Welcome to the Dawn of Open-Source Networking.™

Linux IP Routers

Bob Gilligan
gilligan@vyatta.com

Outline

- About Vyatta: Open source project, and software product
- Areas we're working on or interested in working on
- Some of our performance testing results
- Conclusions

Vyatta - The Service Router, Redefined
Integrated, Yet Open

Web GUI & CLI

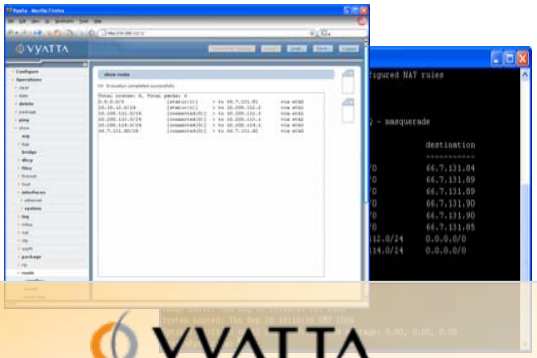

Rich Networking and Security Features

TECHNICAL SPECIFICATIONS

HARDWARE SUPPORT
- 32-bit x86 processors
- PC architecture

INTERFACES
- Linux-supported Ethernet cards (both on the motherboard- and PCI-based cards)
- Sangoma A101 T1/E1 cards
- Sangoma A301 T3/E3 cards

IP AND ROUTING PROTOCOLS
- IPv4
- OSPFv2
- BGP-4
- RIPv2
- Static routes

IP ADDRESS MANAGEMENT
- Static
- DHCP server
- DHCP relay

ENCAPSULATIONS
- Ethernet
- 802.1Q VLANs
- PPP
- Frame Relay
- HDLC

SECURITY
- Stateful inspection firewall
- Network address translation (NAT)
- RADIUS authentication
- Individual user accounts and passwords

HIGH AVAILABILITY
- VRRP (Virtual Router Redundancy Protocol)
- Support for multiple power supplies
- Each protocol sandboxed from others, providing fault isolation

LOGGING AND MONITORING
- Syslog
- SNMPv2c

DEBUGGING AND PACKET SNIFFING
- tcpdump
- Ethereal packet analysis

ADMINISTRATION
- Integrated CLI
- Single configuration file for all protocols and subsystems
- Web GUI
- Telnet
- SSHv2

