QMUL Site Report 2023

Dimitrios Toroz

School of Physical and Chemical Sciences

Department of Physics & Astronomy
+
Department of Chemistry
=
80 Academics, 31 Researchers,10 Technical,
19 Professional Services

Scientific Computing Team

6 Strong but due to staff changes been

2 people down for 18 months.

Dimitrios started in February

New person due to start end of August

- all being well !!! 🤞

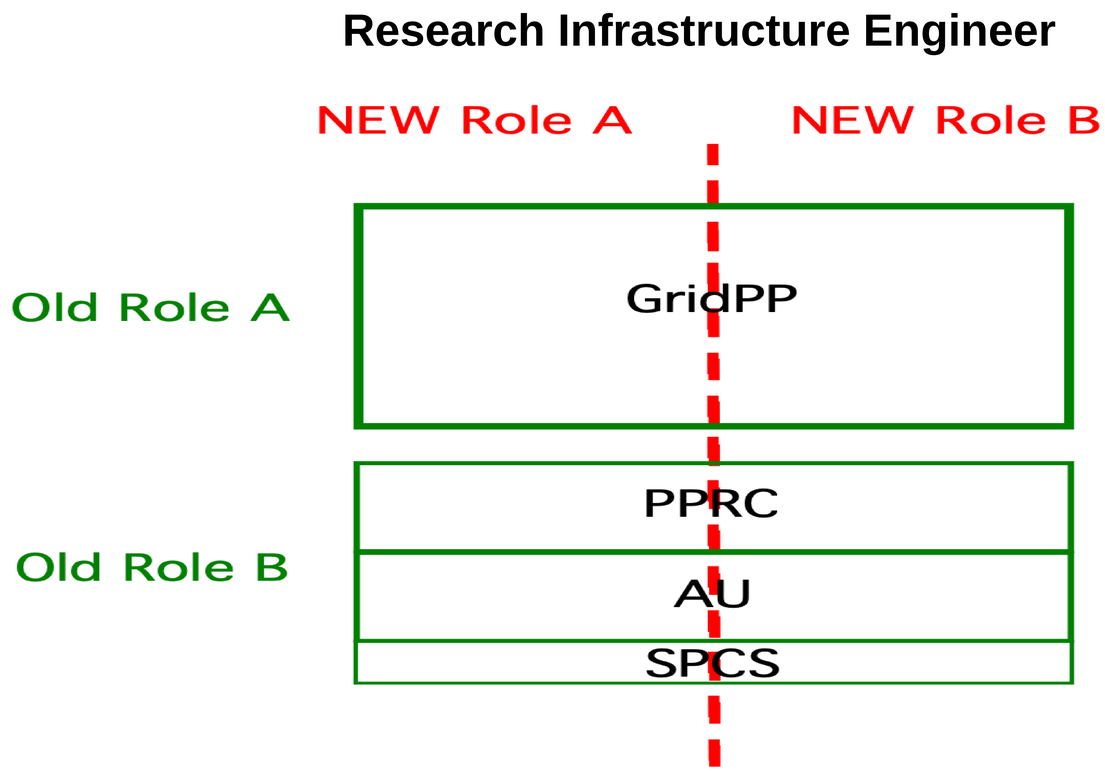

SPCS Computer Room

- 8 Racks, Aircon 3+1 redundancy.
- Campus 100Gb/s Janet Link (redundant pair)
- QMUL Core network 20Gb/s Optical Link to SPCS Computer Room
- 10Gb/s Mikrotik Switches as top of rack switches
- 2x 16A power to Racks each on 3000VA APC UPS
- Dumb PDU + IPMI (a few legacy without IPMI)

- Backlog of commissioning work due to staff changes
- New Proxmox Cluster - DONE(ish)
  3 nodes: Dell R540 40 threads, 96G, 32TB
- New Backup - DONE(ish)
  40 way LTO8 Tape Changer
- Storage Servers
  - ready to commission with new staff
- New School Combined Linux Infrastructure - In planning

Owl for remote meetings

Meeting OWL Pro
owllabs.com

Four of them available in the school.

Popular to academic staff.

