

## How Fujitsu Racing is super-charging its OpenCMS performance



 **Melbourne Business**  
online

Author: Dammian Miller

# About me

- From an island called Australia
- Web developer for 10+ years - Java focus, government, corporate, private, consulting
- OpenCMS since 2005
- Implemented a number of small and large projects using OpenCMS, notably:
  - Australia's leading loyalty program website
  - Australia's leading department store members site
  - Fujitsu Racing V8 Supercars team website
  - And a blatant product placement - my company's fully managed multi-site OpenCMS offering
- OpenCMS evangelist

# Here to introduce the open-source OpenCMS HighPerf module

But before I do, a brief history leading to its creation.

Once upon a time...



# Business Case: Fujitsu Racing V8 Supercars Team Website

## OpenCMS powered V8 Supercars

Complex site comprising many components :

- News
- Events
- Merchandise
- Image Gallery
- Video Gallery
- Form
- Edit areas
- RSS feeds

The screenshot shows the Fujitsu Racing V8 Supercars Team Website. The header includes the Fujitsu Racing logo, the Britek Motorsport logo, and a large image of a blue and white V8 Supercar. The navigation menu includes Home, About, Events, Media, Sponsors, and Merchandise. The main content area is divided into several sections:

- SPONSORS:** Fujitsu, Ford Racing, CAT, Penrite Performance Oils, NGK Spark Plugs, Delphi, WesTrac, and CAT.
- HOME - NEWS:** Two news items: "CONSISTENT RUN FOR ASSAILIT AT WINTON" and "FUJITSU RACING SCORE FIRST TOP-10 OVERALL V8 RESULT...".
- MERCHANDISE:** A grid of t-shirts and caps.
- STANDINGS:** Position: 12, Total points: 964.
- EVENTS:** "OPEN DAY 2008" on Saturday 30 August 2008, 9am - 3pm.
- GALLERY:** A grid of images showing the team and the car.

The footer includes the copyright notice "Copyright © 2008 Fujitsu Racing. All Rights Reserved" and accessibility icons for WCAG 2.0, W3C, and 508.

# The client - delivery requirements

- Trackside update of news, image and video gallery
- Fine-grained user permissions
- Sub-2 second initial browser display - HTML delivery
- Sub 8 second until completion time
- Performance to be delivered at up to 100 concurrent users
- Delivered yesterday

**Heard that before - OpenCMS is my choice, again, however...**

# OpenCMS - highly dynamic software

*With great power comes great responsibility...and system requirements*

- A rule of thumb - the more flexible the software, the slower it will be
  - More calls to data sources
    - Local - datastore - low latency
    - Remote - RSS feeds, web services - high latency
  - Data aggregation and formatting requires additional steps
- The slower a system, the less it will scale as existing resources are unavailable for longer

# *How we made OpenCMS lightning fast*

Decision criteria:

- Local audience - Content Delivery Network (CDN) == overkill
- Time constraint - due in 2 weeks
- Limited budget - more hardware not an option

