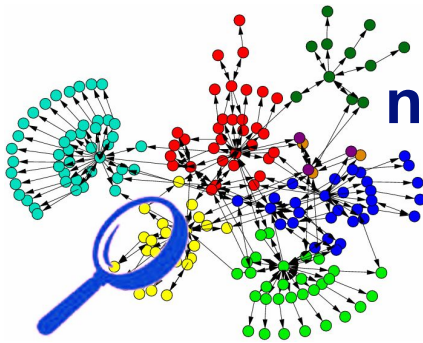


# SWITCH

The Swiss Education & Research Network



## nfdump and NfSen

SWINOG-12  
May 4th 2006, Bern  
Peter Haag



2006 © SWITCH

## nfdump and NfSen

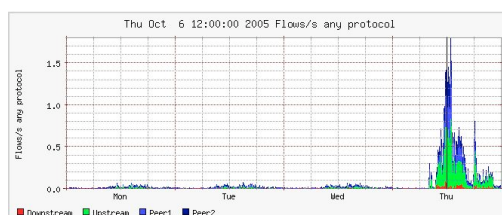
SWITCH

The Swiss Education & Research Network

Some operational questions, popping up now and then:

- Do you see this peek on port 445 as well ?
- What caused this peek on your network graph ?
- How did SoberR spread in your network ?
- Do we have any traffic pattern of this incident ?
- Which host/subnet consumes most of your bandwidth ?
- Which are the top talkers in your network ?
- ...

Sober.R



2006 © SWITCH

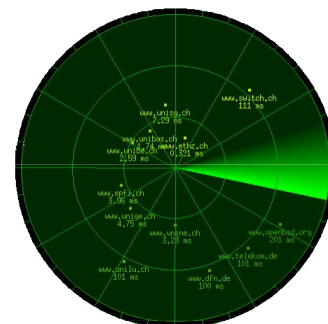


How to find answers for all these questions?

*Netflow turns out to be a good Data Source - although not the only one - for all kind of information and/or events to look at.*

.. in discussions with other teams:

- “Watch your flows for ...”
- “I’ve seen a lot of ... in our flows ...”
- “Host are infected, when you see flows to ...”



What is NetFlow?

**NetFlow is a traffic monitoring technology developed by Cisco Networks. Flows are unidirectional and contain connection related data such as:**

- Source and destination IP address.
- Source and destination port.
- Source and destination AS.
- Level3 protocol, ToS byte, TCP flags.
- Logical input and output interfaces.
- Bytes and packet counters.

Example:

2006-03-30 00:47:33.728 54.971 TCP 172.16.71.66:13599 -> 192.168.10.34:80 .A..SF 215 9890

***Netflow records never contain any user data!***

How to get netflow data and how to look at them?

Routers do provide netflow data ...  
but ...

```
Router# show ip cache flow
```

... seems not to be the solution for every task.

⇒ Tools to collect and look at the netflow data

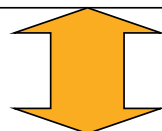
*nfdump and NfSen*



nfdump and NfSen:

NfSen:

- Web based frontend
- Display flows
- Framework to automate tasks

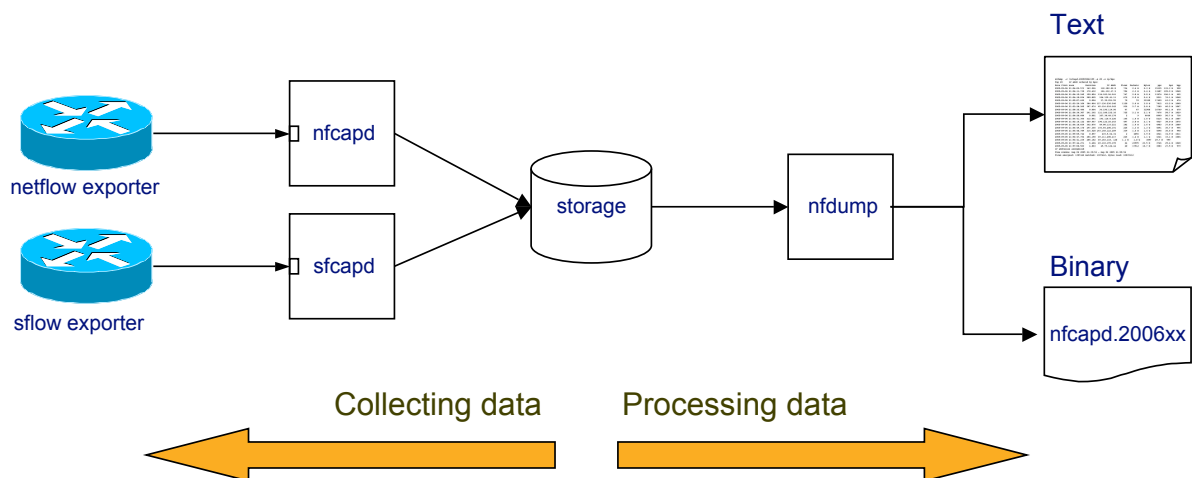


nfdump:

- Collect and store flows
- Process flows on command line



## nfdump overview :



# nfdump and NfSen

## nfdump features:

- CMD line based tool comparable to tcpdump.
- Written in C ⇒ fast.
- Stores netflow data in time sliced files.
- Supports netflow format v5,v7 and v9.
- All processing options support IPv4 and IPv6.
- Powerful pcap like filter syntax:  
( proto tcp and dst net 172.16/16 and src port > 1024 and bytes < 600 )  
or ( bps > 1k and ...
- Flexible flow aggregation.
- Efficient filter engine: > 6 Mio flows/s on 3GHz Intel.
- Lots of fast Top N statistics.
- Anonymizing of IP addresses. ( Crypto-Pan )
- User defined output formats.



**Example:**

List the first 20 tcp flows:

```

forth% nfdump -r /data/rz/nfcapd.200603300150 -K 123.. -c 20 'proto tcp'
Date flow start      Duration Proto      Src IP Addr:Port      Dst IP Addr:Port      Flags      Packets      Bytes
2006-03-30 00:43:40.569 82.880 TCP      130.20.234.125:58035 -> 200.66.27.5:61486 .AP...      673      199208
2006-03-30 00:43:40.569 82.880 TCP      200.66.27.5:61486 -> 130.20.234.125:58035 .A....      421      19674
2006-03-30 00:44:00.082 63.113 TCP      130.20.234.125:55697 -> 159.93.88.3:60454 .AP...      814      1.1 M
2006-03-30 00:44:00.082 63.113 TCP      159.93.88.3:60454 -> 130.20.234.125:55697 .A....      498      25896
2006-03-30 00:45:02.647 0.431 TCP      193.246.238.35:80 -> 192.254.4.182:56547 .A....      1        1500
2006-03-30 00:45:02.647 0.431 TCP      192.254.4.182:56547 -> 193.246.238.35:80 .A....      3        156
2006-03-30 00:45:02.813 0.000 TCP      130.20.234.124:59112 -> 194.50.123.176:45458 .A...F      1         52
2006-03-30 00:45:02.913 0.000 TCP      192.254.4.167:58659 -> 49.20.115.83:80 .A...F      1         52
2006-03-30 00:45:02.913 0.000 TCP      129.66.105.181:11248 -> 192.254.4.183:80 ....S.      1         46
2006-03-30 00:45:02.913 0.000 TCP      192.254.4.183:80 -> 129.66.105.181:11248 .A...S.      1         46
2006-03-30 00:45:02.879 0.000 TCP      129.66.105.181:11247 -> 192.254.4.183:80 .AP...      1        515
2006-03-30 00:45:02.879 0.000 TCP      192.254.4.183:80 -> 129.66.105.181:11247 .A....      1         46
2006-03-30 00:45:02.913 0.355 TCP      214.203.35.177:19027 -> 130.20.234.125:80 .A....      3        156
2006-03-30 01:40:02.347 300.572 TCP      dffe:e6...:199:fd.119 -> dc7e:18...:fe99:2.35541 .AP...      811      66835
2006-03-30 01:40:02.347 300.572 TCP      dc7e:18...:fe99:2.35541 -> dffe:e6...:199:fd.119 .AP.S.      4850      6.9 M
2006-03-30 00:45:02.895 0.000 TCP      192.254.4.183:80 -> 192.254.179.207:56323 .AP...      1        129
2006-03-30 00:45:02.978 0.000 TCP      194.50.123.176:45465 -> 130.20.234.124:55652 ....S.      1         60
2006-03-30 00:45:03.013 0.000 TCP      130.20.234.125:21 -> 50.242.99.240:61288 .A...S.      1         48
2006-03-30 00:45:03.009 0.000 TCP      156.32.82.45:35110 -> 130.20.234.124:25 ....S.      1         60
2006-03-30 00:45:03.041 0.000 TCP      130.20.234.125:80 -> 130.219.188.88:57168 .A....      1         52
IP addresses anonymized
Time window: 2006-03-30 00:40:02 - 2006-03-30 01:49:10
Total flows: 15970 matched: 20, skipped: 0, Bytes read: 838972
Sys: 0.007s flows/second: 2044290.8 Wall: 0.004s flows/second: 3391378.2

```

**Example:**

Show the top 15 IP addresses consuming most bandwidth:

```

forth% nfdump -r /data/rz/nfcapd.200603300150 -K 123... -n 20 -s ip/bps
Top 15 IP Addr ordered by bps:
Date first seen      Duration Proto      IP Addr      Flows      Packets      Bytes      pps      bps      bpp
2006-03-30 00:47:39.999 0.001 TCP      64.132.143.51 2          19      18004      18999      137.4 M      947
2006-03-30 00:45:00.737 0.001 TCP      194.64.105.184 2          20      13600      20000      103.8 M      680
2006-03-30 00:49:16.016 0.001 TCP      163.3.33.241 2          9      12046      9000      91.9 M      1338
2006-03-30 00:49:52.902 0.001 TCP      92.37.170.104 2          10      9208      10000      70.3 M      920
2006-03-30 00:45:06.853 0.001 TCP      214.214.200.81 2          6      6931      6000      52.9 M      1155
2006-03-30 00:46:32.363 0.001 TCP      68.142.57.84 2          10      6720      10000      51.3 M      672
2006-03-30 00:46:30.764 0.001 TCP      151.80.146.115 2          7      6680      7000      51.0 M      954
2006-03-30 00:48:36.966 0.001 TCP      129.4.38.113 2          8      6184      8000      47.2 M      773
2006-03-30 00:49:31.903 0.001 TCP      33.135.213.117 2          6      6104      6000      46.6 M      1017
2006-03-30 01:42:48.834 0.001 TCP      90.38.160.152 2          8      5941      8000      45.3 M      742
2006-03-30 00:48:02.473 0.001 TCP      131.144.55.170 2          6      5608      6000      42.8 M      934
2006-03-30 00:49:29.424 0.001 TCP      24.11.195.220 2          4      4880      4000      37.2 M      1220
2006-03-30 00:48:53.293 0.001 TCP      88.53.69.175 2          6      4721      6000      36.0 M      786
2006-03-30 00:45:41.780 0.001 TCP      49.30.8.60 2          6      3822      6000      29.2 M      637
2006-03-30 01:42:51.618 0.002 TCP      220.24.222.74 2          10      7605      5000      29.0 M      760
IP addresses anonymized
Time window: 2006-03-30 00:40:02 - 2006-03-30 01:49:58
Total flows: 19224 matched: 19224, skipped: 0, Bytes read: 1009920
Sys: 0.046s flows/second: 410112.0 Wall: 0.009s flows/second: 2022089.0

