

bifrost

KTH/CSD course kick-off
Fall 2010

Robert Olsson

What is bifrost?

Small Linux suited for USB

For infrastructure

For research & education

Name of human network and collaboration

Why a new Linux distro? We've been around
For 15 years soon.

What about Red Hat/Debian Ubuntu etc?

For who?

Networking people

Unix/Linux

Needing a small flexible distro

Research & Education

Objectives infrastructure?

Hardware selection – Crucial

Software selection - Crucial

Testing - Crucial

Bug Fixes/Support Crucial

Development - Crucial

Basic functions?

Routing

Firewalling

Login services

Traffic logging

Gateways etc

Virtualization and/or Namespaces (New)

Ipv4, ipv6

Network testing pktgen/netperf etc

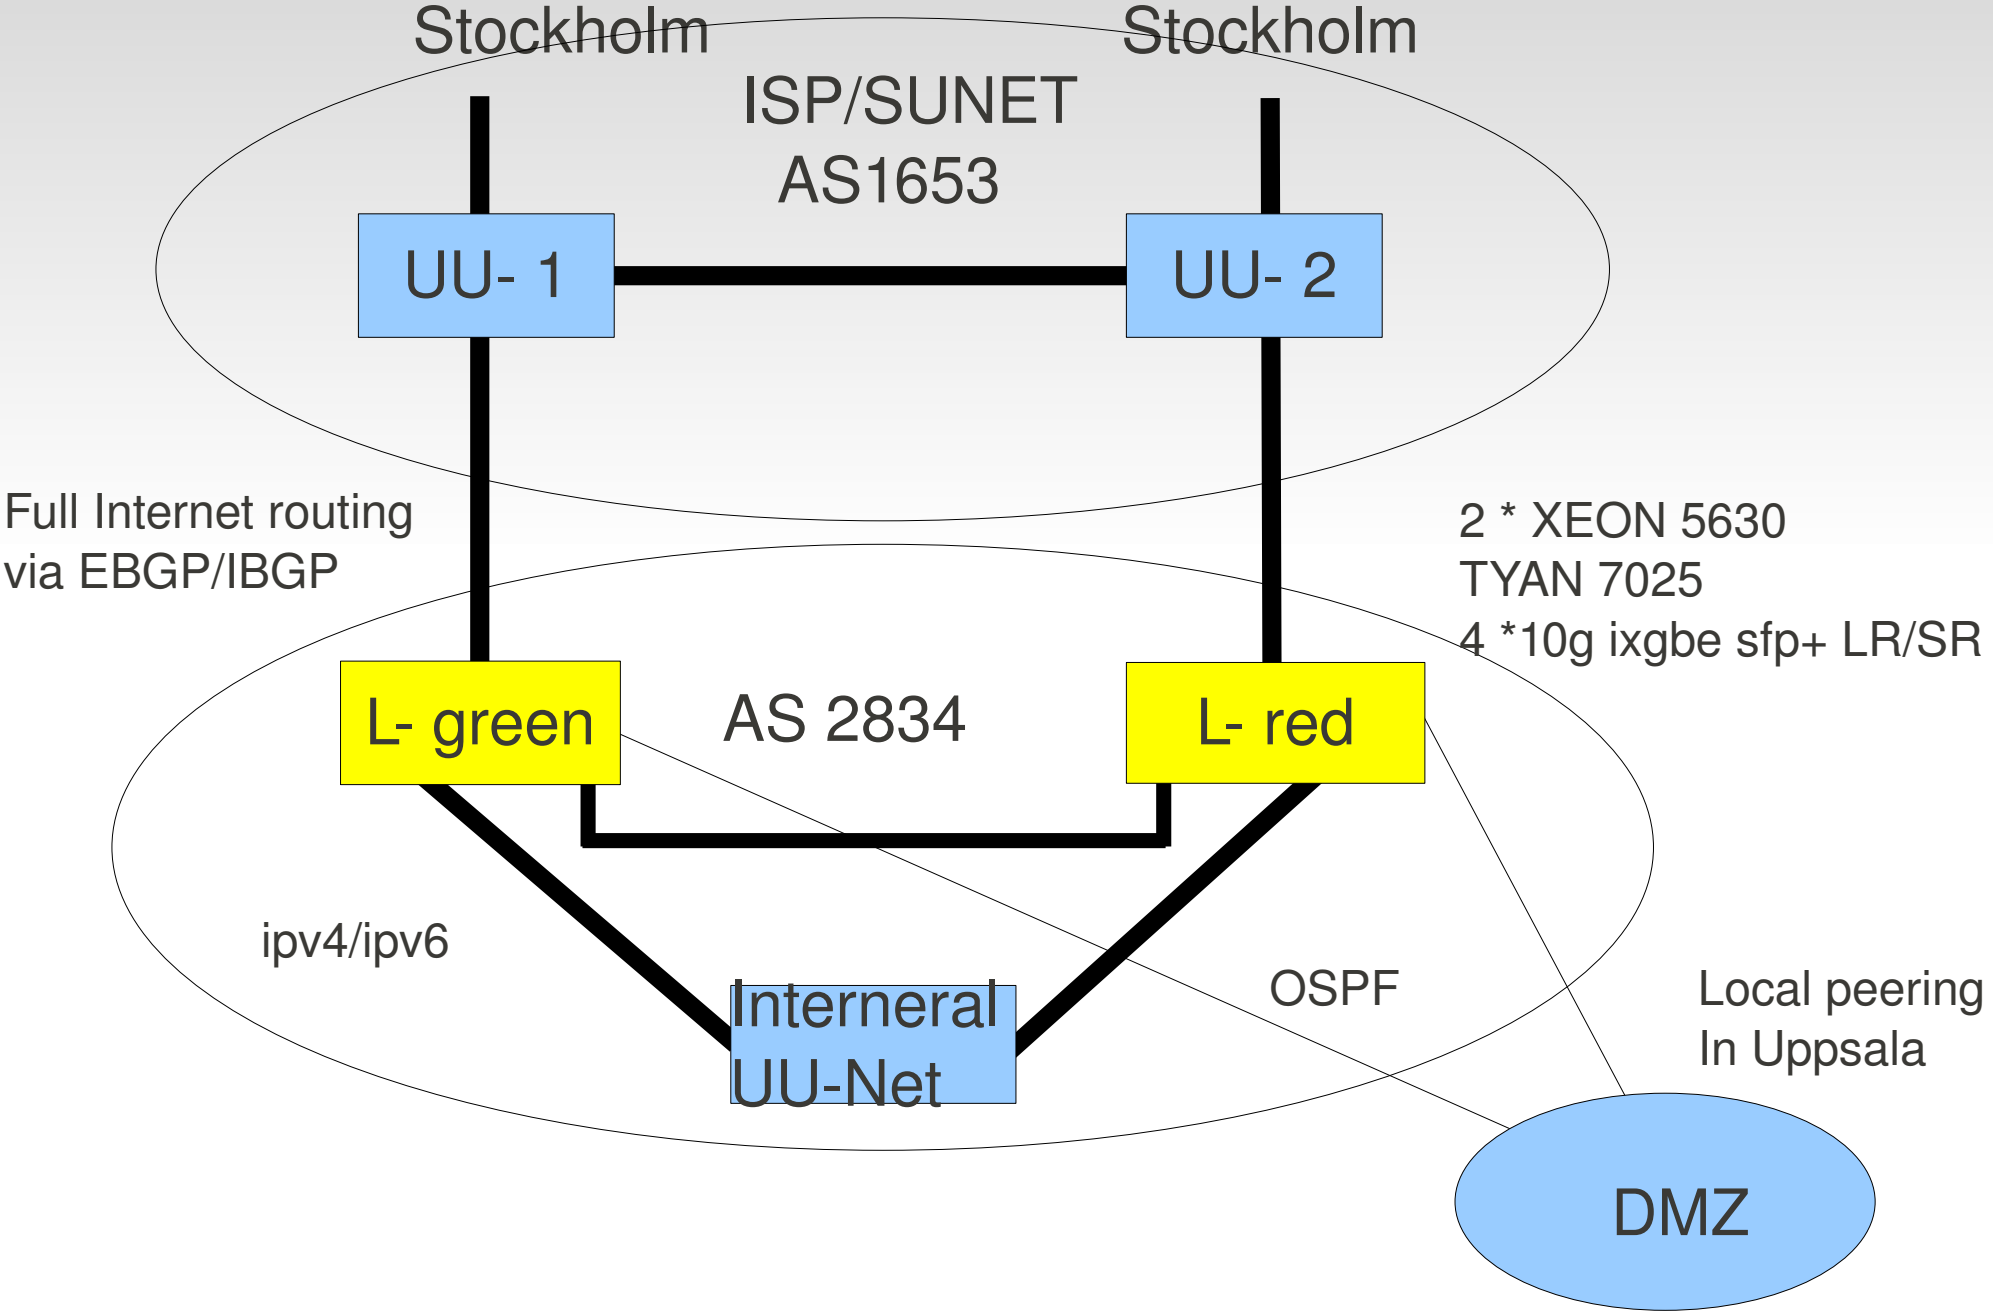
Routing deamons

Packet forwarding is done in Linux kernel
But routing protocols is run by userpace
Deamons

Currently tested versions of quagga
Bgp, OSPF both Ipv4, ipv6
Cisco API

Of course other software can be used
Bird?

