GRF IP Switch

GRF16-PR

Table of Contents

State of the Net
Life on the Information Superhighway

Recent WEB log of a major ISP

- **Chicago 9:20 PM EDT, 6/7/96**
  - "...have seen Chicago lose its link to other routers...determined to be caused by heavy CPU load."
- **Santa Clara 6:10 PM EDT, 6/10/96**
  - "...routers have been crashing due to periods of 100% CPU usage."
- **California 2:40 PM EDT, 6/10/96**
  - "...problems at major peering points...The increased load on our routers caused them to drop their BGP sessions. Even though they were up and reachable...they would not route traffic."
- **East Coast 9:45 PM EDT, 6/13/96**
  - "...Caused not only our peering sessions at MAE-east to drop...While our routers remained up and reachable, the CPU usage was maintaining a 99% usage and would not maintain their BGP sessions. As a result they would route any traffic through or headed to them."
- **Chicago 1:35 PM EDT, 6/20/96**
  - "...Chicago-NAP router has not been able to maintain its connectivity...we cannot continue with the instability of the router presently being used."

ASCEND

State of the Net
The Internet Has Outgrown Current Architecture

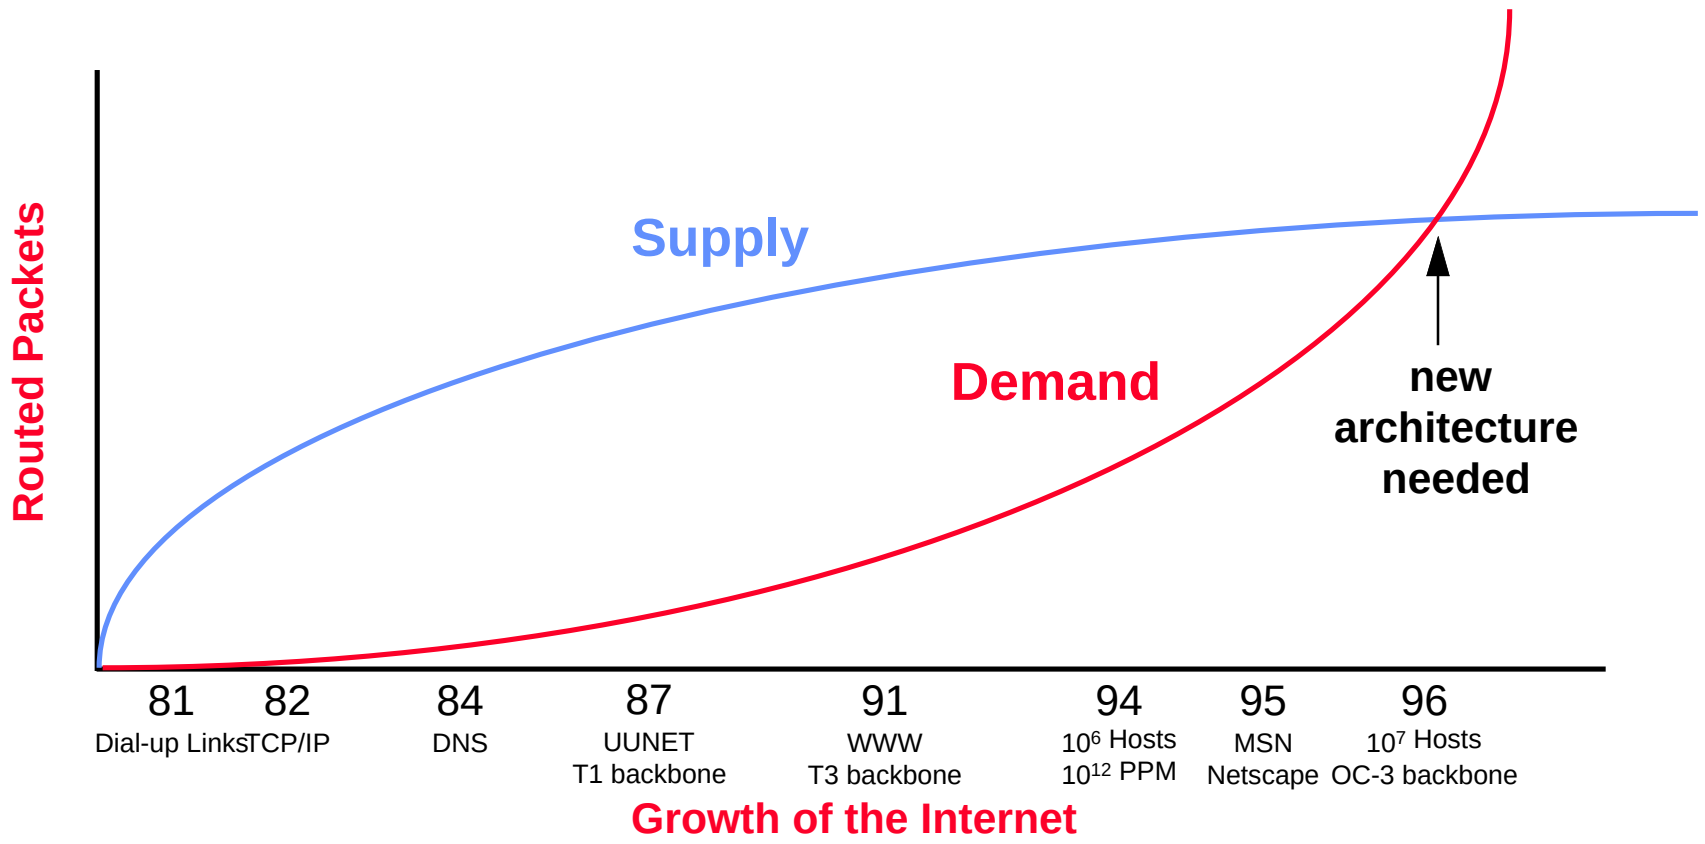

- In early Internet, router supplied more capacity than required
- Since introduction of Netscape and new applications, demand has caught up
- In 1996, demand has outstripped capacity of conventional routers

Product Overview

GRF IP Switch - Designed for Carriers, ISPs and Online Service Providers

- **IP switch**
  - Full Layer-3 routing
  - 16 Gb/s switching fabric
  - Routing functionality distributed across all interfaces
- **High-density, fully redundant modular chassis**
  - GRF 400
    - 4 slots
    - 3U, 5.25 rack mount
  - GRF 1600
    - 16 slots
    - 12U, 21” rack mount
- **Introduces a new level of performance in IP-centric networks**
  - 16 Gb/s aggregate switch bandwidth
  - 10 million packets per second
  - Hardware-assisted, full-route table lookup
- **Open architecture; not tied to specific protocols and WAN/LAN interfaces**