```

## Example:

### Show port scanning candidates:

```
forth% nfdump -r /data/rz/nfcapd.200603300150 -K 123... -A srcip,dstport -s record/packets 'not proto icmp and bytes < 100
and bpp < 100 and packets < 5 and not port 80 and not port 53 and not port 110 and not port 123'
```

Aggregated flows 72506

Top 10 flows ordered by packets:

Date flow start	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Packets	Bytes	bps	Flows
2006-03-30 01:49:23.800	243.842	TCP	83.130.48.231:0	->	0.0.0.0:4899	142172	6.5 M	223891	71151
2006-03-30 01:50:48.603	236.035	TCP	165.17.105.18:0	->	0.0.0.0:5900	34452	1.6 M	56049	17232
2006-03-30 01:52:30.169	143.944	TCP	117.128.149.163:0	->	0.0.0.0:41523	9982	479136	26629	5143
2006-03-30 01:49:22.650	303.173	TCP	221.200.120.170:0	->	0.0.0.0:1433	4638	222624	5874	2319
2006-03-30 01:49:53.945	299.401	TCP	211.135.150.43:0	->	0.0.0.0:135	3273	157104	4197	3273
2006-03-30 01:49:27.196	194.565	TCP	201.143.63.114:0	->	0.0.0.0:139	1845	88560	3641	1613
2006-03-30 01:52:05.471	96.768	TCP	198.246.113.17:0	->	0.0.0.0:445	1678	80544	6658	954
2006-03-30 01:49:54.012	300.038	UDP	210.117.33.36:0	->	0.0.0.0:137	1471	114738	3059	1471
2006-03-30 01:49:22.970	328.838	TCP	164.88.206.114:0	->	0.0.0.0:135	1432	68736	1672	1077
2006-03-30 01:53:00.822	112.524	TCP	24.169.235.184:0	->	0.0.0.0:135	1254	60192	4279	1254

IP addresses anonymized

Time window: 2006-03-30 01:34:53 - 2006-03-30 01:54:57

Total flows: 1178835 matched: 245494, skipped: 0, Bytes read: 57559680

Sys: 0.634s flows/second: 1856716.7 Wall: 0.632s flows/second: 1862657.6

## Example:

### Show the top 15 /24 subnets exchanging most traffic:

```
forth% nfdump -r /data/rz/nfcapd.200603300150 -K 123... -n 15 -A srcip4/24,dstip4/24 -s record/bytes
```

Aggregated flows 7525

Top 15 flows ordered by bytes:

Date flow start	Duration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Packets	Bytes	Flows
2006-03-30 00:41:06.140	4102.844	TCP	130.20.234.0:0	->	130.254.221.0:0	79455	95.1 M	14
2006-03-30 00:42:50.622	4022.361	TCP	130.20.234.0:0	->	194.90.158.0:0	42179	58.2 M	13
2006-03-30 00:40:51.729	4054.221	TCP	130.20.234.0:0	->	220.63.34.0:0	39593	56.0 M	6
2006-03-30 01:41:42.025	443.957	TCP	130.20.224.0:0	->	163.3.42.0:0	30543	43.3 M	7
2006-03-30 00:41:06.140	4102.844	TCP	130.254.221.0:0	->	130.20.234.0:0	60178	29.1 M	14
2006-03-30 01:39:56.087	600.881	TCP	130.20.234.0:0	->	194.84.7.0:0	17836	24.9 M	9
2006-03-30 00:44:39.128	3900.855	TCP	130.20.234.0:0	->	214.124.39.0:0	15912	22.6 M	9
2006-03-30 01:41:01.414	529.568	UDP	130.20.223.0:0	->	130.20.220.0:0	15549	21.4 M	8
2006-03-30 01:41:03.371	329.612	TCP	194.114.160.0:0	->	130.20.234.0:0	14126	20.1 M	4
2006-03-30 01:40:12.986	300.997	TCP	130.20.234.0:0	->	194.168.190.0:0	13101	18.7 M	2
2006-03-30 01:41:24.088	506.896	TCP	130.20.234.0:0	->	24.50.25.0:0	12433	17.8 M	2
2006-03-30 01:43:04.047	300.870	TCP	165.242.80.0:0	->	130.20.234.0:0	9966	14.3 M	1
2006-03-30 00:43:47.441	3935.542	TCP	130.20.234.0:0	->	205.175.61.0:0	10445	13.4 M	15
2006-03-30 00:44:01.619	332.758	TCP	130.20.234.0:0	->	194.61.253.0:0	8973	12.7 M	101
2006-03-30 01:42:43.123	300.860	TCP	130.20.234.0:0	->	69.155.45.0:0	7872	11.3 M	1

IP addresses anonymized

Time window: 2006-03-30 00:40:02 - 2006-03-30 01:49:58

Total flows: 19224 matched: 18797, skipped: 0, Bytes read: 1009920

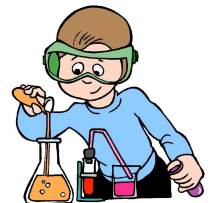
Sys: 0.062s flows/second: 307588.9 Wall: 0.010s flows/second: 1839969.4

## The art of filter design:

- ... depends on your problem you want to look at.
  - Incident Analysis.
  - Host tracking.
  - Port Scanning.
  - Operational issues.
- ... depends on your network.

*nfdump does not do your job, but supports you in doing your job!*

*Use the power of nfdump's filter syntax!*

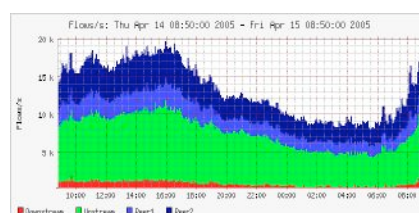


## nfdump is:

- Powerful
- Flexible
- Easy to use
- Fast
- ...

*but ...*

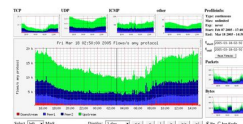
*... don't we all like pictures?*



⇒ NfSen

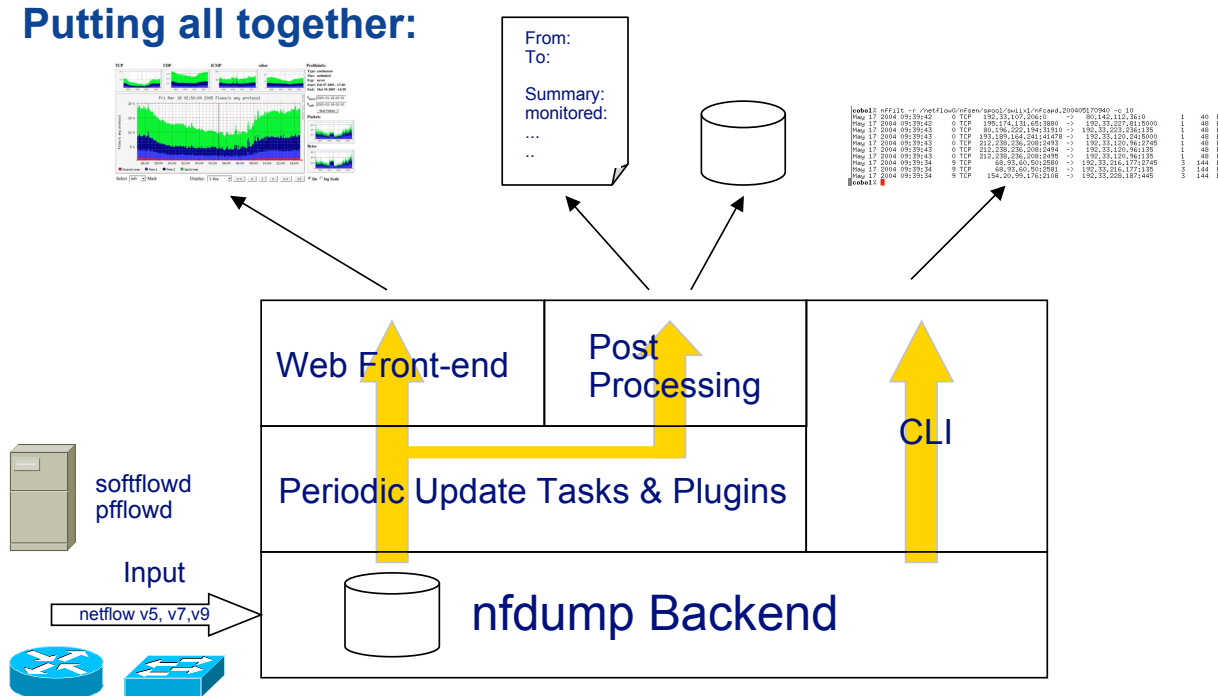
## NfSen features:

- Use the power of nfdump as backend tool. ⇒ modular design.
- Pictures!
- Drill down from overview to the details down to the specific flows.
- Graph current network situation.
- Graph specific profiles.
  - Track hosts, ports etc. from live data.
  - Profile hosts involved in incidents from history data.
- Analyse a specific time window.
- Web based.
- Automatically post process netflow data for reporting and alerting purpose.
- Flexible extensions using plugins.
- Easy to use.
- Auto - Cleanup. Aging data files: max space, max lifetime.



