

July 2021
Geoff Huston

CDNs and Centrality

On the afternoon June 17 of this year there was a widespread outage of online services. In Australia it impacted three of the country's largest banks, the national postal service, the country's reserve bank, and one airline operator. Further afield from Australia the outage impacted the Hong Kong Stock Exchange, and some US airlines. The roll call of impacted services appeared to reach some 500 services.

The problems were first noticed at around mid-afternoon Australian Eastern Time and were not fully restored until four hours later. The question is of course what do these 500 affected enterprises had in common? Why were they fate-sharing in an outage? This was not the result of a deliberate attack. But it was related to the measures many of the Internet's enterprises, both large and small, use to deflect the impact of such attacks. The answer is that they were all customers of Akamai and in particular Akamai's Prolexic DDOS protection system. Akamai's post-event response was:



ABC News Report on Akamai Outage

"A routing table value used by this particular service was inadvertently exceeded. The effect was an unanticipated disruption of service."

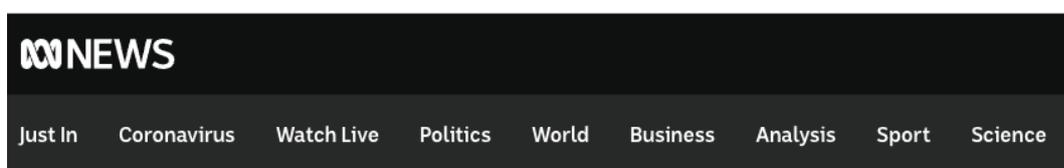
<https://blogs.akamai.com/2021/06/akamai-provides-prolexic-ddos-service-impact-update-status-resolved.html>

Yes, the Internet is still growing, and in the IPv4 network we are closing in on some 1 million routing entries, and in IPv6 the comparable number is now approaching 150,000 entries. Switching elements implement packet decision structures in various ways. Some use Ternary Content-Addressable Memory (TCAM), others use customized ASICs. The common objective here is to make a forwarding decision that is based on a lookup of a best match of the packet's destination address to a routing entry drawn from this pool of some 1 million entries. And this lookup has to be as fast as we can possibly make it. If I am building a very high-speed switch

that can handle the packet load of multiple 100G circuits then I probably would like to have my unit make some 100M or so switching decisions per second. That implies that I need to design a decision system that can perform this lookup across a set of 1 million entries and do so in just 10 nanoseconds. This is a very challenging objective. By comparison the fastest available memory cycle time is a little under 1 nanosecond. However, it's not quite as simple as that. If I am building such a switch element then I need to design it such that it can cope with the peak switching loads in that I can expect in 1, 2 or even 5 years from now. If I get it wrong and do not design sufficient capacity and speed, then my unit will hit obsolescence earlier and this will impose further costs on the network operator. If I over-design, it then I will have spent too much on state-of-the-art high-speed memory and my product will be far more expensive than my competitors. Every device is built with such design trade-offs and the useful life of these devices depends on such indeterminate factors such as the future growth rate of the Internet. The point is that of our networks use such equipment, and the role of the network operator is to constantly upgrade their switching equipment capability to keep just ahead of these demands. And means that it's inevitable that sometimes they will slip up. As did Akamai evidently. We can't just out-design this problem when we are sitting on the very edge of what silicon is capable of.

But this is not an issue that is unique to Akami. All network service providers are chasing a similar objective, and from time to time all of these systems may flip over to an overload state and have a service outage. What was notable about this particular outage was that this outage impacted a large set of online service providers. The business of being the back-end service provider to online platforms is not exactly a highly populated and diverse business environment. In fact, there are very few enterprises that offer such services, and if any of them experience operation outages then the impacts are going to be highly visible.

As it turned out this was not an isolated incident in June 2021. Earlier in the month we saw a set of popular services disappear for an hour or more. The list of impacted services included Twitch, Pinterest, Reddit, Spotify, the New York Times, and the BBC to name just a few. What did all these services have in common? They use Fastly.



ANALYSIS

Fastly global internet outage: why did so many sites go down — and what is a CDN, anyway?

ABC News Report on Fastly Outage

“We experienced a global outage due to an undiscovered software bug that surfaced on June 8 when it was triggered by a valid customer configuration change. We detected the disruption within one minute, then identified and isolated the cause, and disabled the configuration. Within 49 minutes, 95% of our network was operating as normal. This outage was broad and severe, and we’re truly sorry for the impact to our customers and everyone who relies on them.”

<https://www.fastly.com/blog/summary-of-june-8-outage>

Again, this was not an attack. It was more in the nature of a ticking time bomb where a particular customer configuration setting could trigger a failure in the shared CDN platform. And all that needed to happen was for a customer to stumble on that particular configuration setting and the system would fail.

Fastly is built on a configuration management platform called “Varnish”. As Fastly reports:

“Fastly is built on Varnish, which allows for high-powered content delivery including the ability to instantly purge content across its global network. Using Varnish Configuration Language (VCL), Catch can customize their Fastly configurations, resulting in more intelligent caching.

“Setting cache rules based on something as specific as the cookies in the request is something we would normally only consider possible running our own Varnish server. Having the ability to do so at the CDN layer makes our setup much more powerful and streamlined. We are also able to ensure certain file formats are always served from Fastly, and in theory never have to hit our origin more than once.