More than 10 year in production at Uppsala University




R & D related to bifrost

Close collaboration with Linux networking
Developers and industry

NAPI (3 years) now in most network drivers
Pktgen testing in linux,
fib_trie, (routing algo)
routing stats to monitor and understand
network


Etc, etc

Hi-End Hardware



XEON 2 x E5630
TYAN S7025 Motherboard
Intel 82599


Block hw structure



Recent project IIS

- DMA 93 Gbit/s TX
- Separation based HW classifier. BGP and ssh
In control plane (CPU0)


Cache effect/Performance



Overall Effect

- Inelegant handling of heavy net loads
 - System collapse
- Scalabiity affected
 - System and number of NICS
 - A single hogger netdev can bring the system to its knees and deny service to others

Summary 2.4 vs feedback



March 15 report on lkml

Thread: "How to optimize routing performance"

reported by Marten.Wikstrom@framsfab.se

- Linux 2.4 peaks at 27Kpps

- Pentium Pro 200, 64MB RAM


Lab testing HW and SW



Not all were selected...



Hardware - NIC



Intel 10g board Chipset 82599

Open chip specs. Thanks Intel!

KTH/CSD

Current focus

Optical to Open Source Router
Low-Power & Renewable Energy
Virtualization/Namespaces
Performance
Cost

Optical modules

Optical sender and receiver in one module



SFP 1G

XFP 10G

SFP+ 10G

DOM - Optical Monitoring



Optical modules can support optical link monitoring
RX, TX power, temperatures, alarms etc


