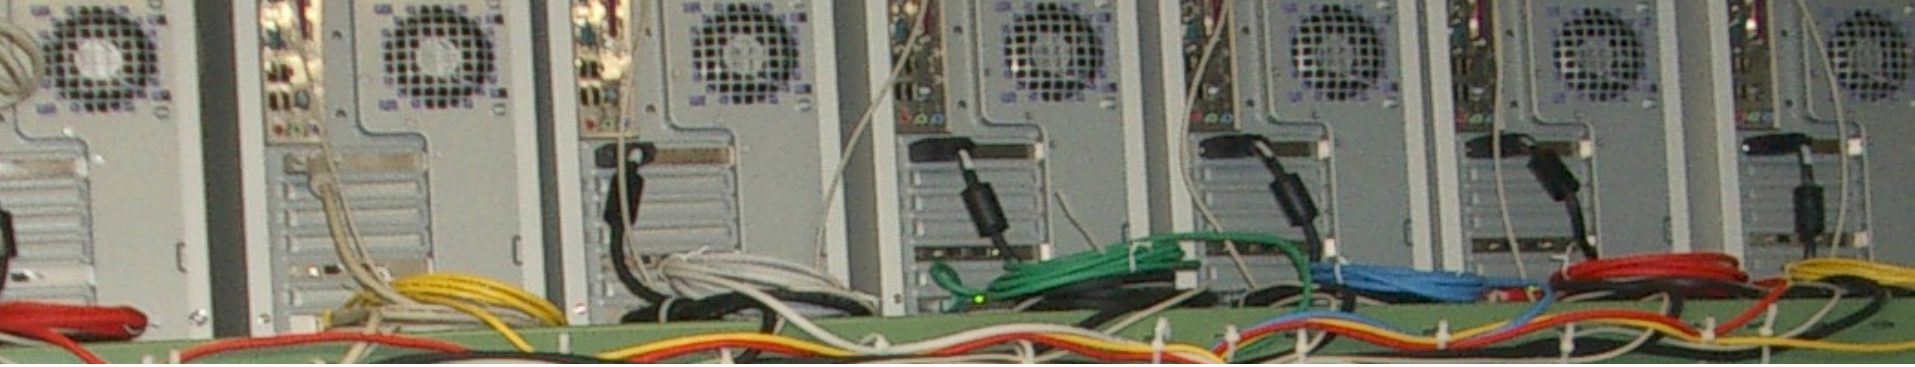

# Inauguration of the TIK Experimental Cluster "Scylla"

January 22nd, 2004

TIK, ETH Zürich, Switzerland

Presentation by:
{duebendorfer, wagner}@tik.ee.ethz.ch

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

# „Scylla“ - TIK Experimental Cluster

**What “Scylla” is:**

An experimental research computer cluster of 22 identical nodes and one additional gateway node.

**Initiators of this cluster:**

- Design, planning, cabling and software installation by Arno Wagner
- Ordering and component tracking by TIK Dienstgruppe
- Node assembly by Alain Breedt
- Financing by TIK and D-ITET

**Cluster administration and support:**

Arno Wagner <wagner@tik.ee.ethz.ch>

# How „Scylla“ was born

## A short historical review

**July 2003**: First ideas for a TIK cluster by the DDoSVax team due to high computing power needs for Blaster/Sobig.F worm traffic analysis

**August 26th, 2003**: Arno Wagner presents a first hardware proposal for a cluster (variants with 11, 14, 19, and 22 nodes).

**Sept. 26th, 2003**: TIK IL-Meeting decides to buy a 12 node cluster

**October 20th, 2003**: KIM accepts cluster project with 22 nodes (36 kFr); recommends it for D-ITET Dep. Conference (DK)

**October 21st, 2003**: Components for cluster are ordered

**Mid November 2003**: Start of cluster assembly, cabling; installation

**January 22nd, 2004**: Inauguration of the cluster

# Scylla's Anatomy

**22 cluster nodes:**

- Athlon XP "Barton" 2.8 GHz
- 1 GB RAM, 120 GB Harddisk
- CD-ROM Drive
- 1 Gbit/s-Ethernet

**One single gateway node:**

- Athlon XP "Barton" 2.8 GHz
- 1 GB RAM
- 2x 200GB Harddisk (200 GB RAID-1 mirrored)
- CDROM
- 1 Gbit/s-Ethernet internal; 100 Mbit/s external

