

Scaling to Olympic Proportions

BBC sets the gold standard for interactive digital Olympic Game coverage

By Diane Burley

The Olympics are a dream and a destination – a culmination of years of planning, hard work and dedication that all plays out on a world stage, obviously for the athletes, and in this case, the British Broadcasting Corporation as well. With London the host of the 2012 Summer Olympic Games, the BBC was the host of coverage – and it wanted its online experience to be as dramatic as Danny Boyle’s opening and closing ceremonies.

We are addicted to, and dependent on, an exponentially increasing amount of data; everything from internet text, search indexes, call records, medical records and e-commerce transactions to sales and marketing data, financial figures, and more. We need it to make decisions fast, deliver effective customer service, and satisfy regulators. In short, we can’t do business without it.

“Our aspiration was that just as the Coronation did for TV in 1953, the Olympics would do for digital in 2012,” explained Phil Fearnley, General Manager, News & Knowledge, BBC Future Media.

The mission was lofty. There were more than 10,000 athletes to cover, a sophisticated audience that was relying more on social media, the need for continuous live coverage – and a demand for multichannel delivery. The steadfast process of Static Publishing, which had been relied on for over 15 years, had to be retired, in favor of a Dynamic Publishing infrastructure.

Dynamic Publishing has evolved to mean much more than to just dynamically serve content from a database and onto a page. Dynamic Publishing means creating a collection of related data

Rank	Country	Athlete	Time	Rank	Country	Athlete	Time	Rank	Country	Athlete
1	BRA	NAZARET MP	8:52.81	5	IRI	HAGHI A	+0:16.72	9	MAR	LAHSAINI M
2	DEN	BAK LY	+0:00.87	6	IRL	MCCANN D	+0:20.64	10	TUR	AKDILEK A
3	SUI	ALBASINI M	+0:01.98	7	VEN	GIL MARTINEZ TA	+0:21.24		COL	DUARTE AREV
4	JPN	BEPPU F	+0:08.13	8	KAZ	BAZAYEV A	+0:24.17		NED	BOOM L

During live-streaming the user could choose different views to appear at the bottom of the iPlayer. Here, athlete information populates the screen.

elements and dynamically serving it as audiences demand. According to Matt Turner, CTO Media, MarkLogic Corporation, “getting the right content to the right user in the right format is key to keeping users engaged and maximizing the investment made in content. Like the BBC site, dynamic publishing can have huge payoffs with record-breaking breadth and depth of coverage.”

In the case of the Olympic coverage, the flow of content was enormous, LIVE-stop, real-time and went across every channel

from web, mobile, tablets and broadcast and included:

- Event statistics and venue data updated in real time
- Athlete and event data from the International Olympic Committee
- 27 video streams from every venues all around the UK
- Feeds from the Press Association and other news organizations
- Stories from BBC journalists
- And of course, endless Social media updates from around the world

One example of dynamic updates was the athlete page on über-swimmer Michael Phelps (one of 10,000 athletes) included a journalist's story about Phelps, as well as dynamically-generated information that updated real-time including a box entitled "Phelps Against the World" (his medal count vs others), Phelps' event performances, where he was swimming next.

It would have been impossible for the team of journalists and editors to maintain this level of detail for the star athletes, and certainly out of the question for every single other athlete that competed in the games – and it was just as impossible for the relational database that had been the workhorse since the 1990s to keep up as well. Dynamic Publishing required a new, flexible architecture that would allow the various data feeds to

- Make Products more relevant with content and data from multiple sources delivered in disparate formats and in real-time
- Make all your content available with asset search and discovery
- Assemble custom content and collections
- Deliver in multiple formats through multi-channels

So what changed? In a November 2012 webcast, Jem Rayfield, Lead Architect, BBC, explained there were two key components critical to the transformation: An enterprise NoSQL content store and a "triple store." The combination allowed an unparalleled level of automation and dynamic delivery. The triple store uses linked-data technology to automate aggregation, publishing and repurposing of interrelated content objects – all driven by an ontological, domain-modeled information architecture. It's an organizational system that conveys that "Michael Phelps" was a member of the 2012 Olympic team, a member of the US swim team, of the men's swim team, of the 4x200-Meter Freestyle Relay, competed in events and heats and won a variety of different "awards." With each medal won, a tote board that depicted Phelps total medal count against all countries automatically updated – real-time. The triple store alone, however, could not process and store

the massive amount of changing data. To handle the volume and ensure ability to scale, the BBC added a MarkLogic to store all assets including stats, tweets, video metadata, images and articles. Video metadata included transcriptions to time-codes, so specific segments of video could be served. MarkLogic is an enterprise NoSQL database that uses XML as its data model. It allows the easy load of multiple data types into a single database – regardless of schema, processes high volumes of content and data in real time and scales as needed with a shared nothing architecture.

The resulting BBC Dynamic Publishing system produced a record number of content pages that maximized the editorial effort to creating the relationships and content and leveraged the system to create the index pages, team pages, schedules and additional content fed to users as they interacted with content. For example, knowing that viewers were often engaged with a "second screen" – a mobile or tablet from which they could actively engage with social media, the BBC also served display-adjusting content to four channels: Interactive television, computer, mobile and tablet. By encouraging second-screen behavior, social media became an integral part of the coverage. The user could customize the feed by choosing content elements that would appear on the BBC's iPlayer. One of the features of the BBC app allows viewers to restart a live feed.

The pas de deux of OWLIM Triple Store and MarkLogic Server served up an astounding:

- 106 million requests for BBC Olympic video content
- 55 million global browsers across the games
- 2.8 Petabytes of Data on a single day
- A daily record of 7.1 million UK browsers

Breaking records for viewership, downloads, and overall user experience – the BBC has set the gold standard for digital coverage. And, with the Olympics behind it, the BBC is looking to revamp its news operations as well. Others broadcast entities are keeping close tabs: The host city for the 2016 Olympics is Brazil, and according to Rayfield, engineers from Globo, the leading Brazilian television network, have reached out to BBC's technical team to get some pointers. With Brazil's population expected to exceed 200 million by 2016 – nearly 3 times the size of the population of the UK, the records set by the BBC – could likely be surpassed.

A former journalist and digital media executive, Diane Burley is the Chief Content Specialist for MarkLogic Corporation.



Find out why organizations choose MarkLogic for their dynamic publishing.

Call +1 877 992 8885

Email sales@marklogic.com

Visit www.marklogic.com

MarkLogic | 999 Skyway Road, Suite 200 | San Carlos, CA 94070