


The 1st International NorNet Users Meeting (NNUW-1)

The NorNet Core Testbed — Introduction and Status

Thomas Dreibholz, dreibh@simula.no

Simula Research Laboratory A/S

16 September 2013



Contents

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Overview: Motivation

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Goals of the NorNet Project

- Building up a **realistic** multi-homing testbed
- Wired and wireless
 - Wired → “NorNet Core”
 - Wireless → “NorNet Edge”
- **Perform research with the testbed!**



This presentation: NorNet Core

How to get a realistic testbed for NorNet Core?

Idea: Distribution of NorNet Core over whole Norway

- **Challenging topology:**
 - Large distances
 - A few “big” cities, many large rural areas
 - Svalbard:
 - Interesting location
 - Many polar research institutions
- **Deployment:**
 - 10 sites in Norway
 - 1 site in Essen, Germany
 - More to come ...




Overview: Concepts

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Idea for NorNet Core: Tunnelling

- Researchers require control over used ISP interfaces
 - Which outgoing (local site) interface
 - Which incoming (remote site) interface
- Idea: Tunnels among sites
 - Router at site A: IPs A_1 , A_2 , A_3
 - Router at site B: IPs B_1 , B_2
 - IP tunnel for each combination:
 $A_1 \leftrightarrow B_1$, $A_1 \leftrightarrow B_2$, $A_2 \leftrightarrow B_1$, $A_2 \leftrightarrow B_2$, $A_3 \leftrightarrow B_1$, $A_3 \leftrightarrow B_2$
 - Fully-connected tunnel mesh among NorNet Core sites
 - Each site's router (called **tunnelbox**) maintains the tunnels
 - Static tunnels
 - NorNet-internal addressing and routing over tunnels



Address Assignment

- NorNet Core address spaces:
 - Private NorNet-internal IPv4 “/8” address space
 - NAT+PAT to public 158.39.4.0/24 range
 - Public NorNet-internal IPv6 “/48” address space
 - 2001:700:4100::/48, reachable from the Internet
- Systematic address assignment:
 - IPv4: 10.<Provider ID>.<Site ID>.<Node ID>/24 per site
 - IPv6: 2001:700:4100:<Provider ID><Site ID>::<Node ID>/64
- NorNet-internal DNS setup including reverse lookup



Make it as easy as possible to keep the overview!

Tunnel Realisation


- **Generic Route Encapsulation (GRE) over IPv4**
 - IETF standard (RFC 2784) → should work in existing network
 - Particularly: firewalls, NAT or even **middleboxes**
 - 20+8 bytes overhead (using GRE key, but no seq. number and checksum)
 - MTU: 1472 bytes
- IPv4 over IPv4 may be simpler, but could be more problematic with future sites
- **IPv6 over IPv6**
 - Very simple, adds just another IPv6 header (40 bytes)
 - IPv6 is still “new”, no need to take care of any “grown infrastructure”
 - MTU: 1460 bytes
- **IPv6 in GRE over IPv4**
 - For all IPv6 relations without IPv6 support by ISPs on both sides

Overview: Hardware

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion


A NorNet Core Site Deployment

A usual NorNet Core site:

- 1x switch
- 4x server
 - 1x tunnelbox
 - 3x research systems
- At least two ISP connections
 - Uninett 
 - Other providers
- IPv4 and IPv6 (if available)



Longyearbyen 78.2°N,15.6°E



NorNet Core på Svalbard

Site Deployment Status (September 2013)

Site	Location	Primary ISP	Second ISP
Simula Research Laboratory	Fornebu, Akershus	UNINETT	Kvantel (Hafslund)
Universitetet i Oslo	Blindern, Oslo	UNINETT	— ⁴
Høgskolen i Gjøvik	Gjøvik, Oppland	UNINETT ¹	— ⁴
Universitetet i Tromsø	Tromsø, Troms	UNINETT ¹	— ⁴
Universitetet i Stavanger	Stavanger, Rogaland	UNINETT ¹	— ⁴
Universitetet i Bergen	Bergen, Hordaland	UNINETT ¹	— ⁴
Universitetet i Agder	Kristiansand, Vest-Agder	UNINETT	— ⁴
Universitetet på Svalbard	Longyearbyen, Svalbard	UNINETT ¹	— ⁴
NTNU Trondheim	Trondheim, Sør-Trøndelag	UNINETT	— ⁴
Høgskolen i Narvik	Narvik, Norland	UNINETT	— ⁴
University of Duisburg-Essen	Essen/Germany	DFN	Versatel ^{2,3}

Some work to do ...

