

Delivering Business Agility with Open Source Solutions on IBM LinuxONE and z Systems



Business agility is the sum of capability and speed

Business agility is the ability of an organization to get to market quickly and effectively in order to solve the business problems it faces. This ability is important to both the success and survival of organizations competing within an environment where innovative and disruptive technologies emerge one after another, driving business transformation by changing and challenging the way people do business.

Agility = Capability + Speed

“Agility is the ability to get to market quickly and effectively to solve the business problems you care about by leveraging best-of-breed capabilities across eco-system, security and management, while benefiting from industry leading scale and performance.”

—Marcel Mitran, Distinguished Engineer, CTO IBM Systems, Software Performance and Linux Ecosystem

For instance, the mobile world and new app economy are driving changes with unprecedented speed and scope—impacting the world’s most powerful economies and developing nations, changing the game for business, industries and society at large, creating new requirements and expectations that have to be addressed.

In a world where consumer expectations continue to grow, the ability to keep pace with new opportunity and to differentiate through more innovation means your IT department is essential to competing in the immediacy of a digital world. The key to success lies in your ability to:

- deliver with lightning-fast response times, all the time;
- provide access any time, anywhere, and from any device, no matter how many users or transactions;
- ensure security for trusted engagements;
- enable personalization through analytics;
- achieve new levels of IT efficiency, agility and responsiveness using a cloud model.

LinuxONE brings to bear combined capability and performance to help drive improved agility. Here are a few examples:

Process more data, get better insights—faster. Discovering business intelligence requires the use of advanced analytics capabilities, such as Apache Hadoop and Apache Spark. Such capabilities are commonly deployed on distributed commodity hardware. However, when these applications are deployed on LinuxONE and z Systems™—a platform that reduces analytics response times by 50 percent compared to competing platforms—given a fixed service-level agreement, you can process more data and receive critical insight faster—an ability that translates into better business agility.

Reduce risk, increase flexibility with diagonal scaling. Database partitioning or “sharding” is a common technique for scaling out a database that has become too large to fit within a single server. But sharding is complex in practice, it carries risks such as higher latency for aggregate queries and a lower level of data consistency, and the size of each shard is limited to the size of the servers. LinuxONE and z Systems allows you to scale out and scale up—increasing the amount of resources available to each shard

server—in an approach termed “diagonal scaling.”¹ Diagonal scaling on z Systems allows you to adapt to changing workloads with industry-leading performance, increased flexibility and reduced risks, and therefore offers better agility.

High-performance compression and security. A third example of improved agility on LinuxONE and z Systems is high-performance secure logging for auditing in the face of growing regulatory requirements. The IBM zEnterprise Data Compression (zEDC) facility allows z Systems to offload main processors while speeding up compression by up to 7.5x. Linux on z Systems offers protected-key function which keeps the encryption key out of main memory and storage, keeping data-at-rest secure, while offering 28x better performance than secure-key function. When combining the speed and capability of zEDC and protected-key function, Linux on IBM z offers unmatched agility for quickly, non-intrusively and securely logging snap-shots of system state such as Docker instances or Apache Spark Resilient Distributed Datasets (RDDs) for auditing purposes.

Open source software programs that run on LinuxONE and z Systems bring a broad set of capabilities and new applications to the platform—new programming languages and run-time environments, relational and NoSQL databases, as well as efficient big data analytics and container technologies. When these capabilities are combined with unique z Systems performance attributes—the ability to run 141 z13 cores at 5.0 GHz and 10 TB of memory in one system, industry-leading I/O bandwidth and data processing throughput, EAL5+ security, and

high-speed data compression and cryptography hardware—your business can obtain better insights, faster response times, and ultimately, improve business agility.

Linux without limits

IBM has officially supported Linux on z Systems since 2000. Red Hat and SuSE ship enterprise Linux distributions for z Systems, and the Debian and Fedora distributions also support the architecture. Throughout its 15-plus year of history, customer usage has advanced to take advantage of the platform’s technical evolution, with Linux representing 27 percent of the current total number of MIPS on z Systems. The platform is mature and production-ready, and is able to run most if not all of the traditional applications you would expect on any Linux server. In fact, most open source products run right out of the box!

