Real Routing at Gigabit Speeds

Debra Deutsch
Product Line Manager, Layer 3 Switching
Prominet Corporation

Prominet

Design Drivers

- Market Requirements
  - » Integrated Layer 2 and Layer 3 product line
  - » Scale from small to high capacity configurations
  - » Low latency for selected traffic; easy to use
  - » Deliver full wire speed performance
- Technology
  - » Division of function among Cajun Switch core ASICs
  - » Non-blocking crossbar switch
  - » Dual-queue, weighted round-robin queuing
  - » Cajun ASICs, Intelligent Cache, undersubscribed switch

This presentation will focus on the areas highlighted in gray

Prominet

System Architecture

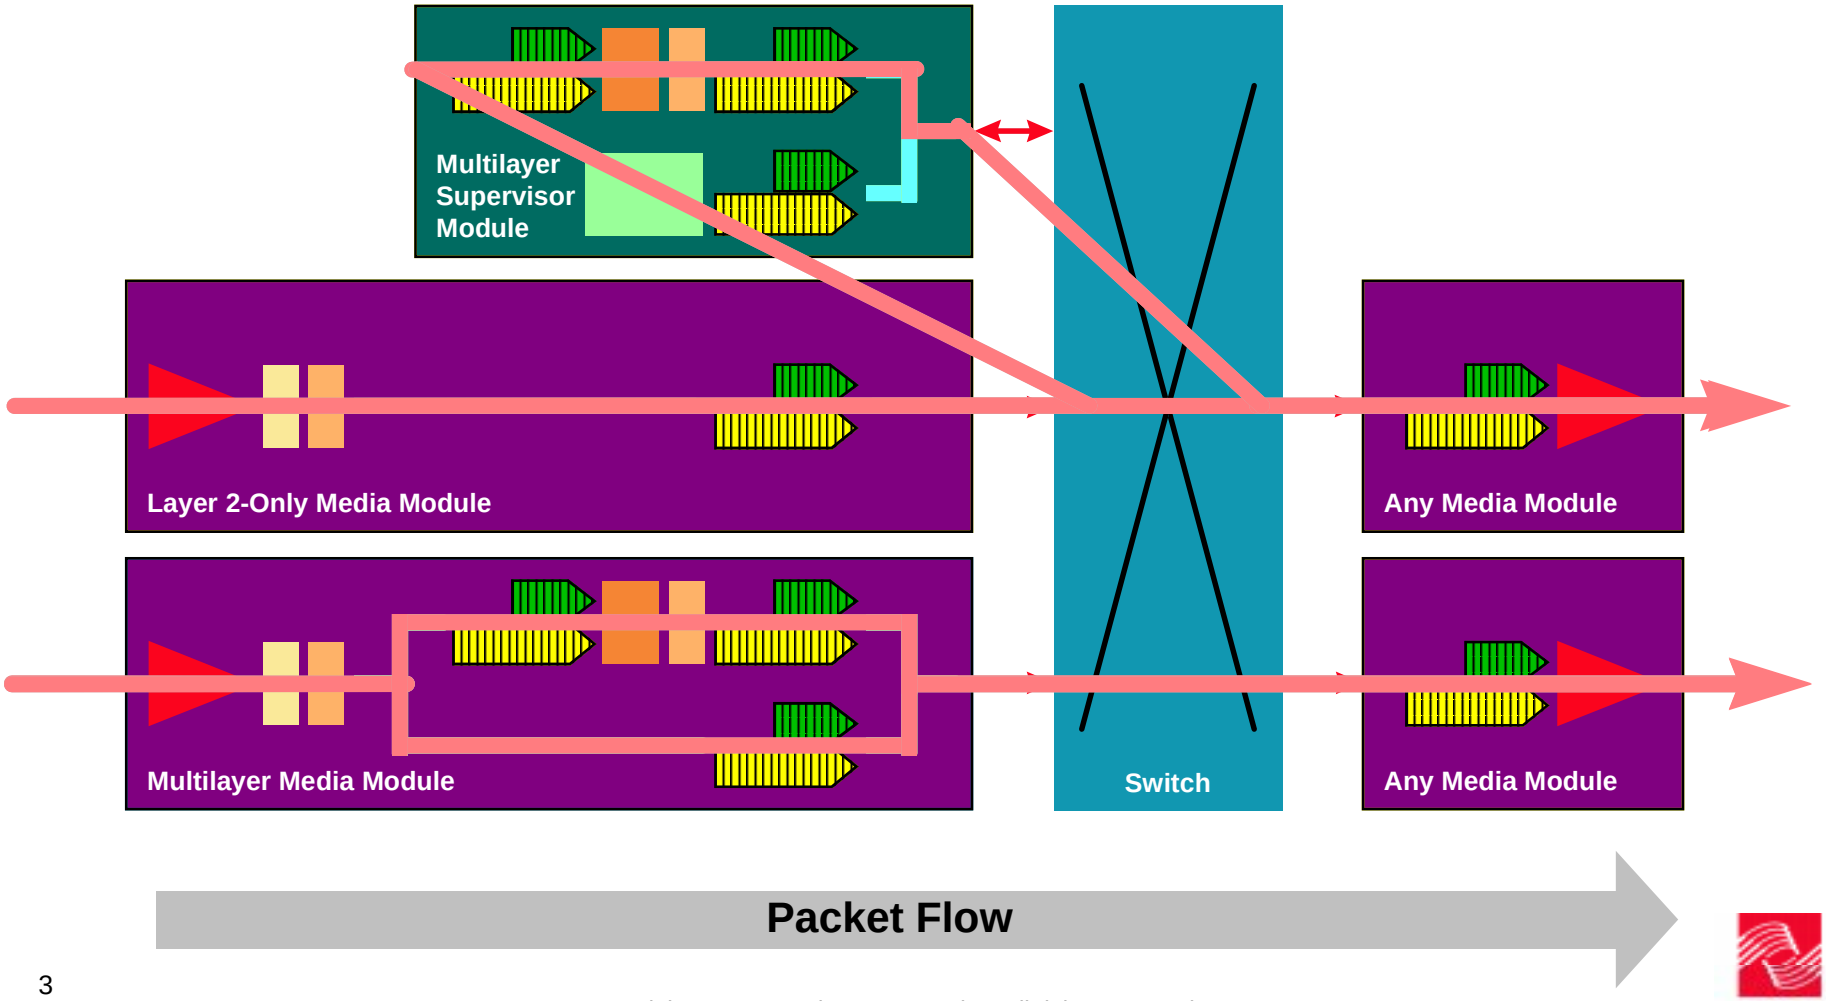

Prominet

Packet Routing ASIC and Address Cache

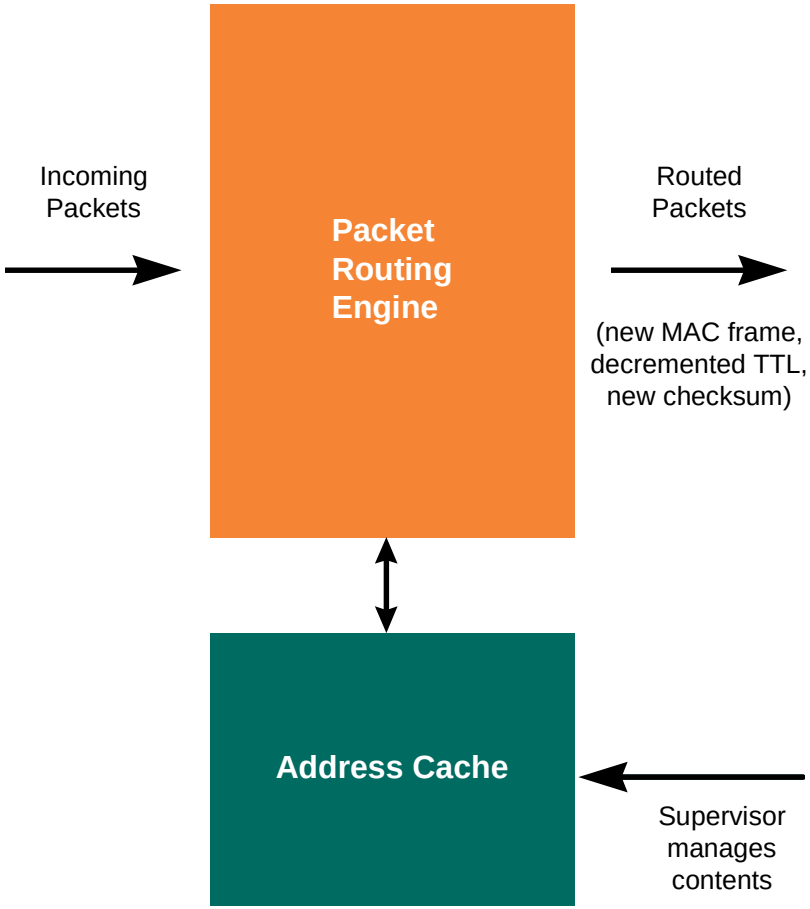

- Complete IP and IPX packet processing
  - » Designed to a cycle budget
  - » Some parallelism
  - » Filtering, prioritization in silicon
- Address cache in SSRAM for speed, modularity, cost
  - » 20,000 fine-grained entries manage RSVP flows
  - » Proprietary fast look-up and management algorithms
    - – chose against longest-prefix matching

Prominet

Interconnects, Queues, and QOS

- Choice of interconnect and mechanisms for QOS are interdependent
- Prominet's solutions
  - » Switching fabric
  - » Support 8 queues in architecture, implement 2
