


Open Source in Network Administration: the ntop Project

Luca Deri <deri@ntop.org>

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 and Linux distro's.
- 2002-03: Version 2.x, added support for commercial protocols (NetFlow v5 and sFlow v2).
- 2004-05: Version 3.x, added RRD support, IPv6 (Loria) and SCSI/FibreChannel (Cisco) support, NetFlow V9/IPFIX (draft), sFlow v5, VoIP.
- 2006-08: ntop consolidation, PF_RING 3.x, n2n 1.x

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 address bar shows 'http://localhost:3000/hostsinfo.html'. The ntop logo is in the top left, and the copyright notice '(C) 1998-2008 - Luca Deri' is in the top right. Below the logo is a search bar and navigation links: 'About', 'Summary', 'All Protocols', 'IP', 'Utils', 'Plugins', 'Admin'. The main heading is 'Host Information'. Below this, there are two dropdown menus: 'Traffic Unit: Bytes' and 'Subnet: All'. The main content is a table with columns: Host, Domain, IP Address, MAC Address, Community, Other Name(s), and Bandwidth. The table lists various hosts with their respective IP addresses and bandwidth usage. Below the table is a 'NOTE' section with bullet points explaining the data. At the bottom, it says 'Report created on Sun May 11 15:47:55 2008 [ntop uptime: 16 sec]'.

Host	Domain	IP Address	MAC Address	Community	Other Name(s)	Bandwidth
192.168.1.81		192.168.1.81				
mi.mirror.garr.it		193.206.139.34				
151.1.245.36		151.1.245.36				
192.168.1.1		192.168.1.1				
jake.unipi.it		131.114.21.22				
192.168.160.11		192.168.160.11	00:16:CB:96:BA:BE			
151.11.185.69		151.11.185.69				
151.11.185.65		151.11.185.65				
192.168.1.1		192.168.1.1	00:1C:A2:37:21:F7			
all-systems.mcast.net		224.0.0.1				
224.0.0.251		224.0.0.251				
83.103.35.4		83.103.35.4				

NOTE:

- You can [define](#) new communities.
- Click [here](#) for more information about host and domain sorting.
- Bandwidth values are the percentage of the total bytes that ntop has seen on the interface. Hover the mouse to see the actual value (rounded to the nearest full percentage point). *The total of the values will NOT be 100% as local traffic will be counted TWICE (once as sent and again as received).*
- The SENT bandwidth is shown as and the RECEIVED bandwidth is shown as

Report created on Sun May 11 15:47:55 2008 [ntop uptime: 16 sec]

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).

Who is using ntop products?



QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

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.

