



F U T U R E

Varnish enables **Future Publishing** to achieve a robust and reliable caching layer for highly available content delivery across its media properties

Case Study:

Future Publishing

Future Publishing uses **Varnish** to achieve robust and reliable caching layers

Background

Future Publishing (Future), self-described as the global platform for specialist media, is home to more than 250 brands, including some of the world's most popular magazines, websites and events. Living under the Future umbrella are renowned titles and websites ranging from Marie Claire, Go.Compare, TechRadar, wallpaper*, to PC Gamer.

As a high-growth publisher with a strong acquisition-based strategy, Future is built around a mix of technologies and infrastructures. The site reliability engineering (SRE) team facilitates ongoing platform development to enable DevOps and Dev teams and ensure that sites are scalable and reliable.

Future's success relies on lean and innovative approaches, and the SRE team often adopts open-source solutions in order to capitalize on the flexibility to experiment and grow services if needed. The open-source approach to caching with Varnish Cache has served the organization well for a number of years. With high growth, however, growing pains were inevitable.

The challenge

Using Varnish Cache as an HTML page cache worked seamlessly when Future served three brands, and even when it served 30 brands. But as the brand expanded to 45 or 50 brands, the struggle for scalability and availability started. Future engineers began discussing how to scale the solution out and cope with resilience problems, such as node failures and their consequences, lost caches and the time required to rebuild.

Their approach yielded what they called Project Panther, which would create a concurrent architecture that would load-balance an incoming request among one of the servers. This load balancer would then "own" the content, and secondary requests would be routed accordingly, and so on. Although this solution worked for up to 60 sites, it became increasingly difficult to debug and diagnose problems, and scaling up the Panther cluster became an issue.

The challenges of doing everything in-house started to mount. "On the one hand, we wanted to make content immediately available, as our editorial staff expect," explained Toby Jackson, Global SRE Lead, Future. "And on the other hand, we want to make sure content is always available while coping with network outages, delays and other challenges to content delivery."

Future Publishing at a glance

Organization

- Future is a global platform specialist media.

Challenges Varnish solves

- Ensure high availability and scalability across multiple media properties
- Reduce number of SRE/support/on-call outage hours
- Slash downtime (e.g., from ave. 3 mins outage/day to 1.7 secs outage/day)
- Add resilience and disaster recovery at origin level

Varnish Enterprise

- Web and API acceleration features, including caching and cache invalidation, VMODs such as MSE and YKey

Moving to Enterprise, VMODs meant that we could use that distributed cache and cluster caching that reduced our load-balancing complexity as well as Ykey, which has optimized performance for cache invalidation and general storage.

-Toby Jackson, Global SRE Lead, Future.

The solution

Varnish Enterprise as robust and reliable caching layer

"Varnish is a key component of our content delivery arsenal," Toby continued. "When we started spending more time than we wanted to make our own solution work, we decided to talk to Varnish to see what technical approach we could take. In three weeks, we went from our four-node cluster with load balancers to the next iteration – what we call

Project Leopard – which is a new four-node Varnish cluster that more or less solved all our problems instantly."

Adopting Varnish Enterprise meant that Future could work with the basics they already knew from Varnish Cache – and simply extend its functionality to tackle some of the deeper scale-up pain points, challenges and bottlenecks they faced.

