SWIN BUR NE

CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Measuring the Performance of NetSniff

Julie-Anne Bussiere

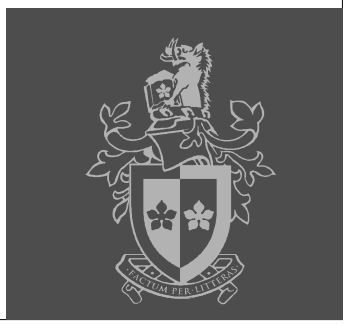

Introduction : NetSniff

- ICE^{3}project (Inverted Capacity Extended Engineering Experiment)
- IP traffic capture and analysis tool
  - PCAP library / tcpdump files
  - Parse traffic to text logfiles
  - Produce statistics
  - Enable data anonymisation (privacy)

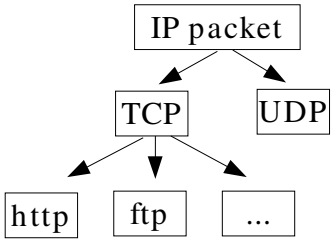

- Figure out IP traffic characteristics, traffic modelisation
- Currently used in low traffic and low bandwidth scenarios : not sufficient for relevant traffic models

Overview

- Purpose of measuring NetSniff performance
- Create traffic
  - Testbed design & configuration
  - TCP tools
- Processing performance
  - Studied traffic
  - Performance results
- Conclusion and future work

Purpose

- **Are there any performance limitations using NetSniff within large network environment ?**

Purpose

- **Are there any performance limitations using NetSniff within large network environment ?**

Not allowed to use existing network traffic at the moment

→ Testbed to generate realistic IP traffic :

- different hosts concurrently generating traffic
- TCP applications: http, https, ftp, ssh, smtp

→ Process on different Hardware boxes

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Testbed design

- Traffic generator tools : no application-layer
- Remote Unix Lab Environment (RULE): 1 physical box, multiple virtual hosts with different IP addresses

+ Local server : http and https, ftp, ssh, smtp

+ Record traffic between virtual hosts and server

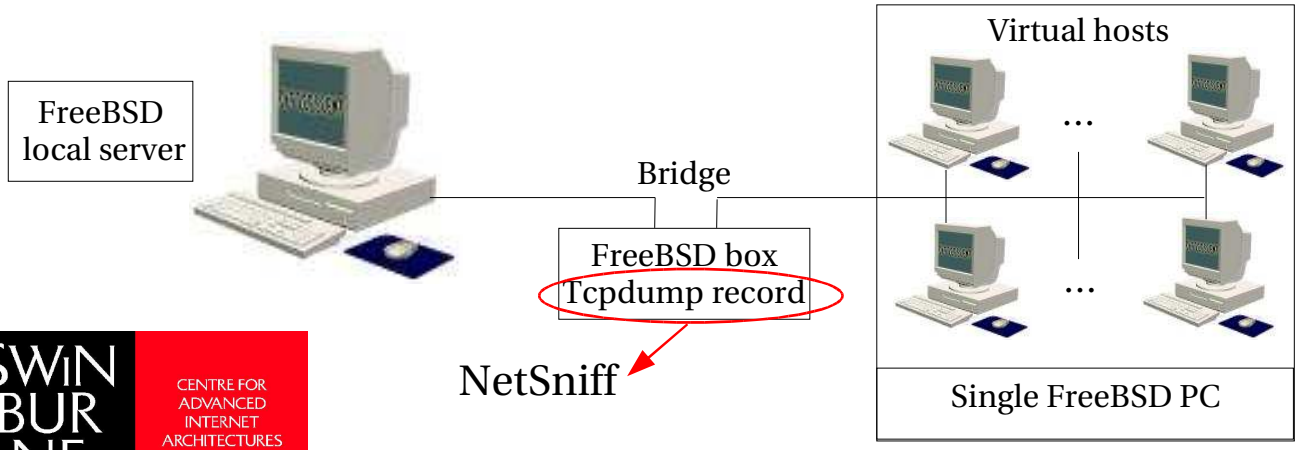

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Testbed configuration

- Server: Apache +mod SSL, enable ftpd and smtpd
- 20 virtual hosts
  - Packages Perl5, wget
  - ssh and ftp with automatic login
- Shell script launching TCP applications successively
- Unix “at” command
- Two ways of using records
  - Live capture : tcpreplay from a remote box
  - NetSniff processes directly the tcpdump files

