



IBM LinuxONE

**A portfolio of hardware,
software, and services for an
enterprise-grade Linux operating
environment**

Marcel Mitran – DE, CTO IBM LinuxONE
mmitran@ca.ibm.com

IBM LinuxONE™





Enterprise-grade Linux



The best of
ENTERPRISE COMPUTING

The best of
LINUX & OPEN

- Dynamic Resource Allocation
- Non-Disruptive Scalability
- Continuous Business Availability
- Operational Efficiency
- Trusted Security
- Data and Transaction Serving

- Freedom & Agility
- Standards Based
- Speed of Innovation
- Developer Productivity
- Community Collaboration
- Quality of Software

2
IBM LinuxONE™

© 2015 IBM Corporation

When you think of Linux you think of open collaboration, fast paced innovation, agility and capability that empowers developers and operators.

When you think of IBM Big Iron, you think of the qualities of service that have defined 5 decades of un-matched resilience, security, dynamic resource allocation, industry leading data serving and scalable transaction processing.

IBM LinuxONE Announcements - August 2015

IBM LINUXONE SYSTEMS

IBM LinuxONE Emperor™

IBM LinuxONE Rockhopper™

IBM LINUXONE SOLUTIONS

Mobile Analytics Cloud DevOps

IBM LinuxONE Elastic Pricing

Open Source & ISV Ecosystem

Distributions	Hypervisors	Languages	Runtimes	Management	Database	Analytics
Supported Versions: ubuntu, SUSE, redhat	KVM, LPAR	python, Ruby, php, Scala, Clojure, JS, OCaml, Java	node.js, JRE, ZF, OpenJDK	docker, CHEF, puppet, Juju, openstack, vmware vRealize, Cloud Manager, urban(code)	MariaDB, PostgreSQL, mongoDB, Oracle, DB2, IBM InfoSphere Analytics	Spark, Hadoop, elasticsearch, Solr, IBM InfoSphere Analytics
Community Versions: debian, SUSE, CentOS						

IBM LinuxONE Community Cloud

Open Mainframe Project

Open Source Contribution

Linux YOUR WAY

Linux WITHOUT LIMITS

Linux WITHOUT RISK

On Aug 2015 we took Linux on z to the next level - we announced IBM LinuxONE. LinuxONE brings together a broad set of new capabilities such as a KVM hypervisor, Ubuntu -- cloud-enabled distribution, and an extreme expansion of the eco-system to include significant set of open source capabilities.

LinuxONE introduces new elastic pricing which represents a new cloud-based pricing model – essentially moving from a capital expenditure to an operational expenditure pricing model.

We introduced the new LinuxONE Community Cloud to enable open communities and open innovation on LinuxONE.

We also worked with the Linux Foundation to start up the Open Mainframe Project which brings together academics, industry players and others to help drive open innovation and collaboration, bringing the best of the mainframe innovation to open communities, while also

bringing mind-share for the mainframe in open communities.

LinuxONE combines the power of open innovation with enterprise qualities of service to take Linux to the next level.



Why Ubuntu OS and Ubuntu OpenStack on LinuxONE?



Providing Significant Time to Value Advantage

Ubuntu on LinuxONE Systems makes it easy to build, model, deploy and manage Enterprise scale out clusters and scalable cloud architectures



4

IBM LinuxONE™

© 2015 IBM Corporation

- Canonical delivers Ubuntu 16.04 support on LinuxONE and IBM z Systems the perfect platform for Hybrid Cloud
- Canonical's Ubuntu provides more flexibility and choice with continued open ecosystem expansion
- Ubuntu allows developers who use Ubuntu on other systems to work in a familiar environment
- Allows CSPs and clients with applications built on Ubuntu to easily extend their environments to include LinuxONE and z Systems
- Lets customers get the performance, scale, reliability and security they need for Linux to run their business, using the tools they know, while getting the same economics they are use to

LinuxONE Elastic Pricing

On Premise Cloud Model

what you pay depends on what you use for your LinuxONE Deployment

Move from CapEx to OpEX

as deployment models change LinuxONE has adapted to the shift

Simple

Simplified monthly or quarterly billing based on your usage

Low barrier to entry

Reduces risk of entry and cost of acquisition by as much as 75%

Metered

Based on real usage data ...