Newly added support to Bifrost/Linux

CWDM MUX/DEMUX 4 Ports

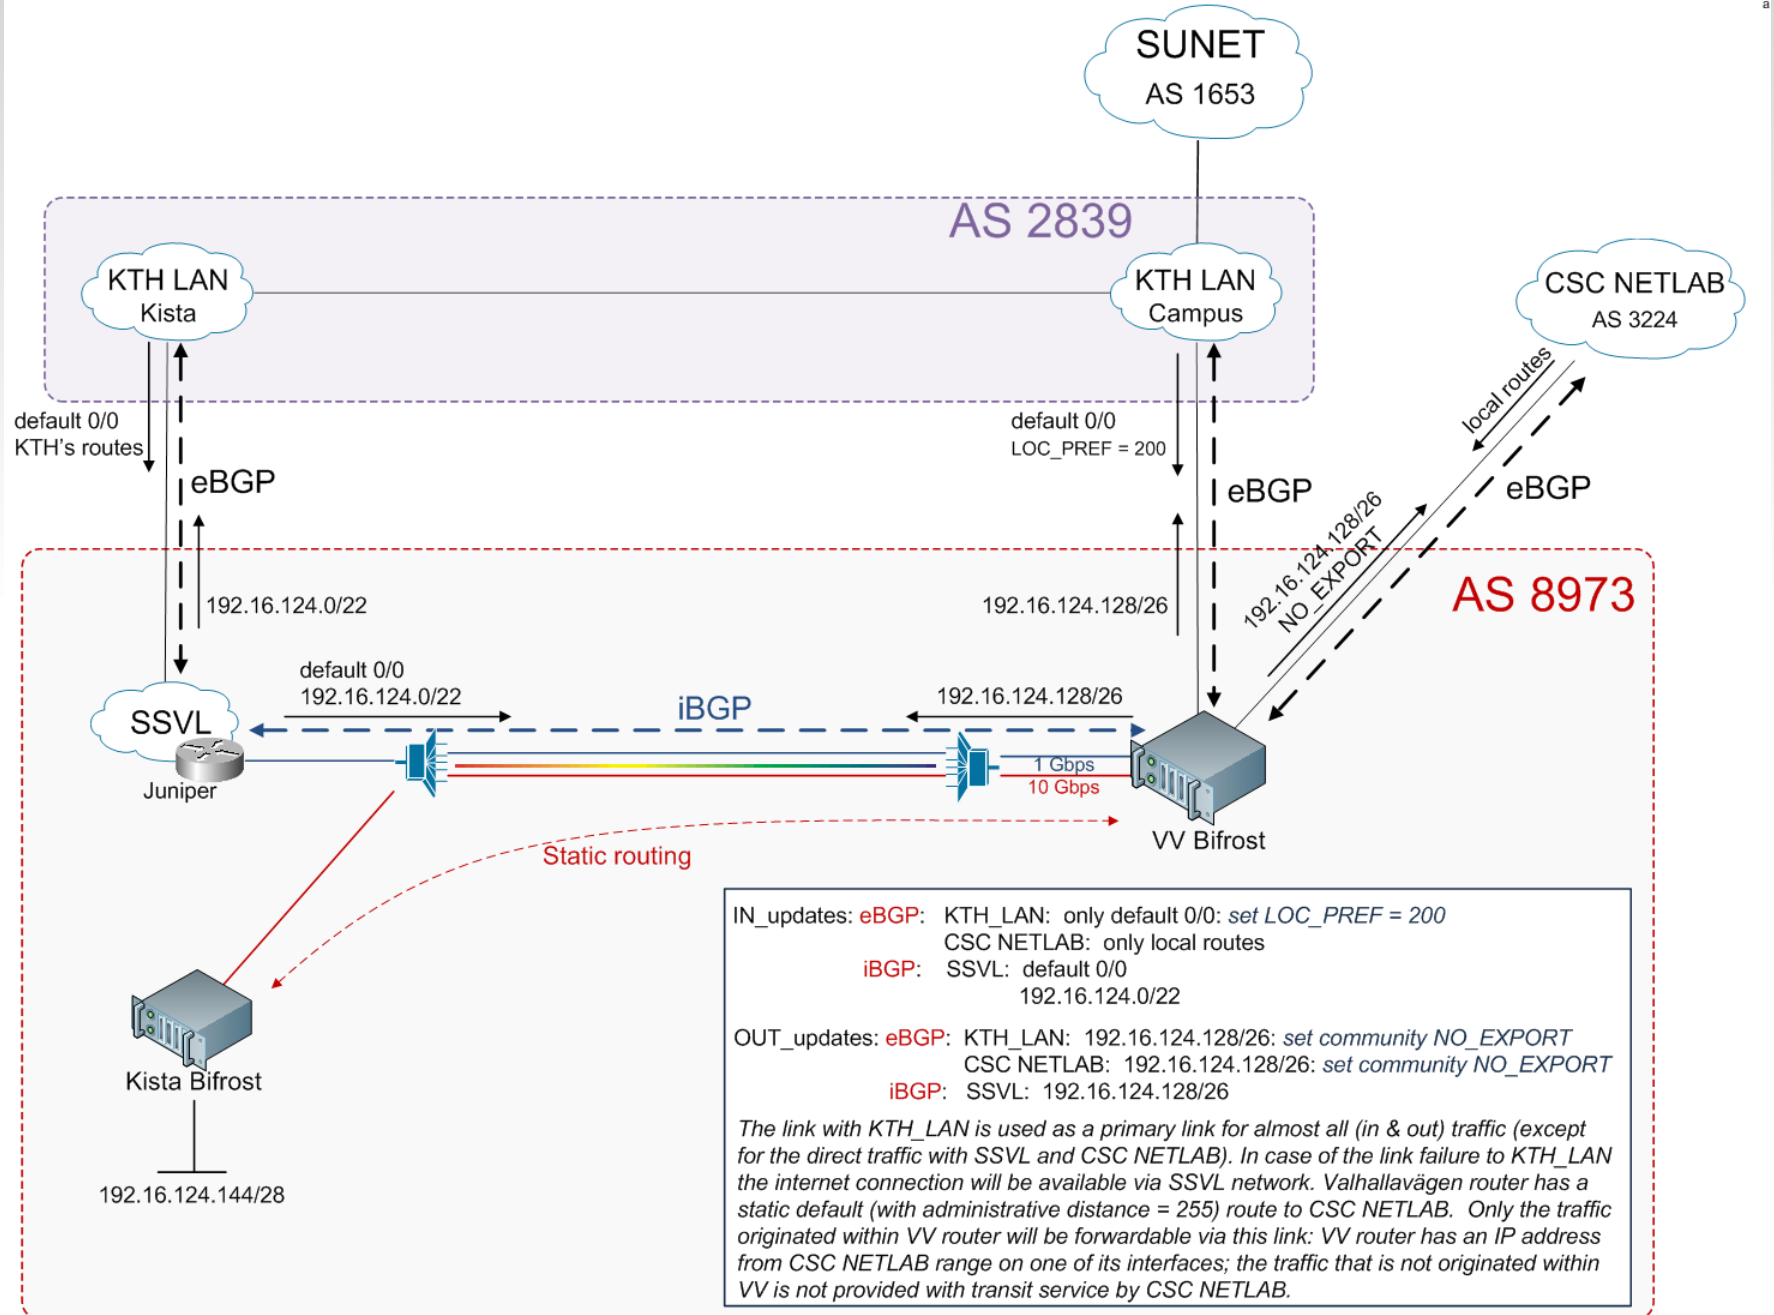


Price idea 600 Euro

CWDM MUX/DEMUX 16 Ports



Usage



That's all

Questions?

Project's ideas

DOM 4-port GIGE card/Linux igb driver

GateWay Optical/WiFi

Energy – Router Power Reduction/Linux

Energy – Ultracaps/Solar or UPS