- 1) IPv6 available from ISP, but not deployed to NorNet Core site
- 2) IPv6 not available from ISP ☹
- 3) Consumer-grade ADSL connection
- 4) Negotiations in progress

Overview: Software

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Idea: *PlanetLab*-based Software for Experiments

- Key idea:
 - Researchers should get virtual machines for their experiments
 - Like *PlanetLab* ...
 - ... but with multi-homing and IPv6, of course
- *PlanetLab* software:
 - Different “stable” distributions: *PlanetLab*, *OneLab*, etc.
 - Current implementation: based on *Linux VServers*
 - Not in mainline kernel
 - Patched kernel, makes upgrades difficult
 - The future: **Linux Containers** (LXC)
 - Active development by *PlanetLab/OneLab*
 - We are involved in testing the LXC software

The LXC-based *PlanetLab/OneLab* Software

- Researchers get container (sliver) inside a Linux environment
- Same kernel, but slivers are separated from each other
- LXC uses *Open vSwitch*:
 - Slivers are connected to a virtual switch
 - Switch is bridged into real network
 - **Own IPv4/IPv6 addresses** for each sliver!
- Fedora Core 18 Linux environment inside the slivers

Details in the tutorial session tomorrow!

Remote Systems


Our servers may be really remote!

The “road” to Longyearbyen på Svalbard, 78.2°N

Virtualisation

“Anything that can go wrong, will go wrong.”
[Murphy's law]

- Experimentation software is experimental
- How to avoid software issues making a remote machine unusable?
- Idea: virtualisation
 - Lightweight, stable software setup:
Ubuntu Server 12.04 LTS
 - VirtualBox 4.2
 - Other software runs in VirtualBox VMs:
 - Tunnelbox VM on physical server #1
 - 2 LXC-based research node VMs on physical servers #2 to #4
 - In case of problem: manual/automatic restart or reinstall of VM



Experiments with Special Requirements

Special requirements for your experiment? Ask!

- **NorNet Core can satisfy special setup requirements for experiments!**
- Example: VMs with **custom operating system**
 - For example: custom Linux, FreeBSD, AROS, ...
 - Currently still requires manual setup, automation as future work
- Other example: VoIP **SIP honeypot**
 - Security project at University of Duisburg-Essen (UDE)
 - Tunnelboxes tunnel SIP traffic to a central honeypot server at UDE site
 - Analysis of SIP attacks tried on the tunnelbox addresses at different sites



Overview: Management

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Software Packages Overview



Server
(base package for physical machine)



Node
(base package for virtual system)



Tunnelbox



File Server



Monitor



Display



Development
(package with useful development tools)

Boot logos to visualise different system roles

Software Packages Contents

- **Server:**
 - VirtualBox host system and VM management (watchdog, backups, ...)
 - Basic server configuration scripts
- **Node:**
 - VirtualBox guest system
 - Basic node configuration scripts (software updates, backups, ...)
 - Service Nodes:
 - **Tunnelbox**
 - **File Server:** a NFS and TFTP file server for NorNet Core
 - **Monitor:** network monitoring software
 - **Display:** network visualisation via web browser
- **Development:** to equip Server or Node installations with development tools