# nfdump and NfSen

## Putting all together:





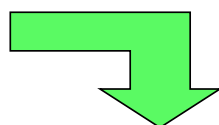


## nfdump and NfSen

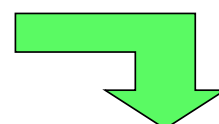
**SWITCH**  
The Swiss Education & Research Network

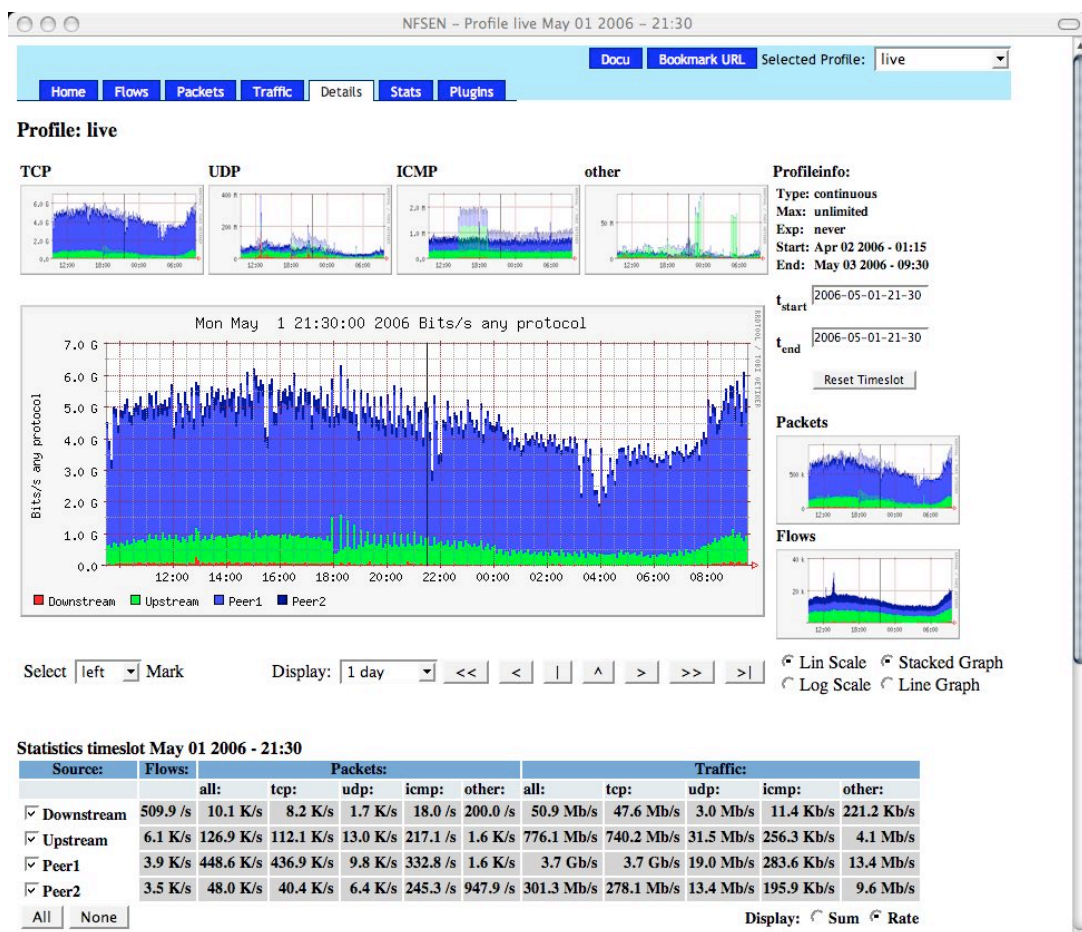
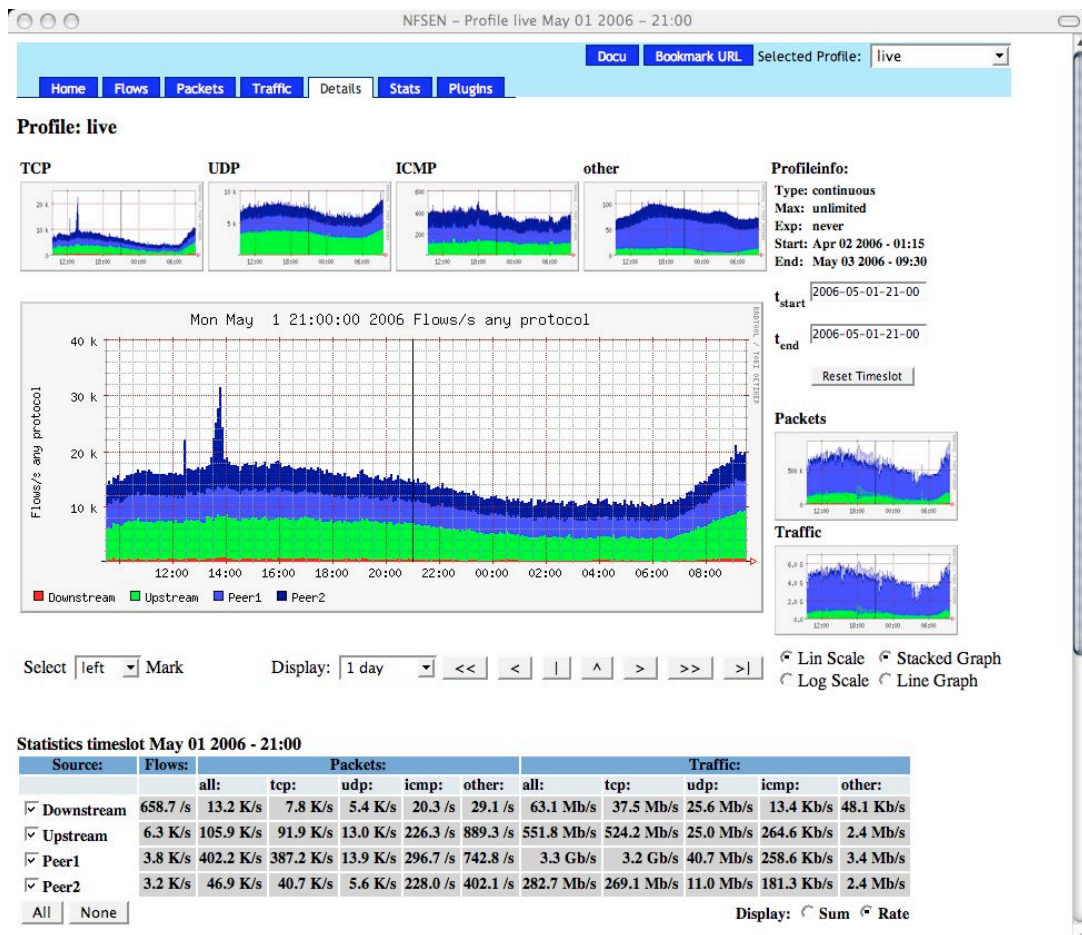


Overview ⇒ Details

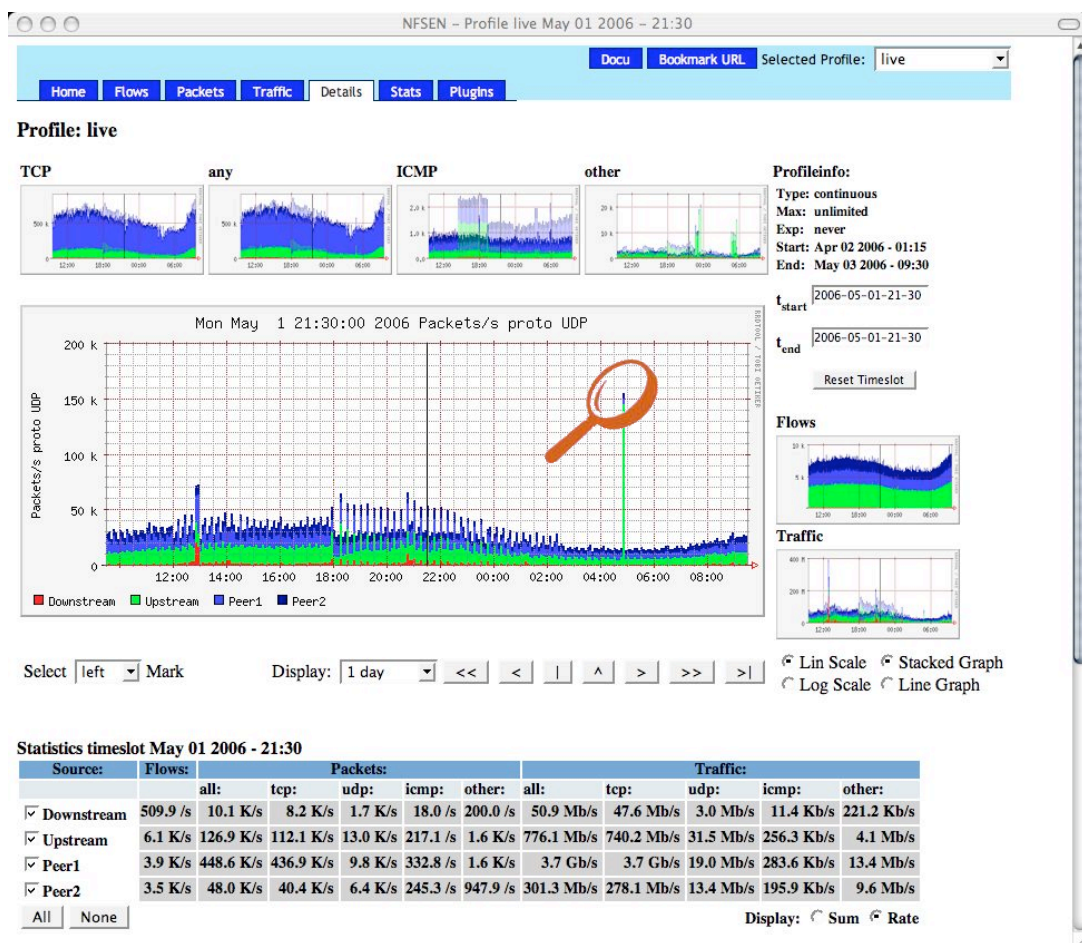
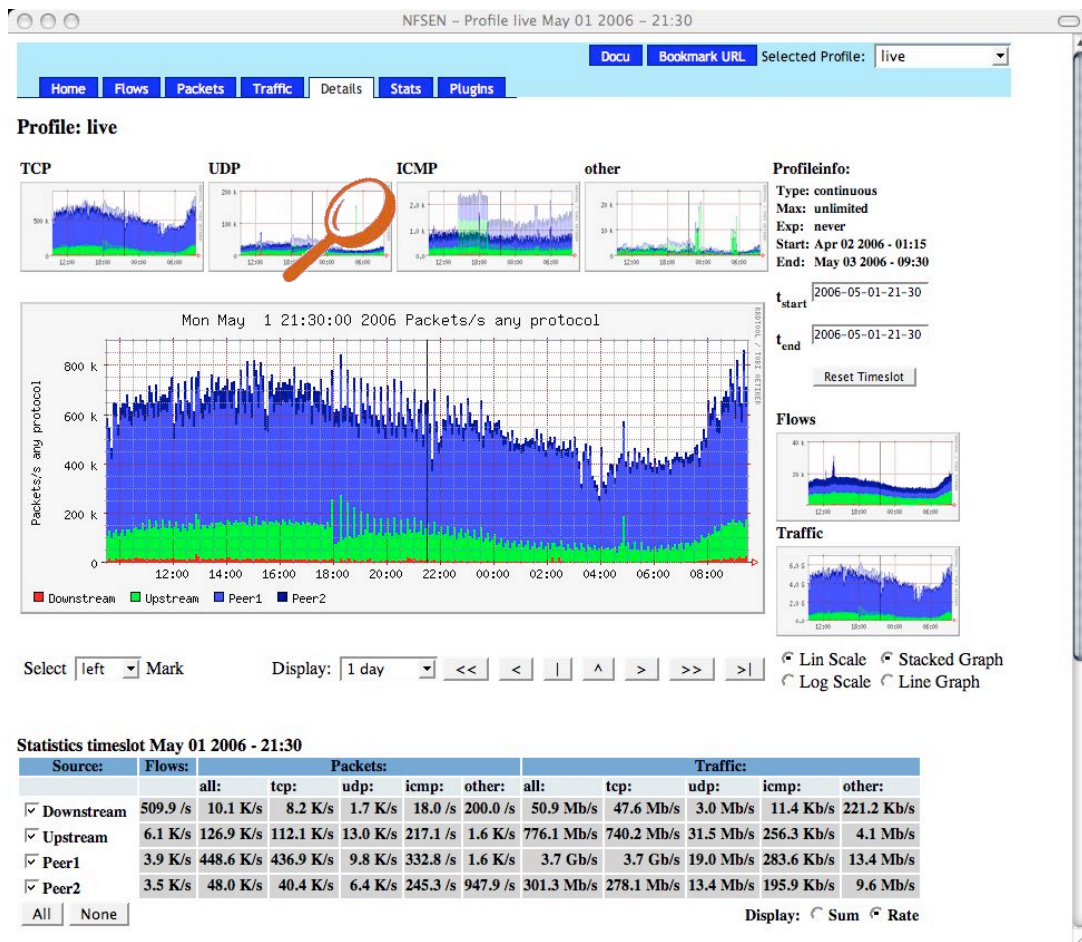