Linux is Linux on any architecture. The IBM Linux Technology Center (LTC), which is dedicated to enabling Linux adoption on IBM platforms, has been contributing continuously to core technologies such as the Linux kernel, glibc and GCC over the years. Overall, architecture-specific code constitutes only a small portion of LTC contributions. For example, less than 2 percent of the kernel is z Systems-specific, with most of the 2 percent related to device drivers. Similarly, platform-specific code makes up about 0.5 percent and 0.3 percent of glibc and GCC, respectively. The majority of these platform details are transparent to application developers and end users.

You do not need to learn any mainframe command to log on to a guest with SSH and be productive in a familiar Bash environment. Graphical user interfaces with remote connection capability, such as X11 and Xvnc, are also available if you prefer a graphical desktop environment. All the essential development tools are available for Linux on z Systems. If some user-space applications are not readily available from distributions, most of them merely need a simple re-compilation to be able to run on the mainframe. Applications written for modern run-time systems, such as Java, Node.js, and PHP applications, simply run out of the box. Because of the similarity and compatibility with other enterprise Linux platforms, you can migrate whole applications stacks from other architectures to z Systems simply and easily, often with just a small amount of work, in hours or days rather than weeks or months.

Attributes of IBM LinuxONE and z Systems

IBM LinuxONE and z Systems platforms are well-known for their reliability and availability. The mainframes ship with resilient and redundant hardware, which ensures continuous operations even in case of hardware faults. The machines boast high MTBFs (mean times between failures), and support hot-swapping of hardware, so they do not typically need to be taken offline for service. The use of IBM GDPS (Geographically Dispersed Parallel Sysplex) helps automate data replication and speeds up recovery from planned or unplanned outages.

IBM z Systems are also equipped with some of the fastest general-purpose processors in the world, such as the recently released IBM z13 (z13) processor. Combined with its large and efficient memory cache hierarchy, the z13 is able to achieve high single-thread performance, which is ideal for scaling up applications to handle the massive volume of transactions that is typical in today's enterprise organizations.

The standard virtualization technologies in the mainframe have been EAL5+ certified—the highest security level and isolation guarantees in the industry. These technologies offer low overhead and higher virtual machine (VM) density compared to other platforms, thanks to the advanced resource over-commit technology in the mainframe. As a result, applications that run on a mainframe can scale out to hundreds or even thousands of co-located virtual machines, as the mainframe essentially becomes a “data center in a box.” What's more, they make the mainframe ready for cloud applications as they support rapid provisioning, multi-tenancy, and capacity scaling on demand.

The co-location of virtualized Linux guests brings significant benefits to applications that run on the mainframe. Using IBM HiperSockets, applications can transfer data from one VM to another in-memory instead of sending the data over the network; the transfer is more reliable (there is no physical connection to lose), more efficient (there is no network latency), and more secure (there is no wire to tap). Co-location is especially important where System of Engagement applications (e.g. mobile, cloud) frequently need to access System of Record data.

The IBM z/Architecture includes hardware support for cryptography, known as Central Processor Assist for Cryptographic Functions (CPACF). In the presence of a suitable cryptographic co-processor, CPACF supports the use of protected keys, which balances the speed of in-memory clear keys and the high security of hardware-backed, tamper-proof secure keys. Linux device driver support for the cryptographic co-processor has existed since 2012, and the OpenSSL and openCryptoki libraries are able to exploit these z Systems features to speed up clear key, protected key and secure key cryptography.²

The IBM zEDC Express adapter is another feature that distinguishes the mainframe as a data-processing powerhouse. It allows applications to off-load zLib-compatible compression work to a hardware co-processor, achieving a good compression ratio without consuming CPU cycles. For databases that employ on-the-fly compression, exploitation of the zEDC feature in Linux is expected³ to improve performance by a factor of five. Users will be able to process more data in the same amount of time, and also save on storage costs.

