Management Concerns with Large Websites

Google is the top one accessed site based on the information collected at web100.com. It's been years since Google morphed into an Internet icon, transforming what was a search engine into the world's information storehouse, answer-finder, and oracle-like source of knowledge ("Web 100", 2009). Google has become an English word used in sentence like "I'm googled" or phrases like "global googlization" or "google universe." The problems Google web administrators are facing are normally the most complicated, and the solutions the Google corporate are deploying are normally the guidance to all practitioners in the profession of web server management.

Part I Technical Management Issues

According to ISO Telecommunications Management Network model and framework for network management, the technical concerns are mainly divided into five categories: faults management, configuration management, accounting management, performance management and security ("FCAPS", 2019).

Faults Management

There are normally two ways to manage faults, hardware and software. At hardware level, fault-tolerant hardware features such as redundant power supplies, redundant array of inexpensive disks (RAID), and high-quality components can be used. However, Google's fault management is mainly provided in software rather than in server-class hardware, which is aligned with their strategy of using commodity PCs to build a high-end computing cluster. Google's web team replicate services across many different machines and automatically detecting and handling failures. Most accesses to the index and other data structures involved in answering a query are read-only. Updates are relatively infrequent and are performed safely by diverting queries away from a service replica during an update, so many inconsistent issues in general-purpose database can be avoided. Accordingly, Google also deploys many small multiprocessors to support the replicating service, so the effect of faults can be contained to smaller pieces of the system. In comparison, large-scale shared-memory machines do not handle individual hardware component or software failures gracefully, with most fault types causing a full system crash. Finally, the Google computing clusters are geographically distributed, while each cluster having around a few thousand machines. This protects Google's web site against catastrophic data center failures, like those arising from earthquakes and large-scale power failures ("WEB SEARCH FOR A PLANET: THE GOOGLE CLUSTER ARCHITECTURE", 2003).

Configuration Management

Configuration management is the practice of handling changes systematically so that a system maintains its integrity over time ("Configuration management", *Wikipedia*). For large web sites like Google, thousands of mid-range PCs are being operated instead of a few high-end multiprocessor servers, which normally incurs significant system configuration, administration and repair costs. However, Google's applications are relatively homogenous, so they can use tools to install and upgrade software on groups of machines. The system configuration is codified in script languages like YAML or Ruby (Heidi, 2019) with tools like Ansible, Chef and Puppet ("Configuration management", *Google Cloud*). The time and cost to maintain 1,000 servers isn't much more than the cost of maintaining 100 servers because all machines have identical configurations. Similarly, the cost of monitoring a cluster using a scalable application-monitoring system does not increase greatly with cluster size. Furthermore, Google keeps repair costs reasonably low by batching repairs and ensuring that the components with the highest failure rates can be easily swap out, such as disks and power supplies ("WEB SEARCH FOR A PLANET: THE GOOGLE CLUSTER ARCHITECTURE", 2003).

Accounting Management

Accounting management is concerned with monitoring, logging and controlling the resource utilization such as disk space, memory, CPU usage, bandwidth etc. Google's computing clusters share machines between applications to increase the utilization of hardware, so interference can occur in any processor resource that is shared between threads of different jobs, such as processor caches and memory access-paths. This interference can negatively affect the performance of latency-sensitive application, but Linux has few defenses against it (Zhang, 2013). To resolve this problem, Google has developed a series of mega-infrastructure-management system from Borg, Omega to CMI[2] (Metz, 2013). The solution of CPI2, uses cycles-per-instruction (CPI) data obtained by hardware performance counters to identify problems, select the likely perpetrators, and then optionally throttle them so that the victims can return to their expected behavior. It lets Google engineers isolate poor performance down to a single task running on a single processor within a cluster of thousands, then drill down to it and select to throttle that task, without causing a CPU overhead of more than 0.1 per cent. It requires no special hardware and its only software dependency appears to be use of Linux. After CPI data is logged and stored offline, along with profiles of antagonist tasks, admins can query it via Google's major internal analysis tool, Dremel, which is used for performance forensics to let Google engineers identify particularly aggressive antagonists. Besides CPI[2], there's a parallel technique called "Google-Wide Profiling" that tracks both hardware and software performance but is only in use on a fraction of Google machines due to concerns about performance (Zhang, 2013).