- How did we arrive at this design?
  - » Technical alternatives
  - » Market and product requirements

Prominet

Shared Memory Interconnect

- A memory pool connects all ports
- Memory bandwidth at least 2x the sum of port bandwidths; controller speed is 2x the sum of port packet rates
- Disadvantages
  - » Memory speed limits ability to scale aggregate bandwidth
  - » Large-scale design not cost-effective in small unit implementation
- Advantages
  - » Minimizes packet copying (low latency)
  - » Controller can be a smart queue manager

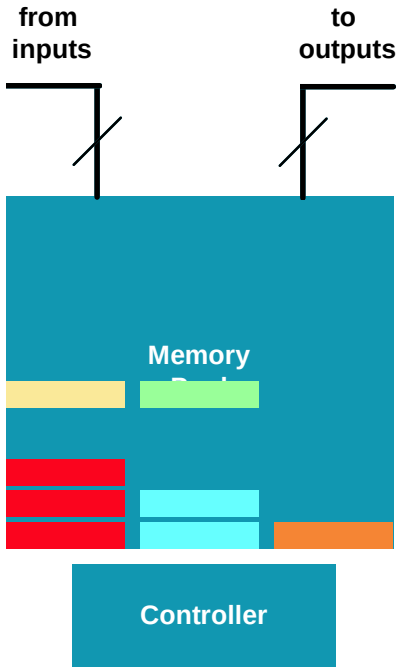

Example: 16 ports @ 1 Gbps
Memory: 4 ns, 128 bits wide
Controller: 48 million packets per second