Packages are conveniently available in an Ubuntu PPA

Monitoring

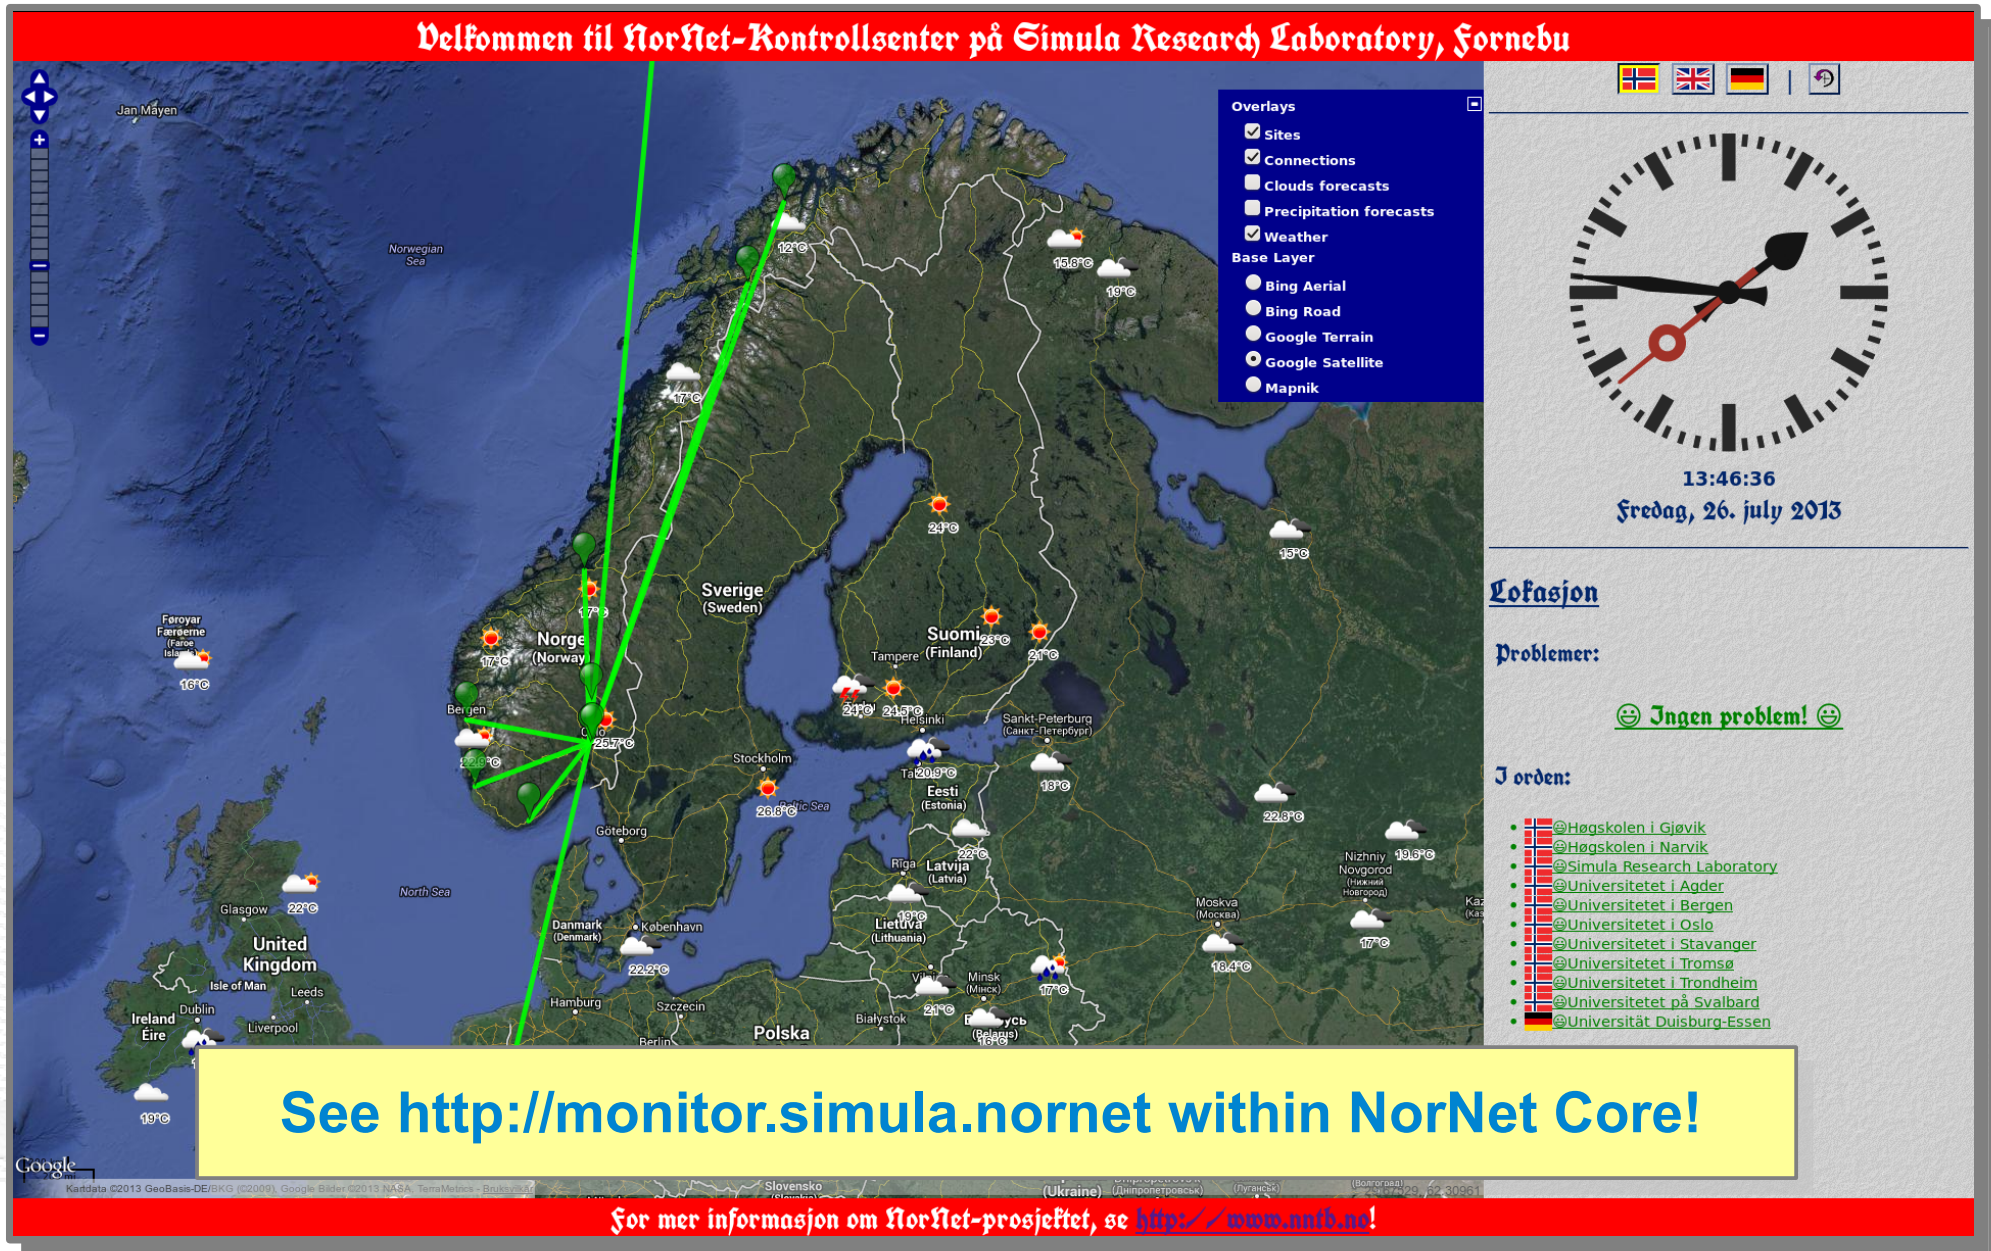
- *PlanetLab*:
 - 575 nodes of 1042 nodes working (September 9, 2013)
⇒ availability ca. 55% ☹
 - NorNet should do much better!
- Direct contact to technical staff/researchers at sites
- Monitoring using Nagios
 - Flexible
 - Extensible by service-specific plug-ins

