Application Intelligence in the Network

Issy Ben-Shaul
CTO, Application Delivery
Cisco Systems, Inc.

The Globally Integrated Enterprise

- Hire globally based on skills, cost, and business
- Operations horizontally and globally integrated
- Strategy and management shaped globally

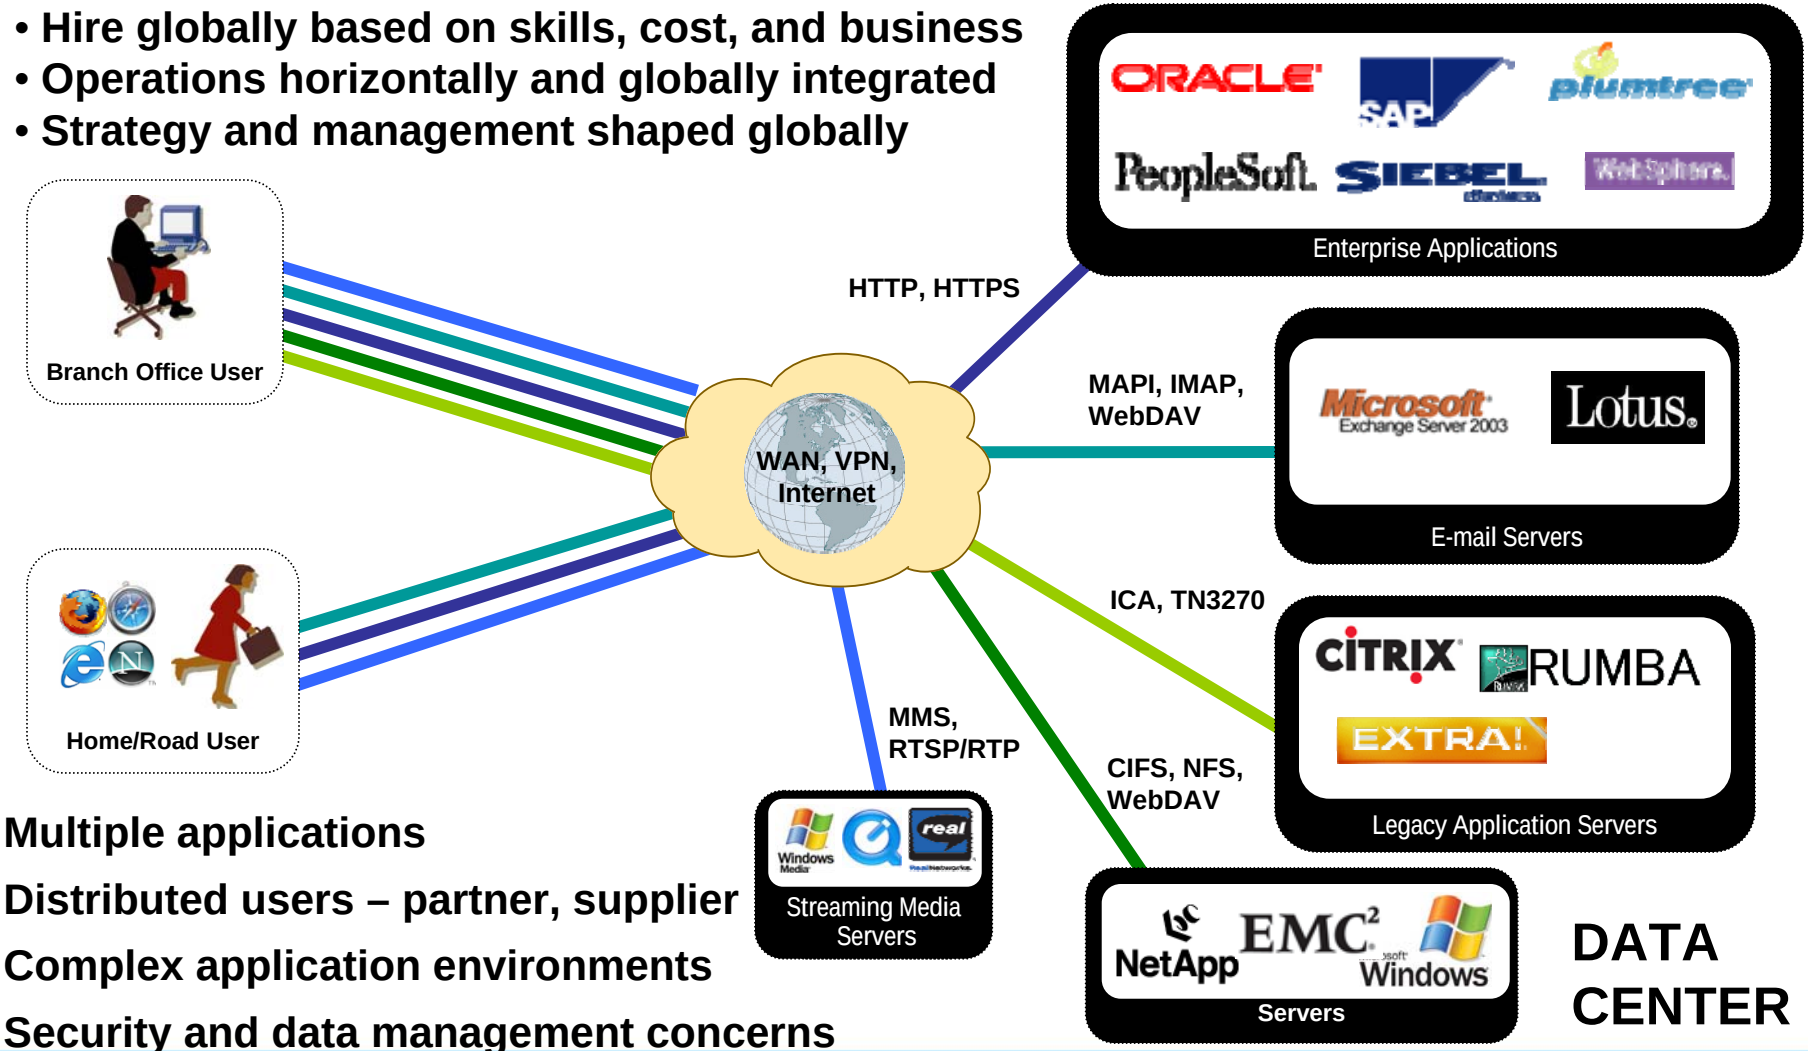

- Multiple applications
- Distributed users – partner, supplier
- Complex application environments
- Security and data management concerns

What is App Intelligence ?

- Passive
  - Classify apps (payload inspection)
- Active
  - Control (police) admission of flows/apps
  - Prioritize resource allocation per app/flow
  - Optimize apps
- Adaptive
  - Learn the characteristics of flows/apps
  - Adjust active policies based on identified characteristics
- **PERFORM ALL OF THE ABOVE at L7/8**

Why App Intelligence: Passive Classification

- App-level visibility
  - monitoring & troubleshooting
  - Traffic engineering: SLAs, bandwidth allocation & latency prioritization
  - App-based optimization
  - Comment: tunneling hinders visibility!
- L7/8 (payload inspection) essential
  - Multiple apps on same L7 protocol/port
  - Dynamic port apps
  - Same app may run on different ports

App Intelligence: L7/8 WAN Optimization

- Business Drivers
  - Branch and data-center consolidation
  - Application Acceleration
  - Bandwidth Savings
- L3/4 optimization is not sufficient due to app latency

