

Monitoring high-speed networks using ntop

Luca Deri <deri@ntop.org>

Project History

- Started in 1997 as monitoring application for the Univ. of Pisa
- 1998: First public release v 0.4 (GPL2)
- 1999-2002: Registered ntop.org, created mailing lists (ntop and ntop-dev) port to several platforms, part of many Linux distributions.
- 2002-03: Version 2.x, added support for commercial protocols (NetFlow v5 and sFlow v2).
- 2004-05: Version 3.x (many parts have been recoded), added RRD support, IPv6 (Loria) and SCSI/FibreChannel (Cisco) support, NetFlow V9/IPFIX draft, sFlow v5.

What is ntop ? [1/2]

ntop is a simple, open source (GPL), portable traffic measurement and monitoring tool, which supports various management activities, including network optimization and planning, and detection of network security violations.

What is ntop ? [2/2]

The screenshot shows the ntop web interface in a browser window. The browser's address bar displays `http://mon03.consiagnet.it/`. The page title is "Global Traffic Statistics". The ntop logo is visible in the top left, and navigation links for "About", "Summary", "All Protocols", "IP", "Media", "Utils", "Plugins", and "Admin" are present. A copyright notice "(C) 1998-2005 - Luca Deri" is in the top right.

Global Traffic Statistics

Network Interface(s)	Name	Device	Type	Speed	Sampling Rate	MTU	Header	Address	IPv6 Addresses
	fxp2	fxp2	Ethernet		0	1514	14	0.0.0.0	
	Consiag	NetFlow-device.2	Ethernet		1	1514	14	192.168.0.0	

Local Domain Name: consiagnet.it
Sampling Since: Fri Apr 29 10:31:51 2005 [11:18]
Active End Nodes: 366

Traffic Report for 'fxp2' [switch]

What can ntop do for me?

- ntop has been created to solve a real monitoring problem (no planning, case studies, market analysis).
- By the time it has been extended to satisfy user requirements.
- Portable and platform neutral: deploy it wherever you want with the same look and feel.
- Minimal requirements to leverage its use.
- Suitable for monitoring both a LAN (default) and a WAN (don't forget to configure ntop properly).

Traffic Measurement

- Data sent/received: Volume and packets, classified according to network/IP protocol.
- Multicast Traffic.
- TCP Session History.
- Bandwidth Measurement and Analysis.
- VLAN/AS traffic statistics.
- VoIP (SIP, Cisco SCCP) Monitoring.

