Scaling network management tools

Olav Kvittem

16th October 2003

Abstract

UNINETT, The norwegian research network, is actively collecting information on it's network in order to asess and maintain the quality of the offered service. We are developing tools to do management and measurements on our own and our customers infrastructure. This talk describes some of the management and measurement activites and the tools.

UNINETT

- The norwegian research network since 1987 with 280 customers in higher education and research
  - 55 people in 4 companies doing internet service, .no names, administrative systems, schools networking advice
- 2.5 Gbps backbone, access to wavelengths/fibre through cooperation with telecom provider
  - local cooperative fibre projects

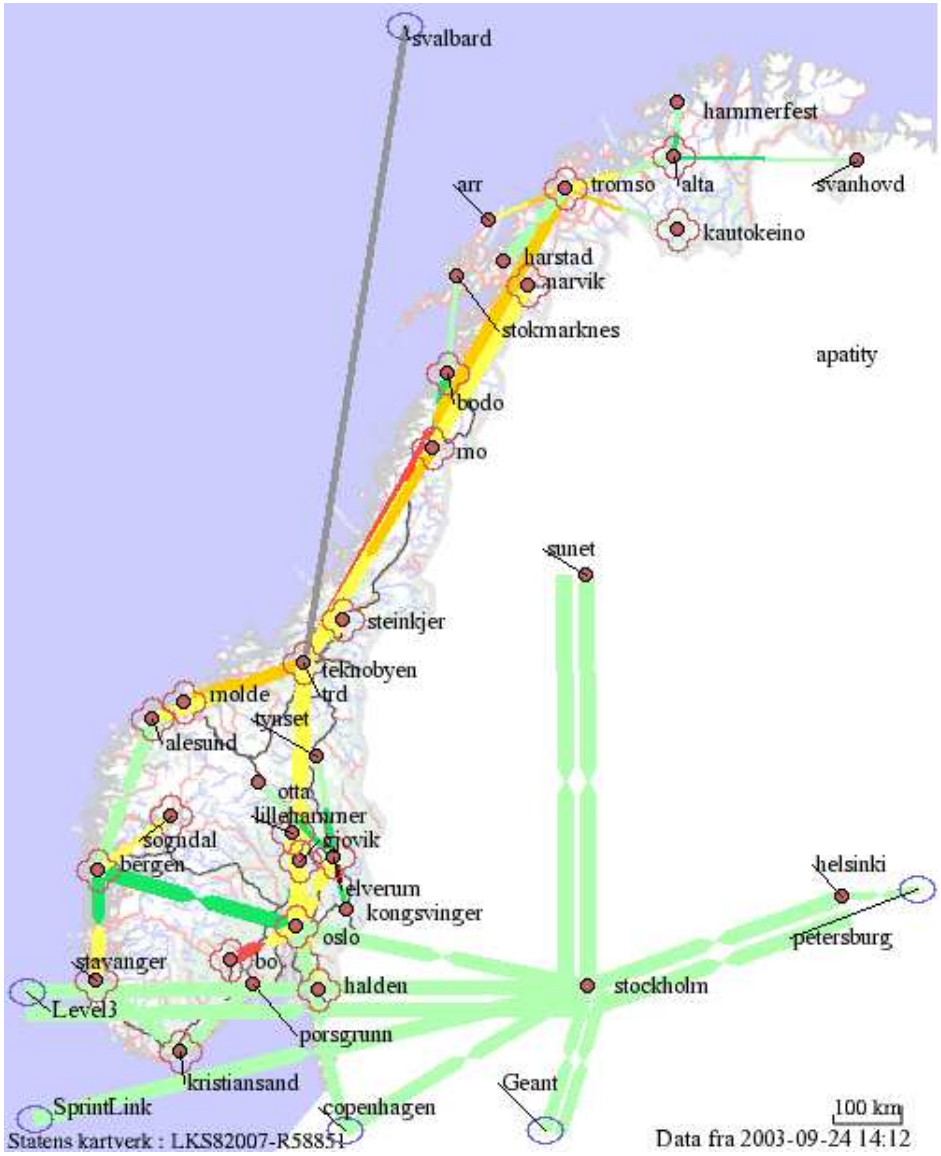
svalbard
hammerfest
arr
tromso
alta
svanhovd
kautokeino
harstad
narvik
stokmarknes
apatity
bodo
mo
sunet
steinkjer
teknobyen
molde
trd
tynset
alesund
otta
lillehammer
sogndal
gjovik
bergen
helsinki
elverum
kongsvinger
oslo
petersburg
stavanger
bo
halden
stockholm
Level3
porsgrunn
kristiansand
Geant
SprintLink
copenhagen
100 km
Statens kartverk : LKS82007-R58851
Data fra 2003-09-24 14:12

Why homegrown

- PSI/Nysernet SNMP from 1990-1998
  - nice but a bit static cumbersome configuration
- Major NM platforms - solved wrong problems
  - centralized operations and intranet-oriented
  - poor history functions
- Resources : Hackers and students and open software => develop tailored tools

Scaling principles

Accessability - UNINETT is distributed in organisation : users, customers, external projects, operations, engineering, services, research, managers, Board of Directors ...

Trends - Proactive better that reactive - se trends in traffic, error rates over days, weeks and years

Numbers - can’t do 100’s of customers, devices and links one by one
=> summaries, thresholds, tables and sorting

Dynamic - config change (SNMP ifIndex)

Usability - easy access to the most important related statistics

Visuality - graphs, maps and animation

Openness - open software : perl, TCL, sh, postgresql, PHP, python, net-snmp, flow-tools, scotty, ...