VYATTA

Graphical User Interface

Command-Line Interface

Routing | Firewall | NAT | VPN | Cfg Mgt

Linux

Easily Extensible

SpamAssassin

Asterisk. The Open Source PBX

ClamAV

Thousands More…

Industry-Standard 32-bit x86

Vyatta: Scalable Software Performance

Twice The Performance

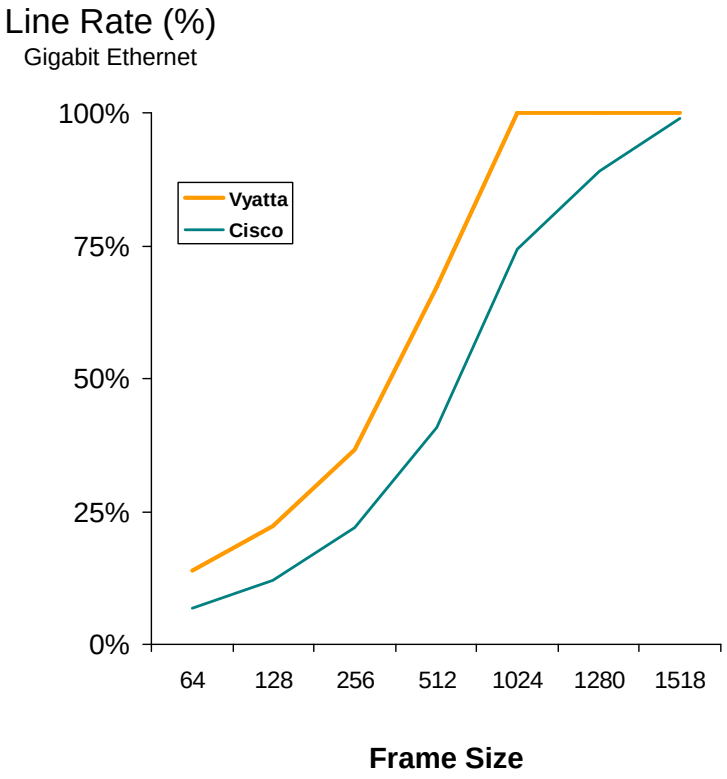

Half The Price

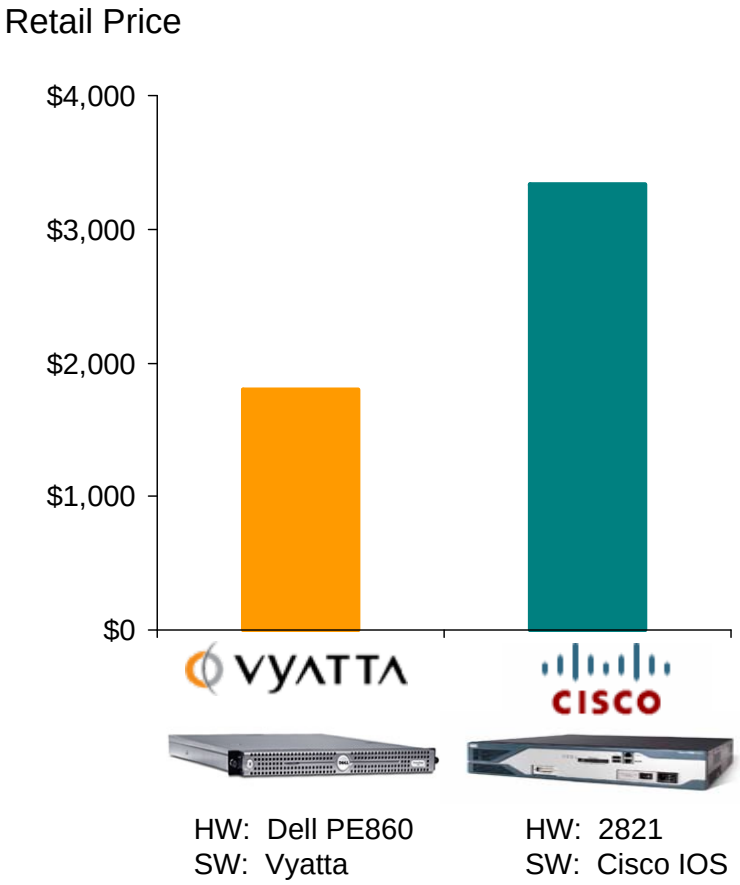

Target Markets: Mid-range Router / Firewall / VPN

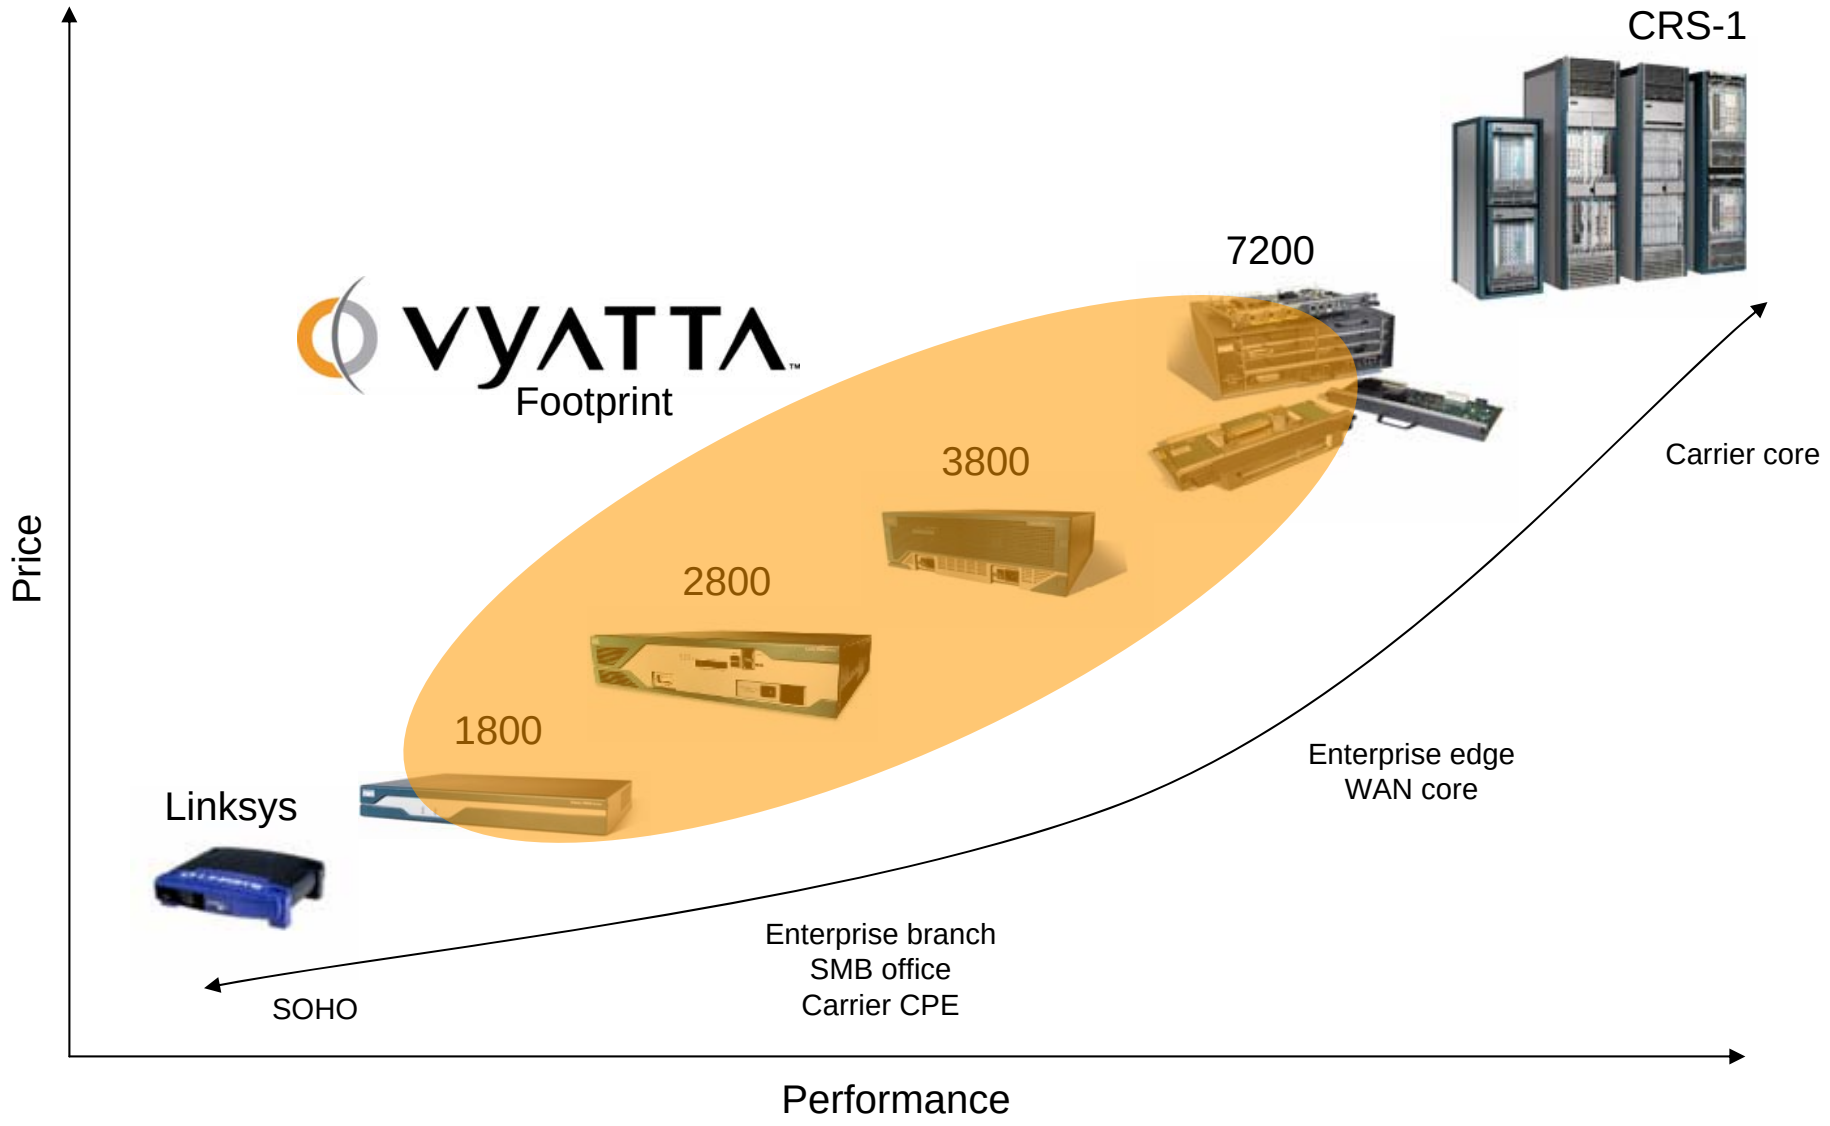

Vyatta Product

- Linux Software Distribution
- Open Source Project and Product
  - Community (free) and Supported (pay) versions available
  - “Appliance” version also available.
  - Open bugzilla bug database, wiki, user group lists, docs
  - Open Git source repository
- Getting Community Version:
  - Start at: http://www.vyatta.com/community/
  - Download and burn live CD: http://www.vyatta.com/download/
  - Full source at: http://suva.vyatta.com/git/

Software Focus Areas

- Interested in working with the community on features relevant to running Linux as a router
  - Router issues not necessarily the same as server or desktop issues
