

# Scalable Network Management and Monitoring

EUNICE 2010

Arne Øslebø  
arne.oslebo@uninett.no





# Ritz/zino

RITZ (nsa.uninett.no) - Remote interface to zino

OpState	AdmState	Age	Downtime	Router	Port	Description
reachable	open	0d 8:44:32.00	0d 0:01:00.00	nlabreistein-gw		
bgp established	working	4d 10:10:49.00		oslo-gw	AS 29517 193.156.90.55	(peer was reset (now up))
port up	closed	0d 12:45:33.00	0d 1:38:57.00	nokut-sw	1	(Lokal link, nokut-sw nokut-FW)
port up	closed	0d 15:15:47.00	0d 0:00:00.00	stolav-gw4	GigabitEthernet2/1.1	(TDCSong 30022048, stolav-nokut-mng, nokut-sw)
port up	closed	0d 15:15:47.00	0d 0:00:00.00	stolav-gw4	GigabitEthernet2/1.2	(TDCSong 30022048, stolav-nokut, nokut-sw)
reachable	closed	0d 15:14:27.00	0d 2:21:34.00	nokut-sw		
port up	closed	0d 15:19:03.00	0d 2:25:53.00	stolav-gw4	GigabitEthernet2/1	(TDCSong 30022048, trunk.nokut-sw)
bgp established	open	0d 13:19:26.00		teknobyen-gw2	AS 64527 158.38.0.218	(peer was reset (now up))
port up	closed	1d 0:58:33.00	0d 0:00:00.00	trd-gw1	ge-1/0/2.0	()
port up	closed	1d 0:57:39.00	0d 0:00:00.00	stjordal-gw	ge-0/0/0.0	()
port up	closed	1d 1:00:59.00	0d 8:28:10.00	trd-gw1	ge-1/0/2	(Stjordalsfiber, trd-stjordal, stjordal-gw)
port up	closed	1d 1:00:29.00	0d 8:27:40.00	stjordal-gw	ge-0/0/0	(Stjordalsfiber, stjordal-trondheim, trd-gw)
bgp established	ignored	77d 20:15:52.00		oslo-gw8	AS 64531 32.1.6.48	(peer was reset (now up))
port up	working	0d 19:36:50.00	0d 0:46:03.00	tromso-gw7	GigabitEthernet2/3	(FringTromso, tromso-krognesvn33, krognesvn33-gw)
reachable	closed	0d 19:38:29.00	0d 0:42:51.00	krognesvn33-gw	hitos.no	
reachable	closed	0d 19:36:15.00	0d 0:31:23.00	krognesvn33-gw		
port up	closed	0d 19:37:32.00	0d 0:46:07.00	bjerkaker-gw	GigabitEthernet1/1/2	(FringTromso, bjerkaker-krognesvn33, krognesvn33-gw)
port down	working	0d 19:25:20.00	0d 19:25:20.00	svalbard-gw2	GigabitEthernet1/1/2	
port down	working	0d 19:42:54.00	0d 19:42:54.00	sarpsborg-gw	GigabitEthernet1/1/2	
bgp established	open	0d 21:16:06.00		oslo-gw	AS 16065 193.156.90.38	(peer was reset (now up))
port up	open	0d 22:41:43.00	0d 0:01:50.00	stakkevollv23-gw	FastEthernet0	(lokal trunk, stakkevollv23-veths)
port up	open	0d 22:47:51.00	0d 0:01:39.00	sandnes-gw	FastEthernet0	(lokal trunk, sandnes-veths)
port up	open	0d 23:17:01.00	0d 0:00:15.00	stolav-gw1	POS12/0	(BaneTele 11310006, oslo-bergen2, bergen-gw)
port adminDown	open	1d 20:09:14.00	0d 0:00:24.00	oslo-gw7	TenGigabitEthernet7/8	(oslo-gw4 te7/8)

# NAV



## Network Administration Visualized

Home

Preferences

Toolbox

Useradmin

Userinfo

Logout arneos

### Velkommen til NAV for UNINETT!

Dersom du mener at du skal ha mer tilgang i NAV enn hva du har nå, ring [drift@uninett.no](mailto:drift@uninett.no).

#### Status Now

3 IP devices down, 0 in shadow

Sysname	IP	Down since	Downtime
uis-gsw.uis.no	128.39.47.122	2010-06-27 00:42	10 hours
weathergoose.uninett.no	158.38.129.146	2010-05-15 04:12	43 days
stokmarknes-gw.uninett.no	158.39.0.48	2010-03-24 23:35	94 days

[Status page](#)

#### NAV links

[OpenStreetMap](#)  
[About NAV](#)  
[Statistics](#)  
[Machine Tracker](#)  
[Reports](#)

#### External links

[UNINETT](#)  
[Ansatt](#)  
[Drift](#)  
[GigaCampus](#)  
[Målepåler](#)  
[Systemdrift](#)  
[Samson Wiki](#)  
[Slashdot](#)  
[VK brukerdok](#)

#### Contact information

UNINETT  
Abels gate 5 - Teknobyen  
NO-7465 Trondheim  
Contact the NAV administrator:  
[drift@uninett.no](mailto:drift@uninett.no)

