AutoFocus: A Tool for Automatic Traffic Analysis

Cristian Estan,
University of California, San Diego

Who is using my link?

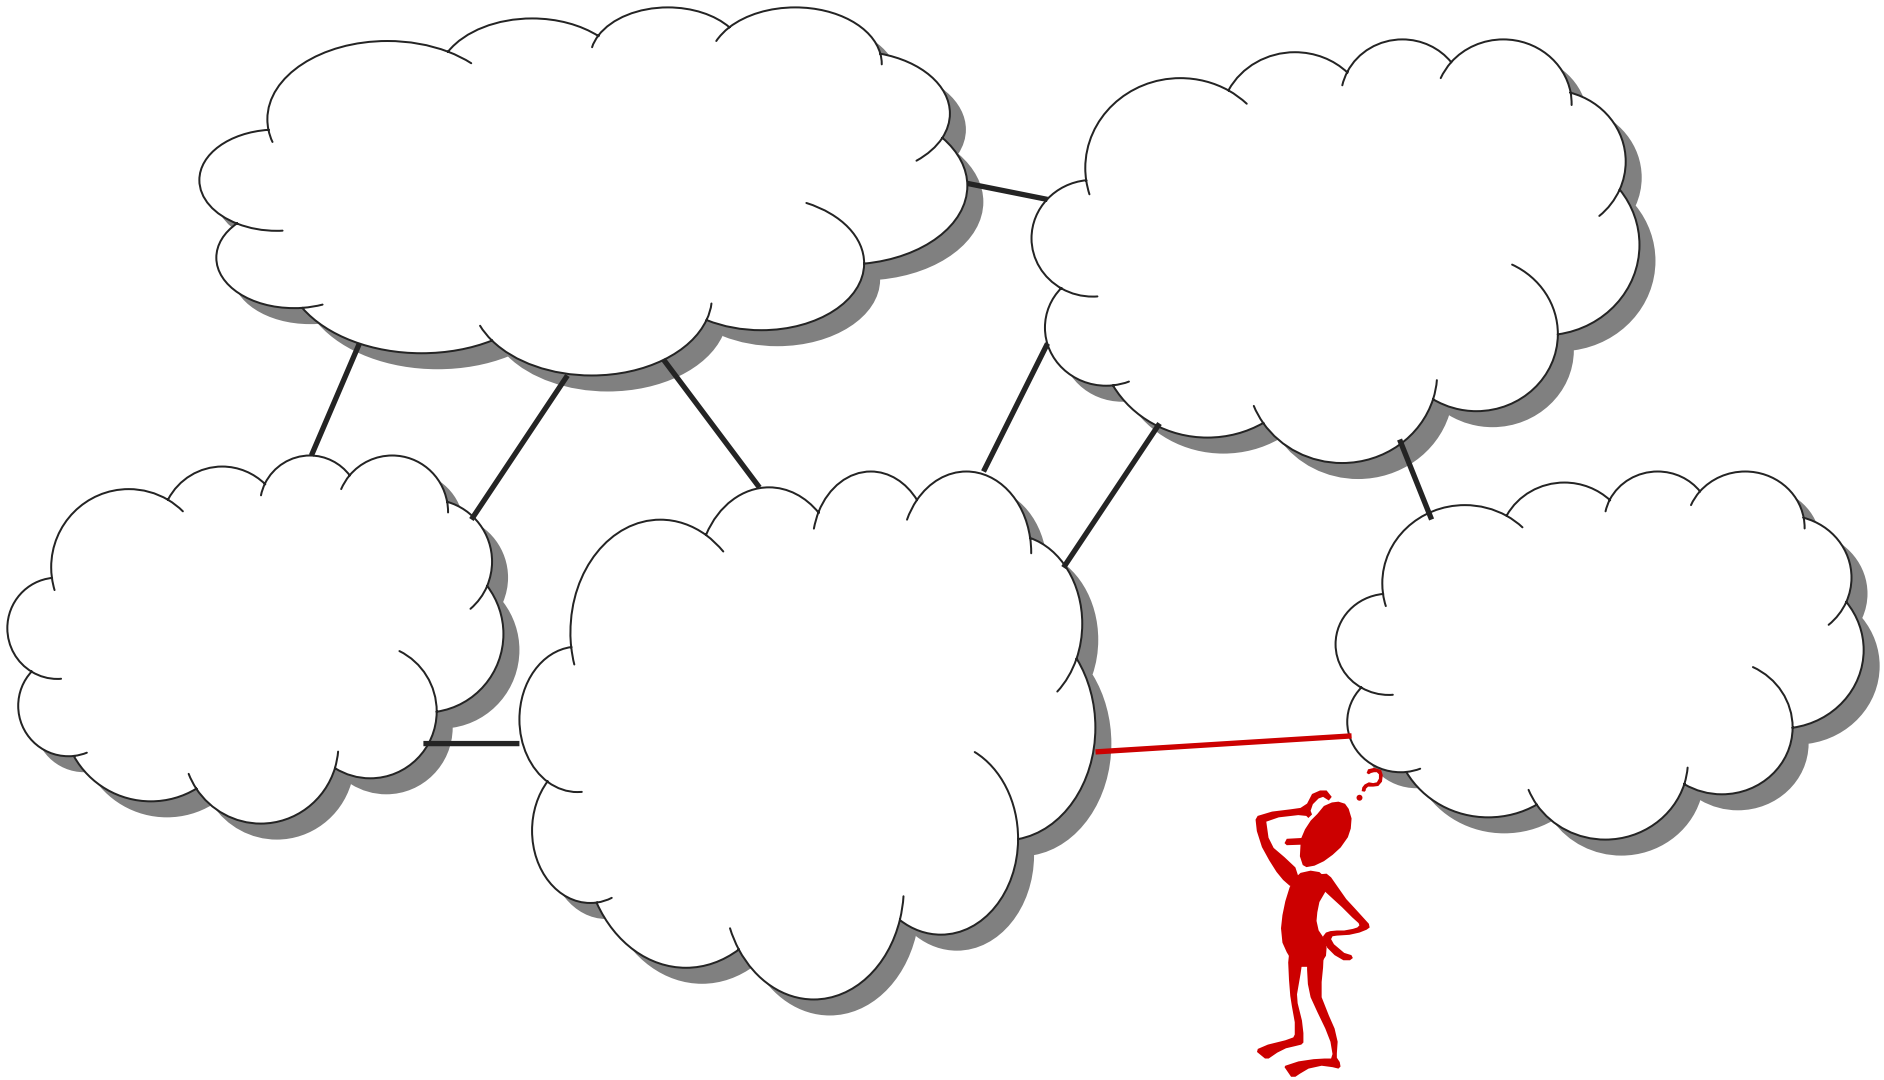

Informal problem definition

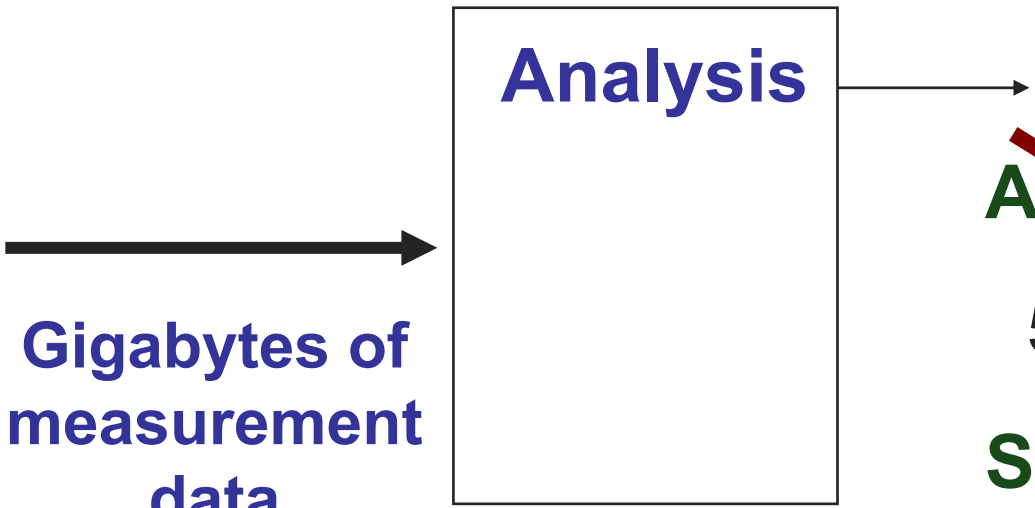

Traffic reports

Applications:

50% of traffic is Kazaa

Sources:

