

DIFFUSE Embedded

Nigel Williams

williams.nigel@gmail.com

Centre for Advanced Internet Architectures (CAIA)
Swinburne University of Technology



Overview



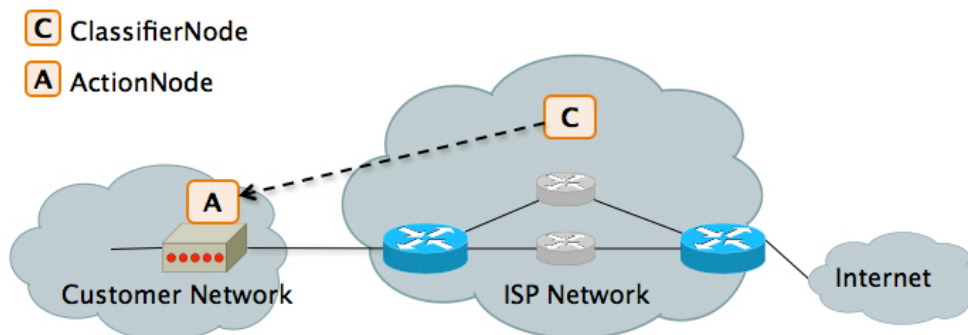
- Can we build and run DIFFUSE on an embedded Linux device?
- Background: DIFFUSE and OpenWRT
- Preliminary Results
- Summary

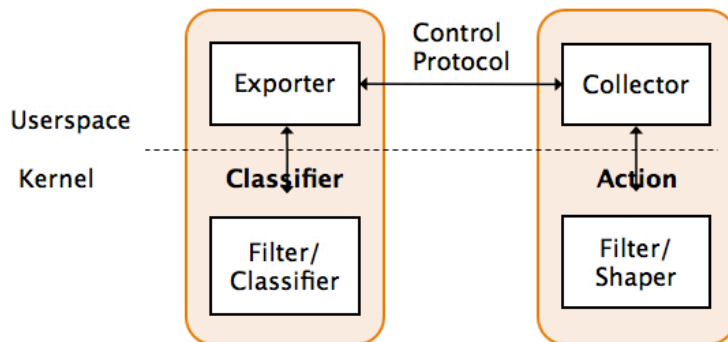
DIFFUSE



- DIFFUSE: Distributed Firewall and Flow Shaper Using Statistical Evidence
 - Automated IP flow classification
 - Measure flow statistics (e.g. packet length) and Machine Learning to make a classification model

DIFFUSE Architecture





DIFFUSE



■ Action Node/Classifier running on CPE

- Can we use existing hardware?
- For home users this would be a consumer-grade router or modem/router combo
- What Operating System?

OpenWRT



- A number of Home Routers already run embedded linux - but customisation is limited.
- OpenWRT can replace factory firmware
 - Flexibility: pre-built packages, sources, cross-compilers
 - Linux/BSD programs
 - Many routers supported: <http://wiki.openwrt.org/toh/start>

OpenWRT



- Image Builder/SDK available from the website:
 - www.openwrt.org
- Used trunk “Attitude Adjustment”
 - Needed SCTP support

TPLink 1043ND



■ Ethernet Switch and Router

- Atheros AR9132 400MHz (MIPS)
- 32MB RAM
- 8MB Flash
- GigE (4xLAN, 1xWAN)
- USB

■ Mid-range specs

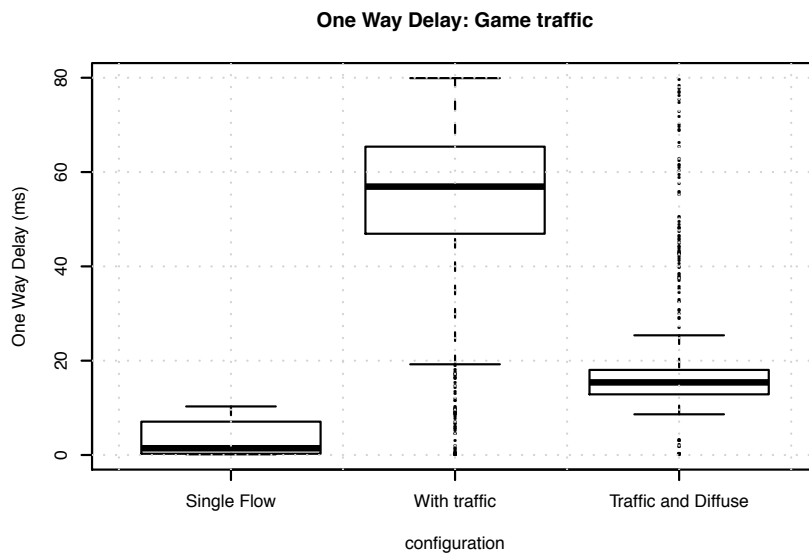
- See OpenWRT Table of hardware: <http://wiki.openwrt.org/toh/start>

