

# DPM Dualstacking Experiences.

## A tale of two SEs.

Matt Doidge (Lancaster), Rob Currie (ECDF)  
10th April 2018

- Stretching this talk beyond a single slide was quite a task.
  - It's tempting to concentrate on the last bit of server-side config, which is a fairly straight-forward procedure.
- But the real meat of the matter is in the preparation, so we'll take a look at that.
- One of the goals is to build confidence though - once you're happy with your IPv6 infrastructure you should be able to dual-stack your SE without worries.
- My apologies if this is a little "for dummies". Write what you know...

# The Lancaster Story:

## Laying the Groundwork with Perfsonar.

---



- The first step was making sure that IPv6 packets could flow through (and from) our network, and that you can get v6 addresses assigned to machines in your networks (this was all done a while back now).
  - Luckily for us our Central IT guys are "v6-keen", and had set up a working DHCPv6 service.
- An invaluable tool was the Perfsonar box - this identified some problems with assymmetric IPv6 routing.
  - Don't assume IPv6 packets will go the same routes as their v4 counterparts!
- Dual-stacking your Perfsonar boxen is a good dry-run for other, job-affecting services.
- Which also allows you to get a feel for the state of the IPv6 infrastucture.

# The Lancaster Adventure:

## Testing the v6 Waters.

---



- Before we dual-stacked a vital site service and threw around more IPv6 packets than Lancaster had ever seen some confidence-building stress-testing was needed to be done.
  - The IPv6 Perfsonar mesh helped us gain confidence that v6 routing was working.
  - We Dual-stacked a handful of DPM disk-pools first which helped build confidence in our firewalls and other configs (and didn't show any sign of breaking anything).
  - To boost the confidence that the network could handle large flows of v6 traffic we set up a dualstacked standalone gridftp server.
  - Remote IPv6 tests were initiated from Ixplus7.
  - Protip from Duncan - direct transfers to /dev/null to stop disks writes slowing down your transfer tests.
- Every step we kept both atlas and the local networking team informed - so nobody had any surprises!

# The Lancaster Tale:

## Dual-Stacking Day.

---

- Configured and enabled v6 firewalls ahead of time.
- We requested v6 addresses on the morning of dual-stacking day.
- As soon as the changes were made we pounced on rolling out the dual-stacking changes, headnode first.
  - Of course we had prewarned our central IT guys that this change was coming.
  - And we are indeed very lucky to be able to get "same day service" for our DNS changes.
- Our strategy was not to rollback when we had problems - we rolled through them.

# The Lancaster Saga:



## It wasn't all coming up roses

---

- Of course things didn't go without a few hitches...
  - At first almost all transfers were failing.
  - It appeared that a lot of DPM intercommunication was preferring v6 (and not falling back to v4).
  - Solved by opening up the v6 firewall to the DPM service ports (DPNS etc), not just the transfer protocol ports.
- We weren't totally out of the woods - jobs were okay but FTS (gridftp) transfers were failing.
  - This was due to not having any contents in /etc/gai.conf - a feature specific to SL6.
  - See our gridpp storage blog post for details.



# v6 Packet Flow.

## Lancaster as a Source/Destination.



# v6 Efficiency.

## Lancaster as a Source/Destination.





# The ECDF Story.

---

- Dual-stacked their Perfsonar first.
- On the Day went for a big-bang upgrade in an attempt to avoid putting in too many requests.
  - Requests went in at the start of the week and the DNS changes went live on the Tuesday evening.
- Transfers to the SE from outside promptly broke.
- Tests showed that the DPM headnode was uncontactable over IPv6 - but then it hadn't been configured to yet.

# The ECDF Story continued.

- Rob, using the site's dual-stacked perfsnar boxes as templates, enabled v6 on the headnode and got it talking to the perfsnar boxes.
  - But external v6 communication to the headnode still failed, despite internal v6 communication working.
  - No good answer was obtained from network support for this - perhaps an incorrectly setup v6 gateway or something with their "experimental" switch config.
- Put in an urgent request to rollback the DNS changes as atlas transfers from dual-stacked sources weren't able to fall back to v4.
- Plan on only having the test storage dual stacked for the time being.

- Dual-stacking is "easy" (with the infrastructure in place), but it can break your SE!
- Transfer protocols (particularly SRM) don't fall back to v4, so interestingly this becomes more of an issue the more sites dual-stack!
- Lay as much groundwork as possible before putting in those v6 address requests
  - Ideally have your v6 requests go "live" on the day - although this might not be an option for everyone.
- Experience gained with a dual stacked Perfsonar box is useful.
- Having a dual stacked local box (perfsonar) and access to dual-stacked external box is vital for debugging. I used Ixplus7, and Rob had his own server.

Dual-stacking your SE needn't be scary, but it does have the potential to break transfers if misconfigured (or even if just left unconfigured).

Like any new technology, gaining experience with IPv6 in a non-production-wrecking capacity is the way to go. The Perfonar box is the best candidate for this.

Some Handy URLs:

[http://psmad.grid.iu.edu/maddash-webui/index.cgi?  
dashboard=UK%20Mesh%20Config](http://psmad.grid.iu.edu/maddash-webui/index.cgi?dashboard=UK%20Mesh%20Config)

[http://gridpp-storage.blogspot.co.uk/2018/01/  
dual-stacking-lancasters-sl6-dpm.html](http://gridpp-storage.blogspot.co.uk/2018/01/dual-stacking-lancasters-sl6-dpm.html)

[https://ggus.eu/?mode=ticket\\_info&ticket\\_id=131610](https://ggus.eu/?mode=ticket_info&ticket_id=131610)

<http://hepix-ipv6.web.cern.ch/>

Matt Doidge (Lancaster), Rob Currie (ECDF)  
10th April 2018

