

Couchio

BBC - CouchDB Replication

Case Study

The BBC has a public obligation to ensure that their main site www.bbc.co.uk is always on, 24/7, no matter what is going on in the world. Even if a catastrophe affects their datacenters they have contingency plans to make sure that their website stays up and is able to provide the latest news to readers wherever they are.

When designing a new foundation for the site, BBC architects evaluated many databases to find the one that best fit their requirements. The primary need was to create a solution that could easily replicate between their two current datacenters, but also that would have the ability to replicate between more datacenters in the future. In addition they wanted to scale the solution horizontally. They also needed the database to write to disk in a consistent and reliable way while handling their large and growing dataset.

Solution

The BBC architects chose CouchDB to create a multi-master multi-datacenter failover configuration. This allows them to use 32 nodes split between two datacenters. Of the 16 nodes in each datacenter, 8 are primary nodes and the other 8 are backup nodes, but the nodes themselves are not aware of the fact that they are designated as a primary or backup node. This works well for the BBC because they can commission more nodes as their need for capacity rises.

The CouchDB solution allows them to do a hot backup while running, and they have never had a CouchDB node fail to restart. They are working with a terabyte of data, serve billions of requests and CouchDB does not have any issues keeping up with any of it.

Results

With their architecture in place, developers at the BBC can focus on building rich internet applications at massive scale. Operations teams can relax and do conservative capacity planning instead of putting out fires left and right. CouchDB doesn't let the BBC and thus the British public as well as the worldwide audience down, even if a datacenter is unavailable.

About the BBC

The British Broadcasting Corporation (BBC) was founded in 1927 and is currently the largest broadcasting organisation in the world.

The BBC's main responsibility is to provide public service radio, television and internet broadcasting within the United Kingdom and to enrich people's lives with programmes and services that inform, educate and entertain.

Quote

Enda Farrell, Software architect and Chief CouchDB Architect at the BBC says, "The BBC chose CouchDB because of its operational stability, its robustness and the N-master replication feature. With CouchDB our developers can release new features to the public sooner, and ensure everything is always available by spreading across multiple datacenters.

CouchDB - we like it."

Who we are



About CouchDB

Apache CouchDB bridges the cloud and the local network by making it simple to write web apps that can scale up to the datacenter and scale down for local use on a smartphone. The synchronization engine (inspired by Lotus Notes) allows sharing of data and applications across ad-hoc groups. Offline replication gives mobile devices better performance (lower latency) while saving battery life by avoiding radio usage. With offline replication solved in a simple way, application developers are free to focus on the code that matters

CouchDB is popular with developers because its document model is a closer fit for many domains, leading to less overall complexity. The JSON HTTP API allows developers to write dynamic database-backed applications without a middle tier: 2-tier CouchApps require only a web-browser and a CouchDB, and can be written by anyone with Ajax / jQuery experience. We firmly believe that empowering people to share data and custom applications at the edge of the network will lead to a more humane web.

Apache CouchDB is currently deployed by the BBC, Meebo, Assay Depot, Engine Yard, (among others) and is an integral part of the Ubuntu operating system.

Couchio

About Couchio

Couchio <http://couch.io> (founded by 3 CouchDB core committers) offers support, hosting, and professional training for CouchDB.

Contact CouchDB

www.couchdb.apache.org

Contact Couchio

www.couch.io
hello@couch.io