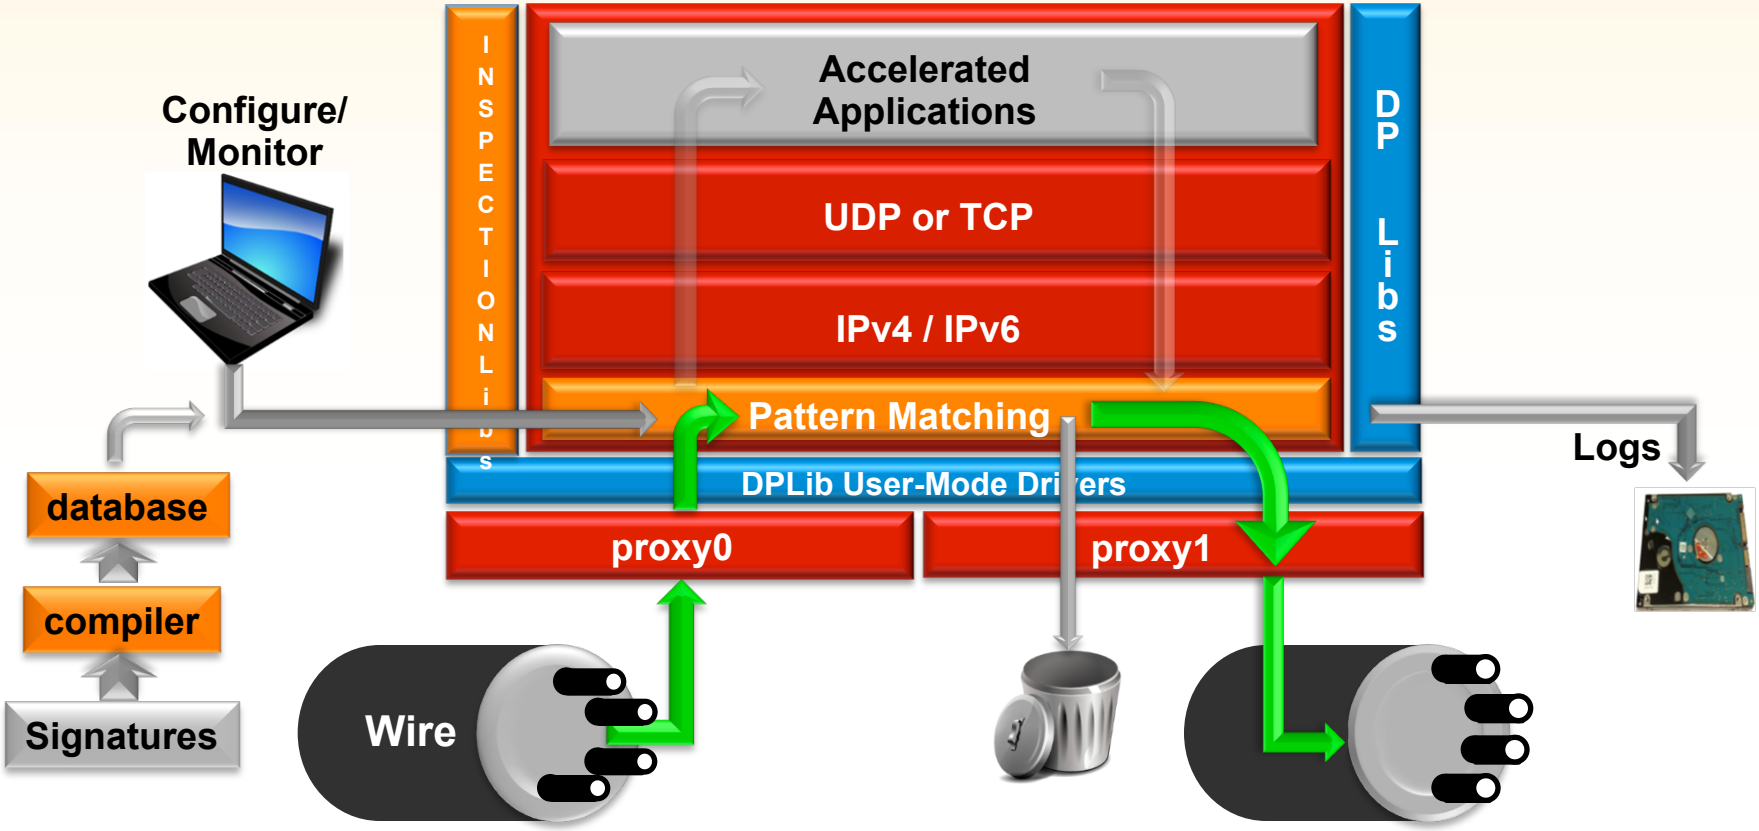

- Pattern Matching (PM) is a **CPU intensive** operation
- High-speed Pattern matching is achieved using the **compiler** and **optimized libraries**

Summary

1. Advanced multicore technologies revolutionizing network elements of all kinds

2. Network acceleration first step in enabling greater intelligence

3. Deep Packet Inspection through software unlocking enormous potential to better analyze and secure data

Thank You

RSA CONFERENCE
CHINA 2012
RSA信息安全大会2012