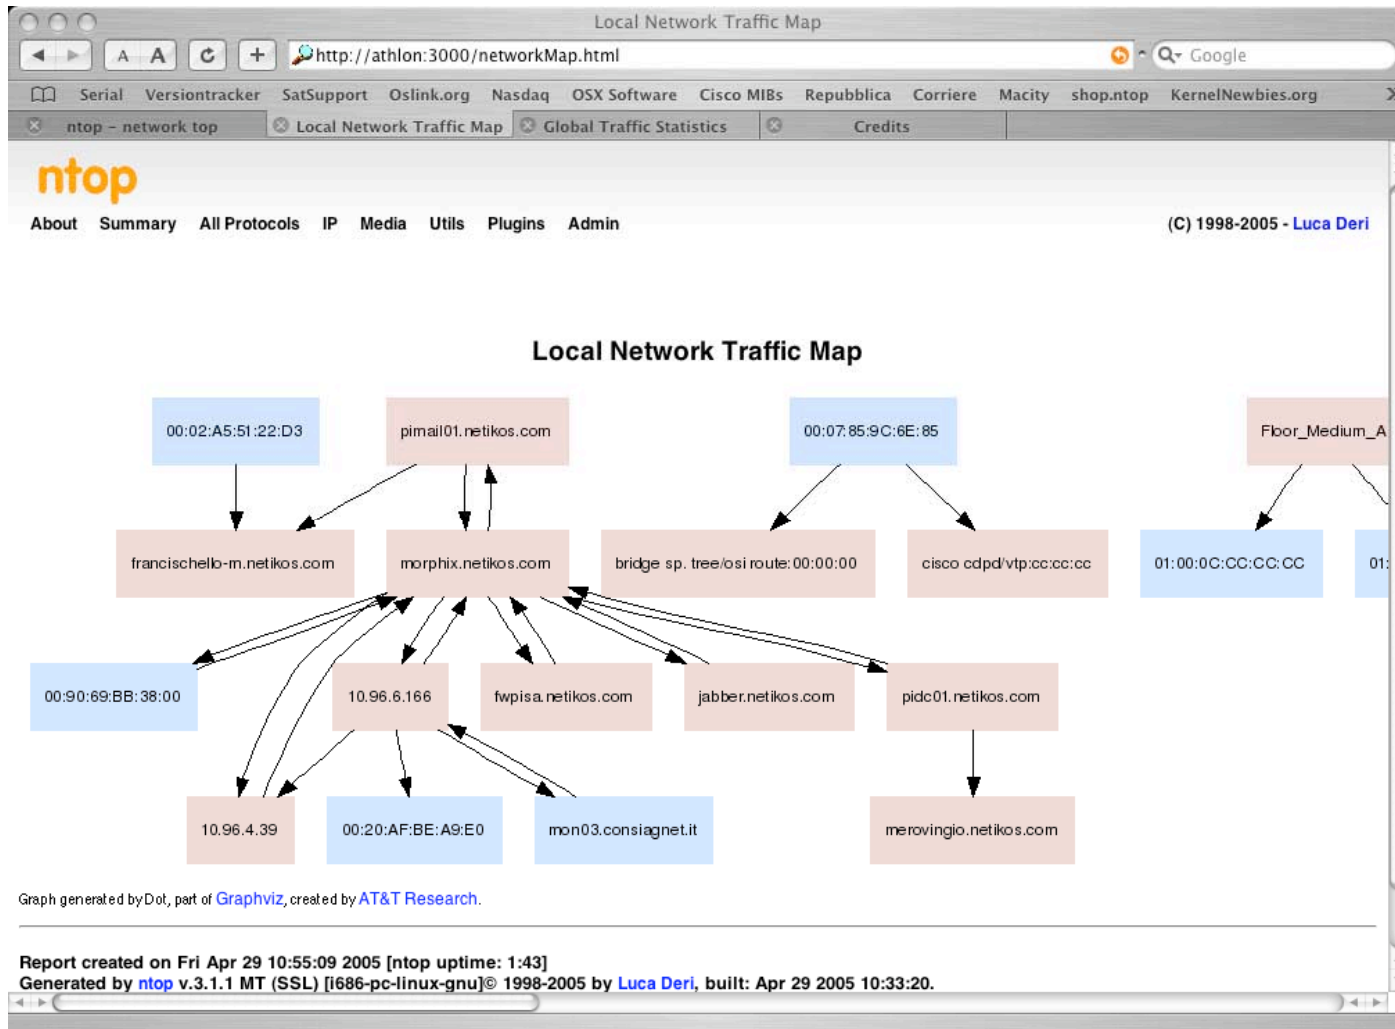
Traffic Characterisation and Monitoring

- Network Flows (user configurable)
- Protocol utilisation (# req, peaks/storms, positive/negative repl.) and distribution.
- Network Traffic Matrix.
- ARP, ICMP Monitoring.
- Detection of many popular P2P protocols (Caida paper)

Network Optimisation and Planning

- Passive network mapping: identification of Routers and Internet Servers (DNS, Proxy).
- Traffic Distribution (Local vs. Remote).
- Service Mapping: service usage (DNS, Routing).
- Network traffic map (Graphwiz)

Network Traffic Map



Network Inventory [1/2]

- Identification of routers and internet servers (DNS, NFS, proxy).
- Resource (Hw Manufacturer), services and OS inventory.
- Unhealthy hosts.

Network Inventory [2/2]

Local Hosts Characterization

http://mon03.consiagnet.it/localHostsCharacterization.html

Serial Versiontracker SatSupport Oslink.org Nasdaq OSX Software Cisco MIBs Repubblica Corriere Macity shop.ntop KernelNewbies.org

ntop - network top Local Hosts Characteriz... Local Hosts Characteriz... Credits

ntop

About Summary All Protocols IP Media Utils Plugins Admin (C) 1998-2005 - Luca Deri

