

Testing Storage Systems: Methodology & Common Pitfalls

John Constable

Wellcome Trust Sanger Institute

@kript / jc18@sanger.ac.uk

#lisa14storage

What I wish I knew when I started testing storage systems three years ago

Covering;
intro, challenges, experiences, lessons learned, tips and tricks, denouement

The Sanger Institute

Funded by Wellcome Trust

2nd largest research charity in the world.
~700 employees (~50 IT)

Large scale genomic research.
Sequenced 1/3 of the human genome (largest
single contributor)





Who *are* all you people?
Why did you come to *this*
talk?

I'm part of the Informatics Support Group
(HPC), also Storage Team

3 years at Sanger, 22 years in IT


What I've learned from evaluating seven
different storage systems in 3 years

Charity (90% funded by Wellcome Trust)

Our users are scientists first, bioinformaticians second, unix geeks third

Cumulative PF Yield(GBases)

■ Cumulative PF Yield(GBases)



Why yes, the data is doubling every 12 months!

Yes, that's faster than 18 month disk/transistor doubling rate

N.B. a kilobase (kb) is a unit of measurement in molecular biology equal to 1000 base pairs of DNA or RNA.

One megabase roughly equiv 1 MB data.

Challenges

Currently 12 staff managing 21PB

Our scientists predict up to 70PB
within the next 5 years.

Its *all* spinning disk

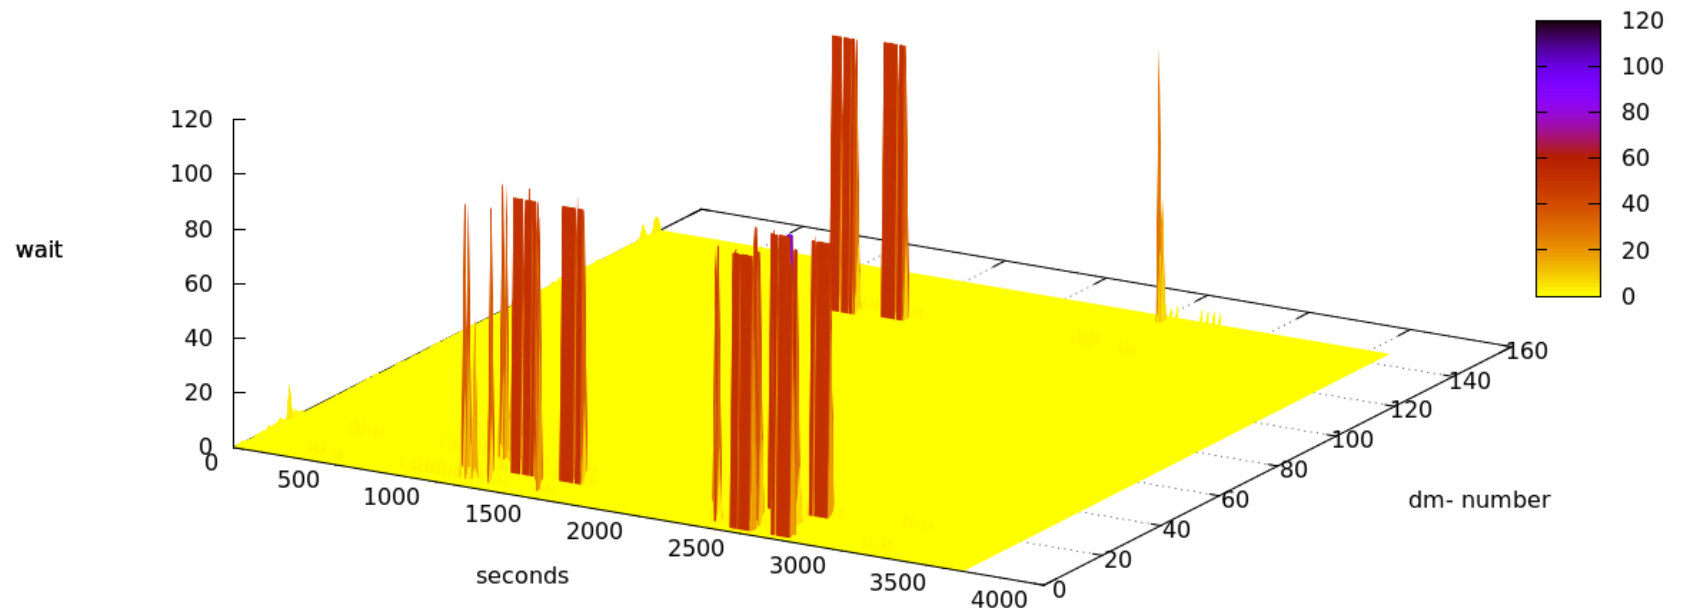
NFS/CIFS (7PB), Lustre (10PB),

iRODS (4PB)