Open Mainframe Project

Create an open source, technical community that industry and community participants may easily participate in and so that they may contribute to the creation of assets and materials that will benefit the development of Enterprise Grade characteristics of Linux such as:

- High Availability and Disaster Recovery
- Reliability, Availability and Serviceability
- Security
- Performance and Scalability

Current Members

MARIST

L3C

Compuware



DATA KINETICS
DATA PERFORMANCE & OPTIMIZATION



University of
Bedfordshire

suse

ca
technologies

East Carolina
UNIVERSITY

VICOM InFINITY

RSM
PARTNERS



ADP

IBM

bmc



Hitachi Data Systems

IBM LinuxONE™

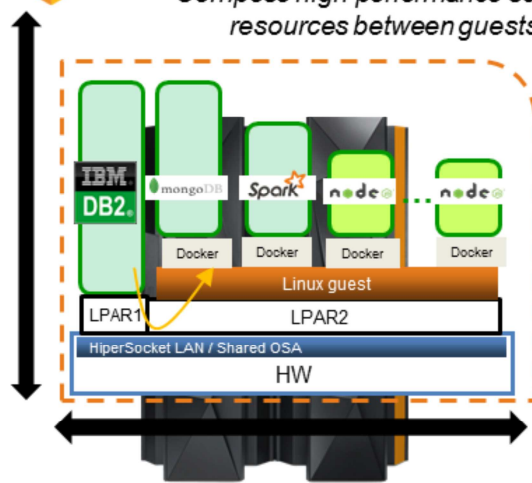
© 2015 IBM Corporation



LinuxONE – Diagonal Scale



Compose high-performance scalable applications. Dynamically and seamlessly re-allocate resources between guests. Provide right-time analytics and powerful engagement



Extreme Virtualization and Scale

- Hypervisor partitioning built into firmware
 - Complete isolation – **EAL5+**
- Supports as many as 85 hypervisor instances – z/VM or KVM
 - **1k** Linux guests/hypervisor
 - **+1 million docker** containers
- Hypervisor communication is via fast, in-memory TCP/IP
 - Hipersockets or Shared-OSA – **3x less latency than discrete servers**
- Massive dedicated I/O – **640 power co-processors**
- **960Meg L4 cache, 5Ghz core, dual-TLBs, crypto acceleration**

Super Elastic System

- Combine horizontal and vertical scaling
- Non-disruptively add/remove resources from Linux guests
- Non-disruptively add/remove Linux guests



Extreme Virtualization with Docker®!



A single LinuxONE Emperor ran more than **1 Million Docker** containers

Workload: busybox httpd server (no NAT)



LinuxONE Emperor runs **2000+ Docker** containers on avg **2.0x** better than a compared Haswell-based system

Workload: Apache Solr

LinuxONE Emperor can host over **10k Docker** containers

Workload: 4k Apache Solr + 6k busybox httpd server (no NAT)

DISCLAIMER:
This chart is based on results from internal lab measurements. Performance results may vary depending on the workload and other factors. Benchmark:
+ Apache Solr search queries driven by Apache Jmeter.
+ LinuxONE Emperor (B21) 1016x CPU on 16 CPU cores with 128GB memory.
+ Haswell-based system is a server (Linux) with 1016x CPU on 16 CPU cores with 128GB memory.
+ Heavy Docker Container Apache Solr v4.10.2 on Haswell Linux v4.10.2-1.
+ Lightweight Docker Container BusyBox.
+ System B21 Docker 1.12.0-rc1 on built storage backend, RAID 10.
NOTE:
+ Each active container is driven by a client thread in Apache Jmeter which keeps sending the same Solr query repeatedly to the container to search documents that contain an arbitrary word in the word list.
+ The Docker runtime was modified to increase a thread count to avoid connection time-out and to expose a Docker instance from a Docker host.
+ A modified Linux 3.10 kernel to support more than 1024 network bridge ports was installed on B21.



The throughput and response-time for a single Linux host running 4096 containers

IBM LinuxONE™

Docker is a registered trademarks of Docker, Inc. in the United States and/or other countries

© 2015 IBM Corporation

Solr is a search engine, used my many web sites to search for terms in a document base

Use case

Greater business agility and lower costs by simplifying IT operations, enabling an organization to go to market quickly and solve business problems effectively.