4

# NAV – autodetecting topology


Physical topology

- ARP
- CAM
- CDP
















Vlan topologi

- trunks
- STP

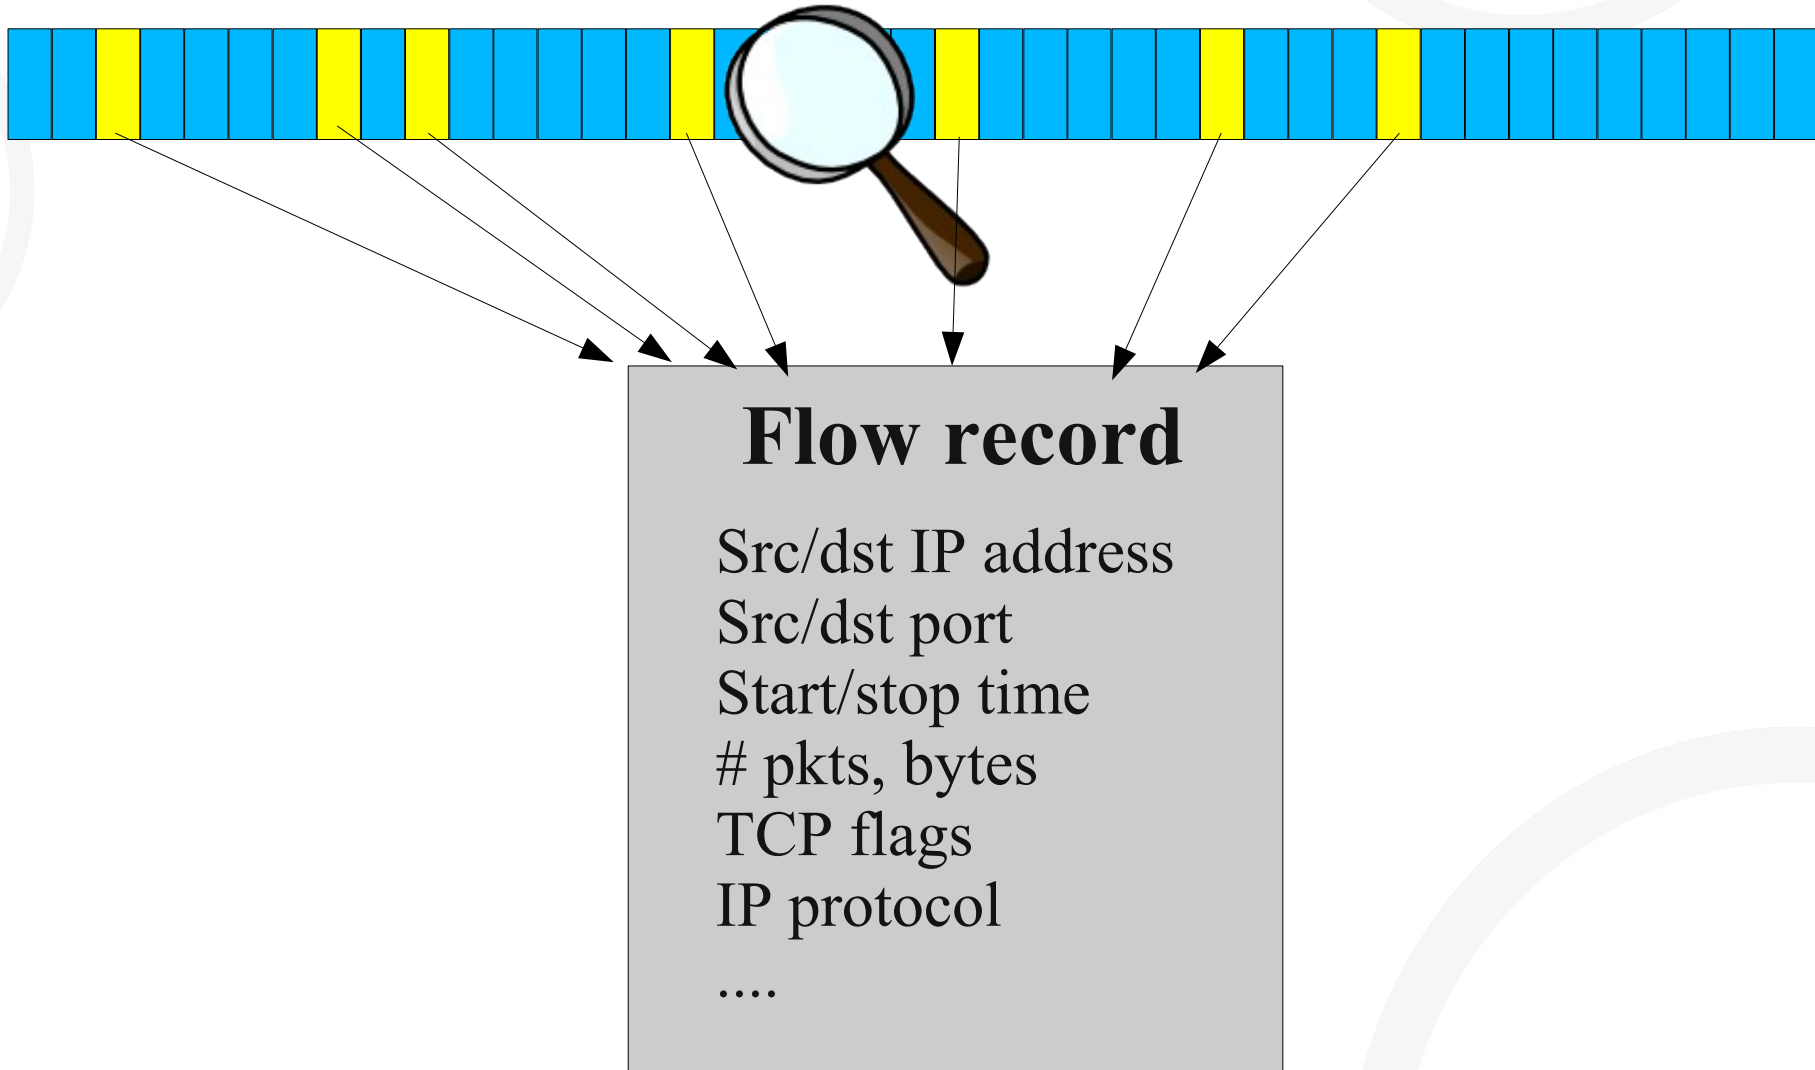
- Trace equipment
- Block equipment



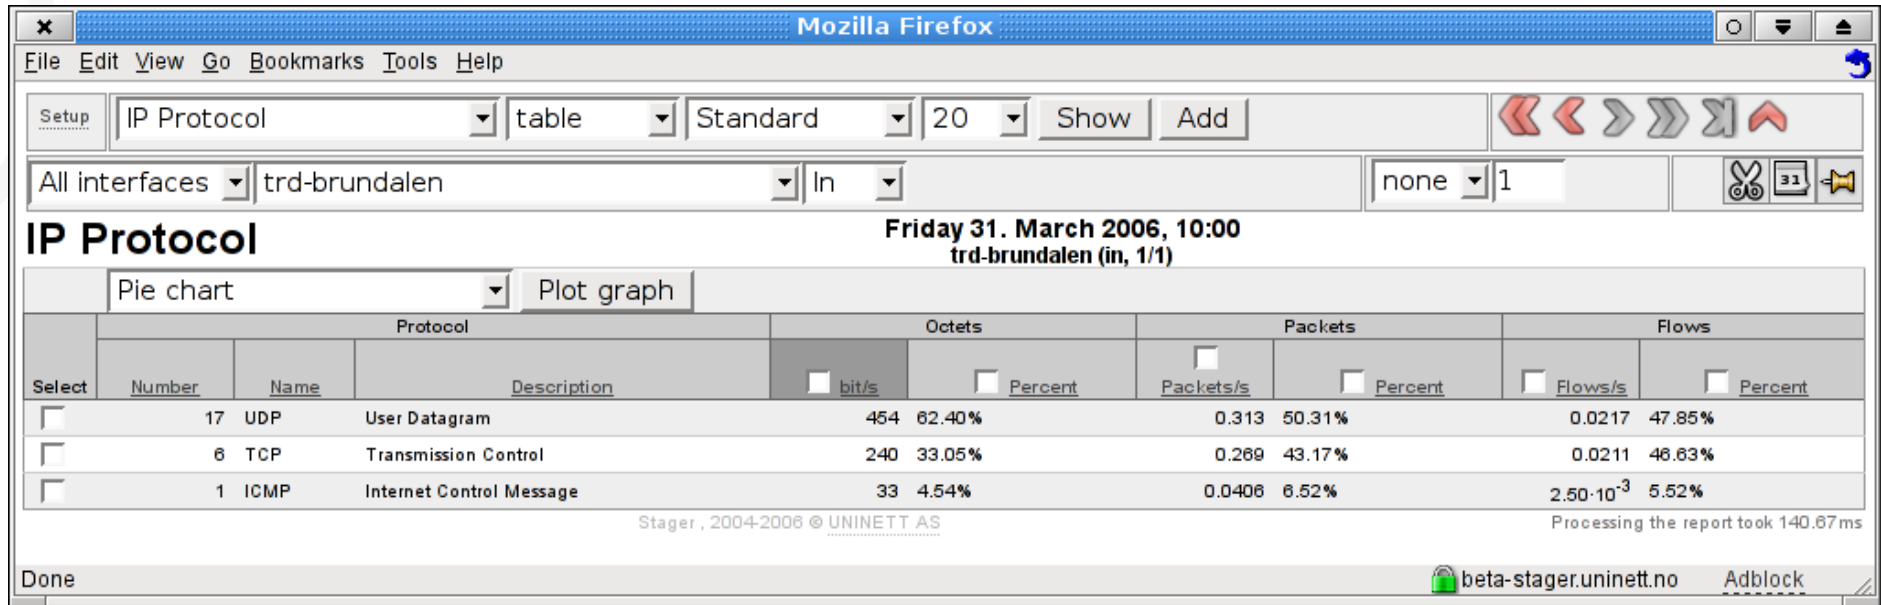
# NAV - toolbox

<b>Netmap</b>  <p>Netmap is a network weather map that displays the layer 3 and layer 2 topology of your network. Current traffic conditions are shown using colored links.</p>	<b>Status</b>  <p>What's going on? Is all network equipment working properly? Click here to find out!</p>	<b>Network Explorer</b>  <p>Explore the network topology as it expands from each router port.</p>
<b>Report</b>  <p>The highly customizable report generator displays information from the NAV-database.</p>	<b>IP Device Info</b>  <p>Displays all recorded information about a given IP device or address. If the device is a NAV registered switch or router a port view with details on status and activity is also shown.</p>	<b>Statistics</b>  <p>Browse statistical data collected by NAV and Cricket. Use Cricket to look at graphs for individual devices and interfaces, and use NAV's ranked statistics to uncover load problems.</p>
<b>Machine Tracker</b>  <p>This tool lets you find MAC addresses that are or have been connected to a given switch port. You can also find the switch port used by a MAC or IP address in a given timerange.</p>	<b>Layer 2 Traceroute</b>  <p>Trace layer 2 network paths between hosts.</p>	<b>Syslog Analyzer</b>  <p>View syslog messages from (Cisco) network equipment.</p>
<b>Hobbit</b>  <p>Hobbit integrates with NAV to monitor your servers.</p>	<b>nfsen</b>  <p>Front end for the nfdump tools</p>	<b>Arnold</b>  <p>Arnold is a switchport blocking tool. Use it to manually block switchports, or to run automated blocking raids. Requires SNMP write access to switches.</p>
<b>Seed Database</b>  <p>NAV does not auto-discover devices, you</p>	<b>Messages</b>  <p>Read and publish operational messages.</p>	<b>Alert Profiles</b>  <p>Setup your alert service subscriptions here.</p>