Local Hosts Characterization

Host	Unhealthy Host	L2 Switch Bridge	Gateway	Printer	NTP/DNS Server	SMTP/POP/IMAP Server	Directory/FTP/HTTP Server	DHCP/WINS Server	DHCP Client	P2P
0.0.0.0	X									
host059-160	X									
host062-160	X									
host053-160	X									
host003-160					X					
host005-160	X									
host029-160	X									
host028-160						X				
dns03.ablia.net	X				X					
dns02.ablia.net	X				X	X	X			
dns01.ablia.net	X				X	X	X			
host119-160	X				X	X	X			
host118-160					X					
host117-160	X					X	X			
host074-160						X				
host073-160						X				
host066-160	X					X				
host069-160						X				
host068-160						X				

Host Fingerprint

The screenshot shows a web browser window titled "Local Host Fingerprints" displaying a table of host data. The table has columns for host identifiers and various fingerprinting categories. Below the main table is a summary table showing the total count for each operating system (OS).

Host	POP	SMTP	Other
host018-156			X
host074-156			
host082-156	mac3aqz07tech04 [POP] mac3aqz07sbel08 [POP] mac3aqz07src10 [POP] mac3aqz07mmar07 [POP] mac3aqz07mmro08 [POP] mac3aqz07mac304 [POP] mac3aqz07rsan11 [POP] mac3aqz07lore13 [POP] mac3aqz07mmur08 [POP] mac3aqz07acas10 [POP]		
freebsd.computerhouseprato.com			
host013-154			X
host019-154			X
host018-154		a.bucciarelli@jumbooffice.it [SMTP]	
host017-154			X
host059-160			X

OS	Total
Windows 2000	122
Windows 9x	28
Linux 2.4.xx	18
Windows 2000 Server	15
FreeBSD 4.7	12
Windows 2000 Server SP4	7
Windows 2000 Pro / XP Pro / 2003 Server	4
Windows 2000 Advanced Server	4
Windows 95	4
Windows 98 SE	4
Windows XP Pro	4
FreeBSD 4.4 / 4.5 / 4.7	4

Based on <http://ettercap.sourceforge.net/>



































Host Health



Data Rcvd Stats	0 %		Rem 100 %
IP vs. Non-IP Rcvd	IP 100 %		Non-IP 0 %
Sent vs. Rcvd Pkts	Sent 51.8 %		Rcvd 48.2 %
Sent vs. Rcvd Data	Sent 33.2 %		Rcvd 66.8 %
Host Type	Name Server		
Historical Data			
Host Healthness (Risk Flags)	1. Unexpected packets (e.g. traffic to closed port or connection reset): [Rcvd: rejected] [Rcvd: port unreact] [Rcvd: hostnet unreact]		

Host Traffic Stats

Time	Tot. Traffic Sent	% Traffic Sent	Tot. Traffic Rcvd	% Traffic Rcvd
11 AM	13.4 MB	74.7 %	26.6 MB	74.0 %
10 AM	4.5 MB	25.3 %	9.3 MB	26.0 %
9 AM	0	0.0 %	0	0.0 %
8 AM	0	0.0 %	0	0.0 %

VoIP Support

Client	Server	Data Sent	Data Rcvd	Note
130.192.225.34    :8000	130.192.225.44    :32854	58.6 KB	70.3 KB	valter called livio
130.192.225.34    :8001	130.192.225.44    :32855	224	146	
stun01.sipphone.com  :3478	130.192.225.34    :47575	216	0	
130.192.225.34    :5060	bill.ipv6.polito.it    :5060	2.8 KB	2.3 KB	valter called livio
130.192.225.44    :5060	bill.ipv6.polito.it    :5060	4.5 KB	5.0 KB	valter called livio
130.192.225.44    :5060	130.192.225.34    :5060	462	361	

Host Type	VoIP Host 
Known Users 	stefano <101> [VoIP]

Integrating ntop into your network

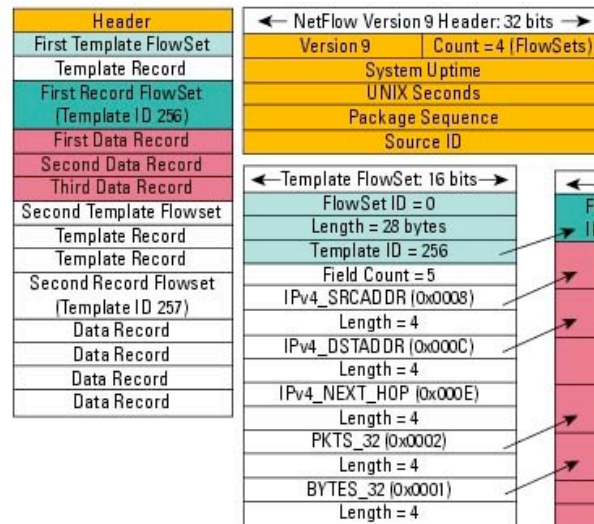
- You can use ntop with as a stand-alone application (via web) or as a traffic measurement server.
- Ntop can export traffic data in several ways:
 - Via the embedded SNMP agent (ntop MIB)
 - XML
 - RRD files
 - PHP/Perl data export
- Ntop, by means of the rrd-alarm companion application, allow users to emit alarms based on some traffic conditions.