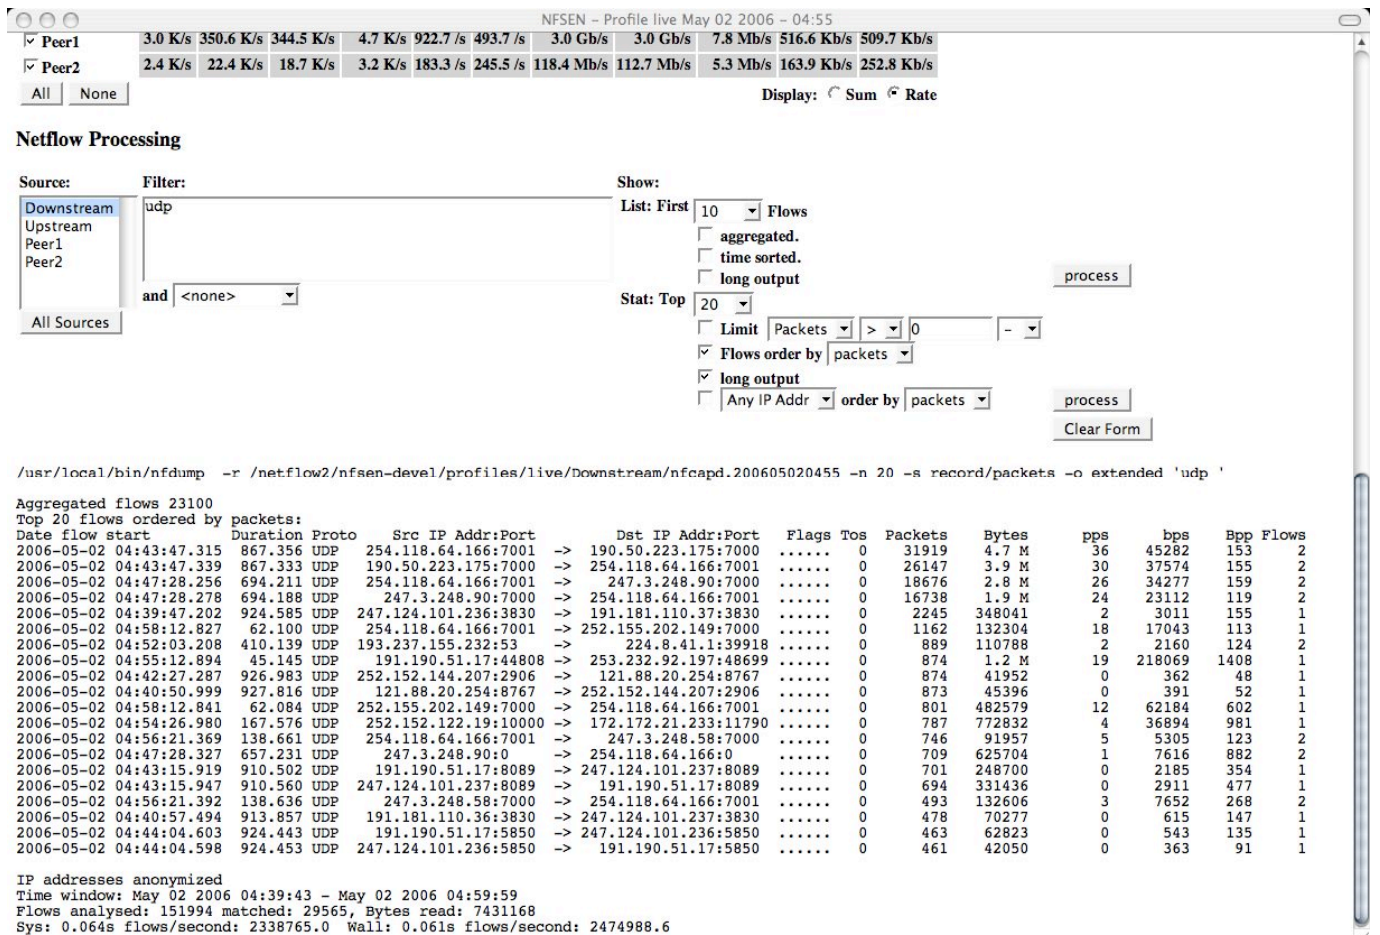
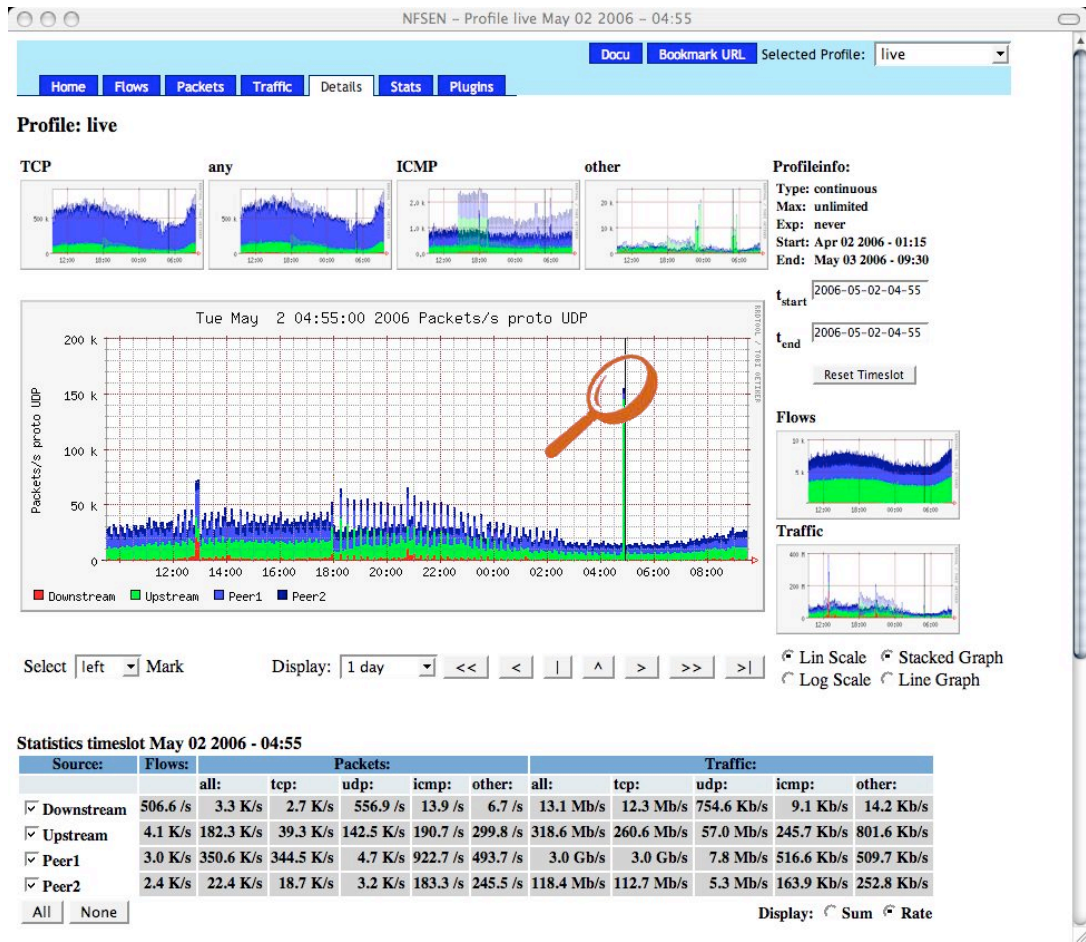
Details ⇒ Flows



