



Evangelos Markatos  
FORTH-ICS  
[markatos@ics.forth.gr](mailto:markatos@ics.forth.gr)

<http://www.ics.forth.gr/~markatos>  
Institute of Computer Science (ICS)  
Foundation for Research and Technology – Hellas (FORTH)

- Motivation
  - Why is network monitoring important?
- Two methodologies for monitoring
  - **Active** network monitoring
    - Examples
  - **Passive** network monitoring
    - Examples
- Introduce the rest of this tutorial
- Summary and Conclusions



# Motivation: Why do we need network traffic monitoring?



*“...for the most part we really have no idea what’s on the network...”*

We ...

*“can’t measure topology in either direction at any layer*

*can’t get precise one-way delay*

*can’t get an hour of packets from the core*

*can’t get accurate bandwidth/capacity information*

*can’t get anything from the core with real addresses in it”*

*...for the most part **we really have no idea what’s on the network...***





K Claffy, CAIDA  
2004

- So, do Network Traffic Monitoring
  - To get a better understanding on “what’s on the network”
- What can you do with monitoring?
  - Know the **state** of the Internet and your network
  - **Capacity** planning
  - **Traffic accounting**
    - Which application generates most traffic?
  - Understand the **performance** of individual applications
    - “why is my application so sloooooooow”?
  - Detect Security threats
    - DoS Attacks, Worm outbreaks



- Inject packets into the network
- Measure their arrival time, loss rate, etc.
- What can you do with it?
  - Measure **delay** (one-way / two-way)
  - Measure **bottleneck link bandwidth**
  - Find **network topology**
  - **What's up?** (in my network)
    - Which nodes are up and running?



```
%> ping www.ist-lobster.org
```

```
PING www.ist-lobster.org (192.87.30.11): 56 data bytes
```

```
64 bytes from 192.87.30.11: icmp_seq=0 ttl=49 time=308.7 ms
```

```
64 bytes from 192.87.30.11: icmp_seq=1 ttl=49 time=307.6 ms
```

```
64 bytes from 192.87.30.11: icmp_seq=2 ttl=49 time=244.4 ms
```

```
--- www.ist-lobster.org ping statistics ---
```

```
3 packets transmitted, 3 packets received, 0% packet loss
```

```
round-trip min/avg/max = 244.4/286.9/308.7 ms
```

```
%>
```

## Traceroute: find all intermediate routers between a source and a destination computer



```
%> traceroute www.ist-lobster.org (from Crete)
```

```
traceroute to www.ist-lobster.org (192.87.30.11), 30 hops max, 40 byte packets
```

```
1 147.52.17.1 (147.52.17.1) 1.050 ms 0.690 ms 0.592 ms
2 olympos-e43.lanh.uoc.gr (147.52.12.1) 1.626 ms 1.033 ms 0.840 ms
3 heraklio-uch-ATM.grnet.gr (194.177.209.141) 129.528 ms 112.644 ms 123.465 ms
4 heraklio2-to-heraklio.backbone.grnet.gr (194.177.209.77) 124.791 ms 116.749 ms 119.965
  ms
5 Syros-to-Heraklio2.backbone.grnet.gr (195.251.27.81) 136.089 ms 104.469 ms 81.116 ms
6 athens3-to-Syros.backbone.grnet.gr (195.251.27.10) 72.664 ms 62.814 ms 67.341 ms
7 grnet.gr1.gr.geant.net (62.40.103.57) 81.392 ms 102.067 ms 79.488 ms
8 gr.de2.de.geant.net (62.40.96.82) 129.641 ms 134.589 ms 144.765 ms
9 de2-2.de1.de.geant.net (62.40.96.54) 139.478 ms 158.336 ms 146.815 ms
10 de.nl1.nl.geant.net (62.40.96.102) 180.696 ms 162.904 ms 173.813 ms
11 surfnet-gw.nl1.nl.geant.net (62.40.103.98) 184.078 ms 158.921 ms 160.933 ms
12 PO11-0.CR1.Amsterdam1.surf.net (145.145.166.33) 145.367 ms 150.166 ms 142.117 ms
13 PO0-0.AR5.Amsterdam1.surf.net (145.145.162.2) 163.605 ms 144.161 ms 177.526 ms
14 145.145.18.46 (145.145.18.46) 178.350 ms 175.365 ms 166.334 ms
15 * * 145.145.18.46 (145.145.18.46) 176.079 ms !X
16 * 145.145.18.46 (145.145.18.46) 171.861 ms !X *
17 145.145.18.46 (145.145.18.46) 192.753 ms !X * 180.104 ms !X
```

VisualRoute 5.0b

File Edit Options Tools Help


Host: <http://www.duma.ru> IP Addresses: 212.11.128.31 Recent Hosts

**Report for www.duma.ru [212.11.128.31]**

Analysis: Node 'www.duma.ru' was found in 22 hops (TTL=107). It is a HTTP server (running Novell-HTTP-Server3.1R1).

