Data-Link Layer Traceback in Ethernet Networks

Michael Snow
Dr. Jung-Min Park

ARIAS Lab
Virginia Tech

Overview

- **Defining the Problem**
- Terminology and Standards
- Related Work
- Operating Environment
- TRACK Overview
- Simulation
- Concluding Remarks

Defining the Problem

- Internet-based attacks are increasingly common
  - Tools are readily available that allow inexperienced users to launch sophisticated attacks
- Spoofing address information is not difficult
  - Often used in Internet-based attacks
  - Common LAN and Internet protocols provide no source address authentication
  - IP and MAC addresses are easy to spoof
- Using spoofed information allows attackers to conceal their actual location

Defining the Problem

- An attacker may further hide by using intermediate hosts to send attack packets
  - Intermediate hosts referred to as “zombies”
  - Attacker sends a control packet to the zombies which in turn launch the actual attack packets
- Need to be able to find the attacker regardless of the source address on the received packets
- Traceback research has provided some solutions
  - Three-stage process

Traceback Process

- Stepping-Stone traceback
  - Attack source (zombie) to attack coordinator
- **Data-Link Layer traceback (DLT)**
  - Edge router of attack source to attack source (zombie)
- IP traceback
  - Victim to attack source network edge

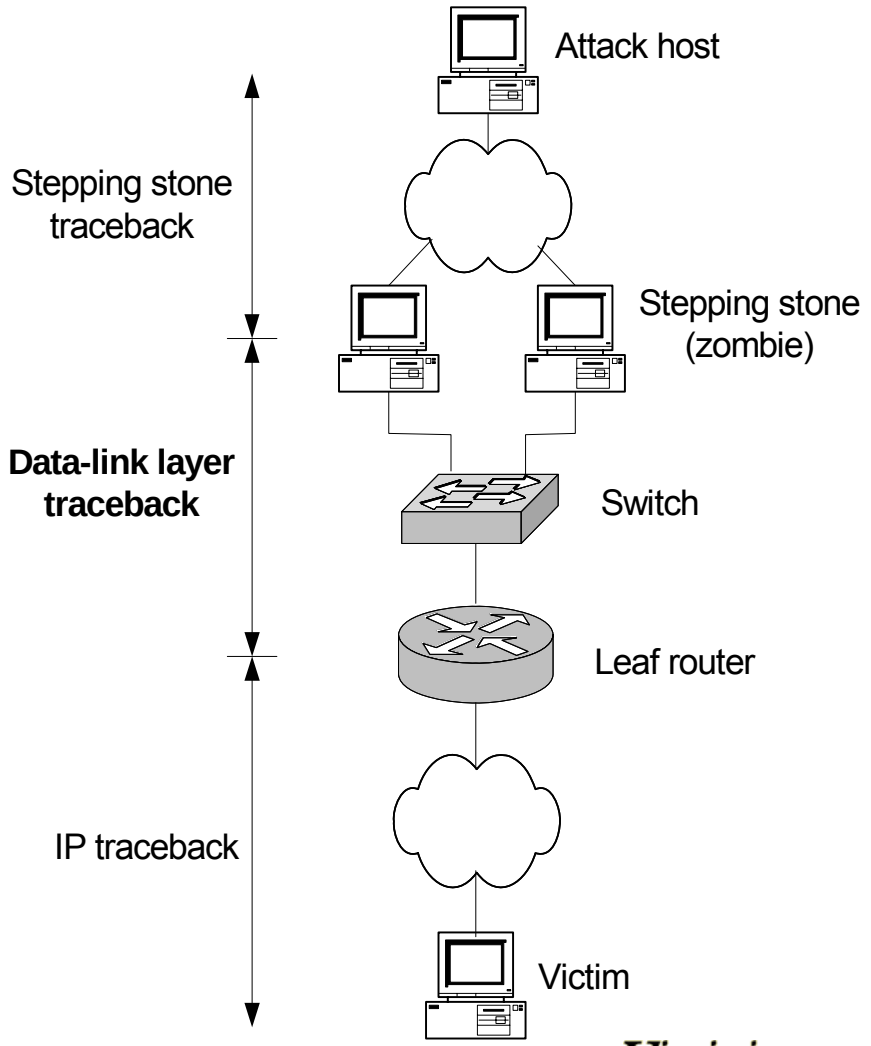

Overview

- Defining the Problem
- **Applicable Standards**
- Related Work
- Operating Environment
- TRACK Overview
- Simulation
- Concluding Remarks