DIFFUSE on the TPLink



■ DIFFUSE in action

- Classify and prioritise Enemy Territory traffic



Performance




- It works, but what are the performance penalties?
 - Is the router still usable under “realistic” conditions
 - Frames per second, throughput, CPU Utilisation
 - Measure CPU
 - Memory
 - Keeping state for concurrent flows

Testbed



- Used Ninjabox and binary search to determine throughput
 - Send UDP packets at different rates and frame size, look for packet loss
 - Not strictly designed for testing devices... requires some manual labour




Testbed



■ Tcpreplay

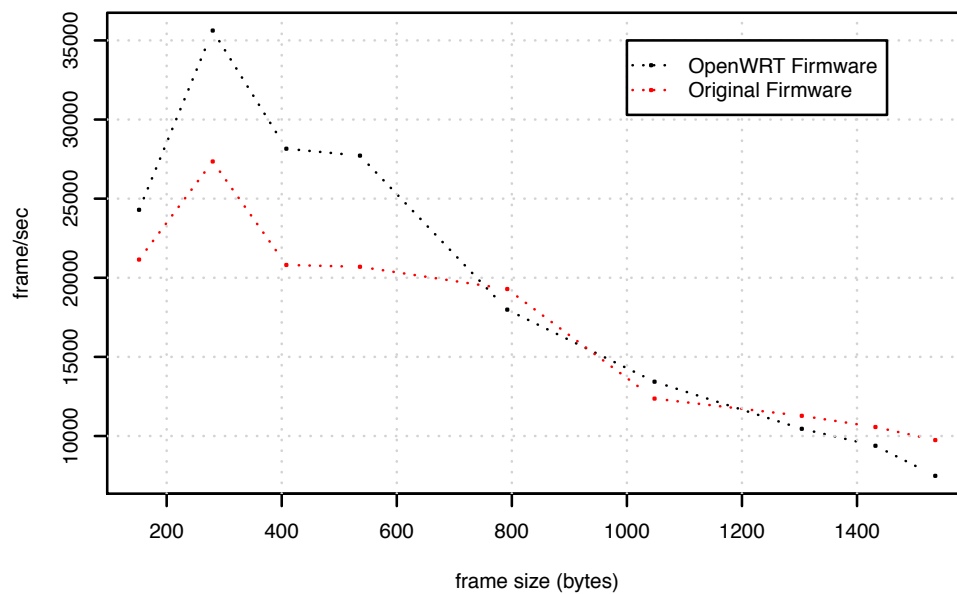
- Generate "other" traffic and send upstream
- Replay game traffic across the router



Throughput



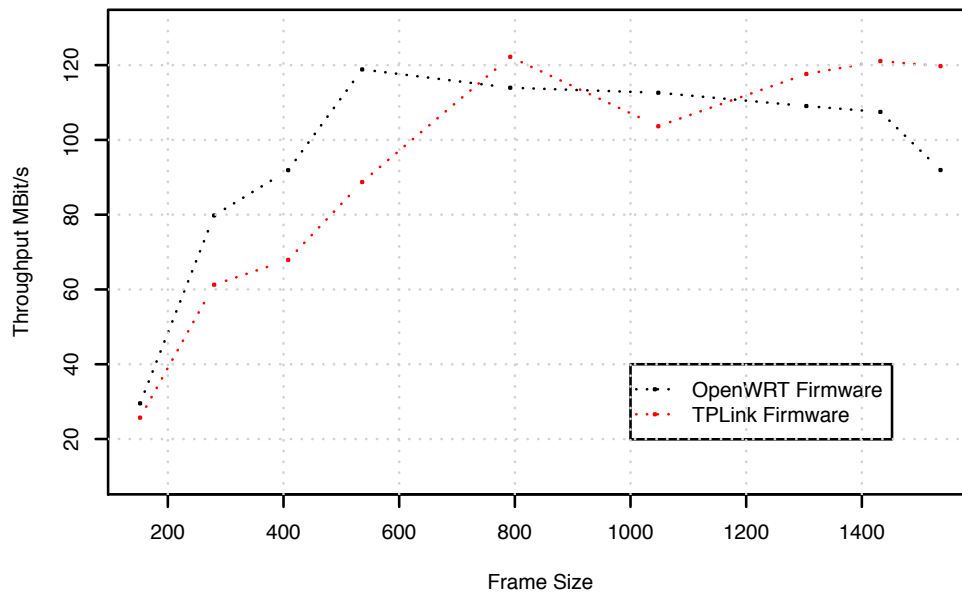
Router Throughput: Frame size Vs Frame rate