- Routing protocol performance: XORP Package
  - Fast convergence large routing tables
  - Software optimization
  - MP scaling
- IP Forwarding performance
  - Performance with large routing tables (> 200,000 routes)
    - Kernel routing table (FIB) hash vs. TRIE tree implementation
  - Performance forwarding min-size (64 byte) packets
  - MP scaling: Efficiently take advantage of dual/quad core processors
    - Most new machines will be dual/quad core
- Scheduling IP forwarding and user-level routing protocols
  - Router runs both; Both are CPU intensive
  - Need to ensure both get adequate CPU under heavy load
  - Efficiently and fairly on MPs

Hardware Focus Areas

- Features NICs should support:
  - PCI-e, especially for serial cards
    - Some new machines support only PCI-e
    - Older serial/WAN NICs are still PCI-X
  - Multiple MAC addresses
    - For MAC-address takeover
    - Used by Virtual Router Redundancy Protocol (VRRP) to provide High Availability
  - NAPI support
  - VLAN/Tagging support
- IPsec performance
  - Raw encryption performance
  - Hw encryption engine performance vs. more cores

Testing results: IP Forwarding Performance

- Two standardized router forwarding performance tests:
  - Zero-Loss Throughput Test
    - Reduce offered rate until all packets get through (higher is better)
  - Packet Loss Test
    - At 100 % offered rate, measure packet loss rate at various packet sizes (lower is better)
  - Both tests defined in the IETF Benchmark Methodology Spec (RFC 2544)
  - Both measure at range of packet sizes (64 bytes – 1518 bytes)
- Test configuration
  - SmartBits 600B network traffic generator
  - 2 GbE links – Bidirectional test

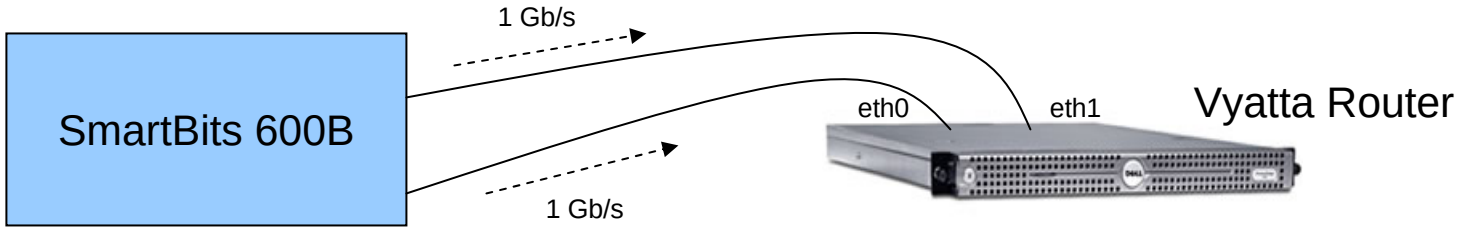

Testing results: IP Forwarding Performance

- Platforms tested:
  1. Dell PE860:
     - Celeron 336: 2.8 GHz CPU, 256 KB L2 cache
     - 533 MHz FSB
     - 2 x On-board BCM 5721 NIC
     - PCI-e x1 lane interconnect to each NIC
  2. SuperMicro PDSM4+ motherboard:
     - Dual-core Pentium-D 935: 3.2 GHz CPUs, 2 MB L2 cache
     - 800 MHz FSB
     - Off-board 2-port Intel 82571 NIC
     - PCI-e x4 lane interconnect to NIC
- Linux 2.6.20 kernel
  - No firewall rules
  - No NAT

