Converging Networks

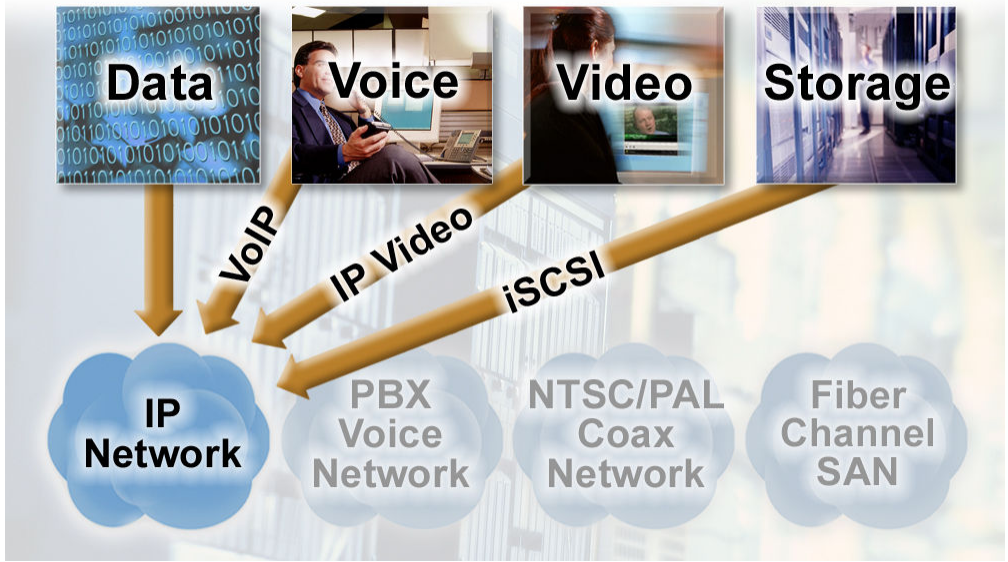

From a Technology View

Günter Honisch
Distinguished Engineer
EMEA Consulting
ghonisch@cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

Agenda

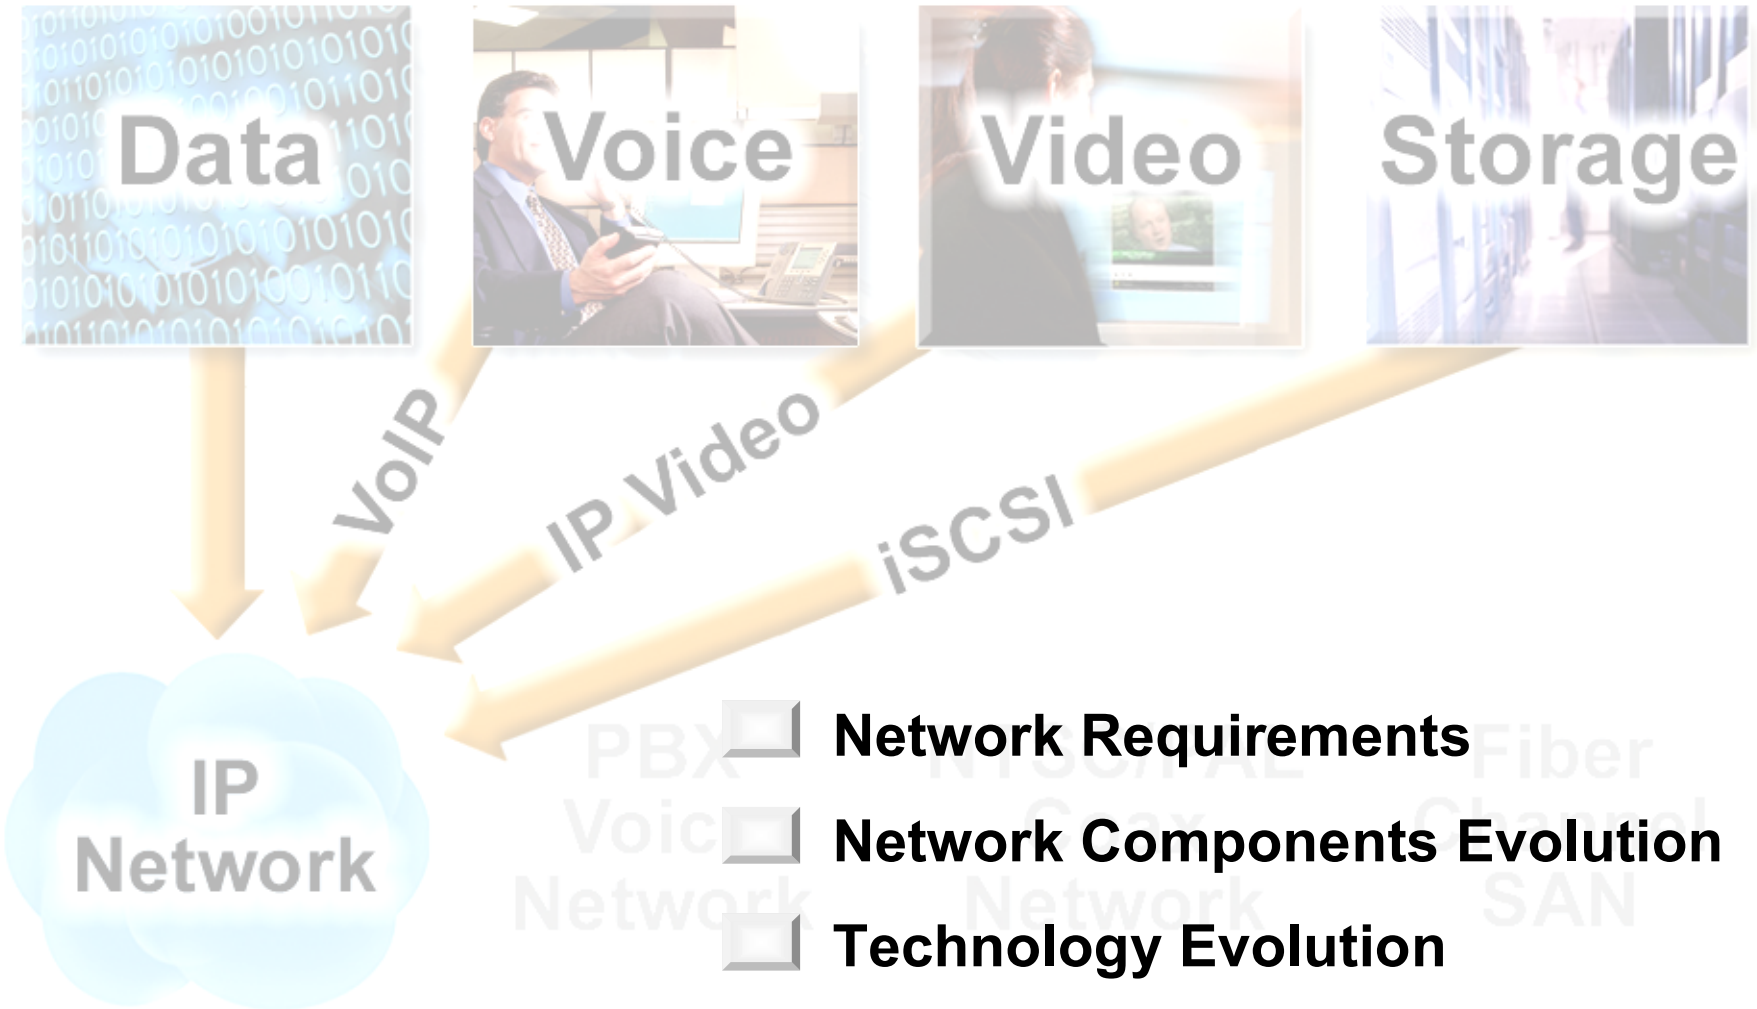

- Network Requirements
- Network Components Evolution
- Technology Evolution

© 2003, Cisco Systems, Inc. All rights reserved.

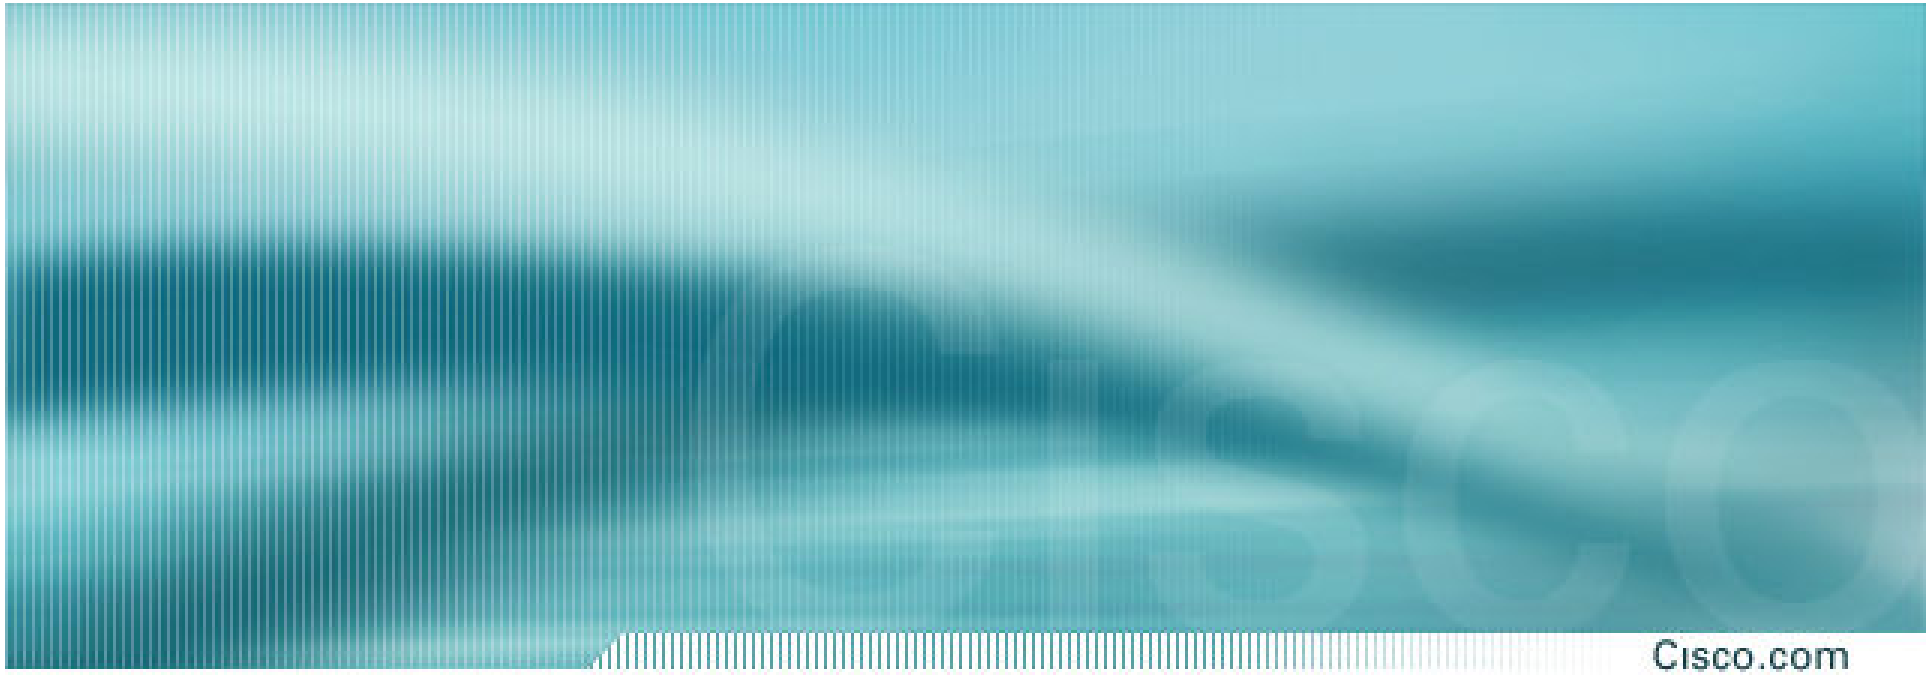

Network Requirements

© 2003, Cisco Systems, Inc. All rights reserved.