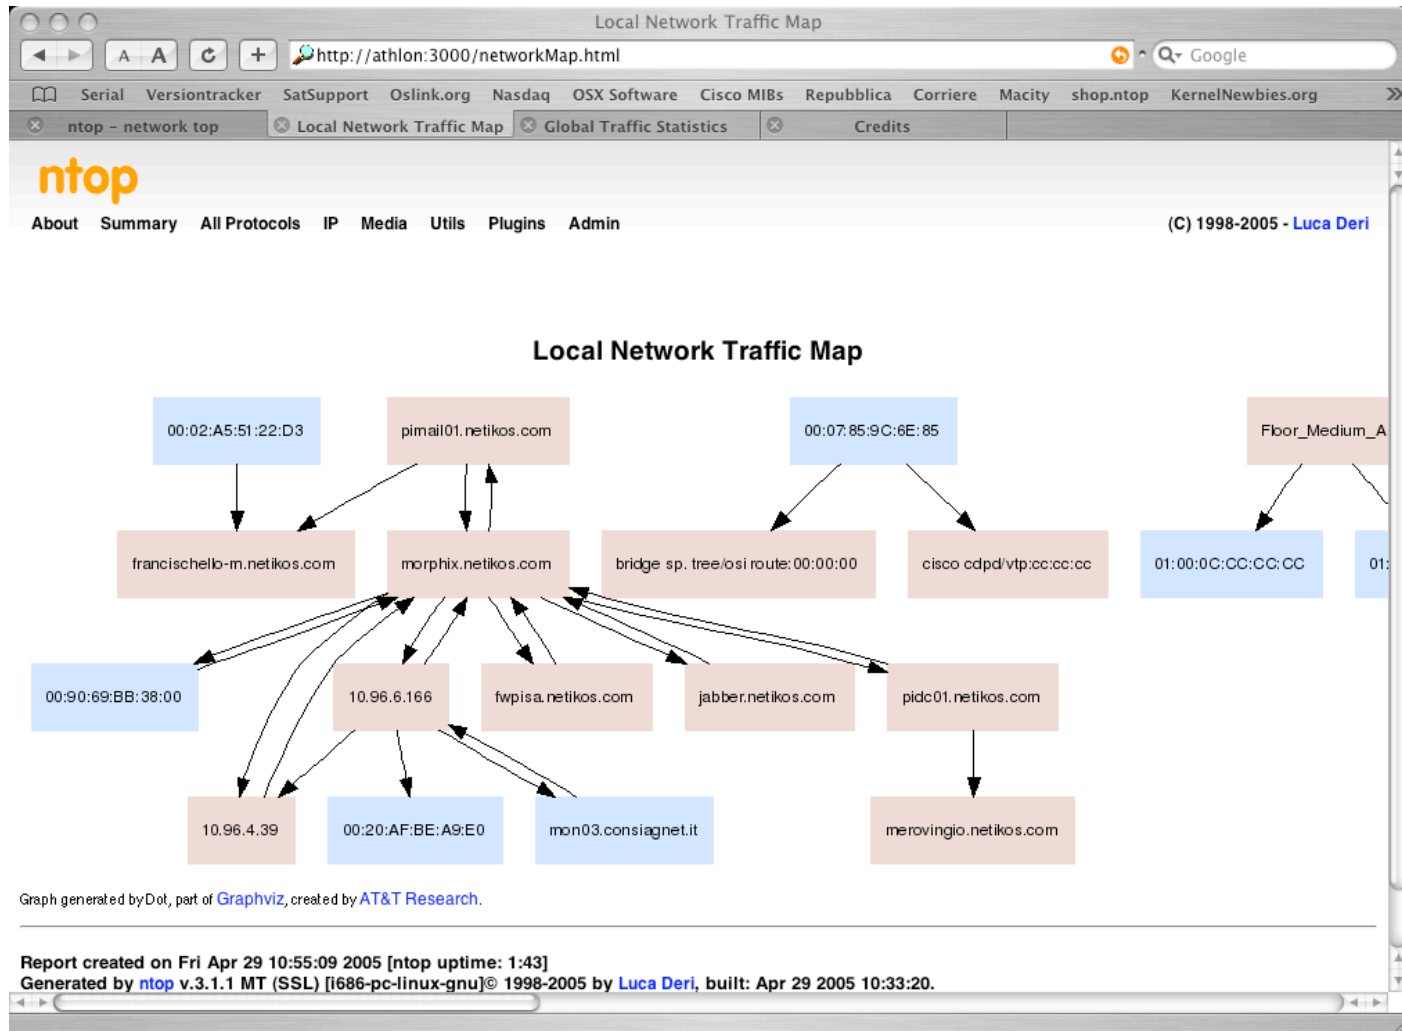
Traffic Characterization and Monitoring

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

Network Optimization 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, 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' with the URL <http://mon03.consiagnet.it/localHostsFingerprint.html>. The browser has several tabs open, including 'ntop - network top', 'Hosts Characterization', 'Local Host Fingerprints', and 'Credits'. The main content is a table with columns for host identifiers and various fingerprinting metrics. Below the table is a summary table showing the total count for each operating system (OS).