Product Overview
End-to-End Networking

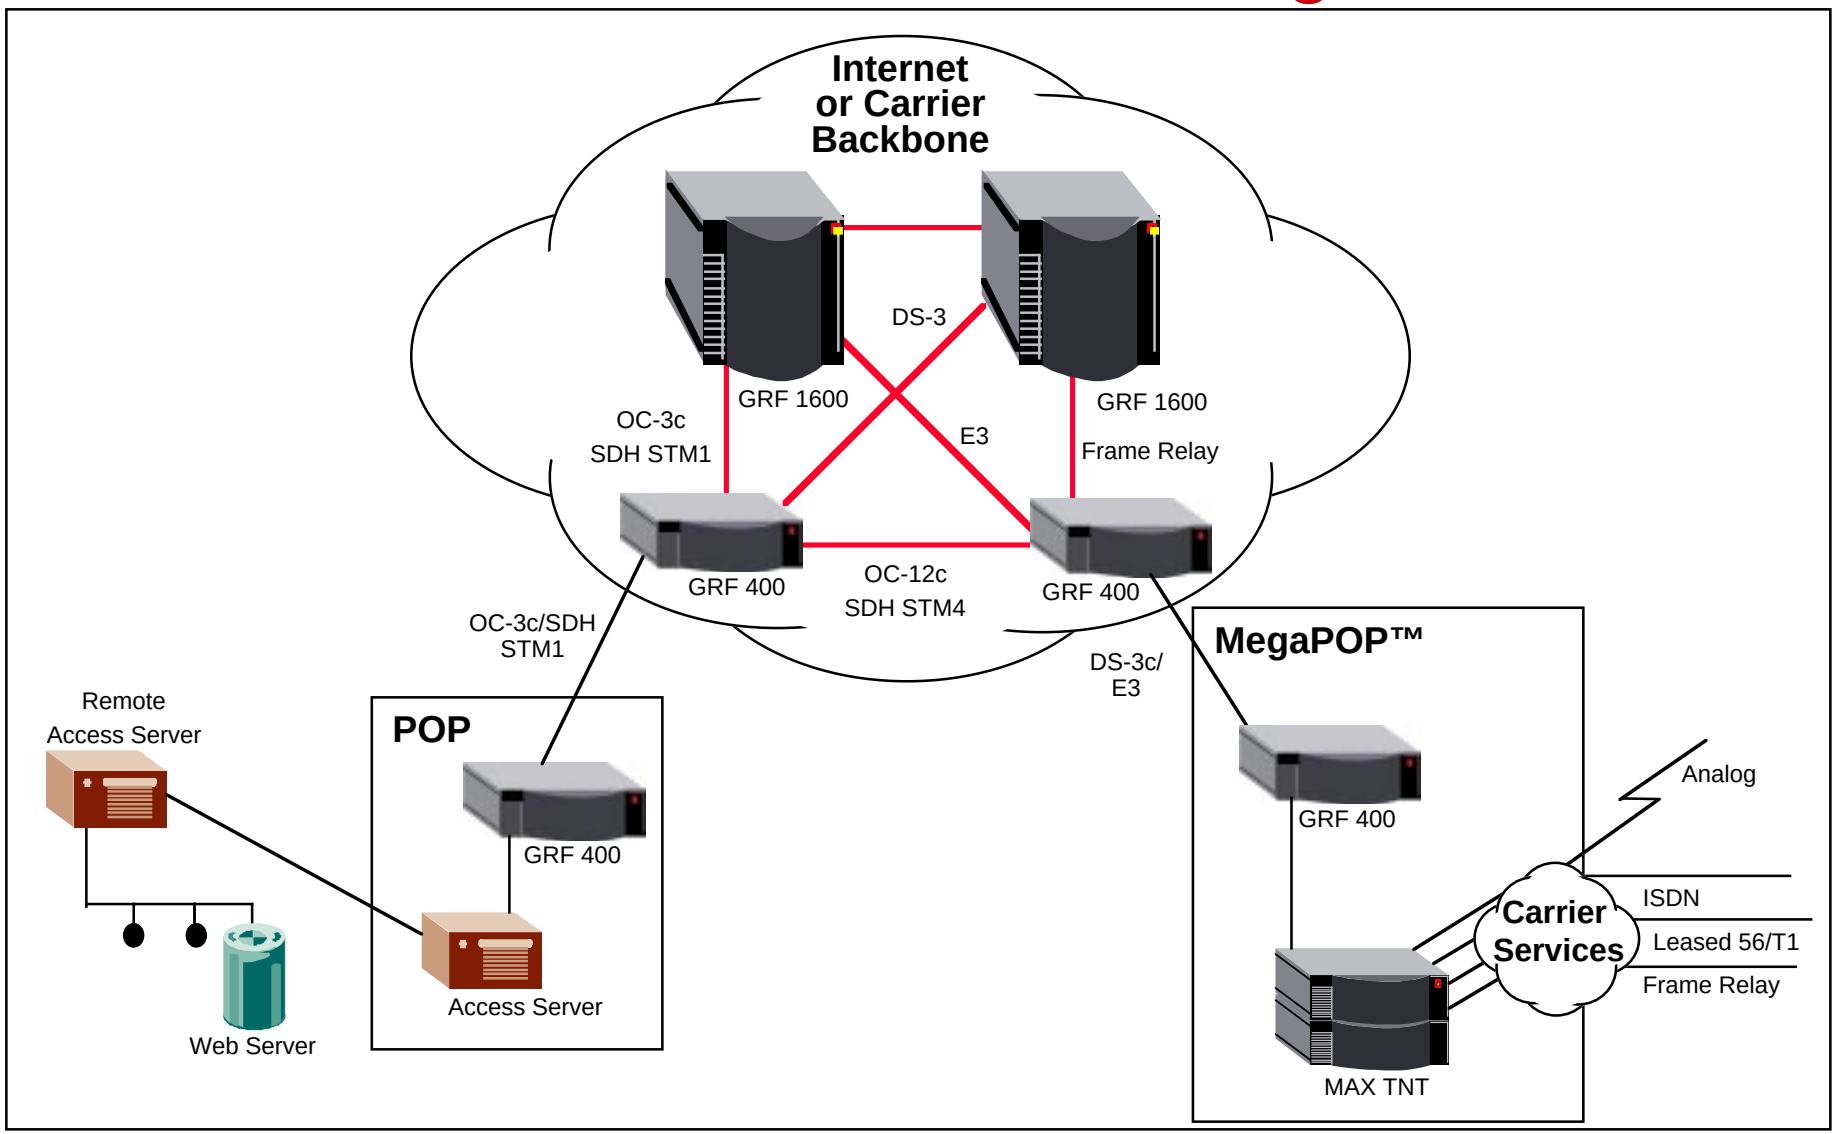

Applications
The POP Today

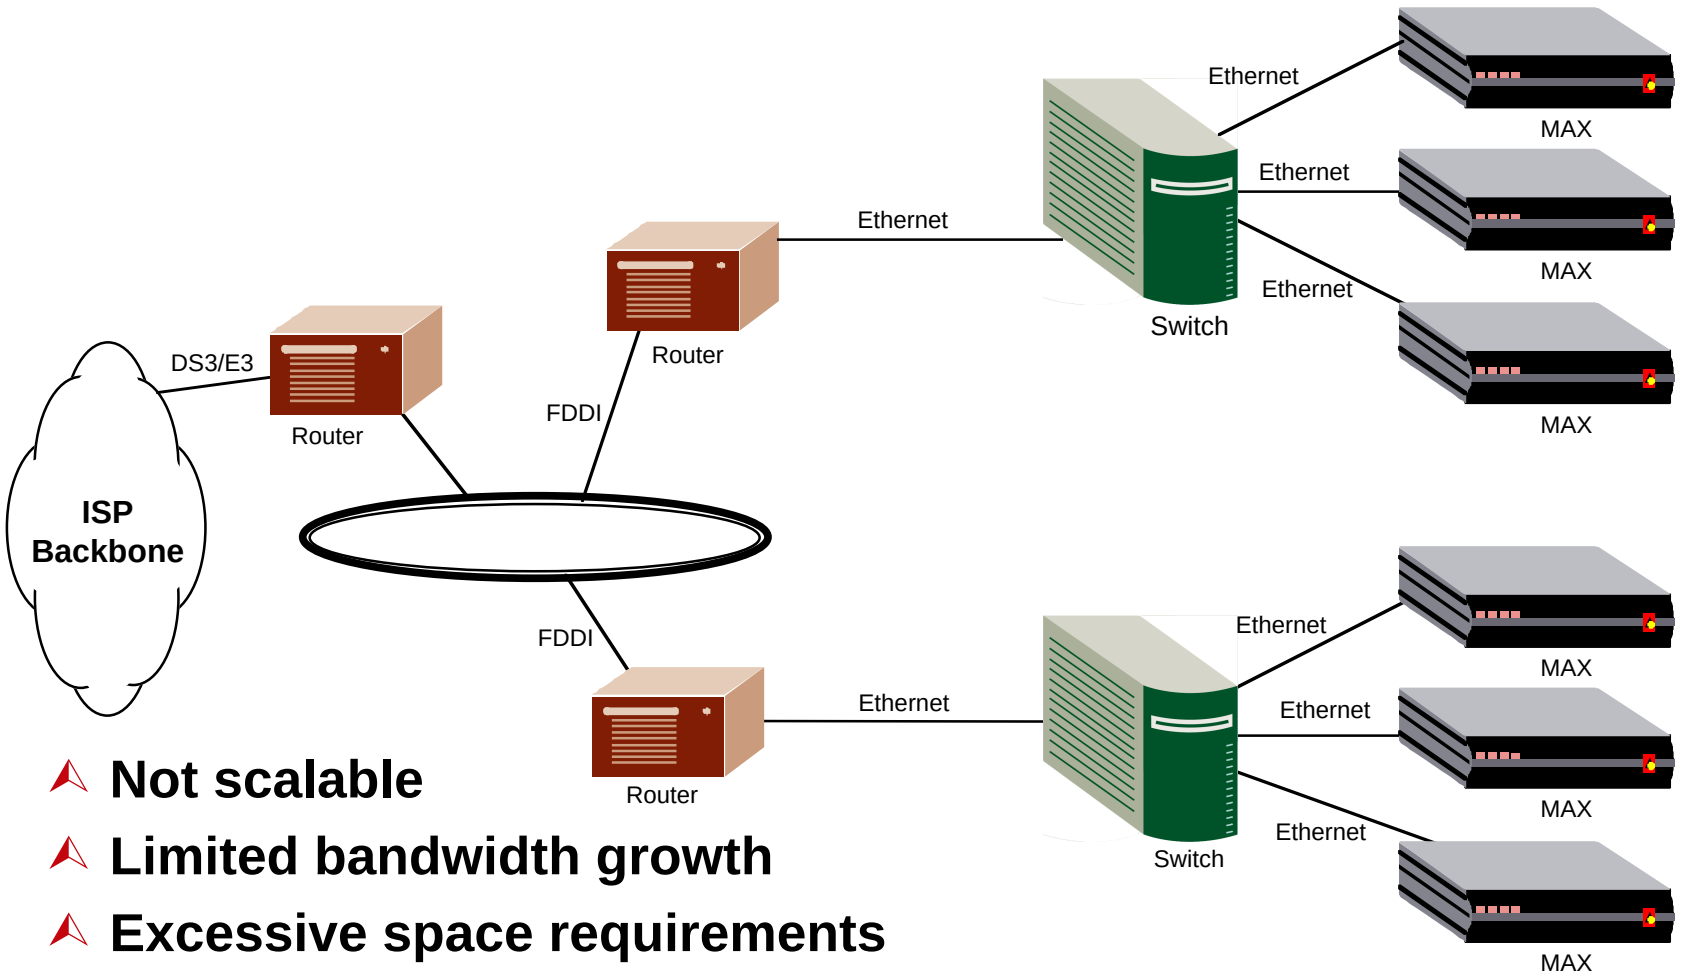

- Not scalable
- Limited bandwidth growth
- Excessive space requirements
- Costly

