

Ellexus Ltd: The I/O Profiling Company

Dr Rosemary Francis, CEO, Good I/O evangelist

I/O Profiling with Ellexus at Diamond Light Source on an IBM Spectrum Scale cluster



The I/O Profiling Company - Protect. Balance. Optimise.

www.ellexus.com

Ellexus Ltd: The I/O Profiling Company

Products: We make system telemetry tools to help you

- improve application performance,
- protect shared storage, and
- manage application dependencies for migration.

Customers include:

TACC



TEXAS
The University of Texas at Austin



diamond



wellcome
sanger
institute


DELLEMC

ANSYS[®]

arm

Mentor[®]


A Siemens Business



Solving the noisy neighbour problem

How we worked with Arm to develop our technology

Example of a rogue job from Arm:



This application is overloading shared storage by putting data in the wrong place.

Temporary data is written to shared storage

Local storage is unused




Ellexus enterprise products

Take control of the way you access your data



Detailed I/O profiling
Application discovery



Dependencies

What do I need to include in my container?

How do I migrate this tool chain?

I/O profiling

What resources do I need to run it?

Debug and triage

Why am I not getting the results I expect?

Why can't I pick up the right licence?


Ellexus enterprise products

Take control of the way you access your data

Live telemetry for on-premises clusters and cloud

Protect storage and find bottlenecks

Cost management and forecasting



Live system telemetry:
I/O monitoring in production

Data that pays for itself

Returns on day one:

- Dependencies for migration
- I/O profiling and sizing
- Job steering
- Telemetry for operations
- Storage protection


Long-term benefits:

- Forecasting
- Optimisation
- Cost-aware scheduling



Tuning and sizing:

How much time are you wasting doing bad I/O?




- Small reads and writes
- Opens on unused files
- Stats on unused files
- Failed filesystem I/O
- Backward seeks
- File Trawls
- Zero byte I/O operations
- Failed network I/O
- Opens on seldom used files
- Stats on used files
- Forward seeks
- Large reads and writes
- Opens on well used files
- Successful network I/O




System telemetry:

How much time are you wasting doing bad I/O?

Traffic Light View - time spent in I/O



Traffic Light View - I/O counts




What is normal?

LSF - System overview ▾

📊 ☆ 🔄 📄 ⚙️ ⏪ 🔍 ⏩ 🕒 Last 15 minutes Refresh every 5s 🔄


Meta data I/O per mount point




File system reads and writes - Throughput per mount point



CPU time



Memory



Ellexus: The I/O Profiling Company
www.ellexus.com



Diamond Light Source

The UK's national synchrotron (particle accelerator)





123

124

125

CHB NLV 053

CHB NLV 054




Beam lines



Deploying Mistral at Diamond Light Source

Beamline data must be processed in real time



We chose an ELK stack data pipeline for Mistral telemetry



First win: eliminating bad meta-data

- Applications were profiled with Mistral
- Some were Stat()ing files many times
- The files did not exist or did not change



Second win: optimizing application deployment

Long pythonpaths cause file system trawls when loading modules.

Removing unused paths improved file system performance and application run time.



Production win: finding noisy neighbours

A single-core job was able to saturate the network of a group of nodes

Mistral located the node and jobID of the offending program



Fast, agile and (hybrid) cloud ready



Ellexus: The I/O Profiling Company
www.ellexus.com

~~Fast, agile and (hybrid) cloud ready~~

Performance, portability and planning



~~Fast, agile and (hybrid) cloud ready~~

Performance, portability and planning

Performance: Understand I/O patterns and requirements

Portability: Understanding application dependencies

Planning: System telemetry

⇒ Capacity planning and charge back

⇒ Make better decisions about how to deploy applications



Thanks for listening

Dr Rosemary Francis
CEO and Director of Technology
rosemary.francis@ellexus.com



Ellexus: The I/O profiling company
www.ellexus.com