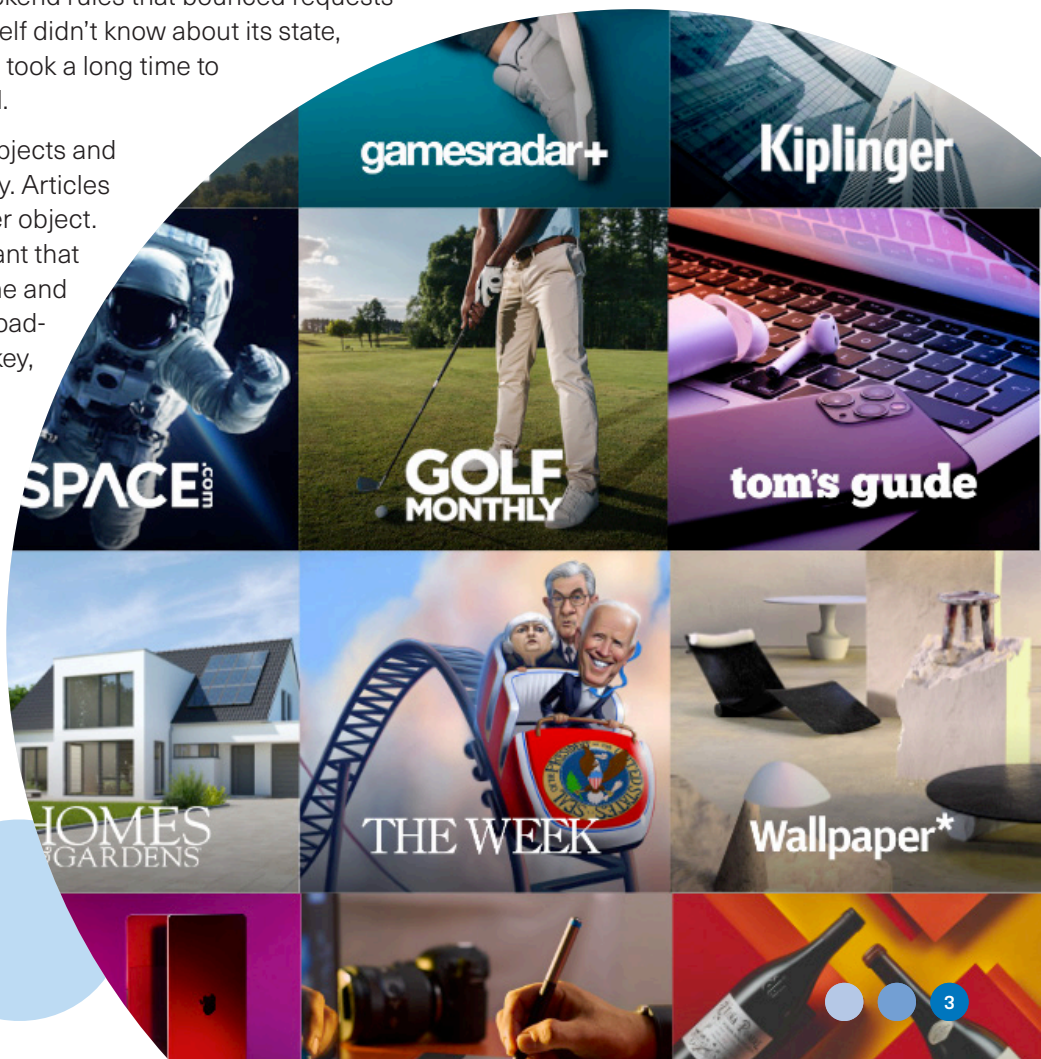
Varnish Enterprise as origin shield

Future's Varnish cluster runs behind their Fastly CDN, their London-based origin servers, and load balancers. Varnish Enterprise effectively shield their Kubernetes apps, which actually render the page content. This is done, Toby explained, "as a disaster-recovery strategy. When Fastly goes down, we can just swing everything to London and serve content directly from Varnish Enterprise. We've hit Varnish with tens of thousands of requests per second without crippling our backend services. It's also useful for observability – Varnish Enterprise allows us to get all the data we need in terms of page content delivery metrics and some cache invalidation tricks."

Varnish Enterprise - Simplifying with VMODs

In addition, the Future team found that having a consistent distributed caching topology baked into the system was considerably eased with Varnish Enterprise and its features/VMODs. "We previously used HAProxy and Varnish Cache along with ACL backend rules that bounced requests around but because the cluster itself didn't know about its state, it was somewhat failure prone and took a long time to recover in an outage," Toby shared.

"We store a lot of relatively small objects and eventually ran into issues with Xkey. Articles on our site could have 60 Xkeys per object. Moving to Enterprise, VMODs meant that we could use that distributed cache and cluster caching that reduced our load-balancing complexity as well as Ykey, which has optimized performance for cache invalidation and general storage. We could also store object caches to persistent disk (MSE) to survive node reboots. Varnish Enterprise has simplified things and added efficiencies."



Results

Because Varnish sits behind a CDN, immediacy and speed was not the primary goal. For Future, distribution and availability mattered more than latency. Launching just before Black Friday, Future decided to roll out the Varnish Enterprise solution just before the busiest period of their calendar year, and had the confidence to do so because “we knew it would work,” Toby revealed.

Increased availability: Downtime reduced from 3 minutes to 1.7 seconds per day

Before Varnish Enterprise, Future achieved approximately 99.8% availability. Since rolling out the new stack, they increased availability to 99.998%. As Toby explained, “It doesn’t sound like much on the face of it, but when you convert the metric, we used to lose two to three minutes and now only 1.7 seconds per day.”

No outages since launch

In the three months the solution has been running, Future has experienced no outages. Before the launch, the team had at least one, possibly two, on-call incidents per week, which each took about three to four hours to resolve. “Within three months, we have already saved at least one work week because the system has not failed once,” Toby exclaimed. “These incidents often involve disruptions for more than one person, and we lose incalculable time from other potential opportunities.”

Savings

Future experienced different quantifiable and non-quantifiable benefits from adopting Varnish Enterprise.

Eliminating failures meant that the team could reduce the size of the backend infrastructure they needed to cope with load.

“Realistically, because we had done so much work to control the blast radius when we had failures, a lot of that was hidden from view,” Toby reasoned. “We were still successful – but at a cost. We were putting a lot of load on backend servers, so a lot of our backend infra was bigger to cope with that recovery time. It does not fail now, so we do not need as much server power. There is a definite side benefit in not having the expense of running that extra infrastructure.”

Functionally, the team found that by having a more robust and reliable caching layer with Varnish Enterprise, fewer requests were made of their core, which in turn meant they didn’t need to have as much core to run. Less is more because the cache is doing a better job.

While less quantifiable, the team, tired of the on-call alerts and knowing they could do better, weighed their options: get help by hiring another full-time employee or buy Varnish Enterprise. What they realized is that in adding another person, the whole team would still experience the disruption caused by incidents, which they felt would not be making the best use of the team they had. Someone would still have to deal with the problem, regardless of whether a new employee joined. Varnish Enterprise would make the problem go away entirely, helping the team use its time, resources and focus more effectively.

“In quite a few cases, we may have challenged the Varnish engineers with some weird questions,” Toby said. “This was reassuring to us because it meant that what we are doing is not simple, and it is worth spending the money on an enterprise solution.”

In addition, Toby described how open, honest communication with the Varnish team gave the team the confidence to try new things, whether they faced challenges or just wanted advice. “We trust the frank honesty and clarity we got alongside responsive support. The value added by Varnish is more than transactional or the cost of a technical license key.”



Los Angeles - Paris - London
Stockholm - Singapore - Karlstad
Dusseldorf - Oslo - Tokyo

FUTURE