20% is from Steve's PC

Informal problem definition

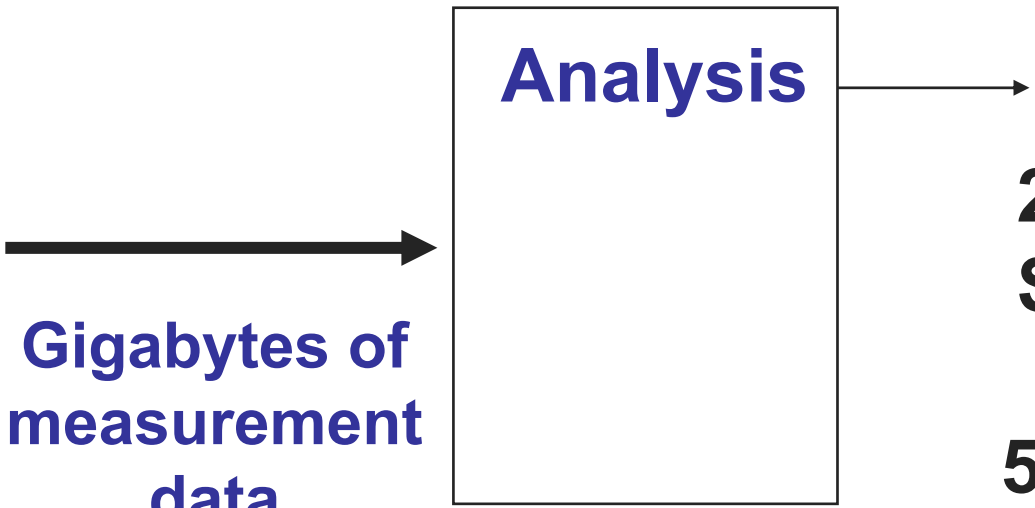

Traffic reports

20% is Kazaa from Steve's PC

50% is Kazaa from network A

AutoFocus: system structure

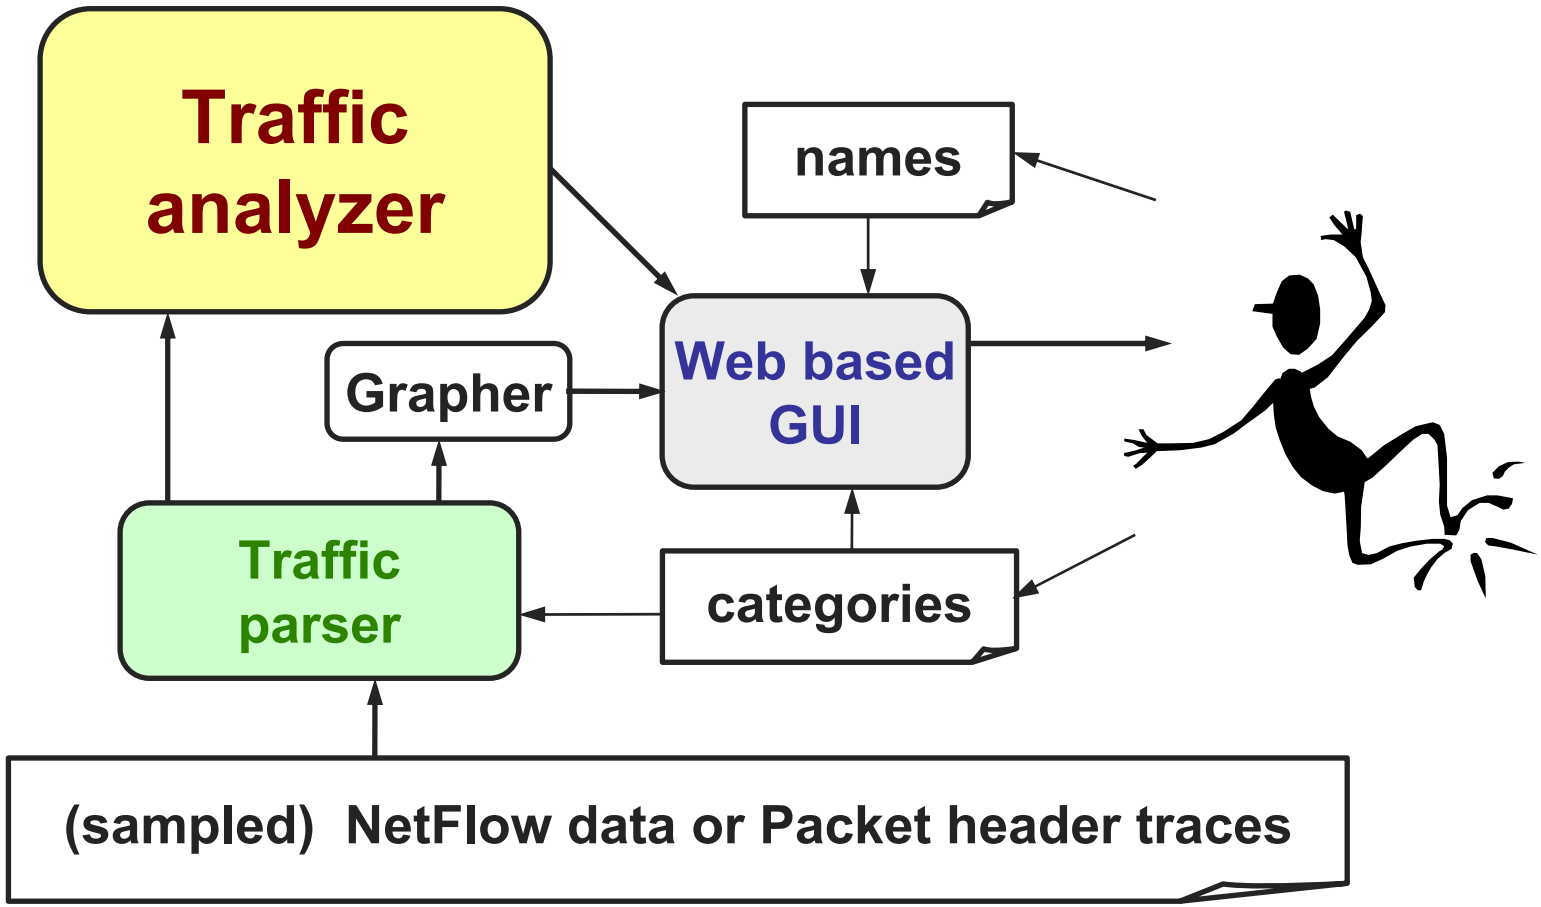

System details

- Availability
  - Downloadable
  - Free for educational, research and non-profit use
- Requirements
  - Linux or BSD (might run on other Unix OSes)
  - 256 Megs of RAM at least
  - 1-10 gigabytes of hard disk (depends on traffic)
  - Recent Netscape, Mozilla or I.E. (Javascript)
  - Needs no web server – no server side scripting

Traffic analysis approach

- Characterize traffic mix by describing **all** important traffic clusters
  - **Multi-field** clusters (e.g. flash crowd described by protocol, port number and IP address)
  - At the the **right** level of **granularity** (e.g. computer, proper prefix length)
  - Analysis is **automated** – finds insightful data without human guidance

Traffic clusters: example

- Incoming web traffic for CS Dept.
  - SrcIP=*,
  - DestIP in 132.239.64.0/21,
  - Proto=TCP,
  - SrcPort=80,
  - DestPort in [1024,65535]

Traffic report

- **Traffic reports** automatically list significant traffic clusters
- Describe only clusters above **threshold** (e.g. T=total of traffic/20)
- **Compression** removes redundant clusters whose traffic can be inferred from more specific clusters