Prominet

Crossbar Interconnect

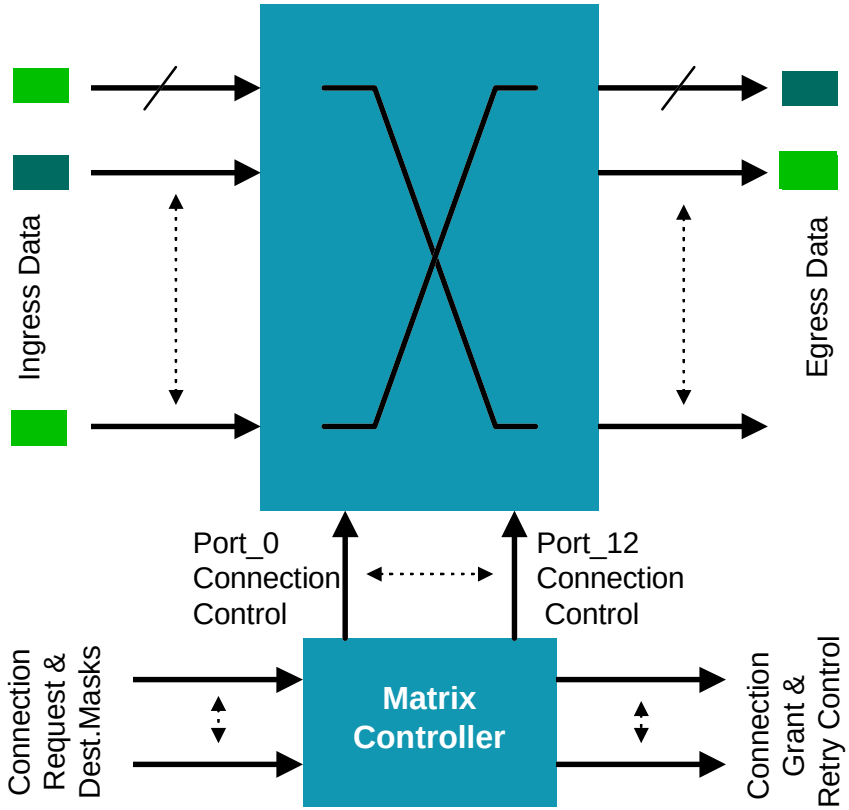

- Crossbar provides network of paths between ports
- Use a combination of input and output queues, run fabric fast to counteract head-of-line blocking
- Disadvantages
  - » Design must address head-of-line blocking
- Advantages
  - » Scales to high aggregate bandwidths
  - » Relative simplicity

Prominet

How Many Queues?