Applicable Standards

- IEEE 802.3: Ethernet
  - Defines layer-1 and layer-2 protocols and formats
  - Media Access Control (MAC) address
- IEEE 802.1D: Media Access Control (MAC) Bridges
  - Frame forwarding
  - Address learning
  - Link management (Spanning Tree Protocol, STP)
- IEEE 802.1Q-2005: Virtual Bridged LAN's
  - Allows the network to be divided into segments independent of physical topology
  - Controls broadcast and forwarding behavior
  - Defines VLAN tag and Multiple Spanning Tree Protocol (MSTP)

Applicable Standards

- IEEE 802.3ac
  - Frame Extensions for Virtual Bridged LAN (VLAN) Tagging on 802.3 Networks
  - Adds 802.1Q tag and Length/Type field to 802.3 frame

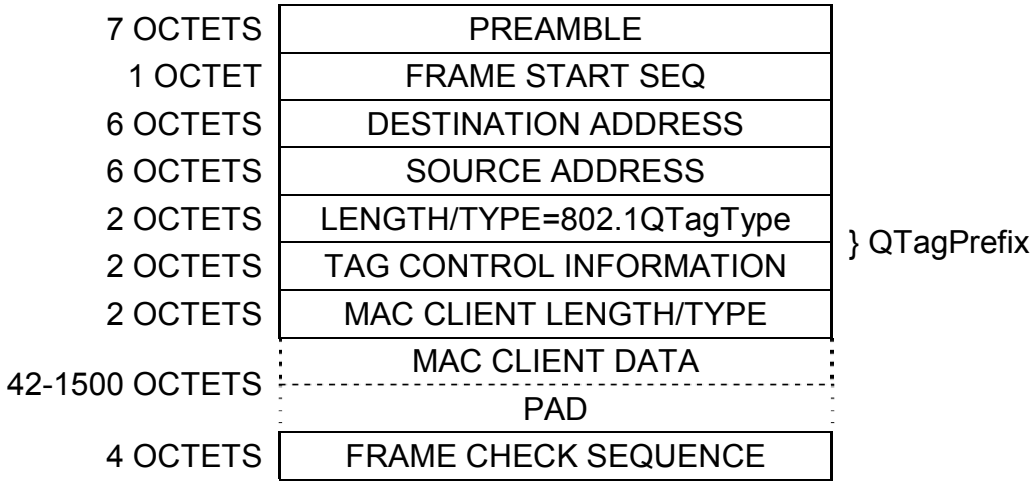

- Simple Network Management Protocol (SNMP) v3
- Bridge Management Information Base (MIB)

Overview

- Defining the Problem
- Applicable Standards
- **Related Work**
- Operating Environment
- TRACK Overview
- Simulation
- Concluding Remarks

Related Work

- Hash-Based IP Traceback
  - Proposed by Snoeren, et. al. in 2001
  - SPIE: Source Path Isolation Engine
    - DGA: Data Generation Agent
      - Generates packet digests and stores in large circular Bloom filter
      - Each traceback-enhanced router has a DGA
    - SCAR: SPIE Collection and Reduction Agent
      - Aggregates data from router DGA's
    - STM: SPIE Traceback Manager
      - Manages traceback requests
      - Builds attack graphs based on subgraphs computed by SCAR's

Related Work

- A Layer-2 Extension to Hash-Based IP Traceback
  - Hazeyama, et. al., 2003
  - Extends SPIE to map packet digests to L2 switch port ID's
    - Allows a particular network switch and port to be identified by a traceback request
  - Extends SPIE DGA to xDGA
    - Adds capability to process and store L2 information
  - Uses SNMP to collect data from switches
    - Forwarding Database (FDB) information: maps a learned MAC address to a port number
    - VLAN/Port mapping: indicates which VLANs are supported by which ports and how default VLAN ID's are assigned