Introduction to Cisco NetFlow

- What is NetFlow? A Cisco-proprietary IP statistics collection feature that collects information on IP flows passing through a router.
- NetFlow Version 9 is a flexible and extensible means to carry NetFlow records from a network node to a collector.

Options Template										
0	1	2	3	4	5	6	7	8	9	10
0	1	2	3	4	5	6	7	8	9	10
FlowSet ID = 1										
Length = 22										
Reserved Template ID = 275										
Option Scope Length = 4 byte										
Option Length = 8 bytes										
Type = 0x0002 (Interface)										
Length = 2 bytes										
Type = 0x0022 (34 decimal) Sampling Interval										
Length = 2 bytes										
Type = 0x0024 (36 decimal) Sampling Algorithm										
Length = 1 byte										
Padding										

Options Data Record										
0	1	2	3	4	5	6	7	8	9	10
0	1	2	3	4	5	6	7	8	9	10
Template ID = 275										
Length = 9 bytes										
Interface Index 2 (Ethernet 1)										
100 (Sampling Interval)										
0x01 Sampling ID Padding										



← Data FlowSet: 32 bits →	
FlowSet ID = 256	Length = 64 bytes
192.168.1.12	
10.5.12.254	
192.168.1.1	
5009	
5344385	
192.168.1.27	
10.5.12.23	
192.168.1.1	
748	
388934	
192.168.1.56	
10.5.12.65	
192.168.1.1	
5	
6534	

Introduction to InMon sFlow

- Ntop is part of the sflow.or consortium.
- Similar to NetFlow: probes send traffic flows to collectors over UDP in sFlow format (RFC 3176).
- A sFlow probe is basically a sniffer that captures packets at X rate (1:400 is default) and sends them coded in sFlow format. The more flows are captured, the more precise are the statistics. Tuning sample rate allows probes to capture at Gb speeds and above.
- sFlow in a nutshell:
 - Embedded in every switch port
 - Monitors traffic flow for all network ports
 - Effective at gigabit speeds
 - Does not impact network performance
 - Continuous monitoring
 - Robust under all network conditions
 - All devices = L2 – L7 flows end-end
 - Real-time and historical, detailed data

Ntop and NetFlow/sFlow


- Ntop supports both NetFlow (v1/5/7/9) and sFlow (v2/5).
- Ntop collects flows on virtual interfaces user-defined.
- Multiple interfaces can be defined independently. Ntop can simultaneously monitor netflow and sflow and pcap in interfaces.
- All the various interfaces have the same look and feel with little differences mainly due to the lack of payload access (NetFlow) hence inability to support packet decode (e.g. for P2P detection).

NetFlow Monitoring: State of the Art

- Cisco NetFlow is a commercial standard for network monitoring and accounting
- Many companies (e.g. Cisco, Juniper, Extreme) ship appliances with embedded NetFlow probes.
- Most commercial probes perform very poorly (~7-10'000 pkt/sec)
- Several collectors available (both commercial and Open Source).
- Very little offering in the probe side.
- NetFlow monitoring cannot cope with Gbit speeds and above hence new mechanisms (e.g. sampled NetFlow) have been used to overcome this problem.

Solution: nProbe+nTop

- The community needed an open source probe able to bring NetFlow both into small and large networks.
- Ability to run at wire speed (at least until 1 Gb) with no need to sample traffic.
- Complete open source solution for both flow generation (nProbe) and collection (nTop)



nProbe: Main Features

- Ability to keep up with Gbit speeds on Ethernet networks handling thousand of packets per second without packet sampling on commodity hardware.
- Support for major OS including Unix, Windows and MacOS X.
- Resource (both CPU and memory) savvy, efficient, designed for environments with limited resources.
- Source code available under GNU GPL (v3) and BSD (v4).
- nProbe v4 (available by the end of spring) new features:
 - Full NetFlow v9 support
 - V9 extensions: payload, network/application latency, RTP.
 - Ability to extend the probe with user-written plugins.