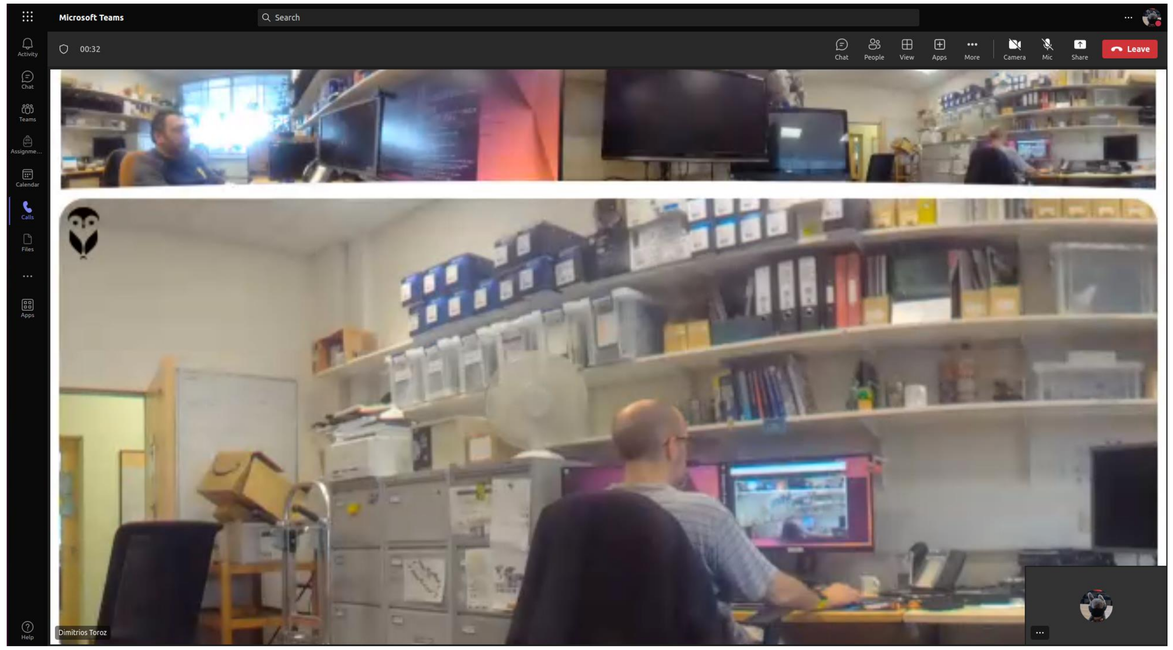

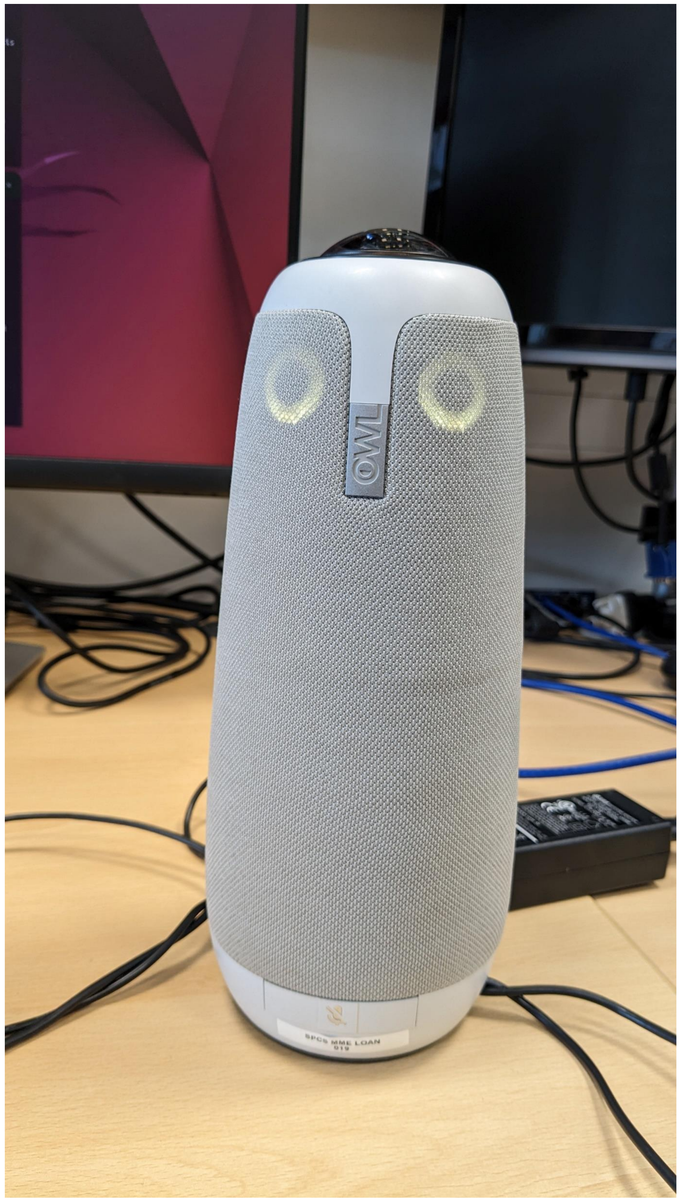

GridPP Resources

- Compute: 2 ARC CEs Total ~17,000 job slots, 4GB RAM per slot
  - 120 * Dell R630 (Intel Gold 6248R)
  - 50 * Lenovo SR570 (intel Gold 6252)
  - 8 Lenovo SR665 AMD EPYC 7402 +2 A100 + 1TB of RAM (One node with 4TB of RAM!)
- Storage: 3 StoRM SEs on top of Lustre (all use hardware raid 6) 13.4PiB
  - 4HGST 60 disk JBODs: total online 1.5PiB
  - 17 Dell 740XD2s 24 disk servers: total online 3.6PiB
  - 8 Lenovo D3284 84 disk JBOD: total online 8.3PiB

GridPP Cluster Basics

- Slurm as the batch system. Lustre for storage.
- Kia DHCP, PowerDNS, ansible + cobbler for deployment.
- Libre NMS monitoring. APC Netbotz for temperature monitoring
- Proxmox for VMs. VM backups onto old NFS server.
- Mellanox 100 G switches (SN2410, SN2010, SN2700) for networking.
- Arc + StoRM for Grid.

Migration Plan

- New 3 node PROXMOX cluster for VMs.
- NIS to LDAP (+Kerberos?)
- Migration Rocky9 for worker nodes.
- Migrate as much other nodes to Rocky 9 (StoRM, Arc, ansible…)
- New switch based on SoniC OS: Dell S5448 (48*100Gb/s + 8*400Gb/s). 100G ports are SFP6-DD, backward compatible with SFP28 (25Gb/s).
- Looking at moving to a layer 3 network across the cluster.
- New 120 TB NVMe based NFS server (will use ZFS) - home + arcce shared session directory + scratch.

DC Refurbishment Plan

- 10KW(average) per rack, 39 racks, 3 rows, dual 32 amps power feeds per rack (x2).

- Chilled water circuit. 4+1 in row cooling per row, 17C water in, 23C out. High level piping. 26C cold / 41C hot aisle air temperature.

- Use heat pumps to reuse heat for district heating system. Reduce costs and carbon footprint.