Related Work

- A Layer-2 Extension to Hash-Based IP Traceback (cont'd)
  - Requires significant storage space and processing power at edge router
  - May fail in the presence of spoofed duplicate MAC addresses
    - FDB data is collected at fixed intervals (though it can be done dynamically)
    - FDB allows only one port per MAC address
    - Attacker sends packets infrequently using another host's MAC address
    - Other host sends packets normally
    - Traceback implicates the wrong host (false positive)

Overview

- Defining the Problem
- Terminology and Standards
- Related Work
- **Operating Environment**
- TRACK Overview
- Simulation
- Concluding Remarks

Operating Environment

- Network Model
  - Switched, wired Ethernet (802.3, 802.1D)
  - Supports VLAN's, frame tagging, MSTP (802.1Q, 802.3ac)
  - Capable of forwarding frames greater than 1500B in length
  - Supports SNMPv3 and the appropriate Bridge-MIB data members
- Traceback Operation
  - Traceback executed on a frame from a wireless host will trace back to the access point
  - Traceback executed on a frame from a host behind a hub or an unmanaged switch will trace to the nearest managed switch
  - Traceback will only be executed from frames that left the network

Operating Environment

- Attacker Model
  - The attacker is capable of spoofing its source MAC and IP addresses
  - The attacker is best served by using the address of another host active on the network

Overview

- Defining the Problem
- Terminology and Standards
- Related Work
- Operating Environment
- **TRACK Overview**
- Simulation
- Concluding Remarks

Tagged Frame Traceback (TRACK)

- Tags data frames
  - Switch ID
  - Port ID
  - Control flags
- Provides basic authentication
- Distributes traceback support load across network
- Decreases false-positive rate
  - Each frame is tagged with entry point
  - Does not depend on MAC address
- Defines two processing entities
  - TRACK Frame Tagger (TFT)
  - TRACK Analysis and Collection Host (TACH)

TRACK

- Tagged Frame Format
  - Based on VLAN-enabled Ethernet standard (802.3ac)

7 OCTETS	PREAMBLE
1 OCTET	FRAME START SEQ
6 OCTETS	DESTINATION ADDRESS
6 OCTETS	SOURCE ADDRESS
2 OCTETS	LENGTH/TYPE=tbt tag
2 OCTETS	802.1Q Tag
13 OCTETS	**TRACEBACK TAG (TBT)**
2 OCTETS	MAC CLIENT LENGTH/TYPE
29-1500 OCTETS	MAC CLIENT DATA
	PAD
4 OCTETS	FRAME CHECK SEQUENCE

TRACK

- Traceback Tag (TBT)
  - FLAGS (4 bits)
  - PORT_ID (12 bits) – the port number on the switch (0-4095)
  - SWITCH_ID (8 bits) – the switch ID (0-255)
  - AUTH (80 bits) – truncated HMAC for authentication

FLAGS	PORT_ID	SWITCH_ID	AUTH
4 bits	*12 bits*	*8 bits*	*80 bits*

- FLAGS field

FLAGS			
PREV	CORE	FAIL	*(reserved)*
bit 7	*6*	*5*	*4*

  - PREV: a previous tag was removed
  - CORE: the frame was not tagged at the network edge
  - FAIL: authentication failed (used for storage)
  - 1 bit reserved

TRACK

- TRACK Frame Tagger (TFT)
  - Exists in each TRACK-enabled switch
  - Responsible for tagging frames as they enter the network