DevOps agility with Docker containers and microservices.

Hybrid Cloud and Choice Require Portability:

Cross Cloud Deployment – move the same application across multiple clouds

Eliminate “lock-in”, become a “Cloud Broker”

Key Messages

Strong I/O subsystem on z13 enables a server to host a large number of containers that generate heavy file accesses.

Superior I/O performance of z means more efficient paging translating to 1.5x better Docker density within a fixed core and memory footprint.

Benchmark

Apache Solr search queries driven by Apache Jmeter

System Stack

z13: Native LPAR on 16 CPU cores with 128G memory

E5-2699 v3: Native Linux on 16 CPU cores with 128G memory

SW: Apache Solr v4.10.0, WebSphere Liberty v8.5.5.2, IBM Java 1.7.0 SR1, Docker 1.6.2
w/ btrfs storage backend, RHEL 7.1





LinuxONE is designed for large, cache-intensive business workloads



Latency introduced by
(1) hop across multiple
L3 rings, and (2) CPU-to-
L3 frequency differences

Haswell chip
E5 2600 V3
(representation)



Latency is minimized –
all cores have equal
access to L3, no
frequency boundaries

LinuxONE chip
(representation)

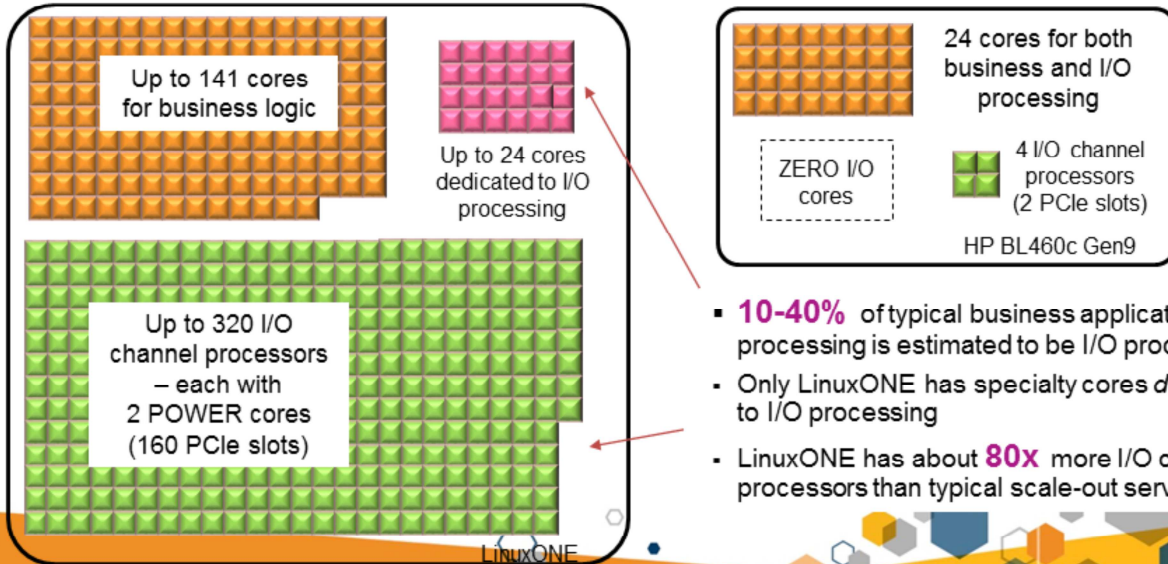
	Total	Per Core	Total	Per Core
L1 cache		64 KB		224 KB
L2 cache		256 KB		4 MB
L3 cache	45 MB	2.5 MB	64 MB	8 MB

LinuxONE also has
960 MB
per drawer off-chip
L4 cache

<http://www.anandtech.com/show/5355/intel-haswell-architecture/10>
<http://www.anandtech.com/show/5353/intel-linuxone-3-sup-to-16-haswell-ep-cores-4>



LinuxONE is designed for high I/O bandwidth business workloads



- **10-40%** of typical business applications processing is estimated to be I/O processing*
- Only LinuxONE has specialty cores *dedicated* to I/O processing
- LinuxONE has about **80x** more I/O channel processors than typical scale-out servers

10

IBM LinuxONE™


LINUXONE

© 2015 IBM Corporation



Database workloads with high I/O bandwidth requirements benefit from LinuxONE architecture

Which platform provides the lowest TCA over 3 years?



