

ESRF MX COMPUTATIONAL AND NETWORK CAPABILITIES

Dr Andrew Peter Hammersley

ESRF

(Note: This is not my field, so detailed questions will have to be relayed to the relevant people. Alexander Popov of the ESRF MX Group is a remote participant.)

ESRF has 6 (or 7) MX End-stations (all can be used simultaneously)

Eiger detectors (1 for MX and 1 for Soft Matter) connected by 10 Gbit Ethernet to “Local Buffer Storage” and on to storage disks and a dedicated MX processing cluster

General Parallel File System (GPFS) is used for high throughput to / from disk

MX Processing Cluster:

Linux 64, Debian 7

mxproc BULLX 9 blades, 108 cores, X5675 3.06GHz 6-core 147Gflops, 2GB RAM per core

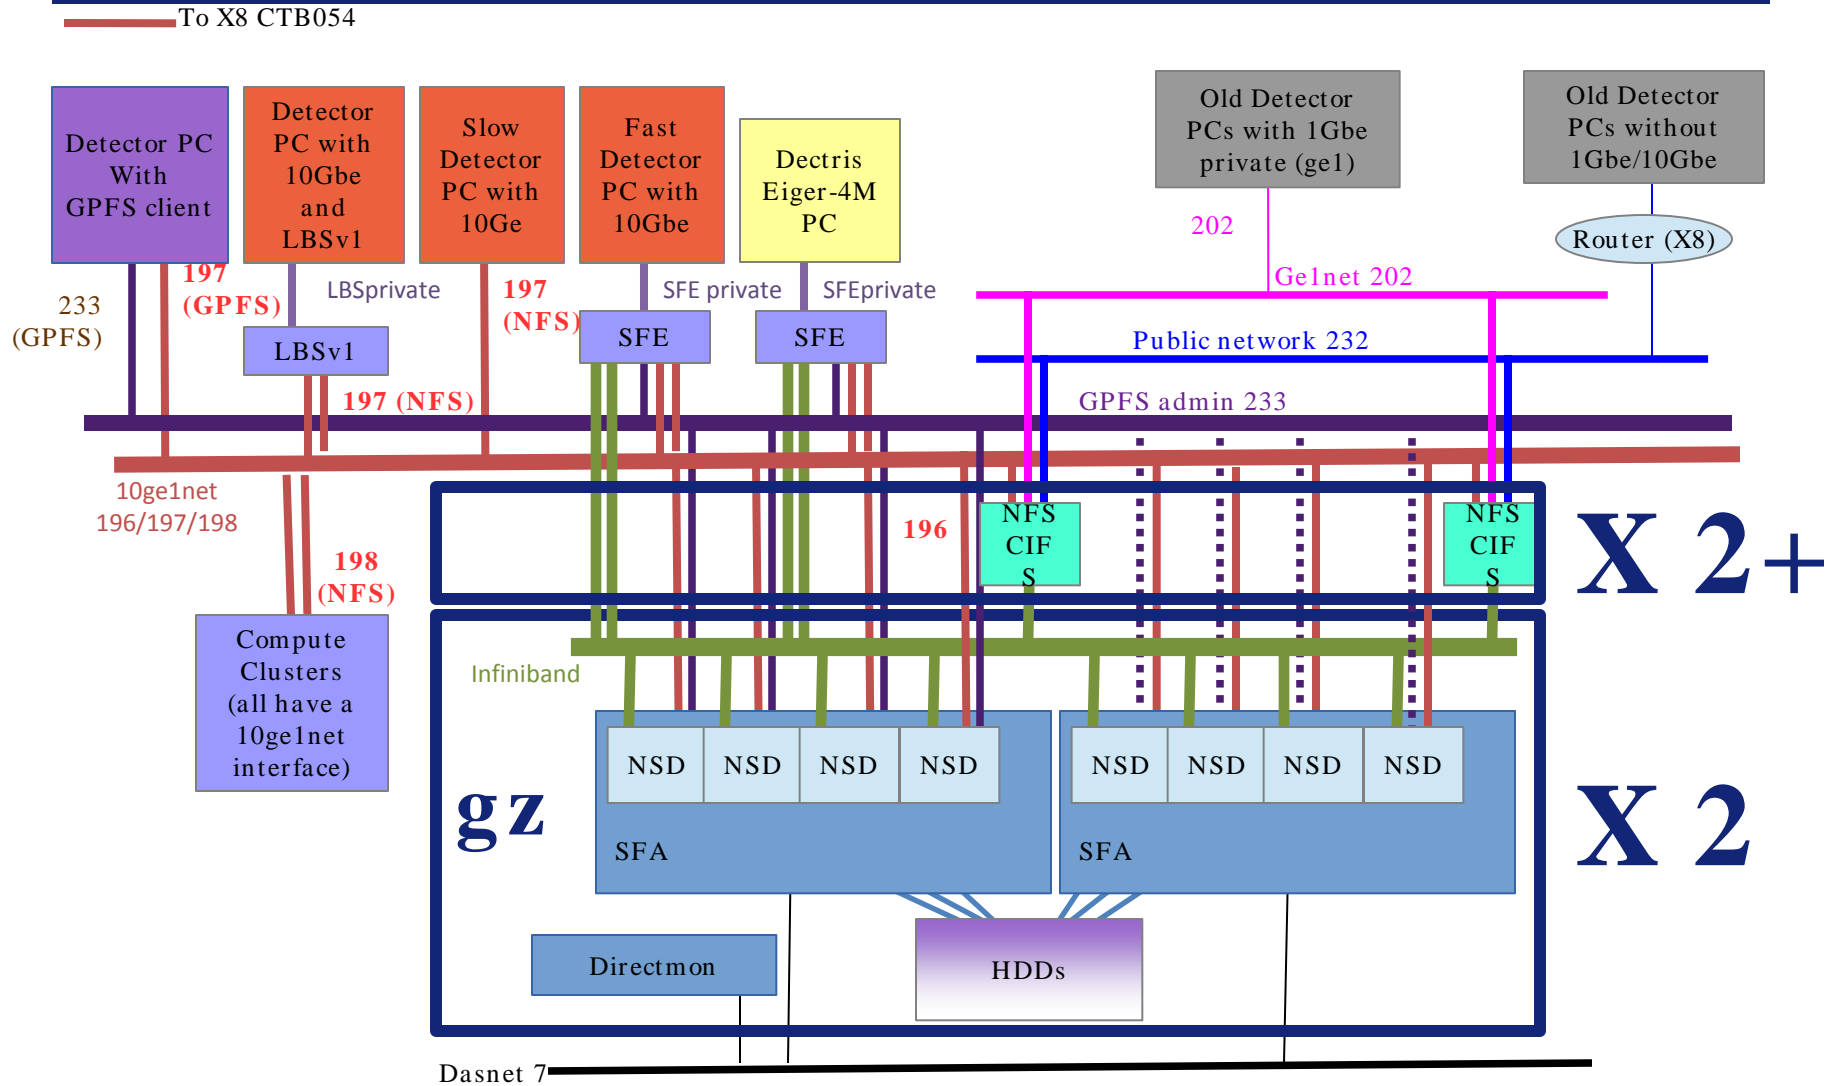
mxnice Dell C8220, 9 blades 180 cores, E5-2680v2 2.8GHz 10-core 227Gflops, 3.2GB 1600MHz RAM per core