Generated traffic

- 1 host successively runs:
- HTTP / HTTPS : download (encrypted) html pages
- FTP : download / upload files of different sizes (KB to MB)
- SSH : login on the server and execute commands
- SCP : download / upload files
- SMTP : telnet to the server on port 25
- ... loop

Records with 1 up to 20 concurrent hosts

TCP tools

- Create new traffic files from recorded ones (more hosts)
- **Tcpslice**

  Part of a file between t1 and t2
- Tcprewrite (tcpreplay)

  Change all IP addresses
- Mergecap (ethereal)

  Join 2 tcpdump files

  with concurrent timestamps

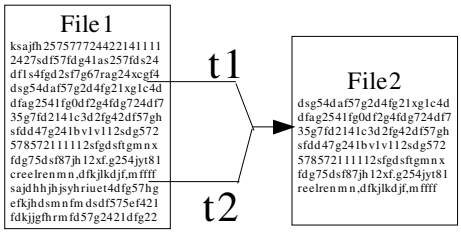

Traffic files of 3 minutes

TCP tools

- To create new traffic files from recorded ones
- Tcpslice

  Part of a file between t1 and t2
- **Tcprewrite (tcpreplay)**

  Change all IP addresses
- Mergecap (ethereal)

  Join 2 tcpdump files

  with concurrent timestamps

192.168.0.101 ... 192.168.0.110

↓ ↓

192.168.25.67 ... 192.168.25.76

Duplicate files changing IP addresses

TCP tools

- To create new traffic files from recorded ones
- Tcpslice

  Part of a file between t1 and t2
- Tcprewrite (tcpreplay)

  Change all IP addresses
- **Mergecap (ethereal)**

  Join 2 tcpdump files

  with concurrent timestamps

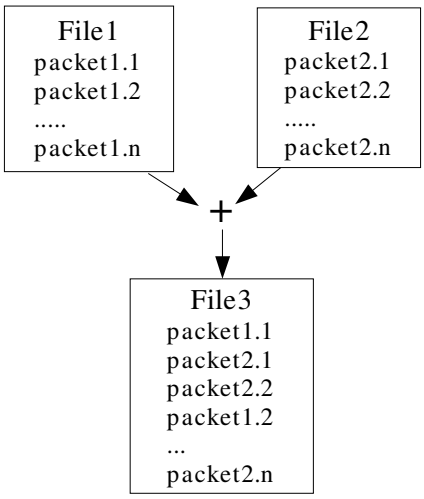

Join original and duplicated files to double the number of hosts

Processing Performance

- Measure the impact of the number of concurrent hosts generating IP traffic :
  - Processing user time
  - CPU use
  - Process size (= Memory usage)
- Traffic files, 1 up to 60 hosts (1 host ◄► 1 application at a time)
- 4 available boxes

Machine	Processor type	CPU (GHz)	RAM (MB)
Box 1	Pentium 4	2.66	512
Box 2	Pentium 4	2.8	512
Box 3	Celeron	2.4	256
Box 4	Pentium 4	2.66	256

Processing Performance

- Studied traffic : files of 3 minutes
- Created files
  - 2 x 15 hosts -> 30 hosts, 2 x 20 hosts -> 40 hosts,

  2 x 30 hosts -> 60 hosts (4 x 15 hosts)
- Recorded traffic / manipulated traffic files
  - Processing user time : 15% / 30% longer
  - CPU use : no significant difference
  - Memory usage : directly related to the original file's process size

