

# IBM LinuxONE for Open-Source Databases

The most scalable, reliable and securable hardware  
and the most popular open-source database  
management systems



for



and



# Based on the DB-Engines ranking for January 2018:

- Relational still rules
- 3 of the top 5, and 6 of the top 10 most popular DBMSs are open-source

All of the top 10 DBMSs *that run on Linux*<sup>®</sup> can run on

**IBM LinuxONE**<sup>™</sup>

## [MongoDB - A Catchy Name To Go With A Catchy Technology](#)

12 January 2018, Seeking Alpha

*As the chief database architect for Synthetic Genomics, I was utterly amazed at the power, flexibility and richness of programming interfaces for Postgres.*

**Mike Sofen**, [Synthetic Genomics](#)

Rank			DBMS	Database Model	Score		
Jan 2018	Dec 2017	Jan 2017			Jan 2018	Dec 2017	Jan 2017
1.	1.	1.	Oracle +	Relational DBMS	1341.94	+0.40	-74.78
2.	2.	2.	MySQL +	Relational DBMS	1299.71	-18.36	-66.58
3.	3.	3.	Microsoft SQL Server +	Relational DBMS	1148.07	-24.42	-72.89
4.	4.	↑ 5.	PostgreSQL +	Relational DBMS	386.18	+0.75	+55.81
5.	5.	↓ 4.	MongoDB +	Document store	330.95	+0.18	-0.96
6.	6.	6.	DB2 +	Relational DBMS	190.28	+0.70	+7.78
7.	7.	↑ 8.	Microsoft Access	Relational DBMS	126.70	+0.82	-0.75
8.	↑ 9.	↓ 7.	Cassandra +	Wide column store	123.88	+0.67	-12.57
9.	↓ 8.	9.	Redis +	Key-value store	123.14	-0.10	+4.44
10.	10.	↑ 11.	Elasticsearch +	Search engine	122.55	+2.77	+16.38

## [PostgreSQL is the DBMS of the Year 2017](#)

2 January 2018, Paul Andlinger, Matthias Gelbmann

## [MongoDB \(MDB\) Stock Rating Upgraded by Zacks Investment Research](#)

23 January 2018, The Ledger Gazette

# Use cases



## Single View

Aggregate structured and unstructured data from disparate data sources, to provide a unified, 360-degree view of enterprise information.

## Internet of Things

Collect and understand monitoring data from devices (e.g., in-store beacons, biometrics), to help users respond quickly (e.g., to market conditions or medical emergencies).

## Combined Analytics

Combine System-of-Record (SOR) data with geospatial and sentiment analyses on news and social media, to achieve deep business insights in real time.

## Mobile

Vertically scale MongoDB to serve millions of users, all within one resilient, high-performance IBM LinuxONE server.

# LinuxONE: Designed for Enterprise Data Serving



**Vertical scale data serving:** up to 320 pairs of dedicated I/O coprocessors, up to 32 TB Redundant Array of Independent Memory (RAIM), the largest cache in the industry by far, allows for larger VMs or containers for vertical scale

**Mixed workloads:** Co-located analytics and system of record opens opportunities for improved insights and better engagement through in-transaction personalization

**Datacenter simplification:** Higher density of VMs reduces the datacenter footprint and administrative overhead, maximizes efficiency

**Trusted operations:** 99.999% availability at near 100% utilization, highest level of security certification, EAL5+

# Use cases



PostgreSQL

## Shift to Open

PostgreSQL uses a stored procedure language (PL/pgSQL) most closely resembling Oracle's PL/SQL, making it a natural choice for Oracle database migration, resulting in savings of 80% or more.

## Co-location

Keep business-critical data in Oracle, but re-host to LinuxONE and migrate the 10's to 100's of DBs on the periphery to PostgreSQL on the same LinuxONE. Shift IT budget away from software licenses and toward innovation.

## Combined DBMS

PostgreSQL is not only an advanced RDBMS, it is also capable of managing NoSQL data. Simplify by using Postgres for both SQL and NoSQL, and running it all on LinuxONE – a single DBMS on a single server, without compromising performance availability or security.

## DBaaS

The cost model for PostgreSQL, combined with its maturity and enterprise readiness, make it an ideal DBMS to make available as a cloud self-service. LinuxONE has a reference architecture for DBaaS, and will soon have one for HyperSecure DBaaS.

# LinuxONE: Designed for Enterprise Data Serving



## High-performance data serving

RDBMSs are typically I/O bound. LinuxONE has massive I/O throughput with up to 24 dedicated I/O offload cores as well as 80X more I/O channel processors than typical x86 servers.

## Better data consistency and reduced overhead

IBM LinuxONE allows PostgreSQL to scale vertically with dynamically allocated resources, instead of horizontally by sharding and replicating the database.

## Lower Total Cost of Ownership

Scaling up means less complexity, less developer time, less administration, less power consumption and cooling, and lower software licensing costs.