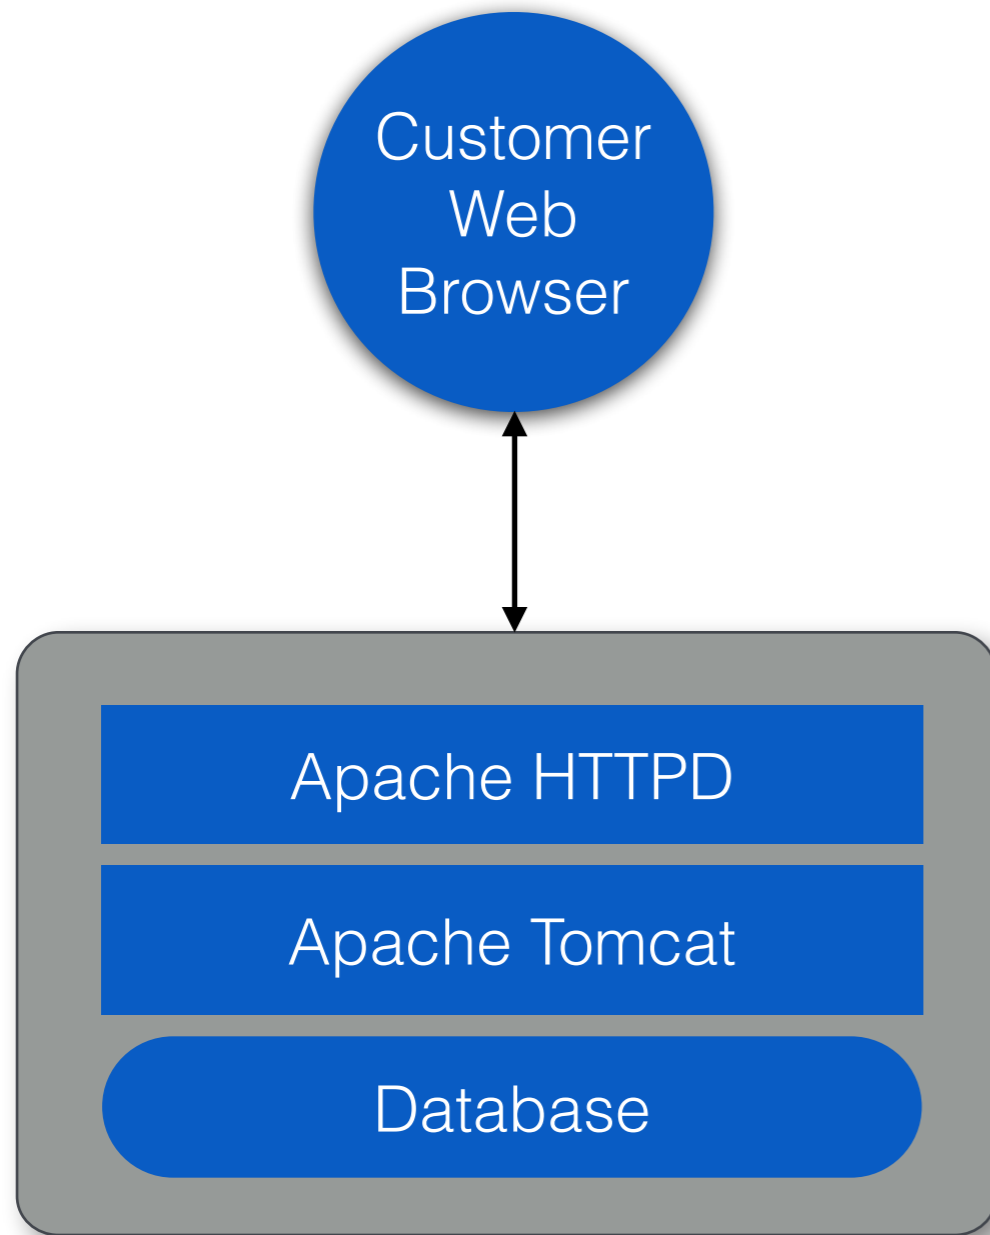
We chose to implement a reverse proxy cache/HTTP cache in front of our Tomcat + Apache web stack.