Automatic cluster selection

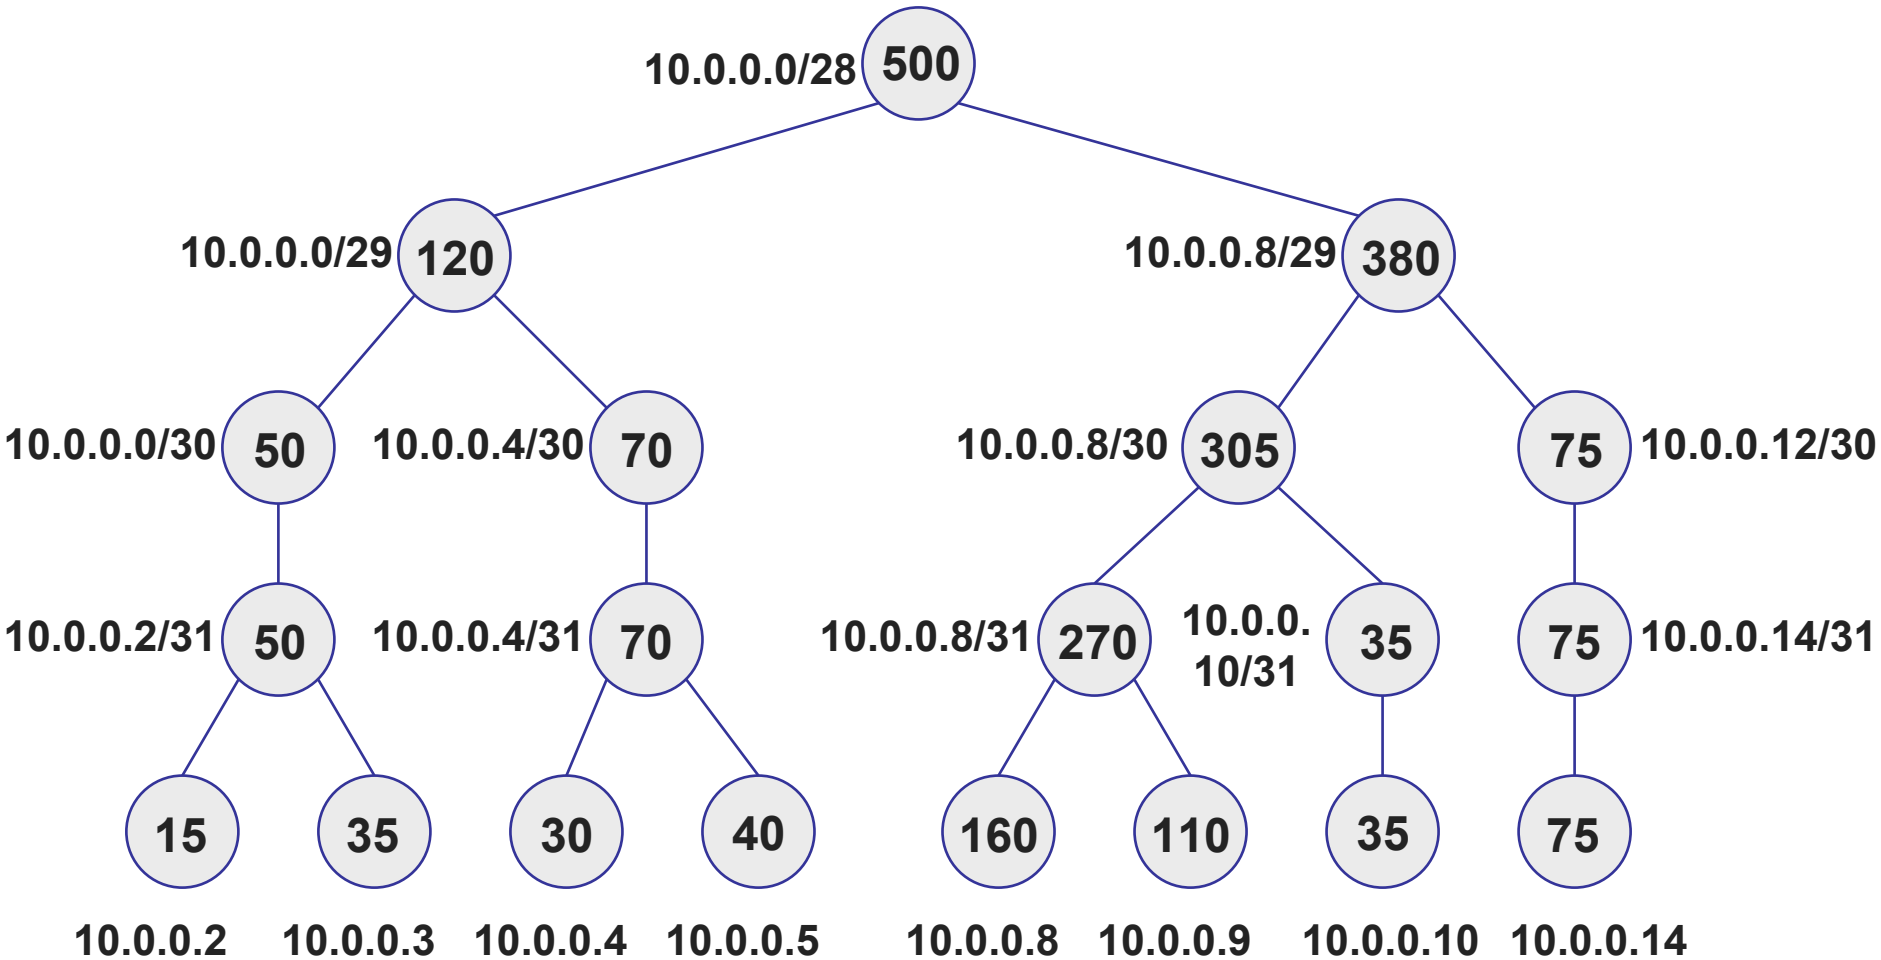

Automatic cluster selection

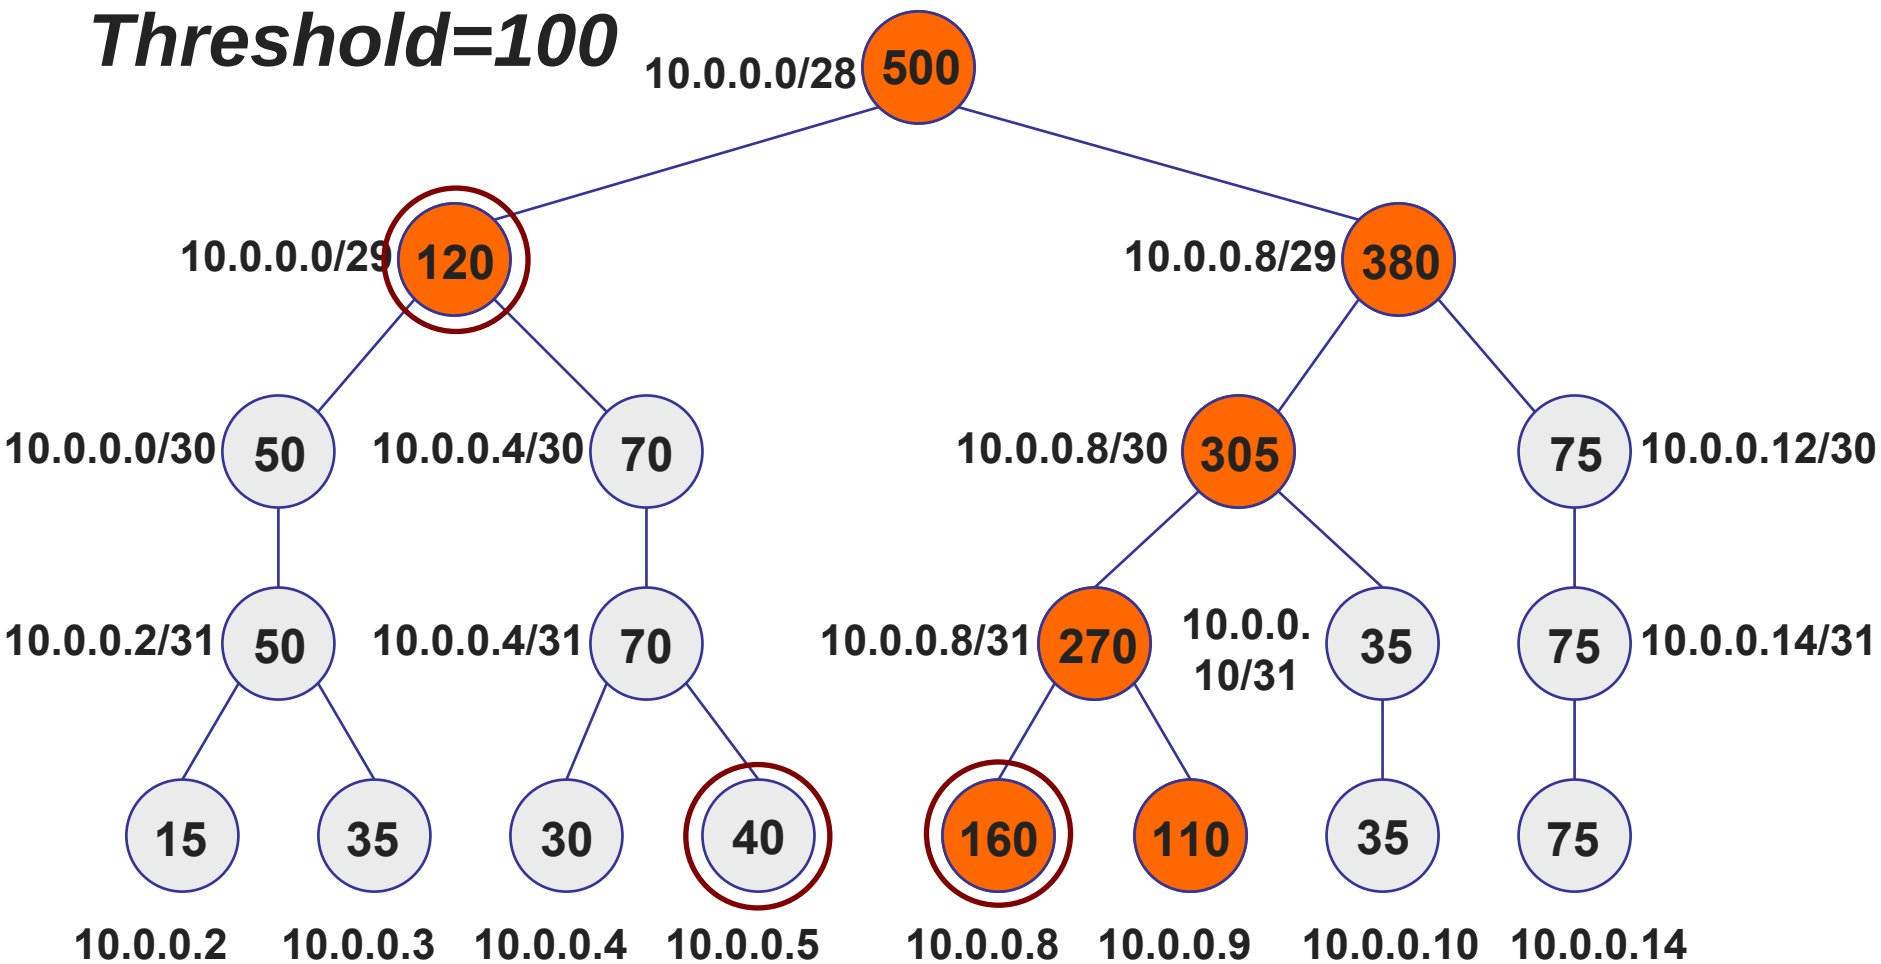

Automatic cluster selection

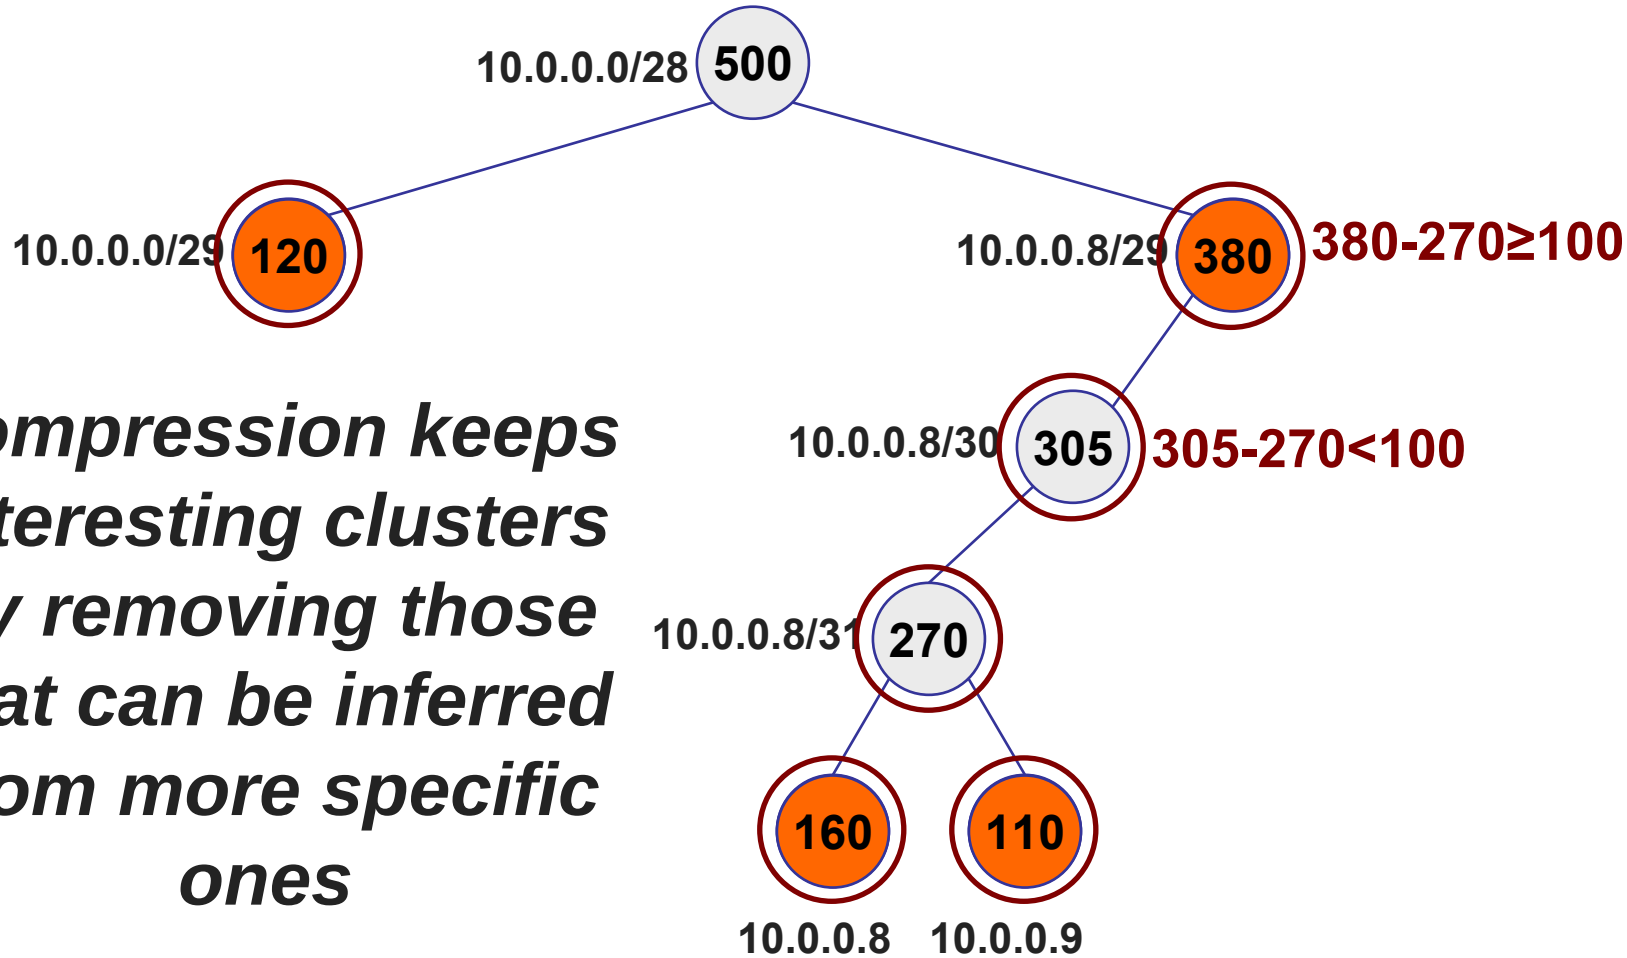

Compression keeps interesting clusters by removing those that can be inferred from more specific ones

Single field report example

Source IP	Traffic pkts.
10.0.0.0/29	120
10.0.0.8/29	380
10.0.0.8	160
10.0.0.9	110

AutoFocus has both single field and multi-field traffic reports

Graphical user interface

- Web based interface
- Many pre-computed traffic reports
- Interactive drill-down
- Traffic categories defined by user