# NetFlow



# Stager



Setup IP Protocol table Standard 20 Show Add

All interfaces trd-brundalen In none 1

### IP Protocol

Friday 31. March 2006, 10:00  
trd-brundalen (in, 1/1)

Pie chart Plot graph

Select	Protocol			Octets		Packets		Flows	
	Number	Name	Description	bit/s	Percent	Packets/s	Percent	Flows/s	Percent
<input type="checkbox"/>	17	UDP	User Datagram	454	62.40%	0.313	50.31%	0.0217	47.85%
<input type="checkbox"/>	6	TCP	Transmission Control	240	33.05%	0.269	43.17%	0.0211	46.63%
<input type="checkbox"/>	1	ICMP	Internet Control Message	33	4.54%	0.0408	6.52%	2.50·10 <sup>-3</sup>	5.52%

Stager, 2004-2006 © UNINETT AS Processing the report took 140.67ms

Done beta-stager.uninett.no Adblock

# Overview report

Protocol   10

All interfaces  In  1

## Protocol

Thursday 15. October 2009, 21:00

All observation points (in)



Pie chart

Other

Plot graph

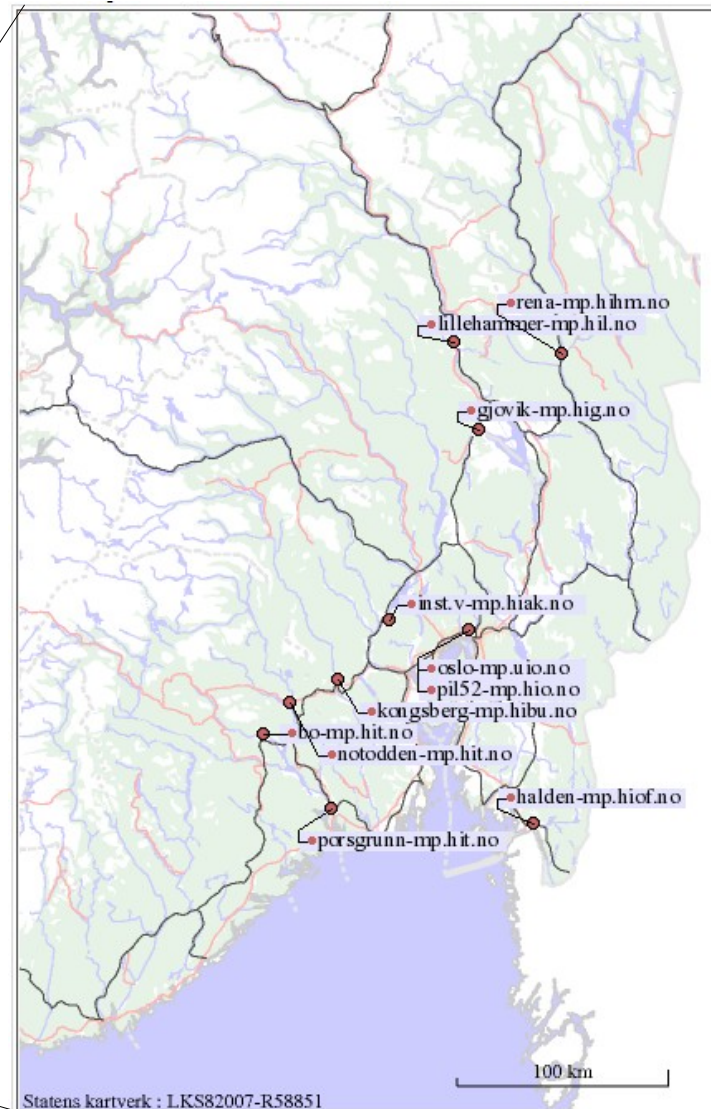
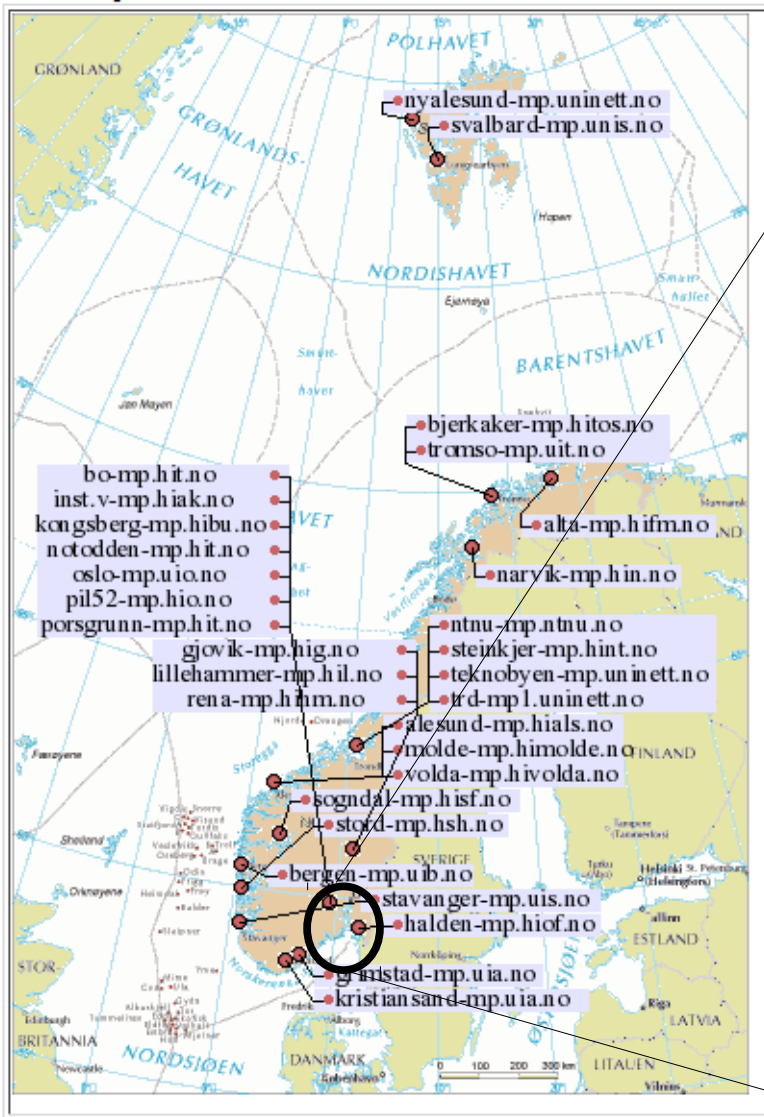
Octets - percent

Select	Observation Point	Octets - percent							
		<input checked="" type="checkbox"/> TCP	<input type="checkbox"/> UDP	<input type="checkbox"/> GRE	<input type="checkbox"/> ESP	<input type="checkbox"/> IPv6	<input type="checkbox"/> ICMP	<input type="checkbox"/> PIM	<input type="checkbox"/> IPv6-ICMP
<input type="checkbox"/>	elverum-gw.elverum-gw2-1	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%
<input type="checkbox"/>	oslo-trd2	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	0.00%	0.00%
<input type="checkbox"/>	trd-tromso2	0.00%	0.00%	0.00%	0.00%	0.00%	75.68%	24.32%	0.00%
<input type="checkbox"/>	elverum-gw.elverum-gw2-2	0.00%	36.36%	0.00%	0.00%	0.00%	63.64%	0.00%	0.00%
<input type="checkbox"/>	hoytek-niva	0.64%	65.75%	0.00%	0.00%	0.00%	33.61%	0.00%	0.00%
<input type="checkbox"/>	trd-niva	3.86%	66.29%	0.00%	0.00%	0.00%	29.84%	0.00%	0.00%
<input type="checkbox"/>	oslo-oslomsh	19.42%	53.90%	0.00%	0.00%	0.00%	26.68%	0.00%	0.00%
<input type="checkbox"/>	teknobyen2-teknobyen	5.89%	3.82%	0.00%	0.00%	0.00%	23.70%	66.59%	0.00%
<input type="checkbox"/>	nhh2-nhh	7.41%	69.95%	0.00%	0.00%	0.00%	22.64%	0.00%	0.00%
<input type="checkbox"/>	nhh-gw2.nhh-gw	7.41%	69.95%	0.00%	0.00%	0.00%	22.64%	0.00%	0.00%