Technology Drivers

- **Converged networks**
  - Voice, Video, Data, Storage
  - Driving requirements for greater intelligence & scaleability
- **Desktop & Server computing power increasing**
  - GE, 10GE fullspeed connections
- **Advanced applications with challenging traffic patterns**
  - Greater reliance on QoS to manage increased bandwidth
- **Mission-critical security**
  - Networks are more open
  - Threats ever-increasing

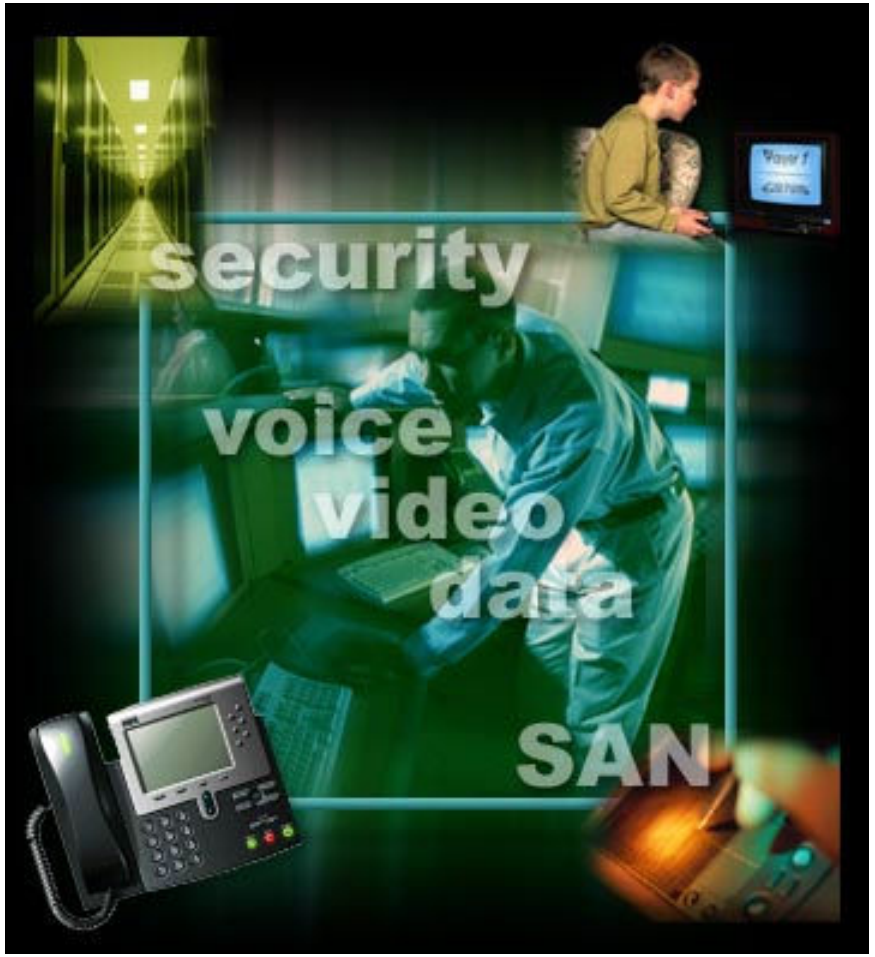

© 2003, Cisco Systems, Inc. All rights reserved.

Network Traffic – what is going on ?

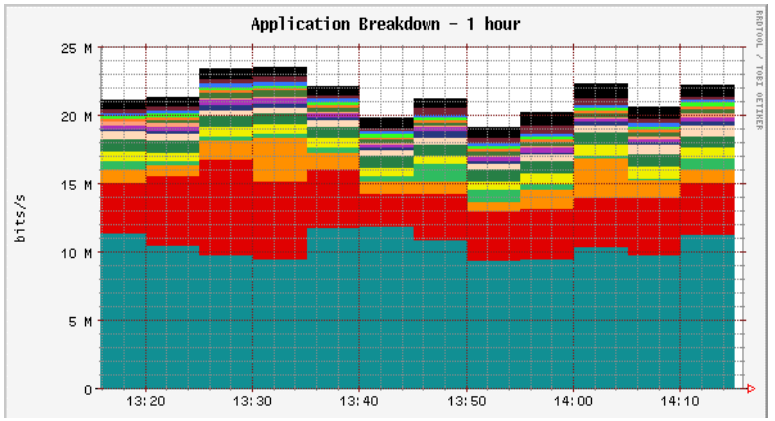

- **>90% of sessions have <10 packets**
  - **Transaction mode (mail, small web page)**
  - **Stresses flowbased mechanisms in the network**
- **>70% of all TCP traffic results from <10% of the sessions, in high rate bursts**
  - **High speed aggressive flows need solid behaviour of forwarding infrastructure (Buffers, Tail drop vs xRED)**
  - **TCP & UDP at high speeds are challenging to integrate**

© 2003, Cisco Systems, Inc. All rights reserved.

Applications Today

Cisco.com

- **What is important:**
  - **Varies by network and population**
- **Key applications:**
  - **The Web+Mail**
  - **Large file transfers**
  - **Interactive access**
  - **Streaming audio/video**
  - **ERP**
  - **Games**

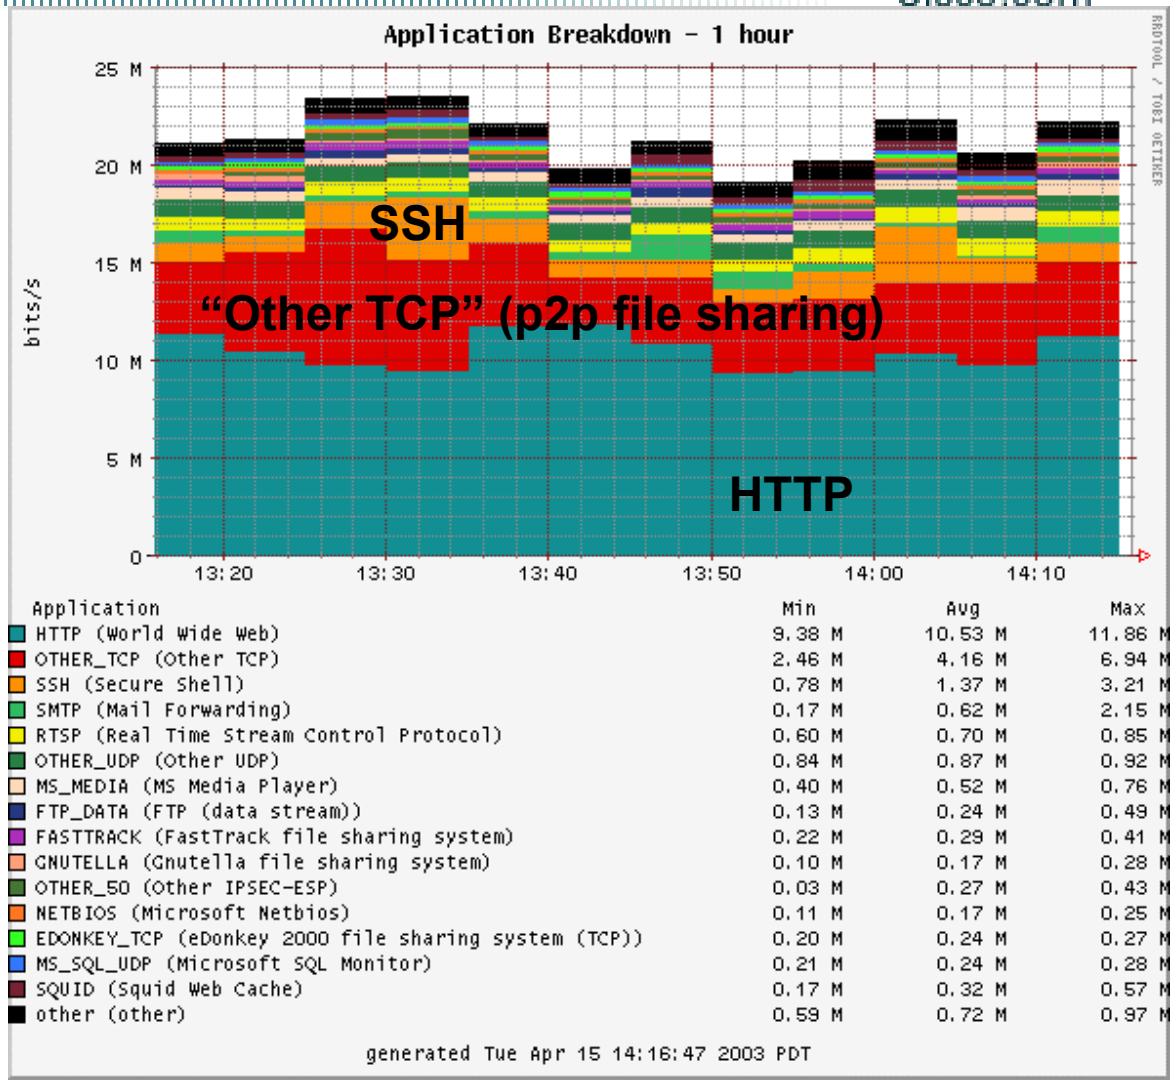

http://www.caida.org/dynamic/analysis/workload/sdnap/0_0_/ts_top_n_app_bytes.html

9:31 AM 16 April 2003

© 2003, Cisco Systems, Inc. All rights reserved.

GE on the Desktop ?

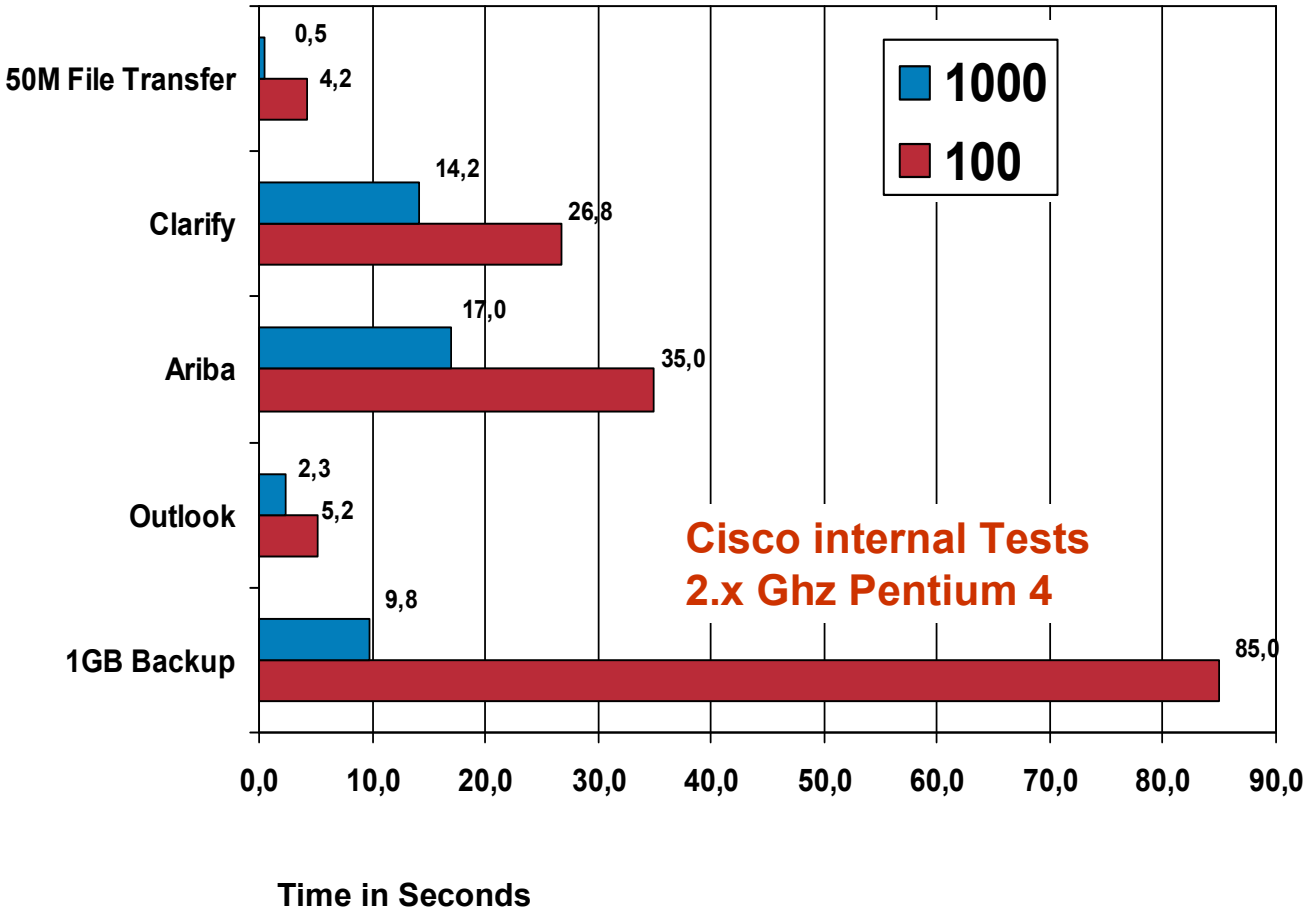

With new Chipsets (Memorybus NIC) observed Bursts up to 950 Mbps

Offloading NICs perform GE speed at very low CPU Utilisation

Stresses buffering and Congestion Handling of NW Components

© 2003, Cisco Systems, Inc. All rights reserved.