The world's leading businesses choose to run mission-critical applications and drive billions of business transactions every day on the mainframe. LinuxONE and z Systems leverage the performance, reliability, security and virtualization capabilities of the mainframe to provide an enterprise-grade Linux platform for running databases (e.g. DB2 and Oracle) and middleware (e.g. WebSphere). The advent of mobile and cloud applications means that IBM z Systems are working harder than ever to drive data transactions on a global scale. As innovations in open source technology gain traction in the market, demand for emerging technologies on the mainframe has increased, creating huge opportunities for open source developers and vendors on this platform.

Expanding the open source ecosystem on LinuxONE and z Systems

IBM has a long history of involvement in open source software development. IBM contributes to many key open source technologies such as the Linux kernel, the Eclipse project, and many Apache projects, including the most recent, Apache Spark. IBM is a member of many open standard organizations and software governance consortia that help shape the future of open source software.⁴

IBM is investing heavily to create a rich open source ecosystem to enable LinuxONE and z System as the premier platform for new application deployment that employs emerging

technologies such as Node.js, MongoDB, PostgreSQL, MariaDB, Docker, Chef, Puppet and Apache Spark. IBM is listening to clients, working with business partners and ISVs, and engaging the open source development community, to bring the most important and most sought-after foundational open source technologies to LinuxONE and z Systems. To enable and encourage software developers to work on the platform, an emphasis is placed on programming languages and run-time technologies. Additional focus is placed on open source applications that can benefit from the strengths of LinuxONE and z Systems and bring business values to users, including database management systems and cloud infrastructure, as well as analytics and big data products. Table 1 lists some of the key open source technologies that have been enabled or confirmed working on the platform; more than 40 packages have been enabled and/or validated so far, and the list continues to grow, as many other packages are currently in the development or validation pipeline.⁵

Table 1: Partial list of popular open source technologies available on Linux on IBM z Systems

Erlang	Apache Cassandra	Chef
GCC Go	Apache CouchDB	Docker
Java (OpenJDK and IBM JDK)	Apache Geode	Docker Compose
LLVM	MariaDB	Docker Swarm
Node.js and V8	MongoDB	OpenStack
OCaml	PostgreSQL	Puppet
PHP	RabbitMQ	
Python	Redis	
Ruby and Rails		

Modern programming languages provide more choice and speed on the platform

Application developers can choose from a variety of modern, popular programming languages and run-time environments, including Java, Node.js (JavaScript), Python, Ruby-on-Rails, Scala, Erlang, and Go—and the list is growing. This variety allows modern applications to be built and run on LinuxONE and z Systems easily.

The most notable of these run-time technologies is Node.js. It is the fastest growing ecosystem for application development. Since 2012 Node.js has grown at a rate three times faster than Java with the number of Node.js modules contributed by developers world-wide surpassing other languages.⁶ As a high-performance, highly scalable, event-driven, server-side JavaScript solution, Node.js has become popular for web application development—the “N” in the term “MEAN stack” stands for “Node.js.”

To meet the high demand for Node.js functionality on z Systems, IBM has ported the open source Node.js code to the platform. IBM released IBM SDK for Node.js 1.2 in early 2015. This SDK is fully compatible with Node.js version 0.12.⁷ The SDK enables developers to write Node.js code on any platform, and test and deploy their applications on Linux on z Systems. Enterprise users of the SDK will enjoy the performance and security afforded by the platform (on the AcmeAir benchmark,⁸ Linux on the z13 out-performs a comparable distributed server with 2.1 times higher throughput in RESTful transactions), as well as the additional monitoring and debugging tools that ship with the SDK. One of the most important benefits from

deploying Node.js on z Systems is the ability to bring the processing of web requests to the platform where the data is hosted (co-location of application and data), which can improve application throughput by two times, and reduce transaction response times by 60 percent. The enablement of Node.js on the platform also means that a large number of open source web application frameworks, such as Express and Sails.js, are now available to mainframe developers.

Node.js relies on the V8 JavaScript just-in-time compiler. For this reason, IBM has also ported V8 to IBM z Systems. The availability of V8 on z Systems has allowed a number of other open source software that requires JavaScript functionality to run on the platform. IBM will be contributing the z Systems port to the V8 project and working closely with the community to further improve the performance of V8 on IBM z Systems.