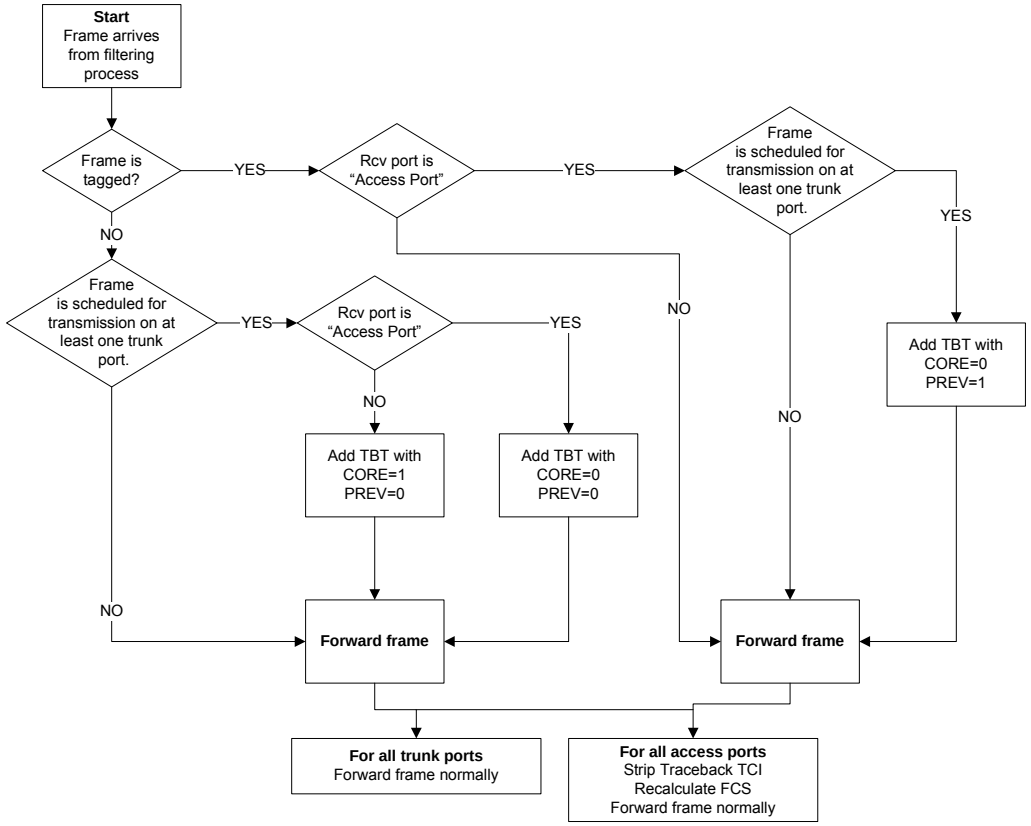

TRACK

- TRACK Analysis and Collection Host (TACH)
  - Stores frame digests as in SPIE
  - Maps frame digest to switch/port ID based on TBT data
    - Host Table and L2 Digest Table
  - Connected at network edge near the router
    - If attached to the root switch or router, frames are mirrored to the TACH for processing
    - If connected between the root switch and router or is internal, the TACH captures and processes every frame
  - Handles L2 traceback requests
    - Initiated by the administrator
    - Forwarded by the IP traceback engine

TRACK

- TACH L2 Digest Storage and Lookup
  - Separate tables for host info and L2 digests
    - Optimized for high-speed operation in the general case
  - Digest table paged periodically
  - Traceback requires an L2 table lookup to get the host table entry

