


XCache Virtual Placement at Birmingham

GridPP Collaboration Meeting, 23rd March, 2022
Mark Slater, Birmingham University

Xcache is a general squid-like caching mechanism based on XRootD and is designed to cache data files from external sources at a location

For the last few years, Birmingham has been using it in place of a proper SE endpoint to aid diskless operation

Effectively, any data files requested from an external SE (Manchester in our case) are cached first and so will be present when requested again



`root://xcache.bham.ac.uk/root://dpm.manc.ac.uk/..`

The hardware that is/was dedicated to Xcache is:

- Basic server (16C, 24GB) to act as redirector
- Two 40TB + One 20TB pool nodes
- ~400 cores for Atlas → 50 simultaneous 8-core jobs
- Pool nodes are connected at 10Gb/s, workers at 1Gb/s.

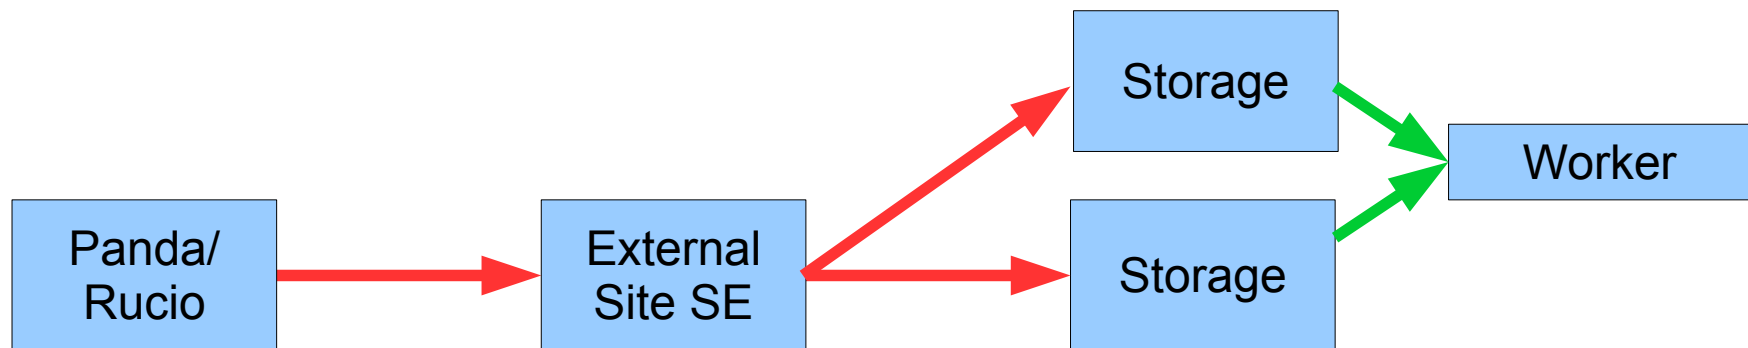
This setup has worked well with minimal intervention for a number of years now

However, the benefits appeared to be limited wrt saving offsite bandwidth

The ‘obvious’ next step for this is to change from being a reactive cache to being prefilled cache (or volatile storage)

Iliya from ATLAS has developed a mechanism to do this, based on Xcache and Xrootd

The setup for this is a little different. Each node runs it’s own Xcache instance and there is no redirector:



In the last couple of weeks I have been attempting to get this system up and running at Birmingham

There are 3 options for setting it up:

Using SLATE:

This means allowing DDM to run Docker containers on the cluster. Pros being you don't have to setup it up or manage it yourself, Cons that you have to allow an outside source root access to machines

Manually run docker images:

The images that are run by SLATE are all publicly available and so you can just run these in your own docker installation with some initial setup. This should give you ~same the configuration as above but with a minimum of setup effort.

Manual install and configuration:

Fundamentally, the VP Xcache setup is a set of package installs, general Xrootd config files and site-specific environment setup. This isn't hard to do by hand but you would need more ongoing maintenance to keep things up-to-date

The images can be found here:

<https://github.com/ivukotic/docker-xcache/tree/master/atlas-xcache>

In order to setup Xcache for VP I went through the following:

Node Setup:

- Blank CentOS 7 node
- Storage mounted as JBOD


Software Setup:

- Enable the OCG Repo
- Install atlas-xcache package
- Create some directories so things are where they should be
- Copy over general config files
- Add an xrootd user
- Create a startup script to launch the xrootd and monitoring daemons/scripts
- Add appropriate firewall holes


Note that this was cribbed from the Dockerfiles and so won't get updated when those do




Network Rate from Nodes



Straight after gets used by workers



%age Filled of each disk on a Single System



My experience with Xcache+VP has been very positive:

Installation went very smoothly ●

Integration into Panda/Rucio also seemed to go very well ●

I haven't had any Xcache related problems in the last week ●

Issues I had previously with Xcache (using my own cert + failing updates) are fixed ●

The only problems that we have encountered were actually (I believe) to do with Bham still using VAC

I haven't found any real downsides on the site end (other than VAC related issues) to this setup so wouldn't have any problems recommending it to others

Many Thanks to Iliya Vukotic, Rod Walker and FaHui Lin of Atlas to helping set all this up!