Web Proxy Auto Discovery for WLCG

Dave Dykstra

CHEP 2016

10 October 2016

Background - Current WLCG Content Delivery Network

- The WLCG CDN is general purpose web caching proxies
  - Driving factor: Frontier Distributed Database for ATLAS & CMS
  - CVMFS spread very quickly in large part because of the existing squid web cache infrastructure deployed for Frontier
  - Many other smaller use cases take advantage of the squids as well, and more would if the squids were easier to find
  - Centralized monitoring is at wlcg-squid-monitor.cern.ch
- It's a pretty big problem that ATLAS & CMS separately maintain Frontier client configurations for squid proxies
  - In very different ways
  - Makes it tough to run opportunistically at each other's sites, and tough for other VOs and applications to take advantage of the caches

Comparison to non-WLCG CDNs

- CDNs on the internet do not depend on smart clients that know how to use forward http proxies
  - Servers are made smarter, and the distributed caching servers (which are often reverse http 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 WLCG 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

- We extend standard for WLCG: if http://wpad/wpad.dat not found, use http://wlcg-wpad.cern.ch/wpad.dat
  - wlcg-wpad.cern.ch WPAD service now available, using site squids registered in ATLAS AGIS and CMS SITECONF for Frontier
    - Cross-checked against squids registered in GOCDB and OIM (cannot simply use all of those because they may not be adequately sized for an application like Frontier)
    - Gives different responses based on source IP of requester matching a squid at the site that is in the same address range according to the Maxmind GeoIP Organization database
      - GeoIP Org DB maps IP address ranges to organization names
    - First user: CMS opportunistic usage of non-CMS sites in the U.S. using one virtual T3 site
  - Service is running on a pair of large 10Gbit/s physical servers (which also support 4 squid proxy services to the internet)
- Large sites will be encouraged to run their own http://wpad/wpad.dat web service to reduce latency and offload CERN servers
  - Especially if they have multiple squid services at their site