'Only' 350TB tape



(gnuplot FTW) iostat info on average queue size



Experiences



Credit: Touring Club Suisse
https://www.flickr.com/photos/touring_club/




Can't make it go fast enough!

Our test or vendors equipment?

Our equipment or vendors' information?

What don't we know?



Scriptable tests are good, but be careful..

25k snapshots in one dir

It didn't break (mostly)!



No fit between use case
and hardware provided

Credit: W Smith:
<https://www.flickr.com/photos/wsmith>




Avoid beta hardware if possible
(Try not to be their test lab)



‘We never thought of that’
or
not able to monitor disk enclosures

Changes in Test Methodology Over Time



Test frameworks



Streaming IO
Random IO

Big Files!
Small Files!
Zero Length files!


Multithreaded
or multi streams.

Software and procedural



Sharing test plan with vendors





Negative testing based on
what went wrong last time
& in production to date


Monitoring of test systems and clients



Project management



Lessons Learned





Good relationship
with vendors crucial

Meet the engineers/developers
if possible

Credit: <https://www.flickr.com/photos/annstheclaf>



Learn new tech

You may fix your existing infrastructure along the way



Credit: Stéfan

<https://www.flickr.com/photos/st3f4n/>

Lessons Learned

Management is not a dirty word




Credit: Gary Burke
<https://www.flickr.com/photos/klignon65/5516037577>

Project management
(This can be just a regular conference call,
and an end date!)



Clear internal communication



Credit: Darren Hester
<https://www.flickr.com/photos/grungetextures/4215941524>

Tips and Tricks



Credit: Rinogmb off for a few day
<https://www.flickr.com/photos/rinogmb>

‘Begin with the end in mind’



Management



Monitoring



Resilience



Handover

Negative Testing





Cable Pulling!

For e.g.

- Network (inc. one pair)
- Storage interconnects
- (FC, IB, Ethernet)
- Chassis interconnects



Failover
(Pref. while under load)



Disk
failures

Useful

Software

Performance deep diving: <https://github.com/brendangregg/perf-tools/>

VD Bench for IO generation on Block and filesystems <http://www.oracle.com/technetwork/server-storage/vdbench-downloads-1901681.html>

SNMP browsing GUI:

<http://ireasoning.com/mibbrowser.shtml>

Books

The Practice of System and Network Administration by Limoncelli, Hogan, Chalup

Systems Performance by Brendan Gregg

Useful Links:

monitoring driven infrastructure:

<https://devblog.timgroup.com/2014/07/30/mdi-monitoring-driven-infrastructure/>

Jason Dixon's talk on The state of open source monitoring;

<https://speakerdeck.com/obfuscure/the-state-of-open-source-monitoring>

JISC Project Management (inc templates):

<http://www.jisc.ac.uk/fundingopportunities/projectmanagement/planning.aspx>

LiquidPlanner.com

https://blogs.oracle.com/brendan/entry/performance_testing_the_7000_series2

A photograph of a subway station with a 'Way out' sign and a list of five items. The sign is yellow with a black arrow pointing up and the text 'Way out'. The background shows a subway track with people walking and a curved tunnel. The list is overlaid on the image.

↑ Way out

1. What you will need to report on when the tickets come in
2. Plan your negative testing
3. Share with your vendor - involve them
4. Test script your frequent management tasks & questions
5. Graph all the things

Thank You!



Acknowledgement: Mr Woobey, Dr Coates, James Beal, Dr Clapham



John Constable
Wellcome Trust Sanger Institute
@kript / jc18@sanger.ac.uk
#keepcalmitsonlystorage