Processing Results

- Processing rate

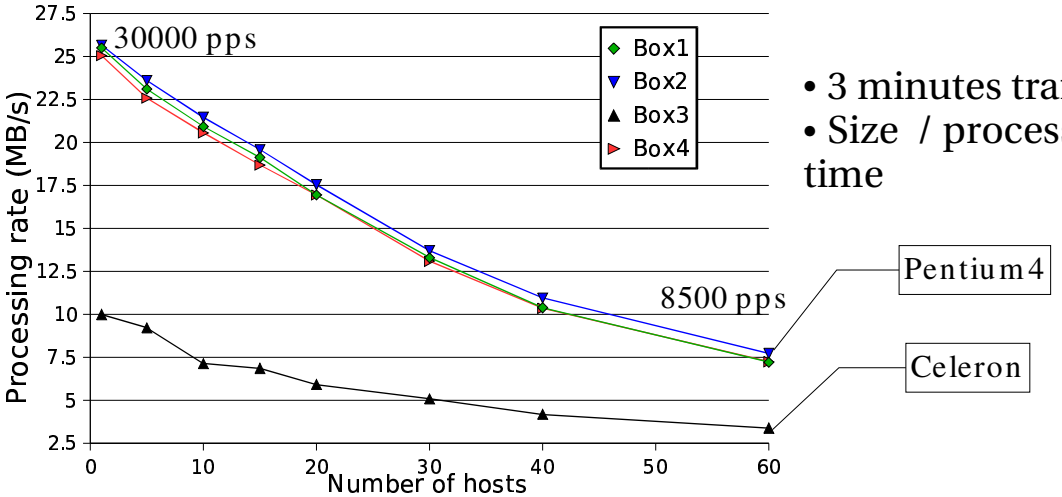

- 3 minutes traffic files
- Size / processing user time

Pentium4

Celeron

→ Concurrent tcpstreams
→ Processor type / CPU clock
→ Expect better perf. from 30 hosts to 60

Processing Results

- Speed ratio : capture time / processing user time

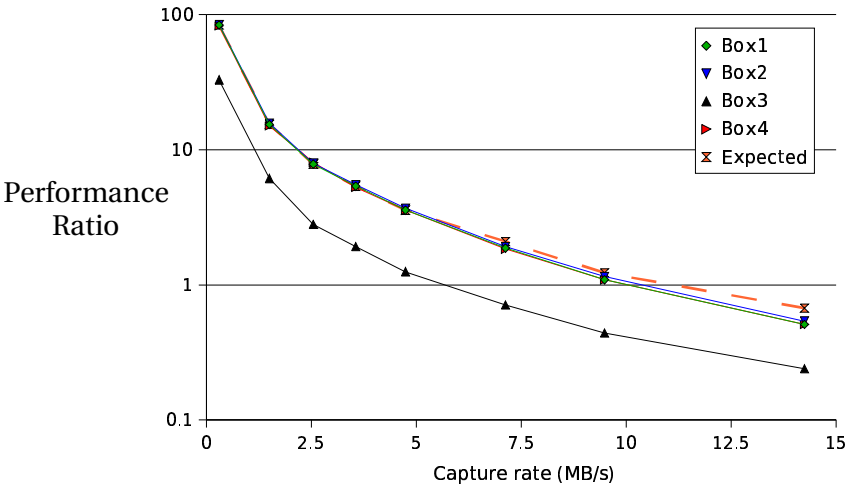

- Logarithmic scale
- Approximation: user time – 15% / 30%

Speed ratio < 1 ◄► packets dropped at live capture ?

- P4 : 10MB/s (~40 hosts), Celeron : 6MB/s (~25 hosts)
- Approx : 11.1 MB/s (~50 hosts)