Applications
Growing the POP

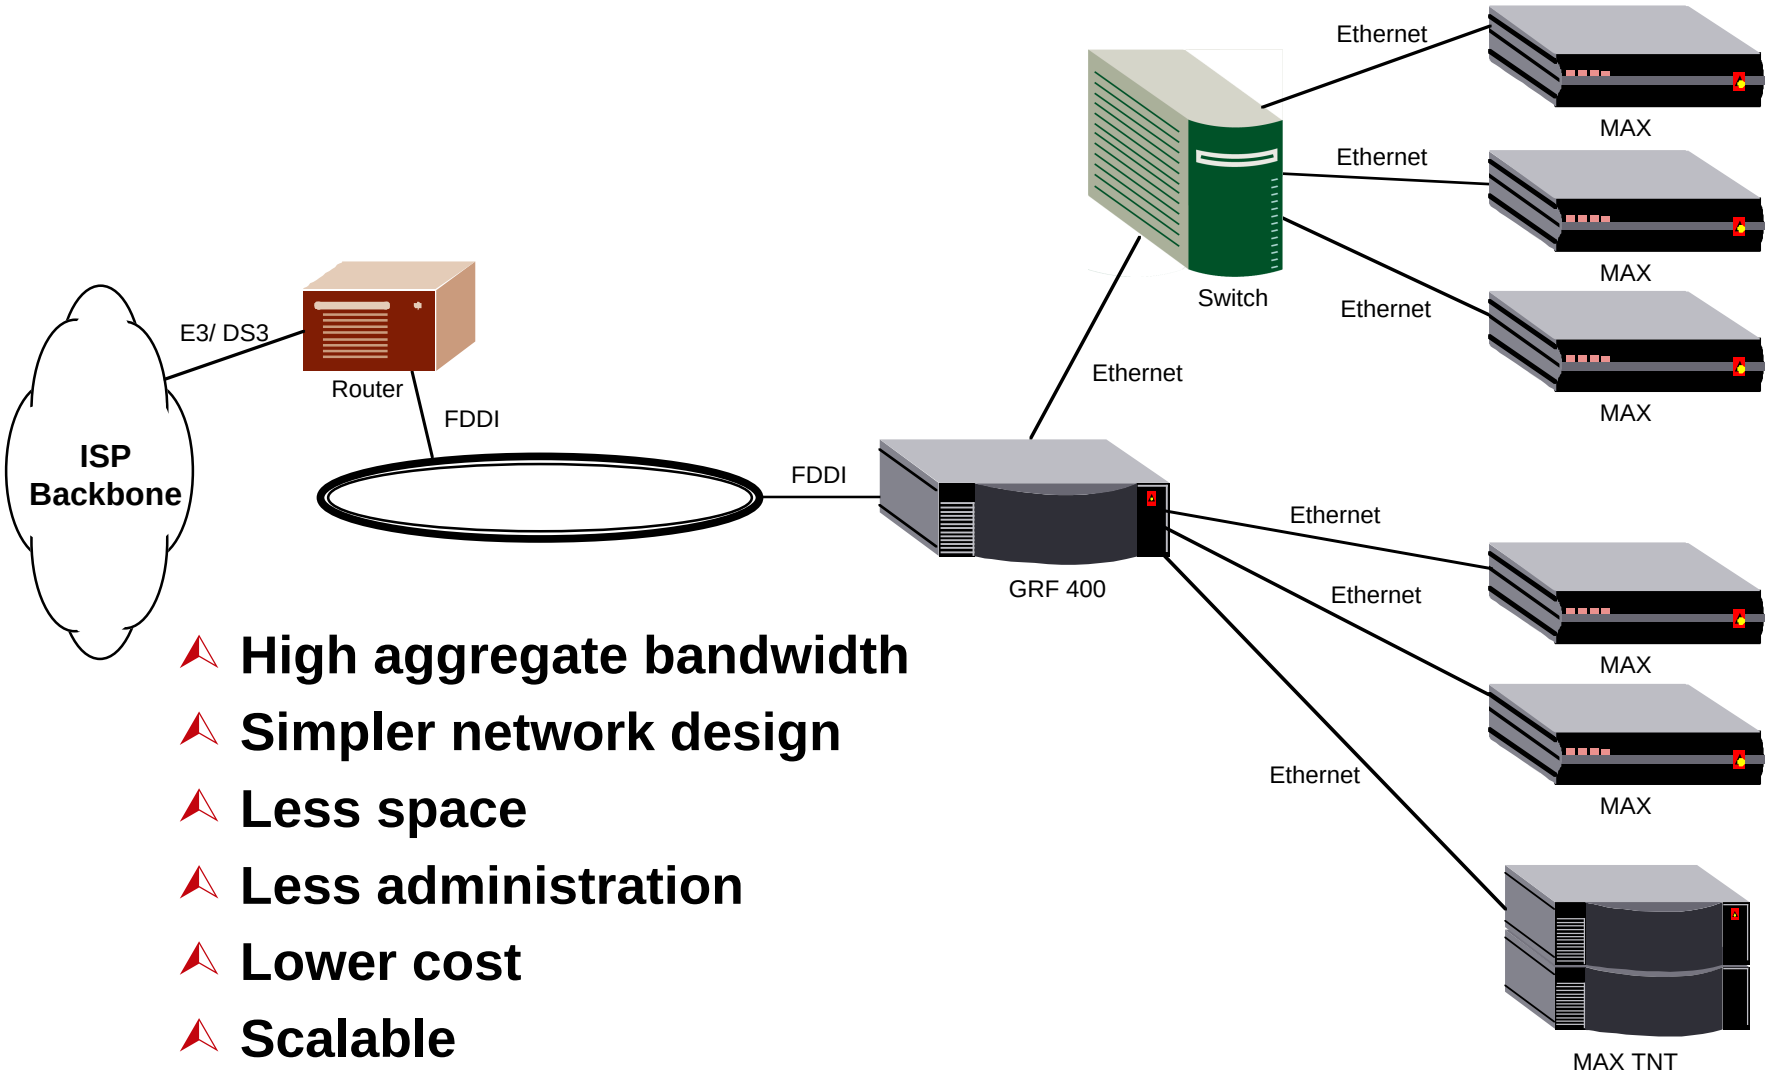

- High aggregate bandwidth
- Simpler network design
- Less space
- Less administration
- Lower cost
- Scalable