Packet Capture: Open Issues

- Monitoring low speed (100 Mbit) networks is already possible using commodity hardware and tools based on libpcap.
- Sometimes even at 100 Mbit there is some (severe) packet loss: we have to shift from thinking in term of speed to number of packets/second that can be captured analyzed.
- Problem statement: monitor high speed (1 Gbit and above) networks with common PCs (64 bit/66 Mhz PCI/X/Express bus) without the need to purchase custom capture cards or measurement boxes.
- Challenge: how to improve packet capture performance without having to buy dedicated/costly network cards?

Libpcap Performance on a Vanilla OS

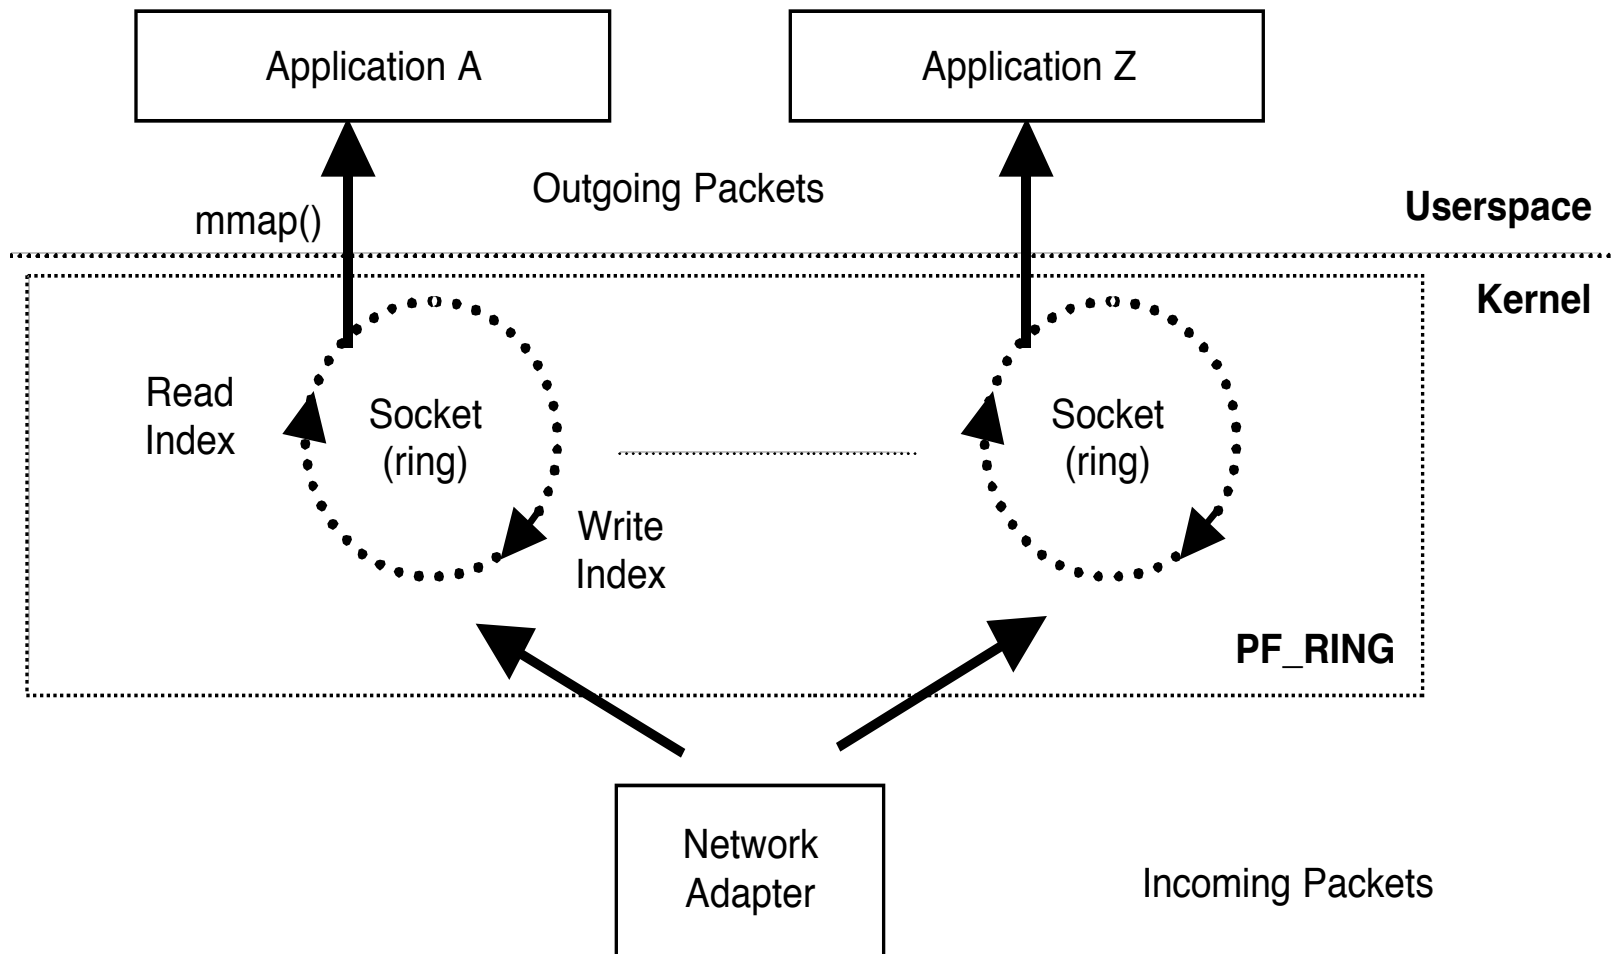
Traffic Capture Application	Linux 2.4.x	FreeBSD 4.8	Windows 2K
Standard Libpcap	0.2 %	34 %	68 %
mmap Libpcap	1 %		
Kernel module	4 %		

