

Architecture for a Hardware Based, TCP/IP Content Scanning System

David V. Schuehler
dvs1@arl.wustl.edu



<http://www.arl.wustl.edu/arl/projects/fpx/>



*Department of Computer Science and Engineering
Applied Research Laboratory*



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Outline

- Problem Statement and Motivation
- Overview of Content Scanning System
- Description of Architecture
- Initial Target Hardware Platform
- Conclusion



*Department of Computer Science and Engineering
Applied Research Laboratory*



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Problem Statement

- Scan TCP data flows at backbone data rates
 - Intrusion Detection Systems
 - Virus detection and elimination
 - Content based routing
 - Extensible networking solutions
- Support enhanced flow manipulation
 - Blocking
 - Unblocking
 - Termination
 - Modification



*Department of Computer Science and Engineering
Applied Research Laboratory*



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Motivation

- Internet has become both dangerous and costly
- Viruses spread to machines world wide
 - Consume computing resources
 - Reduce network throughput
- Email accounts flooded with spam
 - Consume disk space
 - Inconvenience users
- Current prevention techniques inadequate
 - Users fail to apply security patches
 - Anti-virus software doesn't prevent all attacks

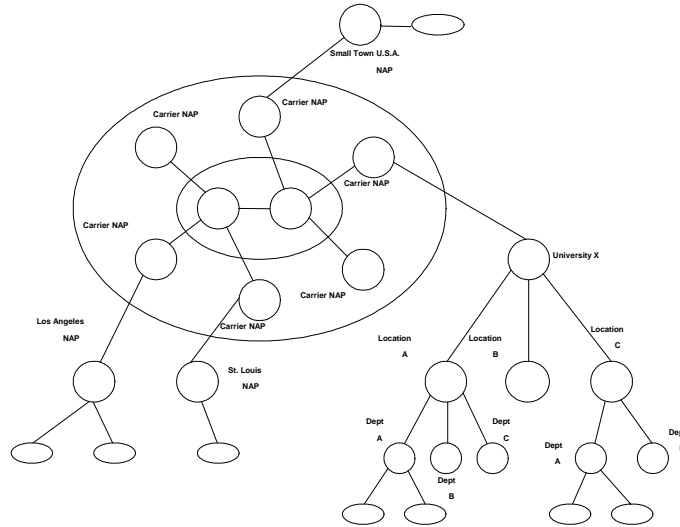


*Department of Computer Science and Engineering
Applied Research Laboratory*



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Widespread Virus Migration




Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Virus Containment



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Solution

A hardware based, TCP/IP content scanner

- Reconstruct TCP byte streams from data packets
- Scan TCP stream data for digital signatures
- Provide enhanced flow management features
 - Quarantine or eliminate viruses
- Operate at Internet backbone data rates
 - Multi gigabit/second data rates




Department of Computer Science and Engineering
Applied Research Laboratory




Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Types of Monitoring Systems

- Out-of-Band Monitor
 - Always a passive solution
 - Limited to detection systems



- In-Band Monitor
 - Able to block or alter transmitted data

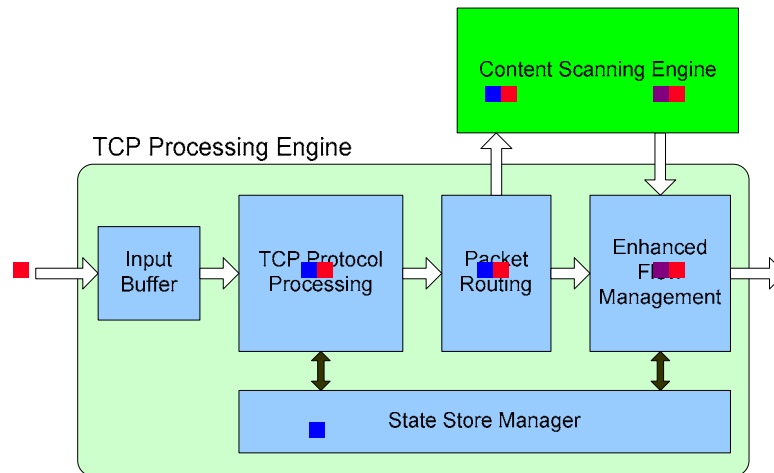


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Content Scanning System



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Environment

- High data rates
 - OC-48 (2.5Gb/s)
 - ~200 ns to process 64 byte packet
 - OC-192 (10Gb/s)
 - ~51 ns to process 64 byte packet
- Backbones support millions of active flows
- Large amounts of per-flow state information