Applications
MegaPOP

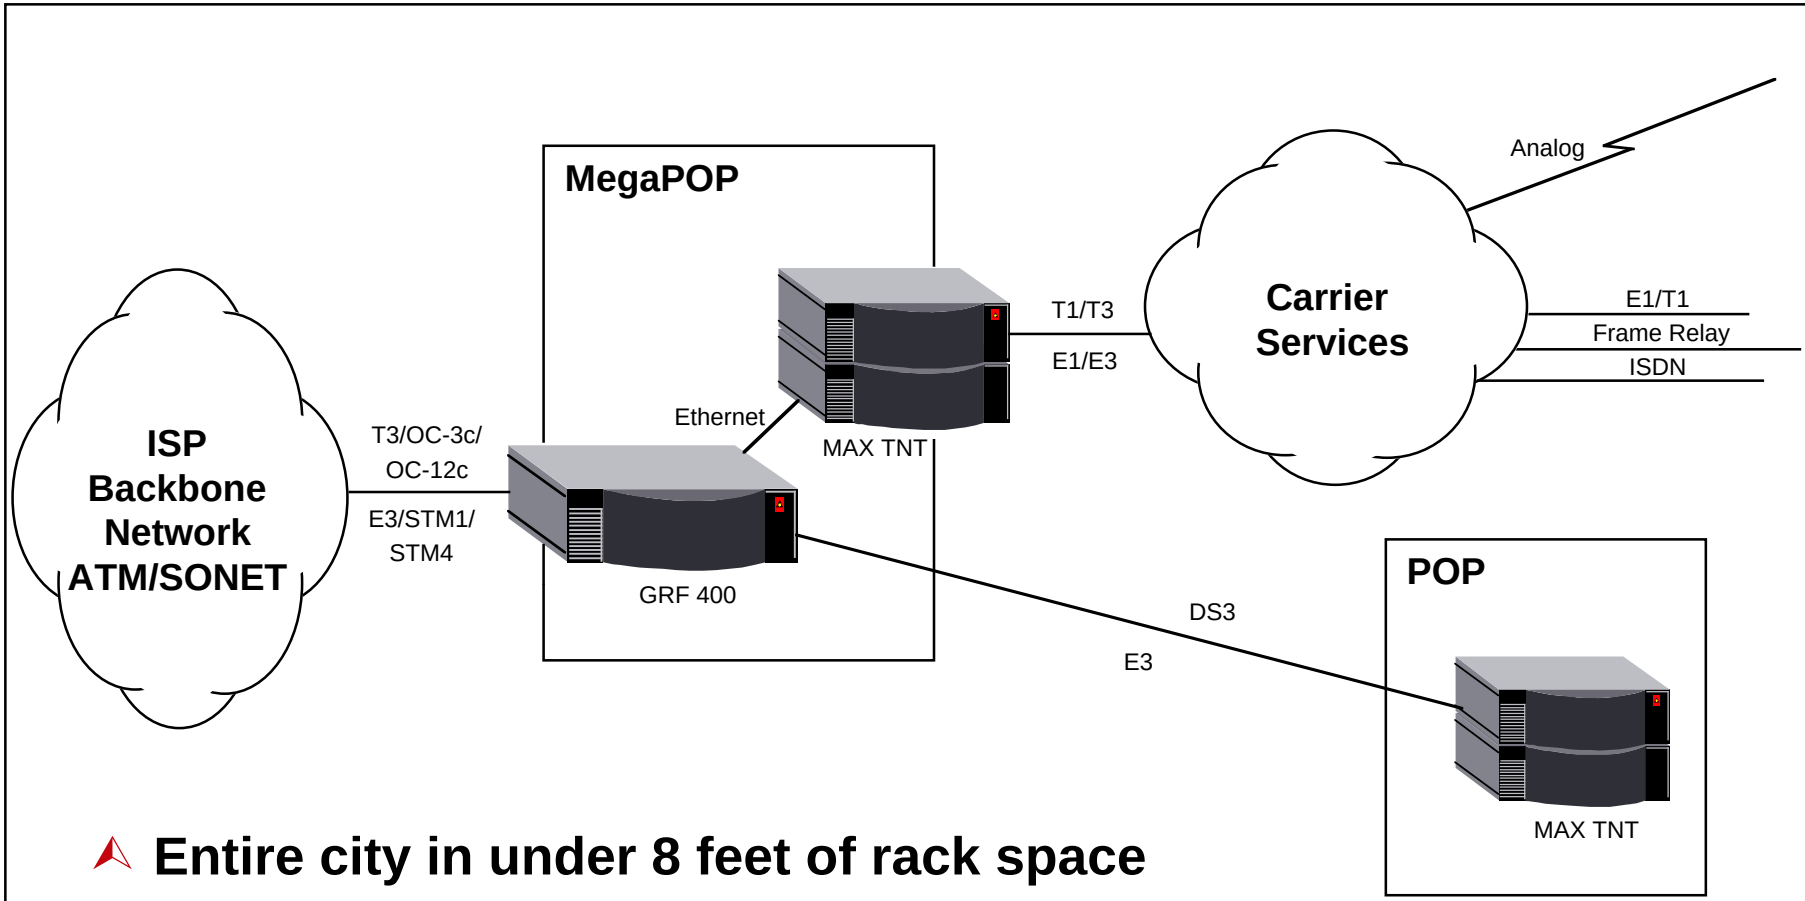

- Entire city in under 8 feet of rack space
- High-speed media to backbone
- ATM or IP over SONET (PPP/Frame Relay)

Applications
GRF in the Backbone

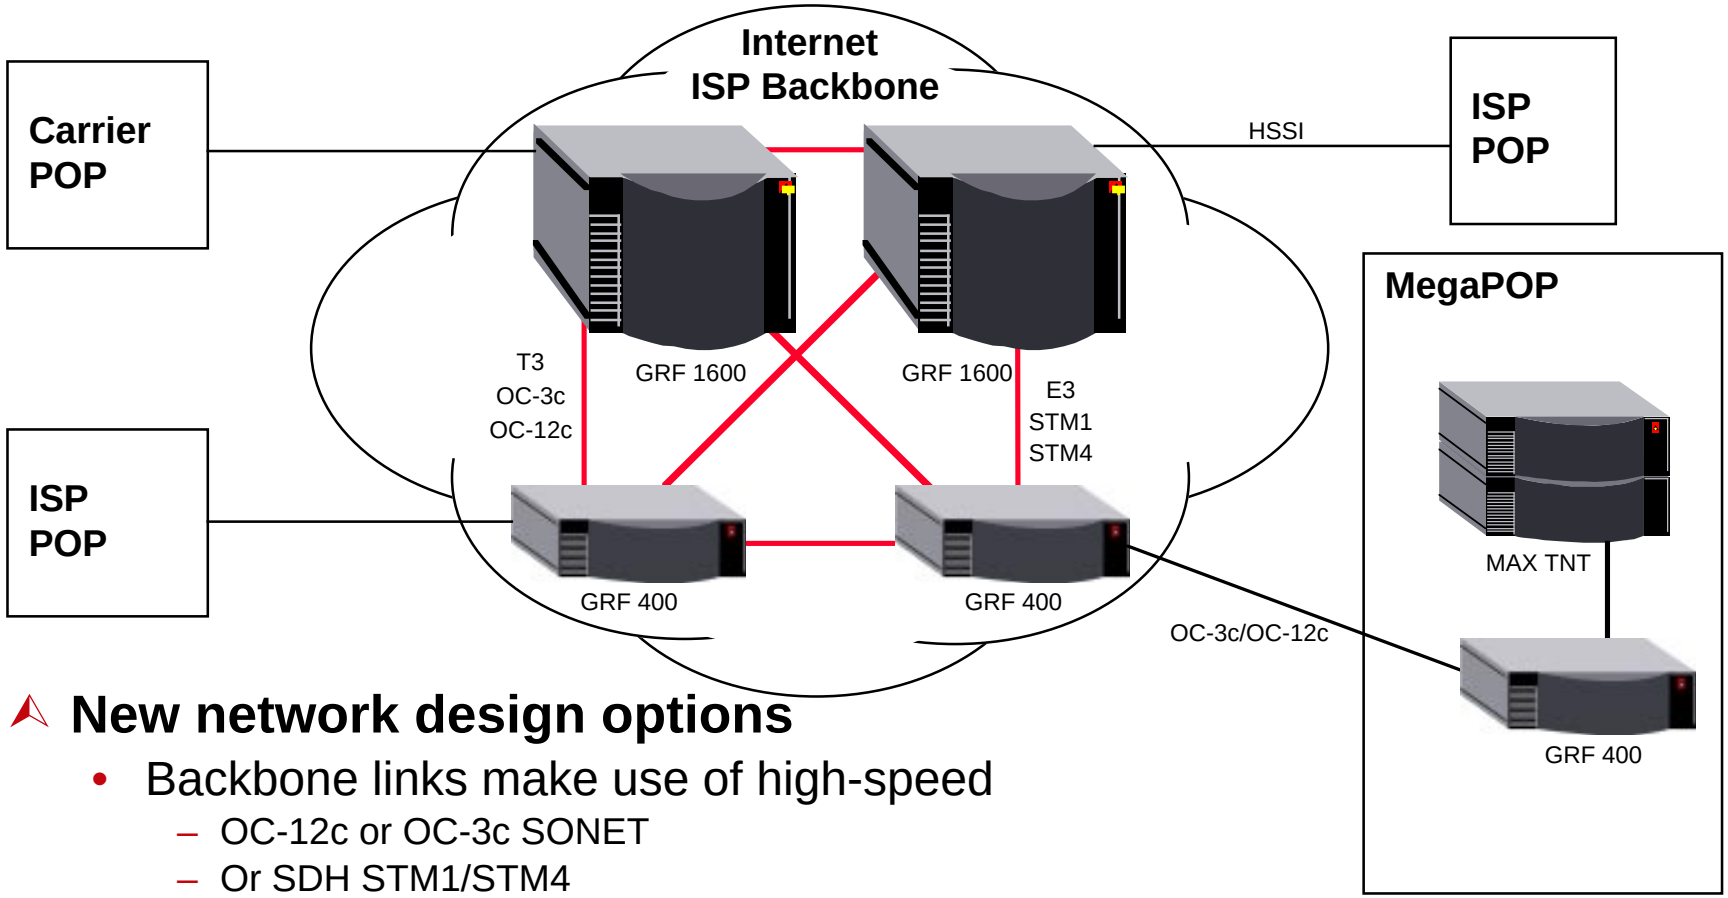

- **New network design options**
  - Backbone links make use of high-speed
    - OC-12c or OC-3c SONET
    - Or SDH STM1/STM4
  - ATM or IP over SONET OC-3c / SDH STM1 options
  - Frame Relay or PPP over SONET OC-3c / SDH STM1 framing options
- **High aggregate PPS**

