



---

## Highlights

- Help open the door to business opportunities that were previously unachievable
  - Power systems of engagement and analytics using Couchbase NoSQL technologies
  - Accelerate business-critical applications with easy-to-deploy IBM® FlashSystem™ all-flash storage
  - Increase the capabilities and value of 21<sup>st</sup> century applications, from online gaming and high-velocity retail to financial transactions and the Internet of Things
- 

# Simply accelerate NoSQL data stores

*IBM FlashSystem accelerates Couchbase NoSQL systems of engagement and analytics*

The business world is undergoing rapid transformation these days, as industry after industry shifts to an economy powered by the Internet and 21<sup>st</sup> century technologies such as the cloud, big-data analytics, mobile and social systems of engagement, and the Internet of Things (IoT). To support this shift, a new breed of data-management tools is being deployed to address requirements of the web-based, mobile and IoT applications used to power the cognitive business.

Chief among these tools are NoSQL databases/stores that provide mechanisms for data storage and retrieval modeled in means other than the tabular relations used in traditional relational databases. Driven by today's new application workloads, various types of NoSQL technologies have increased in popularity in the past 10 years—usually because they can run much faster and are less expensive than traditional relational databases. NoSQL technologies can make systems of social engagement, big-data analytics and real-time web applications faster as well.

NoSQL deployments are normally architected by clustering multiple commodity servers filled with costly DRAM to achieve fast performance. This approach has a number of limitations, including cost, efficiency, resource utilization and energy consumption. To address these challenges, enterprises often integrate flash storage into their larger NoSQL deployments. Flash serves as a lower cost, higher density, persistent extension of



the existing DRAM, providing exceptional performance compared to spinning disk storage technologies. For this reason, NoSQL databases with flash storage configurations are quickly gaining popularity.

A top solution for the NoSQL marketplace is the Couchbase database and IBM FlashSystem storage. Couchbase provides an enterprise-grade distributed NoSQL database engineered for performance, scalability and availability. IBM FlashSystem all-flash storage adds extreme performance, microsecond latency and unmatched operational efficiency. The combination of the Couchbase platform and IBM FlashSystem storage yields a cost-effective high-performance NoSQL solution that opens the door to possibilities that were previously unachievable.

## **The advent of NoSQL**

NoSQL technologies were invented during the past two decades to solve the data management challenges experienced by very large, data-intensive Internet companies such as Google, Facebook and Amazon. Today, however, virtually every enterprise now faces requirements that are most effectively addressed by NoSQL solutions. Three major trends are driving this need:

- The Internet and the cloud have become ubiquitous. Enterprises are moving more and more of their applications online and into cloud-based infrastructures.
- Mobile devices are rapidly becoming the preferred way consumers engage with businesses.
- The amount of data companies must manage has exploded, and solving problems with relational databases involving the ever-increasing data volumes is growing too complex and too expensive.

These trends translate into new requirements that companies didn't face 10 or 20 years ago:

- Large numbers of users and/or data sources, sometimes into the millions at once
- Low latency under high throughput
- High data volumes
- Increasing amounts of unstructured and semi-structured data, not just the rows and columns of structured data that were the norm 10 years ago

Relational databases cannot adequately meet these new requirements. Relational database management systems (RDMS) have served enterprises well for more than 30 years, but they were designed at a very different time and for a different set of requirements. As a result, they struggle to meet the new requirements. For example, RDMS solutions cannot:

- Deliver the highly responsive experience that today's consumers demand online
- Affordably and easily scale to support millions of users. RDMS solutions require enterprises to add expensive hardware to scale up; they do not scale out by deploying lower cost standardized hardware
- Efficiently manage massive volumes of data at high speed
- Handle many different data types, including unstructured and semi-structured data

## **Introducing Couchbase NoSQL technology**

Founded in 2010, Couchbase is already one of the performance and scalability leaders in the NoSQL space. From the beginning, the company focused on innovation and architecting a product to deliver the high performance, availability and scalability required to run mission-critical applications for customers from industries as diverse as equities trading, healthcare, online auctions, manufacturing, retail and many others.

The Couchbase platform—which includes Couchbase Server, Couchbase Lite (the first mobile NoSQL database) and Couchbase Sync Gateway—was purpose-built to solve today's requirements for enterprise-grade web and mobile applications. Specifically, Couchbase delivers:

- **Fast performance at scale**—Couchbase solutions offer sub-millisecond latency to enable highly responsive applications for millions or even hundreds of millions of users.
- **Easy, affordable scalability**—Couchbase provides a distributed database that scales out on commodity hardware with push-button simplicity on-premises or in the cloud.
- **High availability**—The Couchbase platform automatically replicates data across servers, clusters and data centers, so data is always available, 24x7, with no downtime required for maintenance.
- **Data model flexibility**—Couchbase solutions provide outstanding flexibility to handle any kind of data and change the data model on the fly to accommodate new data attributes or new data types. The Couchbase platform eliminates rigid database schemas, enabling developers to build applications faster and easier.