Percentage of captured packets [~80K packet/sec, ~45 Mbit]

Testbed:

- Sender: Dual 1.8 GHz Athlon, 3Com 3c59x Ethernet card
- Collector: VIA C3 533 MHz, Intel 100Mbit Ethernet card
- Network Switch: Cisco Catalyst 3548 XL
- Traffic Generator: tcpreplay (<http://tcpreplay.sourceforge.net/>)

Proposed Solution: Socket Packet Ring (PF_RING)



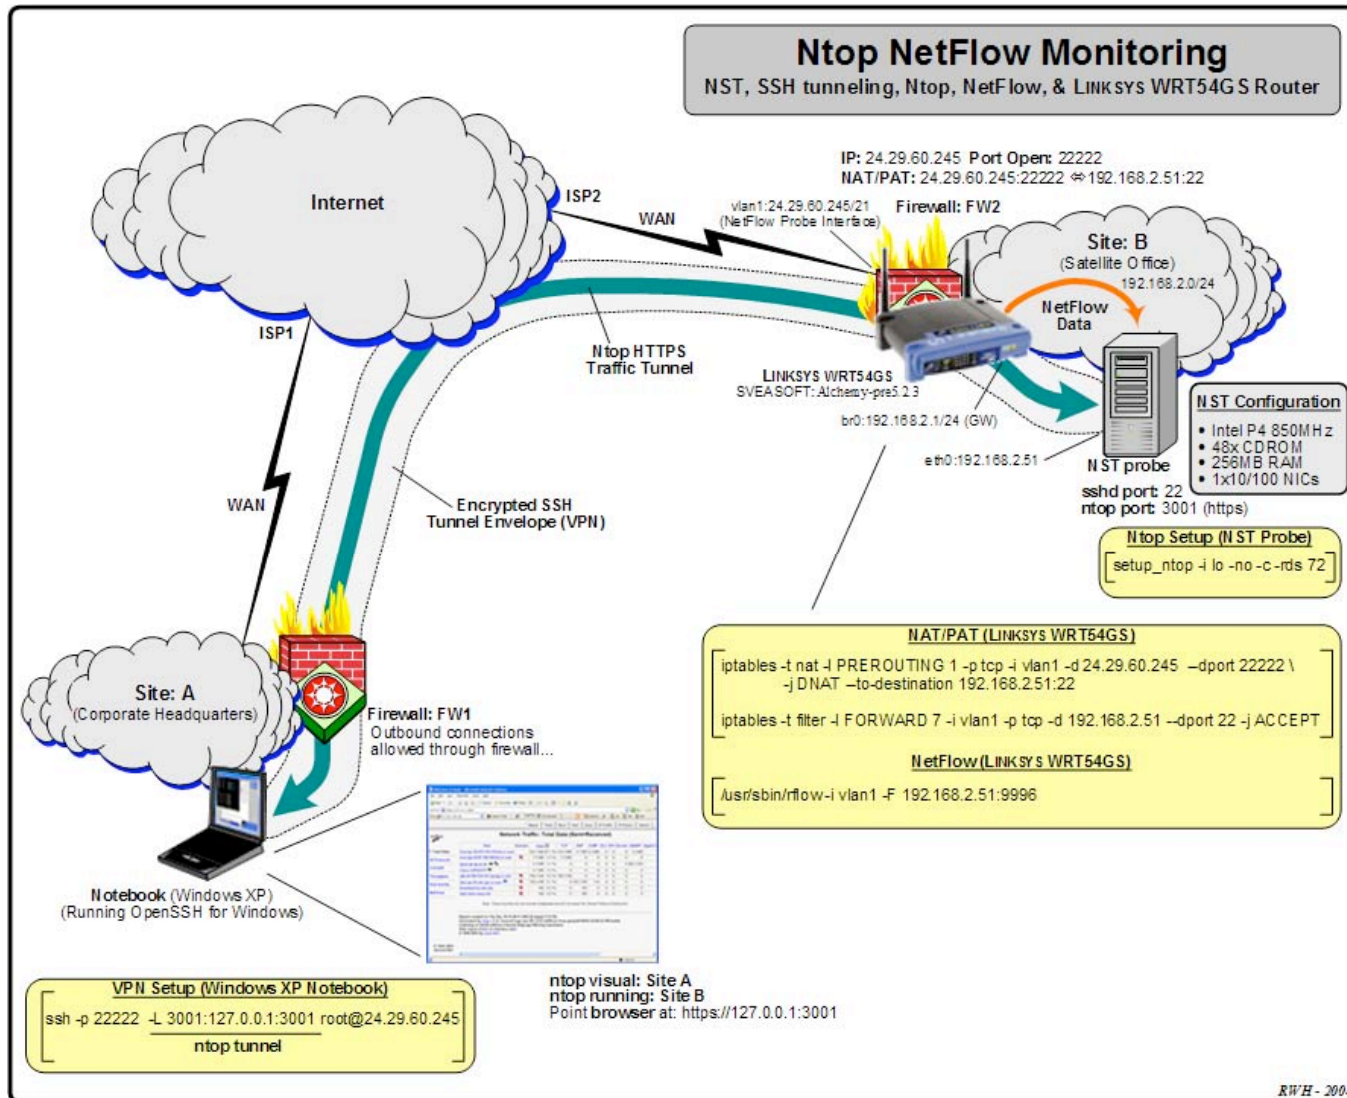
Socket Packet Ring: Packet Capture Evaluation [1/2]

Packet Size (Bytes)	Speed (Mbit)	Speed (Pkt/sec)	Linux 2.6.1 with NAPI and standard libpcap	Linux 2.6.1 with NAPI and mmap()	Linux 2.6.1 with NAPI and Ring	FreeBSD 4.8 with Polling
64	90	175'000	2.5%	14.9%	75.7%	97.3%
512	710	131'000	1.1%	11.7%	47%	47.3%
1500	836	70'000	34.3%	93.5%	92.9%	56.1%