Applications
NAP - New Options with the GRF

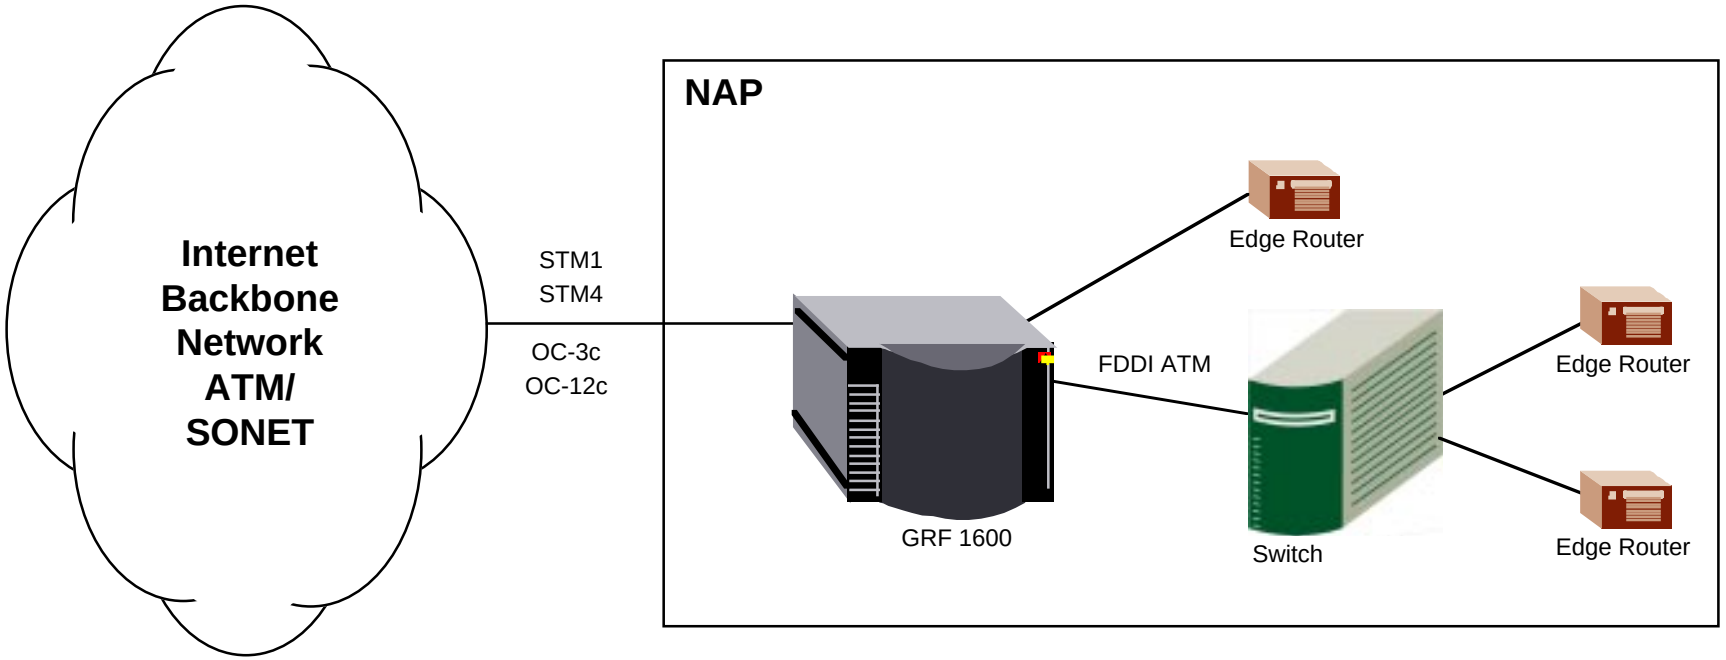

- **Fits into old NAP**
- **Offers new options**
  - OC-12c / STM4
  - Frame Relay or PP over SONET OC-3c / SDH STM1
- **Direct connections to remote peer routers - virtual NAP**
- **No physical co-location needed**

Architecture
Design Objectives

- Compatibility with existing network infrastructures
- Full compliance with industry standards to eliminate need for proprietary gateways or special client software
- IP next-hop address lookup fast enough to take advantage of switch
- Sustainable throughput that is independent of traffic characteristics such as flows and cache hits
- Full wire-speed performance for all external ports

Architecture
Design Objectives – Cont.

- Support for wide variety of popular LAN and WAN media
- Support for ATM without architectural independence upon ATM
- Linear scalability within each IP switch and in a network of IP switches
- Packaging in small chassis to fit into limited space of POP
- Unmatched price/performance

Architecture
GRF 400 Packaging

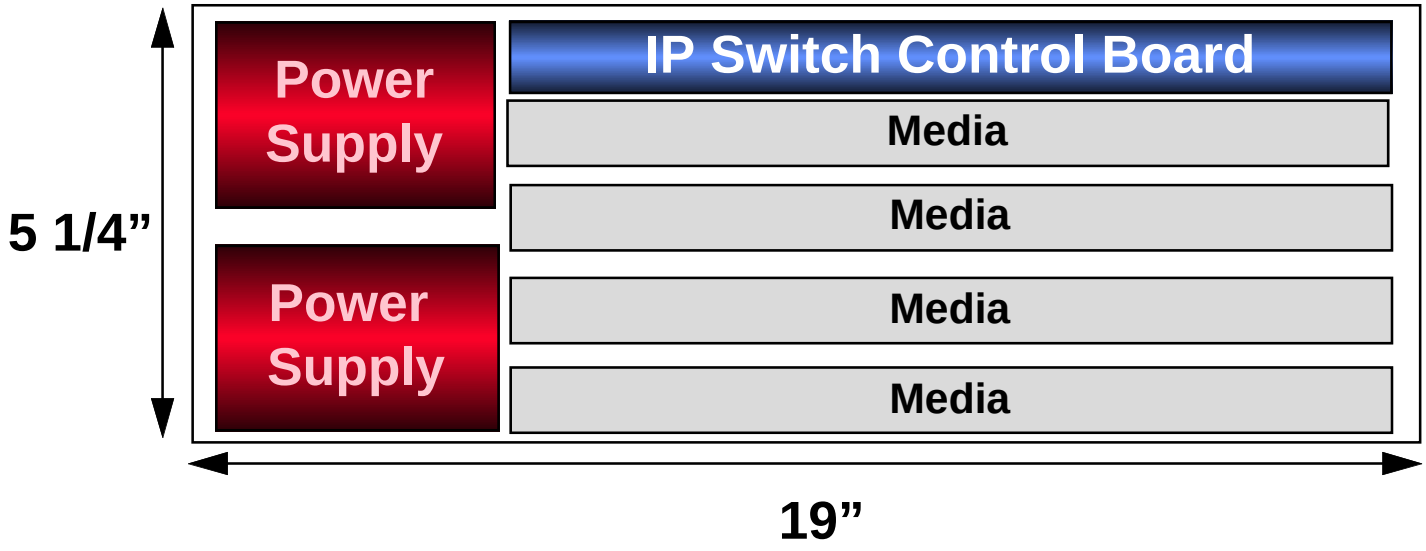

- Up to four IP Forwarding Media Cards
- Dual hot-swappable power supply
- Hot-swappable media cards
- Designed for NEBS compliance

Architecture
GRF 1600 Packaging

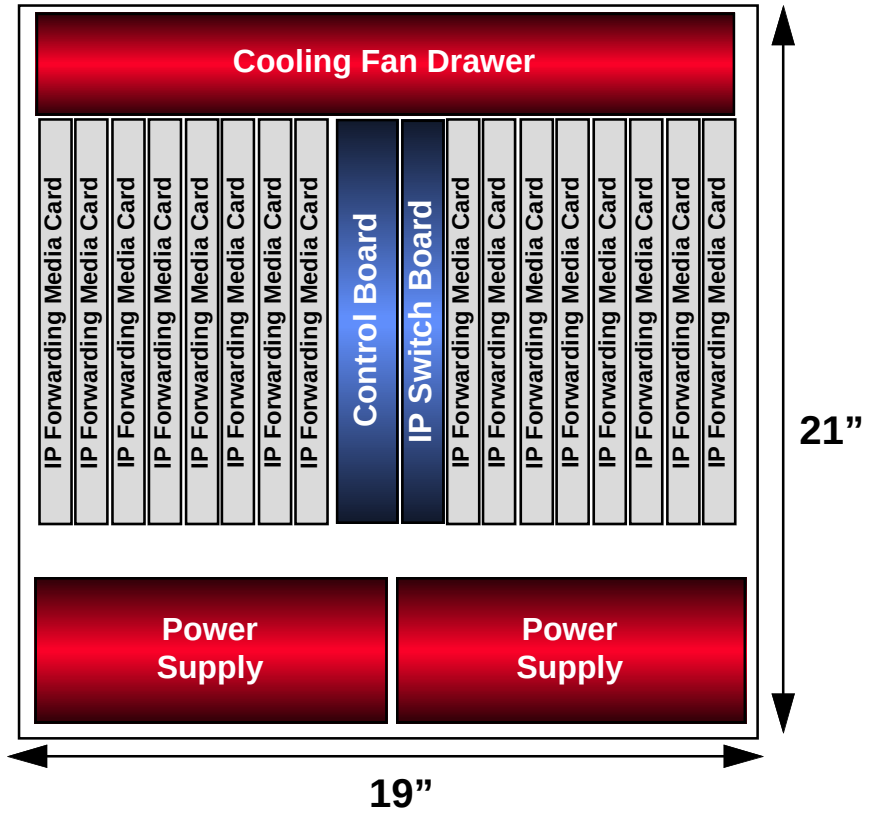

- Up to 16 IP Forwarding Media Cards
- Dual hot-swappable power supply
- Hot-swappable media cards
- Redundant load-balancing fans
- Hot-swappable fan drawer
- Designed for NEBS compliance