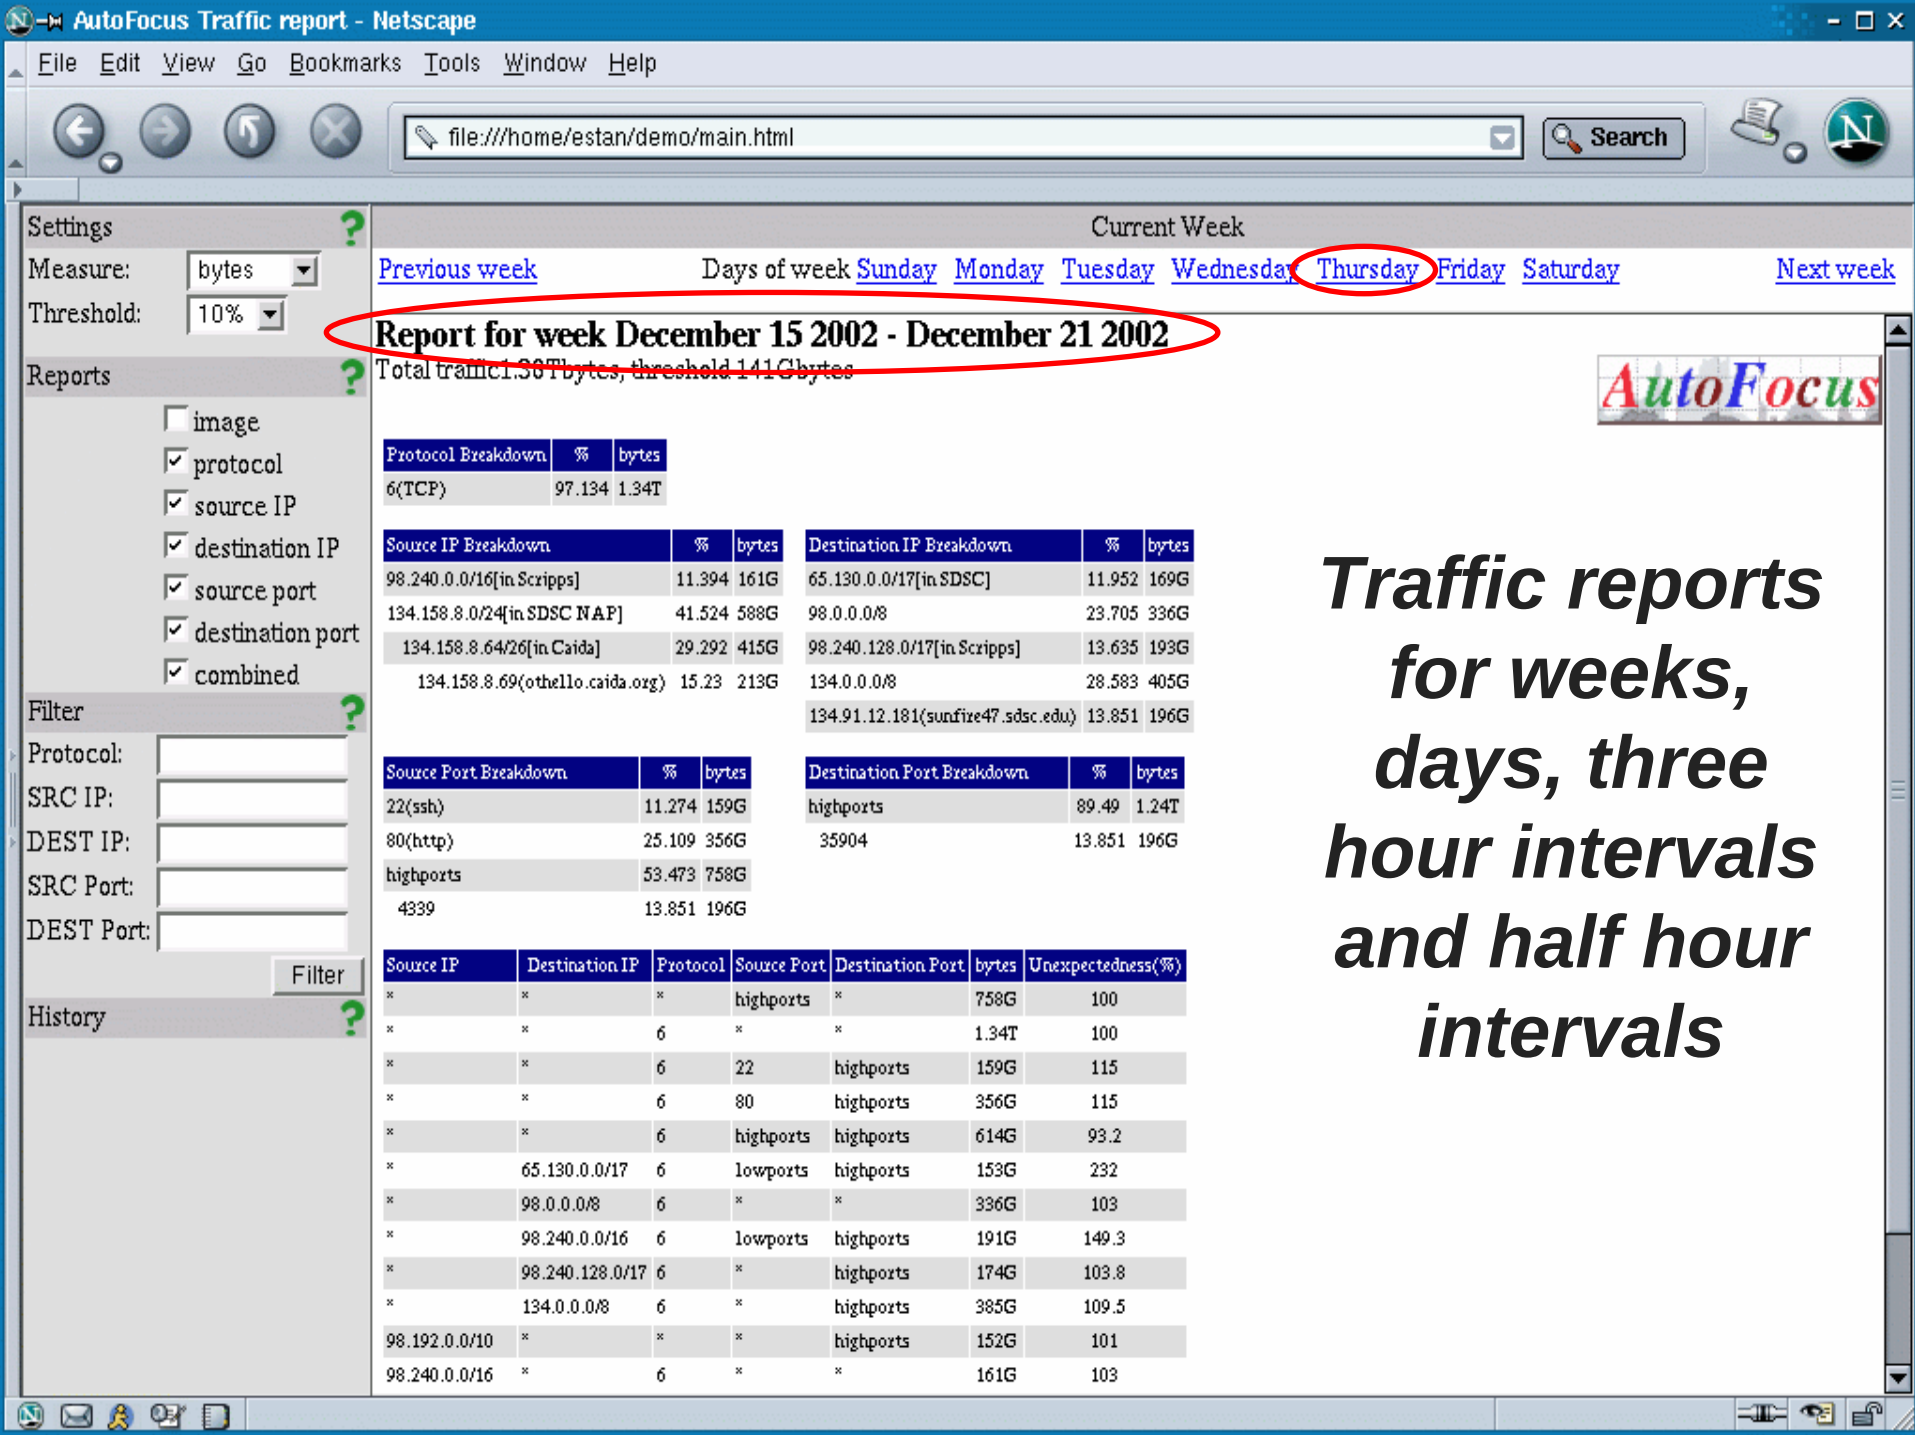

Traffic reports for weeks, days, three hour intervals and half hour intervals

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Threshold: 10%

Reports

- [ ] image
- [x] protocol
- [x] source IP
- [x] destination IP
- [x] source port
- [x] destination port
- [x] combined

Filter

Protocol:

SRC IP:

DEST IP:

SRC Port:

DEST Port:

Filter

History

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

Report for week December 15 2002 - December 21 2002

Total traffic1.38Tbytes, threshold 141Gbytes

AutoFocus

Protocol Breakdown	%	bytes
6(TCP)	97.134	1.34T

Source IP Breakdown	%	bytes
98.240.0.0/16[in Scripps]	11.394	161G
134.158.8.0/24[in SDSC NAP]	41.524	588G
134.158.8.64/26[in Caida]	29.292	415G
134.158.8.69(othello.caida.org)	15.23	213G

Destination IP Breakdown	%	bytes
65.130.0.0/17[in SDSC]	11.952	169G
98.0.0.0/8	23.705	336G
98.240.128.0/17[in Scripps]	13.635	193G
134.0.0.0/8	28.583	405G
134.91.12.181(sunfire47.sdsc.edu)	13.851	196G

Source Port Breakdown	%	bytes
22(ssh)	11.274	159G
80(http)	25.109	356G
highports	53.473	758G
4339	13.851	196G

Destination Port Breakdown	%	bytes
highports	89.49	1.24T
35904	13.851	196G

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	*	highports	*	758G	100
*	*	6	*	*	1.34T	100
*	*	6	22	highports	159G	115
*	*	6	80	highports	356G	115
*	*	6	highports	highports	614G	93.2
*	65.130.0.0/17	6	lowports	highports	153G	232
*	98.0.0.0/8	6	*	*	336G	103
*	98.240.0.0/16	6	lowports	highports	191G	149.3
*	98.240.128.0/17	6	*	highports	174G	103.8
*	134.0.0.0/8	6	*	highports	385G	109.5
98.192.0.0/10	*	*	*	highports	152G	101
98.240.0.0/16	*	6	*	*	161G	103

Traffic reports measure traffic in bytes, packets and flows, have various thresholds