Network maps

3 network load map systems being made in Trondheim

Netmap autozoomed geographical maps (UNINETT)
- menues with URLs
- animation of any point or link load

Zino schematic load maps - tgif (UNINETT/NORDUnet

Nav autodetected campus topology (NTNU)

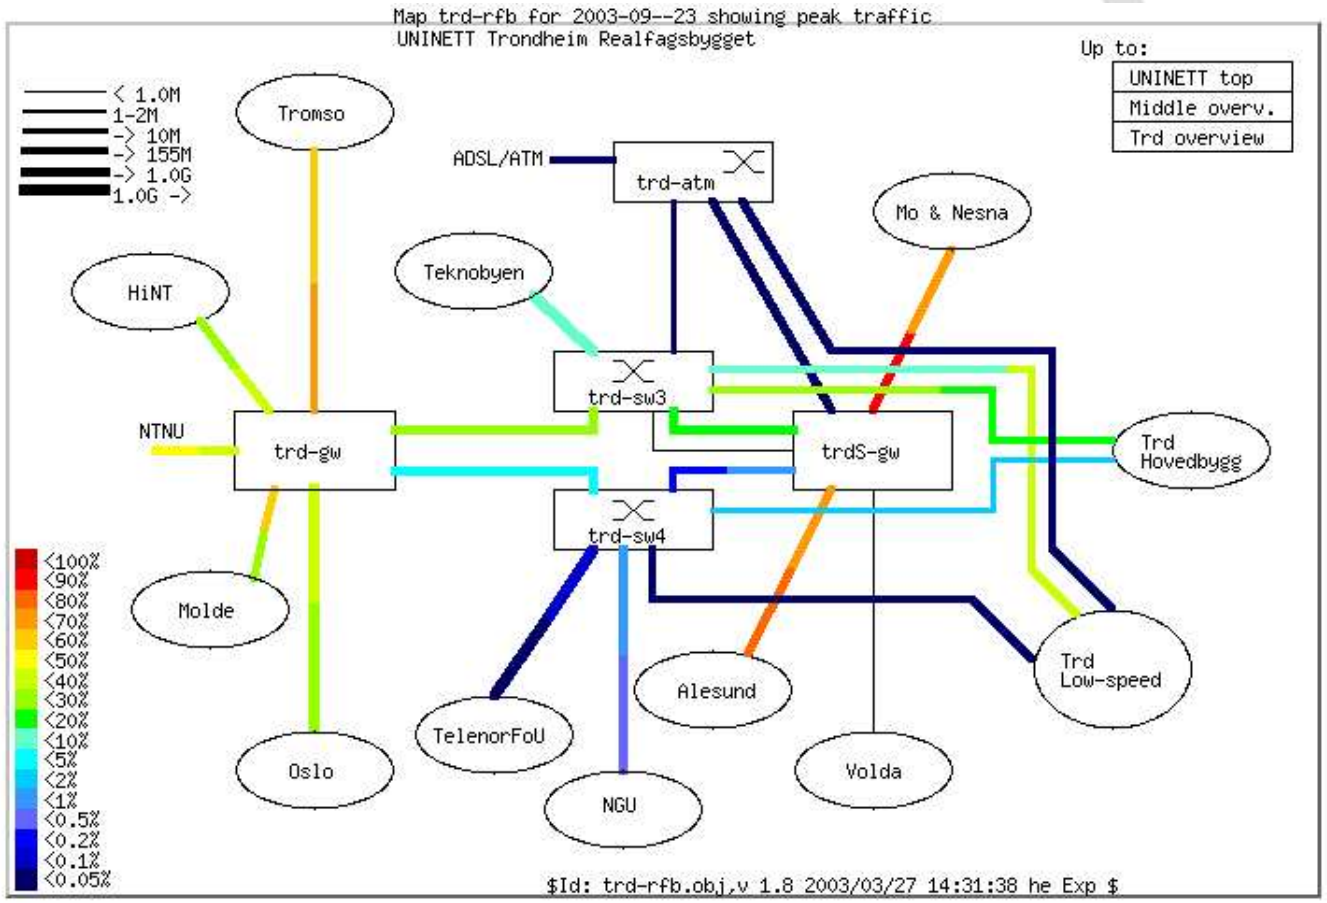
Map trd-rfb for 2003-09--23 showing peak traffic
UNINETT Trondheim Realfagsbygget
Up to:
UNINETT top
Middle overv.
Trd overview
< 1.0M
1-2M
-> 10M
-> 155M
-> 1.0G
1.0G ->
Tromso
ADSL/ATM
trd-atm
Mo & Nesna
Teknobyen
HiNT
trd-sw3
NTNU
trd-gw
trdS-gw
Trd Hovedbygg
trd-sw4
<100%
<90%
<80%
<70%
<60%
<50%
<40%
<30%
<20%
<10%
<5%
<2%
<1%
<0.5%
<0.2%
<0.1%
<0.05%
Molde
Alesund
Trd Low-speed
TelenorFoU
Oslo
Volda
NGU
$Id: trd-rfb.obj,v 1.8 2003/03/27 14:31:38 he Exp $