Architecture
GRF 400 Functional Diagram

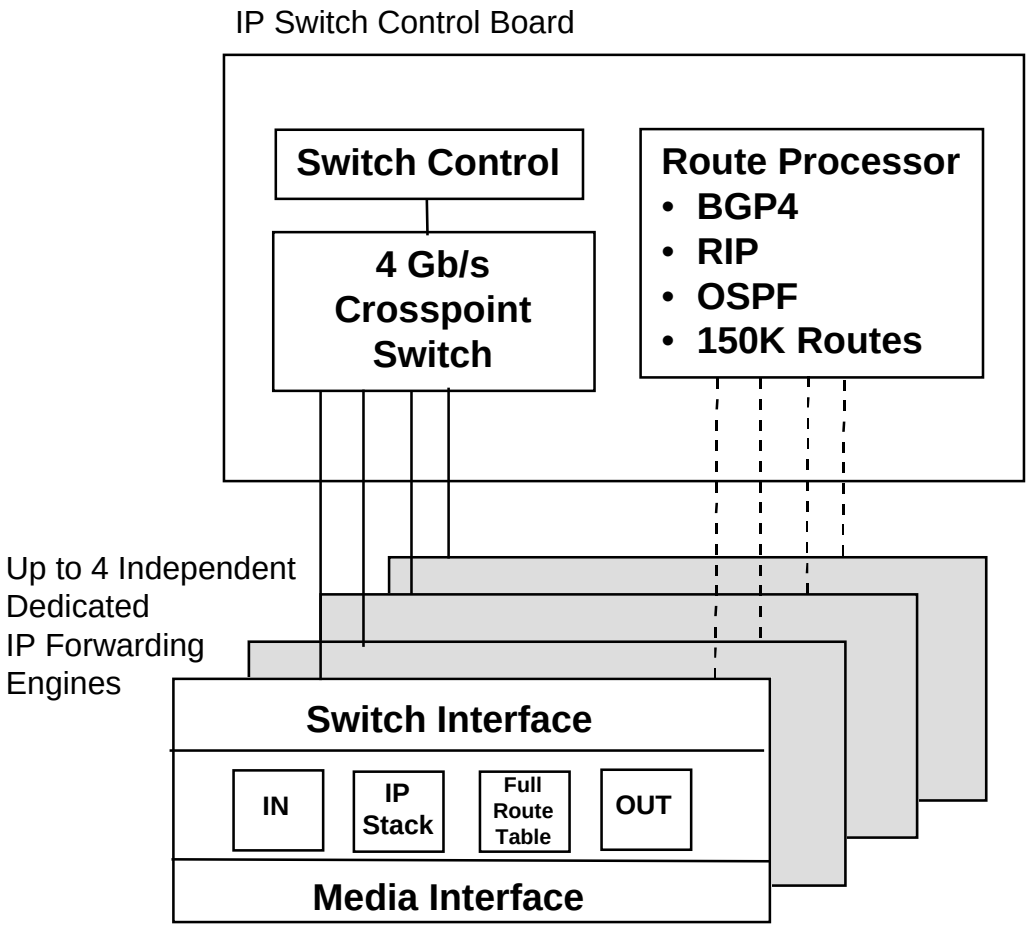

Architecture
GRF 1600 Functional Diagram

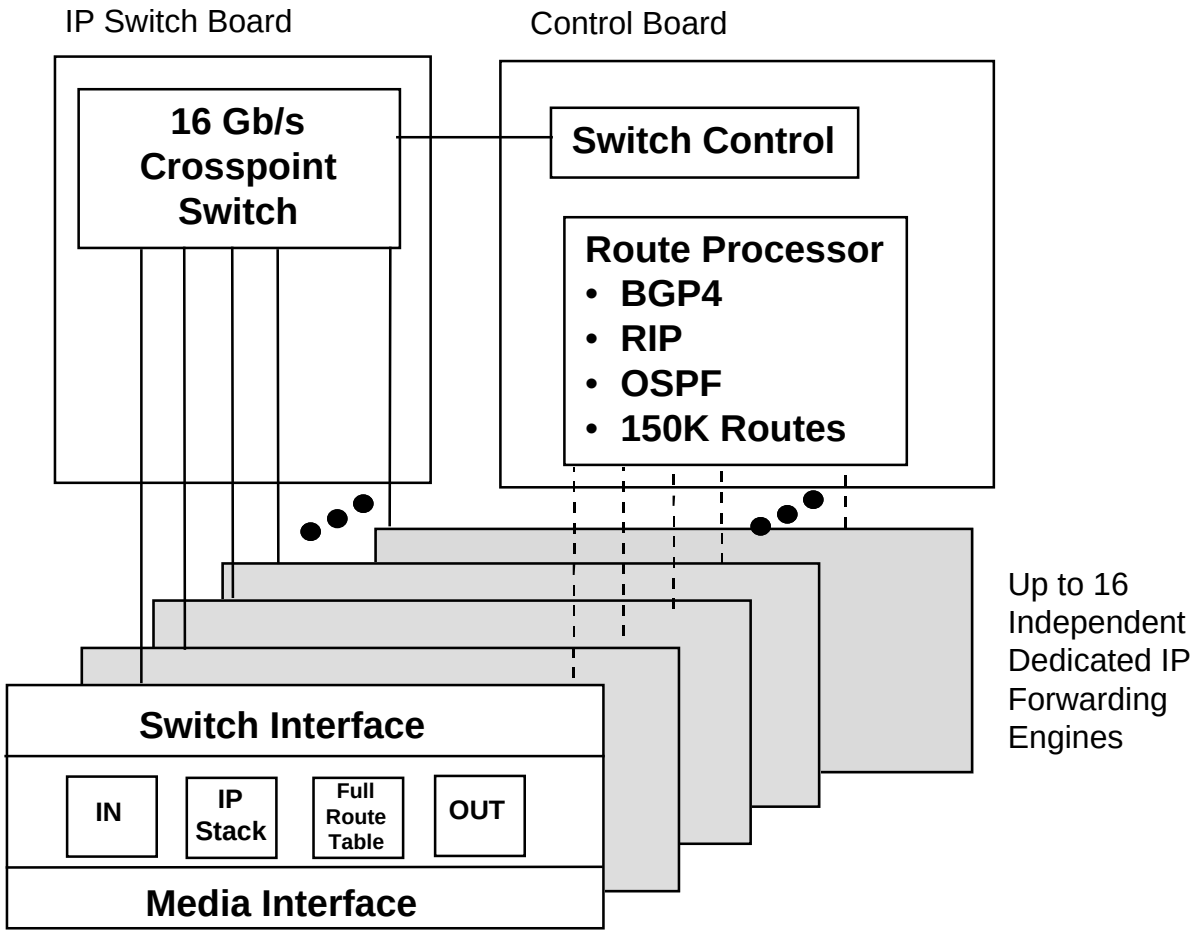

GRF 400 Architecture
IP Switch Control Board

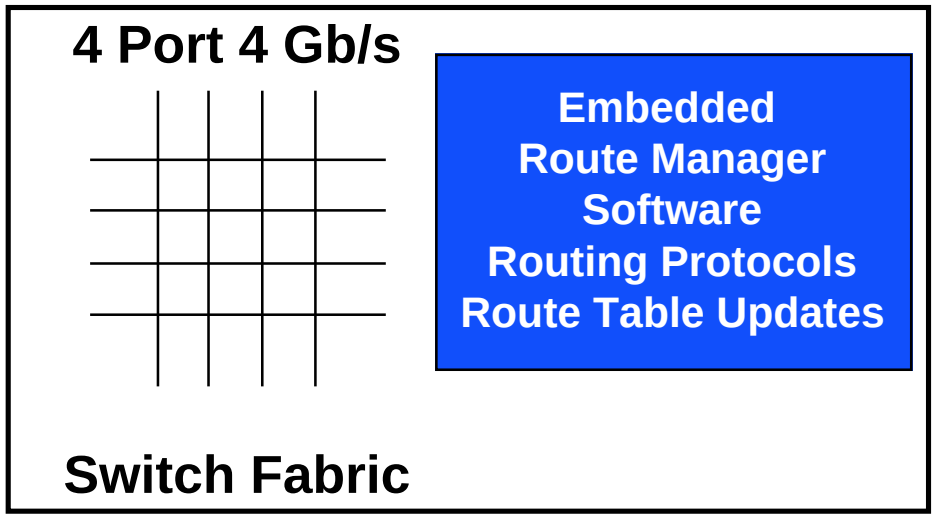

- **Contains 4 Gb/s switch fabric**
- **Network management**
- **Route management software supports routing protocols and route updates:**
  - BGP4
  - IS-IS
  - OSPF
  - RIP
- **Bridging**

GRF 1600 Architecture
Control Board and IP Switch Board

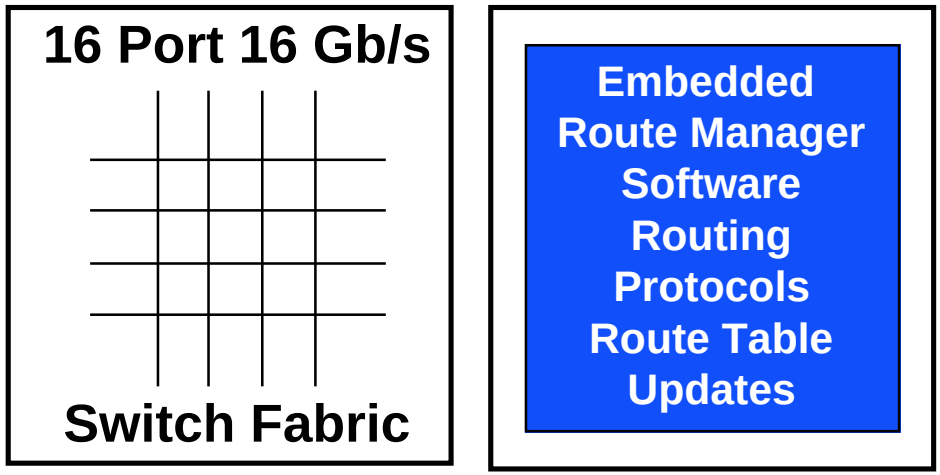

- **IP Switch Board**
  - Contains 16 Gb/s switch fabric
- **Control Board**
  - Network management
  - Route management software supports routing protocols and route updates:
  - BGP4
  - IS-IS
  - OSPF
  - RIP
- **Bridging**