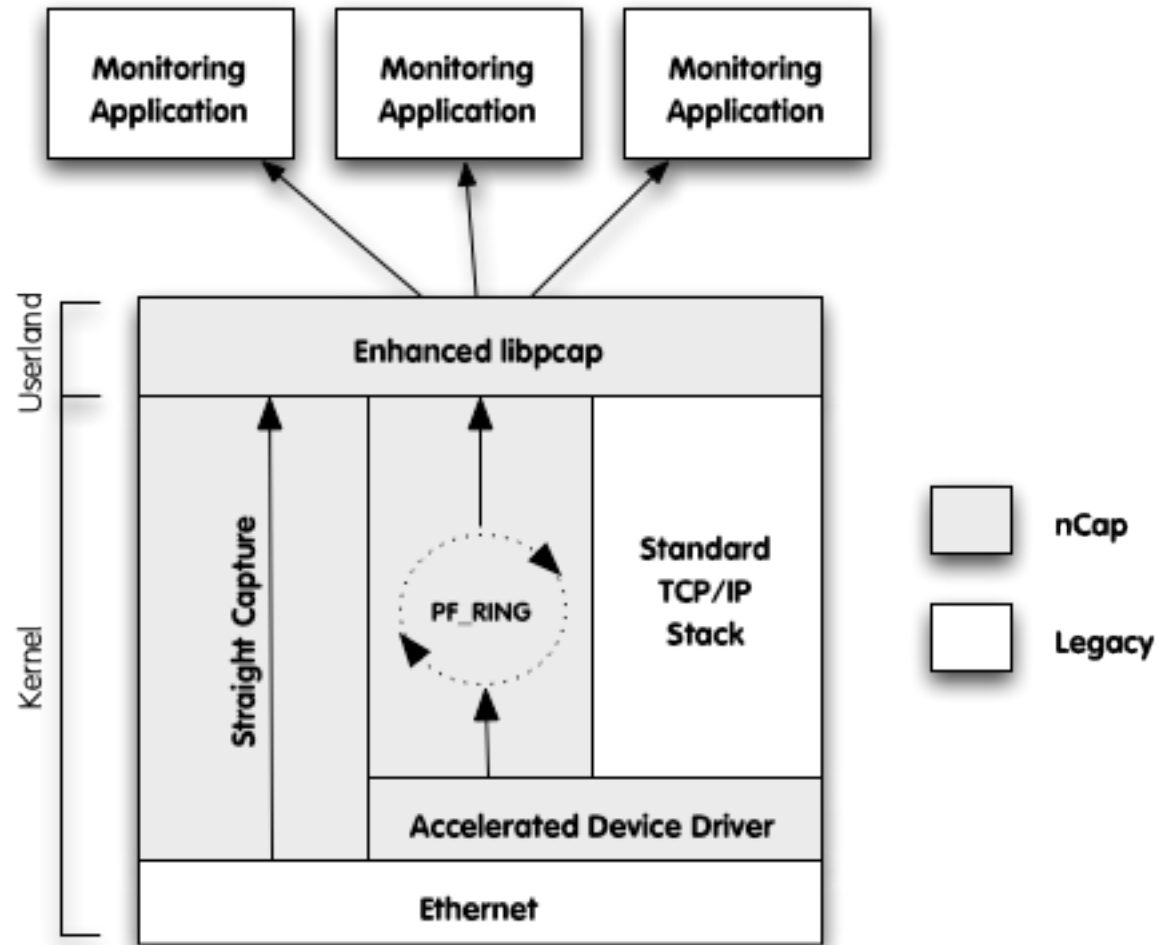
Testbed:

- Sender: Dual 1.8 GHz Athlon, Intel GE 32-bit Ethernet card
- Collector: Pentium 3 550 MHz, Intel GE 32-bit Ethernet card
- Traffic Generator: stream.c (DoS)

PF_RING on Embedded Devices



Welcome to nCap



nCap Features

	Packet Capture Acceleration	Wire Speed Packet Capture	Number of Applications per Adapter
Standard TCP/IP Stack with accelerated driver	Limited	No	Unlimited
PF_RING with accelerated driver	Great	Almost	Unlimited
Straight Capture	Extreme	Yes	One

Further nCap Features

- High speed packet capture: with a P4 HT (3 GHz) you can capture packets at wire speed (1.4 Mpps)
- High-speed traffic generation: cheap trafgen as fast as a hardware trafgen (>> 25'000 Euro).
- Precise packet generation
- Precise packet timestamping (no kernel interaction)
- Enhanced driver currently supports Intel cards (1 and 10 Gb Ethernet).
- Availability (live CD): <http://luca.ntop.org/nCap/>

Conclusions

Over the past 7 years the ntop project has produced:

- Ntop: a mature passive traffic monitoring application able to be integrated into industrial environments.
- nProbe: a fast and extensible NetFlow probe able to use ntop as a central console and to measure traffic using NetFlow even on networks where there aren't NetFlow-enabled routers.
- PF_RING: Linux packet capture acceleration able to run on embedded systems and high-speed SMP servers.
- nCap: wire-speed packet capture and transmission for 1 and 10 Gbit networks.