Customer Database Workloads
each supporting 18.3K tps
Oracle Enterprise Edition
Oracle Real Application Cluster



3 Oracle RAC clusters
4 server nodes per cluster

12 total HP DL380 servers E5-2699v3
2.3GHz 2ch/36co
(432 cores)

\$29.3M (3 yr. TCA)



3 Oracle RAC clusters
4 nodes per cluster

Each node is a Linux guest
LinuxONE with 61 cores

\$13.5M (3 yr. TCA)

54% Lower cost


TCA includes hardware, software, maintenance, support and subscription.
Workload Equivalence derived from a proof-of-concept study conducted at a large Cooperative Bank.



2x Better Data-Serving with LinuxONE




MariaDB 10.1.5



1.8x to 2.1x throughput improvement on Sysbench Benchmark




MongoDB 3.0.4 (WiredTiger, no sharding)



1.8x to 2.1x throughput improvement on YCSB Benchmark

PostgreSQL 9.4



1.8x to 2.2x throughput improvement on pgBench Benchmark

Cloudant DBaaS



1.8x to 2.1x throughput improvement on YCSB Benchmark





Scaling-up with MongoDB on LinuxONE



Announcing MongoDB 3.2

A GIANT LEAP

[LEARN MORE](#)

Single MongoDB node on LinuxONE scales up to **2TBs** with sustained throughput and **response time <5ms**, while supporting **+4 Billion documents, 460,000 reads+writes/second**, with no sharding required!

"We are committed to make MongoDB available on all major platforms and are excited to add support for IBM LinuxONE Enterprise Grade Linux and LinuxONE Platform. This announcement is a leap forward for customers who want to deploy modern, mission-critical applications built with MongoDB and take advantage of the performance, scalability and security of IBM's LinuxONE platform hardware products."

— Eliot Horowitz CTO & Founder, MongoDB

13

IBM LinuxONE™

© 2015 IBM Corporation

NoSQL databases on LinuxOne performs 2x faster than other platforms. Compression hardware to save Spark Resilient Distributed Datasets runs 4.9x faster than other platforms and persists Docker containers 4x faster. LinuxOne can encrypt all the data 28x faster than other platforms. Node.js runs up to 2x faster. Spark Analytics runs up to 3x faster. Who wants to deal with sharded databases when you have a server that can handle TB databases, with over 2 billion documents and 470,000 read/writes per second. And remember all this speed and scalability is for just one LinuxOne server, you can cluster them to scale out as well to have even more scalability. This is a demo that combined unstructured data (tweets and news) with structured data, using real time analytics to provide insight. This demo was developed by 2 people in 4 weeks. They developed the code on their favorite Linux desktops and laptops and deployed it on LinuxOne for the performance and power. Think what you can do with a LinuxOne server!

On average, AcmeAir transaction translates to 1.36x MongoDB operations (Insert, Query, Update, Delete). 460k r/w per sec was based on a fully saturated z13 Hartmut measured, with many instances of both AcmeAir and Mongo. The 27k for the 1TB system is simply based on 16 AcmeAir instances targeting 1 MongoDB instance (with lots of data).

z13 seamless scale up to 350k transaction per sec leading to a 30B proof-point. Workload scales up ~6x higher (thruput) than an alternative server.

460k number was a the outcome of the work done by Hartmut to load the machine to drive

max throughput/box by having many instances of Node/Mongo running

Encryption 28x faster – but for secure key, not for this demo



Scaling-out with Node.js on LinuxONE



"LinuxONE can scale to **up to 30B** RESTful web interactions/day with Dockerized Node.js and MongoDB, driving **over 470K** database read and writes per second, while maintaining **response-times that are 2x better** than alternative platforms."