Zero-Loss Throughput Comparison

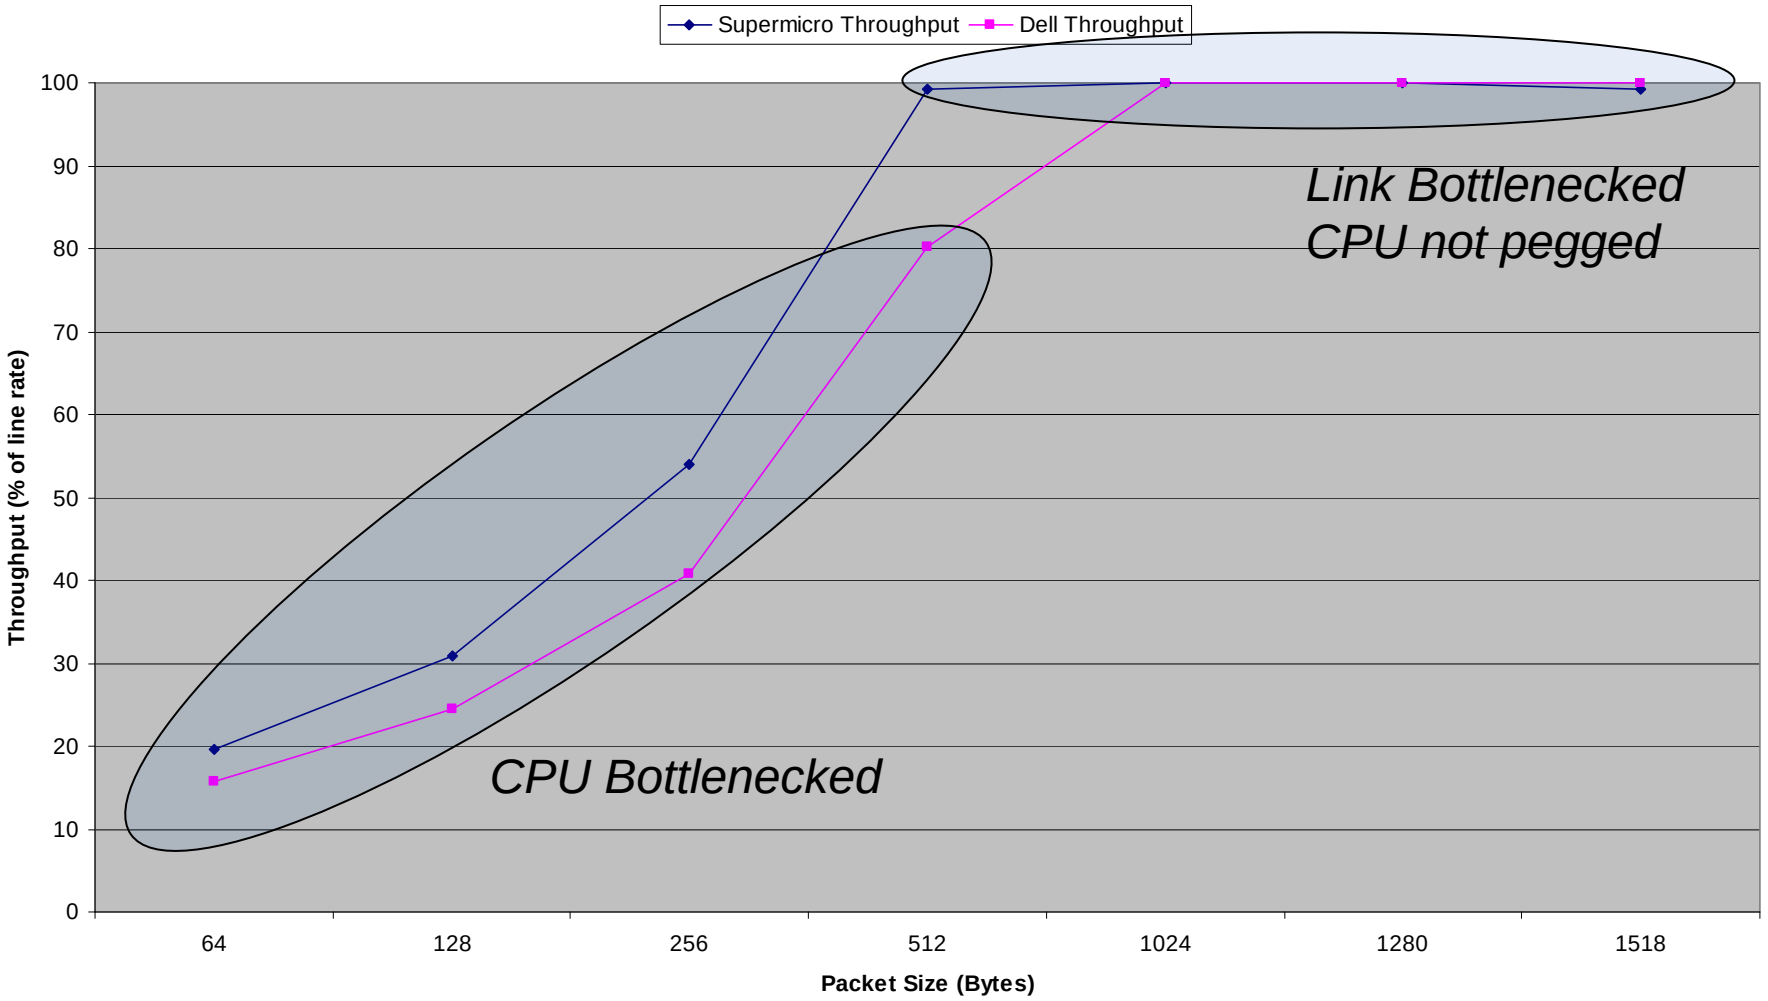

Loss Rate Test Comparison

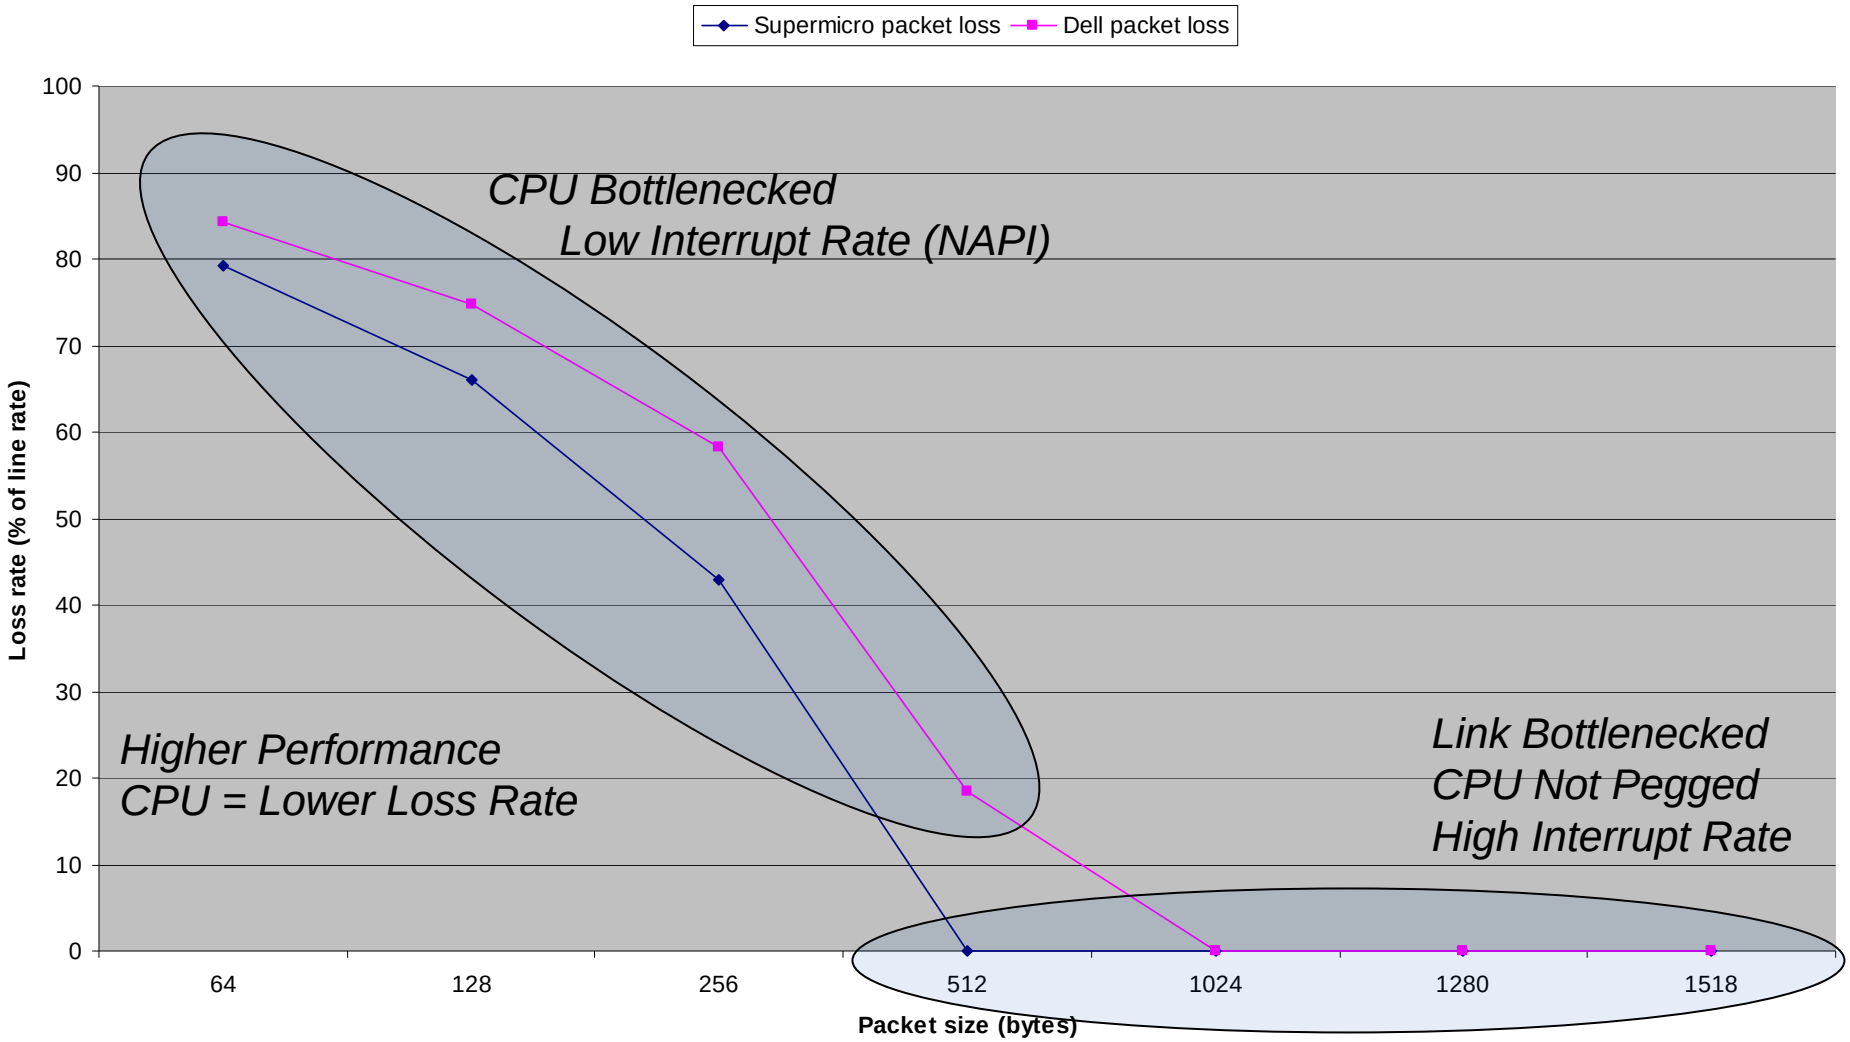

Observations

- “Low end” server platforms deliver excellent IP forwarding performance
- Forwarding performance correlates with CPU performance
  - Higher performance CPU → higher throughput rate, lower loss rate
  - At small packet sizes, when CPU is pegged
- NAPI appears to be working
  - Interrupts moderated when CPU is pegged
- One issue to be investigated:
  - Only one CPU utilized on dual-core Pentium-D platform

Conclusion

- Linux on x86 server platform makes a great IP router!
  - “Twice the performance at half the price”.
- Vyatta is interested in working with the community to improve features relevant to IP routing