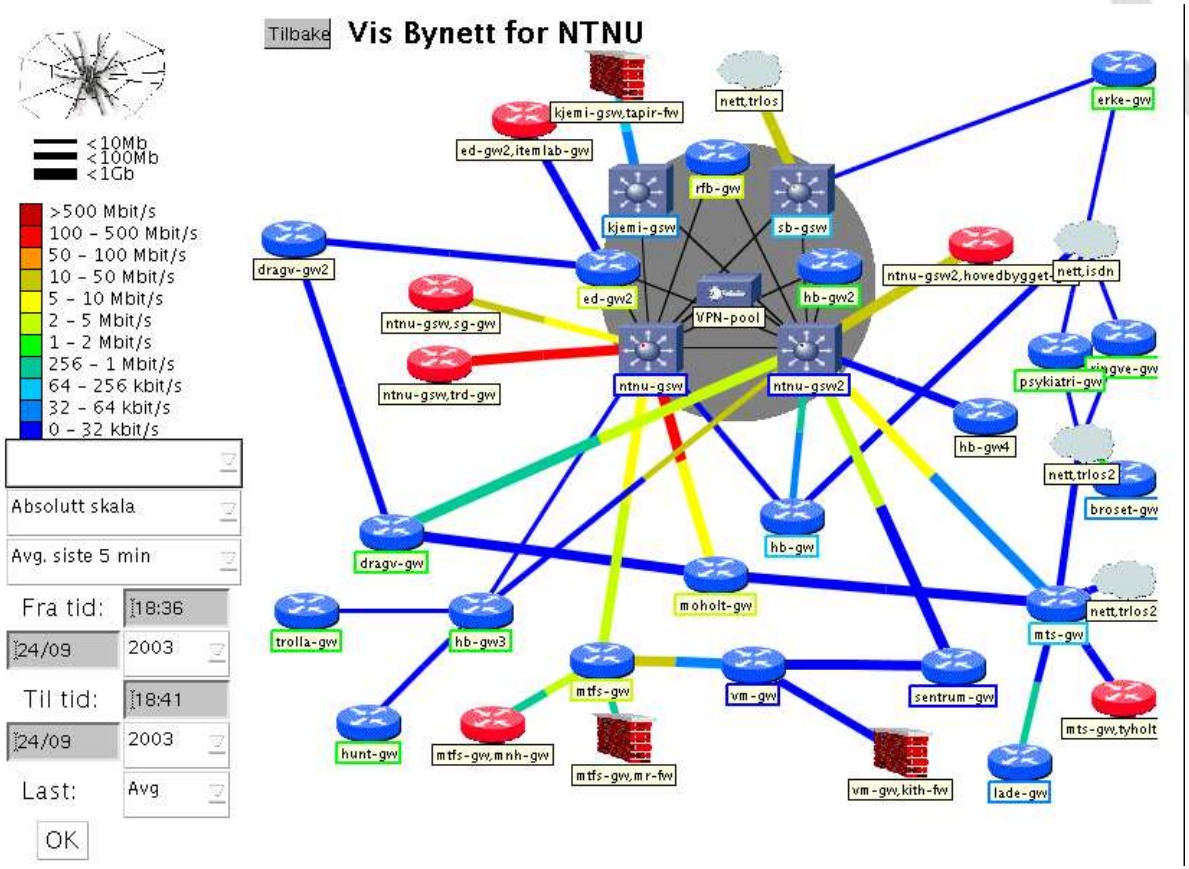
Tilbake
Vis Bynett for NTNU
<10Mb
<100Mb
<1Gb
>500 Mbit/s
100 - 500 Mbit/s
50 - 100 Mbit/s
10 - 50 Mbit/s
5 - 10 Mbit/s
2 - 5 Mbit/s
1 - 2 Mbit/s
256 - 1 Mbit/s
64 - 256 kbit/s
32 - 64 kbit/s
0 - 32 kbit/s
Absolutt skala
Avg. siste 5 min
Fra tid:
18:36
24/09
2003
Til tid:
18:41
24/09
2003
Last:
Avg
OK
kjemi-gsw,tapir-fw
nett,trlos
erke-gw
ed-gw2,itemlab-gw
rfb-gw
kjemi-gsw
sb-gsw
dragv-gw2
ed-gw2
hb-gw2
ntnu-gsw2,hovedbygget
nett,isdn
ntnu-gsw,sg-gw
VPN-pool
ntnu-gsw,trd-gw
ntnu-gsw
ntnu-gsw2
psykiatri-gw
hb-gw4
nett,trlos2
broset-gw
dragv-gw
hb-gw
moholt-gw
trolla-gw
hb-gw3
mts-gw
mtfs-gw
vm-gw
sentrum-gw
mts-gw,tyholt
hunt-gw
mtfs-gw,mnh-gw
mtfs-gw,mr-fw
vm-gw,kith-fw
lade-gw

Network Map System(netmap)

- make information and network maps
- generates suitable clickable WWW-maps from a topology database
- breaks the topology down to suitable maps
- use geographical maps - UTM coordinates
- generate URLs for the maps based on database info

- navigation and menus with a Java-client
- animate - link and cpu load, delays, protocols,..

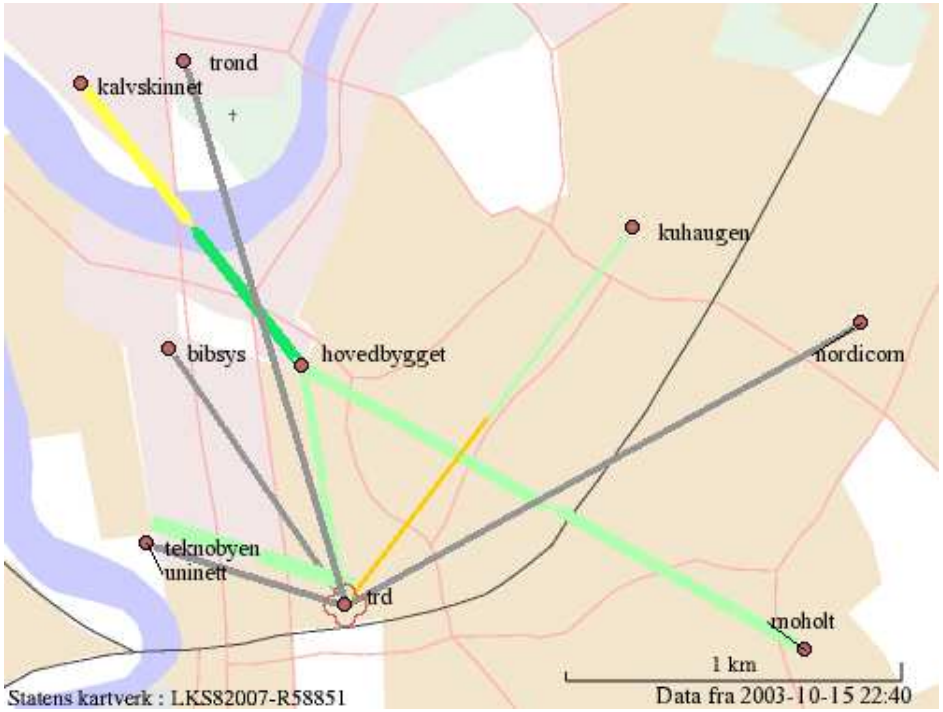

Why measurements

- Problem detection and solving!
- To asess the quality of our service !
- Capacity planning and traffic engineering
- To inform help customers and users on their own via the web
- To assist research that will find interesting phenomena for us...

Partners - arenas

- Work with researchers - offer access to data, being a lab
  - Q2S - Center for Quantifiable Quality of Service (NTNU)
- Actively support student work - projects, thesis
  - Student employees to do programming

International

- International participation in European fora like
  - Terena - TF-NGN - network level experiments and studies(perfmon)
  - EU-projects like Scampi
  - cooperation with measurement activities like CAIDA(AMS)
  - IETF - net management - ipfix, ippm,

Passive activities..