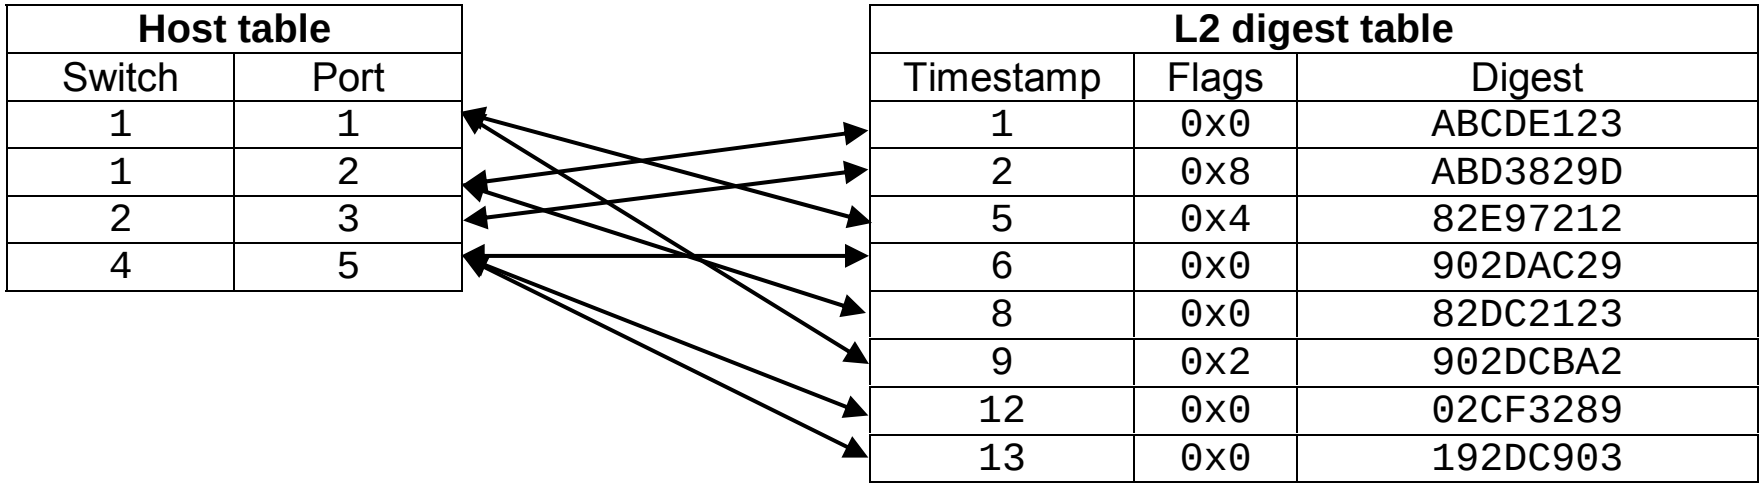

Host table	
Switch	Port
1	1
1	2
2	3
4	5

L2 digest table		
Timestamp	Flags	Digest
1	0x0	ABCDE123
2	0x8	ABD3829D
5	0x4	82E97212
6	0x0	902DAC29
8	0x0	82DC2123
9	0x2	902DCBA2
12	0x0	02CF3289
13	0x0	192DC903

TRACK

- Protocol Security
  - Keyed HMAC (RFC 2401)
    - Truncated result
  - First 32 Bytes of MAC data as input
  - Global or pairwise-shared key
    - Randomly generated or based on administrative password
    - Pairwise keys generated from root key similarly to SNMPv3
    - Can be distributed using SNMP with encryption
    - Should be refreshed periodically

TRACK

- Benefits
  - Allows near-zero false positive rate
  - Removes load from edge router
- Limitations
  - May require separate host (TACH)
  - Requires switch OS or hardware upgrade
  - Increases load on edge switches
  - Requires additional bandwidth
- Can be dynamically enabled and disabled to balance limitations and benefits

Overview

- Defining the Problem
- Terminology and Standards
- Related Work
- Operating Environment
- TRACK Overview
- **Simulation**
- Concluding Remarks

Simulation Framework

- Simulation work performed in OpNET
- Added TBT to Ethernet v2 frame format
- TFT implemented in MAC Relay Entity (MRE) process model
  - Fixed per-frame HMAC calculation time
- Multiple VLANs and redundant links
  - Managed by IEEE MSTP running normally
- TACH is not simulated

Simulation Framework

- Node ID
  - <VLAN ID><Node ID>
- Metrics
  - Ethernet delay
  - MRE queue length
  - Application throughput
- Traffic
  - Light: All hosts HTTP
  - Medium: 6 FTP + 15 HTTP
  - Heavy: 12 FTP + 15 HTTP
- TRACK speed
  - Disabled (baseline)
  - Fast (9.25 µs), Medium (34.83 µs), Slow (143 µs)