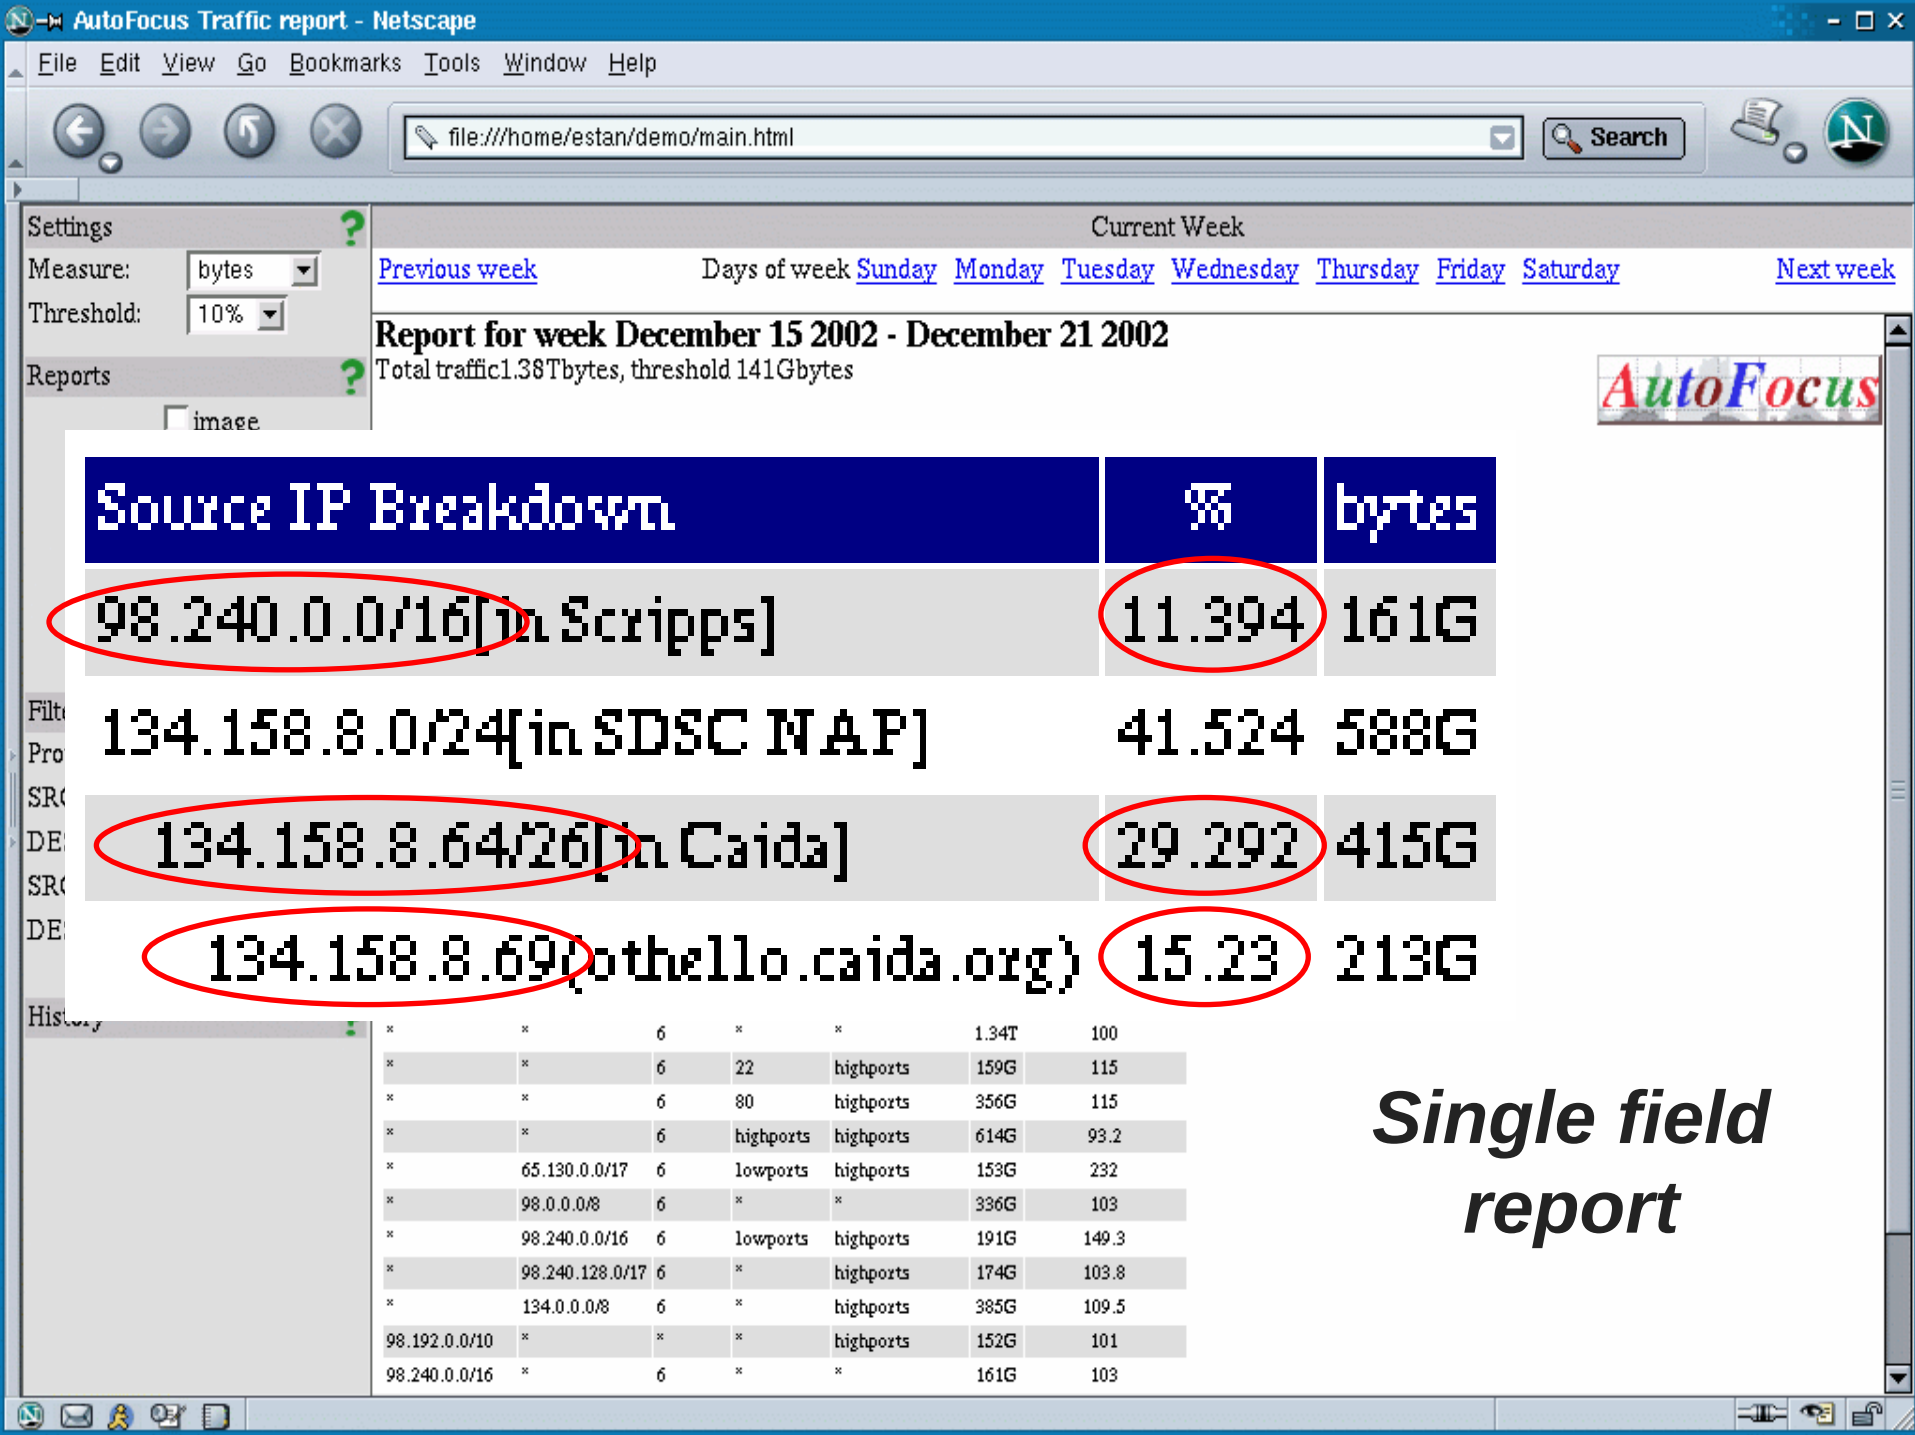
AutoFocus Traffic report - Netscape
File Edit View Go Bookmarks Tools Window Help
file:///home/estan/demo/main.html
Search
Settings
Measure: bytes
Threshold: 10%
Reports
image
Current Week
Previous week
Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday
Next week
Report for week December 15 2002 - December 21 2002
Total traffic1.38Tbytes, threshold 141Gbytes
AutoFocus
Source IP Breakdown
%
bytes
98.240.0.0/16[in Scripps]
11.394
161G
134.158.8.0/24[in SDSC NAP]
41.524
588G
134.158.8.64/26[in Caida]
29.292
415G
134.158.8.69(othello.caida.org)
15.23
213G
* * 6 * * 1.34T 100
* * 6 22 highports 159G 115
* * 6 80 highports 356G 115
* * 6 highports highports 614G 93.2
* 65.130.0.0/17 6 lowports highports 153G 232
* 98.0.0.0/8 6 * * 336G 103
* 98.240.0.0/16 6 lowports highports 191G 149.3
* 98.240.128.0/17 6 * highports 174G 103.8
* 134.0.0.0/8 6 * highports 385G 109.5
98.192.0.0/10 * * * highports 152G 101
98.240.0.0/16 * 6 * * 161G 103
Single field report

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	*	highports	*	758G	100
*	*	6	*	*	1.34T	100
*	*	6	22	highports	159G	115
*	*	6	80	highports	356G	115
*	*	6	highports	highports	614G	93.2
*	65.130.0.0/17	6	lowports	highports	153G	232
*	98.0.0.0/8	6	*	*	336G	103
*	98.240.0.0/16	6	lowports	highports	191G	149.3
*	98.240.128.0/17	6	*	highports	174G	103.8
*	134.0.0.0/8	6	*	highports	385G	109.5
98.192.0.0/10	*	*	*	highports	152G	101
98.240.0.0/16	*	6	*	*	161G	103
134.158.8.0/24	*	6	highports	highports	385G	140.9
134.158.8.0/25	*	6	highports	*	361G	134.9
134.158.8.64/26	*	6	lowports	highports	152G	93.7
134.158.8.69	134.91.12.181	6	4339	35904	196G	35720.5

134.158.8.64/26 * 6 lowports highports 152G 93.7
134.158.8.69 134.91.12.181 6 4339 35904 196G 35720.5

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Threshold: 10%

Reports

- [x] image
- [ ] protocol
- [ ] source IP
- [ ] destination IP
- [ ] source port
- [ ] destination port
- [ ] combined

Filter

Protocol:

SRC IP:

DEST IP:

SRC Port:

DEST Port:

Filter

History

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

Report for week December 15 2002 - December 21 2002

Total traffic1.38Tbytes, threshold 141Gbytes

AutoFocus

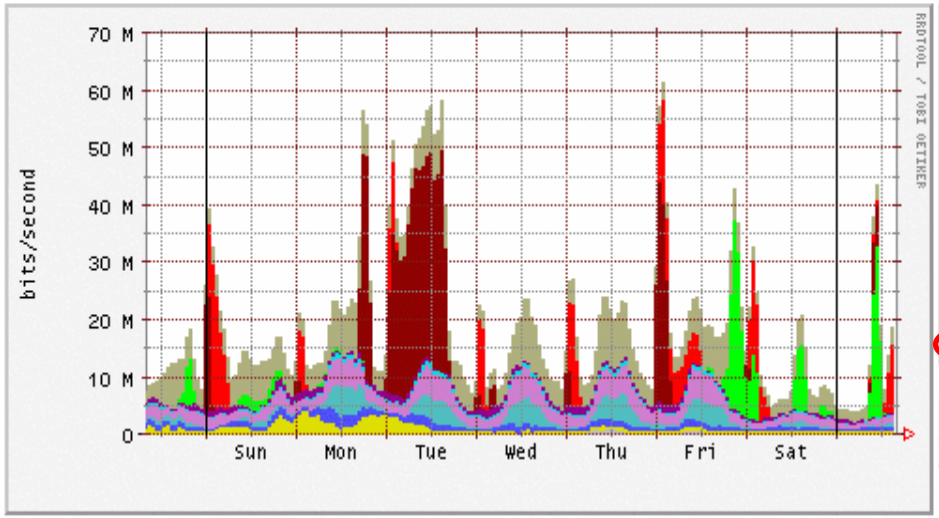

Category Details

- ICMP to Caida
- ICMP from Caida
- Multicast Beacon
- To ssh servers
- From noname Scripps web server
- To Scripps web clients
- Other web traffic
- Traffic towards web servers
- Squid traffic from the NLANR cache
- TCP traffic from boluring.stanford.edu to the Scripps /16
- Backups from Caida to hpss highport to highport
- Backups from Caida to terminator over ssh
- Other traffic