ESRF PROCESSING CLUSTER



Courtesy of E. Eyer, ESRF

ESRF NETWORK SCHEMATIC



User Interface / Control Software: MXCube

Viewing and Quick-Look Quality Control: ADXVIEW, ALBULA for viewing; DOZOR (Alexander Popov) for B-factor. ALSO running on general MX processing cluster

Integration Software: Mainly XDS (MOSFLM used for screening crystals) ~80 – 90% of users use XDS at present

Typically 3 processing pipe-lines are used

Eiger File Conversion: The Eiger HDF5 files must be converted to be input into XDS or MOSFLM. Conversion programs are called as sub-processes to avoid creating extra files. NOTE: The two converter programs are **NOT interchangeable!**

Diagnostics: A test program is run which should take 1 minute at low load.

Current Processing status:

~80% of the time, the one minute program runs in roughly one minute

However in ~20% of the time it can take up to ~10 times longer

This is unacceptable, and can lead users to think that the programs have crashed.

It is thought that **CPU overload** is the main problem

Solution:

A hardware upgrade is in process

Add 8 new blades: 224 cores, E5-2680v4 2.4GHz 14-core, 2.3GB 2400MHz RAM per core

Processing capacity should be doubled

(DIALS should be added as another processing pipe-line)

The upgraded cluster should be available for MX processing in September 2016

The “Local Buffer Store” is no longer necessary as GPFS is performant enough direct to disk. This is due to phased out in ~1 year.

A “Local Preview Processor” may be added instead to provide dedicated processing power for preview / data quality analysis

Up to 750 images per second

Bursts of 1000 – 2000 images


Typical data collection: 180 degrees in 0.1 degree slices; 100 images per second with 18 seconds collection, ~3 Mbytes per image

~300 Mbytes per second

Exceeds 1Gbit Ethernet


Time between experiments: Min ~10 seconds / Typical 5 minutes

Current MX Eiger Set-up



Data Reduction Problem

100 Tbytes Experiment; Reduce (azimuthal integration) by factor 10



Investigating “Remote Direct Memory Access” for fast transparent data access. The data reduction processor will access the detector data automatically.

35 Gbits per second measured

Proof of method by End of June

Concept could be used for MX for previewing and quality control

Project for XFEL type data collection

Embed crystals in gel with “0.5 crystals per shot on average”

Squeeze gel continuously through beam

Collect data at 500 Hz (750 Hz works, but not continuously for long periods of time)

1.8 million images per hour. Roughly 6 Tbytes per hour

Each crystal randomly orientated

Data set 20,000 to 30,000 diffraction images, from ~100,000 images

Integration with CrystFEL and/or cctbx

Where will the processing power come from? (Initially one-off experiments)

(Big ESRF wide problem: Too many files for the file systems!)

THANKS FOR PROVIDING THE INFORMATION:

Olof Svensson: Data Analysis Unit

David Von Stetten: MX Group

Jens Meyer: Beam-line Control Unit

Bruno Lebayle: Systems and Communications

Andy Götz: Software Group