Stager 4.0, 2004-2009 © UNINETT AS

Processing the report took 486.82ms

# Active and passive monitoring infrastructure



# Qstream

Setup Back Multicast TV table Overview 20 Show Add

All UNINETT In none 1

## Multicast TV

Friday 8. January 2010, 11:00  
UNINETT (in, 1/1)

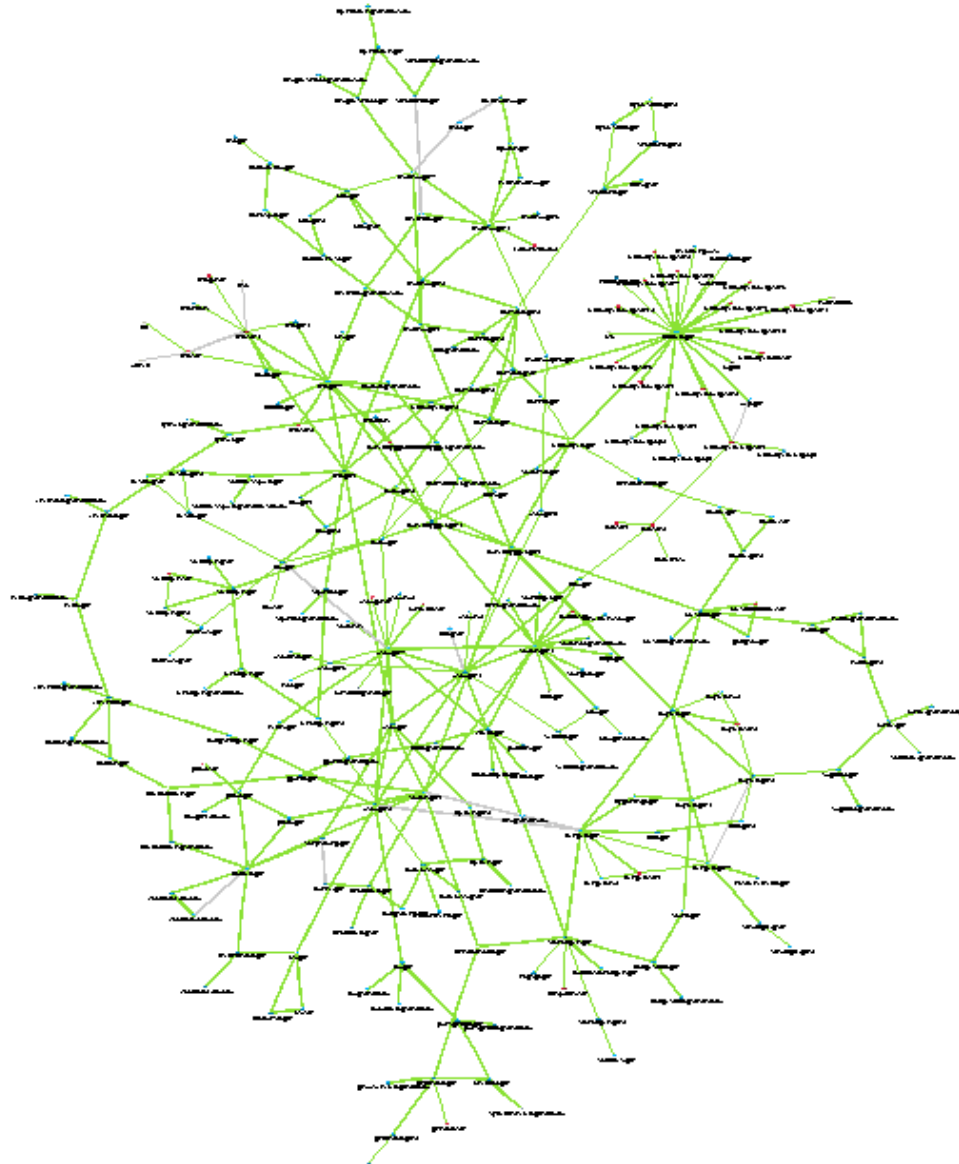
11

Line plot Other Plot graph

Select	Group			# of tests	Average statistics			
	Name	IP	Port		<input checked="" type="checkbox"/> Setup time	Packets/s	<input type="checkbox"/> Bitrate	<input type="checkbox"/> Gap
<input type="checkbox"/>	NRK3	239.193.0.3	1234	6	159ms	560	5.90M	1.88ms
<input type="checkbox"/>	NRK2	239.193.0.2	1234	6	81.5ms	583	6.13M	1.82ms
<input type="checkbox"/>	NRK1	239.193.0.1	1234	6	79.7ms	589	6.20M	1.72ms
<input type="checkbox"/>	Stortinget	224.67.1	1234	6	0.633ms	536	5.70M	1.90ms

Processing the report took 117.2ms

# SNMP scalability issues




# Proprietary SNMP MIBS


- BRIDGE-MIB, CISCO-C2900-MIB, CISCO-CDP-MIB, CISCO-ENTITY-ASSET-MIB, CISCO-ENVMON-MIB, CISCO-ES-STACK-MIB, CISCO-FLASH-MIB, CISCO-HSRP-MIB, CISCO-IETF-IP-MIB, CISCO-MEMORY-POOL-MIB, CISCO-STACK-MIB, CISCO-VLAN-MEMBERSHIP-MIB, CISCO-VTP-MIB, ENTITY-MIB, ESSWITCH.MIB, HP-MIB, IF-MIB, IP-MIB, IPV6-MIB, MAU-MIB, OLD-CISCO-CHASSIS-MIB, OLD-CISCO-CPU-MIB, OLD-CISCO-INTERFACES-MIB, OSPF-MIB, RFC1213, RFC1213-MIB, SECURITY-MIB, SNMPv2-MIB

# SMI: Weak data modeling


Node



List

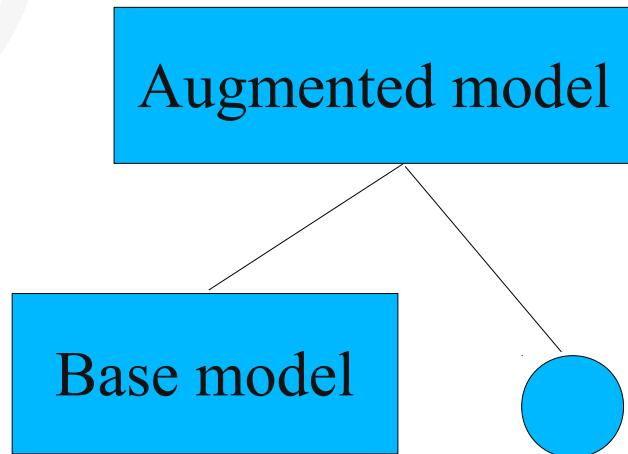


Table

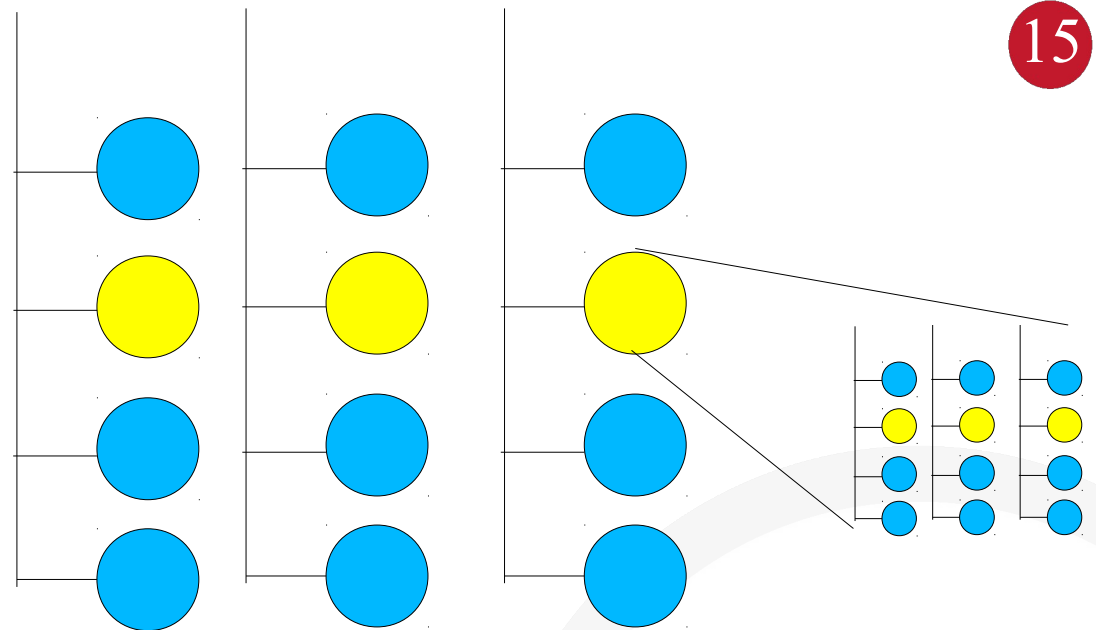


# YANG modeling language

## Augment




## Table in table



# Our NetFlow setup


- 34 routers
- 439 interfaces
- 1/100 sampling rate on most routers
- ~25Gb of compressed raw Netflow data every day
- ~ 1.8 billion flows per day
  - 40-130 million flows per hour