The Impact of Latency

- L1: Network latency
- L4: Transport Latency
- L7: Application latency

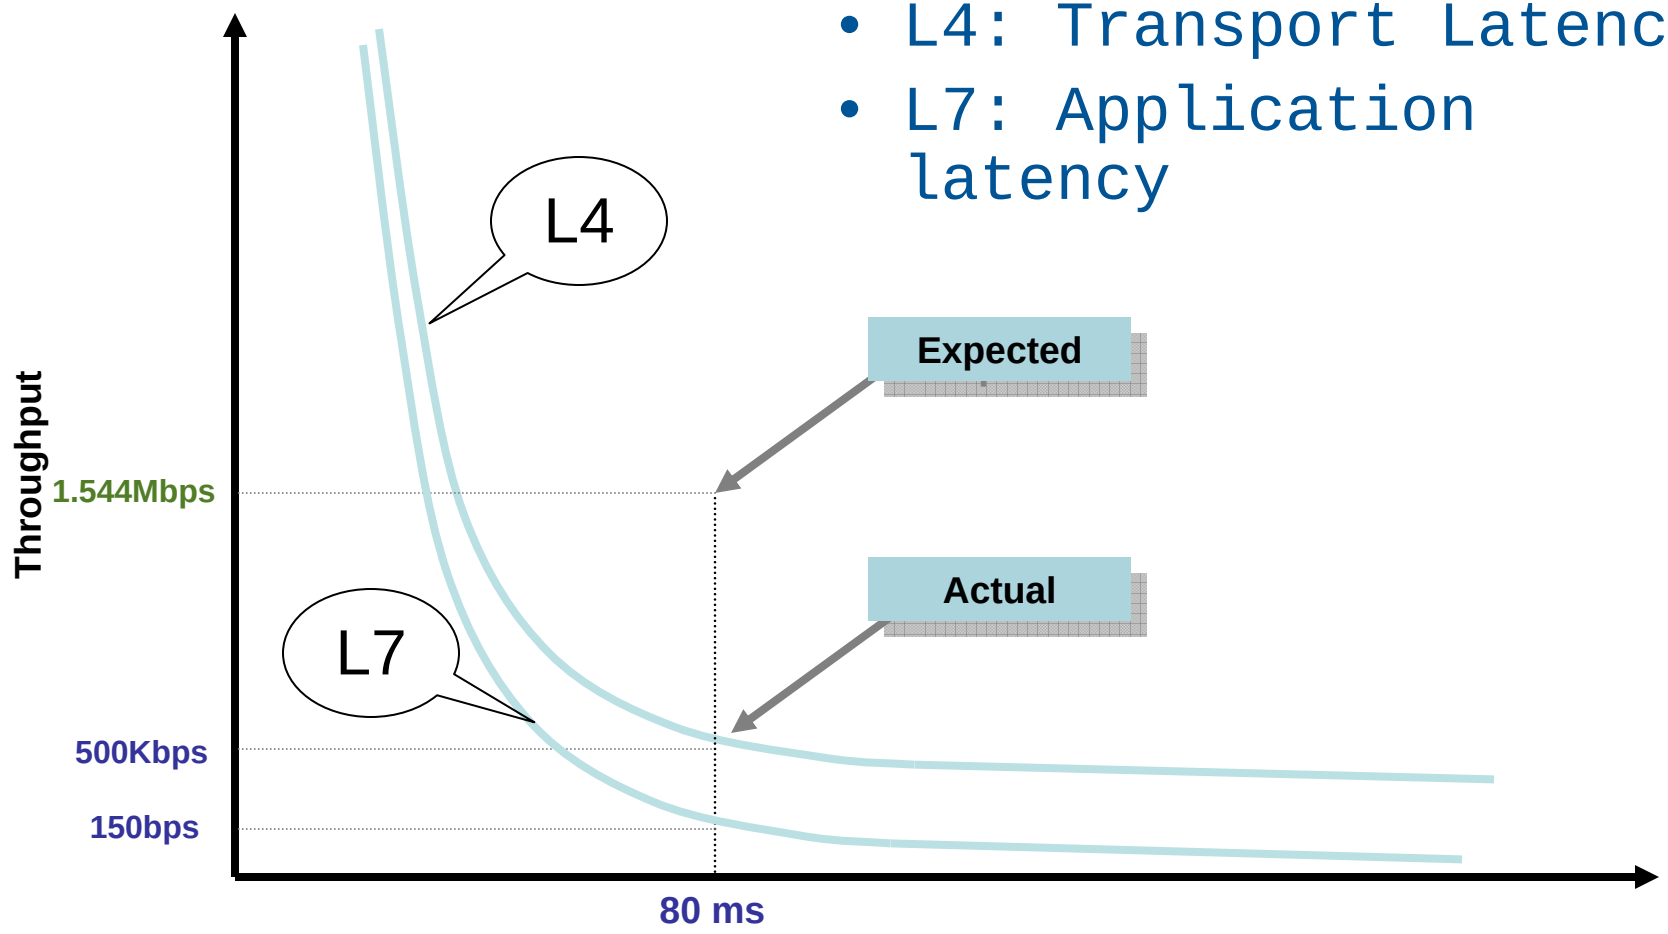

L4 Optimization: Data Redundancy Elimination

- Divide stream to small chunks
- Store in local DB chunk signature and destination peer
- If first time, store and send to peer
- If seen before, send signature only