# What is a reverse proxy cache/HTTP Cache?

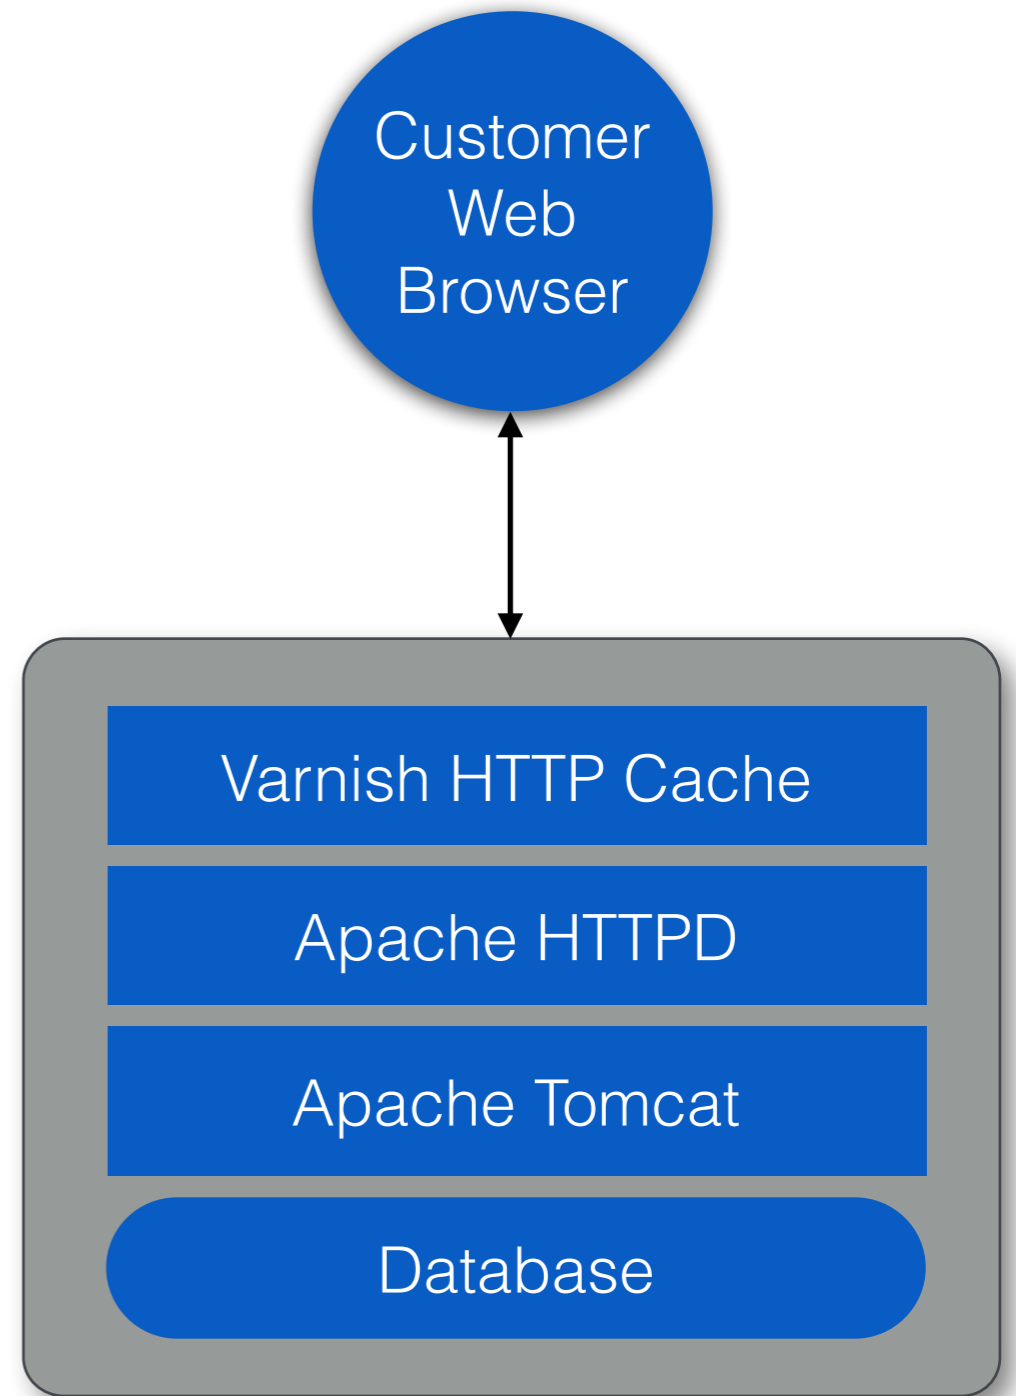
- HTTP Cache sits in front of traditional HTTPD server or the application server taking over as the client connection point
- Populates cache from backend HTTPD or application server directly
- Serves cached pages from memory - optimised
- more efficient with system resource use for each connection - more connections for the same resource usage



## Typical OpenCMS Stack



## Typical OpenCMS Stack with Varnish HTTP Cache



# Why choose a HTTP Cache?

- Non-invasive - no code changes
- Field tested - proven technology
- Works with static *and* dynamic content
- Quick to implement - 30 minutes setup time
- FAST - It is much more efficient to serve pages from optimised cache memory without web or application server overheads
- Efficient - increased system capacity due to efficiencies gained from bypassing web and application server overheads
- Potential for reduced hardware requirements - \$ savings


# Varnish Cache - HTTP cache extraordinaire

- Commercially-backed Free Open-Source Software (FOSS)
- Optimised HTTP cache, not jack of all trades
- Highly configurable via internal domain configuration language - VCL
- VCL supports RegEx., inline C
- Multiple backend support with programmable routing using VCL
- No SSL support
- Supports Edge-Side Includes (ESI) - more on this later