Architecture
IP Forwarding Media Card

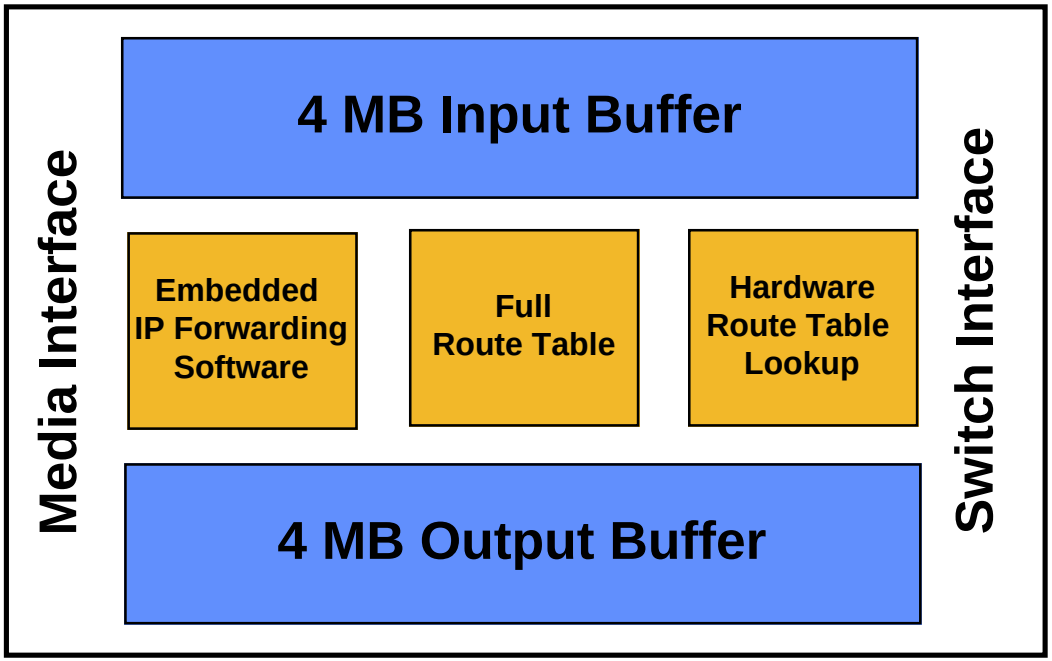

- Up to 4 (GRF 400) or 16 (GRF 1600) IP forwarding media cards
- Complete Layer-3 IP forwarding engine
- Each card has dedicated 1 Gb/s connection to switch
- Full route table up to 150K routes
- Route table hardware lookup next-hop found in under 2.5 microseconds

Architecture
Media Cards

Card	Ports	Speed
Ethernet	8 Ports & 4 Ports	10/100 Mbps (autosensing)
FDDI/CDDI	4 Ports	100 Mbps
HSSI	2 Ports	52 Mbps
OC-3c ATM/SDH STM1	2 Ports	155 Mbps
IP/SONET OC-3c/ SDH STM1	2 Ports	155 Mbps (Frame Relay & PPP Framing)
OC-12c ATM/SDH STM4	1 Port	622 Mbps
HIPPI	1 Port	800 Mbps

Architecture
System Management

- Supports standard and propietary MIBs for puts, gets and traps
- Accessible from SNMP management packages to fit into current management strategy
- Administrative authentication using RADIUS
- Command-line configuration tools

Competitive Analysis

- Overview Conventional Architecture
- Comparing Architectural Options
- Scalable performance
- Questions to Ask Your Router Vendor

Competitive Overview

Limitations of Conventional Router Architectures

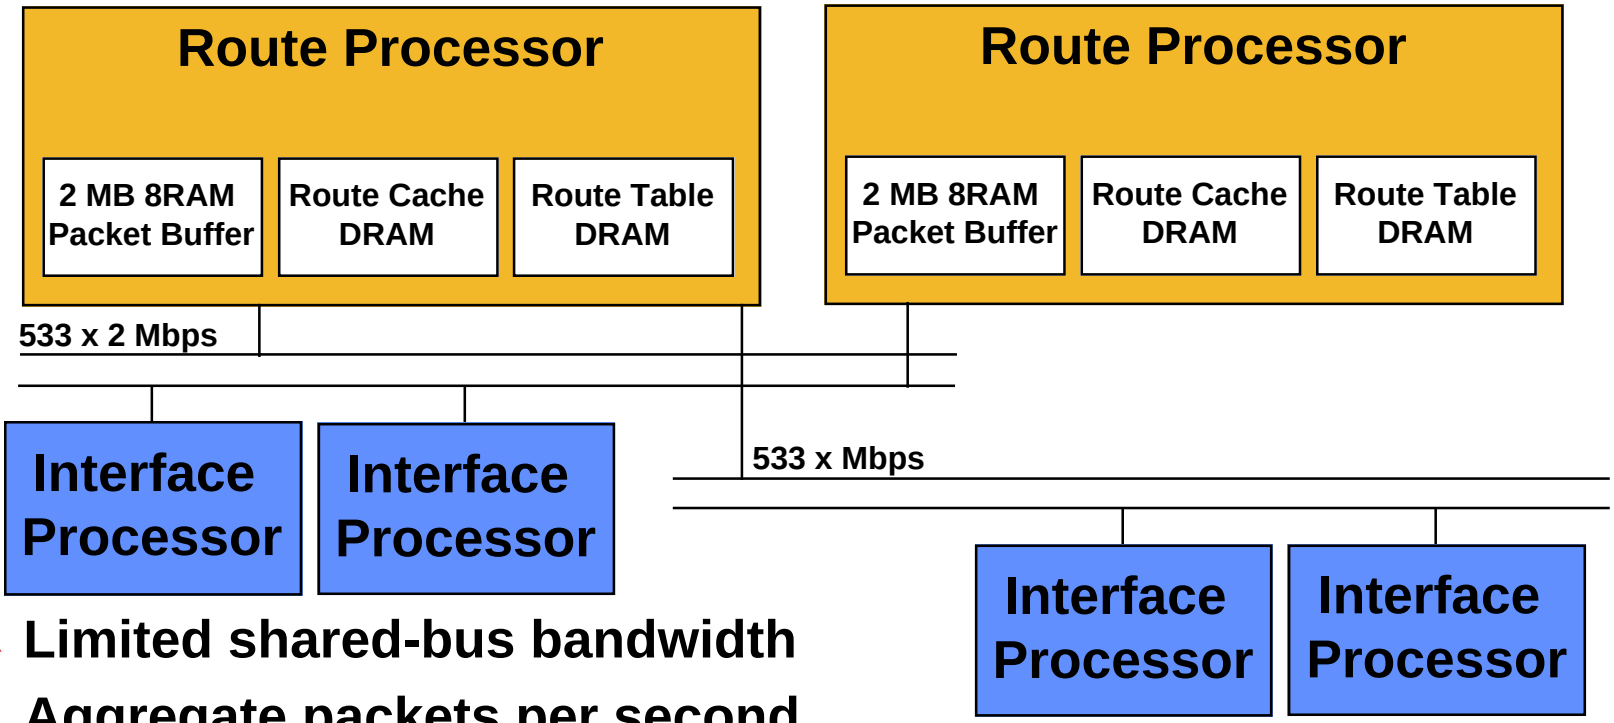

- Limited shared-bus bandwidth
- Aggregate packets per second
- Cached routing
- Software route-table look-up
- Overloaded CPU

Competitive Overview
Comparing Architectural Options

Conventional Router	GRF	GRF Benefit
Media cards depend on central processor for packet forwarding	Each media card is a complete packet forwarding engine	Performance scales linearly
Route caching	Each media card has full route table with all router	Performance remains constant in large dynamic networks
Next hop found by S/W table lookup performed by single shared central CPU	Next hop found by H/W table lookup on each card	- 100 times as fast - Enables use of switching architecture - Multiple CPUs not overloaded
Shared parallel bus aggregate 2 Gb/s	Switch bandwidth aggregate 16 Gb/s	Speed and bandwidth allows line speed packet forwarding using rich Layer-3 header information
Bus Architecture limits bandwidth (PPS)	Support for multiple high-speed media	Multiple OC-12c; OC-3c, FDDI, CDDI, Ethernet, HSSI enabled in one box
Flow characterization assumes well-behaved traffic patterns. Adds demand to CPU. 10% performance improvement	Full route table hardware lookup in 1 microsecond	Performance not dependent upon traffic patterns