Scampi - a EU-project with about 10 participants to develop a free and low cost

- high speed passive measurement platform(10Gbps)
- API with adapted “standard” open software (tcpdump, flow-tools)

SNMP tools 1

Zino SNMP link statistics

- tables, graphs, aggregation and error analysis
- scales by config just by pointing to the router
- map IP-address and link name from description field/ifAlias

SNMP tools 2

Genplot General SNMP statistics package

- collect, aggregate SNMP or other data and present in tables and graphs with zoomed context

JustSNMP - console tool to extract data like links with name and BGP (Scotty)

other tools hw and software version inventory, sw version control

UNINETT Forskningsnettet i Norge

Webkart

CPU and Memory Usage - busyPer

week 30 in 2003

CPU busy percentage in the last 5 second period. Not the last 5 realtime seconds but the last 5 second period in the scheduler.

You may check each line to generate a plot for this variable.

"Raw" plots the raw data. "Other" allows you to plot other variables for the selected line

Filename	Busy value	at	Peak sample	at	Period avg	Plot	
telefonsentral-sw1	38	Tue	39	Tue	38	Raw	Other
158_39_12_2	37	Fri	39	Fri	37	Raw	Other
osalle-gw	31	Mon	61	Sun	31	Raw	Other
bo-gw	35	Thu	92	Sat	30	Raw	Other
lillehammer-gw	49	Wed	100	Wed	22	Raw	Other
halden-gw	31	Wed	74	Wed	20	Raw	Other
molde-gw	24	Sun	63	Sun	19	Raw	Other
breivika-gw	24	Wed	57	Wed	19	Raw	Other
sogndal-gw	22	Thu	68	Wed	19	Raw	Other
borre-gw	26	Wed	44	Wed	17	Raw	Other
supergw	22	Sat	100	Sat	17	Raw	Other
stavanger-gw	19	Tue	46	Wed	16	Raw	Other
elverum-gw	24	Mon	98	Mon	16	Raw	Other
narvik-gw	18	Sun	69	Tue	15	Raw	Other
alta-gw	26	Sun	51	Sun	14	Raw	Other
kautokeino-gw	15	Fri	97	Tue	14	Raw	Other
cadeler30-gw	14	Fri	27	Mon	14	Raw	Other

UNINETT Forskningsnettet i Norge

Webkart

CPU and Memory Usage

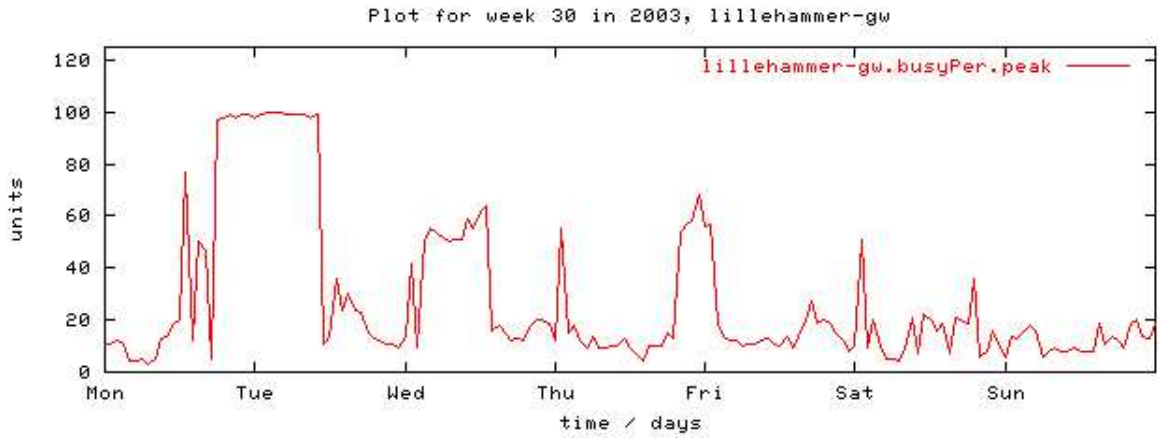

The above image "clickable" for zooming etc. A brief explanation is available.

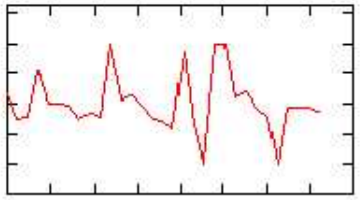

July, 2003

You may view other plots or browse the data for this period

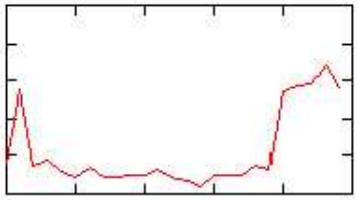

2003-07-24

info@uninett.no

2003-07-11

Zino status monitor

- Polls and handles SNMP events (traps)
- Simple Trouble Ticketing - User authentication
- Fin grained downtime registration (IfLastChange)
- Link identification by Cisco description/ifAlias
- Availability statistics from the logs