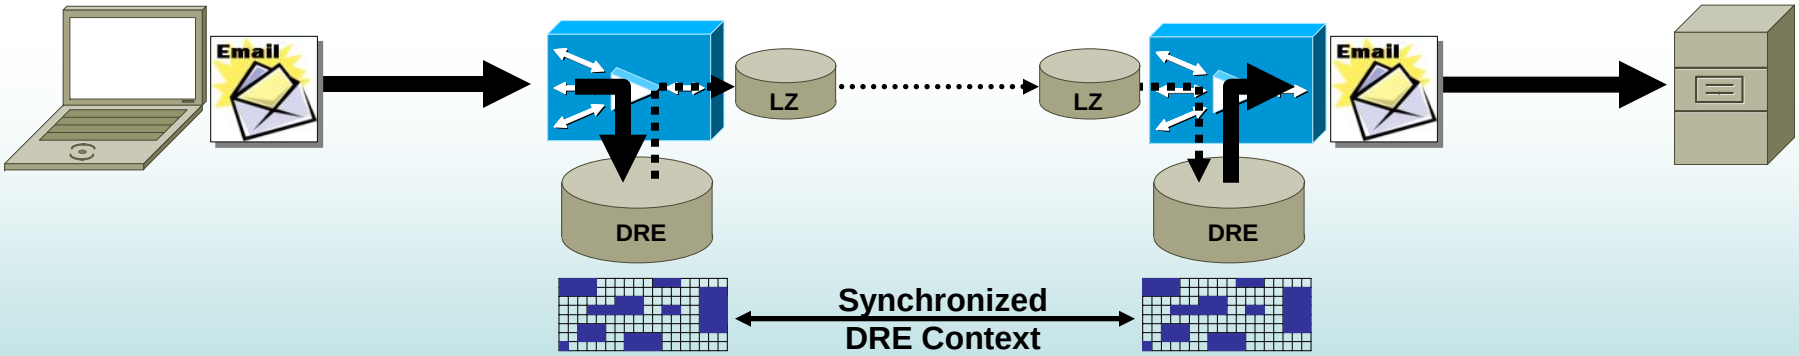

L4: TCP Flow Optimization

- Fill the pipe: Improved TCP window management
  - pipelining (local acknowledgement)
  - Window scaling
- Improved congestion control:
  - Smaller drop, faster climb, fairness
  - Less sensitive to packet loss
- Prioritize processing of latency-sensitive apps
- Flexible policy configuration
- Pacing based on DRE

