The Open High Throughput Computing Content Delivery Network

Dave Dykstra

CHEP 2018

10 July 2018

Background - WLCG Content Delivery Network

- The WLCG Content Delivery Network is based on general purpose squids, primarily for CVMFS and the Frontier Distributed Database Caching system
- Last CHEP I presented a new way to find squids: Web Proxy Auto Discovery
  - At http://wlcg-wpad.cern.ch/wpad.dat and http://wlcg-wpad.fnal.gov/wpad.dat
  - Since then the service has been updated:
    - Based on CMS & ATLAS squid registrations, cross-checked against registrations in GOCDB & OIM, plus a few additional squids manually added
    - Different answers can be given for different address ranges in same GeoIP Organizations -- used to distinguish CERN Meyrin & Wigner, for example
  - Now used in production by U.S. CMS opportunistic computing
  - WLCG standard extended to look first for http://grid-wpad/wpad.dat to provide local override & offload
    - implemented at CERN, including IPv6 support
- Having squids on-site is important for low-latency and performance with many queries and clients, but what about very small sites or extremely distributed resources such as LHC@Home?

Cloudflare CDN

- Cloudflare is a large commercial Content Delivery Network vendor
  - Has easy-to-use web interface
  - Works with any domain when Cloudflare hosts the DNS
  - Caches Http content from any hosted DNS alias in their huge network of distributed caching servers
    - Administrator can choose whether aliases are cached or not cached
  - Works with CVMFS & Frontier
  - Has a free tier that allows unlimited bandwidth
  - Includes DDoS protections, IPv6 support, DNSSEC, and more

AnyCast

- Cloudflare uses a small number of IP addresses automatically routed to their nearest data center with AnyCast
  - No need to abuse DNS caching (as I suggested at the last CHEP) with different responses based on GeoIP
  - Requires a lot of effort on their part to set up, to arrange with many ISPs
    - Not something that could be used to direct traffic to squids at many grid sites; needs one entity controlling the distributed network
- Addresses are shared, but queries are directed to the correct origin server based on the 'Host' header from clients

openhtc.io

- Domain hosted in Cloudflare free tier
- For use when no local squids are available
- “Page Rules” set to Cache Everything, Respect server expiration headers
- Domain is set up for long term sustainability, with multiple people having access
- Initial aliases are made for CVMFS stratum 1s and CMS frontier servers. For example:
  - s1fnal-cvmfs.openhtc.io, s1cern-cvmfs.openhtc.io
  - cms-frontier.openhtc.io
- Extendable to other High Throughput Computing applications if they don’t want to use their own domain

Use cases

- In production use for LHC@Home
- Configured for any Open Science Grid cvmfs installation that sets CVMFS_HTTP_PROXY=DIRECT
- Configured for Frontier for U.S. CMS opportunistic use, when no WLCG squid available to WPAD
  - Hasn't actually run at such sites, but was successfully used to serve a large grid site for a weekend when the local squid was failing
- Planned for LHC OpenData
- Planned for CernVM cvmfs default when nothing found in WLCG WPAD

Performance

- Measurements of CMS@Home have shown an average start time improvement of 295 seconds, almost 5 minutes
  - Average start speedup: 2.2 times [measurements thanks to Laurence Field]

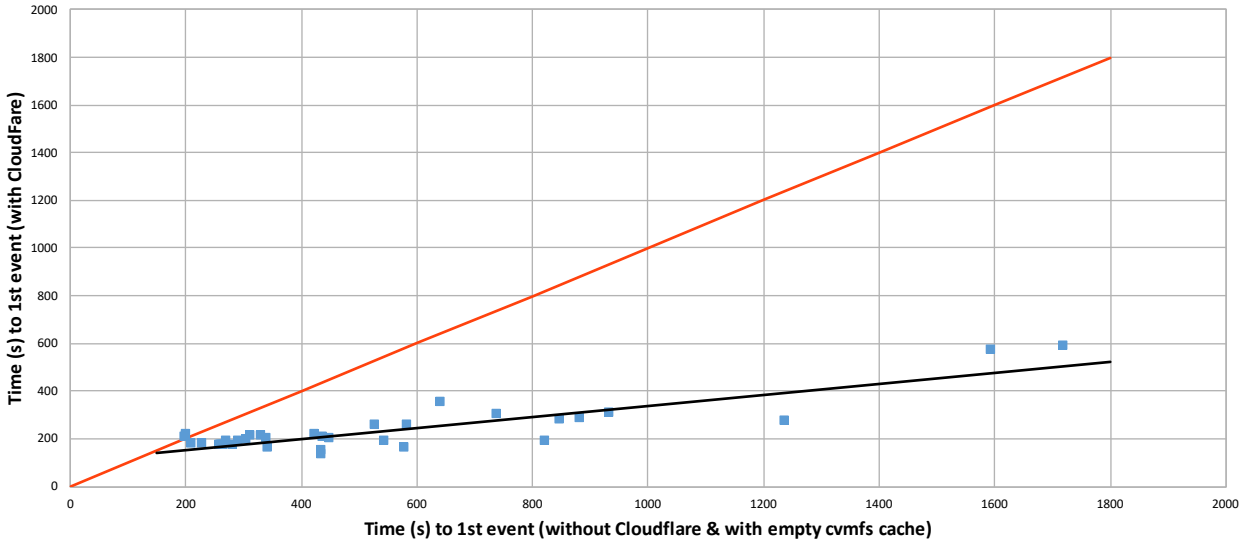

Disadvantages

- Much less detailed monitoring
- Uses more WAN bandwidth, has higher latency than on-site squids
- See more hits on origin servers than using squid alone, but much fewer than the number of clients
- Only caches ports 80 & 8080 (and https on 443)
- Need to disable If-Modified-Since on Frontier, because it is not possible to purge errors
- Cloudflare Terms of Service says they may disable this type of service if they detect it causing problems with their servers
  - I asked, and they said that this is unlikely to happen under the expected usage patterns I described

We can live with these disadvantages

Info

- http://openhtc.io
- How does Cloudflare work?
- How does Anycast work?