Ritz

- Remote Interface To Zino

File Edit View Help

OpState	AdmState	Age	Downtime	Router	Port
port down	working	30d 1:57:57.00	30d 2:00:08.00	tromsoS-gw	Serial6/1/3
port down	working	41d 7:21:18.00	41d 7:25:12.00	altakunde-sw	FastEthernet0/2
port down	working	30d 1:57:55.00	30d 2:00:00.00	bodo-gw	Serial1/6
port up	open	59d 16:05:28.00	0d 1:35:49.00	tromsoS-gw	ATM1/0.25-atm subi
port up	open	59d 16:05:28.00	0d 1:35:45.00	tromsoS-gw	ATM1/0.25-aal5 lay
port up	open	13d 23:58:26.00	0d 1:10:49.00	tromsoS-gw	ATM1/0.27-atm subi
port up	open	13d 23:58:25.00	0d 1:10:48.00	tromsoS-gw	ATM1/0.27-aal5 lay
port up	open	6d 10:01:43.00	0d 0:31:52.00	bergen-gw	POS0/3
port up	open	23d 2:04:23.00	0d 0:12:26.00	tromsoS-gw	Serial6/1/1
port up	open	5d 19:57:38.00	0d 0:05:15.00	trd-gw	GigabitEthernet3/1
port up	open	7d 16:48:48.00	0d 0:04:46.00	oslo-gw2	POS2/0
port up	open	2d 2:53:10.00	0d 0:01:03.00	tromsoS-gw	ATM1/0.28-atm subi
port up	open	2d 2:53:09.00	0d 0:01:02.00	tromsoS-gw	ATM1/0.28-aal5 lay
port up	open	0d 6:51:06.00	0d 0:00:07.00	oslo-gw3	FastEthernet0/2
port up	open	2d 0:52:01.00	0d 0:00:00.00	stolav-gw	POS1/2
port up	open	6d 23:55:18.00	0d 14:17:43.00	oslo-gw6	Serial2/7
port up	open	17d 6:26:59.00	0d 1:30:30.00	oslo-gw6	Serial3/4
port up	open	13d 23:56:15.00	0d 1:09:45.00	tromso-atm	ATM1/0/0
port up	open	48d 4:51:46.00	0d 0:29:57.00	bergen-atm	Ethernet2/0/0
port up	open	15d 6:56:53.00	0d 0:13:50.00	bergen-gw2	Serial2/7
port up	open	10d 20:25:15.00	0d 0:03:31.00	narvik-gw	FastEthernet2/0
port up	open	7d 7:19:01.00	0d 0:03:01.00	mo-gw	Serial4/1
port up	open	0d 0:08:18.00	0d 0:02:30.00	bergen-gw2	Serial2/6
port up	open	6d 3:42:45.00	0d 0:01:57.00	kristiansand-gw	POS6/0

Active measurements - mping

Mping - scaling ping measurements IPv4/IPv6

- polls targets in parallel at controlled rate
- repeated at Poisson based intervals
- statistical analysis - percentiles, distribution
- aggregation with plots, tables, traceroute view
- animate reponse time distribution

Mping table report

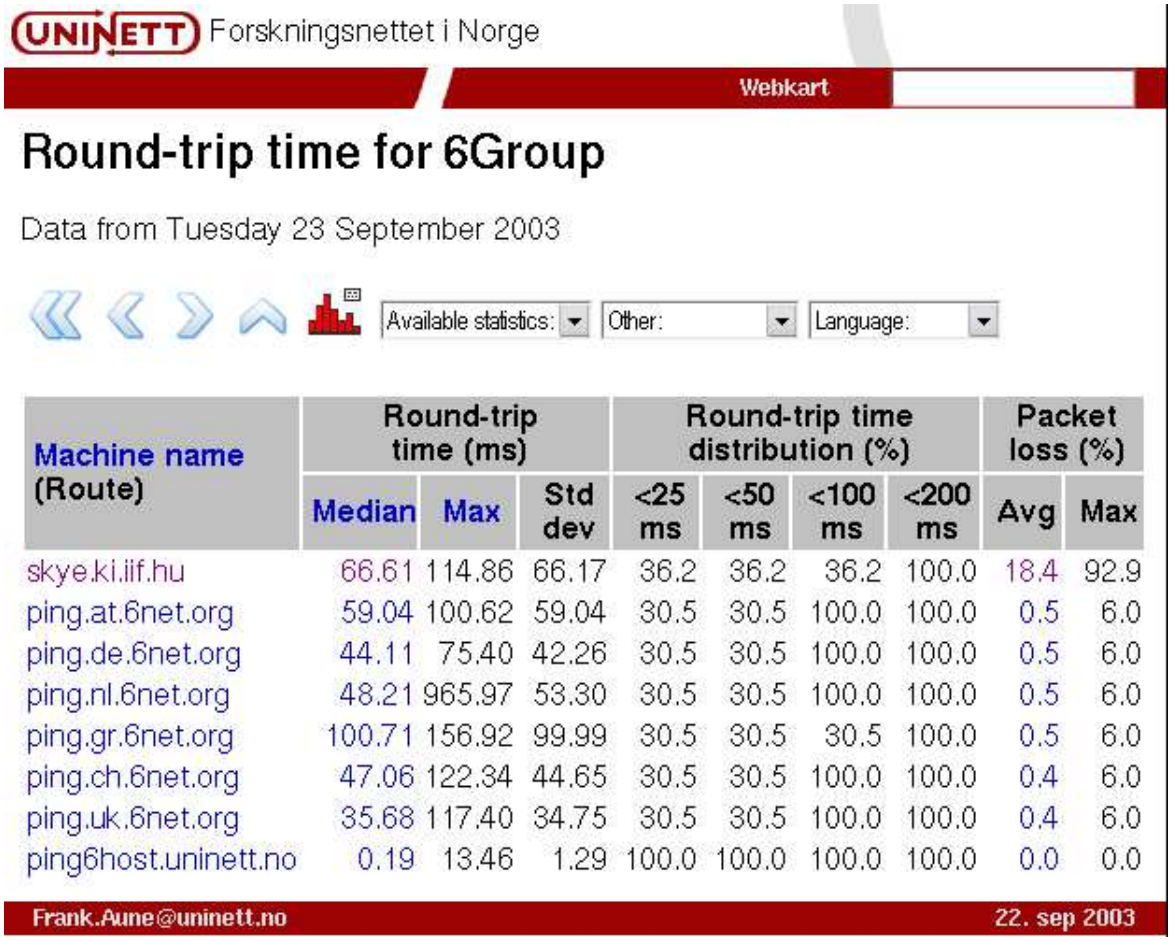

UNINETT Forskningsnettet i Norge

Webkart

Round-trip time for 6Group

Data from Tuesday 23 September 2003

Available statistics: | Other: | Language:

Machine name (Route)	Round-trip time (ms)			Round-trip time distribution (%)				Packet loss (%)	
	Median	Max	Std dev	<25 ms	<50 ms	<100 ms	<200 ms	Avg	Max
skye.ki.iif.hu	66.61	114.86	66.17	36.2	36.2	36.2	100.0	18.4	92.9
ping.at.6net.org	59.04	100.62	59.04	30.5	30.5	100.0	100.0	0.5	6.0
ping.de.6net.org	44.11	75.40	42.26	30.5	30.5	100.0	100.0	0.5	6.0
ping.nl.6net.org	48.21	965.97	53.30	30.5	30.5	100.0	100.0	0.5	6.0
ping.gr.6net.org	100.71	156.92	99.99	30.5	30.5	30.5	100.0	0.5	6.0
ping.ch.6net.org	47.06	122.34	44.65	30.5	30.5	100.0	100.0	0.4	6.0
ping.uk.6net.org	35.68	117.40	34.75	30.5	30.5	100.0	100.0	0.4	6.0
ping6host.uninett.no	0.19	13.46	1.29	100.0	100.0	100.0	100.0	0.0	0.0