16

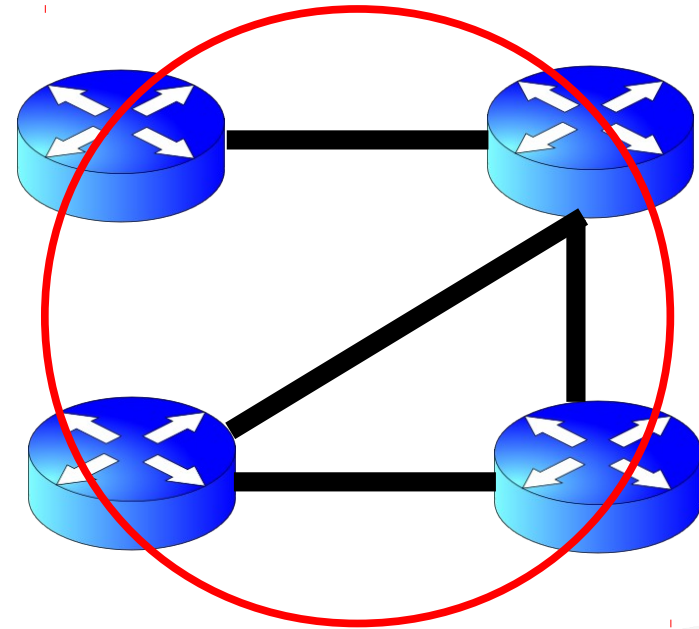



- Anycast
- Database
  - >1TB database size
  - >800 millions entries in a single table