“Varnish is incredibly important to us because we can control the VCL if we want to make changes to the caching layer. With most CDNs, you don’t get that kind of flexibility, but with Fastly you do.”
<https://www.fastly.com/customers/catch/>

Varnish? What’s Varnish when it’s at home?

“Varnish was designed specifically to replace Squid, a client-side proxy that can be adapted and used as a web accelerator. Its main design goal was to increase scalability and capacity of content-heavy dynamic websites as well as heavily consumed APIs. Such sites run on web servers, such as Apache or nginx, primarily origin servers. which create the web content that is to be served. Varnish’s job is not to create content, but to make content delivery lightning fast.”
<https://info.varnish-software.com/blog/history-varnish-cache-10-years>

When the Norwegian online newspaper VG Multimedia reached over 45 million page views every week, it took 12 servers to handle the requests. Anders Berg, a system administrator at VG, believed there had to be a way to reduce the number of servers while also loading the pages faster. He wrote a spec for an open source project dealing with the issue, getting the ball rolling.
<https://www.varnish-software.com/varnish-history/>

So Fastly’s major asset is not some proprietary software system that performs a CDN function uniquely efficiently. Varnish is an open source software tool, like many others. Perhaps what makes Fastly special is its large customer base, and these customers probably use Fastly for their CDN requirements because others have also made the same decision to use Fastly. In other words what makes Fastly so special in the CDN world is its size, and this size is perhaps one of the more compelling reasons why others choose to use Fastly and contribute to its further growth. Size breeds more size.

There are a couple of issues with the CDN market.

Firstly, while its remarkably easy to select one CDN provider and use it for the entirety of one’s online content and service portfolio, it can be more challenging to select two or more such CDN providers and use them together in a self-healing mutual backup setup. For many online service enterprises, it’s a case of “It’s way easier to pick just one CDN. Choose wisely!” From that point on the enterprise is fate sharing with the CDN provider.

Secondly there aren’t that many CDN’s to pick from. If you are after a global footprint, sufficient capability to absorb all but the most extreme DDOS volumetric attacks, or even absorb all such attacks, a functional service interface that allows the CDN to optimise the service delivery yet leaves the customer in control of critical aspects of the security and integrity of the service (such as private keys) then your comparative shopping list is not exactly large. Fastly and Akamai are on most lists of the most popular CDNs, together with Amazon CloudFront, Google Cloud, Microsoft’s Azure, Cloudflare and Limelight. That’s seven. There are more of course, some with more regional focus, some with a particular technical speciality, but those seven enterprises are the core of today’s CDN provider world.

And perhaps that’s the real problem. In the content world it now seems that everyone uses a CDN in one form of another, but there are in reality very few CDNs to pick from. It’s a highly centralised space where volume economics dominates. The larger CDN providers can offer a service level at a price that is impossible for smaller CDN enterprises to achieve. This increases their market share, which of course adds to their size and

further exacerbates the scale of difference in the volume economics between the providers. The provider space bifurcates between the small set of “core” CDN giants and a larger “halo” of much smaller CDN providers who offer various customised approaches that fill whatever specialised gaps are left in the service profiles of the larger providers. The large providers effectively control the growth prospects of the smaller providers because they control what gaps are left in their service offerings. The natural outcome is that large becomes larger, and small tends to get smaller. The inevitable result is that when one CDN service provider experiences a service interruption, and such interruptions have happened in the past and will no doubt happen in the future, then the casualty list of interrupted services can be very large.

Normally, such outages would erode confidence in the service offering and customers would look at alternatives. The result could well be that the investors in the company would get nervous and the stock price would fall. Normally. But these CDN folk appear to be “special”. Fastly’s stock price has increased across the month of June and there is no residual effect of the outage in terms of the stock price.



Fastly Share Price – June 2021

And Akamai’s share price rallied after the 18th June!



Akamai Share Price – June 2021

A bubble is generally characterised by irrational exuberance in the part of investors, where the clamour to get a share of the as-yet-to-be-realised future bounty on offer completely drowns out the more sober commentary on the fundamentals of the market, and also ignores a mundane conversation about the basics of cost and revenue. And that too is one of the issues about a highly centralised market. In a highly centralised market, the rewards on offer for a successful enterprise that completely dominates their market are literally everything on offer in the market, and the prospect of a monopoly premium as a bonus! In the case of such bounties on offer, the market appears to be more than willing to ignore a minor inconvenient truth that the foundations of this CDN technology are not exactly rock-solid and various incidents that impact operations will continue to happen.

Like many other aspects of the Internet, we are seeing a diverse and robustly competitive environment being transformed into a set of highly centralised environments that lead to the formation of incumbent cartels and monopolies. Such a centralised environment creates a set of critical dependences, so that when failures occur, as is inevitable in this space, they are transformed from a minor inconvenience for a few into a major incident that impacts all of us in various ways. Size and centrality do not necessarily only create more robust services. As we’ve seen in these two particular outages in June, and others in the past (DYN in October 2016, for example) size and centrality also poses a greater level of vulnerability for everyone.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net