Gigabit and NO QoS…

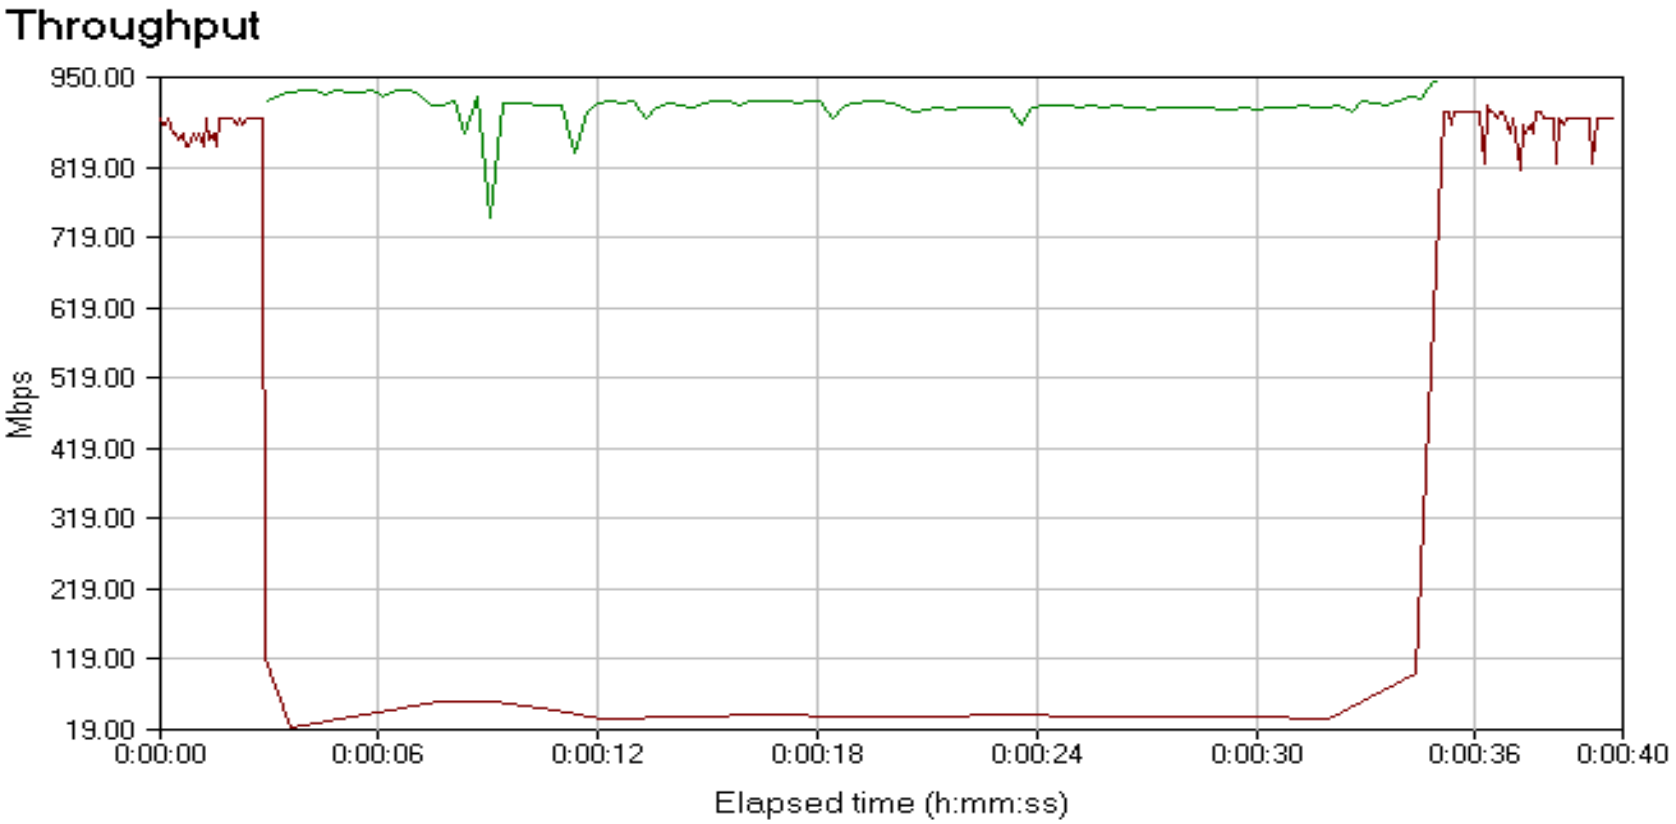

- TCP (red) and UDP (green) Streams with Gig attached hosts
- 1GB data using TCP with 0MB Loss
- 3.7GB data using UDP with 23MB lost
- 15K of 22.5M datagrams lost - 154 max consecutive loss

© 2003, Cisco Systems, Inc. All rights reserved.

Storage over IP

- **Ubiquitous**: Access storage from campus, MAN, and WAN
- **Agnostic**: Single access technology for block and file
- **Scalable**: Distance, node count and performance
- **Economical**: Utilize IP / GE infrastructure and expertise
- **Traffic:** very high Bandwidth up to ~ 10Gbps in the future
- **QOS:** ideally sub Millisecond…

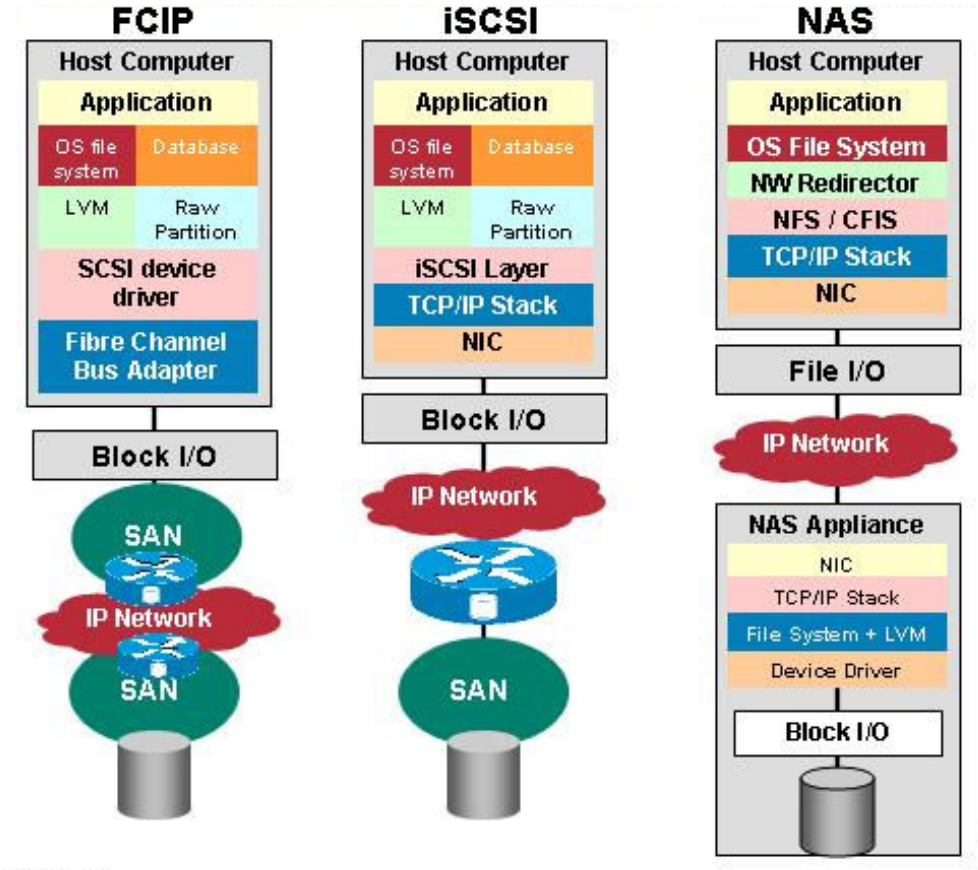

© 2003, Cisco Systems, Inc. All rights reserved.

Increasing Requirements

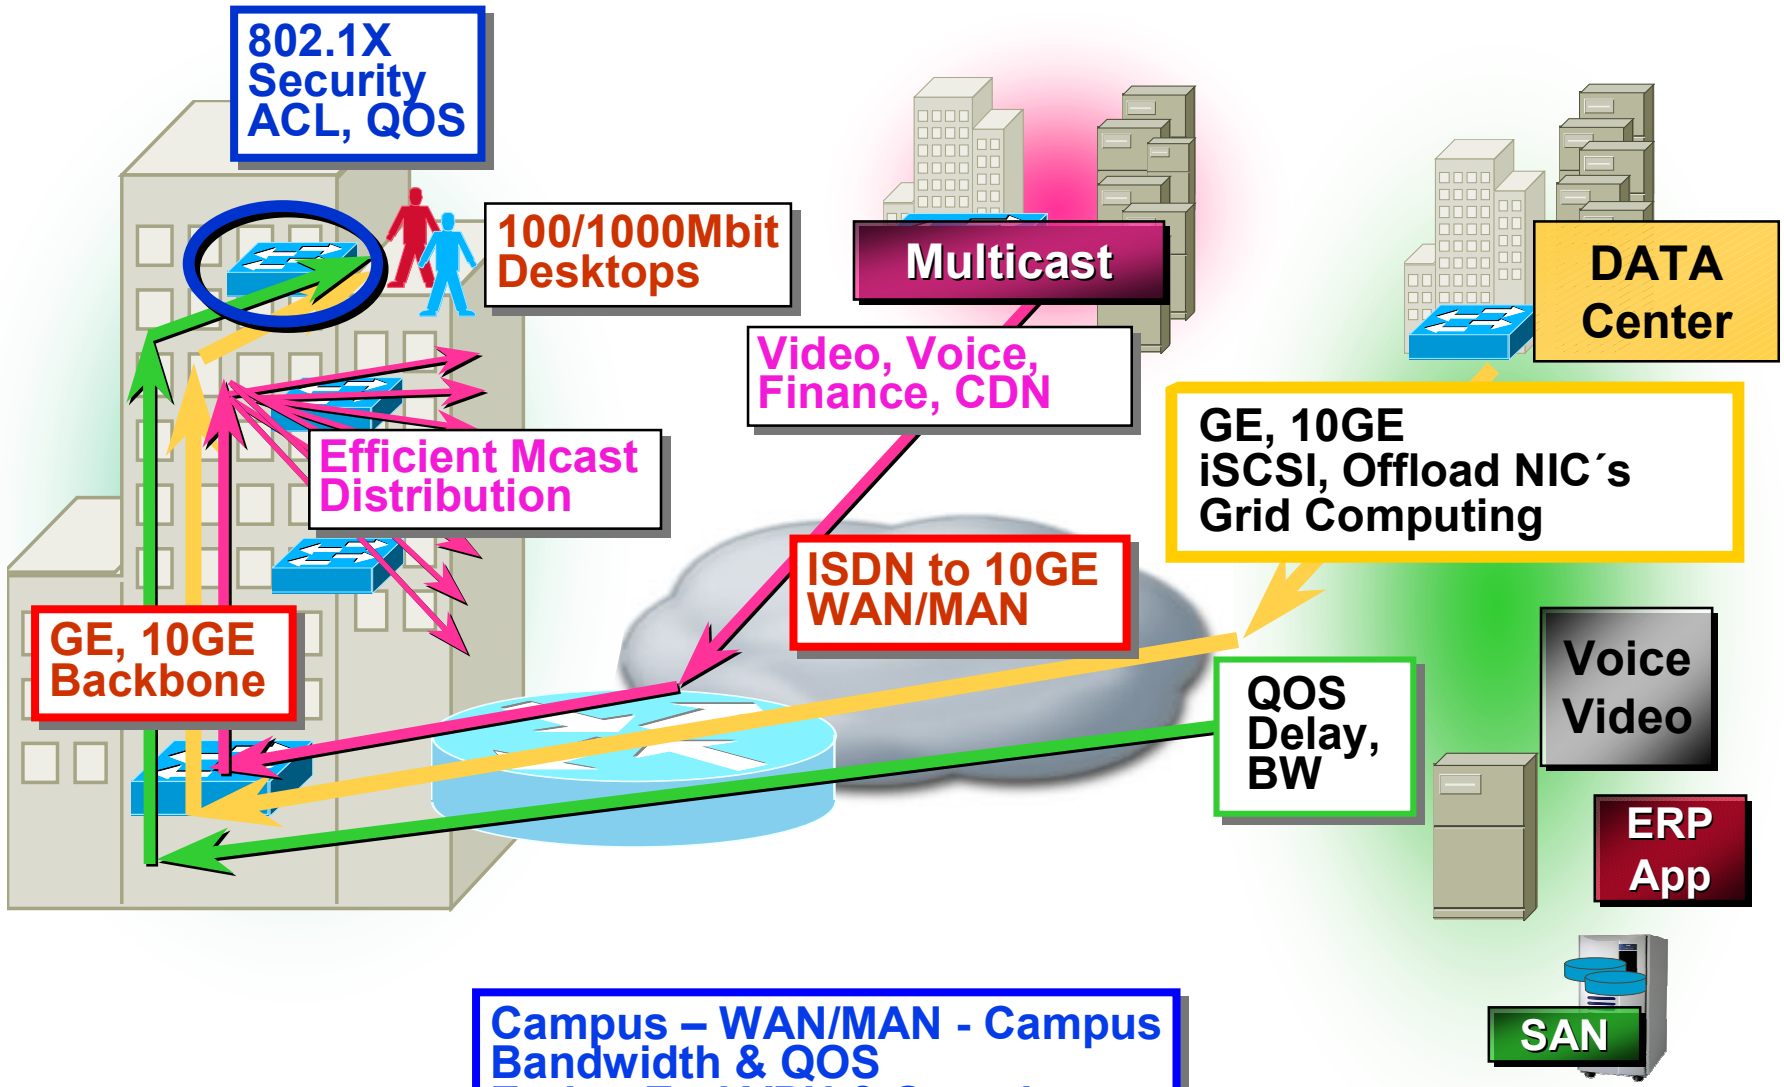

© 2003, Cisco Systems, Inc. All rights reserved.