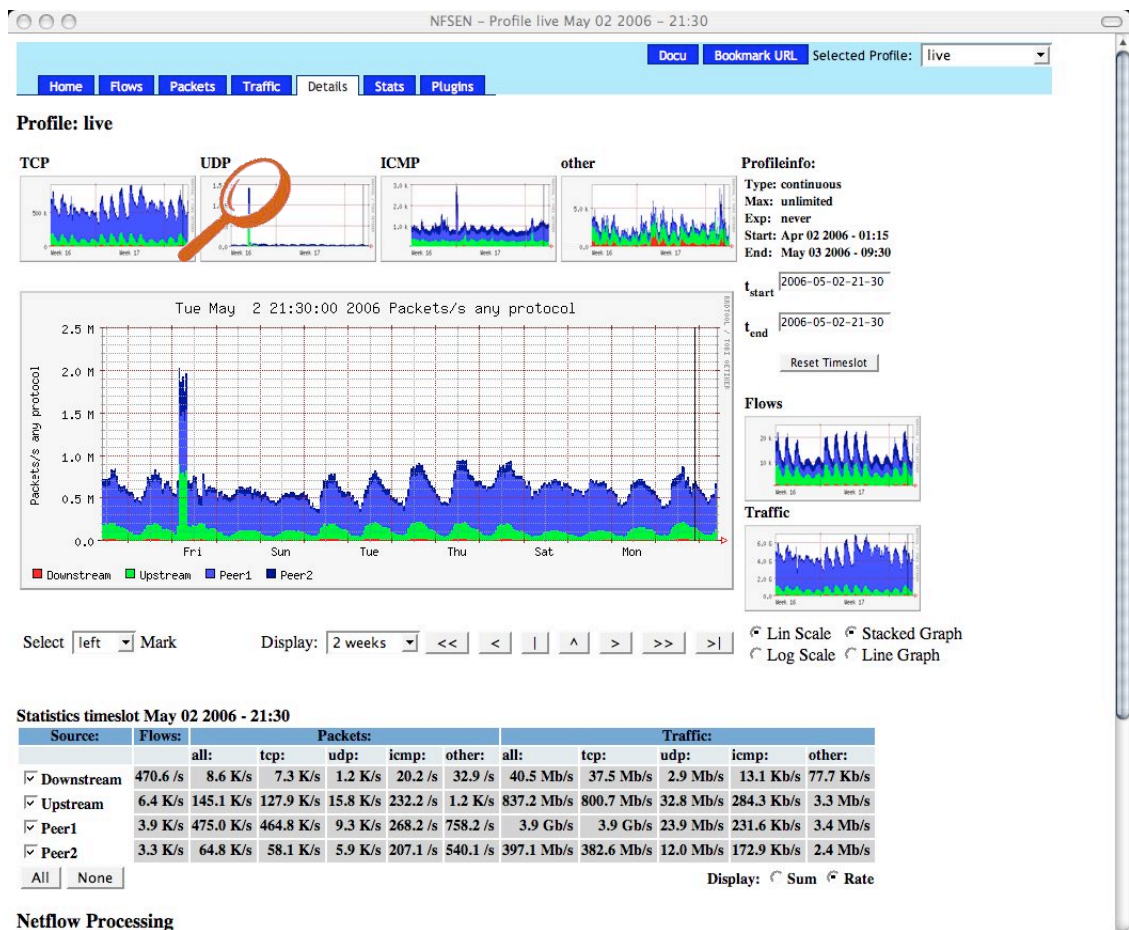
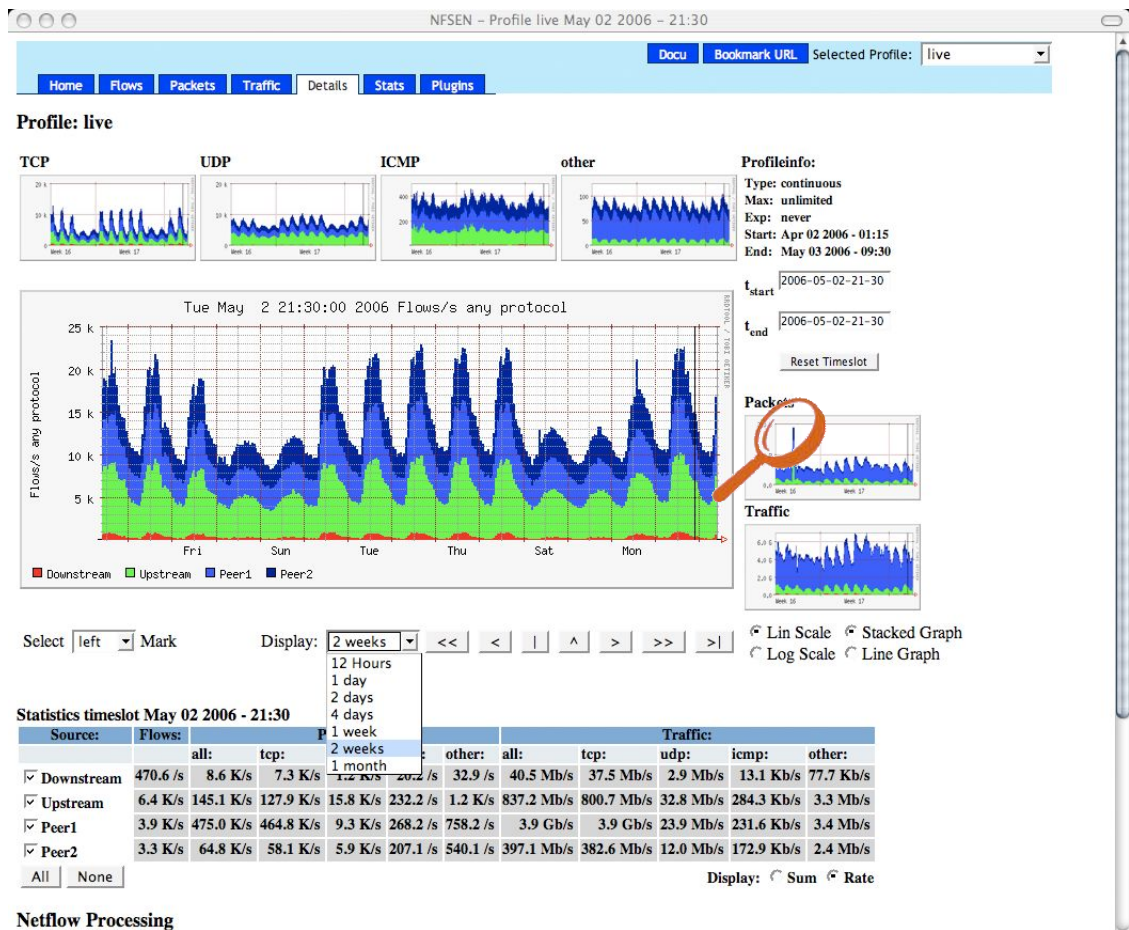