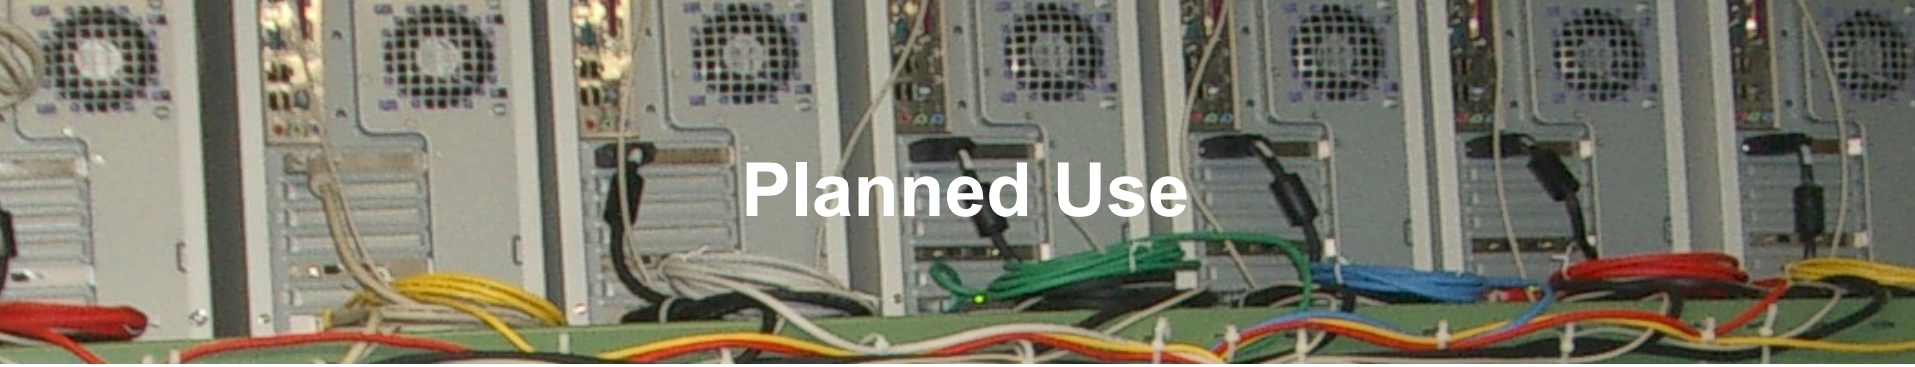

# Planned Use

## Planned use of the cluster

- Large computing power for network traffic analysis (DDoSVax)
- Massive distributed simulations (P2P nodes)
- Experimental kernel installations
- Gigabit/s network testing
- Computer security related research

# Naming

**„Scylla“**

was an attractive nymph and a daughter of Phorcys. One day, the sea-god Glaucis fell passionately in love with her. However, she rejected him.

Therefore, Glaucis asked the sorceress Circe to create him a love potion. Bad enough, Circe instantly fell in love with Glaucis, who rejected her.

Filled with jealousy against Scylla, Circe poured a potion of herbs into the water, where Scylla was bathing in, and then cast her spell. Suddenly, out of Scylla‘s lower body half six monstrous dog heads grew.

**We named the cluster „Scylla“** as the cluster‘s possiblities are fascinating at first, however controlling it efficiently is a delicate task.