Frank.Aune@uninett.no 22. sep 2003

UNINETT Forskningsnettet i Norge

Webkart

Round-trip times and packet loss

For freebsd.ki.iif.hu (skye.ki.iif.hu) on August 2003

Available statistics: | Other: | Language:

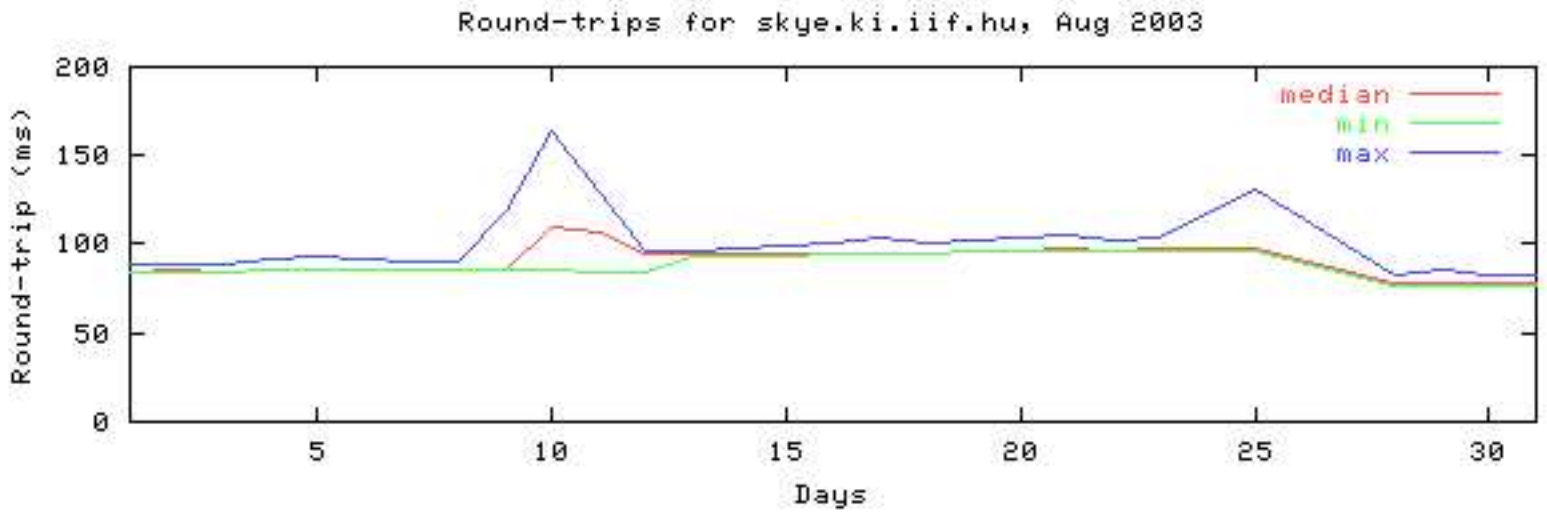

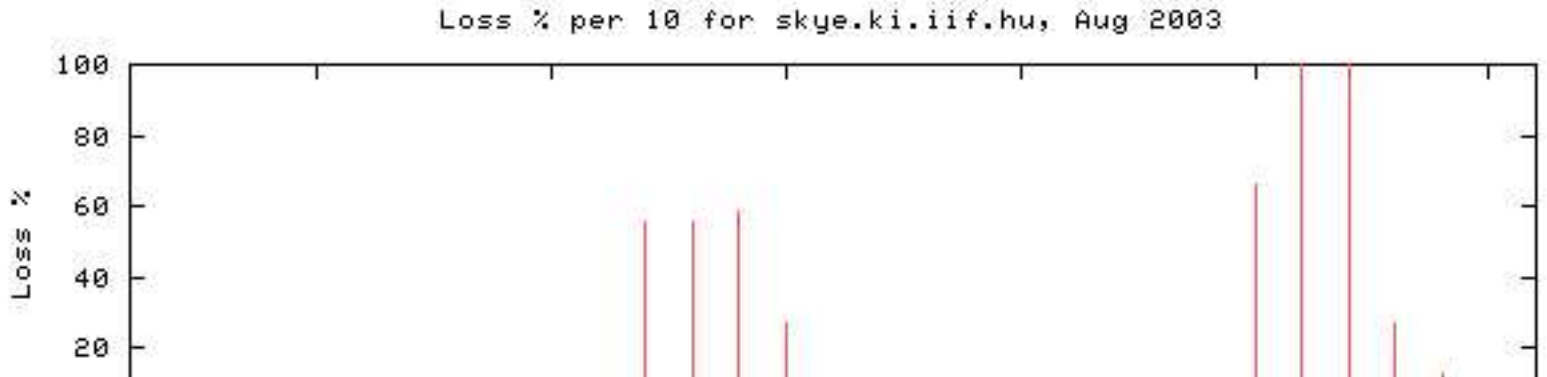

Mping round-trip distribution

UNINETT Forskningsnettet i Norge

Webkart

Round-trip time distribution

For '6Group', Tuesday 23 September 2003

Available statistics: Other: Language:

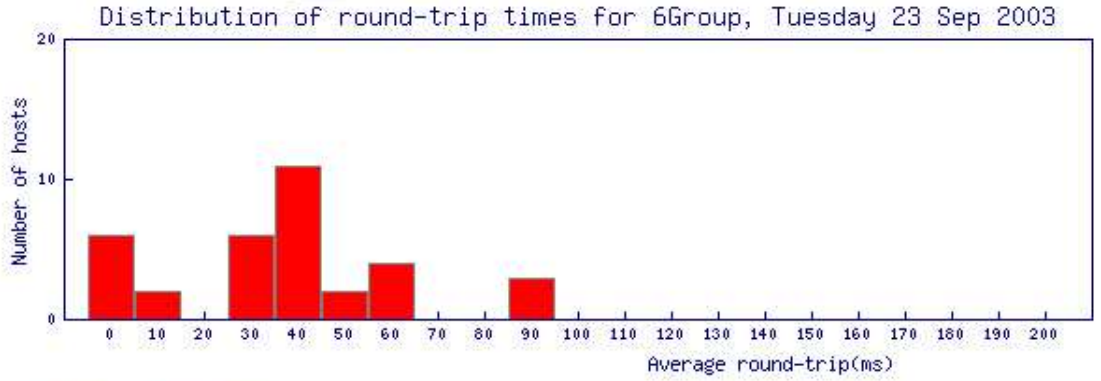

Frank.Aune@uninett.no

micro measurements

micro measurements - what is the short term sub-second load condition on a link

- traditional SNMP statistics poll frequency is in order of minutes
- poll SNMP-agents at sub-second time resolution with interleaving short and long intervals.
- Graph in real time to do immediate diagnostics

- Problem : routers have varying SNMP MIB update interval (1-15sec)
  - puts low priority on answering SNMP

- Extended SNMP-agent for DAG-cards
  - should give at least ms accuracy

- Analysis of DAG-card packet dumps for bursts (talk at NORDUNET 2003)

As short as it makes sense to go

1 ms

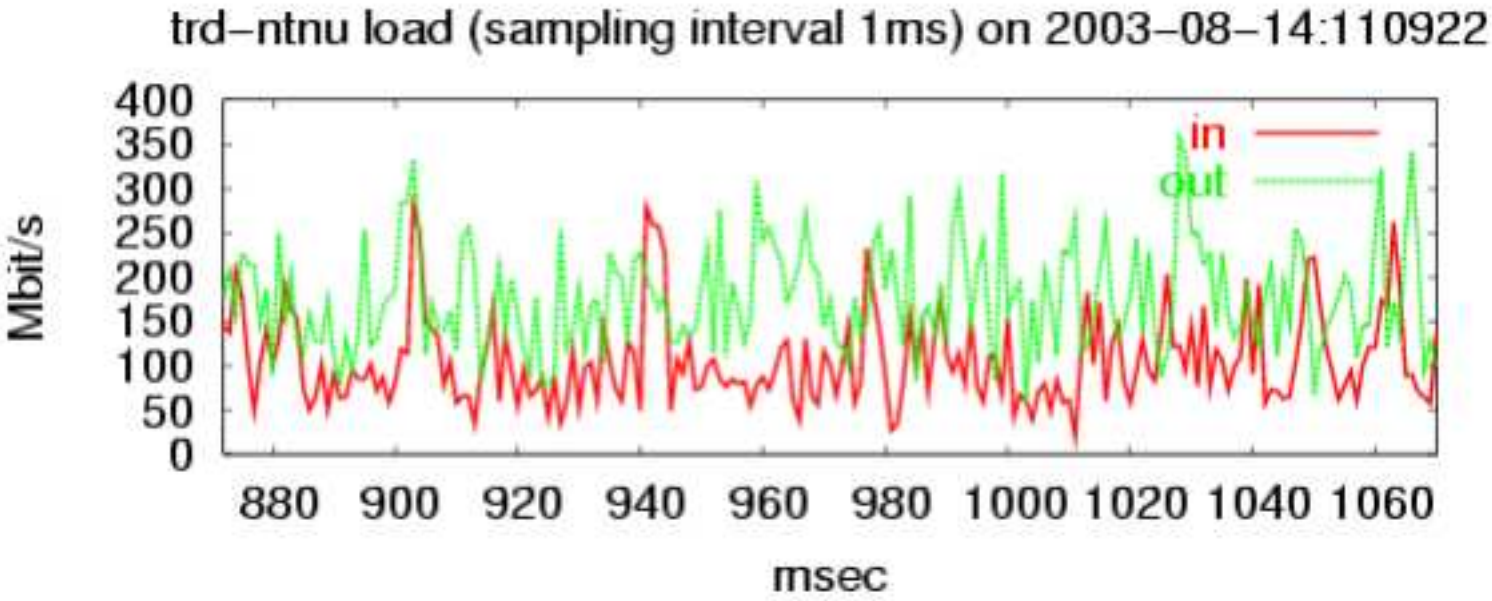

100 µs

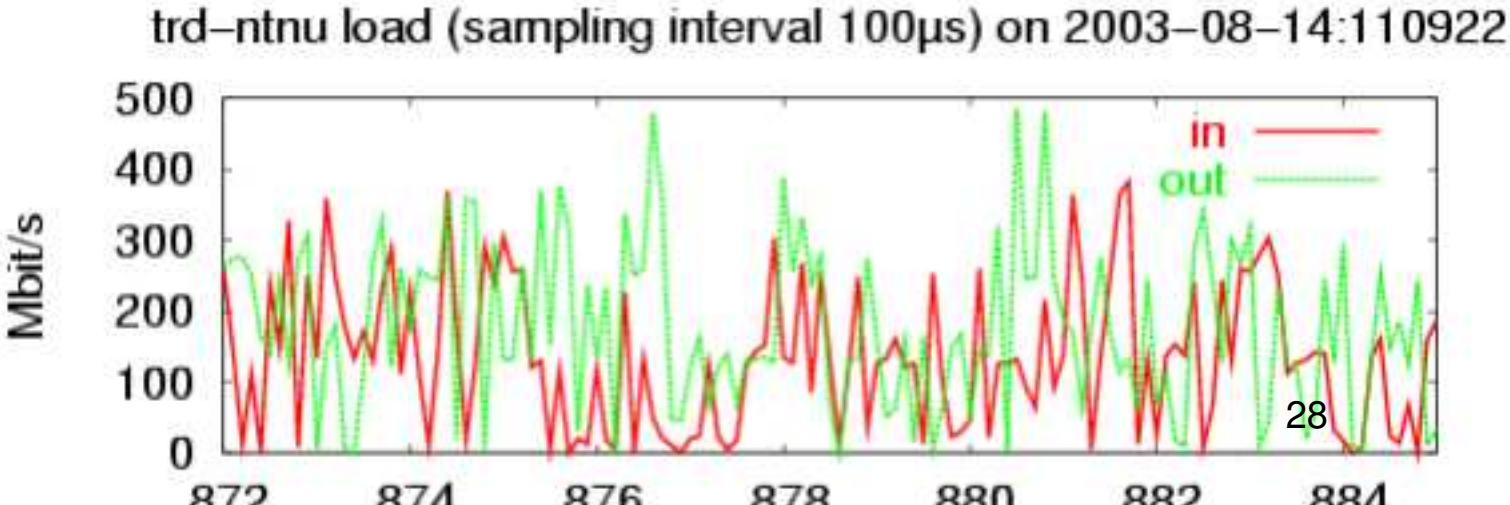

Flow reporting

- Develops IETF ipfix flow generation with passive monitoring cards (DAG, Scampi, Ethernet)