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Competing Solutions

- Firewall
 - Filter based on header fields (address & port)
- Anti-virus software
 - Only protects one system
- Snort
 - Software based detection system
- IDS Systems
 - Traffic rates < 1Gbps
- TCP-Splitter
 - Limited to monitoring



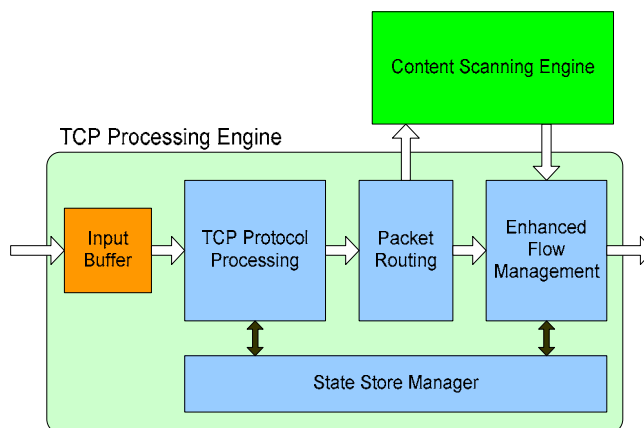
Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Input Buffer

- Buffers packets during downstream processing delays
- Ensures data loss occurs in whole packet increments

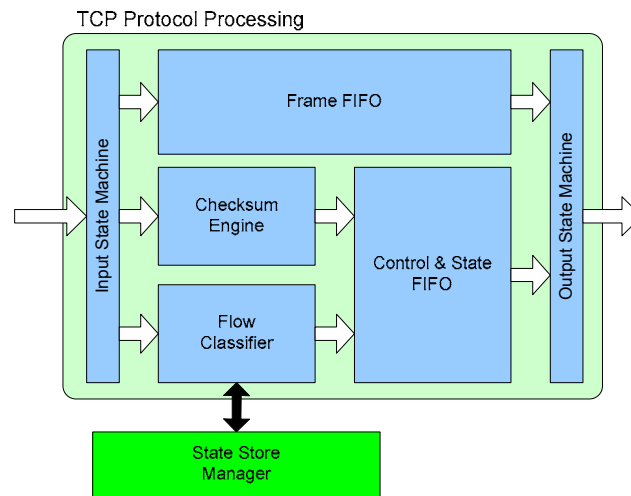


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

TCP Protocol Processing Engine



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Handling Improper Flow Termination

- TCP protocol supports idle flows
- Cannot differentiate between idle flow and abnormal termination
- How to handle idle flows?
 - Timer based termination
 - Least recently used age out policy
 - Random age out policy



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

State Store Manager Features

- Simple interface
- Supports multiple hashing algorithms
- 512 MByte SDRAM module
 - 64 bytes of state per flow
 - 32 bytes used by TCP Processing Engine
 - 32 bytes available for Content Scanner
- 8 million active flows supported