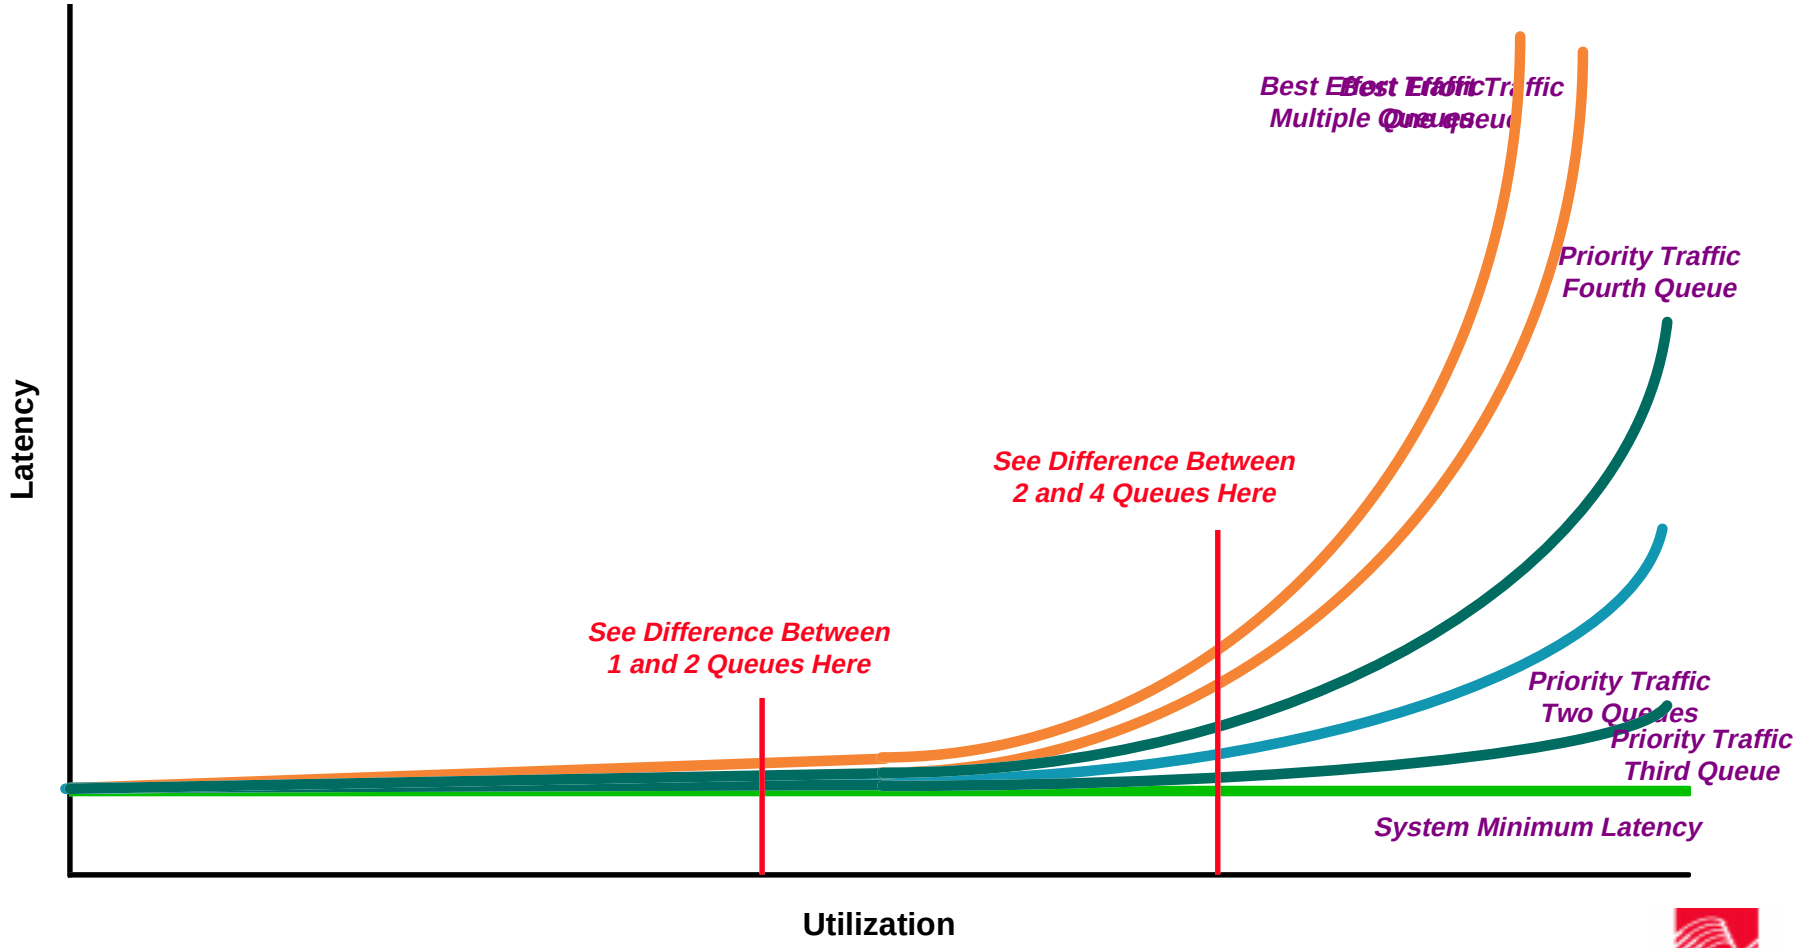

Prominet

LAN Backbone Utilization

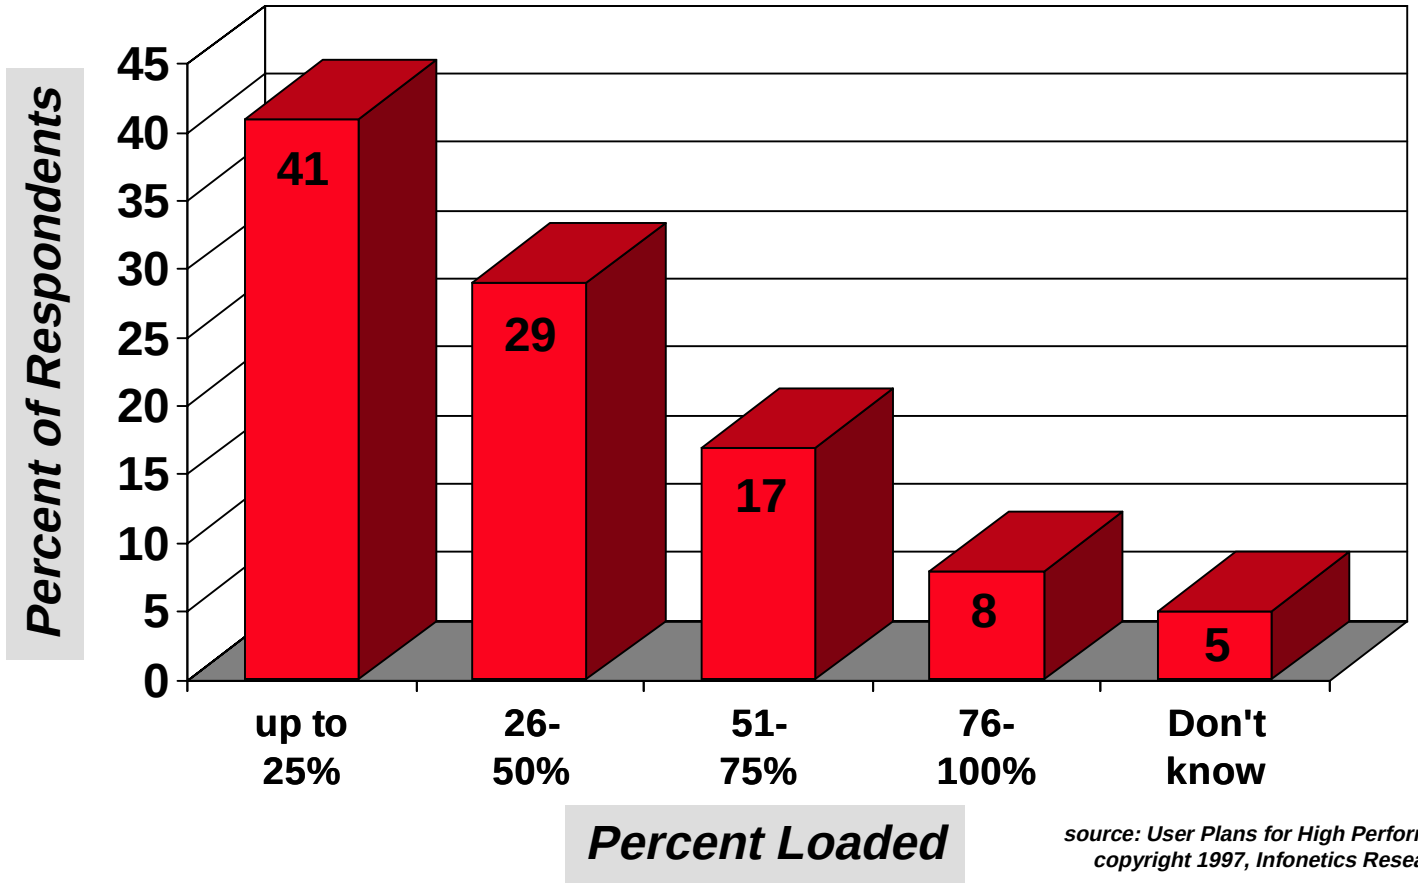

source: User Plans for High Performance LANs
copyright 1997, Infonetics Research, Inc..