Processing Results

- Mean CPU use during processing

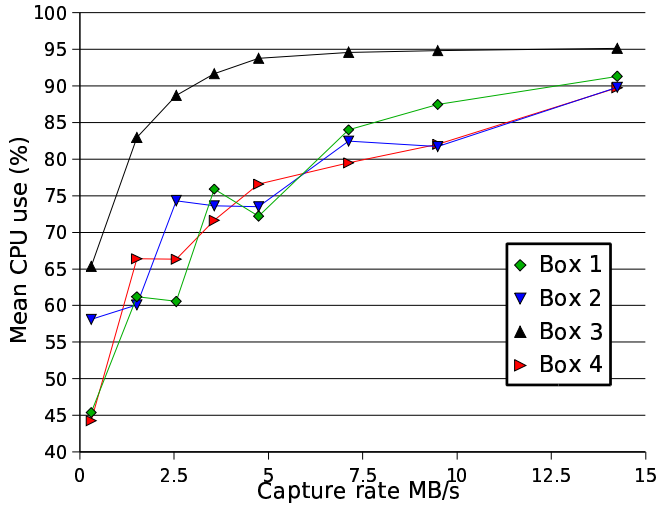

- Average of values taken each second during the processing

→ processor impact
→ high traffic : 90% of cpu

Processing Results

- Memory usage

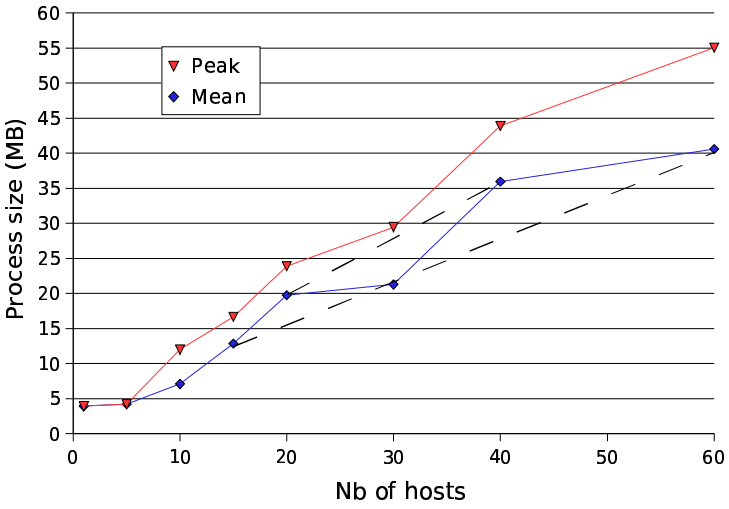

• Values taken each second during the processing

• Mean : polynomial model up to 20 hosts

• Created files : linearly related to traffic files they were created from

Processing Results

- Memory Usage

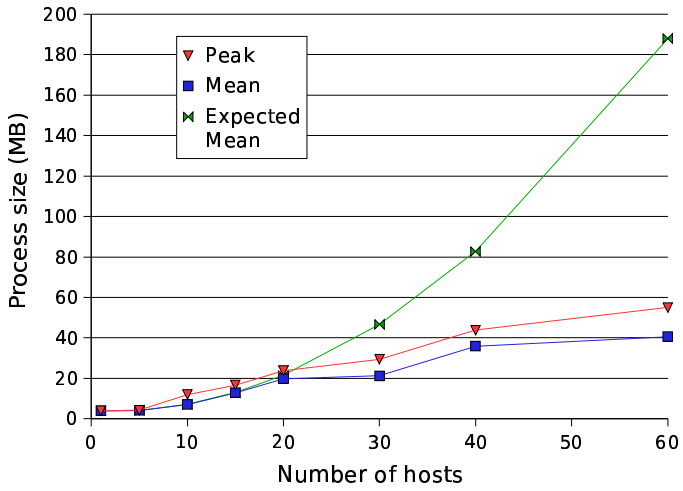

Approximation
Expected mean f(nb hosts)

Up to 190 MB for 60 hosts (14.24MB/s)
→ At least 256 MB RAM

Anonymisation impact

- Privacy reason : eliminate confidential information
- NullIP : set all IP addresses to 0.0.0.0
  - No impact on performance
- Tcpdpriv : prefix-preserving, table based, reversible
  - Default anonymisation mode
  - No significant impact on processing
- CryptoPan : key-based
  - Requires more user time (+ 50%), more CPU use

SWiN BUR NE
CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Conclusion

- Processing performance highly dependent on the number of concurrent flows
  - “tcpstream” parser section = implementation critical point
- Celeron / Pentium 4 : L2 cache memory & bus speed
- CPU clock size measurable
- RAM could constitute a limitation for higher numbers of concurrent tcp streams
- Default mode for anonymising the traffic does not impact performance
- Realistic traffic : realistic packets, but not realistic network behaviour

SWiN BUR NE
CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Future work

- Analysis of live capture performance
  - Replay traffic files (tcpreplay) varying speed
  - Test different HW boxes, network card
  - Analyse % dropped packets with high traffic
- Provide a recommendation for NetSniff processing and live capture performance
  - Hardware
  - Supported traffic

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES

The end

Thank you...

Questions ?

SWINBURNE CENTRE FOR ADVANCED INTERNET ARCHITECTURES