# Anomaly detection




# Per link or network




# Datasets

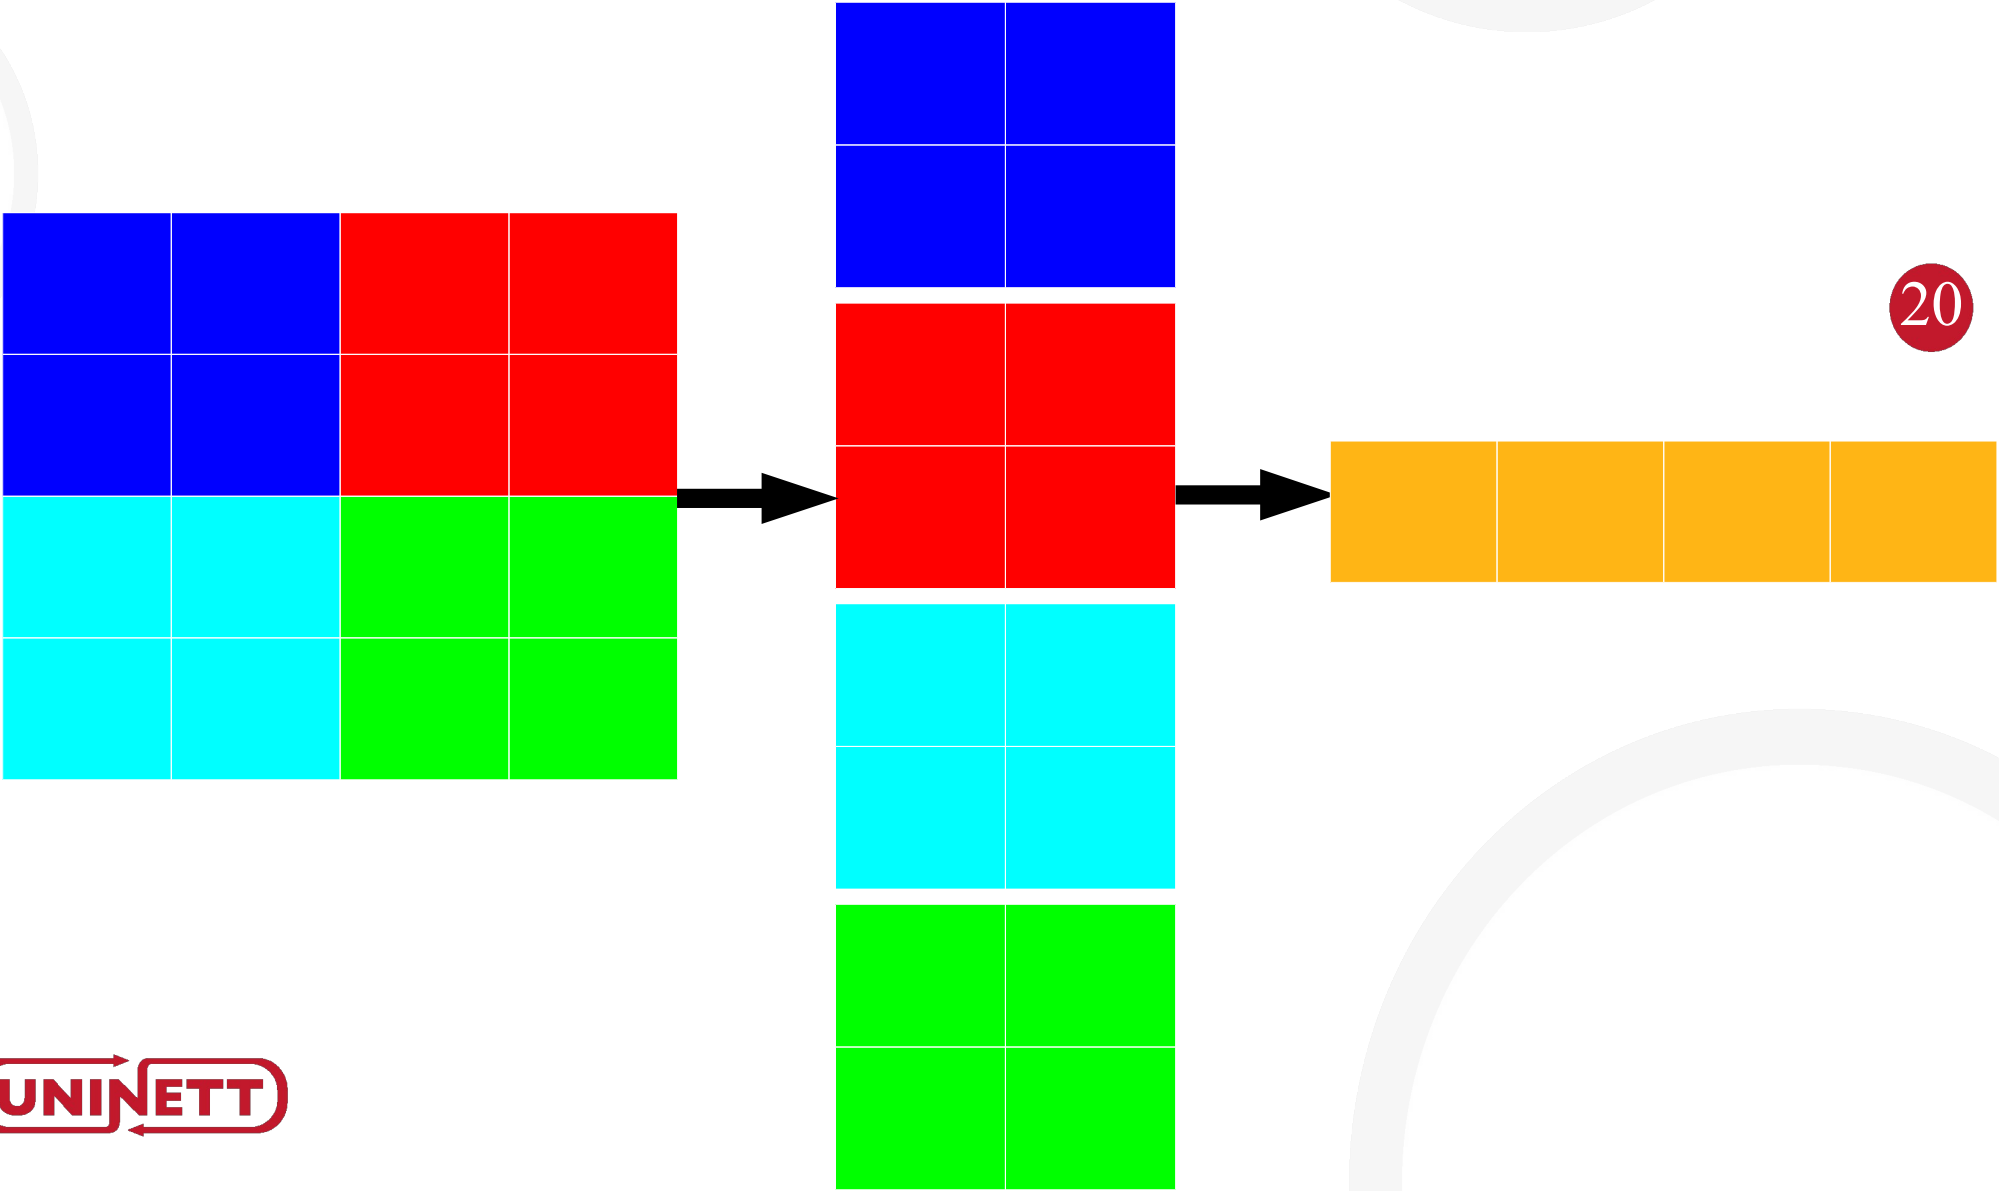
Academic  
world



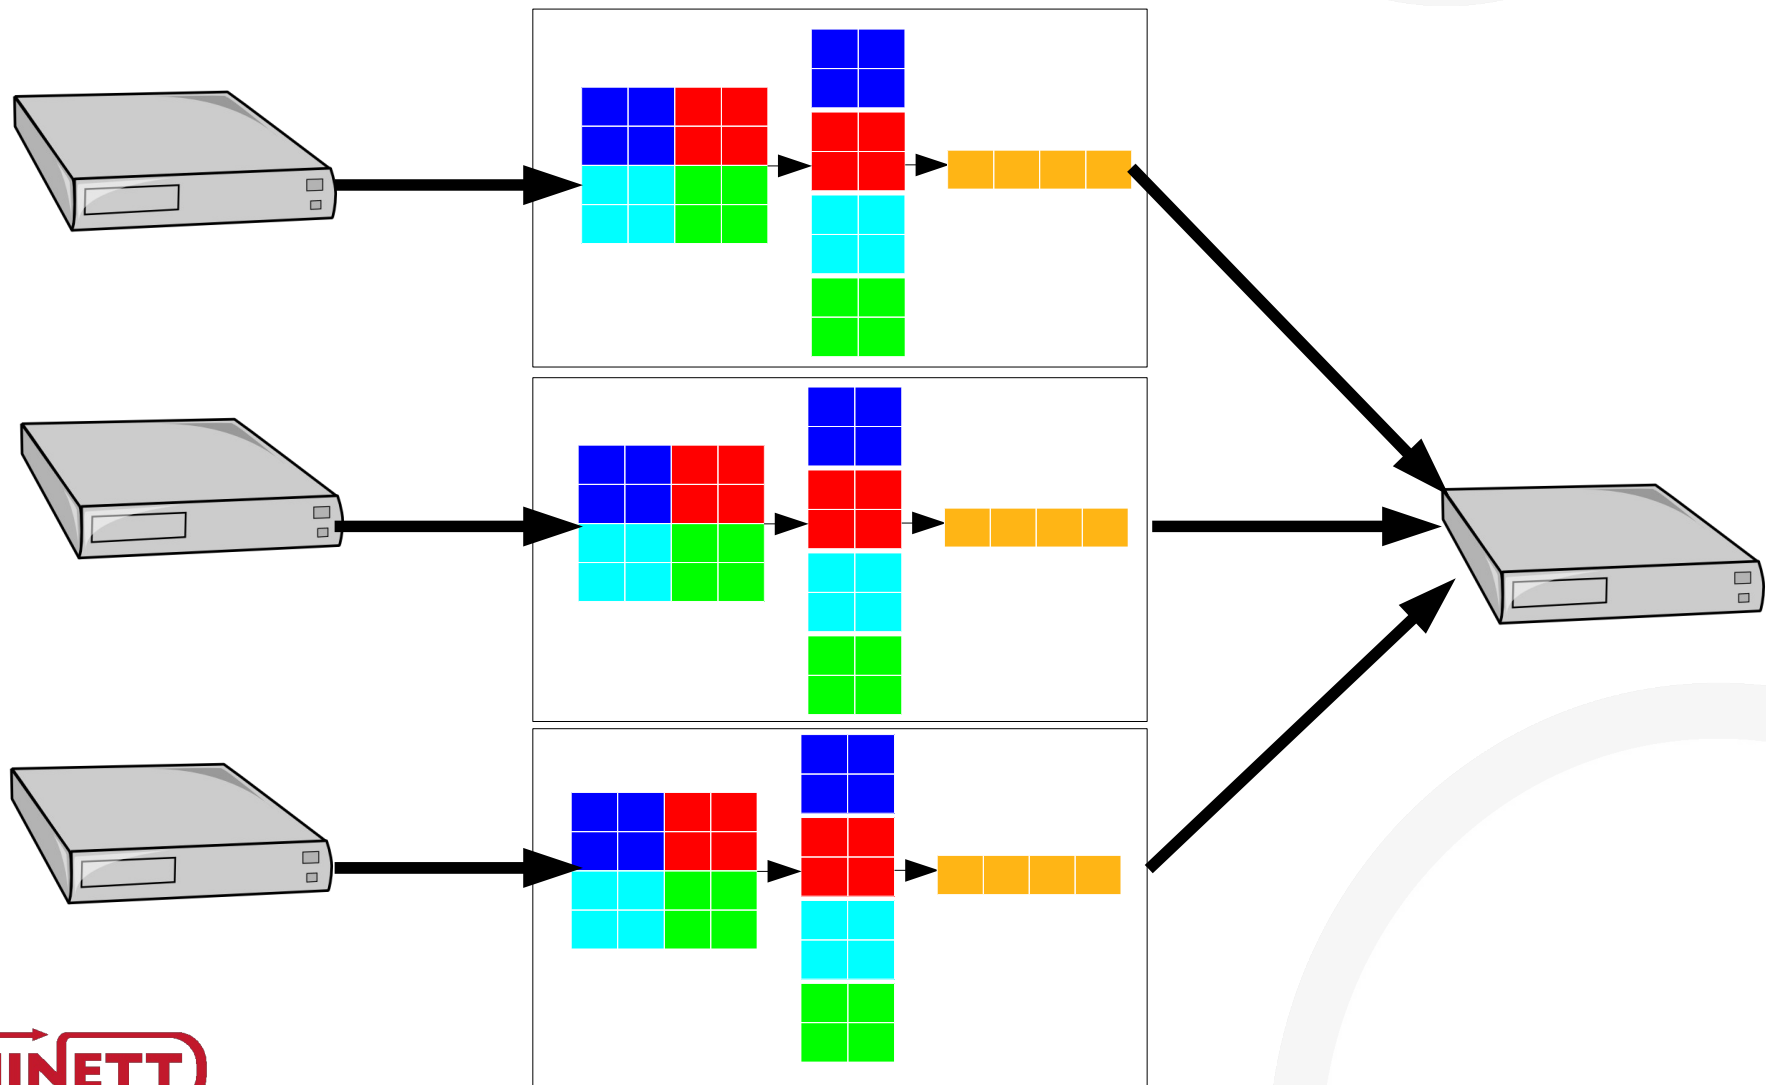