- Flow collection and reporting with “scaling” properties
  - based on flow-tools and will be contributed
  - Postgresql with aggregation and statistics
  - multiuser interface with tables and graphs

SCAMPI FlowRep

UNINETT Forskningsnettet i Norge

Webkart

Report setup > Graph | Source Port | Simple | Limit rows: 25 | OK | [Login]

Week | Month | Day: Monday

Show all | Select Group | uninett-nordunet | Opposite | Direction In Out | Choose interface

Plot type: Area graph | Y-scale: Linear | Plot image size: 800x400 | Omit other | Replot

Source Port

Week 30 2003

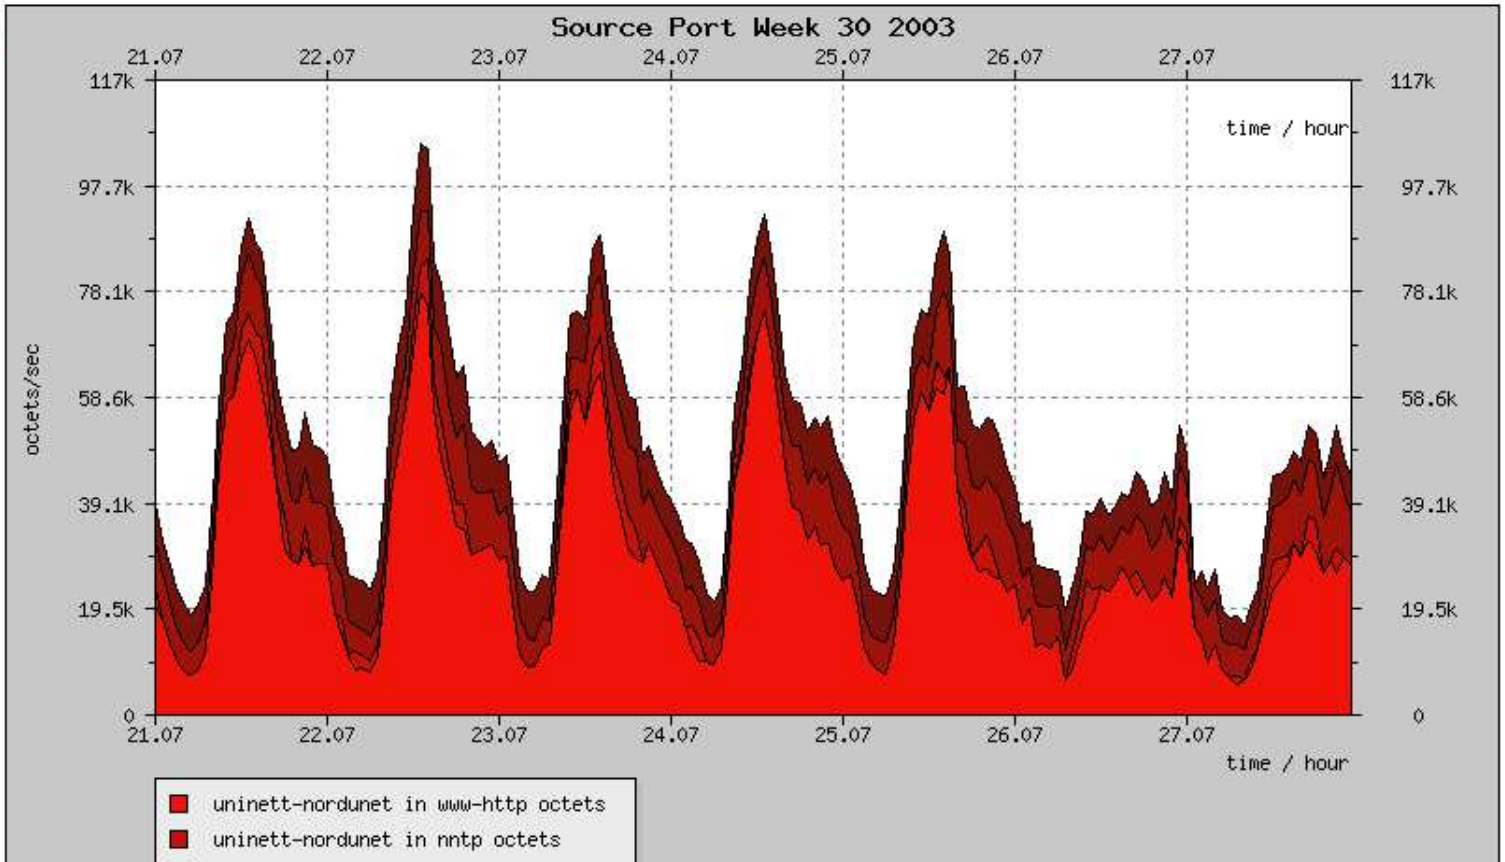

Software availability

Zino status monitor and link statistics with load map - available high quality

mping multi-destination parallel ping with statistics aggregation IPv4/v6 - available

scampi tools flowrep - flow-tools extension - available late this autumn
- other passive monitoring tools - next year

micro-poll available

genplot general SNMP statistics aggregation, reporting av plotting - available

netmap geographic map and menu system - not sufficiently documented

nemo java based netmap client for animation - not documented yet

justnetstat available

Software license

The software is freely available but there is a re-distribution clause:

Corollary 1. *# Copyright (c) 1996, 1997 # UNINETT and NORDUnet. All rights reserved.*

 # Redistribution and use in source and binary forms, with or without

 # modification, are permitted provided that the following conditions ...

more details follow the software

Thank you !

Olav.Kvittem@uninett.no

General : http://www.uninett.no

Statistics : http://drift.uninett.no/index.en.html