Built on Tobi Oetiker's RRD TOOL

Colors – user defined traffic categories

Separate reports for each category

Done

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Threshold: 10%

Reports

- [x] image
- [ ] protocol
- [ ] source IP
- [ ] destination IP
- [ ] source port
- [ ] destination port
- [x] combined

Filter

Protocol:

SRC IP:

DEST IP:

SRC Port:

DEST Port:

Filter

History

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

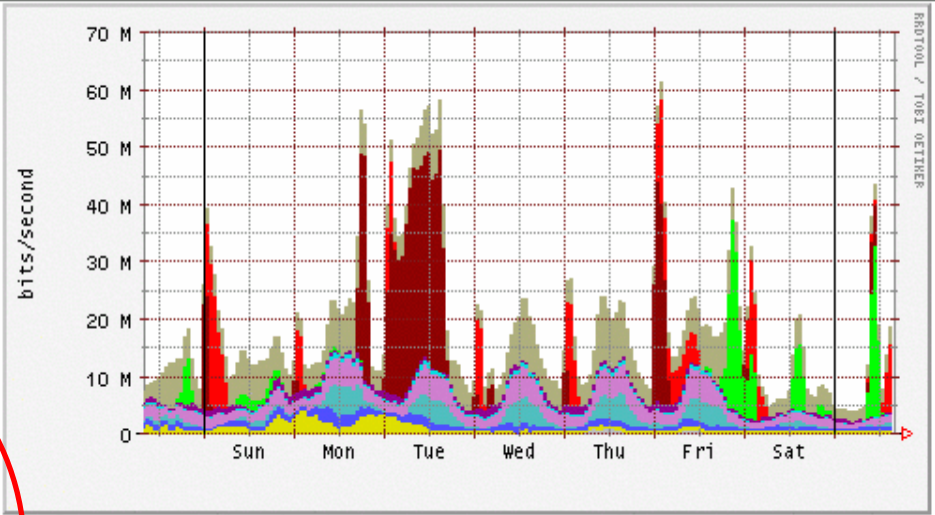

Category Details

- ICMP to Caida
- ICMP from Caida
- Multicast Beacon
- To ssh servers
- From noname Scripps web server
- To Scripps web clients
- Other web traffic
- Traffic towards web servers
- Squid traffic from the NLANR cache
- TCP traffic from bohring.stanford.edu to the Scripps /16
- Backups from Caida to hpss highport to highport
- Backups from Caida to terminator over ssh
- Other traffic

Built on Tobi Oetiker's RRD TOOL

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	*	highports	*	758G	100
*	*	6	*	*	1.34T	100
*	*	6	22	highports	159G	115
*	*	6	80	highports	356G	115
*	*	6	highports	highports	614G	93.2
*	65.130.0.0/17	6	lowports	highports	153G	232
*	98.0.0.0/8	6	*	*	336G	103
*	98.240.0.0/16	6	lowports	highports	191G	149.3
*	98.240.128.0/17	6	*	highports	174G	103.8
*	134.0.0.0/8	6	*	highports	385G	109.5
98.192.0.0/10	*	*	*	highports	152G	101
98.240.0.0/16	*	6	*	*	161G	103
134.158.8.0/24	*	6	highports	highports	385G	140.9
134.158.8.0/25	*	6	highports	*	361G	134.9
134.158.8.64/26	*	6	lowports	highports	152G	93.7

The filter and threshold allow interactive drill-down

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Threshold: 10%

Reports

- [x] image
- [ ] protocol
- [ ] source IP
- [ ] destination IP
- [ ] source port
- [ ] destination port
- [x] combined

Filter

Protocol:

SRC IP: 134.158.8.64/26

DEST IP:

SRC Port:

DEST Port:

Filter

History

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

Report for week December 15 2002 - December 21 2002

Total traffic1.38Tbytes, threshold 141Gbytes

AutoFocus

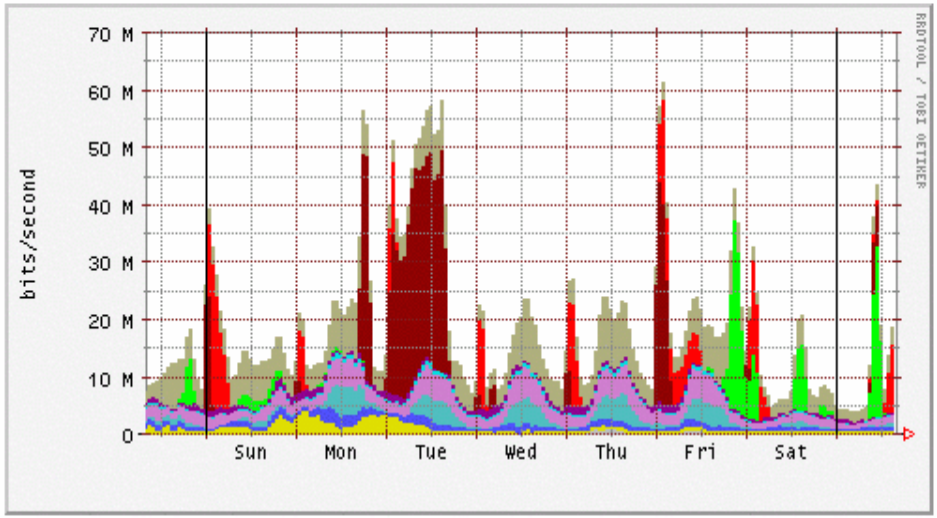

Category Details

- ICMP to Caida
- ICMP from Caida
- Multicast Beacon
- To ssh servers
- From noname Scripps web server
- To Scripps web clients
- Other web traffic
- Traffic towards web servers
- Squid traffic from the NLANR cache
- TCP traffic from boburing.stanford.edu to the Scripps /16
- Backups from Caida to hpss highport to highport
- Backups from Caida to terminator over ssh
- Other traffic

Built on Tobi Oetiker's RRD TOOL

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
134.158.8.64/26	*	6	lowports	highports	152G	93.7
134.158.8.69	134.91.12.181	6	4339	35904	196G	35720.5

The filter and threshold allow interactive drill-down

AutoFocus Traffic report - Netscape

File Edit View Go Bookmarks Tools Window Help

file:///home/estan/demo/main.html Search

Settings

Measure: bytes

Threshold: 2%

Reports

- [x] image
- [ ] protocol
- [ ] source IP
- [ ] destination IP
- [ ] source port
- [ ] destination port
- [x] combined

Filter

Protocol:

SRC IP: 134.158.8.64/26

DEST IP:

SRC Port:

DEST Port:

Filter

History

Dec 15 2002 - Dec 21 2002 (byt

Current Week

Previous week Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday Next week

Report for week December 15 2002 - December 21 2002

Total traffic1.38Tbytes, threshold 28.4Gbytes

AutoFocus

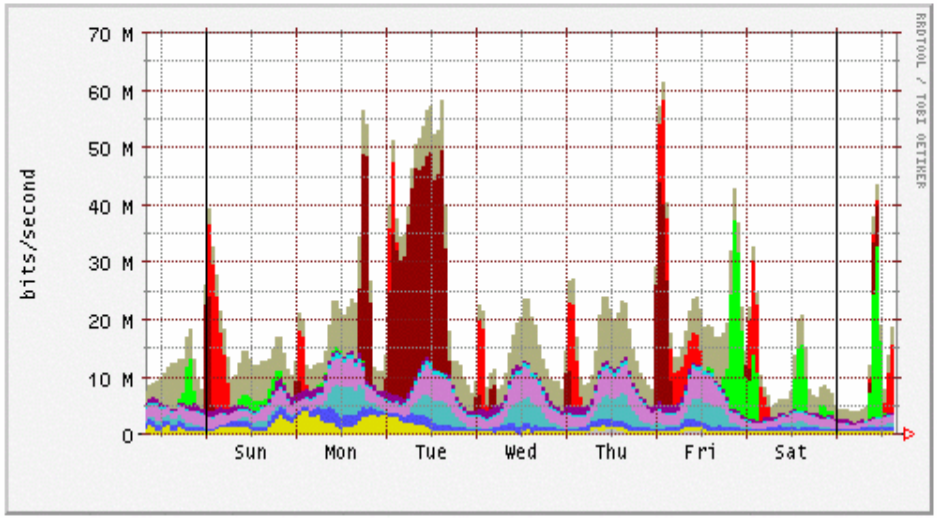

Category Details

- ICMP to Caida
- ICMP from Caida
- Multicast Beacon
- To ssh servers
- From noname Scripps web server
- To Scripps web clients
- Other web traffic
- Traffic towards web servers
- Squid traffic from the NLANR cache
- TCP traffic from boluring.stanford.edu to the Scripps /16
- Backups from Caida to hpss highport to highport
- Backups from Caida to terminator over ssh
- Other traffic

Built on Tobi Oetiker's RRD TOOL

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
134.158.8.64/26	*	6	highports	highports	230G	119.6
134.158.8.64/26	65.130.121.172	6	22	highports	136G	3483.8
134.158.8.64/27	65.130.121.172	6	22	highports	69.3G	2625.9
134.158.8.64/30	65.130.121.172	6	22	highports	36.3G	10634.7
134.158.8.69	134.91.12.181	6	4339	35904	196G	35720.5
134.158.8.80/29	*	6	lowports	highports	29.9G	256.2
134.158.8.96/27	*	6	*	*	134G	103
134.158.8.96/27	*	6	lowports	highports	61.8G	117.2
134.158.8.98	*	6	*	highports	43.4G	67.7
134.158.8.98	19.78.111.89	6	highports	22	30.3G	69383.5
134.158.8.116	65.130.121.172	6	22	highports	30.7G	10634.7

The filter and threshold allow interactive drill-down

Case study : SD-NAP

- Structure of regular traffic mix
  - Backups from CAIDA to tape server
  - FTP from SLAC Stanford
  - Scripps web traffic
  - Web & Squid servers
  - Large ssh traffic
  - Steady ICMP probing from CAIDA
- Unexpected events