Hop	% Loss	IP Address	Node Name	Location	Timezone	ms	Graph	Network
0		128.40.59.193	wolf.casa.ucl.ac.uk					221 University College London
1		128.40.59.245	cisco-2.bartucl.ac.uk			0		University College London (private use)
2		10.0.121.45	-			0		University College London
3		128.40.255.153	-			1		The London MAN
4		194.83.100.62	atmr-ulcc.imn.net.uk	London, UK	0.0	1		University of London Computer Centre
5		146.97.40.85	gl0-0-0-ext-gw6.ja.net			1		University of London Computer Centre
6		128.96.1.15	ten155-gw.ja.net			1		IP allocations for TEN-155 ATM PVC
7		212.1.192.149	janet.uk.ten-155.net	51.50 N, 0.11 W (United Kingdom)	0.0	9		IP allocations for TEN-155 PoP Equipment
8		212.1.193.154	ge.uk40.ten-155.net			0		212.1.197.57
9		212.1.197.57	nl-uk-1.nl40.ten-155.net			10		IP allocations for TEN-155 ATM PVC
10		212.1.192.102	nl-se.se.ten-155.net	59.33 N, 18.05 E (United Kingdom)	0.0	40		IP allocations for TEN-155 external peerings
11		212.1.194.26	stockholm5.se.equip.net			63		Stockholm Interconnect
12		194.68.128.25	stockholm-D00k.ebone.net	Stockholm, Sweden	+1.0	42		Ebone backbone 3
13		195.158.226.77	sesto502-ib-p0-0.ebone.net	Stockholm, Sweden	+1.0	46		Ebone backbone 3
14		195.158.226.54	-			47		Sovam Teleport
15	10	194.186.157.161	cisco0.Moscow.ST.NET	Moscow, Russia		81		Sovam Teleport
16	20	194.186.157.182	cisco02.Moscow.ST.NET	Moscow, Russia		70		Sovam Teleport
17	10	194.67.16.219	ccr-1.Moscow.ST.NET	Moscow, Russia		85		Sovam Teleport
18	10	194.186.0.198	tr-pop-gw.Moscow.ST.NET	Moscow, Russia		88		Glas-Internet Ltd
19	10	195.218.254.83	MOS-gw.glas.net	(Russia)		84		Moscow Mayor's Office
20	10	212.11.128.31	duma.ru	Moscow, Russia		80		Moscow Mayor's Office
21	10	212.11.128.31	duma.ru	Moscow, Russia		80		Moscow Mayor's Office
22		212.11.128.31	www.duma.ru	Moscow, Russia		97		Moscow Mayor's Office

VisualRoute Report for www.duma.ru produced at 18:55 on 24 November, 2000.  
Roundtrip time to www.duma.ru (212.11.128.31): average = 97ms min = 70ms max = 200ms




Nr	Hostname	IP number
10	sl-pp021-stk-1-2.sprintlink.net	144.252.4.76
11	sl-ucberkeley-1-1-0-T3.sprintlink.net	144.228.146.50
12	f5-0.inr-666-eva.berkeley.edu	198.128.16.21
13	f1-0-0.inr-107-eva.Berkeley.EDU	128.32.2.1
14	f8-0.inr-100-eva.Berkeley.EDU	128.32.235.100
15	amber.Berkeley.EDU	128.32.25.12

- What is it?
- Non-intrusive traffic monitoring
  - Much like a **telescope**
  - **Does not inject packets** in the network
- It **passively captures** information from passing packets such as
  - High-level network flows (CISCO Netflow)
  - Network packet headers (NLNR)
  - Entire network packets (incl. payload)
    - if allowed
    - maybe stripped/anonymized (to be shared with a broader audience)



- Traffic Categorization/Accounting:
  - What % of my traffic is due to email?
  - Which subnet generates most outgoing traffic?
- Bandwidth Estimation
  - What % of my bandwidth is available now?
  - What % of my bandwidth is being used?
- Study trends:
  - How does the application mix in the traffic changes with time?
    - ftp in the 80's, www in the 90's, p2p in the 00's
  - How does peer-to-peer traffic changes with time?
- Performance Debugging of individual applications
  - Why is **my** application so sloooow?




