



Customer Case Study

Industry:
Traditional Media

Problem Statement

A traditional media company has experienced tremendous spikes in user traffic after making episodes of some of its most popular TV shows online.

The so called "viral" nature of such content lead to 5-10x increase in user traffic, in span of just few minutes, to company's web servers when more popular episodes were promoted on-air.

The existing setup included a front-end web farm of 14 web servers, running Apache/PHP. The custom code connected to a MySQL 5 DB, using "media ID" as a key to retrieve "media metadata": clip's name, show, data of air, cast and brief description.

An image (frame capture from the video itself) is also displayed. Links were provided to discussion forums dedicated to the show.

After just a few production runs it became clear that the existing setup could not cope with the user traffic. Load averages on Web servers would spike well above 20. Response time would quickly increase to 10-15 seconds and after a short while the web site would become unresponsive, taking down complete "video player" infrastructure. The requests for the dynamic URLs would peak at around 245 req/sec, as reported by load balancers, before the site went down.

An attempt was made to grow the DB side into a cluster of 1 master node and 2 read replicas. However it didn't solve the problem and issues continued.

An engineering estimate was put together. Under consideration was growing the front-end and the back-end farms 3-fold. Outside of significant HW cap-ex, such expansion would've further depleted already limited available space and power at existing datacenter.

Traditional Media Company with extensive online presence.

Hosting infrastructure consists of redundant datacenters, hosting number of websites , utilizing wide range of technologies:

PHP, Java and ASP on Linux and Microsoft Windows platforms.

Some of the content is delivered using a major CDN.

Hardware footprint

In excess of 400 servers.

Typical HW server configuration: 1RU FF server with 2 dual-core Intel or AMD CPUs, 4GB or 8GB of RAM, gigabit attached to Cisco Switches.

The farms are load balanced via dedicated, redundant load balancers and protected via dedicated, redundant firewalls.

aiCache to the rescue

After finding out about web accelerator product from aiCache, a decision was made to arrange for a quick trial. To facilitate a speedy prove-of-concept, total of 2 dual-core servers were repurposed from existing inventory. They already had 64b Linux distribution installed and required no further changes to install aiCache.

aiCache configuration

Upon examining the typical use-cases and HTTP traffic patterns, aiCache was initially configured to cache output of PHP scripts for 10 seconds. That included output that rendered the "What's new" , "Most popular" and show-specific video pages and the page that rendered the meta-data for particular media-ID.

There also was a small number of .js and .css files that were served off the farm - these were configured for 24hr caching, as they rarely change.

The results

As result of aiCache deployment, traffic to the origin server farm was reduced about 24-fold ! The origin web server have become virtually idle - even when the most popular videos went live. The number of connections to the DB cluster has also dropped to single digits.

After 2 weeks of having aiCache in the mix, front-ending user traffic, a decision was made to downsize web farm to just 3 server (down from 14) and to downsize the DB setup to 1 server (there was a pre-existing HA setup).

In light of this remarkable reduction of load on origin servers and overall hosting infrastructure, the Customer saw opportunity to add user-comments to the video pages - re-using some of the servers that were freed as a result of deployment of aiCache.

Additional benefits

aiCache's rich instrumentation was put to good use, allowing for real-time monitoring of user traffic and identification of slower scripts on origin servers.

aiCache Hardware:

2 dual-core 64bit Intel-based 1RU servers, with 16GB of RAM. The rest of configuration is set up using customer's existing standards: 64bit Linux distribution from a well-known vendor, complete with OS support, redundant hard drives with a RAID array. Redundant network connection.

Savings:

- **28 new servers at ~US\$5000 per: US\$140,000.**
- **OS licenses: about US\$20K**
- **DB licenses: about US\$15K**
- **Server install and setup charges: about US\$7K**
- **re-purposing of 13 servers and OS licenses for other applications: US\$75K**

Total CapEx saved: more than US\$200K.

Being able to stay up under even heaviest traffic, while reducing HW footprint, saving space, power and cooling in existing datacenters: priceless.

aiCache's SNMP integration is also utilized - the req/sec and response time is now collected, charted and alerted on by pre-existing SNMP monitoring package.

aiCache's selective log suppression feature is also configured to not log javascript (.js) and content style sheets (.css) requests.

aiCache's on-the-fly compression was turned on for a number of larger .js and .css scripts, resulting in measurable reduction of egress traffic. A significant reduction in requests for auxiliary content was also observed - most likely due to Jxel 's elimination of HTTP response conditionals.