Node.js v0.12


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

High Performance JavaScript for LinuxONE

- Highly scalable, event-driven platform with non-blocking I/O
- Thousands of concurrent connections with minimal overhead
- Up to **2.1x** more RESTful web interactions with AcmeAir in node.js



IBM API Connect



14

IBM LinuxONE™

© 2015 IBM Corporation

Node.js scales nicely in Docker containers

z Systems Connectivity

Co-locate Node.js applications for reduced latency accessing z/OS data/services
Up to **2x** better throughput, **60%** faster response time to DB2 on z/OS*

Security and Dependability

Leverages the trusted environments of z Systems to maximize security and uptime of critical Node.js applications.

Unified Diagnostics Tooling with IBM SDKs v1.2 for Java®

Monitor your application with IBM HealthCenter
Debug your application using Interactive Diagnostic Data Explorer

The IBM SDK for Node.js™ provides a stand-alone JavaScript® runtime and server-side JavaScript solution for IBM platforms. It provides a high-performance, highly scalable, event-driven environment with non-blocking I/O that is programmed with the familiar JavaScript programming language. The IBM SDK for Node.js™ is based on the Node.js™ open source project. It provides a compatible solution for IBM System Z products that require Node.js™ functionality and package management.

Node.js has one of the fastest growing eco-systems out there. More than +93k javascript modules can be quickly and easily downloaded from npmjs.org to build complex and sophisticated javascript applications with ease and speed.

According to <http://www.modulecounts.com>, Node.js was the fastest growing eco-system amongst languages – the rate of modules contributed was around 3x that of any other language in August, 2014.

Acme is a different benchmark than the one we measured 63% on in the chart, that one is a web-serving benchmark. Acme is an airline reservation system.

acme airlines is a Node.js benchmark that a (now ex-) IBMer wrote <http://spyker.blogspot.ca/2013/05/announcing-acme-air-performance.html>. We ran it on Intel and zEC12. It was 29% faster on the latter.

AcmeAir Benchmark (<https://github.com/acmeair/acmeair>) shows an implementation of a fictitious airline called "Acme Air". The application was built with the some key business requirements: the ability to scale to billions of web API calls per day, the need to develop and deploy the application in public clouds (as opposed to dedicated pre-allocated infrastructure), and the need to support multiple channels for user interaction (with mobile enablement first and browser/Web 2.0 second).

The AcmeAir benchmark was installed and measured:

Intel Xeon E5-2670 @ 2.60 GHz using the 32-bit Joyent 0.10.30 Node.js
IBM zEnterprise EC12 using the 31-bit IBM SDK for Node.js V1.1

Server side scripting

-Scripting to next level, write MW in script language

-Cloud foundry written in Ruby, so no longer use scripting lang to shift files but use them for more interesting tasks

-For node, used for web serving, web pg serving, data processing

-No differ betw java and java script

-Server side: SW being used to run on a server, so this is why perf is critical cause servers are expensive, so people pay attention to perf, security, so need to build the IT infra, want to do more with less – do transaction processing, data serving

-Lang that allows u to write apps so they run on servers as opposed to old way where scripts to automat something were on the client side. On the server side, say on bank app, goes to a server transaction with some QoS e.g has to be secure, so bank processing billions a day – infra in place with HW, O/S and MW that runs on top of it e.g a resilient database like DB2, IMS, then surface it to the mobile app in a real way – so when build app to run on server, the language is differ cause it biz logic for the app and the infra that runs this biz logic to manage that transaction envr

-Net is its goes from simple things to now more complex things

-Now enable yourself to do this on a server side

-In past scripting didn't have to work all the time but in server envr has to work all the time, has to scale, be fast, be efficient – so scripting used to do something in a different way

The 2x perf is mainly due to the IO capability of the platform

Module counts - # modules in different repositories that cover differ languages, sharing libraries that you built that others can use

Support: we r building assets with intent to support, on a customer by cust baiss, then eventually have a PID in place to have a chargeable service