Throughput



Router Throughput: BW Vs Frame Size



CPU Utilisation




- CPU measurements were noisy - see Tech Report
- How does DIFFUSE affect the CPU Load
 - As a Classifier Node and as Action Node with IPFW
 - Used a number of methods
 - mpstat tool
 - CPU counters from /proc/stat

CPU Utilisation



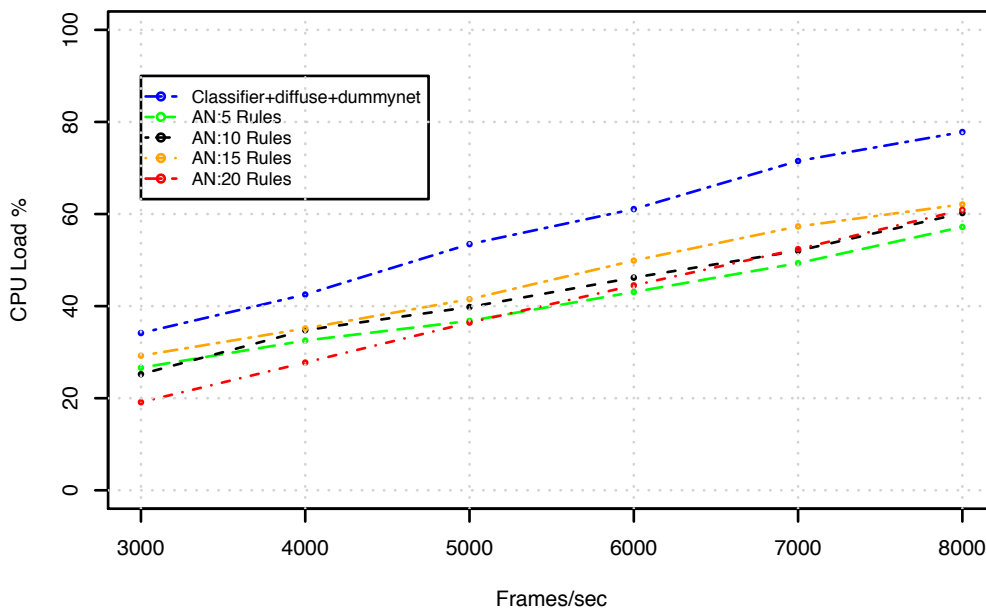
CPU Load Vs Packet Rate, 512 Byte packet



CPU Utilisation



AN CPU Load and Number of IPFW Rules



Memory



- Number of concurrent connections limited by RAM
 - Need to maintain an entry for each flow
 - TPLink firmware and OpenWRT defaults:
 - TPLink firmware: 5,100 flows
 - OpenWRT: 16,364 flows
 - Measuring maximum number of flow:
 - Initiating more and more new TCP connections across router
 - Approximately 32,000 connections

Summary



- Built DIFFUSE for OpenWRT on TPLink
- The router was able to act as a classifier node and as an action node
 - Firmware image, packages for TPLink 1043ND
- There is a performance hit on throughput and number of connections
 - 8000pps, ~32Mbit/s
 - There is enough RAM to handle several thousand concurrent flows
 - And routers are getting more RAM and faster CPUs...

Future Work



- Encountered a number of problems that limited testing time
 - Software bugs, BSD to Linux glue code
- More testing scenarios
 - 100's of concurrent flows, with many 'real-time flow' matches
 - Native linux queuing implementation for Action Node