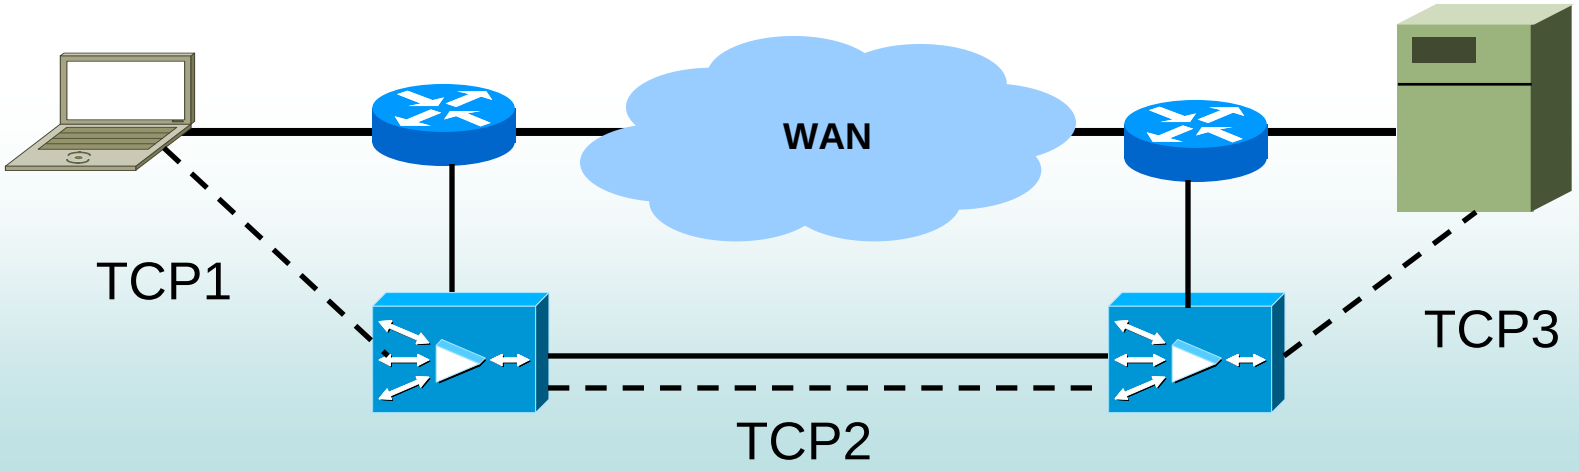

Example: CIFS Latency

- Open a 2MB Word File: ~1000 sequential messages
- L4 optimization does not address app latency