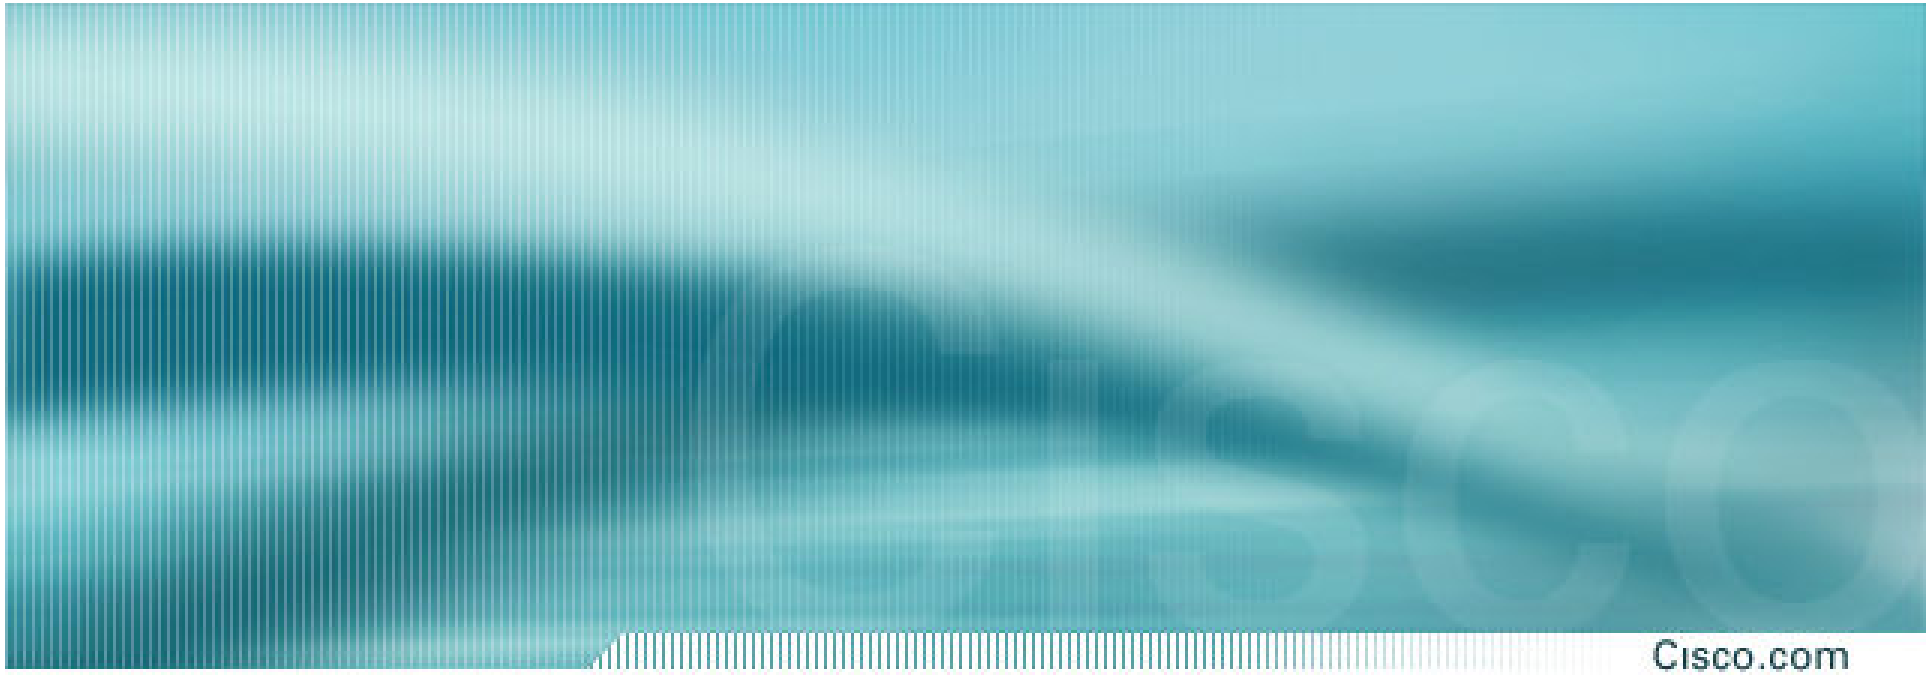

Network Component Evolution

© 2003, Cisco Systems, Inc. All rights reserved.

Building Network Components- some Decisions…

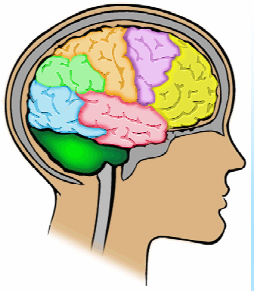

Controlplane

Congestion Control

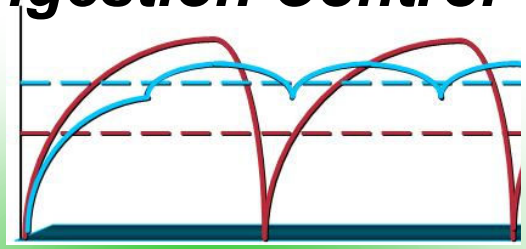

QOS

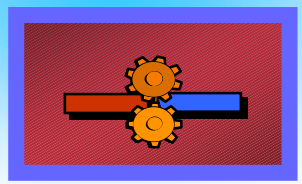

Dataplane

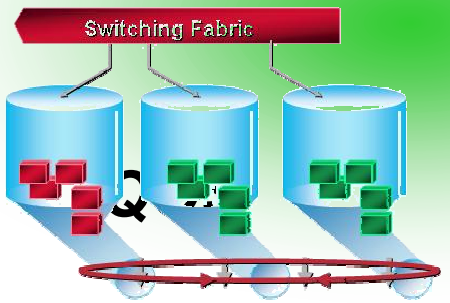

Buffers

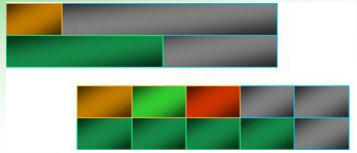

Architecture

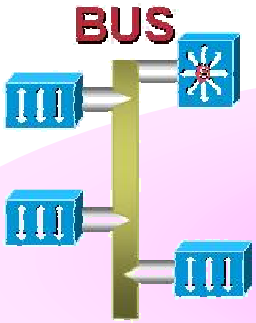

Fabric

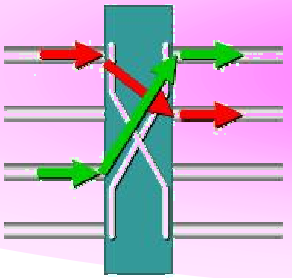

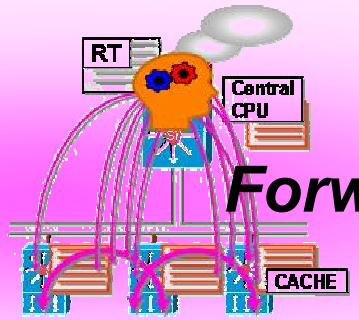

Forwarding

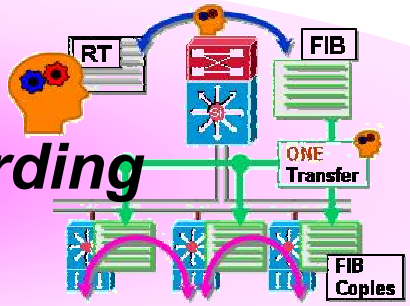

Flow

Topology

© 2003, Cisco Systems, Inc. All rights reserved.

It is all about Brains…
Controlplane - Dataplane

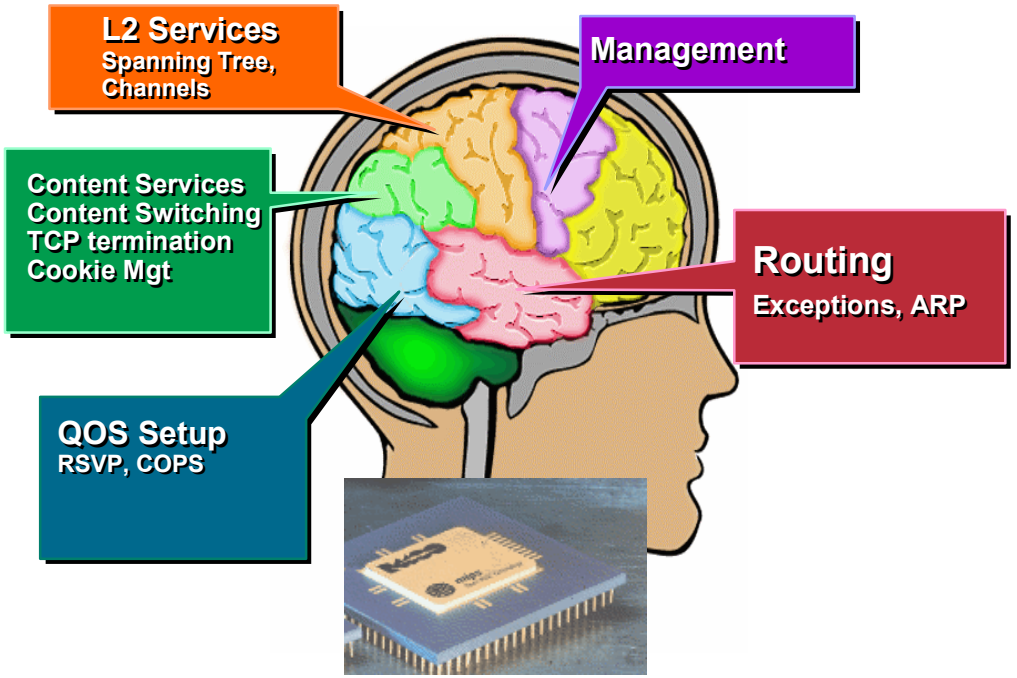

Control-Plane „SW“

- RISC/CISC CPU based
- Routing
- Spanningtree
- Housekeeping
- Management

Dataplane „HW“

- Switching in ASICS or NW Proc
- forwards & manipulates Data
- handles Access Control, QOS

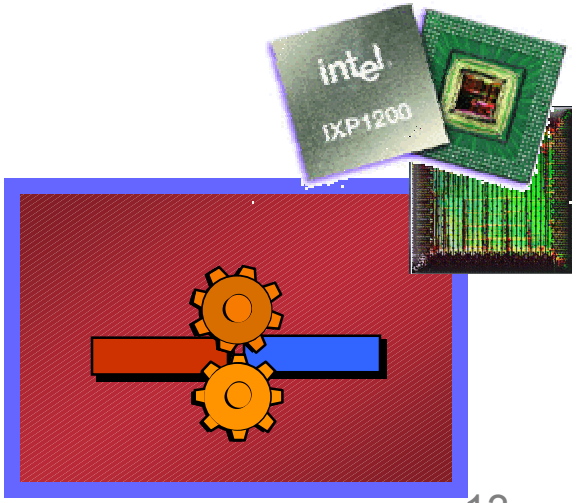

© 2003, Cisco Systems, Inc. All rights reserved.

True Scaleability ?