SQL to NoSQL databases benefit from LinuxONE and z Systems

Open source relational database management systems (RDBMSs) such as PostgreSQL serve a great amount of data in many organizations. First released in 1996, PostgreSQL is a well-established enterprise-grade database system that powers many web sites and government agencies.⁹ PostgreSQL 9.4 already runs well on z Systems, and it is able to capitalize on the strengths of the platform as a data-serving platform. According to a report by the PostgreSQL consultancy firm 2ndQuadrant,¹⁰ PostgreSQL can offer better per-core performance (1.6 to 2.2 times better on the pgbench benchmark) and scalability on the mainframe compared to distributed systems, as shown in Figure 1.

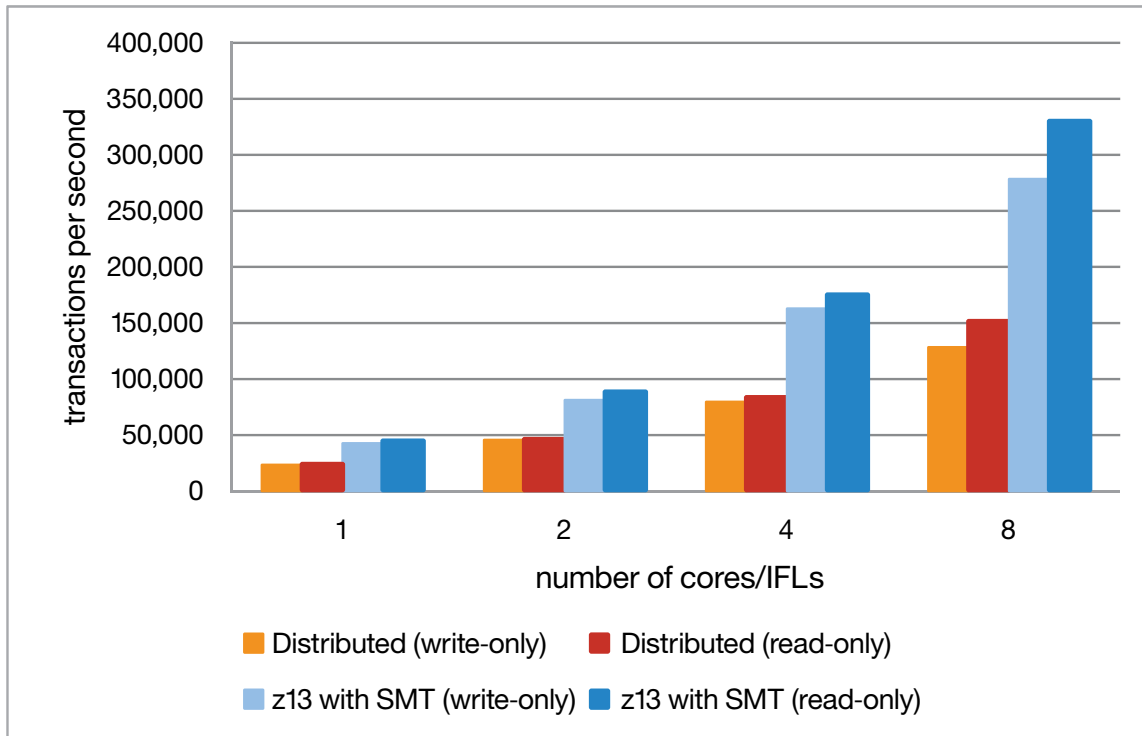


Figure 1. Pgbench results (read-only and write-only, in-memory workloads) on IBM z13 vs. distributed systems

Another notable RDBMS is MariaDB, which is being adopted by many Linux distributions as the replacement for MySQL. It has also shown a substantial throughput advantage on z Systems (1.8 to 2.1 times better than distributed systems) on the Sysbench OLTP benchmark, as shown in Figure 2.