Information flow

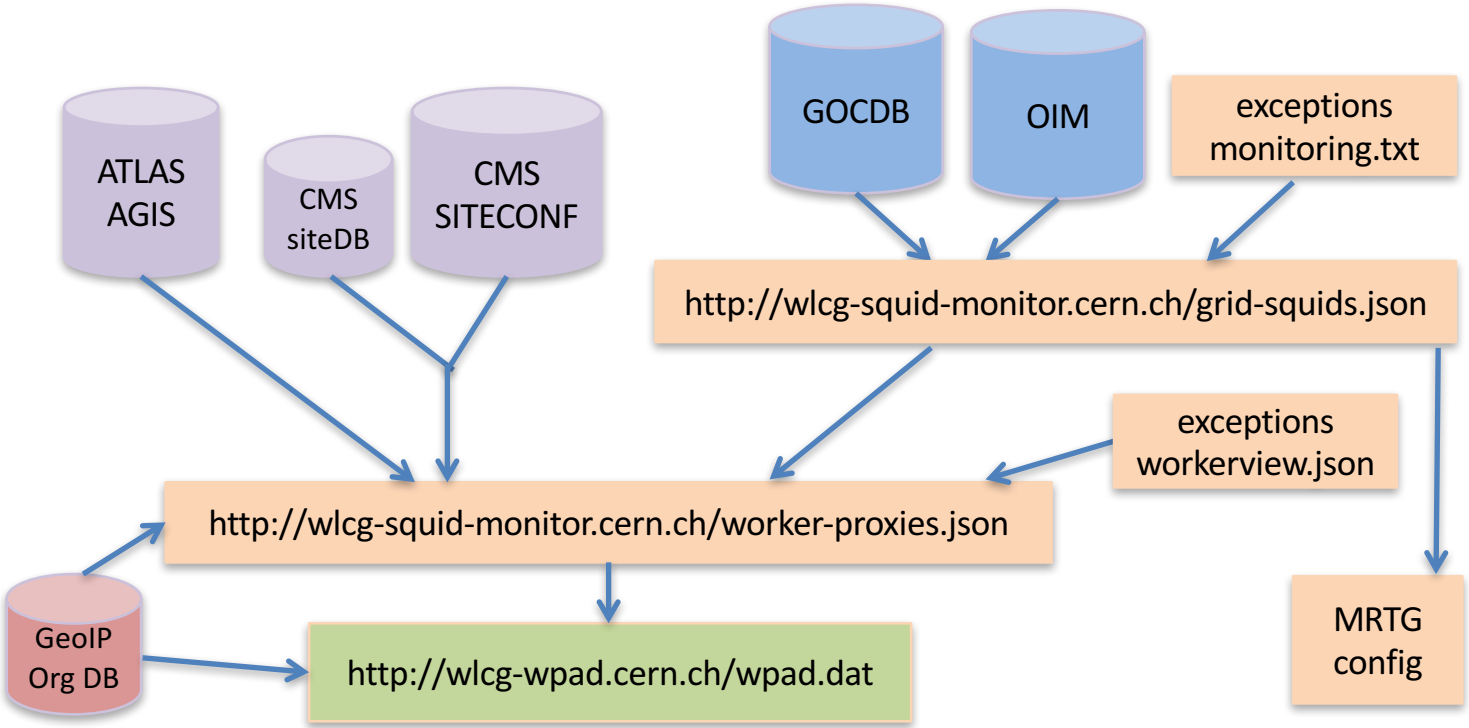

Special cases

- Proxies from ATLAS & CMS don't all map to a GeoIP Organization, because some are on a private network
  - Public name of squid from GOCDB/OIM for same site, if it is registered, is used instead to identify the GeoIP Org
- Can set or replace proxies for any GeoIP Organization in an exceptions json file on wlcg-squid-monitor
  - Fixing information sources is preferable of course
- There are cases of multiple site names mapping to the same GeoIP Organization with different proxies
  - Will need to distinguish by source IP address ranges, preferably coming from information sources, not exceptions file (only the latter implemented so far)
- Sites can have different squid services for different purposes
  - Selected by shortcuts (e.g. atlas, cms, cvmfs, frontier; defined by config) in exceptions file
- Client-based load-balancing has to be approximated

Separate services + client load-balancing

```
// For USCMS-FNAL-WC1; CMS: T1_US_FNAL
functionFindProxyForURL(url, host) {
    if (shExpMatch(url, "*cmsfrontier*.cern.ch*") || shExpMatch(url, "*/cvmfs
/cms*.cern.ch*")) {
        ran = Math.random();
        if (ran < 0.25) return "PROXY http://cmsfrontier2.fnal.gov:3128; PROX
Y http://cmsfrontier3.fnal.gov:3128; PROXY http://cmsfrontier4.fnal.gov:3128;
PROXY http://cmsfrontier1.fnal.gov:3128";
        if (ran < 0.5) return "PROXY http://cmsfrontier3.fnal.gov:3128; PROXY
http://cmsfrontier4.fnal.gov:3128; PROXY http://cmsfrontier1.fnal.gov:3128; P
ROXY http://cmsfrontier2.fnal.gov:3128";
        if (ran < 0.75) return "PROXY http://cmsfrontier4.fnal.gov:3128; PROX
Y http://cmsfrontier1.fnal.gov:3128; PROXY http://cmsfrontier2.fnal.gov:3128;
PROXY http://cmsfrontier3.fnal.gov:3128";
        return "PROXY http://cmsfrontier1.fnal.gov:3128; PROXY http://cmsfron
tier2.fnal.gov:3128; PROXY http://cmsfrontier3.fnal.gov:3128; PROXY http://cm
sfrontier4.fnal.gov:3128";
    }
    return "PROXY http://squid.fnal.gov:3128";
}
```

LHC@Home WPAD

- Also http://lhchomeproxy.cern.ch/wpad.dat
  - For use by LHC@Home (BOINC) clients
  - A DNS alias for the same server, but it behaves differently for the different requested server (in the Host http header)
    - Returns a list of externally-accessible proxies that are sorted by GeoIP longitude/latitude relative to source IP
    - Currently only includes the proxies configured on the same pair of machines and on lhchomeproxy.fnal.gov
    - Other sites are being recruited to run additional squids for this service around the world

Future work

- Add retrieving source IP address range from information sources
- Add IPv6 support
- There may be a need to add more distributed servers
  - Would be quite easy to add it to existing similar machines at Fermilab
- At some point there will probably be good reasons to extend the WPAD service to dynamically started proxies in clouds or as grid jobs
  - Probably integrate with Shoal

Conclusion

- There's now a standard way for any client that needs web caches to find them wherever they run on the WLCG
- Links:
  - https://twiki.cern.ch/twiki/bin/view/LCG/HttpProxyDiscoveryTaskForce
  - http://wlcg-wpad.cern.ch/wpad.dat
    - Optionally add ?ip= to change IPv4 address looked up
  - http://wlcg-squid-monitor.cern.ch/worker-proxies.json
  - https://github.com/pacwget/pacwget