



The Value of Content Distribution Networks
Mike Axelrod, Google
axelrod@google.com

Introduction

- Well understood facts:
 - Fast is better than slow but it costs more to be fast
 - Network has to be fast and the content has to be local and
 - Demand for BW continues to grow faster than revenue
- As we know, bandwidth demand surpassing supply
 - more users...
 - at higher speeds...
 - downloading ever richer content...
 - onto a greater variety of devices
- Inefficient to send a same copy of media multiple times over Trans-Pacific/Atlantic/Sat links
- Broadband access drives rich content - which in turn drives demand on the backbone – never ending cycle



Introduction

- YouTube video, consumes lots of traffic at Google peering points.
- Solutions are needed to push content to the edge of the network to improve performance
- Google peering policy .
 - present at most of IXs in Europe, Asia and US.
 - desire to supplement peering with a more efficient way to distribute content.
- Serving content from the edge of the network closer to the user improves performance, user happiness and reduces service provider's demands for the long haul BW.



Google Global Cache beta

- Today each request for Google content, such as YouTube video or map-tiles results in a unique copy send to the user.
- Users are directed to the nearest Google data centre, but there is room for improvement and efficiency gains for partner ISPs
- Hosting all content everywhere is inefficient, it makes sense to cache popular content and serve it locally.
- Google Global Cache (GGC) allows a large portion of these requests to be served from a small node located inside the service provider's network serving regional users
- Significant performance improvements for proxied connections over high latency links. (20%)

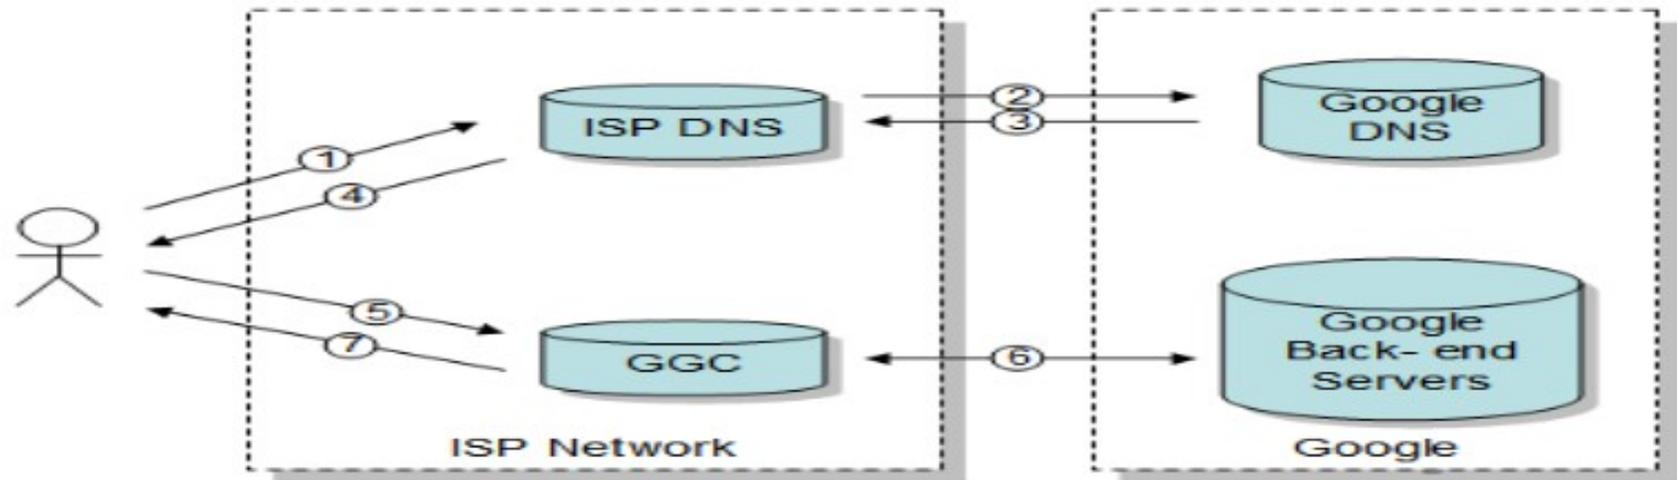


Google Global Cache beta - How it works

- Based on user location and the type of content requested Google DNS will direct user to the best Google node
- If given node already has the requested content in its local cache, it will serve the content directly to the end user.
- When content is served from the node end user experiences improved performance and the service provider does not incur the expense of carrying the request and service traffic across their peering and/or transit links.
- If the content is not already stored in the node, it will retrieve it from Google and store it for future requests.



Google Global Cache beta – How it works



- User resolves *content_host.google.com*. (example only)
- If the ISP DNS doesn't already know the IP address, it queries Google DNS for the IP address of *content_host.google.com*.
- Google DNS server knows that this ISP has a Google Global Cache node that should be able to service the request, so it replies with that IP address.
- IP address is returned to the user.
- The user requests the desired content from the IP address, which happens to be a GGC node on their ISP's network instead of a server on the Google network.
- If the content is cacheable and not already on the GGC node, it requests the content from Google on behalf of the user and caches it for future requests.
- Once the GGC node has the content, it can serve it to the requesting user as well as subsequent users looking for the same content.

Google Global Cache beta – How it works

- Key is to select Google server in close proximity to the requesting client
- Most CDN systems today use the DNS to make such server selection decisions. However, DNS provides only the IP address of the client's local DNS server to the CDN, and not the user's IP address.
- Implicit assumption that clients are close to their local DNS servers could lead to suboptimal node selections
- The only way to improve the situation is to ensure that users point to right resolvers
- IP-2-GEO techniques used to make sure that DNS IP and user IP map to proper geography where node is located



Google Global Cache beta - How it works

- GGC gives control to ISPs by collecting IP prefixes using BGP from ISPs hosting GGC - to determine which users should be served by a given node.
- Proxy mode: GGC pre-establishes connections to Google back ends, terminates user TCP session locally and is capable of caching static content.
- Significant speed improvements over the high latency links due to the elimination of at least 3 rrt (syn/syn-ack/ack) and reducing server time-outs.



Google Global Cache beta – deployment

- GGC runs on rack mountable servers 4-6 in each cluster
- Each server provides back-up to others in a cluster
- Deployment in ISP networks that have substantial traffic to Google and are interested in saving BW
- Needs
 - rack space 8-12 RU
 - power ~ at 1800-2700W
 - IPs and Ethernet ports



Google Global Cache beta – Summary

- Compliments open peering model
- Improves performance for ISP customers
- Saves BW and money for ISPs that serve YouTube
- Scalable long term solution for edge content distribution
- Caching and localization are vital for optimization of network resources



Thank You.



Mike Axelrod
Google, Inc.
axelrod@google.com