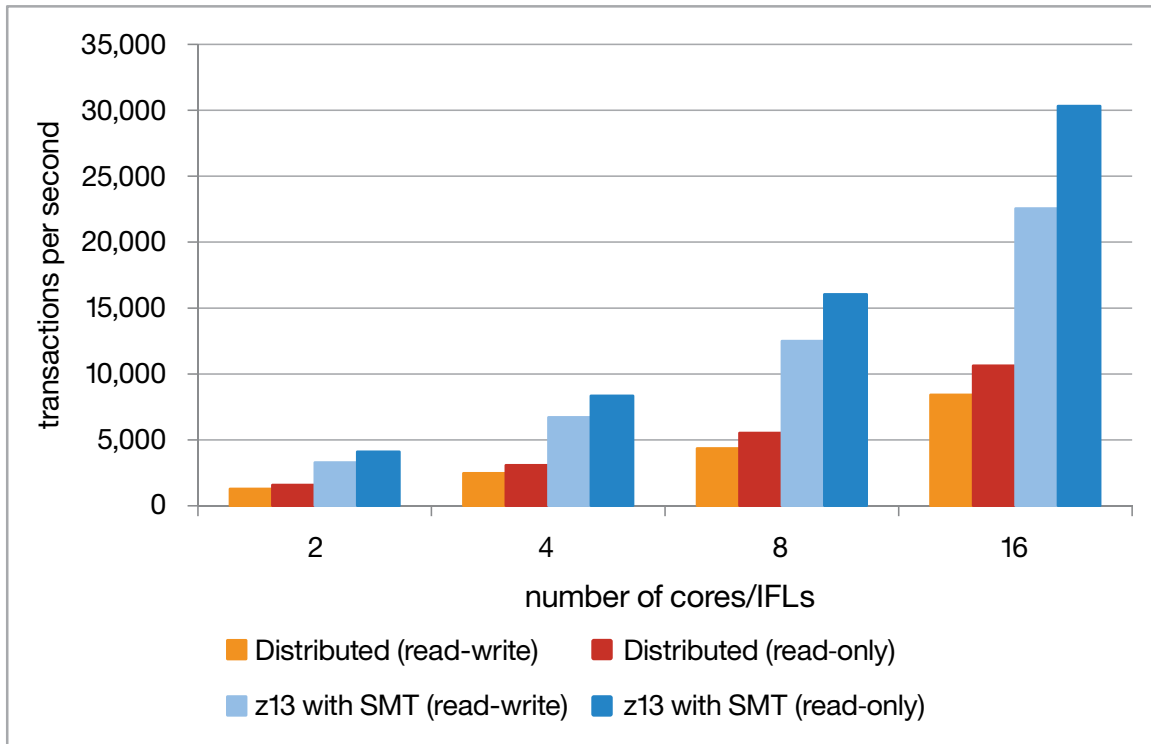


Figure 2. MariaDB OLTP results (read/write and read-only, in-memory workloads) on IBM z13 vs. distributed systems

Besides relational databases, NoSQL databases are increasingly being adopted in big data and real-time analytics applications. Some of the popular NoSQL databases include MongoDB (the “M” in “MEAN stack”), Apache Cassandra, Redis, Apache CouchDB, and Apache Geode. IBM has shown that all these databases run well on Linux on z Systems, and that a healthy variety of choices exist for NoSQL databases on the mainframe.

In the case of MongoDB, some porting effort was required because it uses little-endian byte order pervasively, and the mainframe is a big-endian architecture.¹¹ Fortunately the byte order issue is well understood and eliminating unnecessary assumptions of little-endian byte order in the code is often straightforward. Users do not need to care about the native byte order of most programs, and it is easy to write portable code that can function correctly on either type of architecture.¹² In the end only 0.14 percent of the code had to be modified to enable MongoDB to run on Linux on z Systems.

The superior single-thread performance and large memory capacity of the IBM z13 allows MongoDB to scale up very well. On various YCSB (Yahoo Cloud Service Benchmark) workloads, MongoDB running on Linux on the z13 is able to achieve 1.9 to 2.2 times better throughput than the latest distributed platforms, as Figure 3 shows. Because a z13 LPAR can support a maximum of 8TB of memory, a single Linux system can accommodate a larger database that would have required sharding to multiple MongoDB servers. This capability affords users the flexibility to reduce or avoid the risks and costs associated with sharding.