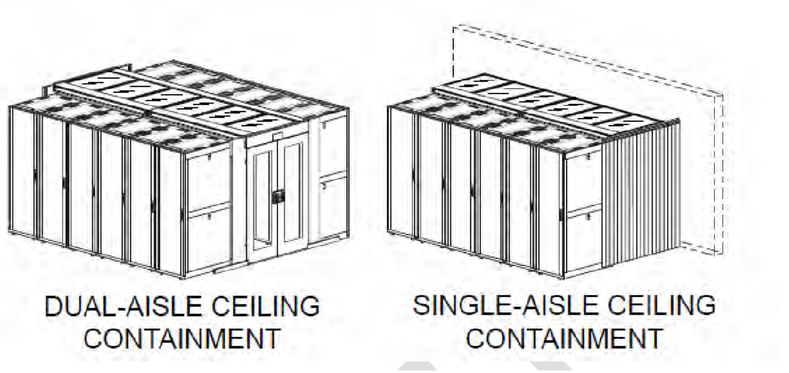

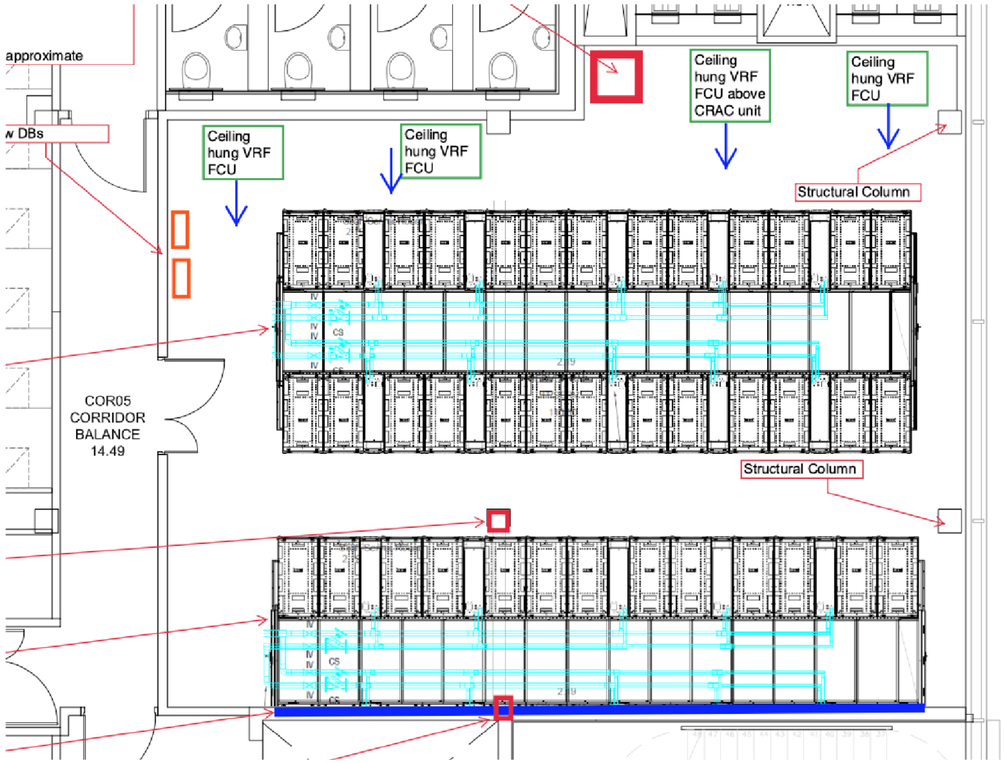

“Heating” System

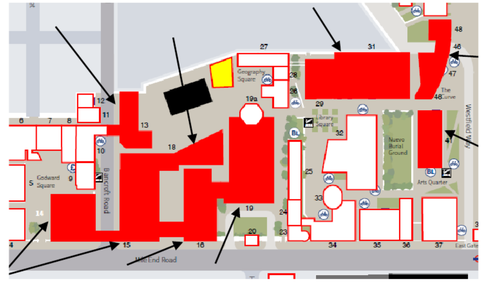

Now Fully Funded

District heating

Hot water store

75C

Cluster

26C

Heat pump

Heat exchanger

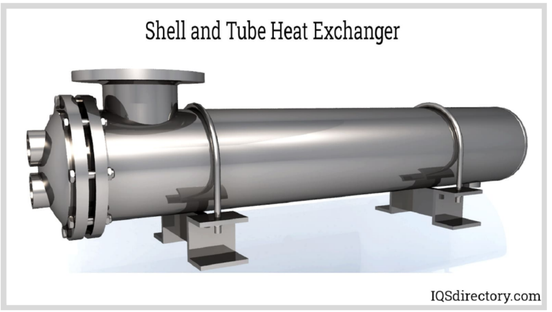

17C

17

75C

Cold water store

Hot water store

Dry air cooler

Monitoring system Deployment

LibreNMS Overview

SNMP-based auto-discovering network monitoring System

PHP web application

Uses MySQL/MariaDB

Includes a wide range of hardware

Routers, Switches, Access points, Security gateways

Features

Create alerts

Auto discovery

Multiple authentication methods (MySQL, HTTP)