### Innovation in multiple dimensions

The latest version of the Couchbase platform, Couchbase Server 4.0, combines the powerful query capabilities of a relational database with the performance, scalability and flexibility of a NoSQL database. With Couchbase Server 4.0 you can use a NoSQL database to query billions of documents and meet the performance and scalability requirements of enterprise web, mobile and IoT applications. Couchbase Server 4.0 enables companies to develop with agility and deploy high-performance applications at virtually any scale.

Couchbase Server 4.0 includes significant advances in both architecture and features, including multidimensional scaling (MDS), *Non-1st Query Language* (N1QL) and the Couchbase ForestDB storage engine. These capabilities dramatically extend the scalability and performance advantages of Couchbase Server, as well as its ability to support an even broader set of use cases:

- N1QL combines the expressive power of SQL with the flexibility of JavaScript Object Notation (JSON) documents. N1QL enables queries to Couchbase Server with standard database drivers, preferred enterprise or web frameworks, Representational State Transfer (RESTful) web services or simple application programming interfaces (APIs).
- MDS provides the ability to manage core database functions—query, index and read/write—as individual services. Different services can be enabled on different nodes to optimize the resources for each, enabling you to run query services on nodes with faster processors, scale a single service as needed and add more query or index resources without adding more read/write capacity.
- Global Secondary Indexing (GSI) offers the ability to manage indexes on dedicated nodes. You can store the index on as little as one node while storing the data on many nodes, providing an efficient solution to the problem of querying at scale.

ForestDB is a key value storage engine developed by the Couchbase Caching and Storage Team. Its main index structure is built from hierarchical B+ Tree-based Trie. Compared with traditional B+ Tree-based storage engines, ForestDB shows significantly better read and write performance with less storage overhead. ForestDB has been tested in various server operating system environments (including Centos, Ubuntu, Mac OS x and Microsoft Windows) and mobile operating systems (including Apple iOS and Google Android). ForestDB is explicitly designed to be used with flash storage.

### Multiplying use cases

Enterprises are implementing Couchbase solutions across a wide range of use cases, each with its own set of requirements such as very high performance, very high availability and extreme data model flexibility. Major enterprises across many different industries are adopting Couchbase NoSQL technology to meet their data management needs:

- Consumer electronics and technology companies such as eBay and Cisco
- Retail companies such as Walmart and Tesco
- Financial services companies, including VISA and Wells Fargo
- Telecommunications companies, including Verizon, AT&T and Vodafone

### Introducing IBM FlashSystem 900

In order for NoSQL technologies such as those provided by Couchbase to deliver the speed and performance advantages over conventional relational database solutions that enterprises now require, high-performance non-volatile storage is necessary to reduce the reliance on expensive, volatile DRAM and essentially eliminate any need for slow disk-based storage. IBM FlashSystem 900 provides such an industry-leading solution.



Figure 1. IBM FlashSystem 900

Powered by IBM FlashCore™ technology, IBM FlashSystem 900 delivers the extreme performance, enterprise reliability and operational efficiencies required by the Couchbase platform. For storage environments already equipped with robust storage management capabilities such as those provided by Couchbase solutions or by IBM Spectrum Virtualize™ or IBM Spectrum Scale™ software, the ultra-low latency, market-leading performance and ease of implementation offered by IBM FlashSystem 900 make it an ideal choice for transforming business-critical applications into engines of business growth and competitive advantage.

IBM FlashSystem 900 is composed of up to 12 massively parallel IBM MicroLatency® modules that provide extremely high storage density with ultra-low latency in the 100 microsecond range. IBM FlashSystem 900 can scale usable capacity from as low as 2 TB to as much as 57 TB in a single system. The MicroLatency modules also support an offload AES-256 encryption engine, high-speed internal interfaces and full hot-swap and storage capacity scale-out capabilities, enabling enterprises to achieve lower cost per capacity with the same enterprise reliability.