MongoDB achieves up to 2x better YCSB throughput on IBM z13 over alternative platforms

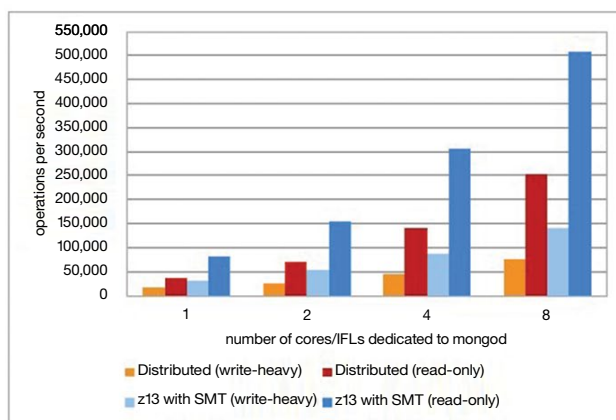


Figure 3. MongoDB YCSB (update-heavy) throughput on IBM z13 vs. distributed systems

Faster analytics equals better business insight

Big data and real-time analytics are key to success to many enterprises. Apache Spark is an open source cluster computing framework that allows user programs to load data into a cluster's memory and query it repeatedly. It is designed to perform both

batch processing and new workloads, such as streaming, interactive queries, and machine learning. IBM has tested with the Spark-Perf benchmark suite on z Systems and demonstrated that the platform can process 1.5 times more data (see Figure 4) than distributed systems for model building, leading to real-time insights with higher accuracy.




Figure 4. Spark-Perf “Spark” benchmark speed-up on z13 vs. distributed systems

Improve DevOps agility with Docker containers

Docker is an open, portable, light-weight run-time and packaging tool for Linux containers. Docker has gained a lot of momentum since its first release in 2013. Linux containers provide operating system-level virtualization for running multiple isolated applications on the same host. Using containers, it is easy to build and ship complex applications that have dependencies on a deep software stack or other infrastructure, without having to worry about interference from other applications that

might have conflicting requirements. Compared to a standard VM, containers are much faster to boot, more efficient to run, and offer higher application density.

IBM has built Docker on Linux on z Systems.¹³ It is straightforward to set up a base image and create Docker containers based on the image. Advanced application composition and container clustering on the platform are possible with the use of Docker Compose and Docker Swarm. In addition to the main Docker tools, IBM is working to provide Dockerfiles that allow users to create containerized applications of their choice to run on the mainframe.

The virtualization technology in z Systems and LinuxONE complements Docker by offering flexibility in the level of security isolation. For example, while developing or testing an application, you can obtain increased application density by deploying containers directly on one logical partition (LPAR), which eliminates the need for VM provisioning. On the other hand, when deploying the application in production, you can put the containers inside second-level guests to get better isolation. For maximum security, you can run individual containerized applications inside separate LPARs, which give isolation on a bare-metal level unmatched by other platforms. See Figure 5 for an illustration of how containers can run on LPARs as well as VMs. In these capacities, using Docker on IBM z Systems can help to solve real DevOps problems for enterprise applications.¹⁴