Performance Management

Performance management may be the most important part in the five categories, and its success depends on the success of all of the other categories. A single query on Google reads hundreds of megabytes of data and consumes tens of billions of CPU cycles. Supporting a peak request stream of thousands of queries per second requires an infrastructure comparable in size to that of the largest supercomputer. Therefore, Google chooses to run queries parallelly on more than 15,000 commodity-class PCs rather than a smaller number of high-end servers. In another word, Google looks for the highest performance/price, not just performance.

When a user enters a query to Google, A DNS-based load-balancing system selects a cluster by accounting for the user's geographic proximity to each physical cluster. The load-balancing system minimizes round-trip time for the user's request, while also considering the available capacity at the various clusters. Thereafter, A hardware-based load balancer in each cluster monitors the available set of Google Web servers (GWSs) and performs local load balancing of requests across a set of them. In the indexing phase, Google divides the index with many terabytes of data into pieces called index shards, each having a randomly chosen subset of documents from the full index. Each shard is served by a pool of machines, and each request in the shard chooses a machine within a pool using an intermediate load balancer. In the document-serving phase, the similar strategy of "distributing to shards backed by multiple server replicas through a load balancer" is used. By parallelizing the search over many machines, Google reduces the average latency necessary to answer a query, dividing the total computation across more CPUs and disks. Because individual shards don't need to communicate with each other, the resulting speedup is nearly linear.

At hardware side, Google exploits both simultaneous multithreading (SMT) and chip multiprocessor (CMP) architectures for thread-level parallelism. The thread-level parallelism allows near-linear speedup with the number of cores, and a shared L2 cache of reasonable size speeds up interprocessor communication. As Google has partitioned index data and computation to minimize communication, the computation-to-communication ratio is high and expensive large shared-memory can be avoided ("WEB SEARCH FOR A PLANET: THE GOOGLE CLUSTER ARCHITECTURE", 2003).

Security Management

The security of the infrastructure is designed in five progressive layers starting from the hardware infrastructure, to service deployment, storage services, internet communication to operation security ("Google Infrastructure Security Design Overview").

At hardware layer, Google builds its own data centers which incorporate multiple layers of physical security protections, including biometric identification, metal detection, cameras, vehicle barriers, and laser-based intrusion detection systems. Google also custom-designed the server boards, the networking equipment and the security chips to authenticate Google devices. Moreover, Google uses cryptographic signatures to validate low-level components like the BIOS, bootloader, kernel, and base operating system image.

At service deployment layer, Google uses cryptographic authentication and authorization at the application layer for inter-service communication rather than network segmentation or firewalling. Besides, a variety of isolation and sandboxing techniques are used to protect a service from other services running on the same machine.

At data storage layer, Google configures the storage services to use keys from the central key management service. Before a decommissioned encrypted storage device can physically leave, it is cleaned using a multi-step process that includes two independent verifications. Devices that do not pass this wiping procedure are physically destroyed.

At internet communication layer, Google uses an infrastructure service called Google Front End (GFE). The GFE ensures that all TLS connections are terminated using correct certificates and following best practices such as supporting perfect forward secrecy. It additionally applies protections against Denial of Service attacks.

At operation security layer, Google applies central source control and two-party review features. Beyond that, Google introduces libraries that prevent developers from introducing certain classes of security bugs, such as XSS vulnerabilities. Automated tools like fuzzers, static analysis tools, and web security scanners.are used to detect security bugs.

Part II Non-Technical Management Issues

Google also face non-technical management issues in its 20-year history. Those issues can be also put into five categories: social, political, legal, privacy and Ethical.

Social Issues

Although Google is criticized for its tax avoidance strategy, the issue itself is not directly associated with its web sites. A more related example is the criticism for the high amount of energy used to maintain its servers, which was estimated as a million (Strand, 2008). In 2011, the Dalles plant of Google was expected to demand about 103 megawatts of electricity – enough to power 82, 000 homes. Besides using the technical approach to reduce power usage, which also increases net profit for Google, Google also commits to shift some of their electricity demands to renewable sources like wind, solar energy and hydroelectric power (Glanz, 2012). It installed solar panels on the roofs at its Mountain View facilities and in 2010 it invested $39 million in wind power ("Criticism of Google", 2019).