Instead of commodity hardware and solid-state drives (SSDs), IBM FlashSystem 900 leverages the purpose-engineered advantages of IBM FlashCore technology to provide two independent dimensions of data protection. System-level RAID is implemented across all the MicroLatency modules in every IBM FlashSystem array. In addition, patented IBM Variable Stripe RAID™ is implemented within each MicroLatency module. Variable Stripe RAID monitors the health of the flash media, detects and manages flash failures efficiently and optimizes utilization of all flash resources.

## The value of Couchbase and IBM FlashSystem working together

The typical NoSQL implementation model relies on replication of costly, low-density, volatile DRAM and locally attached SSDs for NoSQL data repositories. Instead, the Couchbase ForestDB storage engine takes advantage of IBM FlashSystem 900 as a persistent, ultra-low latency, highly available, all-flash storage pool that can be shared across multiple Couchbase application servers. This architecture offers many benefits, including:

- *IBM MicroLatency.* IBM FlashSystem 900 provides multiple advantages over direct-attached SSDs. Because its flash resources and capacity are consolidated in a single array, performance, latency and system resiliency can all be optimized far beyond the capabilities of distributed SSDs. Also, the ratio of controllers to flash capacity is maximized within IBM FlashSystem 900, eliminating the throughput and performance bottlenecks often encountered by commodity server/SSD architectures. Using IBM FlashSystem 900 instead of hard disk drives for Couchbase persistent back-end storage dramatically increases system performance, as confirmed by IBM test results.
- *Density.* IBM FlashSystem provides up to 57 TB of usable capacity in only 2U as compared to standard server configurations, which are typically under 500 GB of DRAM.
- *Reliability.* IBM FlashSystem uses enterprise-class, two-dimensional flash RAID technology, supporting both Variable Stripe RAID and system-level RAID 5. Variable Stripe RAID maintains system performance and capacity in the event of partial or full flash chip failures, helping reduce downtime and forestall system repairs. System-wide RAID 5 also helps prevent data loss and improves availability. Together, these two dimensions of reliability and data protection are so effective that IBM offers a flash wear guarantee as part of the standard product warranty.

- *Easy integration.* IBM FlashSystem 900 deploys as a simple, very fast, highly resilient, microsecond-latency storage resource. It integrates easily with the Couchbase platform, allows integrating data volumes for each Couchbase server/instance directly from IBM FlashSystem 900, and is simply managed by Couchbase and standard operating systems such as Red Hat Linux.
- *Low cost.* Deploying Couchbase solutions with IBM FlashSystem 900 helps reduce costs, because using much lower-priced flash storage decreases the amount of DRAM required to support high-performance NoSQL implementations. In fact, with IBM FlashSystem 900, the multiple server nodes required to provide enough memory and storage capacity for larger Couchbase databases can be eliminated, thus significantly cutting overall implementation costs.

## Demonstrating value

An IBM test lab conducted benchmark testing using a Couchbase NoSQL database configuration supported by a single IBM FlashSystem 900 array. The primary aim of the exercise was to drive as much input/output (I/O) through the test bed as possible to demonstrate the high-performance capabilities of Couchbase and IBM FlashSystem working together. The test bed consisted of:

- One Lenovo Flex System x440 Compute Node (server blade) with 256 GB memory
- One IBM FlashSystem 900 array
- Red Hat Enterprise Linux v6.7 operating system
- Couchbase ForestDB NoSQL data store
- ForestDB-Benchmark, using APIs generated by Couchbase Server<sup>1</sup>

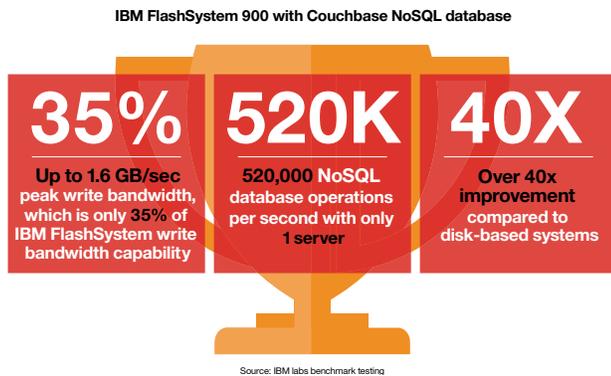


Figure 2. Couchbase test results with IBM FlashSystem 900

As highlighted in Figure 2, the tests achieved 520,000 NoSQL database operations per second, more than 40 times the performance previously seen from disk-based systems. The peak write bandwidth driven by the benchmark was 1.6 GB per second, which was only around 35 percent of the write bandwidth capability of the IBM FlashSystem 900 array. The performance of the benchmark was bottlenecked by the CPU capacity on the single server blade, where CPU utilization peaked at around 89 percent.