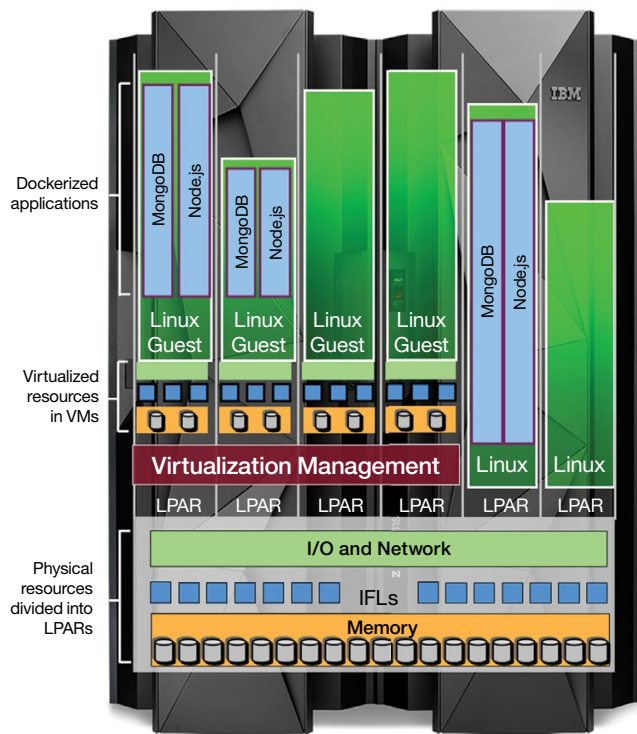


Figure 5. Different levels of isolation—LPARs, VMs and Docker containers

How to gain access to IBM LinuxONE and z Systems

IBM is committed to working with open source development communities to enable them to develop, test, and continuously integrate code on IBM LinuxONE and z Systems. IBM offers access to LinuxONE and z Systems hardware in a number of ways.

- IBM PartnerWorld members can sign up for the IBM Systems Application Advantage (also known as “Chiphopper”) program, which helps partners port and support existing Linux applications on to IBM z Systems, for a limited time, at no cost. For more information, visit https://www.ibm.com/partnerworld/wps/servlet/ContentHandler/isv_com_dvm_techval_chiphopper.
- Open source developers can get access through the Community Development System for Linux service. The access is for a limited time, but renewals are possible depending on the situation. Visit <http://www.ibm.com/systems/z/os/linux/support/community.html> for details.
- IBM is working with partnering universities on Open Access Mainframe Community Cloud, which will provide LinuxONE and z Systems resources to research projects and open source development communities.

IBM LinuxONE and z Systems: business agility through capability and speed

Open source software that run on LinuxONE and z Systems have brought a broad set of new capabilities to the platform. When these capabilities are combined with unique z Systems performance features, your business will obtain better insights, faster response times, and ultimately, improved business agility. As the open source software ecosystem grows on z Systems, and mainframe hardware becomes more accessible to open source developers, the platform is more suitable than ever for enterprise Linux application development.

For more information

To learn more about Delivering Business Agility with Open Source Solutions and Linux on z Systems, please contact your IBM representative or IBM Business Partner, or visit the following website: ibm.com/developerworks/community/groups/community/lozopensource



© Copyright IBM Corporation 2015

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
August 2015

IBM, the IBM logo, ibm.com, and z Systems 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

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. A claims author certification form is being completed for each of the claims mentioned in the paper. 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.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation. Statements regarding IBM’s future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

- <http://highscalability.com/strategy-diagonal-scaling-dont-forget-scale-out-and>
- https://www-01.ibm.com/support/knowledgecenter/linuxonibm/com.ibm.linux.z.xci/lxci_c_geninfo.html
- Linux kernel support for the zEDC Express adapter will be available in the near future.
- IBM is a member or sponsor of: The Linux Foundation, The OpenStack Foundation, The Node.js Foundation, The OpenJDK Governing Board, OASIS, and The Apache Software Foundation, among others.



Please Recycle

⁵ The central list of ported open source software can be found at our developerWorks community: <https://www.ibm.com/developerworks/community/groups/community/lozopensource>. Feel free to raise requests in the community forum for open source software that are not yet available on the platform.

⁶ <http://www.modulecounts.com/>

⁷ <http://www.ibm.com/developerworks/web/nodesdk/>

⁸ <https://github.com/acmeair/acmeair-nodejs>

⁹ <http://www.postgresql.org/about/users/>

¹⁰ Mark Wong (2ndQuadrant Ltd.), *OLTP Performance Benchmark of PostgreSQL 9.4 on IBM z Systems*, <http://2ndquadrant.com/en/support/support-ibm-z-systems/performance-analysis/>

¹¹ To try out the z Systems port of MongoDB, see <https://github.com/linux-on-ibm-z/docs/wiki/Building-MongoDB>.

¹² Rob Pike, *The Byte Order Fallacy*, <http://commandcenter.blogspot.ca/2012/04/byte-order-fallacy.html>

¹³ <https://www.ibm.com/developerworks/linux/linux390/docker.html>

¹⁴ <https://www.youtube.com/watch?t=705&v=MHJmNZSRve0>