Structure of regular traffic mix

- Backups from CAIDA to tape server

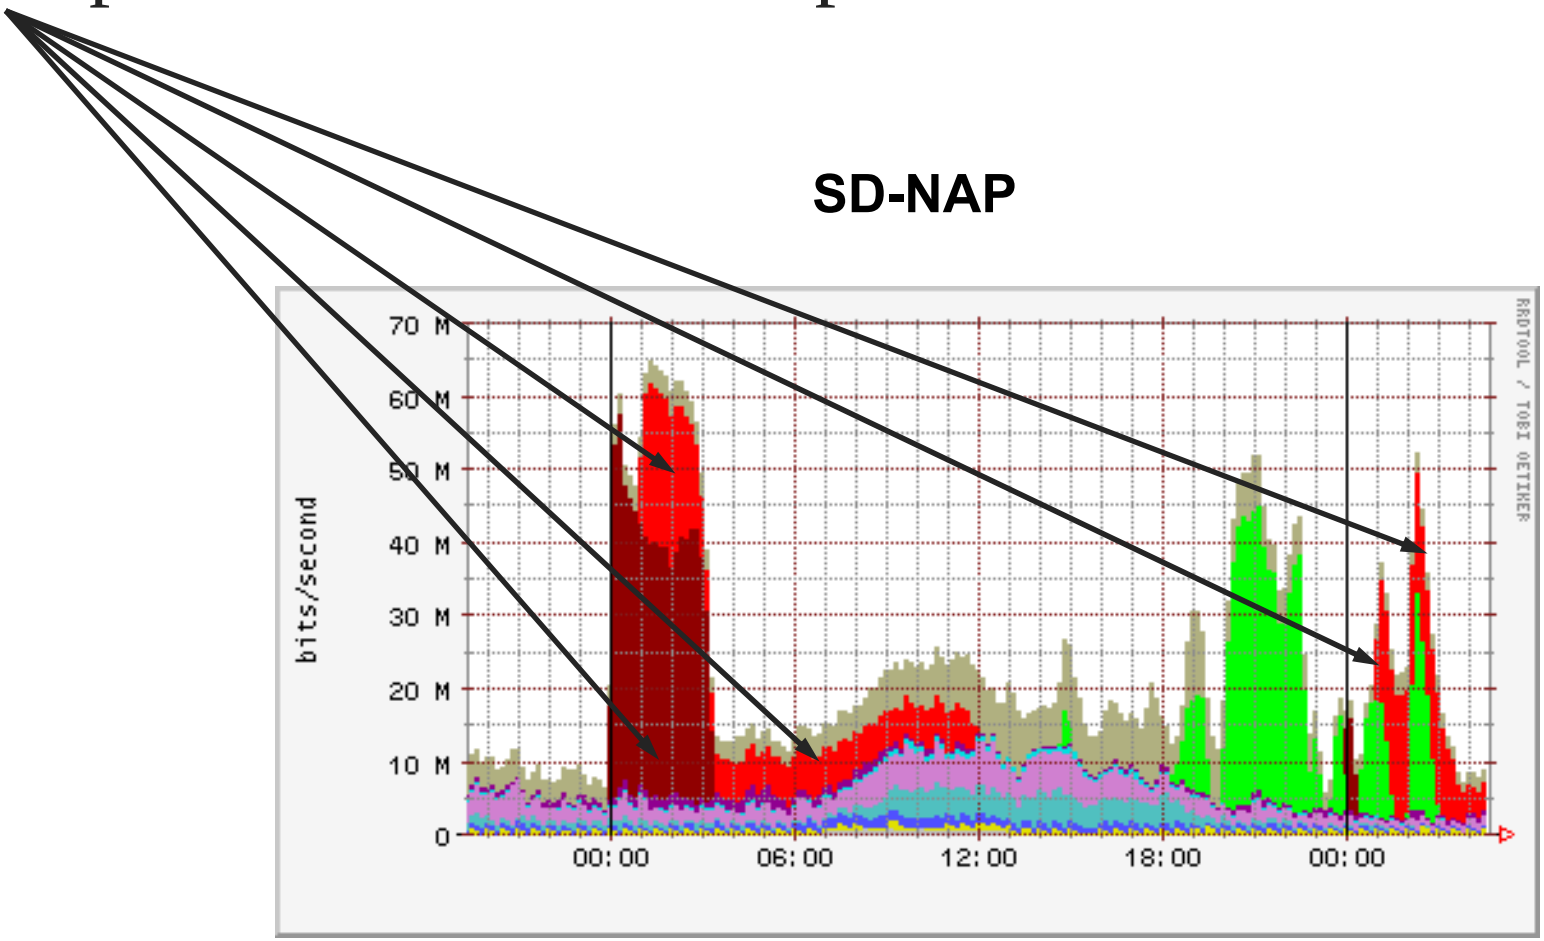

Structure of regular traffic mix

- Backups from CAIDA to tape server
  - Semi-regular time pattern

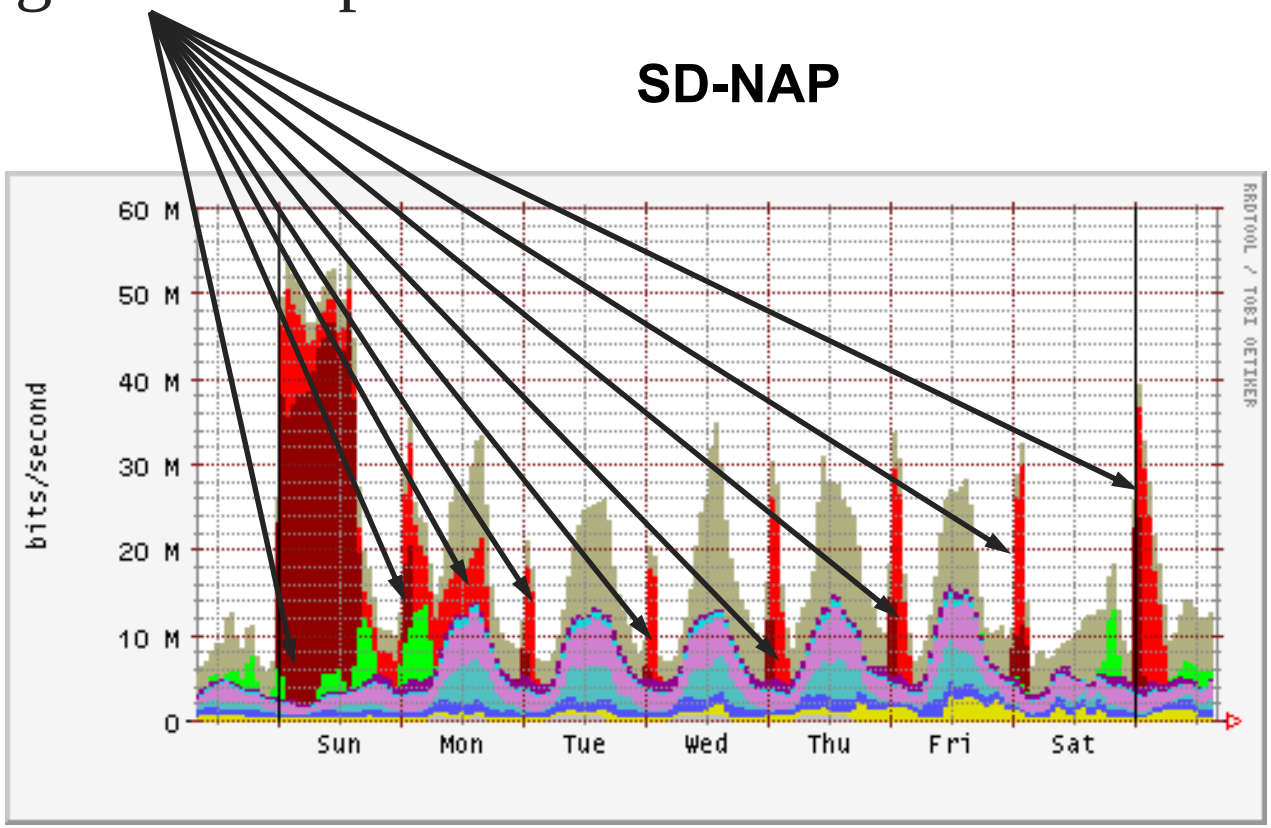

Structure of regular traffic mix

- Steady ICMP probing from CAIDA

The flow view highlights different traffic clusters

SD-NAP

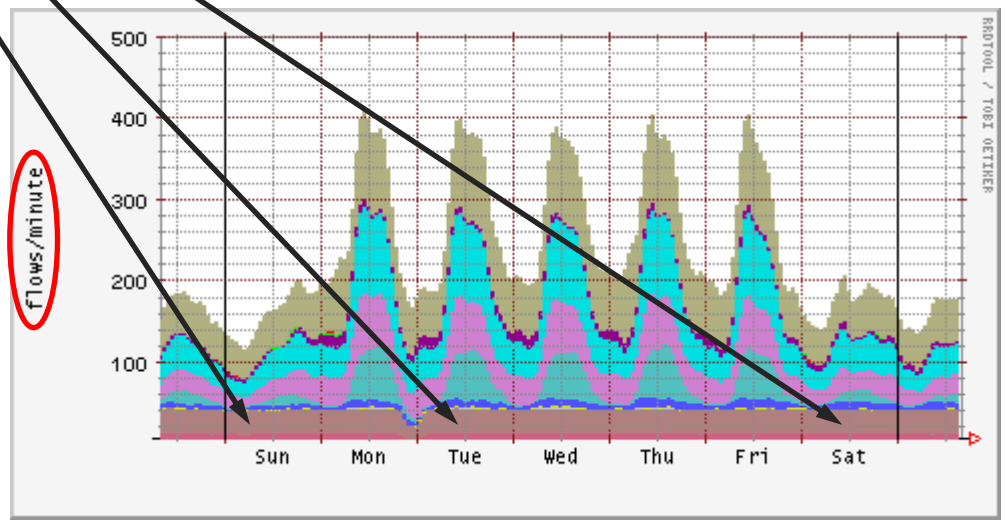

Analysis of unusual events

- Sapphire/SQL Slammer worm
  - Find worm port & proto automatically

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	6	highports	highports	827M	77.7
*	*	17	highports	1434	10.5G	112.6
*	152.249.0.0/16	*	*	*	604M	100
138.0.0.0/9	*	*	*	highports	3.66G	99.4
138.0.0.0/10	*	*	highports	*	3.68G	99.9
138.54.3.58	*	17	3341	1434	2.14G	672.5
138.54.11.4	*	17	7062	1434	950M	1551.3
152.249.56.0/22	*	*	highports	highports	723M	103.4
152.249.191.120	*	17	1959	1434	1.78G	810.0
152.249.191.121	96.0.0.0/8	17	1531	1434	645M	39523.7
152.249.210.3	*	17	4315	1434	2.36G	609.5
152.249.254.152	*	17	3787	1434	1.53G	941.8

Analysis of unusual events

- Sapphire/SQL Slammer worm
  - Can identify infected hosts

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	6	highports	highports	827M	77.7
*	*	17	highports	1434	10.5G	112.6
*	152.249.0.0/16	*	*	*	604M	100
138.0.0.0/9	*	*	*	highports	3.66G	99.4
138.0.0.0/10	*	*	highports	*	3.68G	99.9
138.54.3.58	*	17	3341	1434	2.14G	672.5
138.54.11.4	*	17	7062	1434	950M	1551.3
152.249.56.0/22	*	*	highports	highports	723M	103.4
152.249.191.120	*	17	1959	1434	1.78G	810.0
152.249.191.121	96.0.0.0/8	17	1531	1434	645M	39523.7
152.249.210.3	*	17	4315	1434	2.36G	609.5
152.249.254.152	*	17	3787	1434	1.53G	941.8

How can AutoFocus help you?

- Understand your regular traffic mix better
  - Better planning of network growth
  - Better traffic policing
- Understand unusual events
  - More effective reactions to worms, DoS attacks
  - Notice effects of route changes on traffic

Benefits w.r.t. existing tools

- Multi-field aggregation
- Automatically finds right granularity
- Drill-down
  - Per category reports
  - Using filter

Thank you!