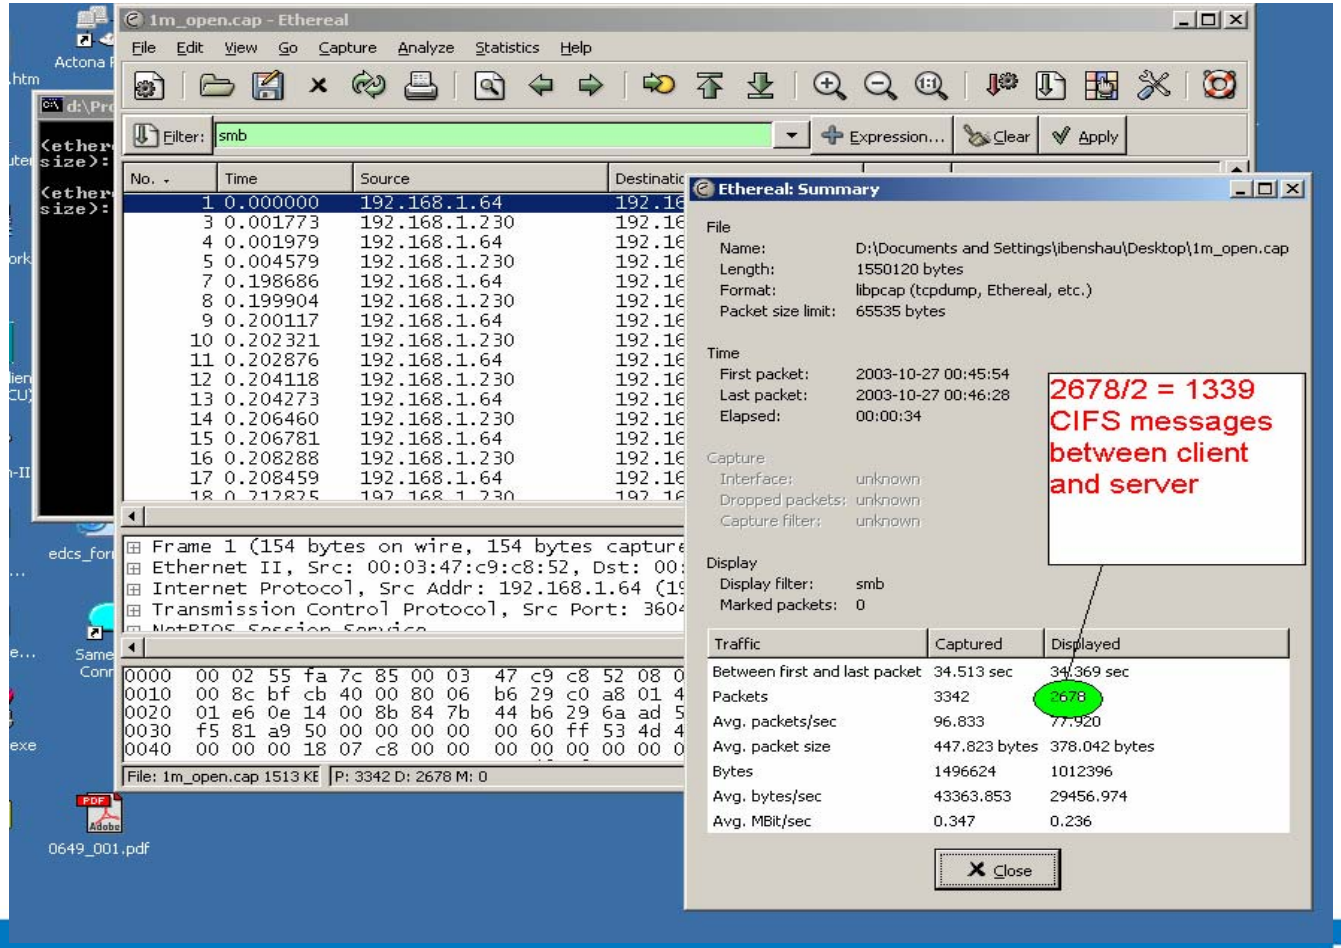

L7: CIFS Optimizations

- Understand protocol to perform:
  - Safe object and meta-data caching
  - Read ahead
  - Write behind (flush on close)
  - Local handling of operations
  - Scalable: 10x offload of core devices
- Bottom line: > 90% WAN Latency reduction

