

# Evolution of the HEP Content Distribution Network

Dave Dykstra

CernVM Workshop

6 June 2016

# Current HEP Content Delivery Network

- The HEP CDN is general purpose squid proxies, at least at all WLCG sites that support ATLAS or CMS
  - Driving factor: Frontier Distributed Database
  - CVMFS spread very quickly because of the infrastructure
  - Many other smaller use cases take advantage of them as well, and more would if they were easier to find
  - Centralized monitoring at [wlcg-squid-monitor.cern.ch](http://wlcg-squid-monitor.cern.ch)
- It's a pretty big problem that ATLAS & CMS separately maintain configurations for proxies
  - In very different ways
  - Makes it tough to run opportunistically at each other's sites, and tough for other VOs to take advantage of them

# Comparison to non-HEP CDNs

- CDNs on the internet do not depend on smart clients that know how to use forward proxies
  - Servers are made smarter, and the distributed caching servers (at least some are reverse proxies) all need to be configured to know about the ultimate source of data
  - Often DNS caching is abused to point to different servers in different areas
- It works when all the servers in a CDN are controlled by one entity, but that's not the grid model
  - The hit rates of the HEP applications are so high that it usually makes sense to have caches on the premises, especially when they're as easy to maintain as squid
  - There's probably a way to distribute configurations that would work with decentralized control, but there's an easier way

# Web Proxy Auto Discovery

- There's a de-facto internet standard for finding web proxies
  - Supported by all major web browsers
  - Clients try <http://wpad/wpad.dat> to read a Proxy Auto Config format file, a javascript subset, for example:

```
function FindProxyForURL(url, host) {  
    return "PROXY http://squid.aglt2.org:3128"  
}
```
  - Can select different values based on destination url or source ip address
  - Open source pacparser library available to interpret, supported by both Frontier and CVMFS
  - Also supported by an open source wget wrapper that I wrote, pacwget

# WLCG WPAD

- For WLCG, if <http://wpad/wpad.dat> not found, use <http://wlcg-wpad.cern.ch/wpad.dat>
  - Initial version based on squids registered in GOCDDB & OIM was made available last Friday
    - Gives different answers based on source IP matching GeoIP organization
    - For testing, add “?ip=N.N.N.N” to URL to change IP
  - Running on a pair of large 10Gbit/s physical servers that also support 4 external squid proxy services
  - Large sites will be encouraged to run their own wpad web service to reduce latency and offload CERN servers
    - Especially if they have multiple squid services

# Customizing the response

- Soon will be able to give custom responses based on a configuration file on the servers
  - Set or override proxies for any GeoIP organization
  - Option to return different answers based on destination URLs
    - Shortcuts in configuration file – ATLAS/CMS, CVMFS/FRONTIER
  - No current plans to separate source address ranges

# LHC@Home WPAD

- Also <http://lhchomeproxy.cern.ch/wpad.dat>
  - For use by LHC@Home (BOINC) clients
  - An alias for the same server, but it behaves differently for the different Host header
    - Returns a list of externally-accessible proxies that are sorted by GeoIP longitude/latitude relative to source IP
    - Currently the only proxy configured is on the same machine, but lhchomeproxy.fnal.gov is coming very soon and other sites are also being recruited around the world

# External proxies at CERN & FNAL

- CERN has had a pair of 10gbit/s machines for several months and FNAL has almost identical machines available by next week
  - Owned by CMS both places
  - Reduced performance because of high latency
  - Four proxy services on different TCP ports:
    - CMS Frontier backup, for failing site squids
      - Monitored for failover
    - CMS Opportunistic, for sites without squids
    - CVMFS backup, for failing site squids
      - Plan to configure clients when both sites up, and monitor for failover
    - LHC@Home proxies
  - No one service can use all the bandwidth
  - Each service independently monitored

# IPv6 support

- frontier-squid-3 supports IPv6
  - Has been available about a year
    - Since squid-3.5 when Collapsed Forwarding was finally implemented
      - squid-3 was complete re-write in C++
    - Not default because of one remaining serious bug affecting Frontier, not CVMFS
      - Collapsed forwarding broken with If-Modified-Since
    - Contractor hired to fix bug, estimating 11 July delivery
  - Installed on the External proxy machines
- wlcg-wpad service still needs some work for IPv6

# Other future work

- At some point there will probably be reason to extend the WPAD service to dynamically started proxies in clouds or as grid jobs
  - Perhaps integrate with Shoal
- There may be a need to add more distributed servers