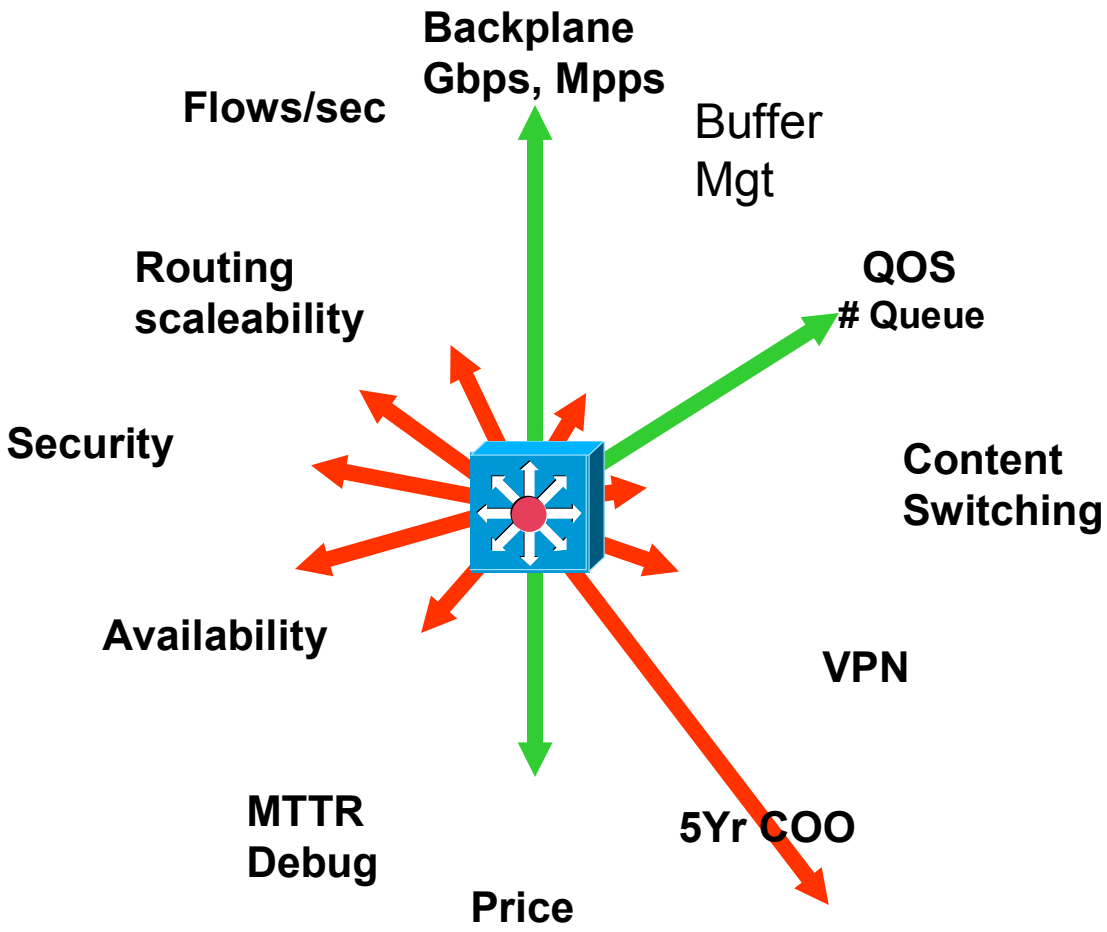

© 2003, Cisco Systems, Inc. All rights reserved.

True Scaleability !

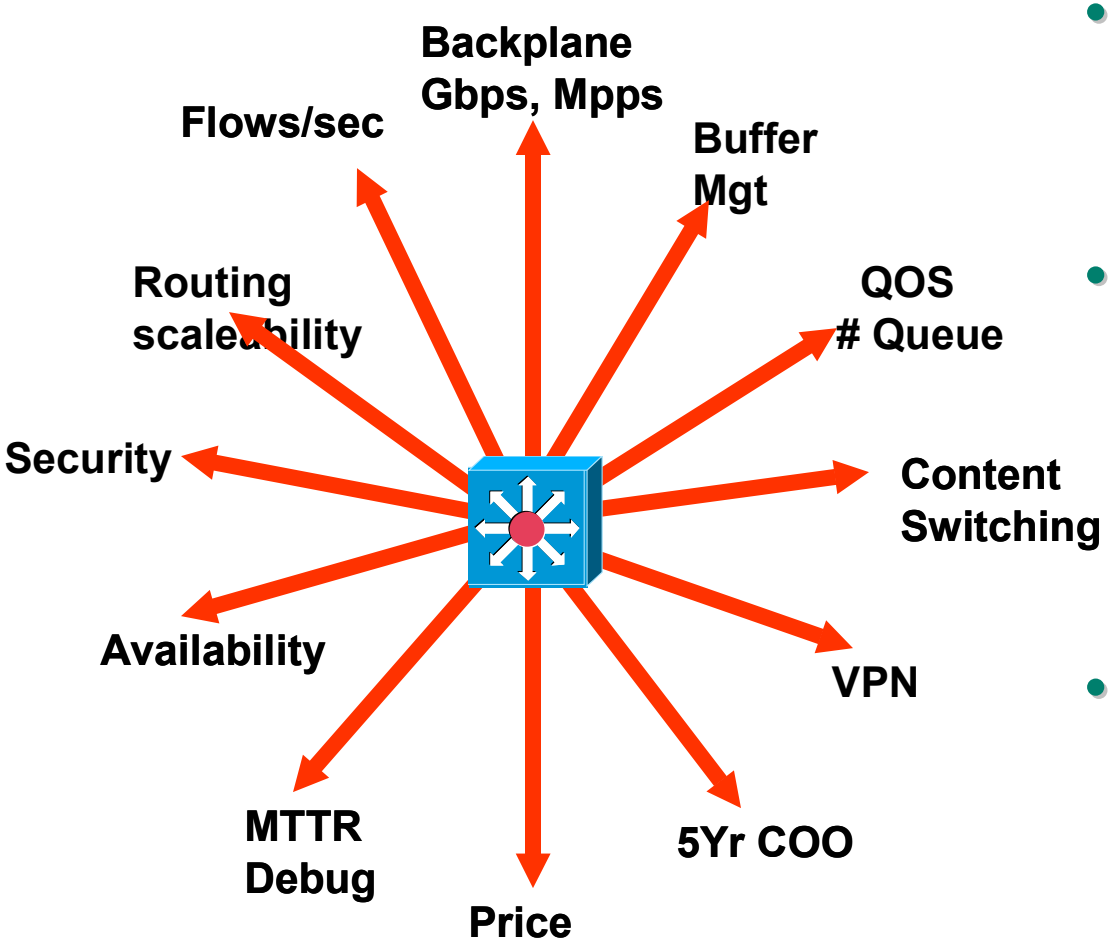

- Scaleability deals with all networking layers
- It is important to understand the different components of End to End Performance
- Focus on just one of the many variables does not help in real life situations

© 2003, Cisco Systems, Inc. All rights reserved.

Network Processor

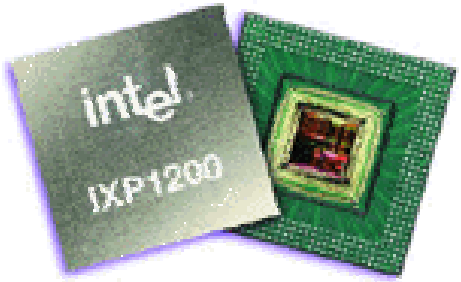

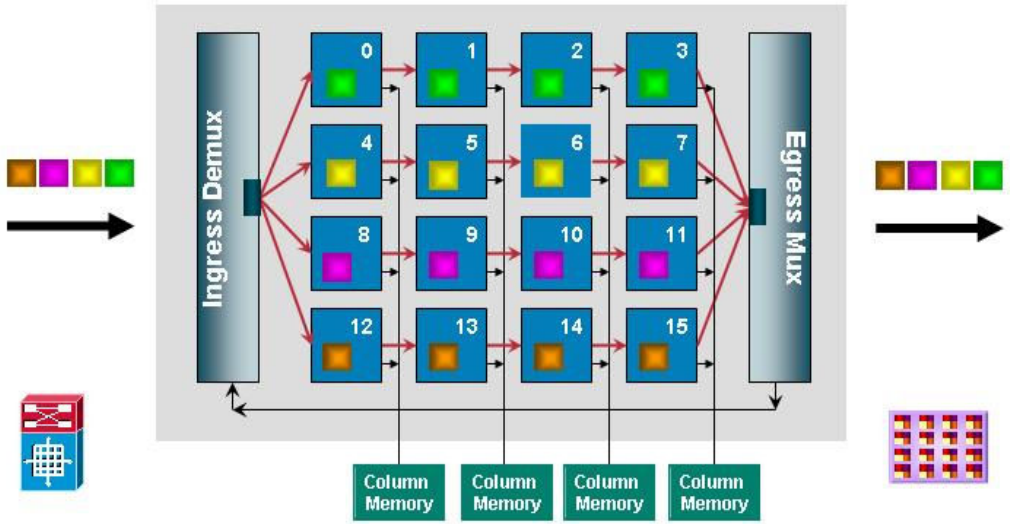

Shorter Development Cycles
Changes possible
More difficult to predict performance

© 2003, Cisco Systems, Inc. All rights reserved.

ASIC Technology

Catalyst FFE
Technology = 0.18u
Die size = 13.4 x 13.4
Transistors = 91M
Pin Count: 1188

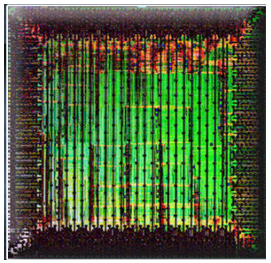

	Feature Size/ # Metal Layers	# Transistors (Million)	Chip Size (mm2)
intel. Pentium4	0.18µ/6M	42	217
Cisco Systems FFE	0.18µ/6M	91	180

Long Development Cycles
Difficult to change
Very high Performance
Costefficient in high numbers

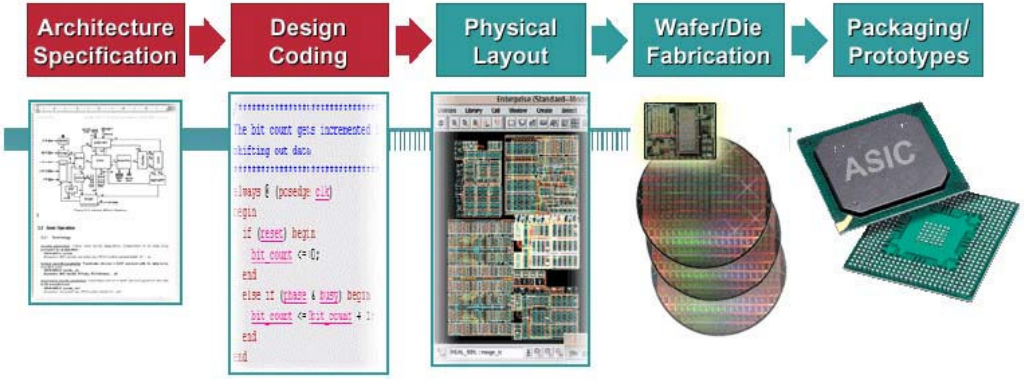

© 2003, Cisco Systems, Inc. All rights reserved.

Efficient Lookup – Ternary CAM

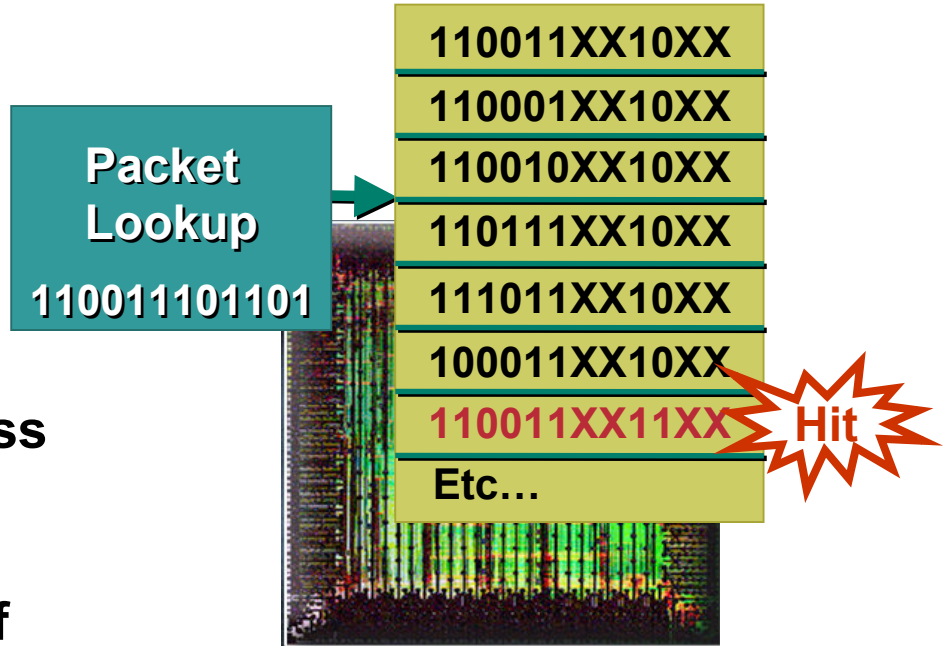

- Used for Routing Decisions, Access Control Lists, QOS, L2 Tables
- All entries are checked in parallel Same performance independent of number of entries
- Advantages
  - Longest match lookup
  - One lookup—Fixed latency
  - flexibility to ignore fields

© 2003, Cisco Systems, Inc. All rights reserved.