Optimized CIFS:
only small fraction with WAN delay

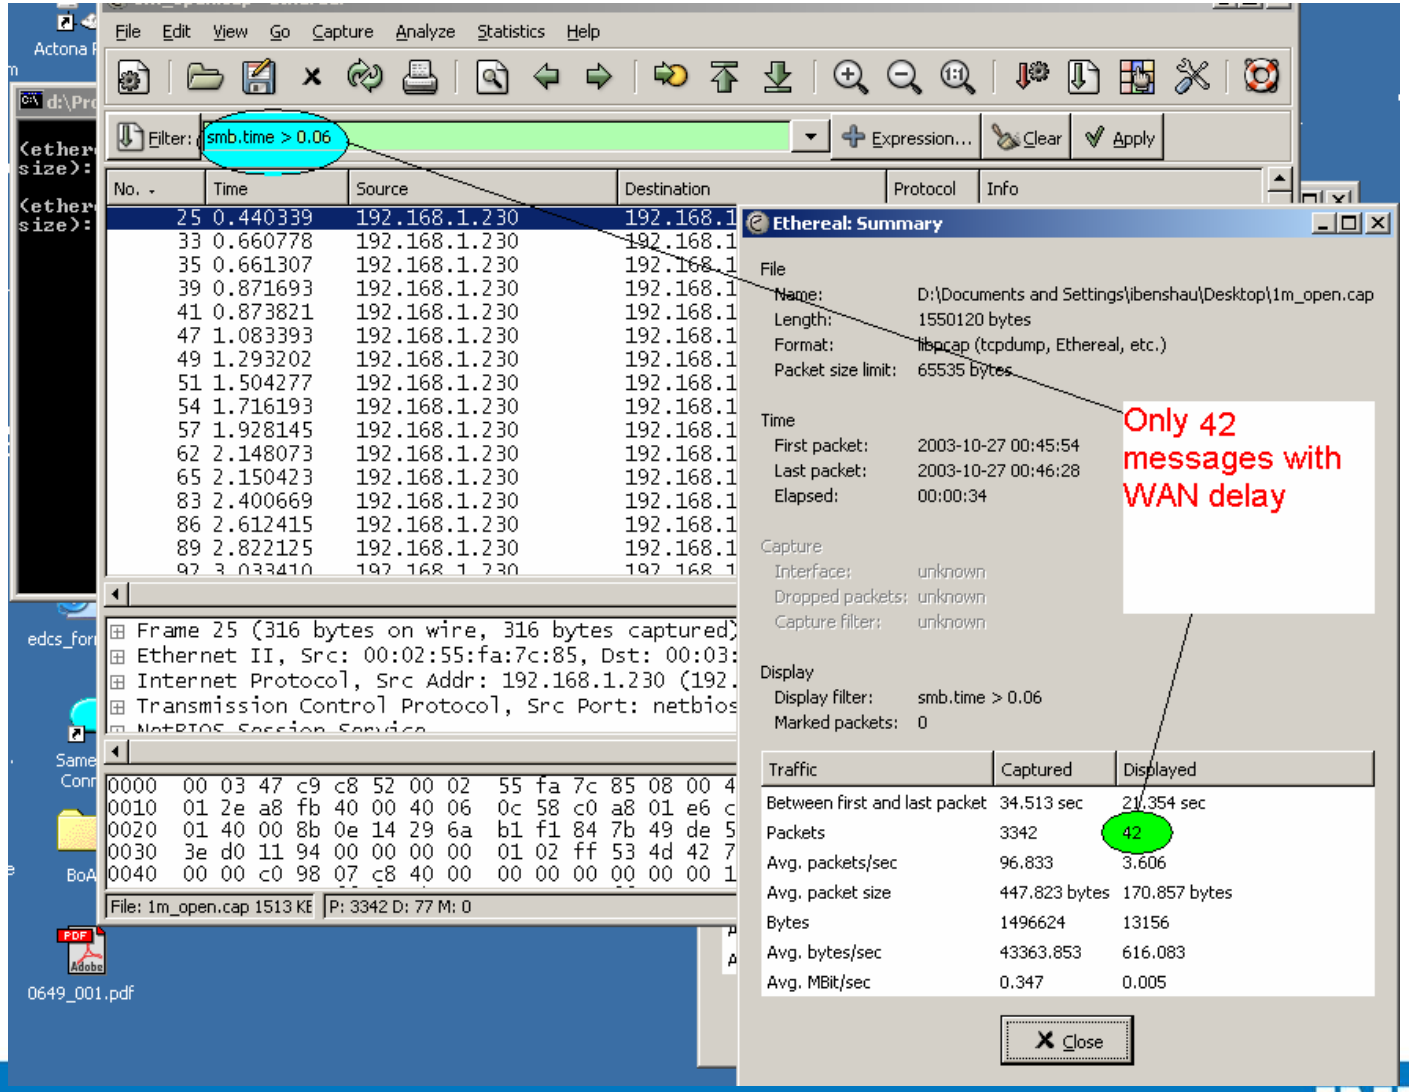

Example: L7/8 HTTP Optimizations

- Pre-fetch embedded, related URLs
- Reuse Connections
- Cache objects based on app-knowledge