Real  
world



# Parallelizing

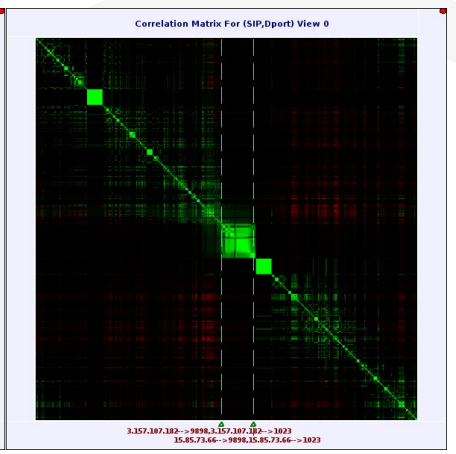
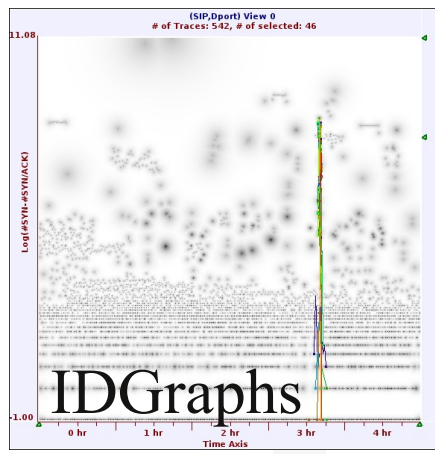
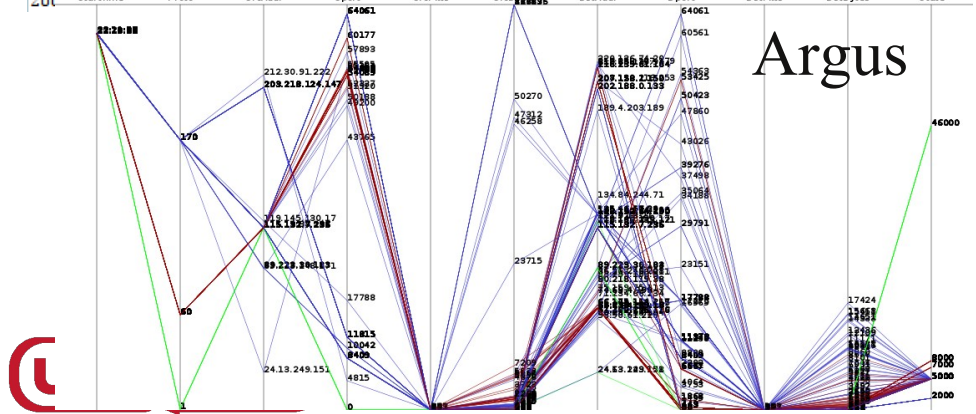