Beta version of AutoFocus downloadable from

http://ial.ucsd.edu/AutoFocus/

Any questions?

Acknowledgements: Stefan Savage, George Varghese, Vern Paxson, David Moore, Liliana Estan, Mike Hunter, Pat Wilson, Jennifer Rexford, K Claffy, Alex Snoeren, Geoff Voelker, NIST,NSF

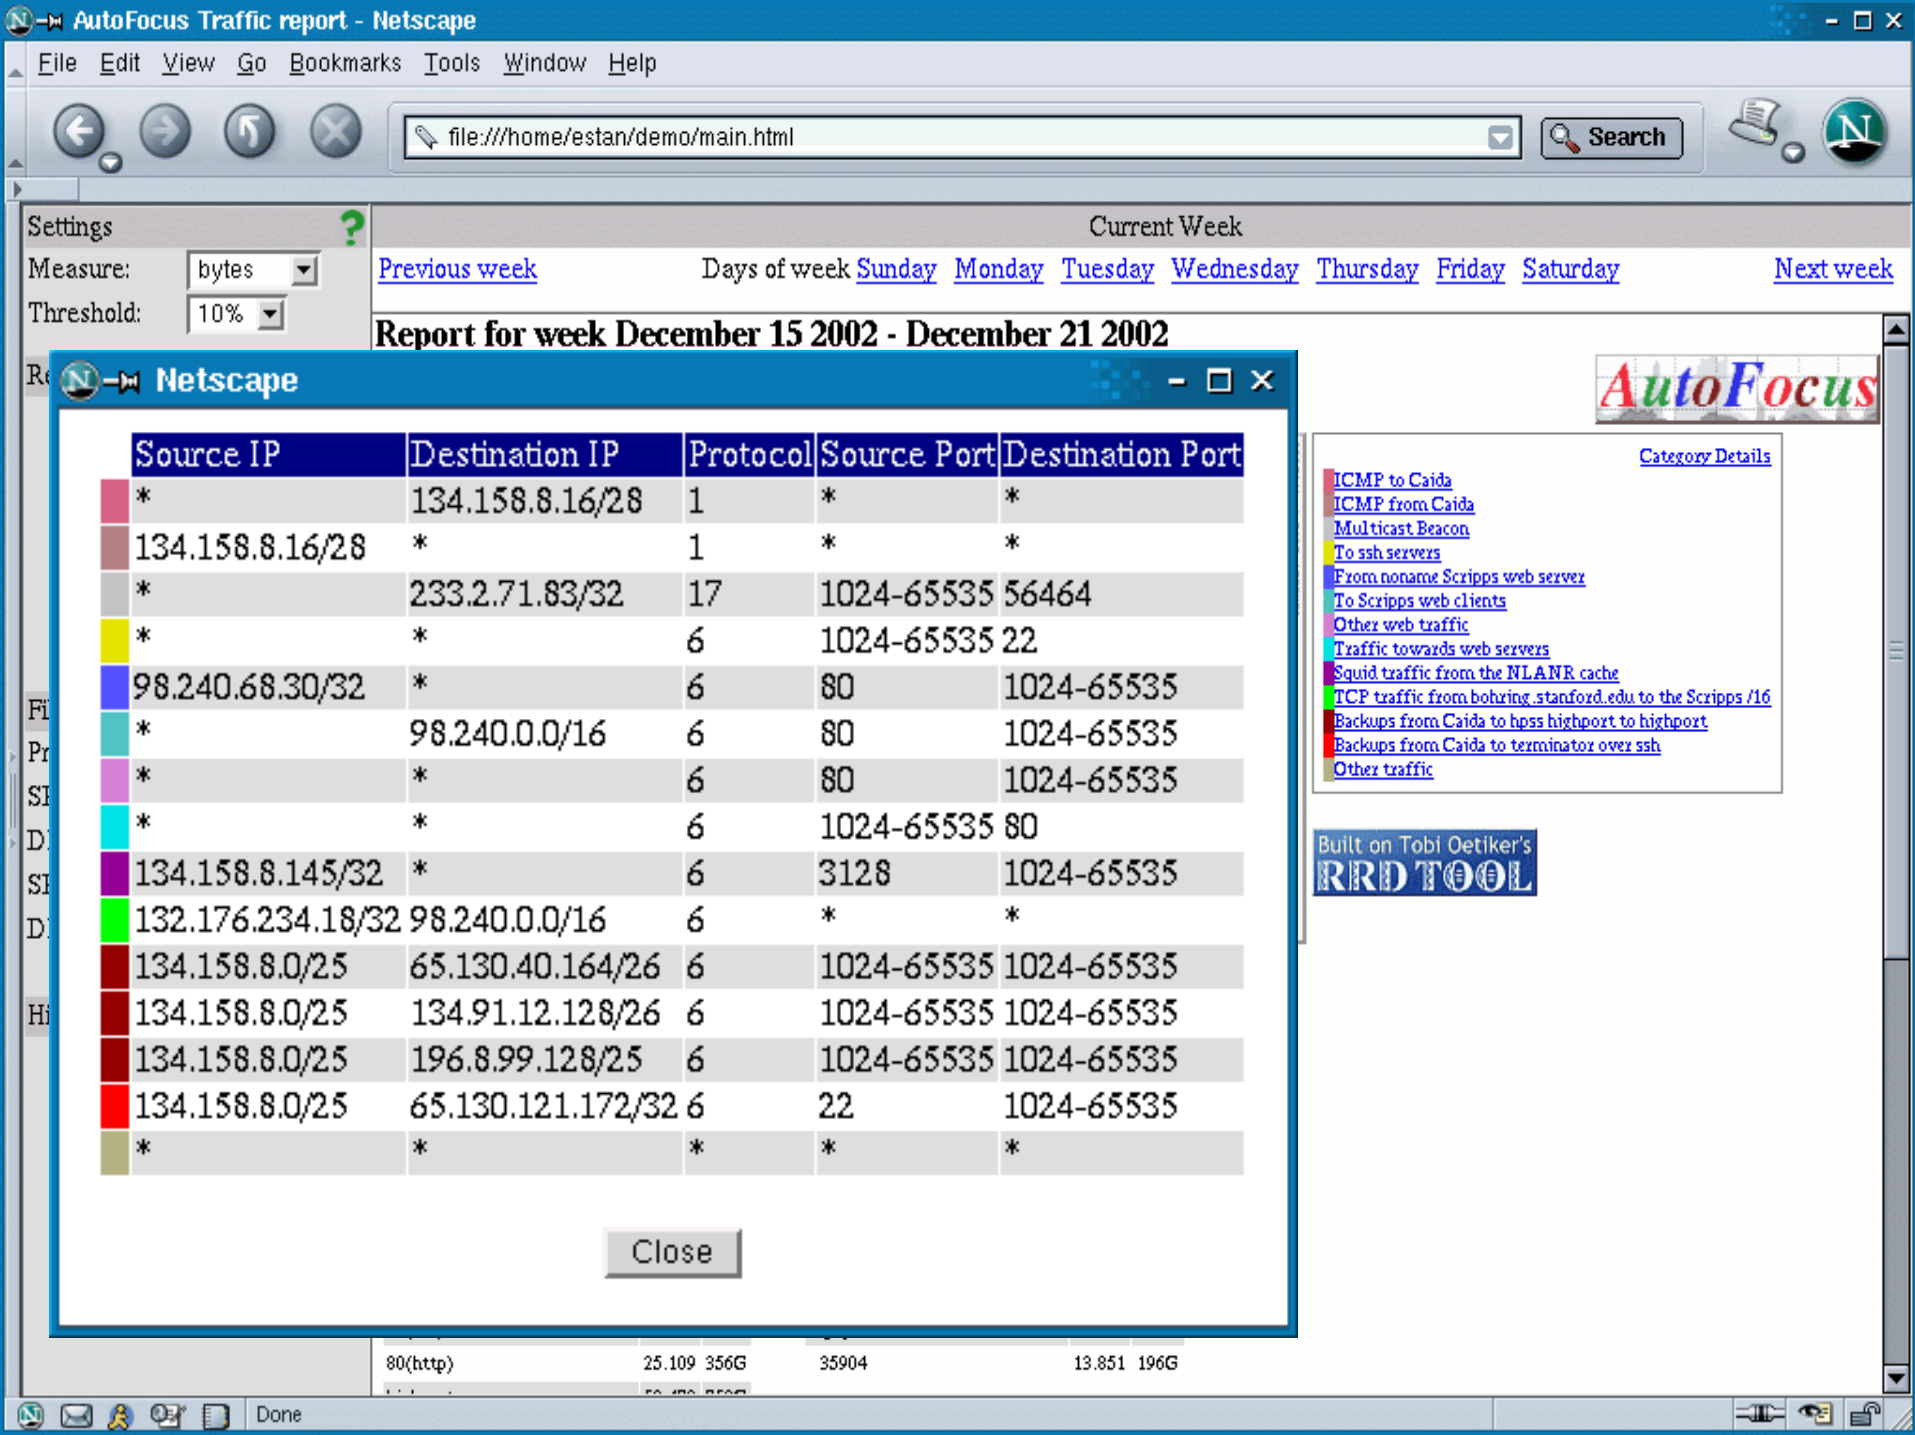
AutoFocus Traffic report - Netscape
File Edit View Go Bookmarks Tools Window Help
file:///home/estan/demo/main.html
Search
Settings
Measure: bytes
Threshold: 10%
Current Week
Previous week
Days of week Sunday Monday Tuesday Wednesday Thursday Friday Saturday
Next week
Report for week December 15 2002 - December 21 2002
Netscape
Source IP | Destination IP | Protocol | Source Port | Destination Port
* | 134.158.8.16/28 | 1 | * | *
134.158.8.16/28 | * | 1 | * | *
* | 233.2.71.83/32 | 17 | 1024-65535 | 56464
* | * | 6 | 1024-65535 | 22
98.240.68.30/32 | * | 6 | 80 | 1024-65535
* | 98.240.0.0/16 | 6 | 80 | 1024-65535
* | * | 6 | 80 | 1024-65535
* | * | 6 | 1024-65535 | 80
134.158.8.145/32 | * | 6 | 3128 | 1024-65535
132.176.234.18/32 | 98.240.0.0/16 | 6 | * | *
134.158.8.0/25 | 65.130.40.164/26 | 6 | 1024-65535 | 1024-65535
134.158.8.0/25 | 134.91.12.128/26 | 6 | 1024-65535 | 1024-65535
134.158.8.0/25 | 196.8.99.128/25 | 6 | 1024-65535 | 1024-65535
134.158.8.0/25 | 65.130.121.172/32 | 6 | 22 | 1024-65535
* | * | * | * | *
Close
AutoFocus
Category Details
ICMP to Caida
ICMP from Caida
Multicast Beacon
To ssh servers
From noname Scripps web server
To Scripps web clients
Other web traffic
Traffic towards web servers
Squid traffic from the NLANR cache
TCP traffic from boburing.stanford.edu to the Scripps /16
Backups from Caida to hpss highport to highport
Backups from Caida to terminator over ssh
Other traffic
Built on Tobi Oetiker's RRD TOOL
80(http) 25.109 356G 35904 13.851 196G
Done

Definition: unexpectedness

- To highlight non-obvious traffic clusters by using **unexpectedness label**
  - 50% of all traffic is web
  - Prefix B receives 20% of all traffic
  - The web traffic received by prefix B is 15% instead of 50%*20%=10%, unexpectedness label is 15%/10%=150%