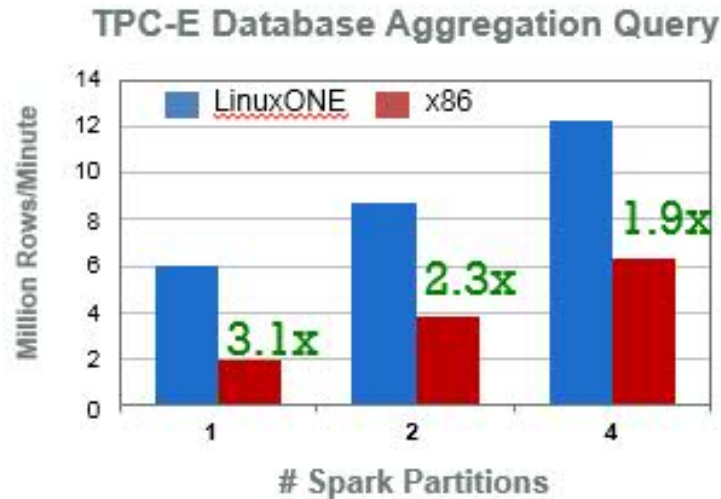
## Uncompromised Enhanced Security

LinuxONE achieves EAL5+ security certification, with no compromised performance, using in-core cryptographic accelerators and optional high-speed PCIe cryptography cards. Its HSM is certified at FIPS 140-2 Level 4.

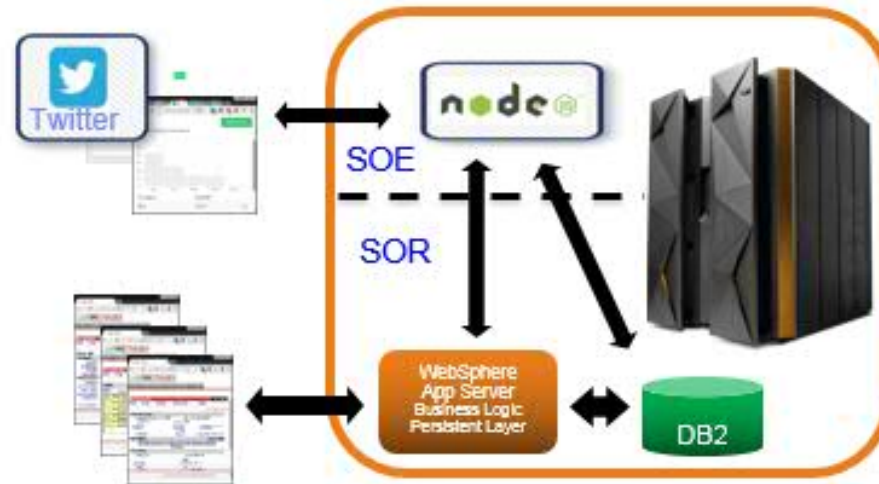
# Faster Data means Better Analytics

## Co-locate for Right-Time insights

- IBM **HiperSockets**™ technology eliminates the network latency for collocated VMs
- Up to 3x more analytics performed with Spark – better insights in less time!
- 60% faster response time by co-locating Node.js with the data
- No ETL needed – critical enterprise data never have to leave the LinuxONE



Apache Spark co-located on LinuxONE drove **up to 3x** faster than Spark running off- platform on x86



Co-locating Node.js on LinuxONE vs. x86 results in **60% Faster** Response Time **2.5x better** Throughput

# When to think of LinuxONE

**You want the option of deployment of on-premises and/or in the cloud** that provides high availability, data consistency, and qualities of service

**You have high operational costs** to manage a large distributed cluster with multiple points of failures when data grows exponentially

**You are concerned about potential security exposures** when enterprise data is sent through the network across distributed clusters

**You have run-time overhead** in aggregate queries caused by collecting results from multiple shards via high-latency network links, and shard-balancing due to data migration between shards

**Weak consistency and durability guarantees** when updating conflicts between shards creating “split brain” situations

**Your infrastructure lacks scale up capability** to complement scale out, which diminishes the ability to handle demand peaks with large volumes of data



# Summary

IBM LinuxONE is a highly engineered system for secure data serving:

- Huge processor cache
- Large memory
- Massive I/O capability
- Highly available
- Most secure platform on the market
- Able to run multiple and mixed workloads (VMs and/or Docker containers) without compromise

MongoDB is the most popular NoSQL database manager, and on LinuxONE it can have individual instances as large as 17TB, and take advantage of LinuxONE pervasive encryption

PostgreSQL is the DB-Engines 2017 DBMS of the Year, and is an enterprise class relational database worthy of business critical workloads. With LinuxONE's massive I/O capabilities and its huge cache, Postgres has almost no bounds.

On LinuxONE, both MongoDB and PostgreSQL can be deployed on native Linux logical partitions (LPARs), as virtual machines, in Docker containers, or in any combination. Later in 2018 you can deploy them in LinuxONE Secure Service Containers.

LinuxONE provides accelerated data encryption and compression, 50 years meantime before failure (MTBF), and meets the highest levels of security certifications.

# Let's get started!

1. Watch the [YouTube video](#), then try a database on the [LinuxONE Community Cloud](#)
2. Schedule a half-day workshop to determine scope, requirements, success criteria and time line for your proof of concept
3. Conduct your PoC; review results
4. Schedule full-day workshop to identify and size your target production workload and environment – which LinuxONE model, which database manager(s), timing, IBM Lab Services



# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

Db2*	IBM (logo)*
HiperSockets*	LinuxONE
IBM*	
ibm.com	

\* Registered trademarks of IBM Corporation

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other product and service names might be trademarks of IBM or other companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at [www.ibm.com/systems/support/machine\\_warranties/machine\\_code/aut.html](http://www.ibm.com/systems/support/machine_warranties/machine_code/aut.html) ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.