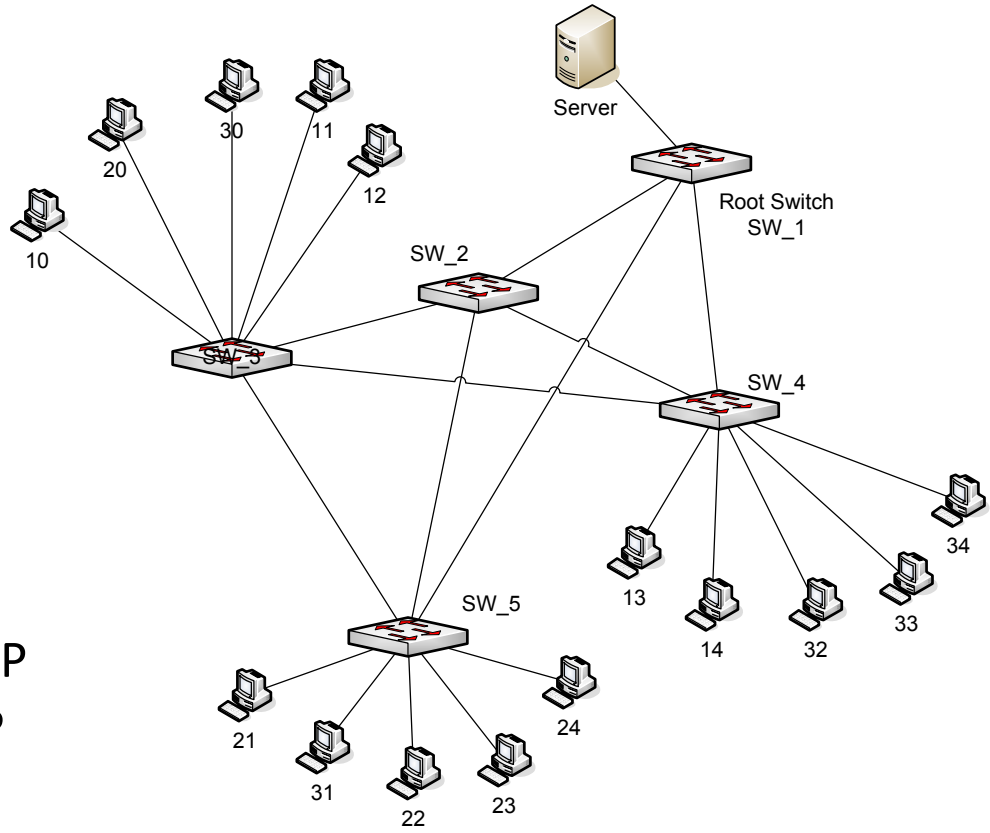

Simulation Results

- Ethernet delay, medium load

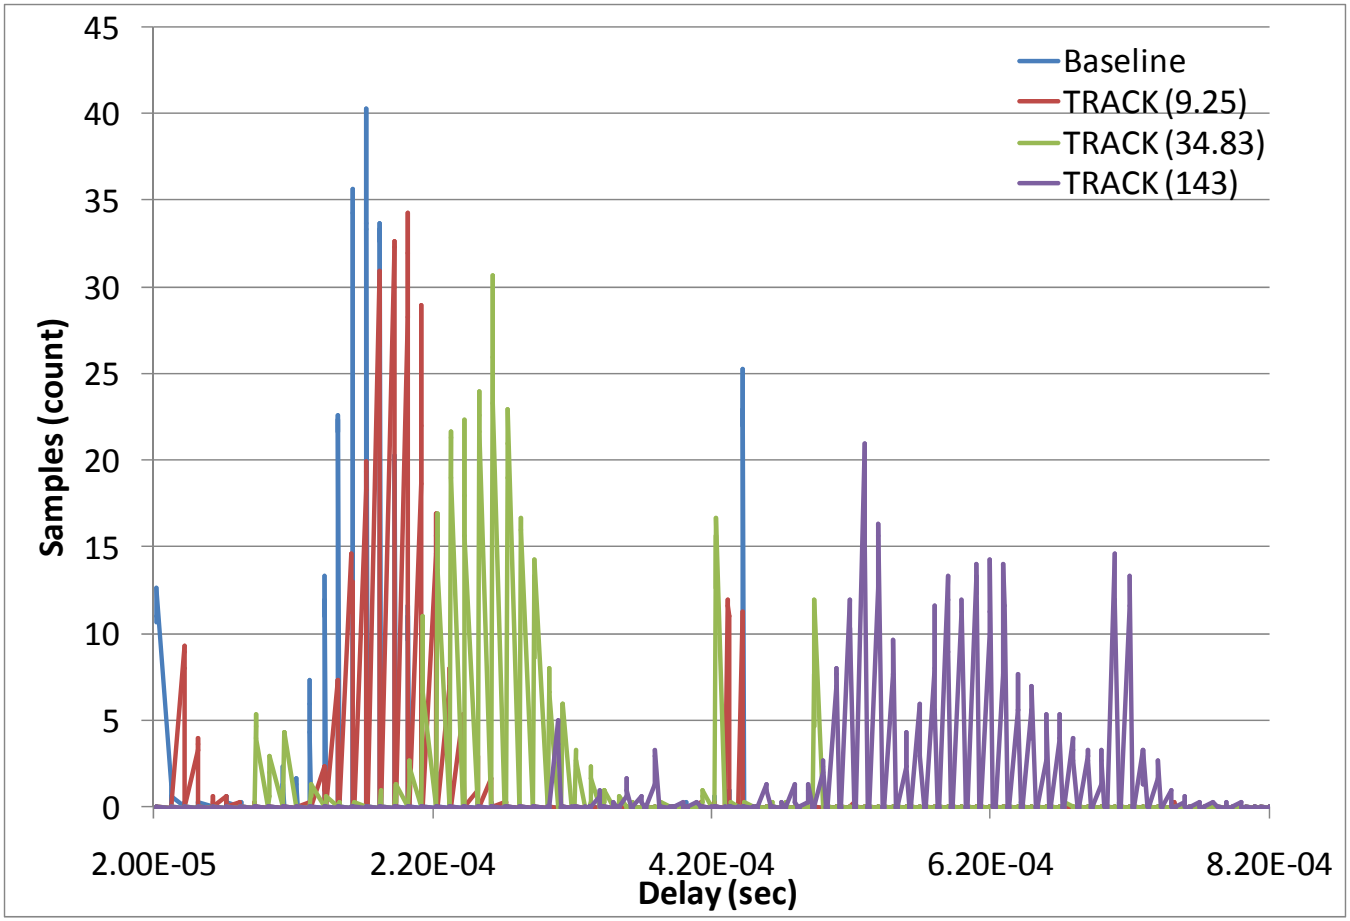

Simulation Results

- MRE Queue Length, Medium Load

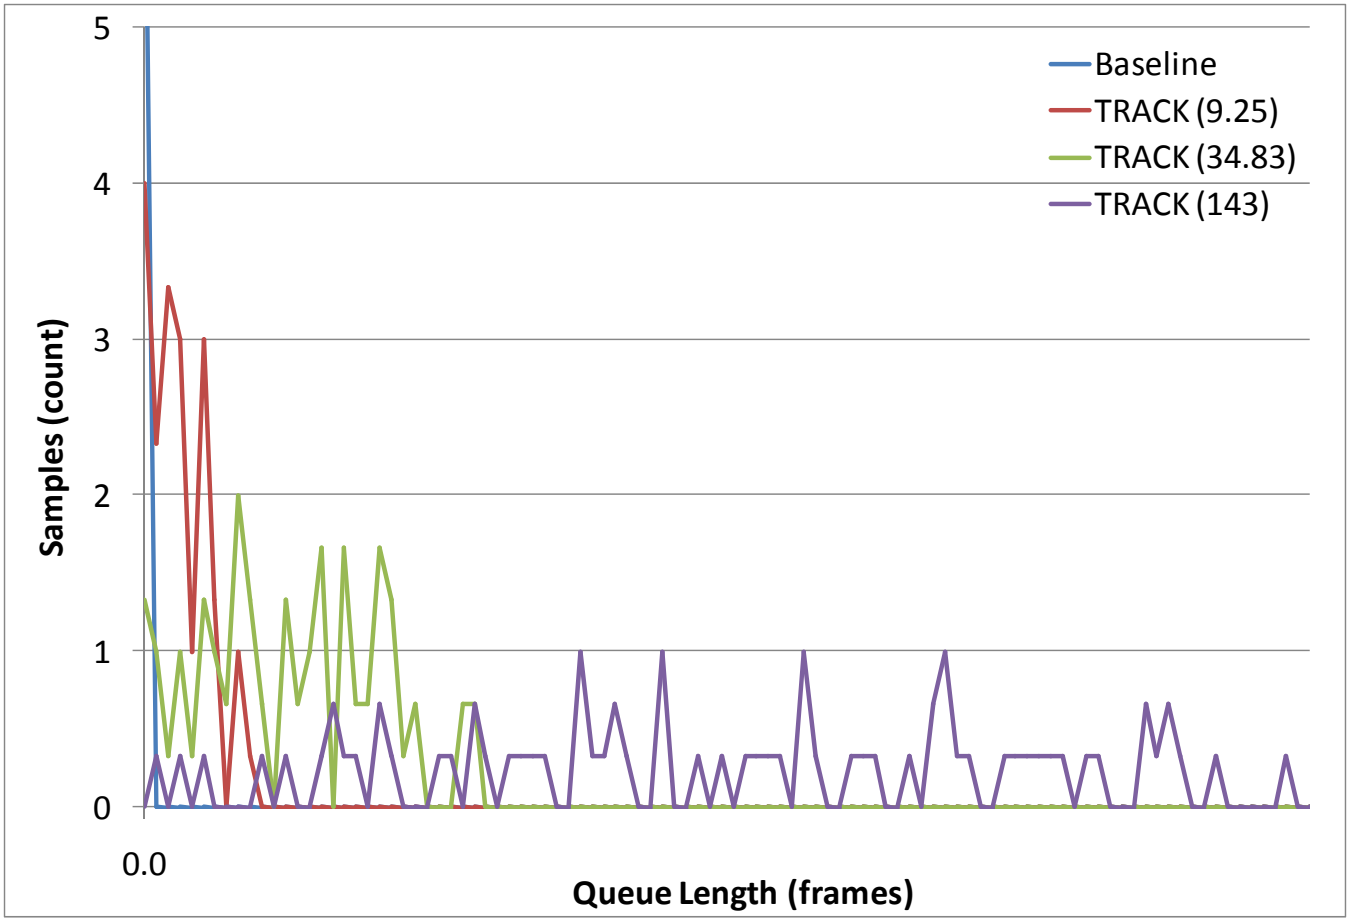

Simulation Results

- Ethernet Delay, Medium Load

	Baseline	Fast	Medium	Slow
Value	199.0µs	214.5µs	278.5µs	616.5µs
Diff		7%	33%	102%

- MRE Queue Length, Medium Load

	Baseline	Fast	Medium	Slow
Value	2.91×10^{-4}	1.32×10^{-3}	5.06×10^{-3}	2.61×10^{-2}
Diff		128%	178%	196%

Overview

- Defining the Problem
- Terminology and Standards
- Related Work
- Operating Environment
- TRACK Overview
- Simulation
- **Concluding Remarks**

Concluding Remarks

- L2 traceback is a difficult area lacking much research
- Only one major published scheme
  - Places additional load on edge router
  - Susceptible to well-crafted attacks
- TRACK – proposed frame tagging scheme
  - Improves traceback reliability
  - Reduces edge router load and storage requirements
  - Does not significantly impact application performance
- Future Work
  - Dynamic TRACK – pay attention to significant events
  - Better tag integration scheme

Questions?