Google's street view is also criticized for providing information that could potentially be useful to terrorists In the United Kingdom during March 2010, Liberal Democrats MP Paul Keetch and unnamed military officers criticized Google for including pictures of the entrance to the British Army Special Air Service (SAS) base, stating that terrorists might use the information to plan attacks. Google responded that it "only takes images from public roads and this is no different to what anyone could see travelling down the road themselves, therefore there is no appreciable security risk." However, Google was subsequently forced to remove images of the SAS base and other military, security and intelligence installations, admitting that its trained drivers had failed to not take photographs in areas banned under the Official Secrets Act ("Criticism of Google", 2019).

Political Issues

In a 2018 American documentary called *The Creepy Line*, Google's influence on public opinion was explored, together with its power that is not regulated or controlled by national government legislation ("The Creepy Line", 2019). One of the examples featured in the film is that during the 2016 presidential election, search engine data are studied by using a "Nielsen-ratings-type network of confidants." The result showed a significant pro-Clinton bias in Google search results, which Epstein estimated could have influenced millions of voters in her favor (Jekielek, 2018). Although Google said it didn't seek to manipulate political sentiment, or make "ranking tweaks" to search results, it is still questioned whether Google should have the power.

Matt Zook, a collaborator of Graham's who teaches at the University of Kentucky, demonstrated another example about what happens when someone performs a Google search for abortion: the user is led not just to abortion clinics and services but to organisations that campaign against it. "There's a huge power within Google Maps to just make some things visible and some things less visible," he notes (Grabar, 2015).

Although the EU has taken some very harsh positions on Google, and some of the regulations they implemented will be very damaging, tech moves infinitely faster than government. Tech giant like Google can easily move outside of the regulations by changing some tech specs. Will it be a Sisyphus's curse?

Legal Issues

For large websites like Google, its profit source and its dominating power comes from monopoly to some degree. Accordingly, they have to face anti-trust law in US and Europe.

In Europe, the European Commission has pursued several competition laws cases against Google. One of the complaints from European Commission is that Google abuses its position as a dominant search engine to favor its own services over those of competitors. Other comparison sites complained of a precipitous drop in web traffic due to changes in the Google search algorithm, and some were driven out of business. The investigation began in 2010 and concluded in July 2017 with a €2.42 billion fine against the parent company Alphabet, and an order to change its practices within 90 days (Kelion, 2017).

In US, Google reached an advertising agreement with Yahoo!, which would have allowed Yahoo! to feature Google advertisements on its web pages. The alliance between the two companies was never completely realized because of antitrust concerns by the U.S. Department of Justice. As a result, Google pulled out of the deal in November 2008 ("Criticism of Google", 2019).

Ethic Issues

The most sensitive ethic issues may be the privacy concerns. In the famous book *Who Owns the Future*, the 2014 Goldsmith Award winner, Lanier Lanier calls Google "Siren Servers," alluding to the Sirens of Ulysses. Instead of paying each individual for their contribution to the data pool, the Siren Servers concentrate wealth in the hands of the few who control the data centers. For example, Google's translation algorithm, which amalgamates previous translations uploaded by people online, giving the user its best guess. The people behind the source translations receive no payment for their work, while Google profits from increased ad visibility as a powerful Siren Server ("Who Owns the Future?", 2018).

Sometimes the ethic issues may evolve to rules such as the right to be forgotten enshrined by European (Chee, 2019). People can ask search engines like Google to remove inadequate or irrelevant information from web results appearing under searches for their names. As a response, says it has since received 845,501 requests to remove links, and removed 45 percent of the 3.3 million links it was asked to scrap. However, the balance between privacy and free speech is tricky. After France's privacy watchdog CNIL fined Google 100,000 euros for refusing to de-list sensitive information from search results globally upon request, Google took its fight to the French Council of State and finally won the case at the European Court of Justice. Google will not have to apply Europe's "right to be forgotten" law globally.

Summary

Both technical issues and non-technical issues are equally important to manager large web sites like Google. While technical issues are more important for a start-up company to survive in the market, the non-technical issues may cost much more than expected to companies managing large web sites. However, the solutions to those non-technical issues may not be so straightforward as the solutions to the technical issues. In some cases, maybe time is the only solution.

Reference Material

Chee, F. Y. “You have the right to be forgotten by Google - but only in Europe.” *Reuters*. Sep. 24, 2019. Retrieved from: https://www.reuters.com/article/us-eu-alphabet-privacy/google-wins-in-right-to-be-forgotten-fight-with-france-idUSKBN1W90R5. Accessed on Oct. 19, 2019.