Host	OS	Other	Total
host018-156		X	
host074-156			
host082-156			
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



































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 unrec] [Rcvd: hostnet unrec]	

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/Python/JSON 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	11	12	13	14	15
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	11	12	13	14	15
Template ID = 275															
Length = 9 bytes															
Interface Index 2 (Ethernet 1)															
100 (Sampling Interval)															
0x01 Sampling ID Padding															

Header	
First Template FlowSet	Template Record
First Record FlowSet (Template ID 256)	First Data Record
Second Data Record	Third Data Record
Second Template Flowset	Template Record
Template Record	Template Record
Second Record Flowset (Template ID 257)	Data Record
Data Record	Data Record
Data Record	Data Record

← NetFlow Version 9 Header: 32 bits →	
Version 9	Count = 4 (FlowSets)
System Uptime	
UNIX Seconds	
Package Sequence	
Source ID	

← Template FlowSet: 16 bits →	
FlowSet ID = 0	Length = 28 bytes
Template ID = 256	Field Count = 5
IPv4_SRCADDR (0x0008)	Length = 4
IPv4_DSTADDR (0x000C)	Length = 4
IPv4_NEXT_HOP (0x000E)	Length = 4
PKTS_32 (0x0002)	Length = 4
BYTES_32 (0x0001)	Length = 4

← 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

sFlow

Ntop and NetFlow/sFlow


- Ntop supports both NetFlow (v1/5/7/9)/IPFIX 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)