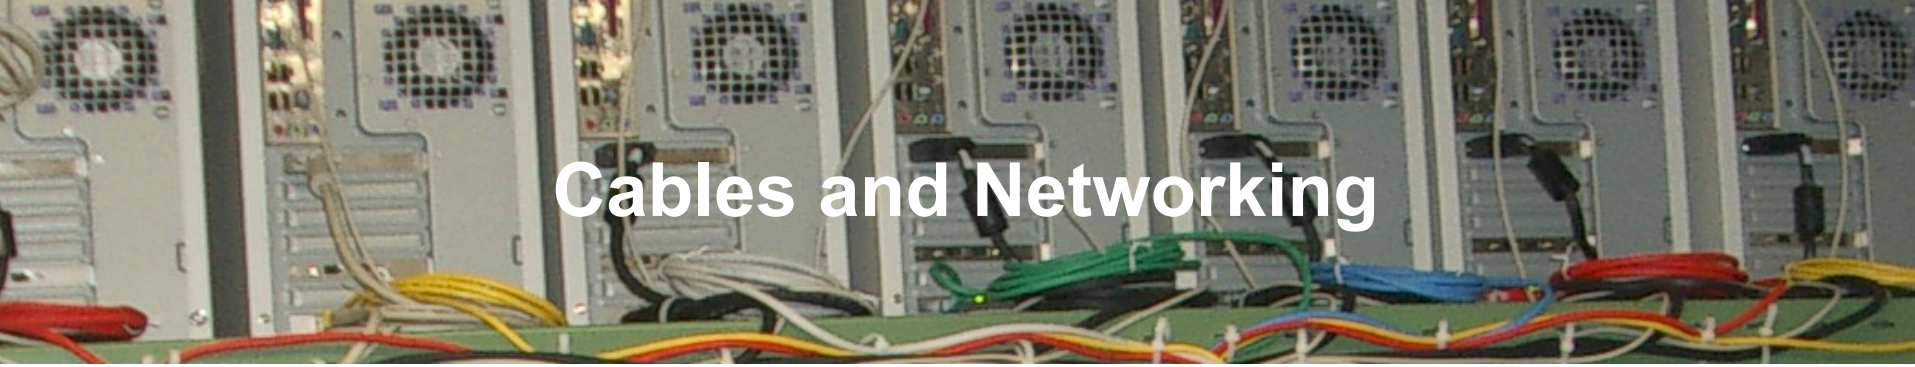

# Cables and Networking

**Cables:**

In total ca. 260 m, thereof:

70 m of Ethernet Cat5e Cable

130 m of Video and keyboard cables

60 m of power cables

And about 200 cable ties.

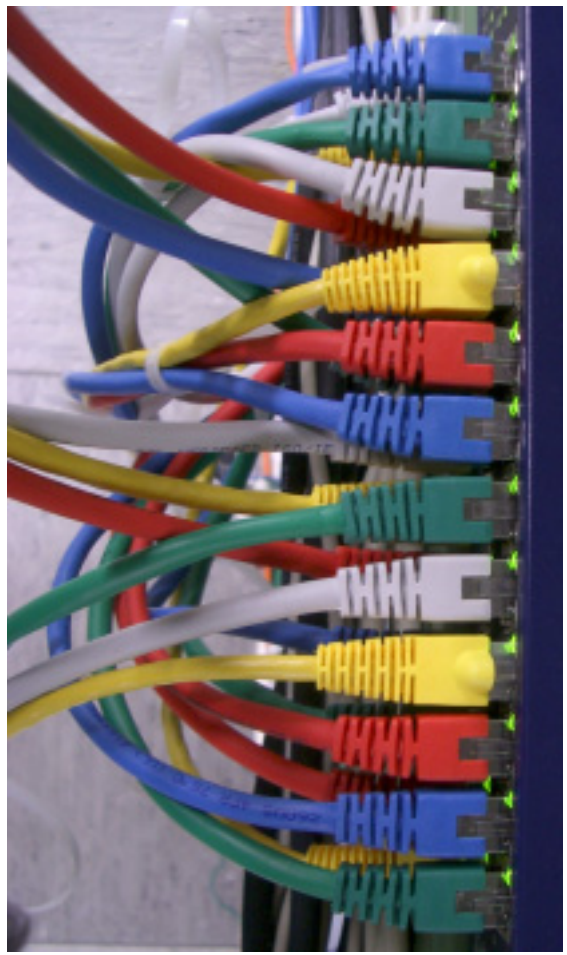

**Networking:**

- 24 port switch full-duplex at 1 GBit/s;
  20 Gbit/s aggregate bandwidth
- Cluster nodes have private IPs 10.0.0.0/24;
- Gateway node: scylla.ethz.ch (external interface)

# Software Installation

**Nodes n01-n22:**

Debian Linux 2.4.22 with the

openMosix cluster software extension

**Gateway:**

Debian Linux 2.4.24; role as Debian-testing mirror for nodes

**Software distribution to nodes (from gateway)**

FAI (Fully Automatic Installation): Boot via Etherboot (floppy), kernel is fetched via tftp from gateway.

→ Allows simple and individual change of boot kernel of the nodes from the gateway.

# Electrical Power Consumption

**One node:**

Approx. 150W power consumption

**Full cluster** (gateway + 22 nodes):

Approx. 3.5 kW power consumption
(and heat generation!)