# Distributed processing

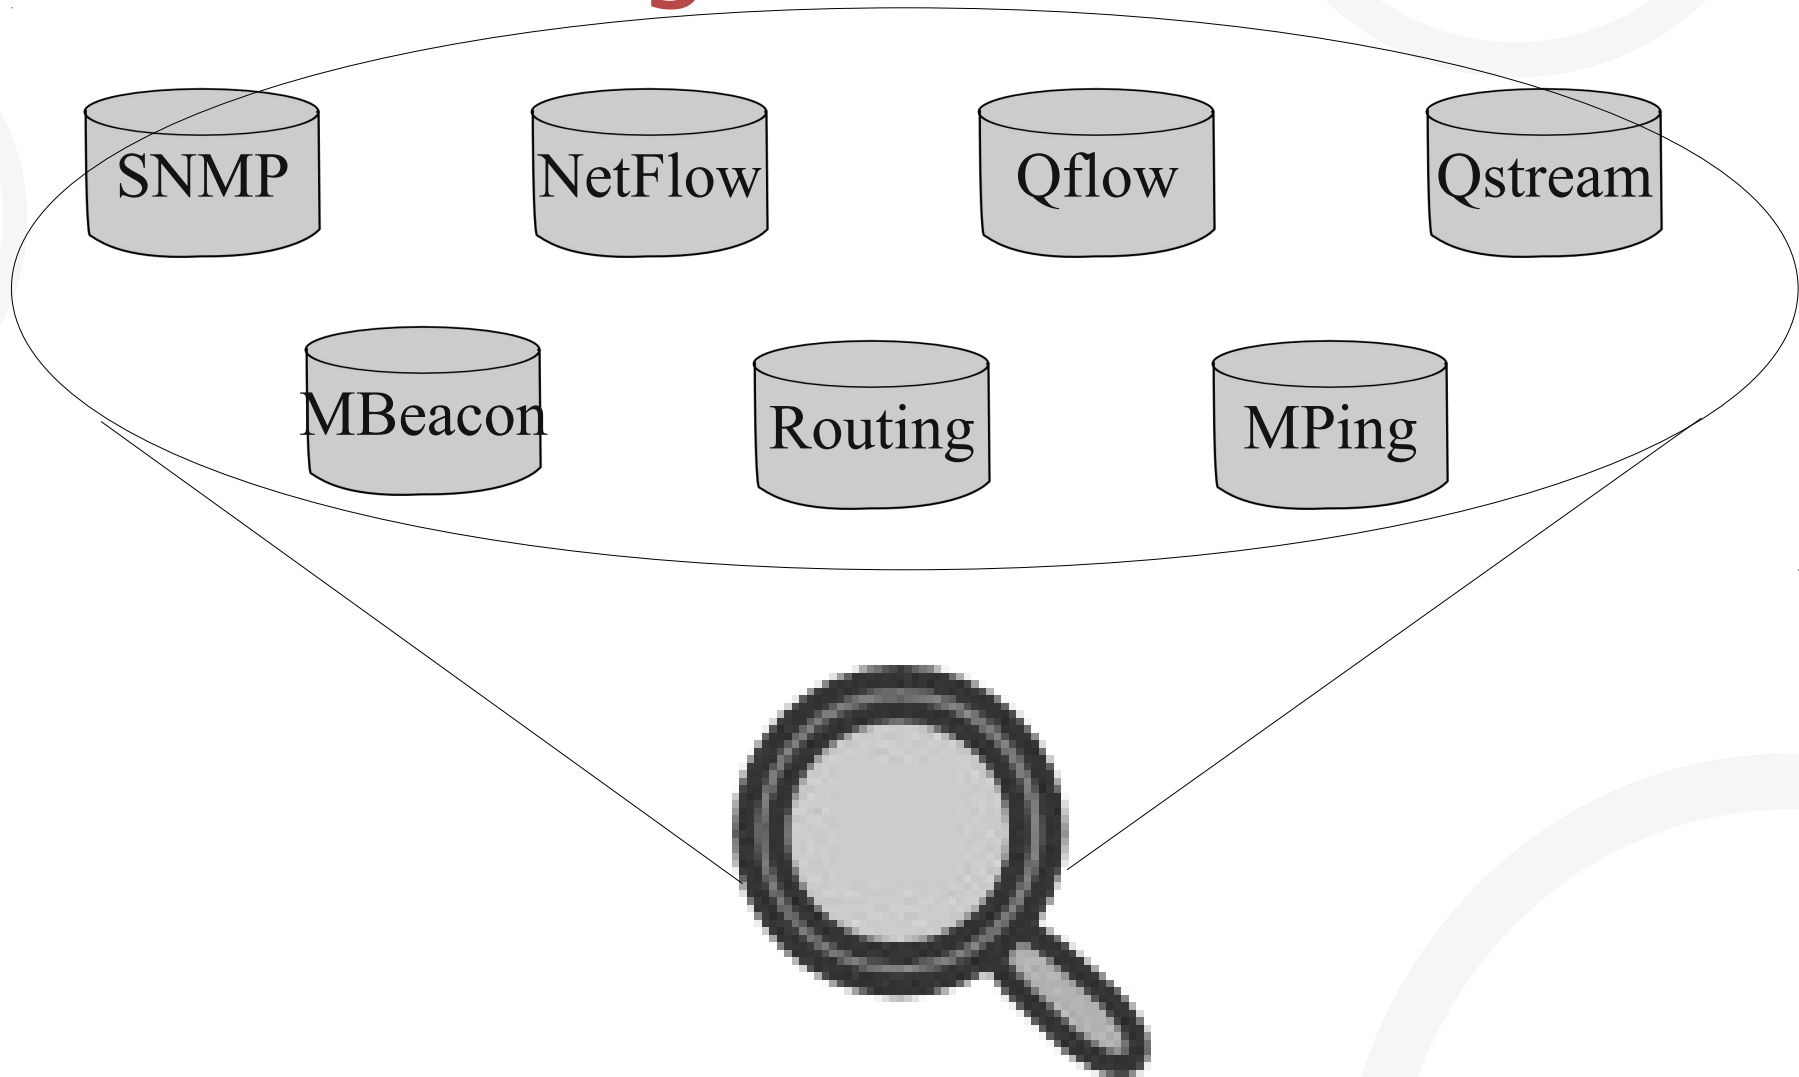


# Visualization


2009-08-24 22:49:18.056	0.000	TCP	174.98.75.63:80	->	97.227.177.67:3062	1	150
2009-08-24 22:49:10.501	0.000	TCP	174.98.75.63:80	->	119.219.244.75:2280	1	4
2009-08-24 22:49:44.676	0.010	TCP	174.98.75.63:80	->	97.225.243.155:2603	3	4501
2009-08-24 22:49:47.032	0.000	TCP	174.98.75.63:80	->	97.225.243.155:2615	1	1351
2009-08-24 22:49:08.899	0.000	TCP	174.98.75.63:80	->	58.190.234.182:58053	1	4
2009-08-24 22:49:34.551	0.000	TCP	174.98.75.63:80	->	97.227.242.138:62000	1	1501
2009-08-24 22:49:41.108	0.000	TCP	174.98.75.63:80	->	97.227.242.138:62025	1	1501
2009-08-24 22:48:59.869	0.000	TCP	174.98.75.63:80	->	97.227.242.138:61804	1	1501
2009-08-24 22:49:26.659	0.000	TCP	51.80.70.7:443	->	118.137.67.101:52092	1	1501
2009-08-24 22:49:28.304	0.000	TCP	118.137.49.192:52495	->	46.181.211.7:25	1	61
2009-08-24 22:49:01.831	0.036	TCP	118.137.49.192:39131	->	58.211.1.57:25	3	1721
2009-08-24 22:49:01.605	0.000	TCP	118.137.49.192:39131	->	58.211.1.57:25	1	61
2009-08-24 22:49:07.978	0.000	TCP	118.136.7.199:80	->	45.136.171.120:1759	1	4
2009-08-24 22:49:17.460	0.000	TCP	118.136.7.199:4899	->	94.242.141.79:2817	1	4
2009-08-24 22:48:58.143	51.810	TCP	151.231.10.6:53323	->	118.137.177.63:18860	13	531
2009-08-24 22:49:15.994	0.482	TCP	46.33.253.254:80	->	58.190.66.168:2949	2	481
2009-08-24 22:49:39.283	6.142	TCP	163.209.187.248:52190	->	118.138.225.238:2993	2	1151
2009-08-24 22:48:51.316	52.424	TCP	154.223.202.193:27243	->	97.227.15.220:55678	9	13421
2009-08-24 22:48:53.924	0.000	TCP	118.137.52.7:443	->	190.215.116.255:1588	1	4
2009-08-24 22:48:52.339	0.000	UDP	118.137.52.7:20684	->	42.156.135.226:6269	1	31
2009-08-24 22:48:53.476	0.000	UDP	118.137.52.7:20684	->	171.190.208.53:48099	1	4
2009-08-24 22:49:05.444	0.000	UDP	118.137.52.7:20684	->	174.213.20.160:48269	1	4
2009-08-24 22:49:21.778	0.000	UDP	118.137.52.7:20684	->	175.38.30.208:23363	1	4
2009-08-24 22:49:06.298	0.000	UDP	118.137.52.7:20684	->	54.255.250.157:22346	1	31
2009-08-24 22:49:26.148	0.000	TCP	118.137.52.7:20684	->	121.53.105.39:17704	1	361
2009-08-24 22:49:41.203	0.000	UDP	118.137.52.7:20684	->	181.4.90.143:49377	1	4
2009-08-24 22:49:20.706	0.004	TCP	172.55.32.7:65037	->	118.201.118.184:80	2	81
2009-08-24 22:49:20.595	0.000	TCP	172.55.32.7:65040	->	118.201.118.184:80	1	4
2009-08-24 22:49:29.566	2.034	TCP	172.55.32.7:65047	->	118.201.118.184:80	2	801
2009-08-24 22:49:29.884	0.000	TCP	172.55.32.7:65049	->	118.201.118.184:80	1	401
2009-08-24 22:48:57.610	53.110	UDP	191.86.64.6:6491	->	97.226.81.118:12133	18	10461
2009-08-24 22:48:51.876	55.722	UDP	167.3.162.0:3672	->	118.137.161.252:54867	11	10601
2009-08-24 22:49:06.482	32.979	UDP	165.71.162.62:38801	->	118.137.123.177:28881	2	3071
						1	481



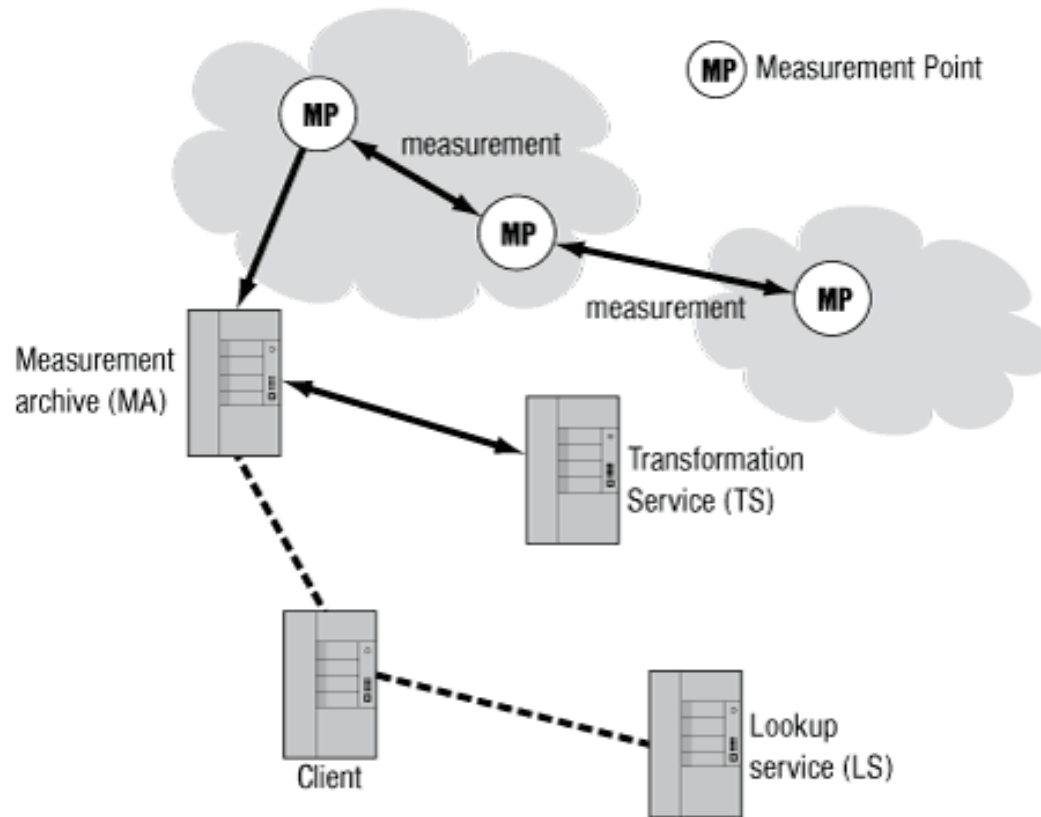
# Correlating data



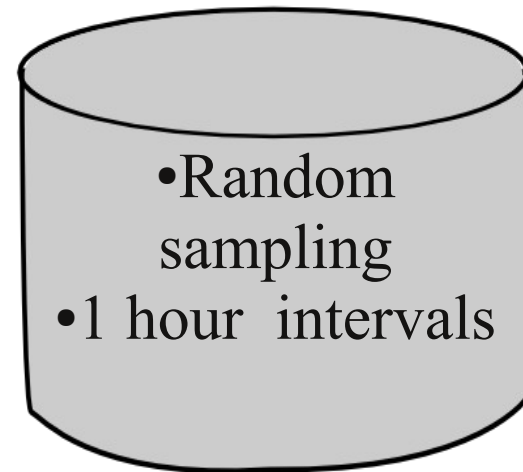


# Multidomain monitoring



# perfSONAR




# Normalizing data



How do you collect and compare values for the last 15 minutes?

# Standardized access control



# End user experience



28

# Summary

- Need multiple technologies for proper network management and monitoring
  - SNMP
  - Netflow
  - Passive and active monitoring
- Challenges
  - Scalable anomaly detection
  - Good visualization
  - Distributed processing
  - Multi-domain monitoring
  - Translate monitoring results into end user experiences