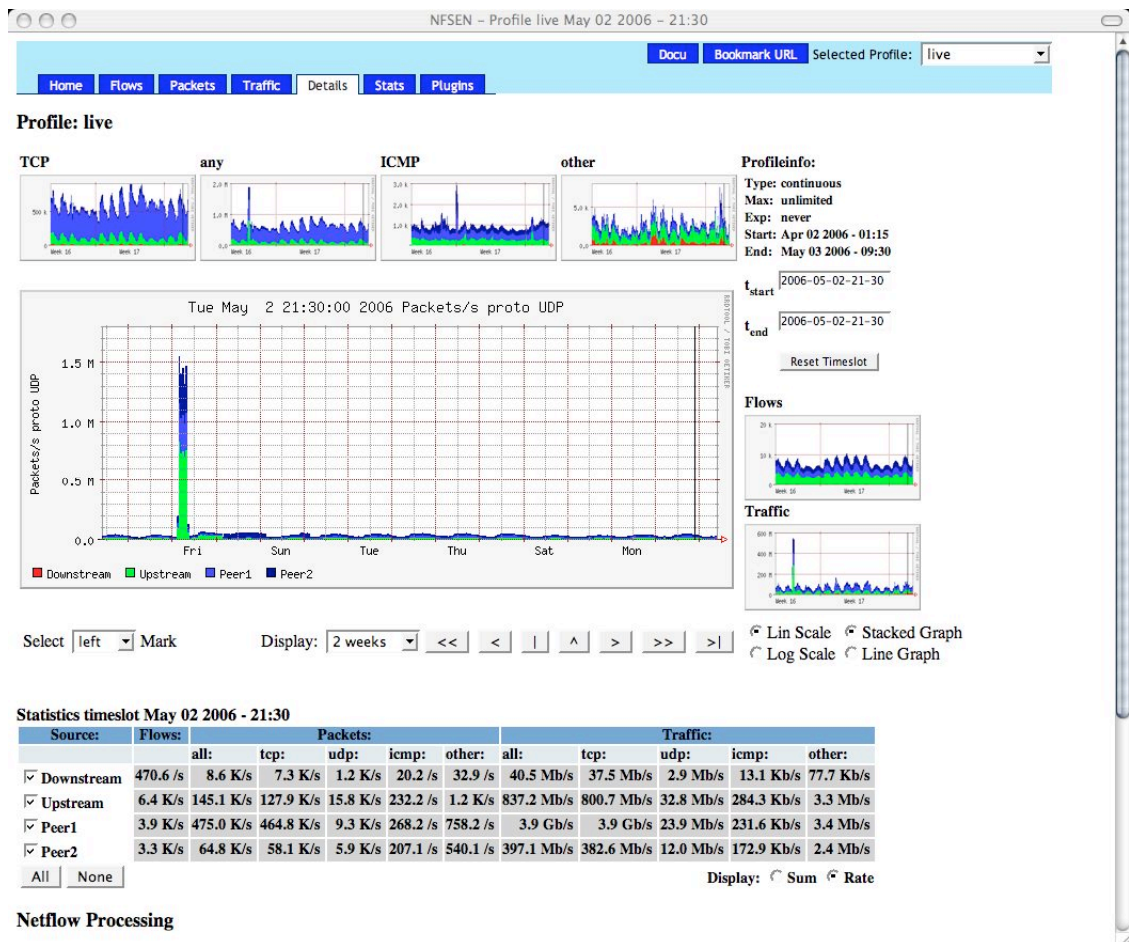




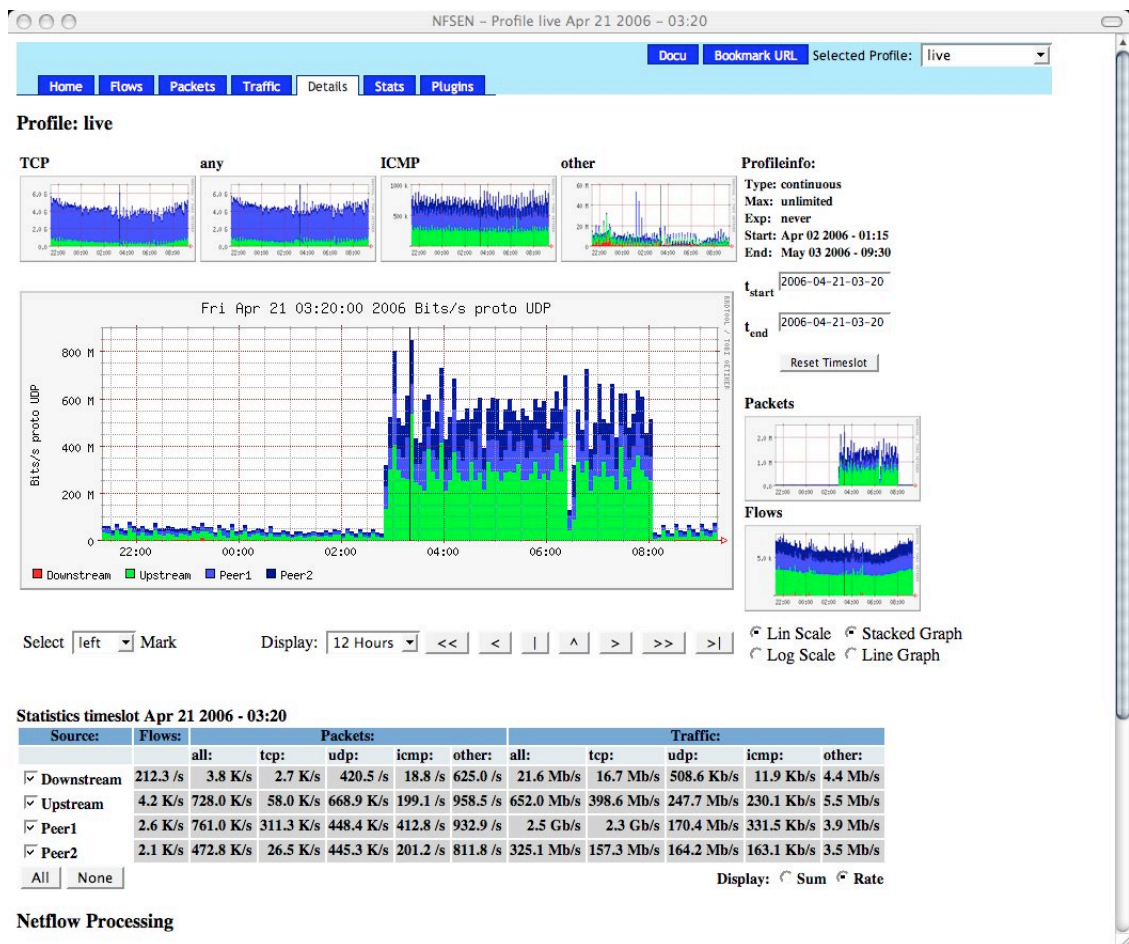
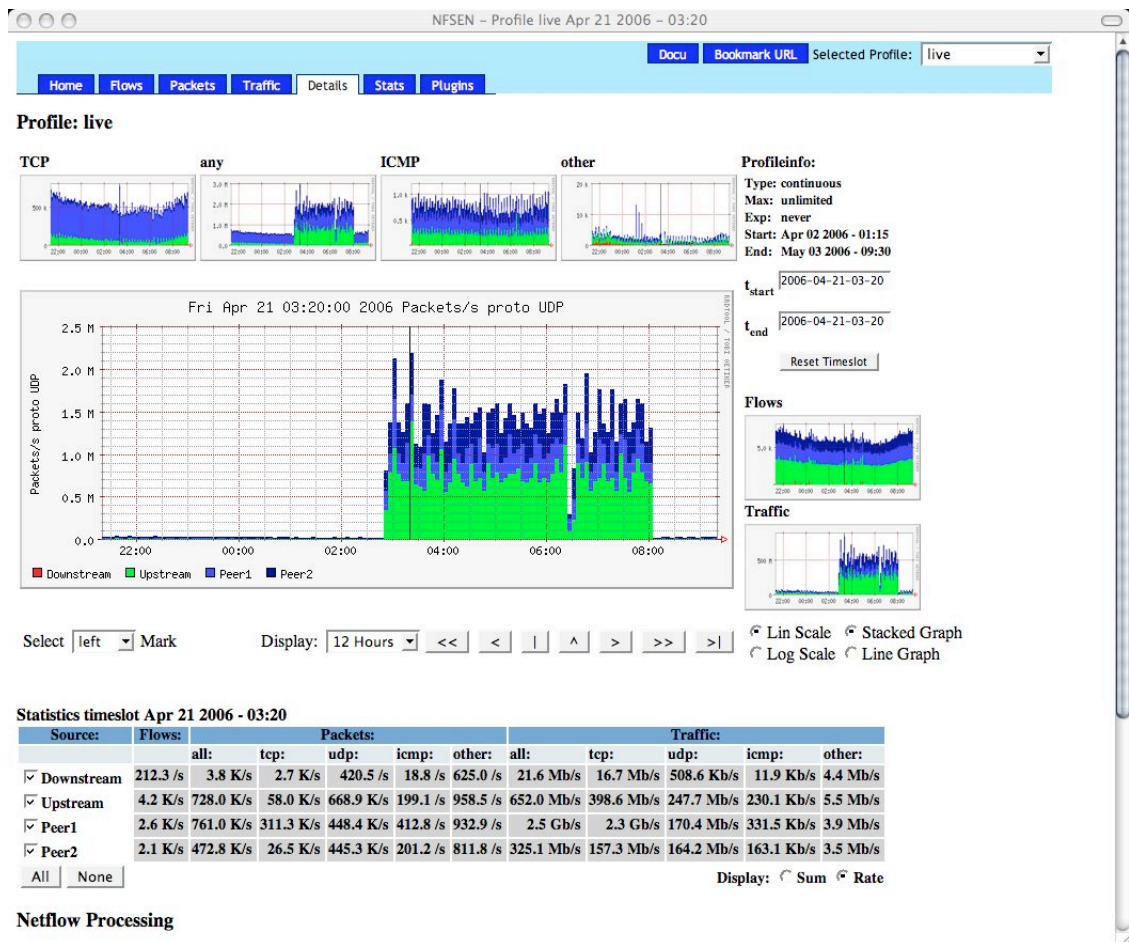












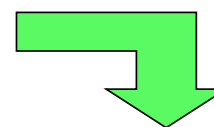




The Swiss Education & Research Network



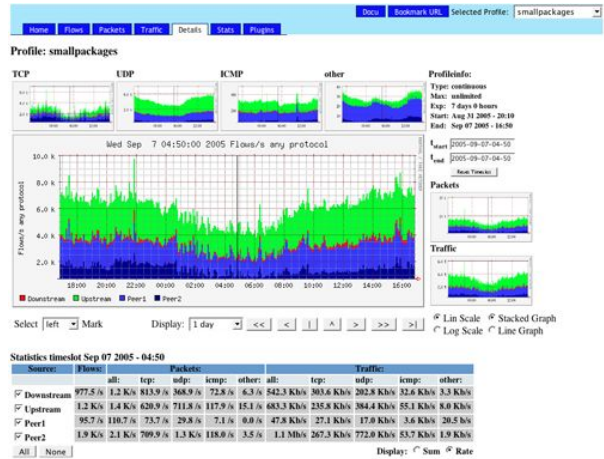
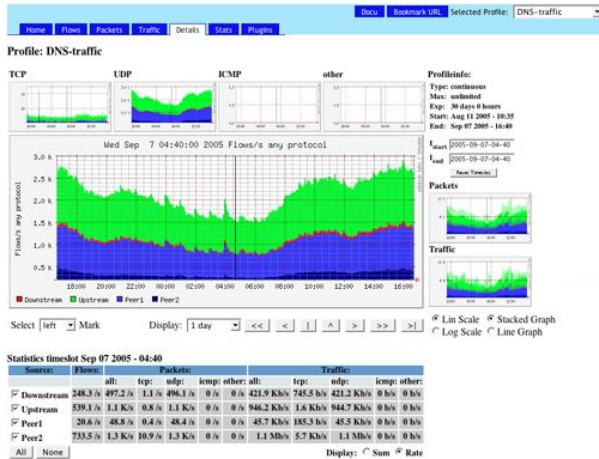
## The Swiss Education &amp; Research Network



## Example Profiles:

Filter: '( udp or tcp ) and port 53'

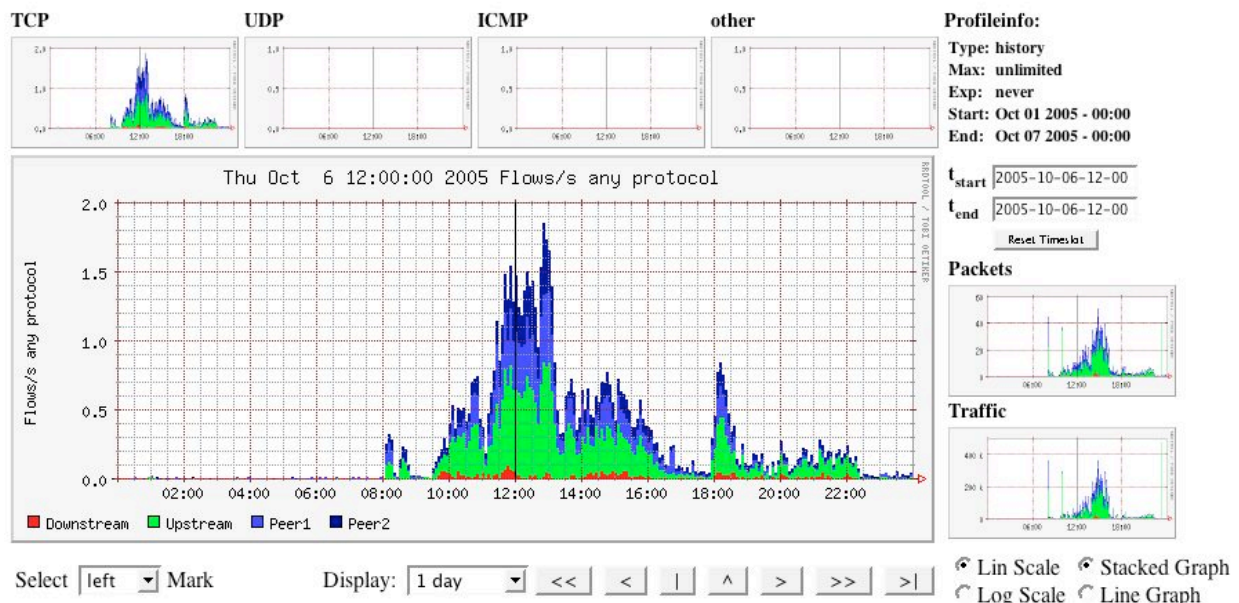
Filter: 'bytes < 100'