TCAM HW CEF Implementation

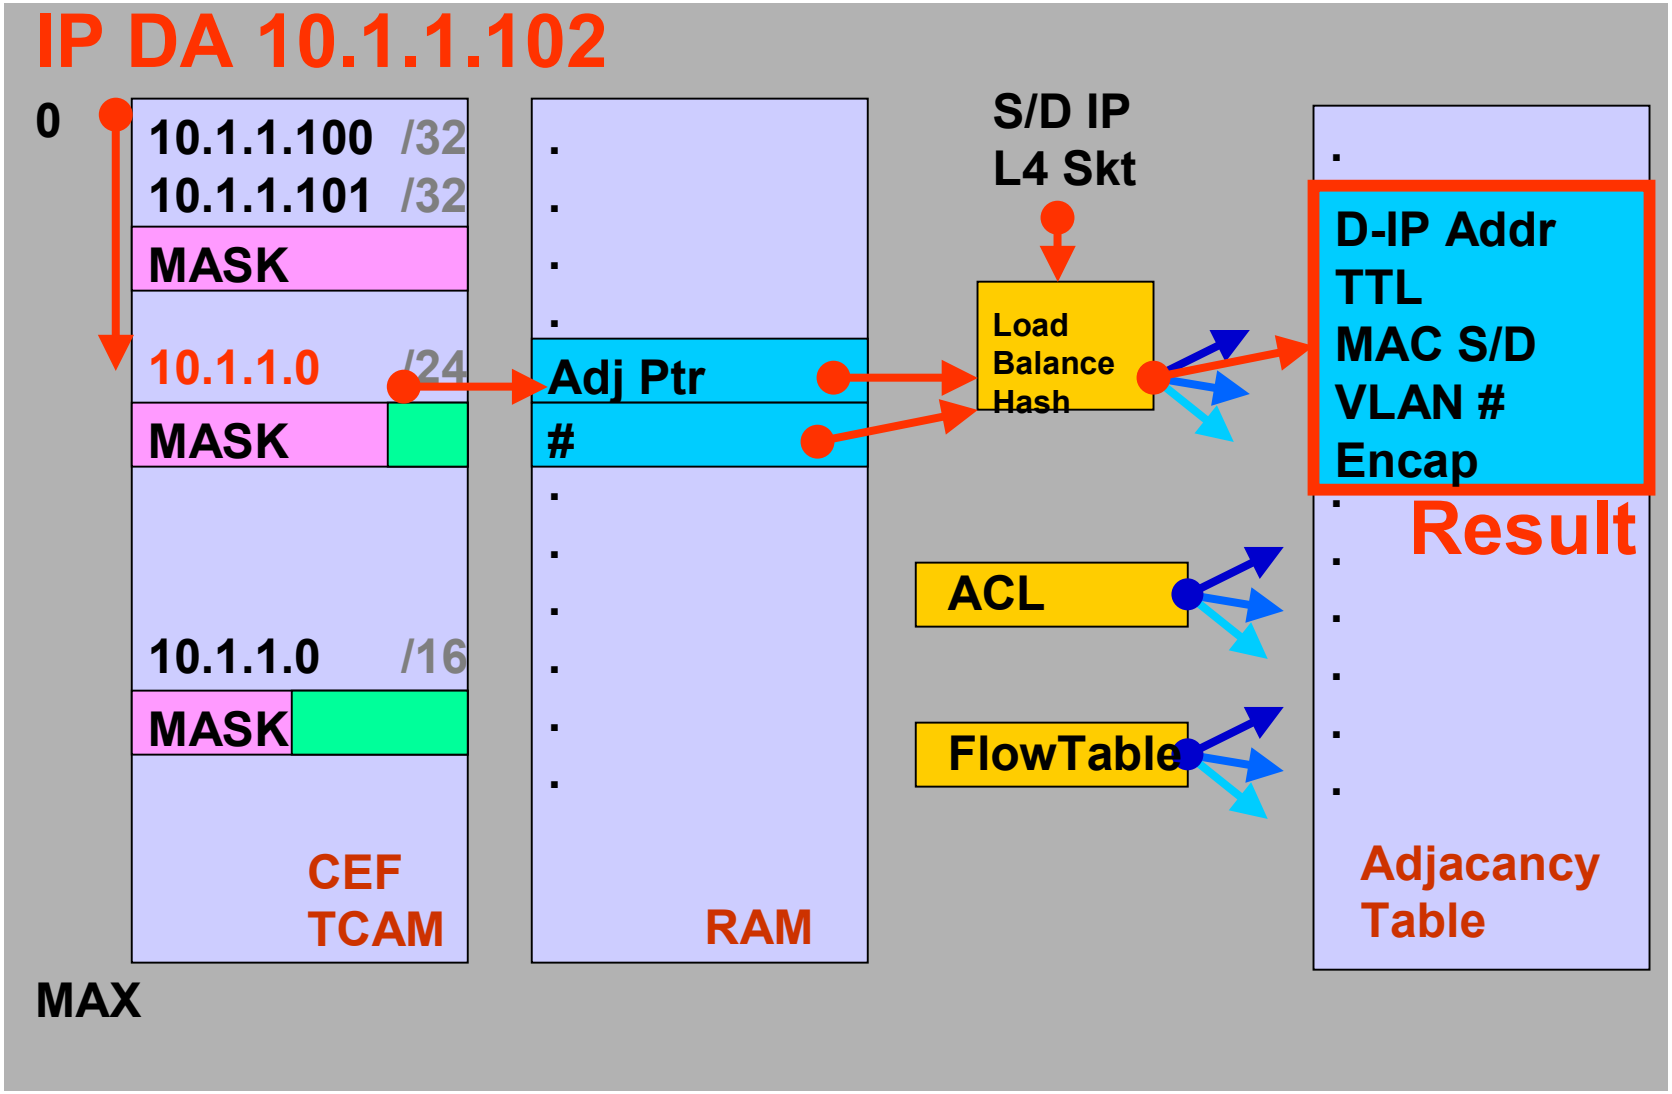

© 2003, Cisco Systems, Inc. All rights reserved.

SCALEABLE Forwarding
why CPU Flow based mechanisms do not scale

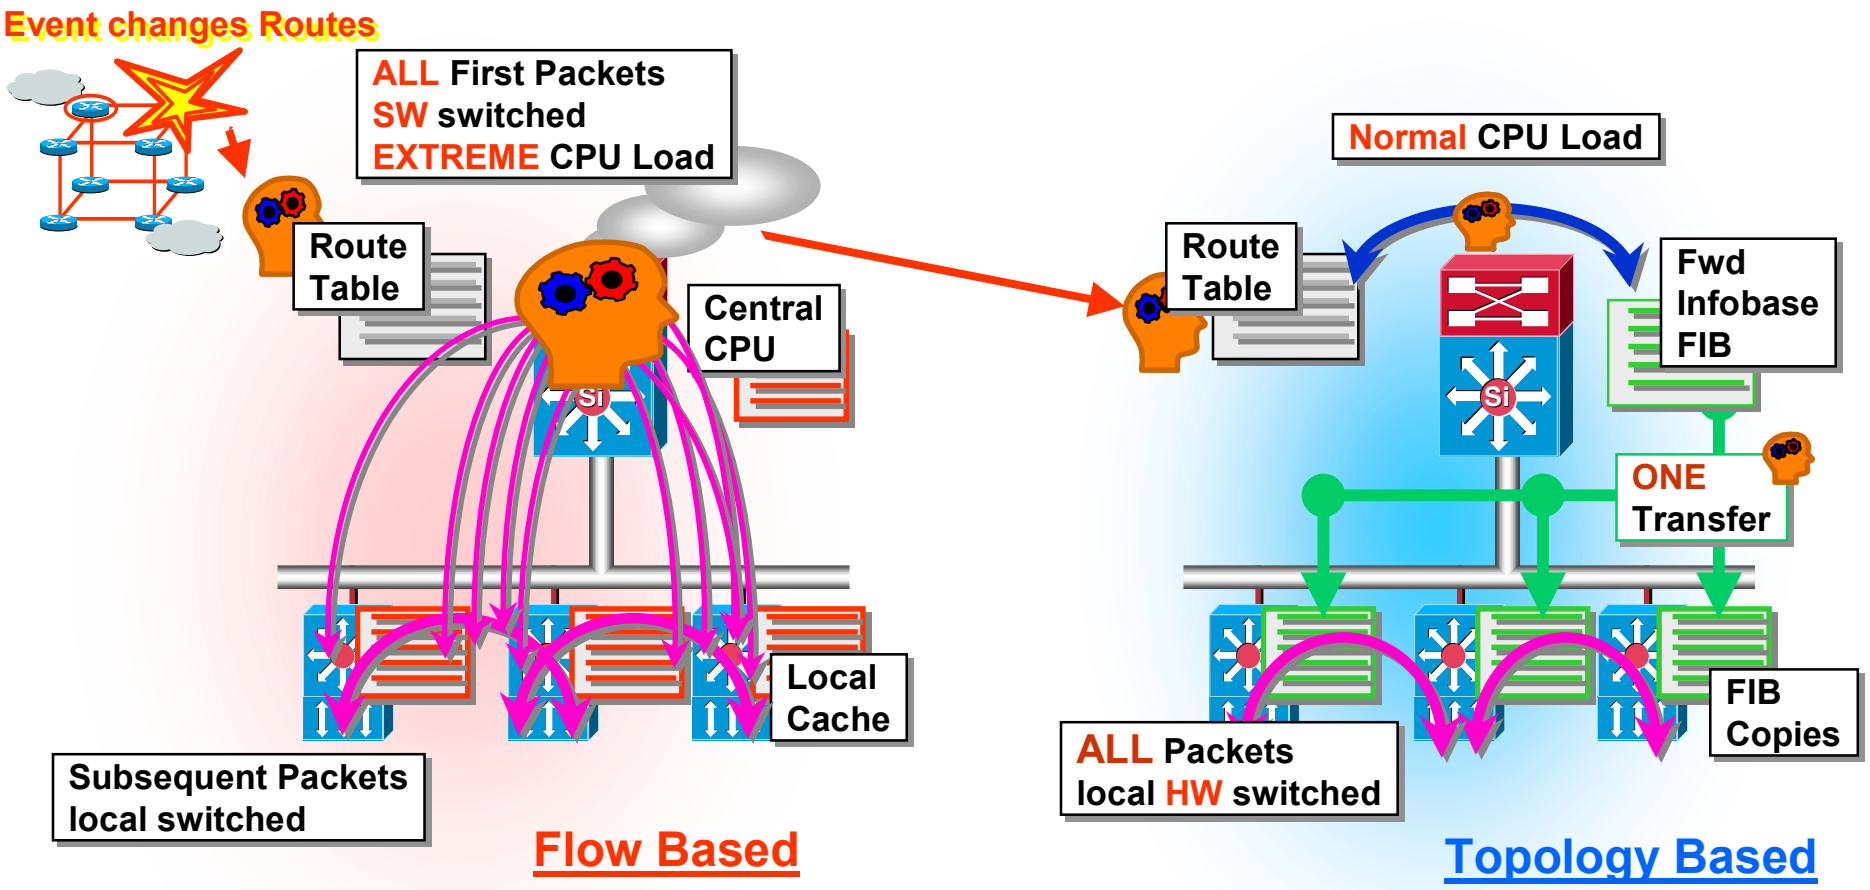

Distributed FLOW based Forwarding will never scale to Real World Traffic Patterns
I.e. 100 Mio PPS would require 2 MPPS CPU performance for Flows with average 50 Pkts
MOST Flows are much shorter - ALL CPUs are slower today...

© 2003, Cisco Systems, Inc. All rights reserved.

The Weak Link in the Chain
Flow Based Models

- **The "First" packet in a flow must be presented to the CPU where a forwarding decision is made & subsequent packets are handled in hardware.**
- **CPU performance is small compared to ASIC system.**
- **Problems dealing with Network Reconfigurations due to building Route Tables, all new flows & cache purge all at the same time**
- **Many new flows may overwhelm processor, causing packet loss of data as well as control plane traffic (such as routing protocols)**
- **To ensure data integrity, a certain amount of maintenance "cache churn" is expected to keep the state of the flow cache current. Reflected as CPU utilization.**

© 2003, Cisco Systems, Inc. All rights reserved.