# Varnished OpenCMS Fujitsu Racing - the numbers

- Simple indicative test - not real-world
- UAT system
- LAN-based test - not real-world
- *Synthetic* benchmarks using Apache Bench
- `ab -c 100 -n 1000`
- Test scenario - 100 concurrent users repeated 1000 times

# Supercharging...start your engines!





# So we have our speed, but now we need to pay the price - content freshness - dynamism

- HTTP cache is unaware of when backend content is updated
- Content refresh is time-based
- Different content types or paths can have different expiry times. Vendor-dependent implementation and capabilities

# How we made OpenCMS lightning fast - *without* sacrificing dynamism

## OpenCMS HighPerf - PURGE on publish

- Vendor independent HTTP PURGE request type. Tells HTTP cache to purge requested resource from cache
- Triggered on publish event
- Next resource request results in a cache miss - HTTP cache will request and return fresh resource from web/application server
- Enables HTTP cache benefits without sacrificing content freshness with time-based purge/refresh policies.
  - Memory-based content lookup and delivery - fast
  - Optimised system resource use
  - Increased system capacity

# OpenCMS HighPerf information

- OpenCMS self-contained module including all dependencies
- LGPL license
- Source code contained within module including all dependency modifications
- Self-signed certificates need to be added to local JSSE certificates store
- Extra parameter to tell OpenCMS HighPerf to allow self-signed certificates - not recommended for production environments
  - `-Dau.com.melbournebusinessonline.opencms.publish.allowSelfSigned`

# Future Directions - Enterprise capabilities

## Cluster awareness

- Enable purging to multiple backends behind a single domain
  - We can achieve this using the OpenCMS site aliases.  
Limitations?
- Purge piping via a single HTTP cache in front of a cluster of HTTP caches?

# Future Directions - Edge Side Includes (ESI)

## ESI - Edge Side Includes

- Ratified HTTP cache specification - vendor independent
- Enables HTTP cache to cache individual page components
- assemble page from components
- Only purge components that have changed
- Fine granularity compared to page-level caching
- Supports compression



# Future Directions - Edge Side Includes (ESI) - page components

The screenshot shows the Fujitsu Racing website layout with the following components labeled on the left and pointed to by red arrows:

- Header:** Points to the top navigation bar containing the Fujitsu Racing logo, the Britek Motorsport logo, and a large image of a blue and yellow racing car.
- Navigation:** Points to the menu items: Home, About, Events, Media, Sponsors, and Merchandise.
- Race schedule:** Points to a row of 14 track icons labeled R1 through R14.
- Advertising:** Points to the 'SPONSORS' section, which lists Fujitsu, Ford, CAT, Penrite, NGK, and Delphi.
- Video Gallery:** Points to a video player showing a V8 Supercars Australia Championship Series race.
- News:** Points to the 'HOME - NEWS' section, which features two news articles with images and text.
- Merchandise:** Points to a grid of merchandise items, including t-shirts and caps.
- Driver Info:** Points to the 'DRIVER' section for Jason Bright, featuring his name and a photo.
- Series Standings:** Points to the 'STANDINGS' section, which shows Jason Bright's position (12) and total points (964).
- Newsletter Form:** Points to the 'NEWSLETTER' section, which says 'coming soon'.
- Events:** Points to the 'EVENTS' section, which features an 'OPEN DAY 2008' event with a date and time.
- Image Gallery:** Points to the 'GALLERY' section, which shows images of the racing car and two women.
- Footer:** Points to the bottom of the page, which contains copyright information and accessibility icons.

# Future Directions - Edge Side Includes (ESI) - Implementation

Initial investigation has revealed the following implementation methods:

- Implement via template mechanism
  - Potentially no OpenCMS changes required
  - Requires explicit re-coding of site template to utilise ESI
  - Relies on ESI support or fails?
- Implement via OpenCMS FlexCache
  - Requires OpenCMS modifications
  - Does not rely on ESI enablement - only enables if ESI support present
  - Transparent enabling for all existing site templates
- Other methods - ADE?

# OpenCMS HighPerf - where to get it?

Get OpenCMS HighPerf from the following URL. Be sure to check back for updates and future modules!

- <http://MelbourneBusinessOnline.com.au>

# Questions & Answers, Feedback & Suggestions

- Dammian Miller
- Melbourne Business Online

**Thank you for attending**