Competitive Overview
The GRF Delivers Scalable Performance

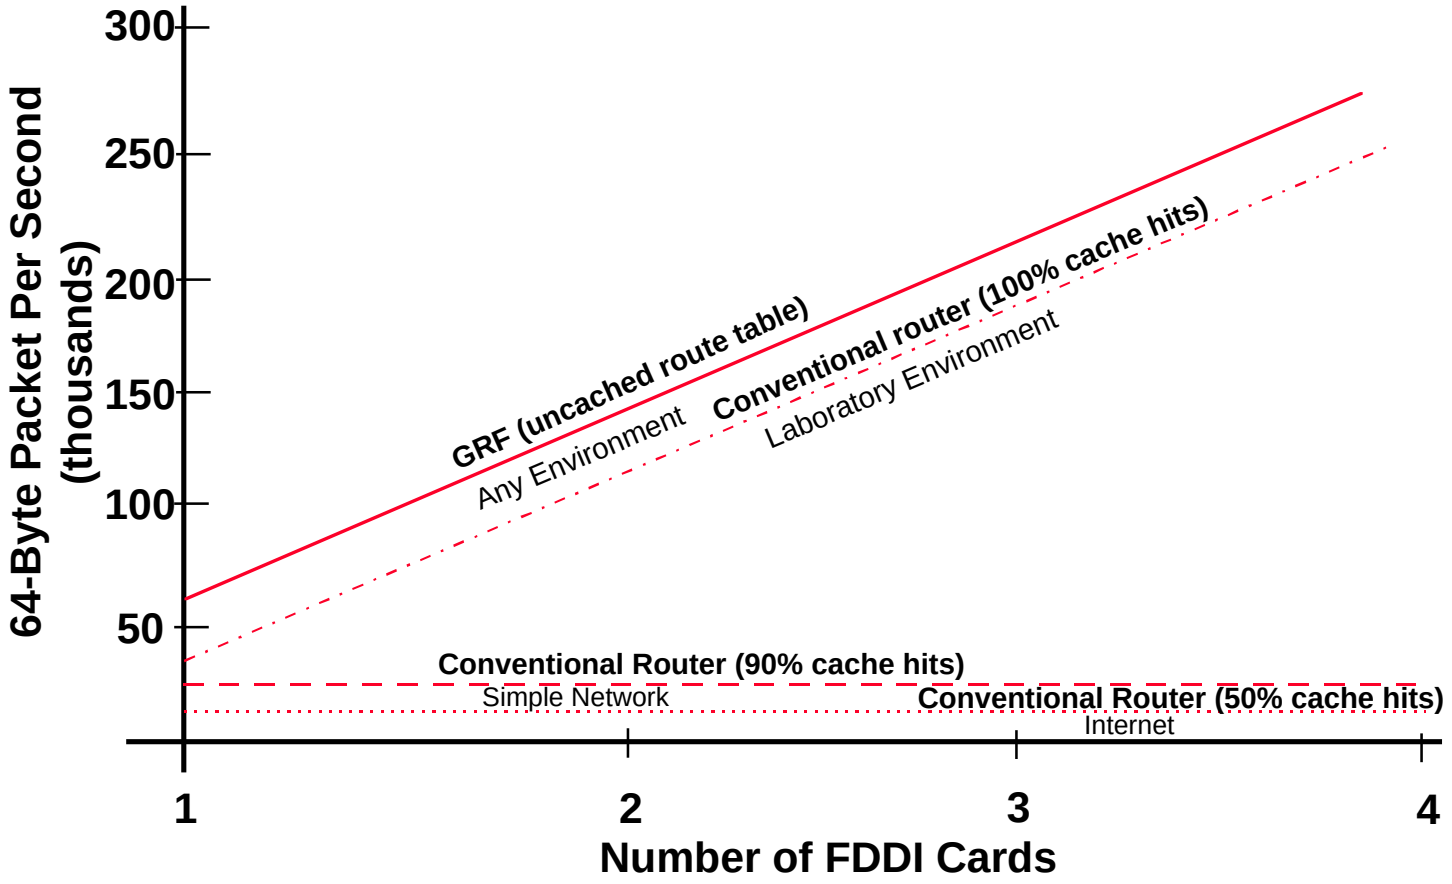

Competitive Overview
Questions to Ask Your Router Vendor

- **What happens when there are random IP destination addresses and associated cache misses?**
  - Their answer - Performance drops to less than 10% of advertised performance
  - Ascend's answer - No effect
- **What conditions will saturate CPU utilization?**
  - Their answer - 2 cards, 30K pps incoming or caching
  - Ascend's answer - No effect
- **What happens when CPU is saturated?**
  - Their answer
    - Peering sessions are dropped
    - Packets dropped
    - Aggregate performance drops to 1% of advertised
    - Console locks up
    - Keyboard locks up
    - Router Panics and Reboots
  - Ascend's answer - No effect

Competitive Overview
Ipsilon

- **Cell is not an IP packet**
  - 5.3 MCS is really 500 KPPS
- **PC routers are slow and vulnerable to overload**
- **BGP4 is not supported**
- **Special proprietary SW is required**
  - Flow Management SW in nodes
- **Special gateware required for Ethernet and FDDI**
- **Only works with ATM**
- **Not scalable**
- **Overall performance is limited by Amdahl's law**

Amdahl's Law

Combining Slow and Fast Processes to Do the Work

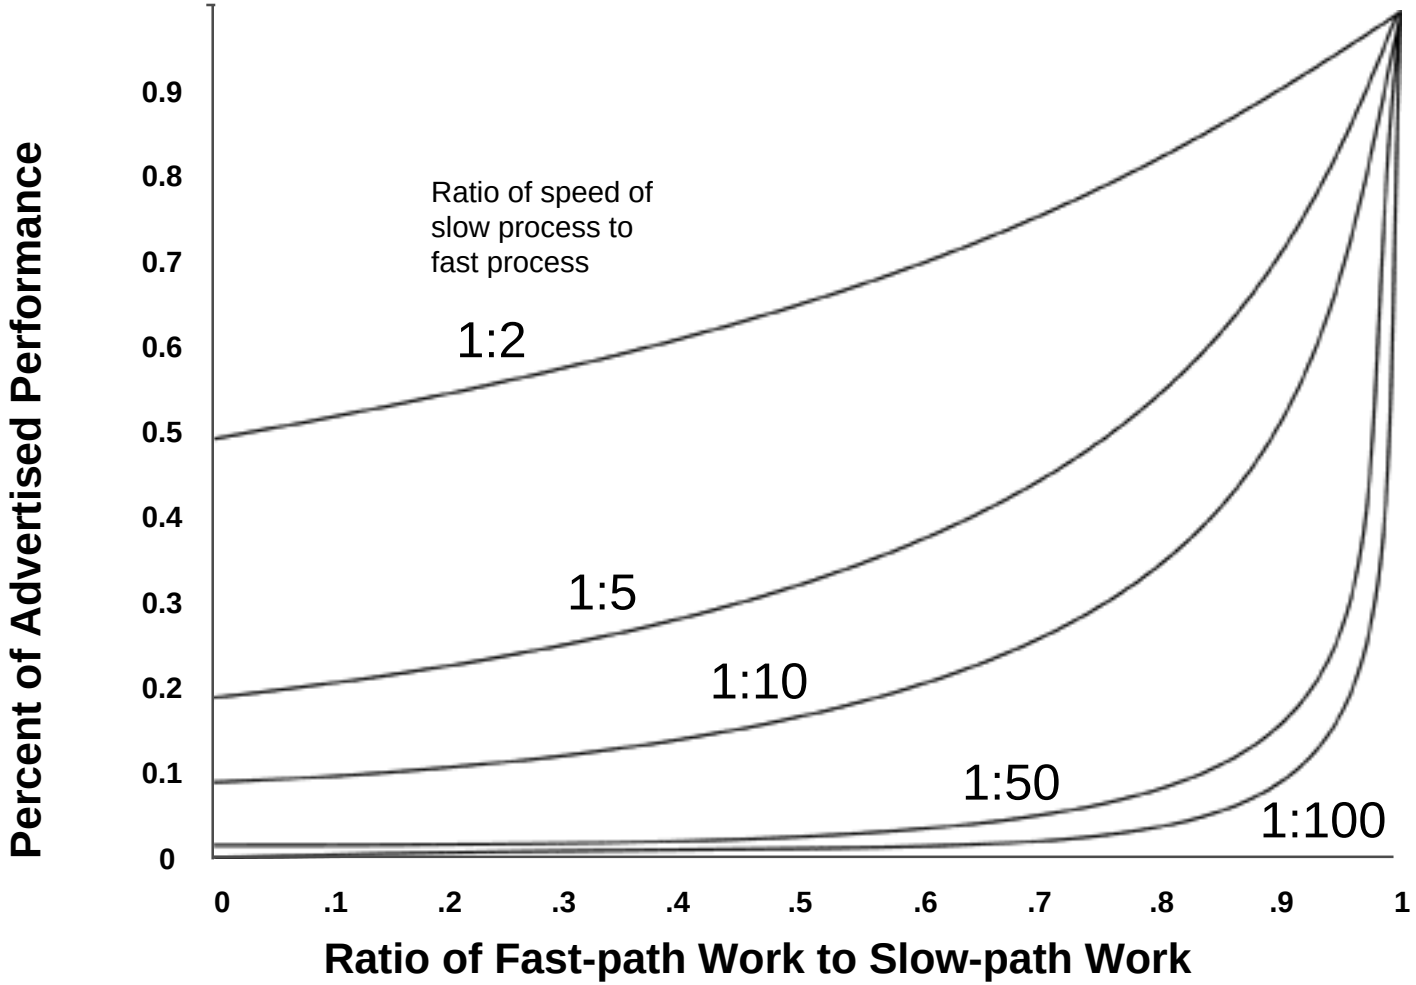

Summary

We look at every packet.

Summary

Because we can.