“Criticism of Google.” *Wikipedia*. Oct. 11, 2019. Retrieved from: https://en.wikipedia.org/wiki/Criticism_of_Google. Accessed on Oct. 19, 2019.

“Configuration management.” *Google Cloud.* Retrieved from: https://cloud.google.com/solutions/configuration-management/. Accessed on Oct. 18, 2019.

“Configuration management” *Wikipedia.* Oct. 13, 2019. Retrieved from: https://en.wikipedia.org/wiki/Configuration_management. Accessed on Oct. 18, 2019.

“FCAPS.” *Wikipedia*. Oct. 10, 2019. Retrieved from: https://en.wikipedia.org/wiki/FCAPS. Accessed on Oct. 18, 2019.

Glanz, J. “Online Cloud Services Rely on Coal or Nuclear Power, Report Says.” *The New York Times*. Apr. 17, 2012. Retrieved from: https://www.nytimes.com/2012/04/18/business/energy-environment/cloud-services-rely-on-coal-or-nuclear-power-greenpeace-says.html. Accessed on Oct. 19, 2019.

“Google Infrastructure Security Design Overview.” *Google Cloud*. Retrieved from: https://cloud.google.com/security/infrastructure/design/resources/google_infrastructure_whitepaper_fa.pdf?utm_medium=et&utm_source=google.com%2Fcloud&utm_campaign=multilayered_security&utm_content=download_the_whitepaper. Accessed on Oct. 19, 2019.

Grabar, H. “Cracks in the digital map: what the 'geoweb' gets wrong about real streets.” *The Guardian*. Jan. 8, 2015. Retrieved from: https://www.theguardian.com/cities/2015/jan/08/digital-map-what-geoweb-gets-wrong-about-real-streets. Accessed on Oct. 19, 2019.

Heidi, E. “An Introduction to Configuration Management.” *Digital Ocean*. May. 17, 2019. Retrieved from: https://www.digitalocean.com/community/tutorials/an-introduction-to-configuration-management. Accessed on Oct. 19, 2019.

Jekielek, J. Su, N. “The Creepy Line’ Documentary Shows Tech Giants Influencing Voters.” *The Epoch Times*. Sep. 26, 2018. Retrieved from: https://www.theepochtimes.com/the-creepy-line-shows-tech-giants-influencing-voters_2671853.html. Accessed on Oct. 19, 2019.

Kelion, L. “Google hit with record EU fine over Shopping service.” *BBC News*. Jun. 27, 2017. Retrieved from: https://www.bbc.com/news/technology-40406542. Accessed on Oct. 19, 2019.

Metz, C. “Return of the Borg: How Twitter Rebuilt Google's Secret Weapon.” *Wired*. May. 3, 2013. Retrieved from: https://www.wired.com/2013/03/google-borg-twitter-mesos/. Accessed on Oct. 19, 2019.

Strand, G. “Keyword: Evil.” *Harper’s Magazine*. Mar. 2008. Retrieved from: https://web.archive.org/web/20120627035844/http://harpers.org/media/slideshow/annot/2008-03/index.html. Accessed on Oct. 19, 2019.

“The Creepy Line.” Wikipedia. Oct. 16, 2019. Retrieved from: https://en.wikipedia.org/wiki/The_Creepy_Line. Accessed on Oct. 19, 2019.

"Web 100." *Web100.com.* Jan. 21, 2009. Retrieved from: http://www.web100.com/category/web-100/. Accessed on Oct. 18, 2019.

"WEB SEARCH FOR A PLANET: THE GOOGLE CLUSTER ARCHITECTURE." *IEEE 2003, 0272-1732/03, pp. 22-28*. 2003. Retrieved from: http://static.googleusercontent.com/media/research.google.com/en//archive/googlecluster-ieee.pdf. Accessed on Oct. 18, 2019.

"Who Owns the Future?" *Wikipedia*. Aug. 12, 2018. https://en.wikipedia.org/wiki/Who_Owns_the_Future%3F. Accessed on Oct. 19, 2019.

Zhang, X. et al. "CPI^2: CPU performance isolation for shared compute clusters." *ACM 978-1-4503-1994-2/13/04*. pp. 379-391. Apr. 15-17, 2013. Retrieved from: http://static.googleusercontent.com/media/research.google.com/en/us/pubs/archive/40737.pdf. Accessed on Oct. 19, 2019.