Prominet

Prominet P550 Queuing Performance

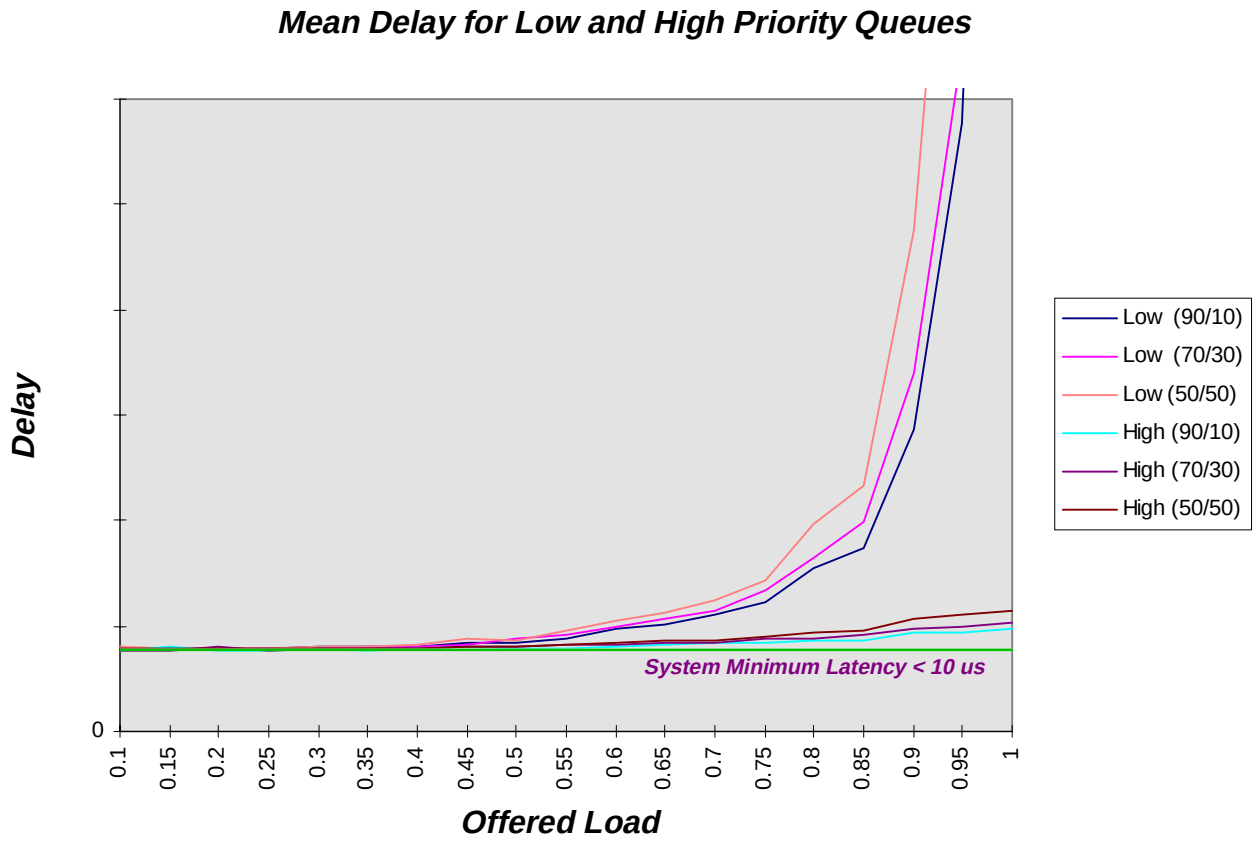

Prominet

Going Even Faster

More ports

- Consider size of master routing table
- No change to packet path on media modules and supervisor
- Use a larger crossbar
  - » datapath speed unchanged
  - » more ports
  - » faster controller to allocate more ports in same cycle time

Faster ports

- Need larger address tables at routing ASICs
- Parallelism can get us close to 10 Gbps today
- We are using “sweet spot” clock rates
- Semiconductor speeds double every 3 years
- Prominet would likely use an architecture very similar to what we have today!

Prominet

Conclusions

- Gigabit-scaled routing is simply a matter of great engineering!
  - » Many feasible design choices
    - – Work around memory speed limitations
  - » Latitude to match technology to market requirements
- Network managers should
  - » Plan to use routers where routing is the right answer
  - » Expect interoperable, cost-effective systems
  - » Take advantage of inexpensive capacity to simplify campus network management

Prominet