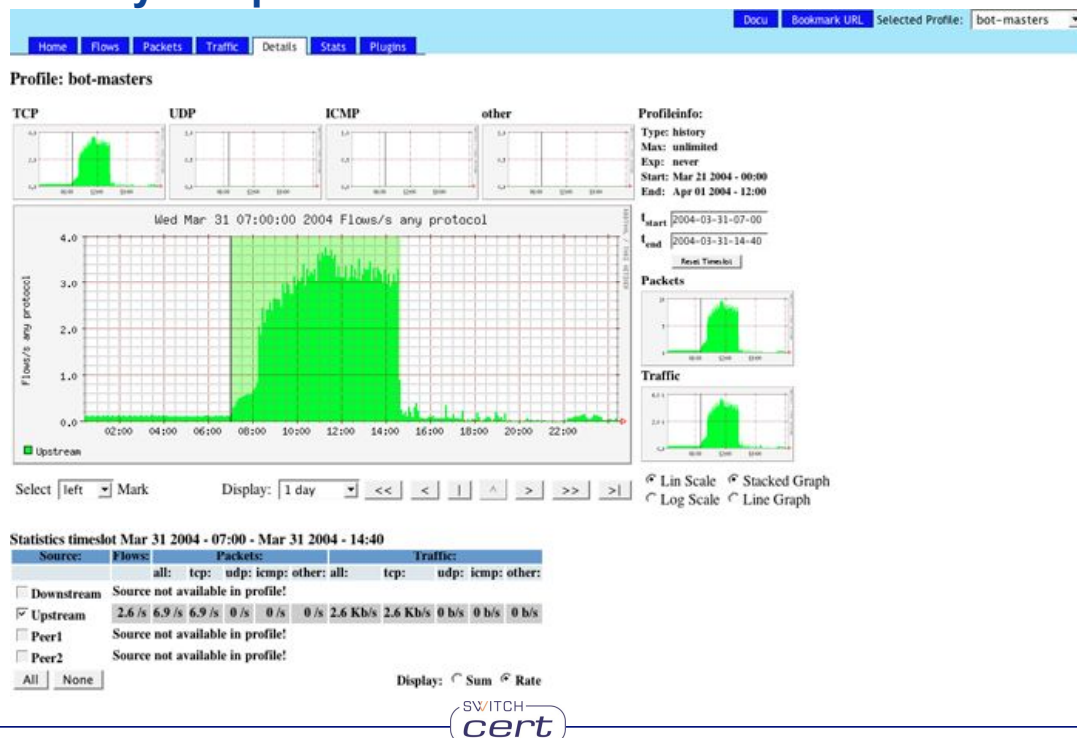
Filters may be as complex as the the filter syntax of nfdump allows.  
Example: '((src net 172.16/16 and src port > 1024 ) or dst host 192.168.16.17 and dst port 80) and packets > 1000 and pps > 150'

## NfSen/NFDUMP

## SoberR: 'tcp and dst port 587'



## Incident analysis - profile a hacked host:



## Plugins - what for?

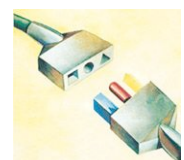
- Monitoring and alerting.
- Track for known botnet masters and send notifications.
- Track for possible scanners or DoS attacks, not necessarily visible in the graph.
- Port Tracking.

## Backend Plugins are:

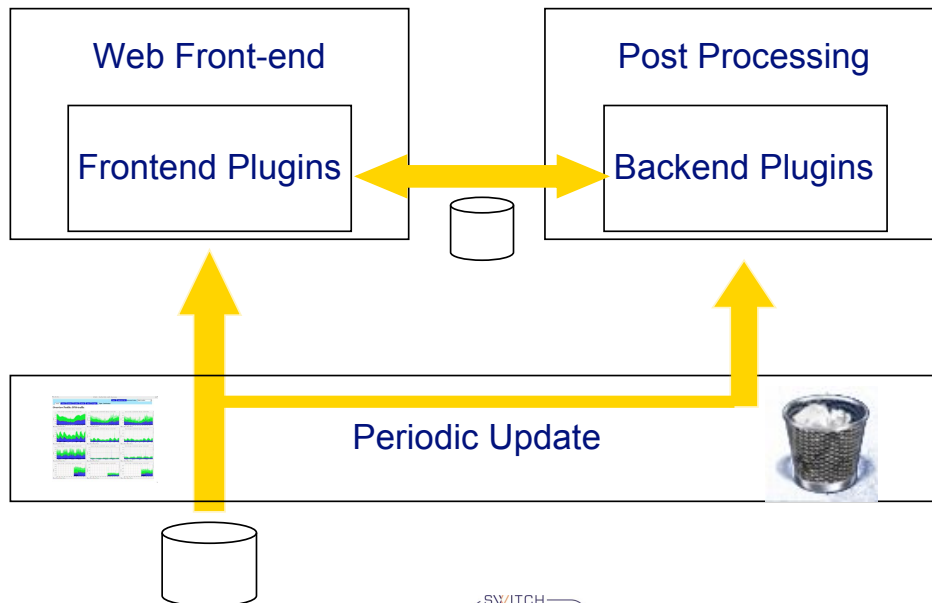
- Simple Perl modules hooked into the NfSen backend.
- Automatically called at regular 5 Min intervals.

## Frontend Plugins are:

- Simple PHP modules hooked into NfSen frontend.
- Called by selecting the tab.

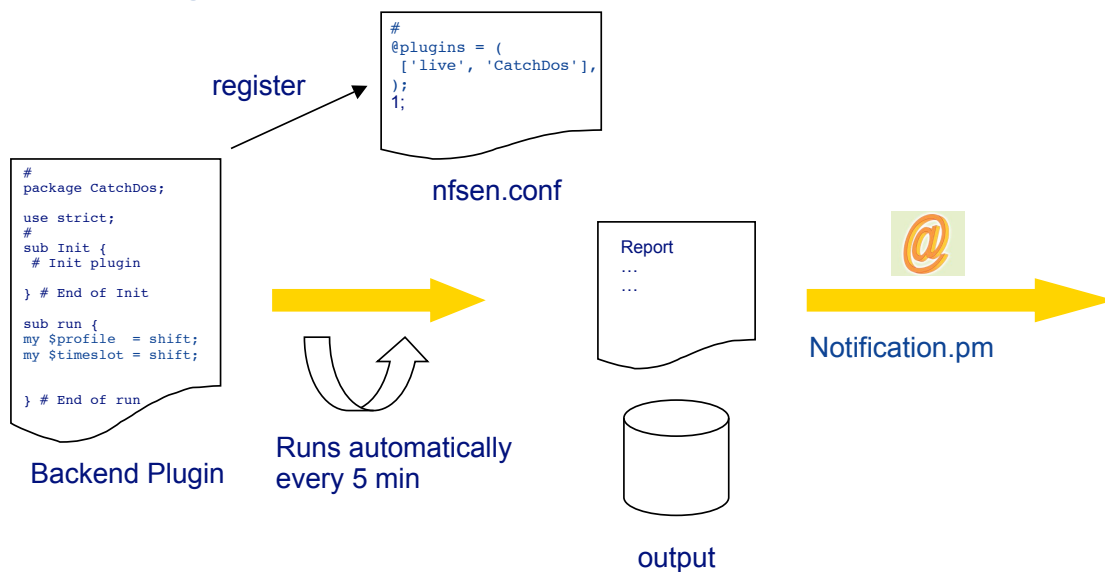


## NfSen Plugins:



# nfdump and NfSen

## Backend Plugins:



## Writing Backend Plugins:

- Select a plugin name: **MyPlugin**
- Create a Perl module named **MyPlugin.pm**
- Write your code.
- Try/debug your module offline using **\$BINDIR/testPlugin:**  
**./testPlugin -p MyPlugin -P live -t 200603140800**
- Store the file **MyPlugin.pm** in the directory **\$BACKEND\_PLUGINDIR**  
( e.g. **/data/nfsen/plugins** )
- Register the plugin in **nfsen.conf** for the profiles in question:

```
@plugins = ( # profile    # modul
             # [ '*',    'demoplugin' ],
             [ 'live',   'MyPlugin'],
);
```

```
#!/usr/bin/perl

package MyPlugin;

use strict;
# Periodic function
# input:  profilename
#        timeslot. Format yyyyymmddHHMM e.g. 200503031200
sub run {
    my $profile = shift;
    my $timeslot = shift;
    # Called at every cycle
    # Your code goes here

}

sub Init {
    # Do module init staff here

    # return 1 on success - module successfully loaded
    # return 0 on failure - module disabled
    return 1;
}

sub BEGIN {
    # Standard BEGIN Perl function - See Perl documentation
    # Called on loading the module
}

sub END {
    # Standard END Perl function - See Perl documentation
    # Called on unloading the module
}

1;
```



## Example Candidates for scanning activities:

```
...
#
#
# Define a nice filter:
# We like to see flows containing more than xxx packets
my $limit = 6000 ;
my $nf_filter = 'duration < 3500 and packets < 3 and bpp < 100 and src as 559';

# Periodic function
# input:  profilename
#        timeslot. Format yyyyymmddHHMM e.g. 200503031200
sub run {
    my $profile = shift;
    my $timeslot = shift;

    syslog('debug', "CatchScanners run: Profile: $profile, Time: $timeslot");

    my %profileinfo = NfSen::ReadProfile($profile);
    my $netflow_sources = $profileinfo{'sourcelist'};

    #
    # process all sources of this profile at once
    my @output = `nfdump -M $PROFILEDATADIR/$profile/$netflow_sources -r nfcapd.$timeslot -a -A srcip,dstport -l $limit -f $nf_filter`;

    #
    # Process the output and notify the duty team
}
...
```

( IP addresses anonymised )