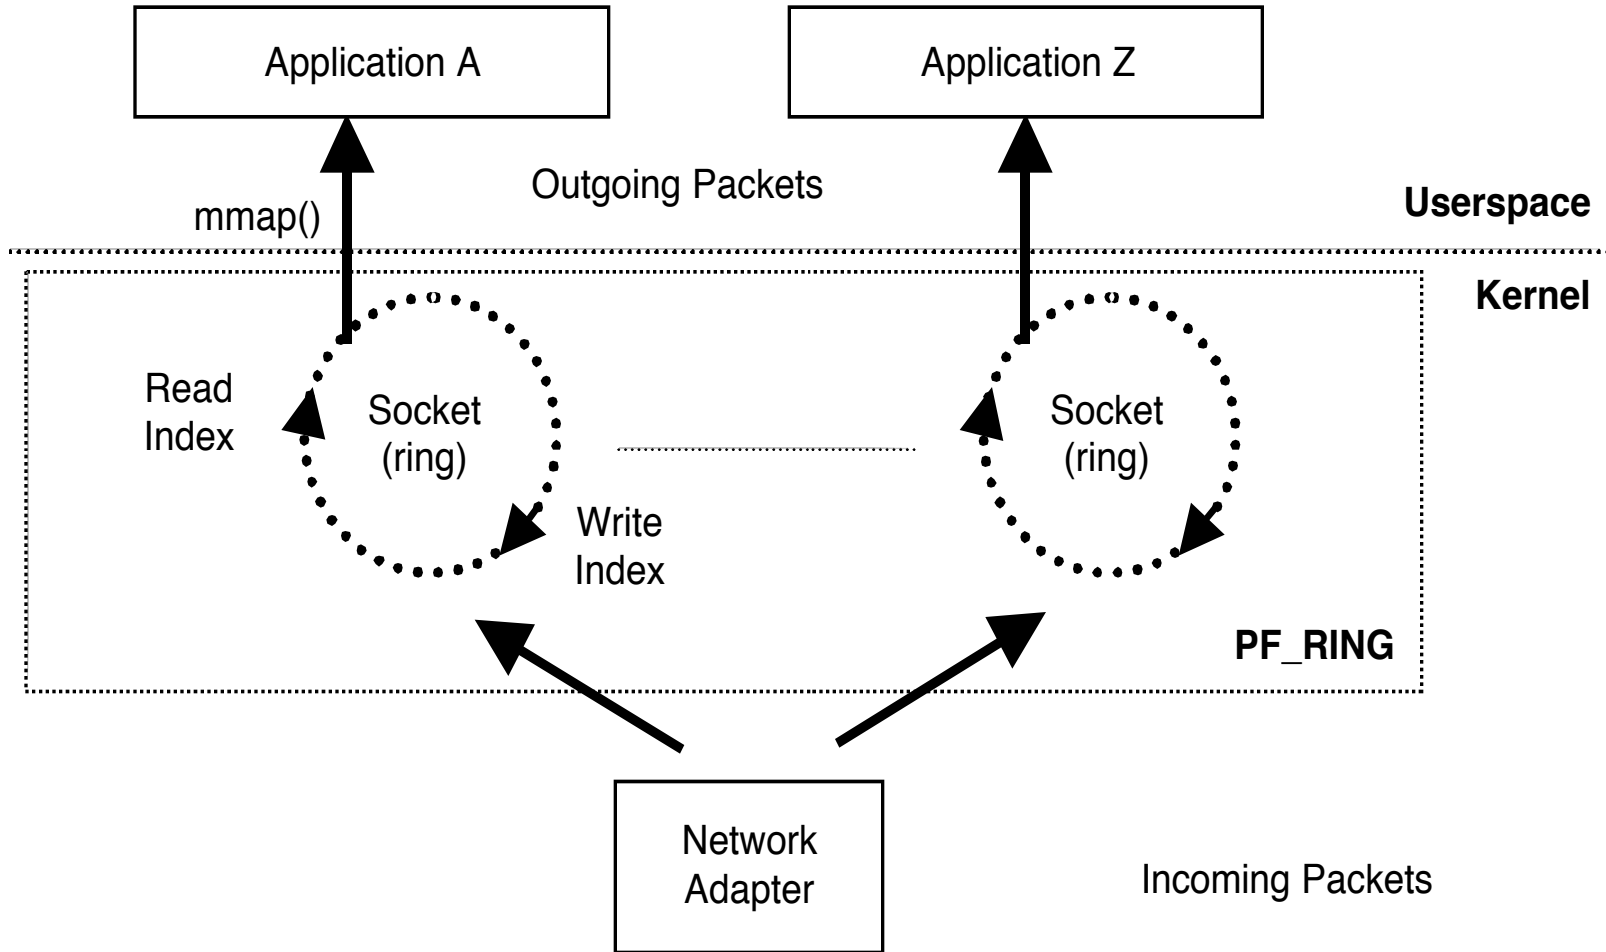
nProbe: Main Features

- 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.
- nProbe v4 new features:
 - Full NetFlow v9/IPFIX support
 - V9 extensions: payload, network/application latency, SIP/RTP.
 - Ability to extend the probe with user-written plugins.
- nProbe v5 will be released later this summer.

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?

Packet Filter Ring (PF_RING)



PF_RING: Benefits

- It creates a straight path for incoming packets in order to make them first-class citizens.
- No need to use custom network cards: any card is supported.
- Transparent to applications: legacy applications need to be recompiled in order to use it.
- No kernel or low-level programming is required.
- Developers familiar with network applications can immediately take advantage of it without having to learn new APIs.