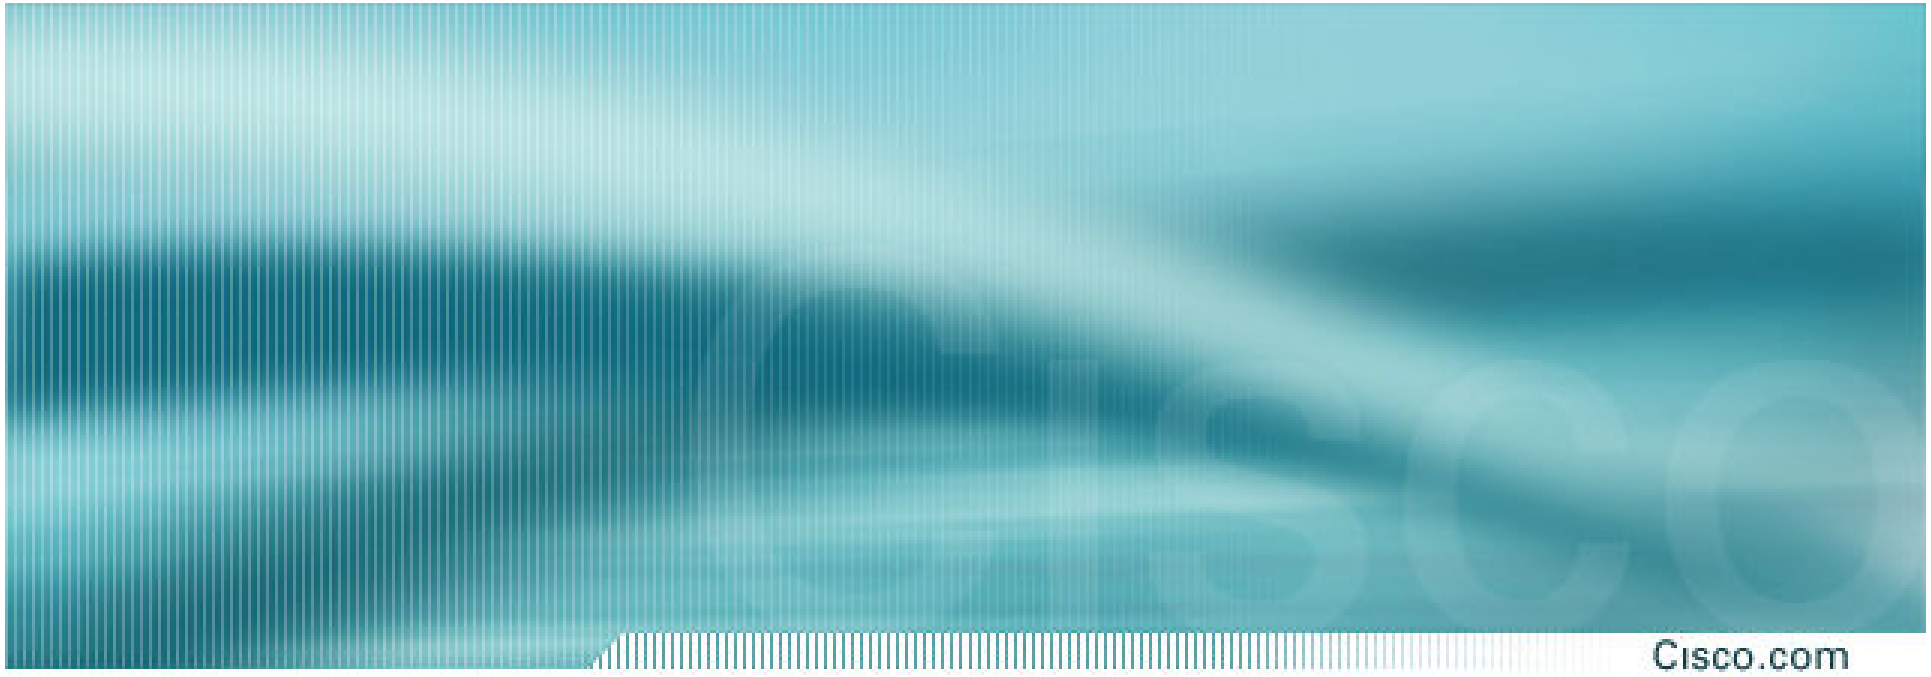

Technology Evolution

© 2003, Cisco Systems, Inc. All rights reserved.

Ethernet Is Extending Reach

- Price/Performance
- Consistency
- Services ubiquity
- Perceived Simplicity
- BUT
- Ethernet is just a Frame Format…
- SP-Class Ethernet in WAN/MAN is different than simple Campus
- Fundamental Issues like RTT size Buffers, WAN Class QOS still apply

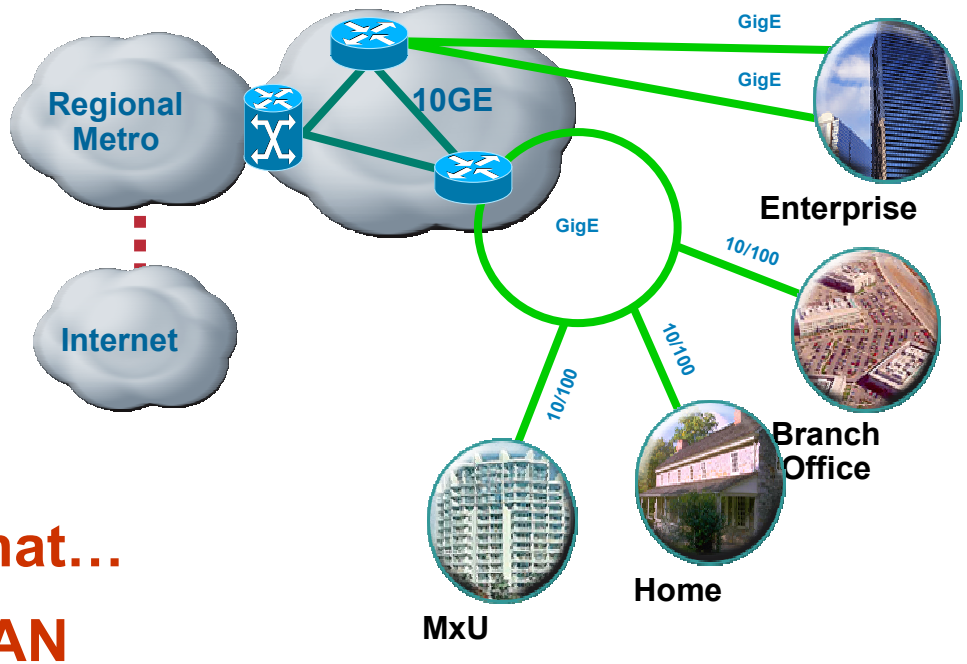

© 2003, Cisco Systems, Inc. All rights reserved.

What is Quality of Service?

The Pragmatic Answer:

QoS is managed fairness of resources

The Technical Answer:

Set of techniques to manage Delay, Jitter, Packet Loss, and Bandwidth for Flows in a Network

© 2003, Cisco Systems, Inc. All rights reserved.

The IP QoS Pendulum

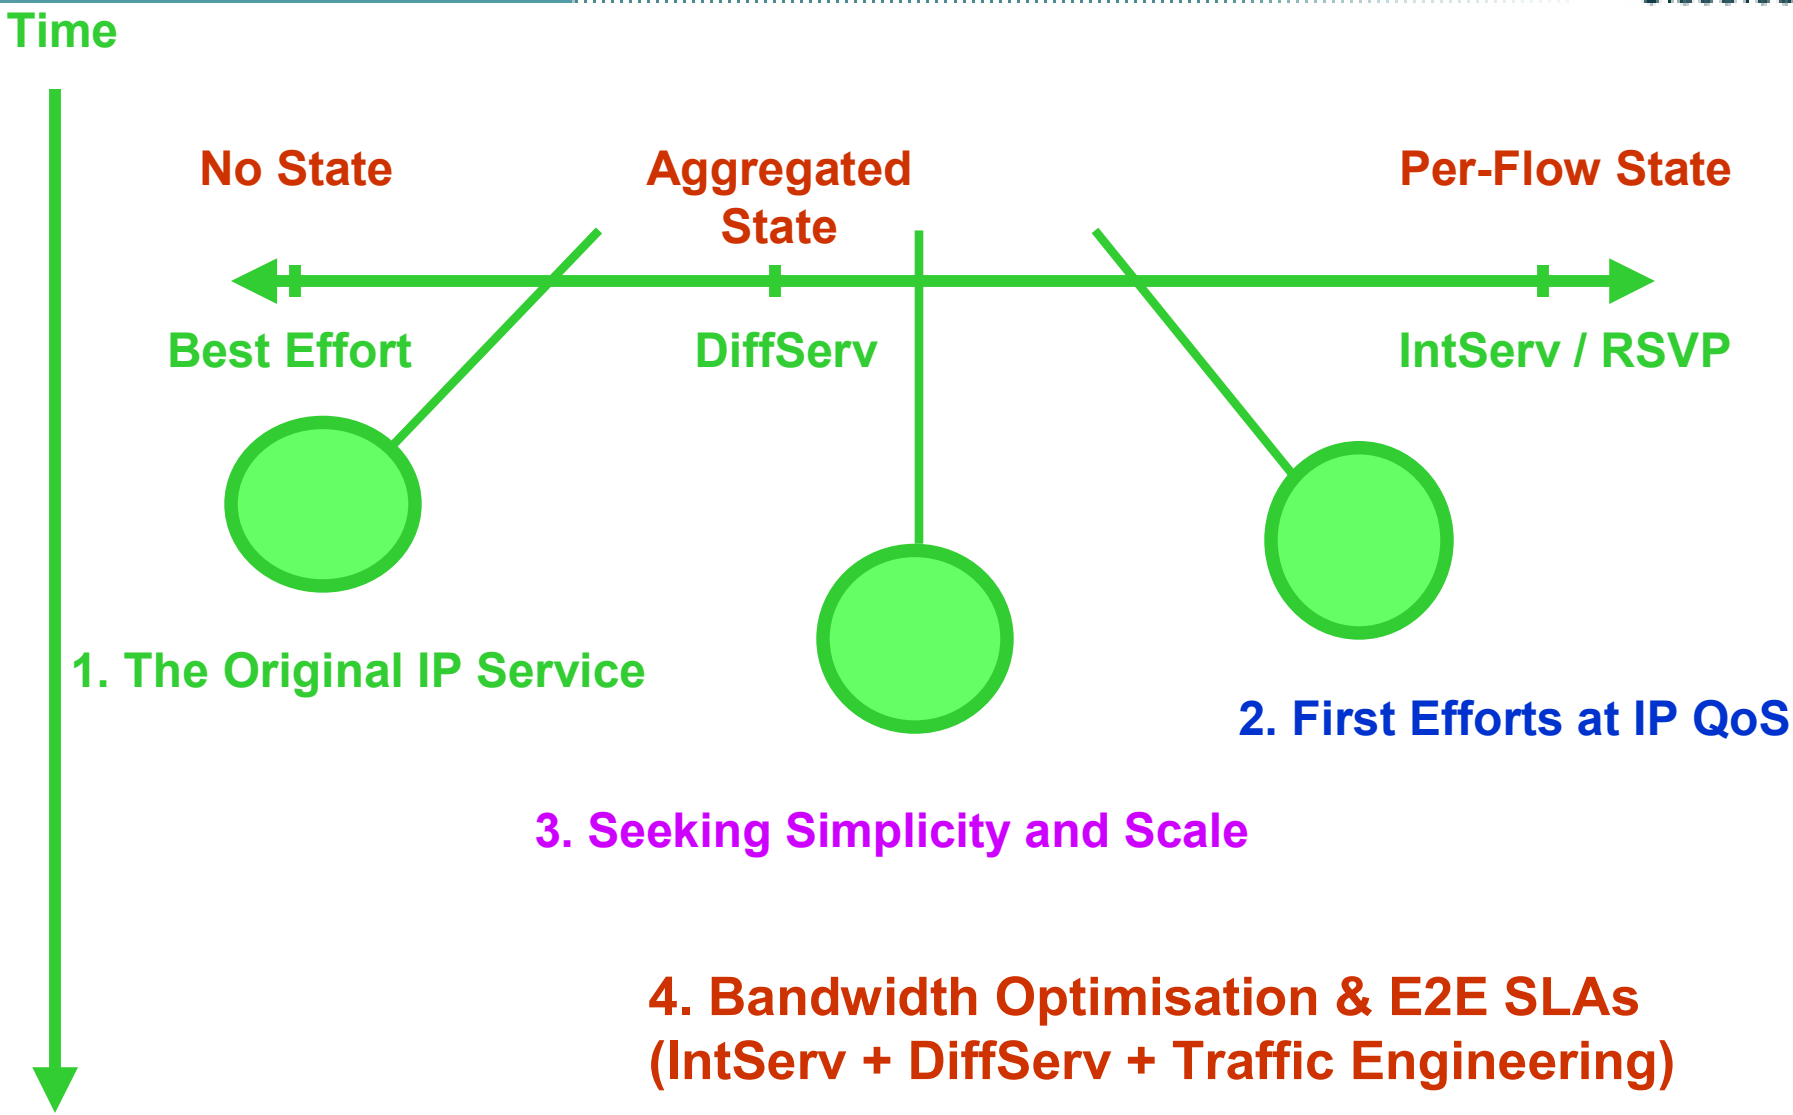

© 2003, Cisco Systems, Inc. All rights reserved.

Tight SLA – QOS

VOIP Class SLA

- Max latency 15 ms and no drops when rate <= policy bandwidth

Business Latency-Sensitive Class (EF) SLA

- latency below 30ms when no overload
- If other classes idle, take 100% of the bandwidth
- Maximum latency when this class is overloaded: < 30ms

Business Throughput-Sensitive Class (AF) SLA

- minimum bandwidth guarantee when link is congested

Default Class (Best Effort) SLA

- quantum and bandwidth at remaining percentage

© 2003, Cisco Systems, Inc. All rights reserved.

Tight SLA – Fast Convergence

- **Eliminate transient loops and black holes**
- **High availability requirements**
  - **99.999% per day ⇔ 0.9 sec of downtime**
- **VoIP requirements**
  - **40msec Loss of Connectivity : glitch**
  - **1-2 sec Loss of Connectivity : call drop**
- **Improvements**
  - **SPF (OSPF, IS-IS) Optimisation**
  - **IP Event Dampening**
  - **Multicast Sub-Second Convergence**
  - **BGP Optimisation**
  - **MPLS Fast Re-Route for Sub-100ms Restoration**

© 2003, Cisco Systems, Inc. All rights reserved.