Customisable dashboards

Distributed Polling

Device Backup integration (Oxidised)

Installation

https://docs.librenms.org/Installation/

Manually (RHEL/Centos), Installed with nginx instead of Apache

Add a device

• CLI
./addhost.php yourhostname [community] [v1|v2c] [port] [udp|udp6|tcp|tcp6]

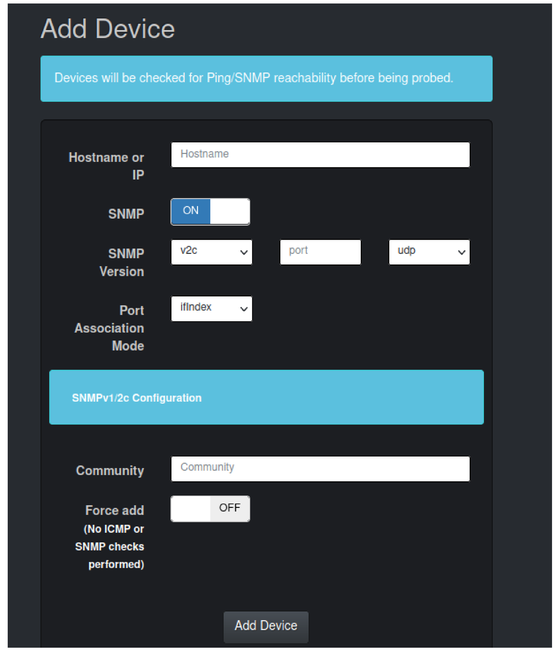

• WebUI
• Ping Only Device
• Auto-discovery
• API

Setup

• Several Devices added so far.
• Include servers, switches, PDUs, storage arrays etc.
• Divided in different device groups.

Hints Challenges

- Applications

Librenms Monitoring itself

Monitoring GPUs performance

SNMP extend:

Edit your snmpd.conf file : extend mysql /etc/snmp/mysql

Configuring nginx:

Editing the librenms.conf file.

```
server {
 listen      80;
  server_name “your server”;
 location /nginx-status {
     stub_status on;
     access_log   off;
     allow “server ip”;
     deny all;
 }
```

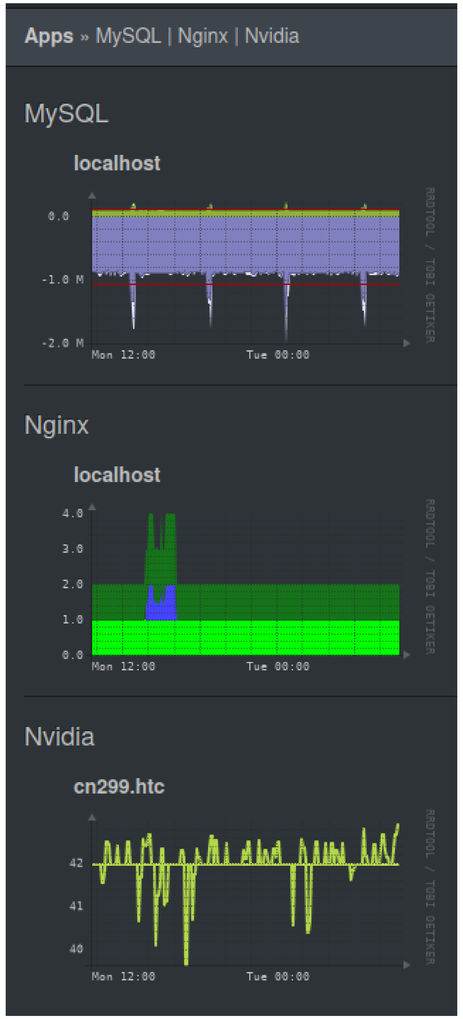

Future Use

- **Customisable alerting**
  Flexible alerting system, notify by email
- **Oxidized integration**
  Automatically extract, store, and compare configuration files from network devices.
- **SmokePING**
  Tool to keep track of network latency.
- **Network-Weathermap**
  Capability to build network maps to help visualize network traffic flow rates.