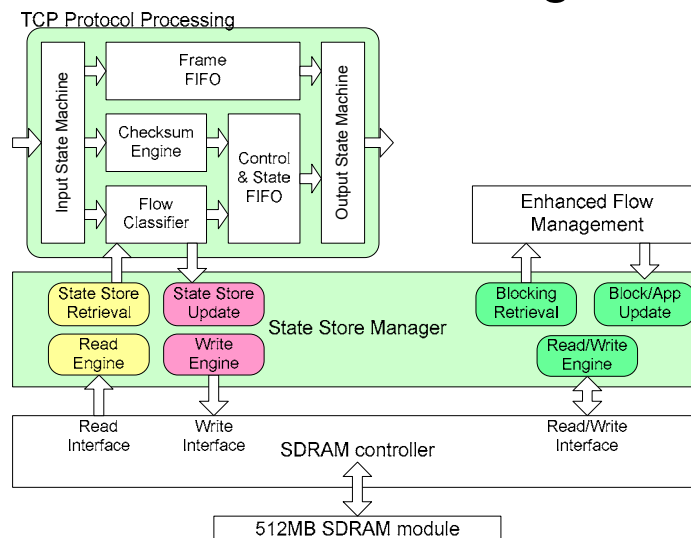


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

State Store Manager



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Per-Flow State Store Record

63	32	31	0
Flow Id	Hash Value		
Flags/Next Flow Id	Source IP Addr		
Sequence #	Dest IP Addr		
Blocking Sequence #	TCP Ports		
Content Scanner Context	Content Scanner Context		
Content Scanner Context	Content Scanner Context		
Content Scanner Context	Content Scanner Context		
Content Scanner Context	Content Scanner Context		



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Hash Implementation Tradeoffs

- Unlimited hash entry chaining
 - Pro: Best option for fully monitoring all flows
 - Con: Excessive time required to perform lookup
- No hash entry chaining
 - Pro: Easy to implement
Fast
 - Con: Potential for incomplete monitoring of flows
- Limited hash entry chaining
 - Pro: Bounded time to perform lookup
 - Con: Potential for incomplete monitoring of flows
Excessive time required to perform lookup

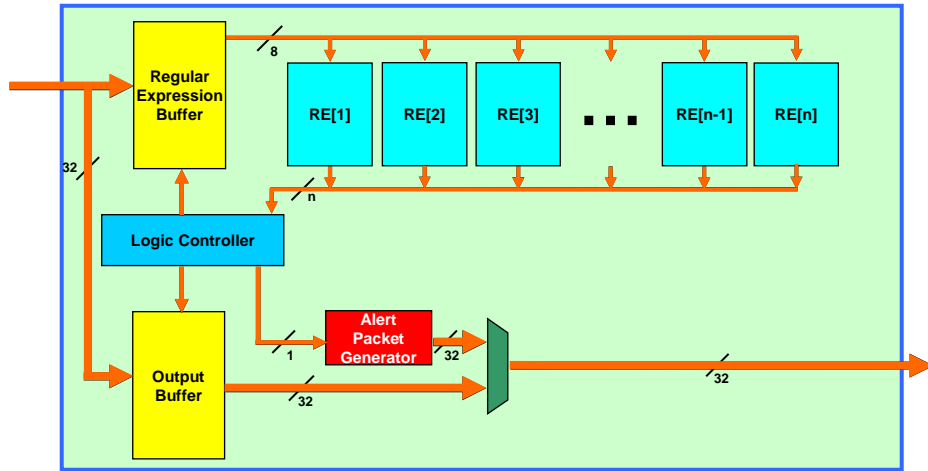


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Single Content Scanning Module