If the benchmark had been scaled out to more server blades, a much higher rate of NoSQL database operations could have been achieved, even if those server blades had been attached to the same IBM FlashSystem 900 array. Even direct-attached SSDs cannot match this performance because of bottlenecks in the server and SSD controllers. An enterprise would need to

buy additional servers just to alleviate the controller bottleneck in order to get remotely close to the half million operations per second achieved by the benchmark.

In the case of a Couchbase cluster—the standard configuration—with each server in the cluster directly attached to the same IBM FlashSystem 900 array, there would be plenty of storage capacity available in that one array. Plus the cluster would be so fast that each of its servers could function well with less DRAM. Such a configuration would offer much less complexity, lower cost, higher density and much higher performance than a configuration using disks or SSDs.

### Real-world value

Many different real-world use cases could benefit from this level of NoSQL performance. For example, massive online gaming applications with thousands of concurrent users could easily generate a half million Couchbase operations per second and greatly benefit from the levels of performance offered by IBM FlashSystem 900. Such an application would be characterized by multiple database indexes seeing a high mutation rate under a very heavy workload. The application would utilize this particular configuration to ensure that the concurrent online users would all experience the best possible customer service during peak traffic hours. Online retail enterprises experience a similar scenario on peak shopping days.

Every online game also needs to maintain a dashboard that records user activities, user rewards obtained, battle history, win rate, special equipment and more. The dashboard must be updated frequently to ensure the best user experience and high customer satisfaction. In order to achieve this, online gaming companies and similar enterprises would configure their NoSQL database indexes and back-end data storage in much the same way as in the IBM benchmark tests.

The Couchbase configuration with IBM FlashSystem 900 is used in real-world scenarios to ensure that applications experiencing extremely heavy types of workloads can still function during cache failures. This is very important in a highly distributed computing environment. When thousands, perhaps even millions, of users are accessing an application online simultaneously, the database persistent store is the authoritative store of information for cache. During any cache failure, all data requests from the application will instead go directly to the persistent store. The configuration used for the IBM lab tests ensures that the persistent store can handle all data requests during cache failures, avoiding unacceptable user experiences and low customer satisfaction.

Thus, the IBM benchmark testing confirmed that Couchbase NoSQL environments powered by IBM FlashSystem 900 can fully support the heavy write workloads that a wide range of 21<sup>st</sup> century real-world applications demand—from online gaming and high-velocity retail to financial transactions and the IoT.

## **Accelerating the future of business**

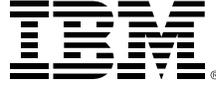
Enterprises as diverse as large research institutions, energy utilities and online dating sites have been turning to NoSQL technology to handle the rapid-fire, high-volume, variable application I/O from their systems of engagement. But IT infrastructure supporting these applications has, until now, been expensive, unreliable and inefficient. Couchbase technologies combined with IBM FlashSystem storage changes the equation for NoSQL systems, providing a framework for cognitive enterprises to effectively integrate the benefits of NoSQL data management into their business models, lowering costs while enabling ever higher levels of innovation.

## **For more information**

To learn more about IBM FlashSystem 900, please contact your IBM representative or IBM Business Partner, or visit: [ibm.com/systems/storage/flash/900/](http://ibm.com/systems/storage/flash/900/)

To learn more about Couchbase solutions, visit: [couchbase.com](http://couchbase.com)

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: [ibm.com/financing](http://ibm.com/financing)



---

© Copyright IBM Corporation 2016

IBM Systems  
Route 100  
Somers, NY 10589

Produced in the United States of America  
February 2016

IBM, the IBM logo, ibm.com, IBM FlashSystem, IBM FlashCore, MicroLatency, IBM Spectrum Virtualize, IBM Spectrum Scale, and Variable Stripe RAID are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided. Couchbase Server, Couchbase Lite, Couchbase Sync Gateway, and Couchbase ForestDB are not IBM products or offerings. Couchbase Server, Couchbase Lite, Couchbase Sync Gateway, and Couchbase ForestDB are sold or licensed, as the case may be, to users under Couchbase’s terms and conditions, which are provided with the products or offerings. Availability, and any and all warranties, services and support for Couchbase Server, Couchbase Lite, Couchbase Sync Gateway, and Couchbase ForestDB is the direct responsibility of, and is provided directly to users by, Couchbase.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

<sup>1</sup>“Benchmark program for ForestDB, Couchstore, LevelDB, RocksDB, and WiredTiger,” *GitHub, Inc.*, Accessed January 13, 2016.  
<https://github.com/couchbaselabs/ForestDB-Benchmark>



Please Recycle