In the Beginning:
Optical OXC promised elimination of Layers

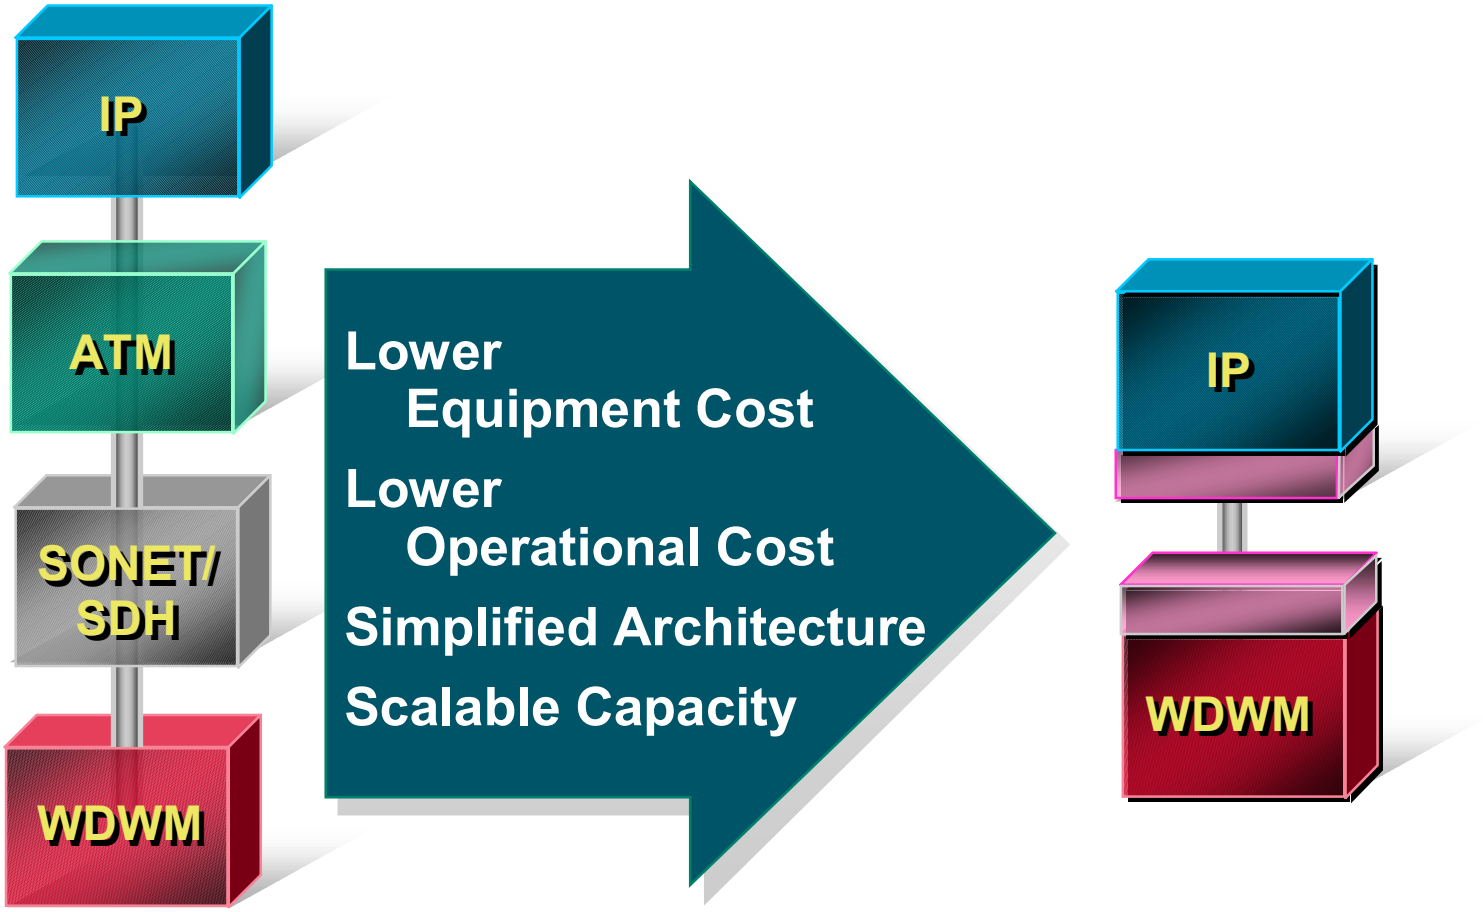

© 2003, Cisco Systems, Inc. All rights reserved.

The Promise of a Next Gen Optical Control plane:

- **Protocols to control "optical" transport networks**
  - Leverage new advancements in optical technology
  - Data optimized architectures
    - handle unpredictable data traffic
    - economical for low–revenue services
  - Future-proofed, open architectures
- **Protocols that better integrate IP and Optical**
  - Integrated control across data and transport
  - New differentiated service and business models
  - Simplified, automated operational processes & systems

© 2003, Cisco Systems, Inc. All rights reserved.

Extending MPLS Protocols for the Optical/Unified Control Plane

Forwarding Plane Extends MPLS Labels

- DWDM
- Fiber
- TDM

Control Plane Extends MPLS-TE

- OSPF-TE & IS-IS routing
- RSVP-TE for signaling
- Adds in-band & out-of-band control channel support
- LMP- Link Management Protocol

Mgmt & Control Address OTN specific needs

- Physical vs. Logical
- Transport reqs (e.g., Protection & Restoration, Explicit Interfaces, etc.)

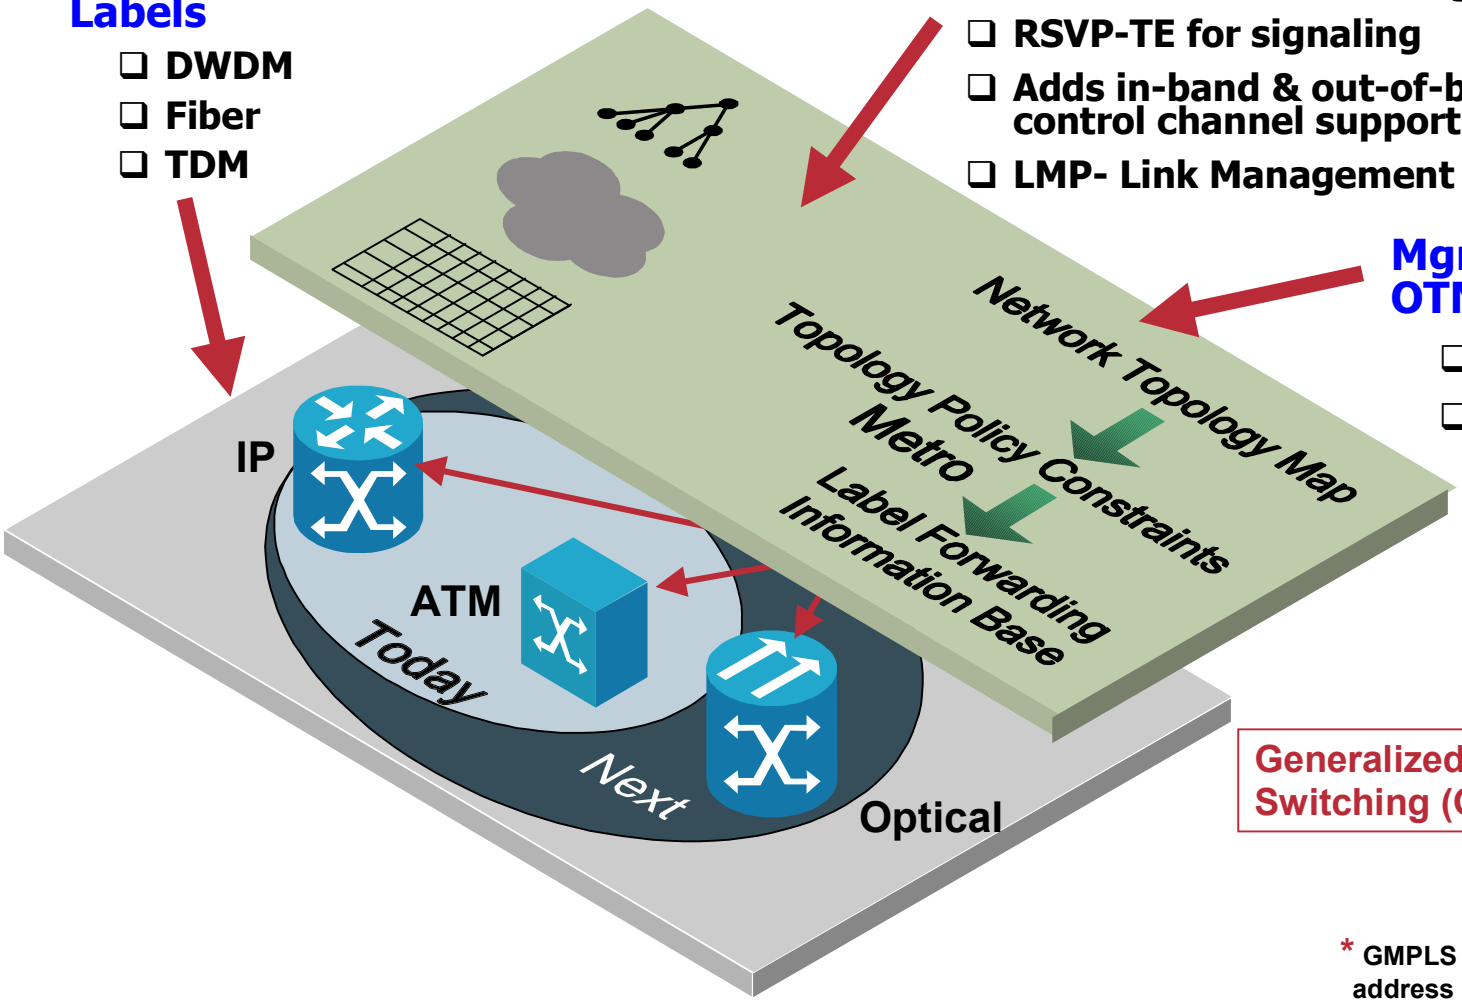

Generalized Multi-Protocol Label Switching (GMPLS)*

* GMPLS name is misleading as it address specific optical transport control issues.

© 2003, Cisco Systems, Inc. All rights reserved.

Ongoing Challenges to UCP Evolution

1. Excess Bandwidth in Network Core

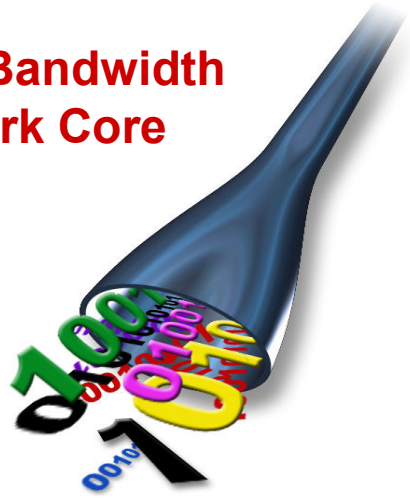

2. Telecom Bankruptcies &/or Business Restructuring = Disruption

3. Lack of Capital, Reduced R&D Budgets

4. Focus on near term Revenue; Leverage current assets

5. Concerns w/ Vendor Longevity

© 2003, Cisco Systems, Inc. All rights reserved.

So What Are the Characteristics of the Ideal Converged Next-gen Network?

Fusing the Best Properties of Today's Networks onto a Common Lowest Cost Infrastructure

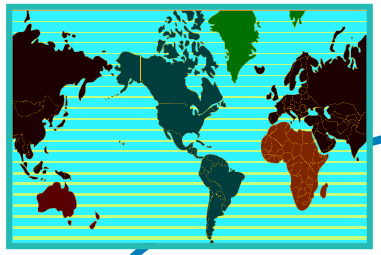

Ubiquity/Reliability of the PSTN

Mobility of the GSM Network

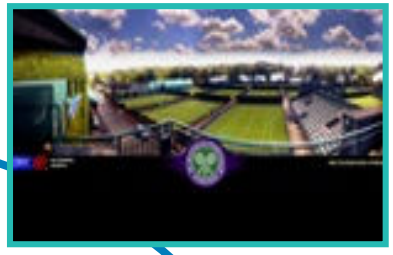

Content Richness of Cable/Television

"Next-gen Network"

Bandwidth of an Optical Network

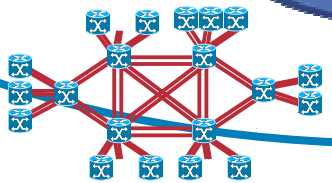

Latency Control of an ATM Network

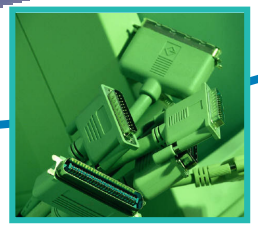

Operational Ease of Ethernet

Flexibility of the Internet

© 2003, Cisco Systems, Inc. All rights reserved.

Q & A

© 2003, Cisco Systems, Inc. All rights reserved.
© 2002, Cisco Systems, Inc. All rights reserved.