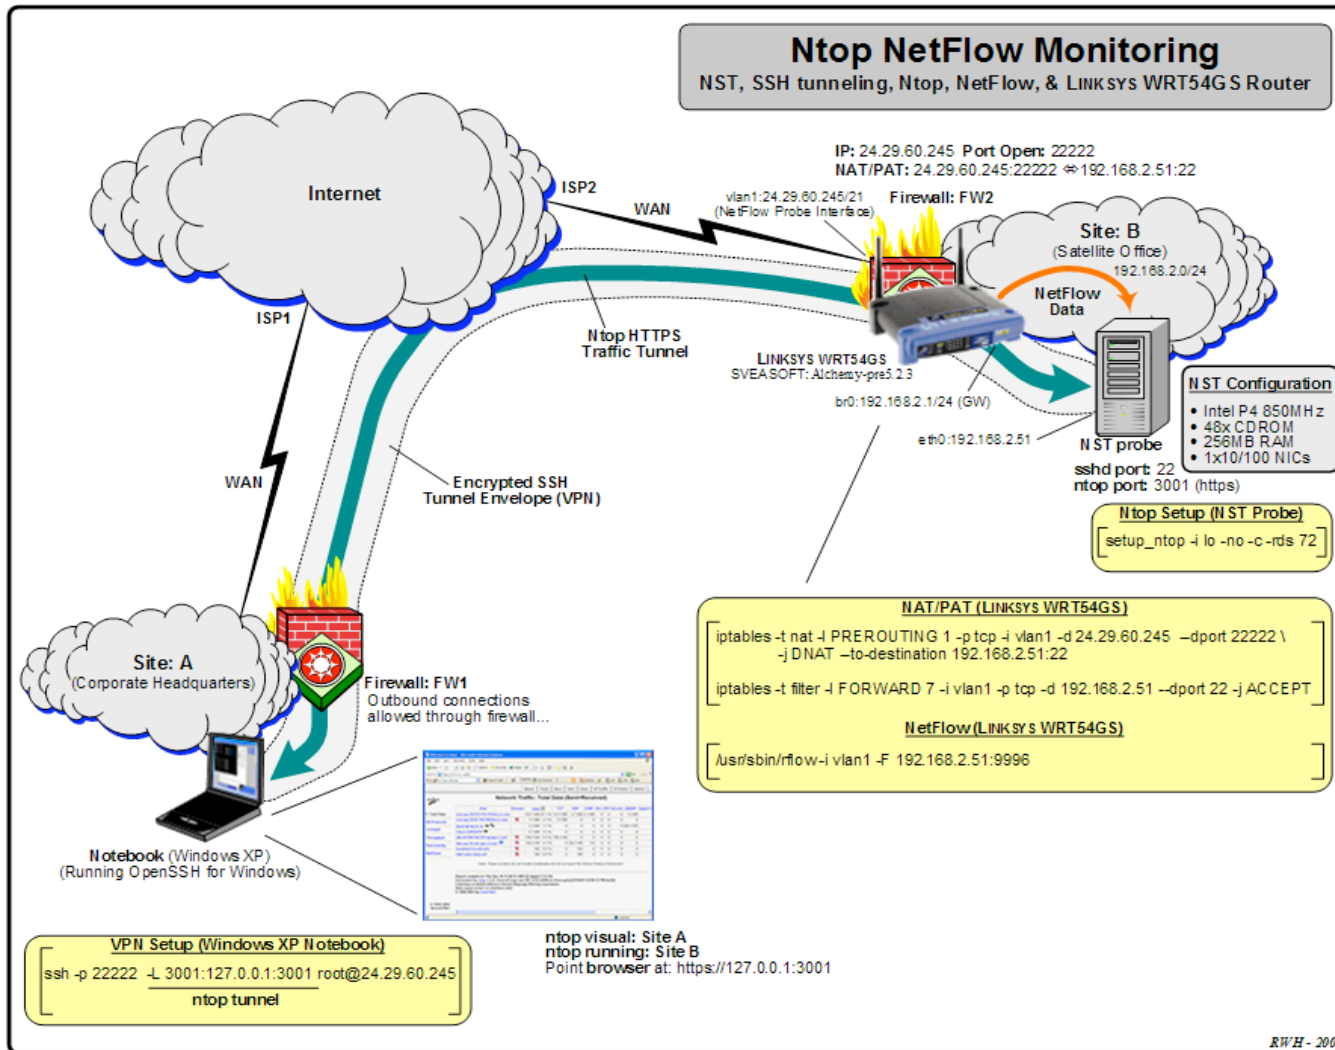
PF_RING: Performance Evaluation

Pkt Size	Kpps	Mpps	% CPU Idle	Wire-Speed
250	259.23	518	> 90%	Yes
250	462.9	925.9	88%	Yes
128	355.1	363.6	86%	Yes
128	844.6	864.8	82%	Yes

Test setup: pcount, full packet size, 3.2 GHz Celeron (single-core) - IXIA 400 Traffic Generator