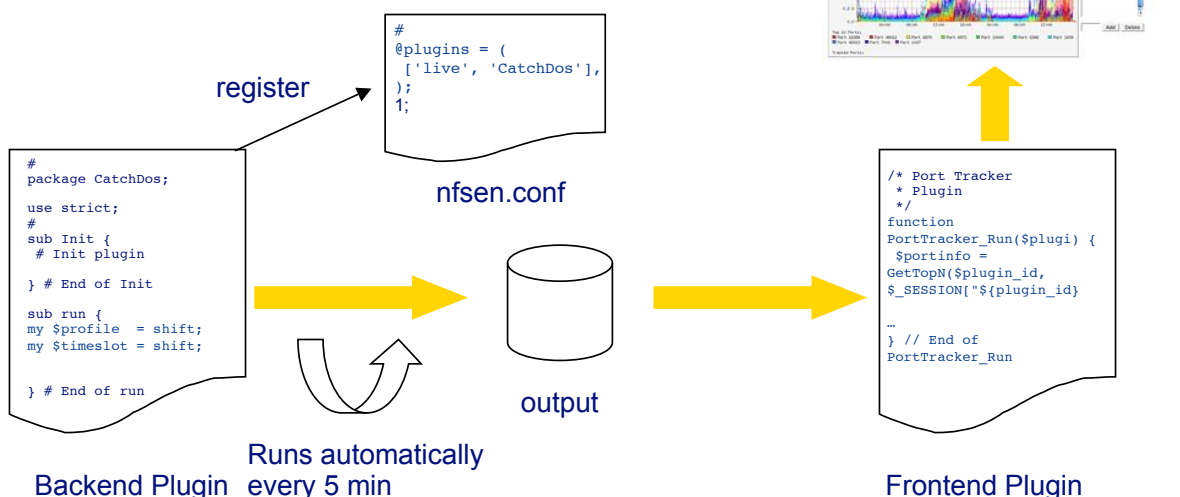
2006 © SWITCH



43

# nfdump and NfSen

## Backend / Frontend Plugins:



2006 © SWITCH



44

## Writing Frontend Plugins:

- Write the Backend plugin: **MyPlugin**
- Create a PHP module named **MyPlugin.php**
- Write your code.
  - Must have 2 well defined functions:
 

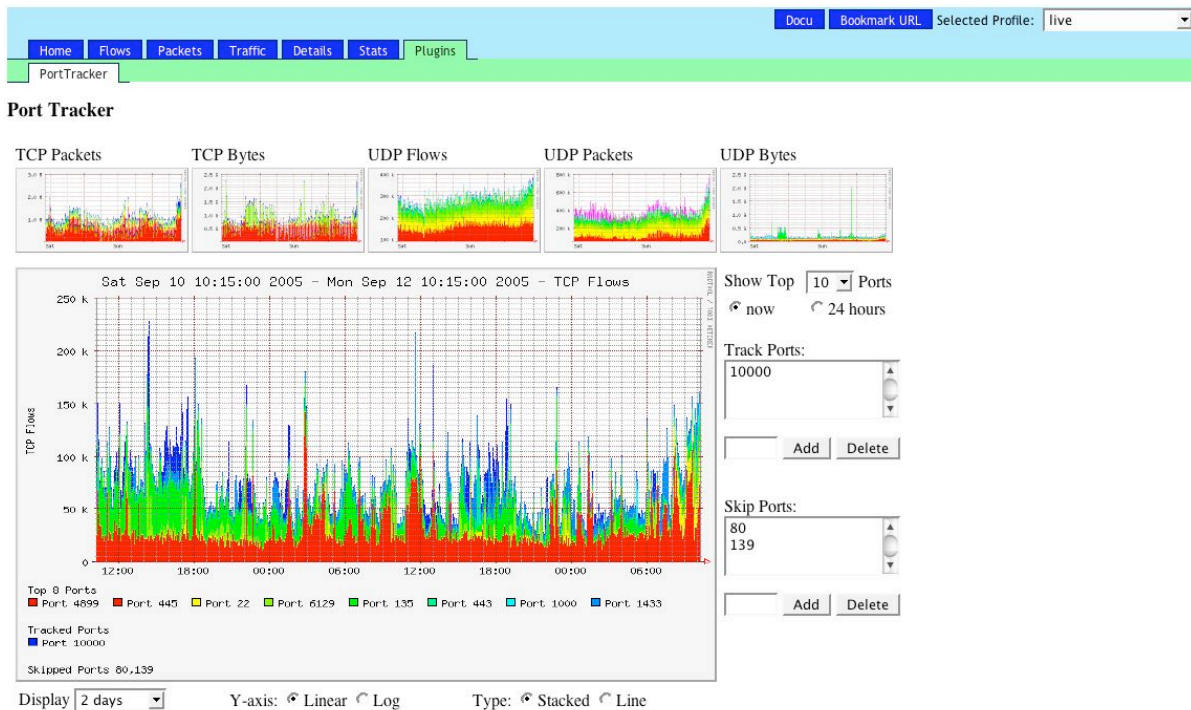
```
function <plugin_name>_ParseInput( $plugin_id )
function <plugin_name>_Run( $plugin_id )
```
  - Have each a unique plugin ID: **\$plugin\_id**
  - Run at any time the user selects the plugin.
  - Profile read only information available in **\$\_SESSION['profileinfo']**  
example: **\$\_SESSION['profileinfo']['name']**
- Store the file **MyPlugin.php** in the directory **\$FRONTEND\_PLUGINDIR** (e.g. **/var/www/htdocs/nfsen/plugins**)
- Reload **nfsen-run** background process: **\$BINDIR/nfsen reload**
- Check correct load of module in syslog file.

```
<?
/*
 * MyPlugin plugin
 */

// Required functions
/*
 * This function is called prior to any output to the web browser and is intended
 * for the plugin to parse possible form data. This function is called only, if this
 * plugin is selected in the plugins tab
 */
function MyPlugin_ParseInput( $plugin_id ) {
    if ( isset($_POST["${plugin_id}_variable"]) ) {
    } else {
        $_SESSION['warning'] = "Warning ...";
        $_SESSION['error'] = "Error ...";
    }
} // End of MyPlugin_ParseInput

/*
 * This function is called after the header with the navigation bar have been sent to the
 * browser. It's now up to this function what to display.
 * This function is called only, if this plugin is selected in the plugins tab
 */
function MyPlugin_Run( $plugin_id ) {
    $_SESSION["${plugin_id}_variable"] = ...
} // End of MyPlugin_Run

?>
```



Top 10 Statistics

Rank	TCP						UDP					
	Flows	Packets	Bytes	Flows	Packets	Bytes	Flows	Packets	Bytes	Flows	Packets	Bytes
1	80	176537	80	3278204	119	1142939174	1434	112023	53	252282	6970	89666470
2	4899	41763	119	807279	80	327540402	1027	88515	6970	140122	1434	45257351
3	445	31484	22	784231	40173	289081589	53	74135	1434	112024	1027	43540859
4	139	30747	1521	313884	4662	233088855	1026	57346	1027	88696	1026	27285365
5	22	26020	4662	276854	22	223138454	123	17664	1026	57501	0	23636673

## nfdump and NfSen

SWITCH

The Swiss Education & Research Network

### Planned Plugin: Host behaviour based worm detection:

Result of a PhD network security research work in the context of the DDoSVax project at **Swiss Federal Institute of Technology Zurich**:  
<http://www.tik.ee.ethz.ch/~ddosvax/>

**Idea: Infected hosts show a different behaviour and can be put into different classes:**

#### „Traffic“ class:

Worm infected hosts tend to send considerably more traffic than they receive.

#### „Responder“ class:

If many more hosts start to answer requests, they probably are victims of a worm.

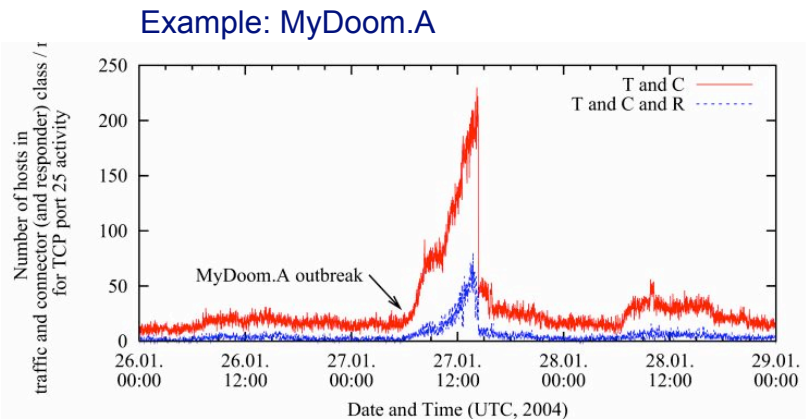
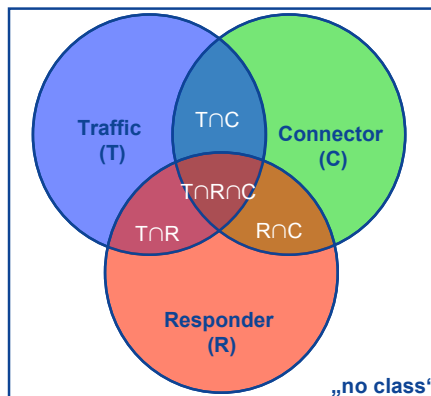
#### „Connector“ class:

Worm infected hosts typically open many connections.





## DDoS Vax : Host behaviour based worm detection:

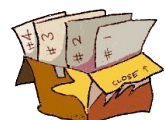


Most interesting for worm detection are cardinalities of class combinations.

# nfdump and NfSen

## Figures @ SWITCH:

- Server: 2 x 3GHz 2GB Ram. Debian Linux Kernel 2.6.10
- 3TB ( 2TB + 1TB ) AXUS Disk Raid
- XFS file system.
- Gigabit Ethernet interfaces.
- 5min workload avg. ca. 5%.
- 25GB Netflow data / day.
- About 41 days of netflow data available.



## Next Steps - Todo list - a lot of work:

### NfSen:

- more plugins ..
- Improved profiles.
- Improved interface.
- ...

### nfdump:

- **Related filters: 'Worm Footprint Tracking'**  
`first { dst ip <A> dst port 445 bytes > 600 }`  
`then { src ip <A> and dst ip 172.16.17.18 and dst port 80 }`
- Include more v9/sflow data in capture files.
- Realtime flow processing.
- Nesting directory levels for data organisation
- ...



## Summary:

- **Powerful and flexible tools for all sort of netflow tasks.**
  - Network monitoring.
  - Incident Handling.
  - All sort of tracking ...
- **Open Source under BSD License.**
- **Cmd line tool: nfdump**
  - Written in C. Runs on most \*nix.
  - Tested on Linux Kernel 2.4.\* and 2.6.\*, FreeBSD, OpenBSD, Solaris.
  - Available at <http://nfdump.sourceforge.net>
- **Web based frontend: NfSen**
  - Written in PHP and Perl.
  - Extendable using plugins.
  - Available at <http://nfsen.sourceforge.net>





Thank you for your  
attention.  
Any Questions?