Cluster has its own dedicated
3 * 25A power line and circuit breaker

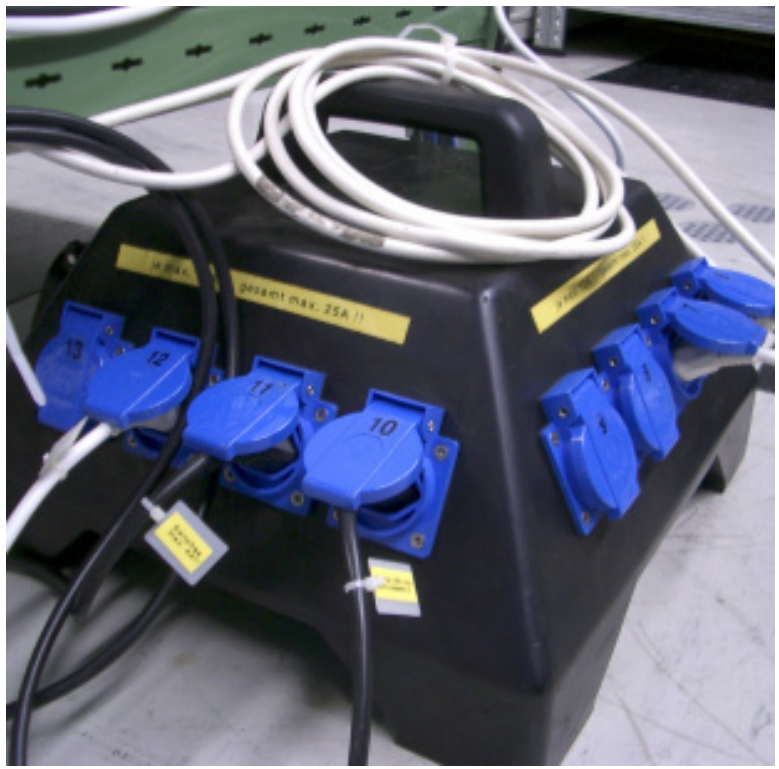

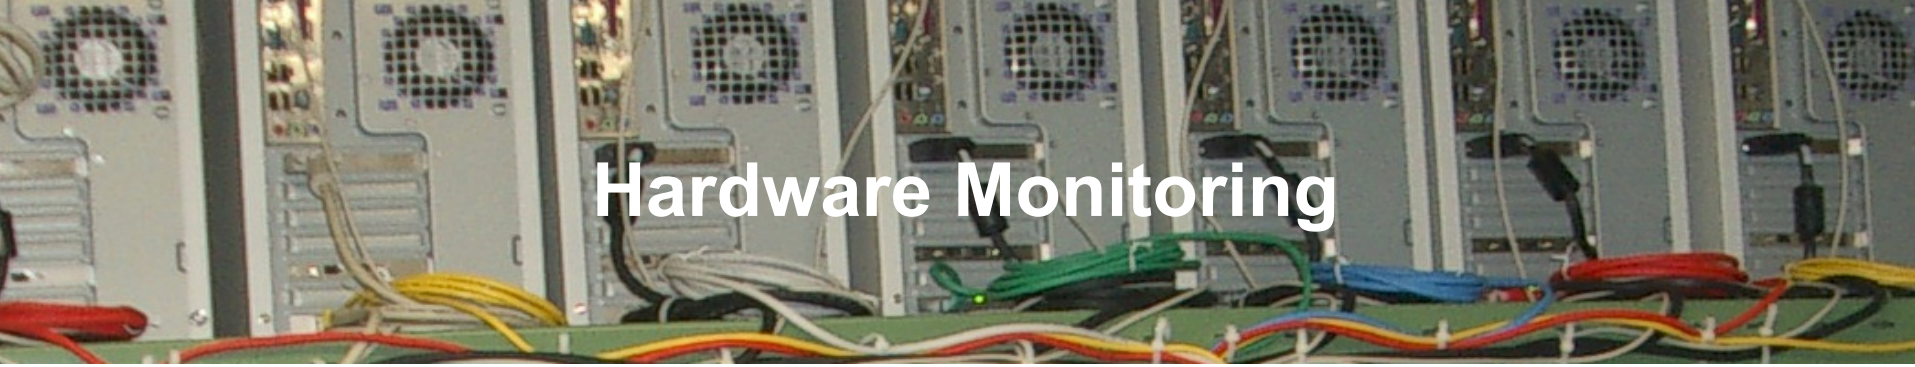

# Hardware Monitoring

## Harddisks failure

SMART monitoring of the harddisk drive in each node by the gateway; detects HDDs about to fail.

## Temperature alerts

Temperature monitoring of each node by the gateway, and automatic shutdown of individual nodes or the whole cluster in case of overtemperature.

# Varia

- Gateway: scylla.ethz.ch
- Nodes: n01-n20 (n21,n22 in PlanetLab temporary)
- n01 is Test node with remote power switch and serial console
- “Experimental” cluster means that OS installation and software configurations can be changed
- Complete new installation with FAI in less than 10 min.
- 100 GB persistent data storage/node; a new installation only changes 10 GB system/swap space partition

# Thanks for your attention!

Feel free to ask any questions.