SOR+SOI+SOE-in-a-Box with LinuxONE

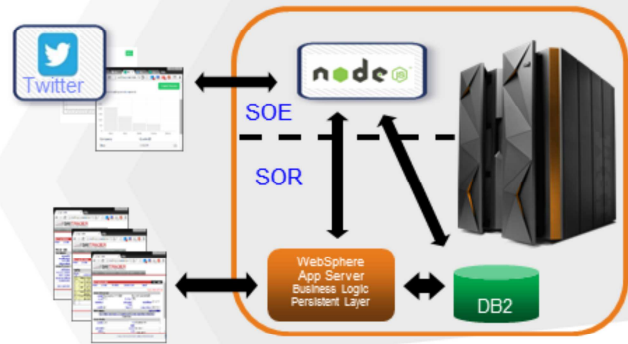


Co-located SOR, SOI, and SOE for right-time insights and powerfull engagement

TPC-E Database Aggregation Query





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




Statistical models show how consolidating workloads drives up CPU utilization



- Service-level-guarantees require provisioning for **peak utilization**
- But costs are inversely proportional to **average utilization**
- Consolidation leads to higher average utilization levels and lower cost/workload



Variable workloads consolidated on LinuxONE have lower TCO



17

IBM LinuxONE™

Performance comparison based on IBM internal tests comparing IBM LinuxONE private cloud with one comparable configured private cloud running an aggregation of light, medium and heavy workloads designed to replicate typical IBM customer workload design in the marketplace. System configurations are based on actual hardware installed within IBM internal testing. These comparison estimates based on a 3-year Total Cost of Ownership (TCO) using publicly available data and prices, including a 10% discount for volume and currency of January 1, 2015. LinuxONE and TCO estimate includes data on infrastructure (system memory, storage, virtualization, OS, cloud management, middleware, power, floor space and labor). Results may vary based on actual workload, system configurations, customer applications, workload and other variables in a production environment and may produce different results. Use of this document should reflect the applicable data for their specific environment.


© 2015 IBM Corporation



A breakdown of TCO shows how software costs soar on the x86 platform



Case Study: 123 Workloads (219 VMs)



For LinuxONE, **78%** fewer cores results in **72%** lower middleware costs

18

IBM LinuxONE™



© 2015 IBM Corporation



In x86 environments, high priority workloads must have dedicated servers...


IBM

...to avoid the effects of 'noisy neighbors' – causes increased proliferation of x86 servers



Virtualized on 3 Intel 40-core servers (120 cores total - Linux)

\$13.7M (3 yr. TCA)



32 cores on LinuxONE

\$5.77M (3 yr. TCA)

58% Lower cost

Consolidation ratios derived from IBM internal studies. LinuxONE numbers derived from measurements on similar server. Workloads running WebSphere 8.6 ND, DS210 AESE, and Monitoring software. Results may vary based on customer workload profiles/characteristics. Prices will vary by country.



LinuxONE is designed to ensure highest availability and lowest downtime



Comprehensive, multi-layered strategy includes...

Error Prevention

- Hardware and firmware designed to protect against outages
- Built-in redundancy eliminates single points of failure
- Extensive testing and failure analysis at every level



Error Detection and Correction

- Error detection embedded in components
- Built-in automated diagnostics; problem determination and isolation
- Non-disruptive installation, upgrades and maintenance avoids outages


Error Recovery

- Automated failover to speed recovery and minimize system impact
- Business continuity and disaster recovery solutions – GDPS, HyperDispatch, Call Home, etc.




LinuxONE is designed for minimal unplanned downtime

Unplanned Downtime of >four (4) hours on each server hardware platform (2015)



ITIC survey shows IBM LinuxONE had no – 0% – unplanned system downtime due to inherent flaws in hardware

Recent IDC study concludes clients who leverage LinuxONE can virtually eliminate lost productivity caused by downtime



Source: IDC 2016

21

IBM LinuxONE™

Source: ITIC 2016 - 2016 Global Server Hardware, Server OS Reliability Survey

© 2016 IBM Corporation

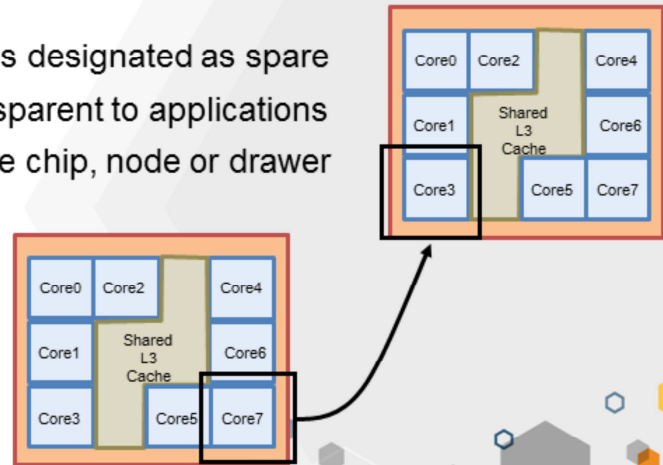


If a core fails, a spare can be “turned on” without system or program interruption



