


# The 2nd International NorNet Users Meeting (NNUW-2)

## The NorNet Core Testbed — Introduction and Status in August 2014

**Thomas Dreibholz, [dreibh@simula.no](mailto:dreibh@simula.no)**

Simula Research Laboratory

**28 August 2014**



# Contents

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

# Overview: Motivation

- Motivation
- Concepts
- Hardware
- Software
- Users and Research
- 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?




# Overview: Concepts

- Motivation
- Concepts
- Hardware
- Software
- Users and Research
- 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-internal address spaces:
  - Private NorNet-internal IPv4 “/8” address space (NAT to outside)
  - Public NorNet-internal IPv6 “/48” address space
- Systematic address assignment:
  - IPv4: 10.<Provider ID>.<Site ID>.<Node ID>/24 per site
  - IPv6: 2001:700:4100:<PP><SS>::<NN>/64  
(PP=Provider ID; SS=Site ID; NN=Node ID)
- 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
- **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


**Information on routes? Yes → talk by Forough Golkar**

# Overview: Hardware

- Motivation
- Concepts
- Hardware
- Software
- Users and Research
- 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)

Additional researcher-provided sites:

- Varying configurations
- VM setups, powerful servers, “retro-style” PCs ...



**UNIS**

Longyearbyen 78.2°N,15.6°E



NorNet Core på Svalbard

# Site Deployment Status (September 2013)

No.	Site	ISP 1	ISP 2
1	Simula Research Laboratory	Uninett	Kvantel
2	Universitetet i Oslo	Uninett	-
3	Høgskolen i Gjøvik	Uninett	-
4	Universitetet i Tromsø	Uninett	-
5	Universitetet i Stavanger	Uninett	-
6	Universitetet i Bergen	Uninett	-
7	Universitetet i Agder	Uninett	-
8	Universitetet på Svalbard	Uninett	-
9	Universitetet i Trondheim	Uninett	-
10	Høgskolen i Narvik	Uninett	-
11	Universität Duisburg-Essen	DFN	Versatel

There was a lot of work to do ...


IPv4 and IPv6  
 ISP negotiation in progress

IPv4 only (ISP without IPv6 support ☹)  
 IPv4 only (site's network without IPv6 support)

# Site Deployment Status (August 2014)

No.	Site	ISP 1	ISP 2	ISP 3	ISP 4
1	Simula Research Laboratory	Uninett	Kvantel	Telenor	PowerTech
2	Universitetet i Oslo	Uninett	Broadnet	PowerTech	
3	Høgskolen i Gjøvik	Uninett	PowerTech		
4	Universitetet i Tromsø	Uninett	Telenor	PowerTech	
5	Universitetet i Stavanger	Uninett	Altibox	PowerTech	
6	Universitetet i Bergen	Uninett	BKK		
7	Universitetet i Agder	Uninett	PowerTech	–	
8	Universitetet på Svalbard	Uninett	Telenor		
9	Universitetet i Trondheim	Uninett	PowerTech		
10	Høgskolen i Narvik	Uninett	Broadnet	PowerTech	
11	Høgskolen i Oslo og Akershus	Uninett	–		
12	Karlstads Universitet	SUNET			
13	Universität Kaiserslautern	DFN			
14	Universität Duisburg-Essen	DFN	Versatel		
15	Hainan University	CERNET	China Unicom		
16	The University of Kansas	KanREN			

A significant progress!

 IPv4 and IPv6

 ISP negotiation in progress

 IPv4 only (ISP without IPv6 support ☹)

 IPv4 only (site's network without IPv6 support)

<https://www.nntb.no/pub/nornet-configuration/NorNetCore-Sites.html>

# Some Site Statistics (August 2014)



Active Sites	16
Distinct ISPs of Active Sites	13
Distinct Countries of Active Sites	5
Total IPv4 Interfaces	34
Total IPv4 Tunnels	561
Total IPv6 Interfaces	22
Total IPv6 Tunnels	231
Inactive Sites	0

<https://www.nntb.no/pub/nor-net-configuration/NorNetCore-Sites.html>

# Next Steps for Deployment

- More IPv6 connectivity
  - RFC 2460 soon celebrates its 16<sup>th</sup> anniversary
  - Providers must support it, of course
  - May be ask sites for tunnel to site's IPv6 connection?  
(temporary fix until native connectivity is deployed in site's network)
- Some more ISPs
  - Diversity: cable TV, satellite, ...
- Some more sites
  - Hosted by interested researchers in other countries



- What about your country?