Towards Adaptive Network Management

- Behavioral Classification: discover application characteristics (chattiness, compressibility)
- Prioritize/optimize based on application type and importance
- Dynamically allocate bandwidth per app based on compressibility

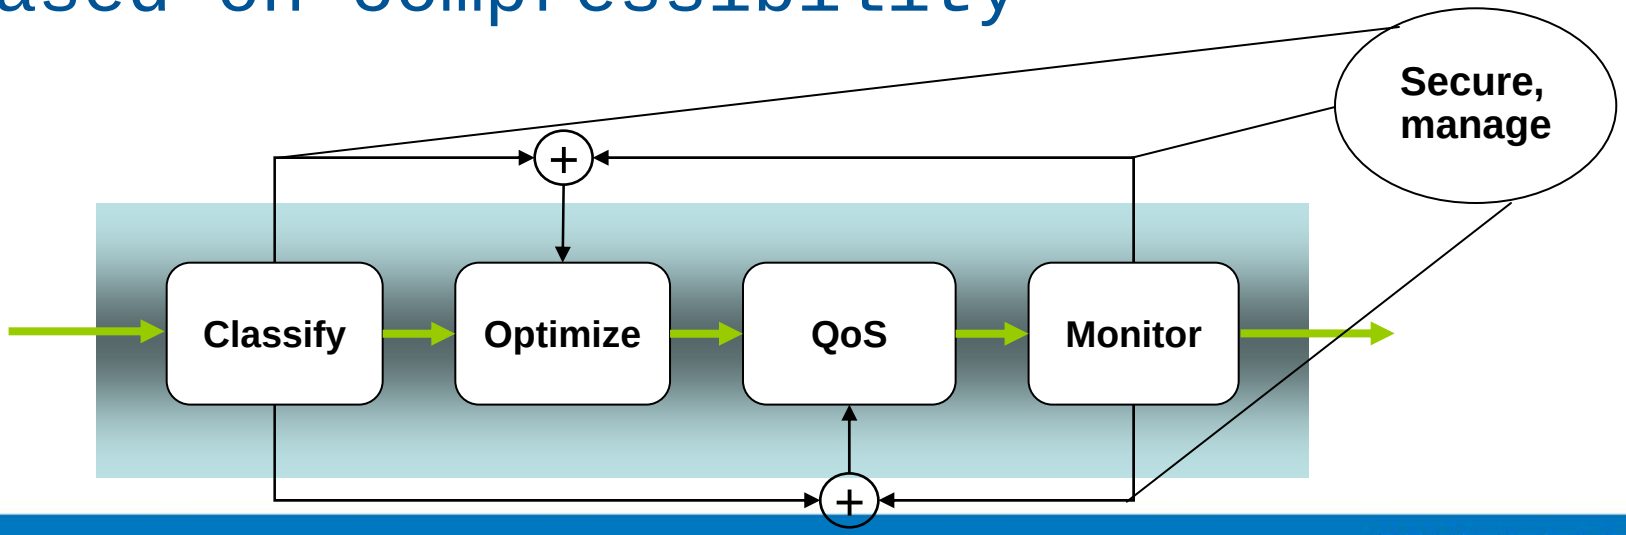

Summary

- Most enterprise apps are network-based
- The network is the only common platform for all apps
- App-intelligence is essential to
  - Monitor, troubleshoot, operate
  - Optimize, Accelerate
  - Control, Secure
  - Prioritize based on app importance
  - MATCH APP and CUSTOMER NEEDS with available network resources