- Each LinuxONE server has two cores designated as spare
- Core failover (called sparing) is transparent to applications
- Spares need not be local to the same chip, node or drawer
- Any core can failover to a spare

x86 servers
do not have core
sparing





Achieve 100% continuous availability

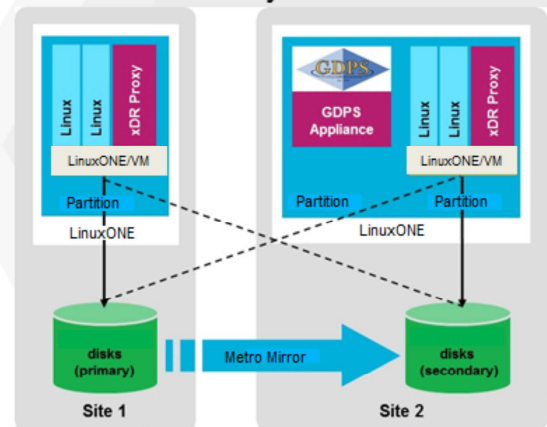


Geographically Dispersed Parallel Sysplex –

A fully integrated software solution providing continuous availability / disaster recovery for LinuxONE virtual machines

- Applications restarted in another site when disaster event occurs
- Single point of control and automation, reliable and rapid recovery
- Keeps data available and consistent
- Protects against disk subsystem failures

Recover from outages
in **6 seconds** instead of 2+ hours



23


IBM LinuxONE™

GDPS Virtual Appliance requires LinuxONE/VM and the HyperSwap. Optionally, HyperSwap is not currently available for bare partitions or for KVM on LinuxONE.


© 2015 IBM Corporation



Advanced cryptography is handled at multiple levels depending on business requirements



On-chip cryptography



- Each core has its own cryptographic co-processor
 - Optimized for encryption functions
- Crypto Express5S PCIe card (optional) adds additional crypto capability
 - Elliptic Curve, SHA3, Visa FPE, etc.
- Meets FIPS, ANSI, PKI, and DK standards



Optional Crypto Express5S for additional advanced cryptography features


24

IBM LinuxONE™

© 2015 IBM Corporation



Top to bottom security is built in, not bolted on



Resource Access Control Facility (RACF) is the backbone of LinuxONE security

- Access control to all classes of resources
- Integrated into LinuxONE/VM
- Supports cryptographic services
- Supports digital certificates


Enables application and database security without modifying applications

Reduces security complexity and expense:

- Central security process that is easy to apply to new workloads or as user base increases
- Tracks activity to address audit and compliance requirements



Average security incursion costs on LinuxONE are much lower than other platforms



Comprehensive Security:

Only LinuxONE natively covers multiple dimensions of security; other platforms require augmentation and additional expense (Solitaire)

Customer surveys show LinuxONE average incursion costs are a fraction of distributed platforms



IBM LinuxONE is a better choice if...



- ...You're running very large **I/O- or cache-intensive** workloads like database applications or transaction processing
- ...Your **insights and engagement** are time-sensitive
- ...You're looking to **consolidate** large numbers of servers
- ...Your **software license costs** are unsustainable
- ...Even the smallest **system failure** is unacceptable
- ...A **security** breach would be catastrophic to your business
- ...You're looking to reduce overall **operational expenses**



IBM LinuxONE™

IBM LinuxONE™

© 2015 IBM Corporation



Thank You

Marcel Mitran – CTO IBM LinuxONE
mmitran@ca.ibm.com



IBM **LinuxONE**

IBM LinuxONE™

© 2015 IBM Corporation