# Passive Monitoring: What can it be used for? Security




- **Intrusion Detection**
  - Are any of my computers compromised?
  - Is there any attacker trying to intrude into my network?
- **Large-scale Attack Detection** – Detection of Epidemics
  - DoS Attack detection
    - e.g. Detect sharp increases in TCP/SYN packets
  - Zero-day worm detection
    - e.g. Detect lots of identical packets, never seen before, from several sources to several destinations
    - e.g. Detect worm characteristics
      - such as NOP sleds: long sequences of executable code
- **Network Telescopes**
  - They monitor unused IP addresses (“dark matter”)
  - Ordinarily, unused IP addresses should not receive traffic
  - Observe victims of DoS attacks
    - “back-scatter” traffic
  - Observe infected hosts
  - Port scans



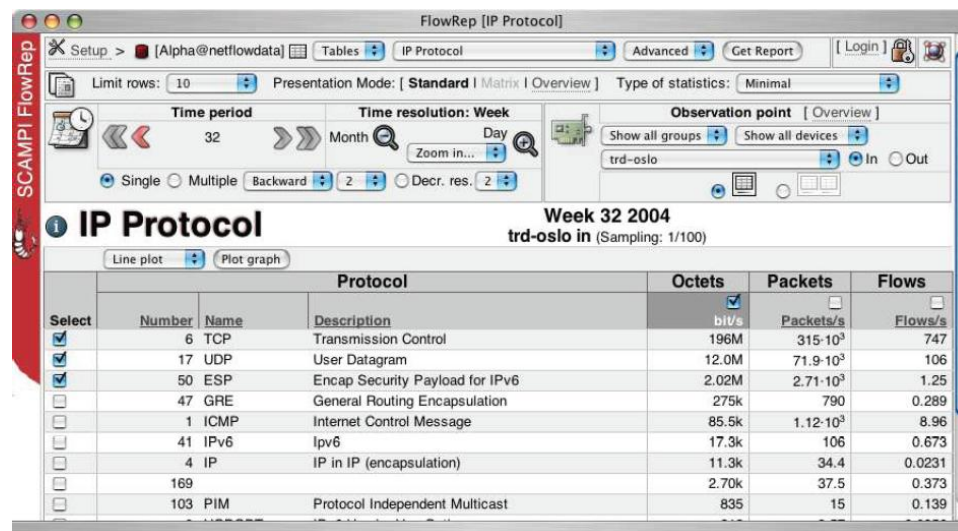
- Equipment varies from
  - low-end (for low-speed networks) to
  - Sophisticated equipment (for high-speed networks)
- Low-end passive monitors (100Mbps – 1Gbps)
  - An ordinary PC
  - An ordinary network Interface (i.e. an Ethernet card)
    - put in promiscuous mode
  - Mirror all packets from a router to a port connected to the above PC



- High-end passive monitors (1Gbps – 10Gbps)
  - High-end computer
  - Specialized network interface
    - DAG Cards (Endace)
    - Combo cards (SCAMPI project)
    - Hardware-based filtering capabilities
      - Process packets at line speeds



- pcap: packet capture library from Berkeley
- MAPI: Monitoring API (Application Programming Interface)
  - Developed within the IST SCAMPI project
    - co-funded by EU
- Net-flow-related tools
  - Graphical interfaces



- Current solutions for packet capture
  - pcap (packet capture library)
- New approaches for packet capture
  - MAPI (Monitoring API)
  - FFPF (Fairly Fast Packet Filters)
- Hardware Support for 10 Gbps links
  - Combo cards
  - DAG cards

- Flow-level passive monitoring
  - General introduction
    - NetFlow and IPFIX
  - NetFlow-based applications
    - Stager (UNINETT)
    - NERD (TNO)
- Summary, Conclusions, and Future Trends

*“...for the most part **we really have no idea what’s on the network...**”*

K Claffy

- Traffic Monitoring help us understand what’s on the network
- Passive Network Traffic Monitoring applications:
  - Performance
    - traffic accounting/categorization
    - Performance debugging
  - Security applications
    - DoS attack detection,
    - Internet epidemics
    - Intrusion Detection



Evangelos Markatos  
FORTH-ICS  
[markatos@ics.forth.gr](mailto:markatos@ics.forth.gr)

<http://www.ics.forth.gr/~markatos>  
Institute of Computer Science (ICS)  
Foundation for Research and Technology – Hellas (FORTH)