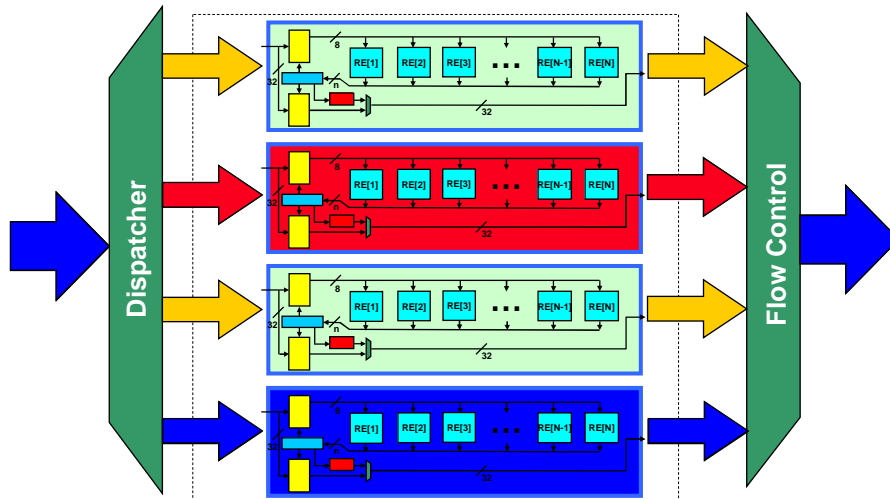


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Four Parallel Scanners

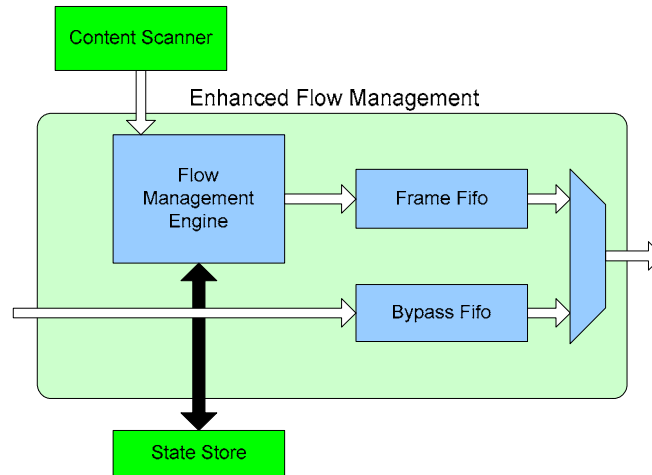


Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Enhanced Flow Management




Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Target Hardware Platform Field Programmable Port Extender (FPX)



- Oscillators
- Static Ram
- RAD (XCV2000E)
- NID (XCV600E)
- PROM




Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

FPX Based Network Appliance



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Conclusion

- Architecture for TCP/IP content scanning system
- Hardware based solution
 - OC-48 line rates (target clock rate 75–80 MHz)
- Scalable design
 - Should support OC-192 rates using latest technology
- Regular expression based content matching
- Large capacity state store
 - 8 million active flows
 - 64 bytes of per flow storage (32 bytes for scanning engine)
- Enhanced flow manipulation features
 - Flow blocking, unblocking, termination and modification



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Architecture for a Hardware Based, TCP/IP Content Scanning System

David V. Schuehler
dvs1@arl.wustl.edu



<http://www.arl.wustl.edu/arl/projects/fpx/>



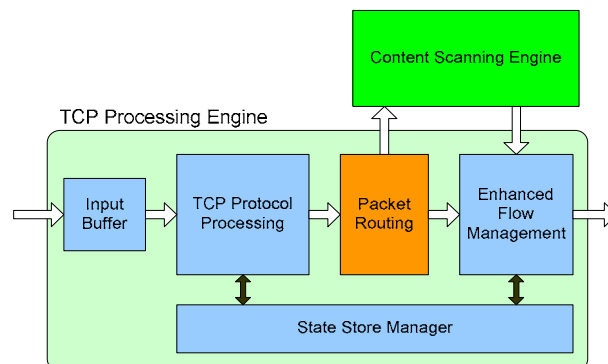
Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Packet Routing

- Scanning Engine (normal data flow)
- Outbound Processing (bypass scanning engine)
- Drop Packet (ensures in order flow of stream data)



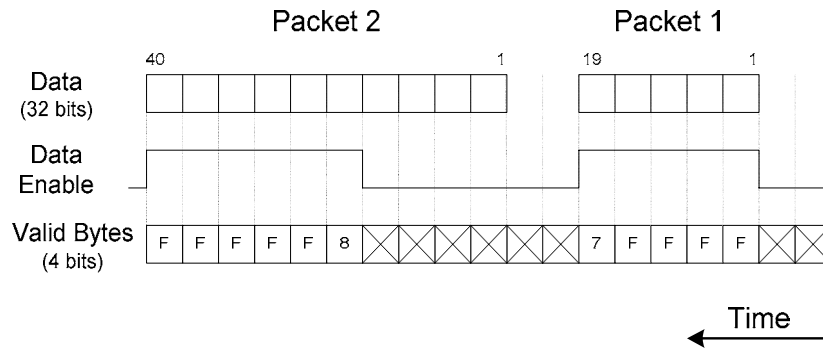
Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE

Overlapping Retransmissions

- Ensure client application receives a consistent byte stream
- Toggle TCP data enable and valid bytes signals



Department of Computer Science and Engineering
Applied Research Laboratory



Washington University in St. Louis
SCHOOL OF ENGINEERING & APPLIED SCIENCE