# Overview: Software

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



# 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 4.3
  - 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



# *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 developing and 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!**

# Development Status 2013

- Customised Ubuntu Server 12.04 LTS for physical machine installations
- Using *PlanetLab* LXC software distribution for research nodes
  - Based on Fedora Core 18
  - From upstream project's nightly builds at *OneLab*
  - URL: <http://build.onelab.eu/lxc/>
- Extensions based on tags in PLC configuration
  - Python-based management scripts
- Custom additional software packages
  - NetPerfMeter, SubNetCalc, RSPLIB, tsctp, etc.

Inflexibility by dependency  
on upstream builds

# Development Status 2014

- **Customised “NorNet” distribution** of *PlanetLab*'s “LXC” distribution
  - Git forks of upstream *PlanetLab* LXC repositories
  - NorNet-customisations, including **kernel with MPTCP support**
  - Cooperation with *PlanetLab/OneLab* for development!
- **Own “build and test” infrastructure**
  - Build servers: *queenstown*, *arrowtown*, *cromwell*  
(to compile the full distributions “NorNet” and “LXC”)
  - Test master: *earnslaw*  
(for automated tests of the builds)
  - LXC test (for PLC): *wakatipu*
  - KVM test (for research nodes): *bjordammen*
  - Publication server: *benlomond*

Flexibility plus  
automated test runs

**Nightly builds of “NorNet” and “LXC”:** <http://benlomond.nntb.no/>

# Next Steps for Development

- Research software refinements
  - Add NorNet features to PLC web interface
- VPN access to NorNet Core network
  - More convenient access (not just via SSH gateway or local switch)
- Other ideas
  - KVM-based virtualisation
  - OpenStack?
  - Direct support for booting custom kernels
  - Testbed federation

**Let us discuss your ideas and suggestions!**

# Overview:

## Users and Research

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



# Users and Research

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

**[James P. G. Sterbenz]**

- We already got some users!
- Examples:
  - Shared Bottleneck Detection (UiO+Simula)
  - VoIP Misuse Detection (UDE)
  - Multi-Path Transport (Simula, UDE, UiO, HU, etc.)
  - Balia Congestion Control (Bell Labs in South Korea)
  - IPv4/IPv6 Performance Comparison (Simula)
  - ...

Some interesting talks  
at the NNUW-2!

See <https://www.nntb.no/projects/> for further projects using NorNet!

**Next step: get even more users!**

# The “NorNet World Tour 2014”



- 01/2014: Centre for Advanced Internet Architectures (CAIA) at Swinburne University  
Melbourne, Victoria/Australia
- 05/2014: Polytechnic School of Engineering at New York University (NYU)  
Brooklyn, New York/U.S.A.
- 05/2014: PlanetLab Consortium at Princeton University  
Princeton, New Jersey/U.S.A.
- 05/2014: University of British Columbia (UBC)  
Vancouver, British Columbia/Canada
- 09/2014: Kungliga Tekniska högskolan (KTH Royal Institute of Technology)  
Stockholm/Sweden
- 10/2014: Academy, Industry and Government of the Hainan Province  
Haikou, Hainan/China
- 10/2014: Tsinghua University [planned]  
Beijing/China
- 12/2014: NorNet demo presentation at the IEEE GLOBECOM  
Austin, Texas/U.S.A.
- 01/2015: ... [planned]  
[planned]/Australia



**Interested in a NorNet presentation? Just ask!**

# Collaborations

- PlanetLab/OneLab
  - Development and testing of the research software
  - URLs: <https://www.planet-lab.org>, <https://www.onelab.eu>
- RIPE Atlas
  - Connectivity and reachability measurements
  - URL: <https://atlas.ripe.net>
  - Node deployed at site in Longyearbyen
- Seattle
  - Open Peer-to-Peer Computing, project at NYU
  - URL: <https://seattle.poly.edu>
  - Running inside NorNet Core slice
- ToMaTo
  - Topology Management Tool
  - URL: <http://tomato-lab.org>
  - Part of the G-Lab testbed



You can use NorNet Core as well, of course!

**Join our tutorial session!  
Here at the NNUW-2!**

- 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?

The 5th International Workshop on Protocols and Applications with Multi-Homing Support

# PAMS 2015

**See <https://simula.no/pams-2015!>**

**[Submission deadline: October 1, 2014](#)**

**March 2015, Gwangju/South Korea**

**In conjunction with the 29th IEEE AINA**

# Overview: Conclusion

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

“NorNet wants to be a building block of the railroad to heaven” ...



... and not be another unused testbed that paves the road to hell!

Any Questions?

N  RNET

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