The Nagios logo, featuring the word "Nagios" in a bold, black, sans-serif font. The letter "N" is stylized with a horizontal bar extending to the left. A registered trademark symbol (®) is located at the top right of the "s".

Nagios[®]

How to visualise NorNet Core?

“Kontrollsenteret”



Overview:

Research and Users

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Research and Users

“The road to hell is paved with unused testbeds.”

[James P. G. Sterbenz]

- Of course, NorNet does **not** intend to be another unused testbed!
Goal: “NorNet wants to be a building block of the railroad to heaven.”
- NorNet will be open for all interested researchers!
 - Similar to *PlanetLab* ...
 - ... but with higher node availability and tighter monitoring
 - ... and, of course, **multi-homing** and **IPv6**
- **Particularly, it can also be used by you!**

More details here at the NNUW-1!

How to use NorNet Core?

**Join our tutorial sessions!
Here at the NNUW-1!**

- Contents:
 - Get access to NorNet Core
 - User and slice management
 - Access to slices
 - Using and configuring slivers with own software
 - How to make use of multi-homing?

Overview:

Conclusion

- Motivation
- Concepts
- Hardware
- Software
- Management
- Research and Users
- Conclusion

Conclusion and Future Work

- NorNet is progressing!
 - Initial site deployment completed
 - Software setup ready to start experiments
- Future work:
 - Make sites multi-homed
 - Some additional sites
 - Improve/refine management software
 - Get more users

And, of course, do some research!

Any Questions?

N  RNET

Visit <http://www.nntb.no> for further information!