PF_RING on Embedded Devices

<http://nst.sourceforge.net/nst/docs/user/ch09s02.html>



n2n: Private Overlay for Nw Administration

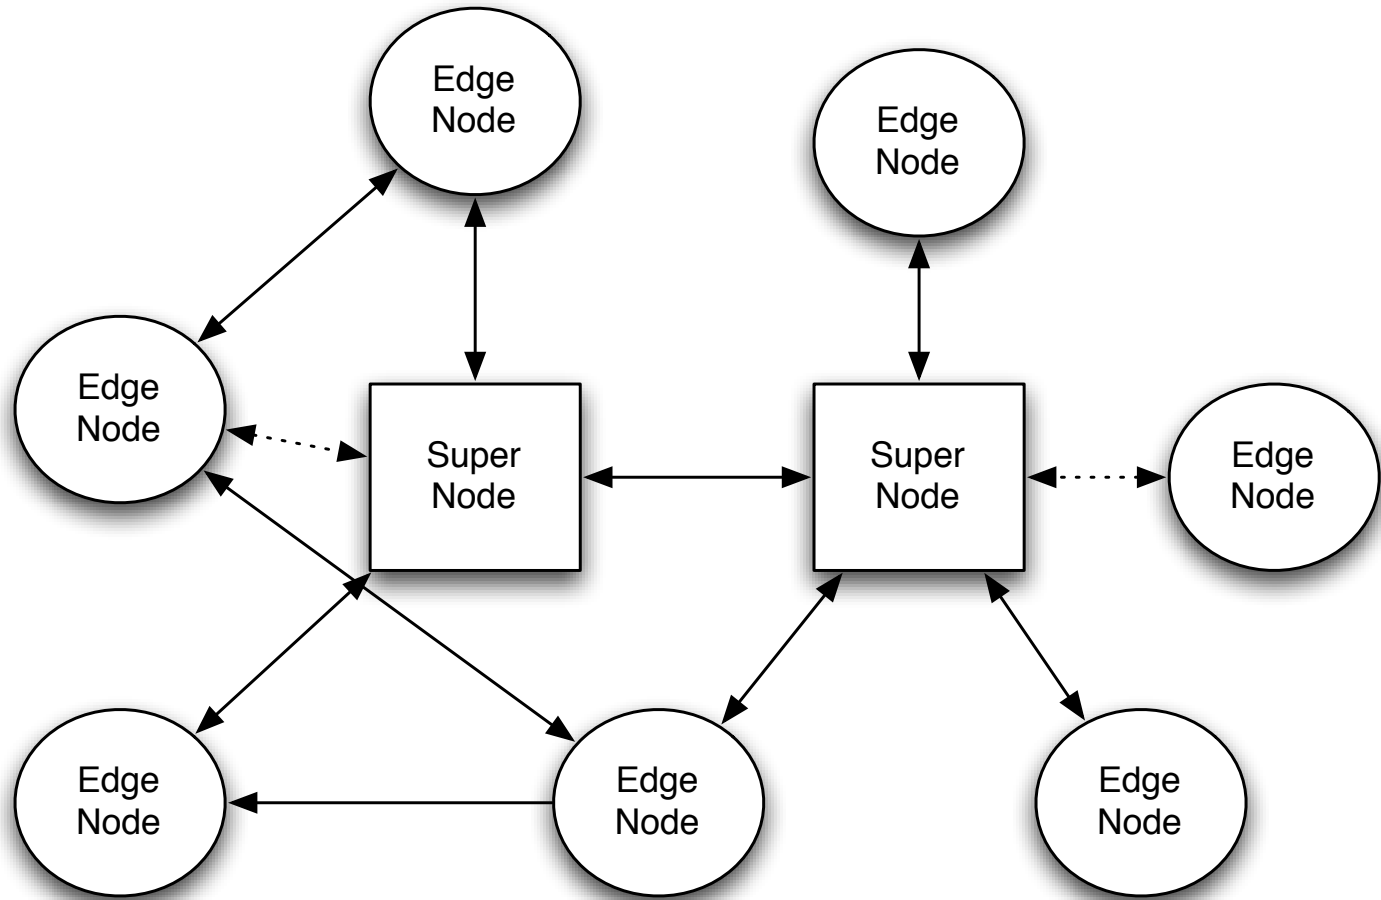
Motivation

- NAT devices mask the user's IP identity and limit peers accessibility.
- No control over the connection configuration, totally managed by ISPs.
- Firewall greatly reduce the possibility of a user being contacted by a direct session opened elsewhere over the Internet.

Vision

- The internet should be a "transparent" IP-based transport for users, not a geographical/ISP constrain.
- Users should control/create their community networks (today network administrators do).

N2N Architecture



N2N Features

- A n2n network is an encrypted L2 P2P-VPN.
- Unlike Skype/Hamachi, encryption is performed on edge nodes using open protocols with user-defined encryption keys.
- n2n users can simultaneously belong to multiple networks.
- Ability to cross NAT and firewalls in the reverse traffic direction (i.e. from outside to inside) so that n2n nodes are reachable even if running on a private network.
- n2n networks are not meant to be self-contained, but it is possible to route traffic across n2n and non-n2n networks.

Conclusions

Over the past 10 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.
- n2n: layer 2, peer-to-peer VPN for remote system connectivity and administration.