Nines all the Way Down: Resiliency for Linux on IBM zSystems and LinuxONE

Wilhelm Mild

IBM Executive IT Architect - Integration Architectures for Containers, Mobile, IBM Z and Linux Infrastructure

wilhelm.mild@de.ibm.com

Steven Cook

GDPS Software Designer & Team Lead, Certified IT Specialist, Client Technical Advocate, IBM Z Infrastructure

cooksd@us.ibm.com

Agenda

Why is Resiliency?
(At a business level, what is this all about?)

What is Resiliency?
(High Availability, Business Continuity, RTO… Huh?)

What can LinuxONE do for me?
(How does resiliency function, from hardware to containers?)

Sample Resiliency Framework
(Multi-site Build for Maximal Uptime)

Agenda

Why is Resiliency?
(At a business level, what is this all about?)

What is Resiliency?
(High Availability, Business Continuity, RTO… Huh?)

What can LinuxONE do for me?
(How does resiliency function, from hardware to containers?)

Sample Resiliency Framework
(Multi-site Build for Maximal Uptime)

Business Continuity: How are you doing?

- **What level of Business Continuity do you require, and do you fulfill that requirement end-to-end?**
- Can you remove people as Single Points of Failure?
- Can you reduce recovery times using accurate and automated processes?
- Can you stress test innovations at production scale?
- Can you optimize recovery processes?
- How do you verify accuracy and viability of your recovery?
- What does it cost you to fail?

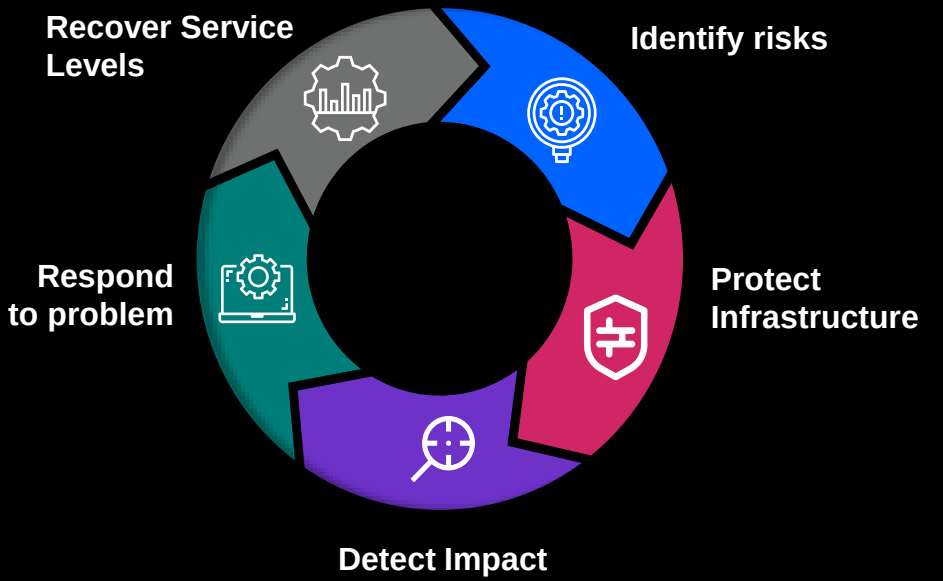

How much is your infrastructure costing you per year?

Sample downtime costs / hour excluding fines and penalties

IBM LinuxONE Design Point

	99%	99.9%	99.95%	99.99%	99.99999%
$150,000+ / hr (98% respondents)	$8,766,000	$876,600	$438,300	$87,660	$88
$300,000+ / hr (88% respondents)	$26,000,000	$2,640,000	$1,315,000	$262,300	$263
$1,000,000+ / hr (40% respondents)	$87,660,000	$8,766,000	$4,383,000	$876,600	$877

https://www.ibm.com/downloads/cas/DV0XZV6R

Hybrid Cloud differentiators with IBM Z & LinuxONE

IBM Z Benefits

Low Latency and Large Volume Data Serving and Transaction processing

Enterprise class infrastructure – Elastic, Scalable, Available and Resilient

Highest levels of Security and Compliance

Adoption Patterns

Enterprise Hybrid Cloud
2.4M containers-per-box

Colocation & Modernization with z/OS
7x shorter batch windows

Secure Cloud Native
Confidential computing

Extreme Consolidation and scalable Data Serving
69% lower Total Cost of Ownership

99.99999% *system availability*

4:1 *better data-center footprint*
2:1 *lower power envelope*

3.8x better *Java throughput*

Enterprise grade. **Open** by design. **Secured** by IBM Z.

Agenda

Why is Resiliency?
(At a business level, what is this all about?)

What is Resiliency?
(High Availability, Business Continuity, RTO… Huh?)

What can LinuxONE do for me?
(How does resiliency function, from hardware to containers?)

Sample Resiliency Framework
(Multi-site Build for Maximal Uptime)

Why does system downtime happen?

Planned outages

- Test phase for new applications
- System maintenance for older hardware or operating systems
- Standard service that requires a reboot or re-IPL

Unplanned outages – system failures

- Hardware, firmware, or OS failures
- Network outages, or Service outages from cloud providers
- Application failures or configuration problems

Unplanned outage due to human errors

- Lack of skills, poor documentation, heterogenous environments

Unplanned outages -- disasters

- Security threats
- Natural disasters

What it means: High availability

- High Availability (HA) – Provide service during defined periods, at acceptable or agreed upon levels, and masks *unplanned* outages from end-users. It employs Fault Tolerance; Automated Failure Detection, Recovery, Bypass Reconfiguration, Testing, Problem and Change Management

- Business Continuity (BC) -- Continuously operate and mask *planned* outages from end-users. It employs Non-disruptive hardware and software changes, non-disruptive configuration, software coexistence.

- Continuous Availability (CA) -- Deliver non-disruptive service to the end user 7 days a week, 24 hours a day (there are no planned **or** unplanned outages).

The goal is to strive to provide Continuous Availability.

Key Performance Indicators for HA: RTO, RPO, and NRO

Business Resiliency Plan

Recovery **T**ime **O**bjective (**RTO**)

What time difference can be between Failure and a total productional run level ?

Recovery **P**oint **O**bjective (**RPO**)

What is the toleration for data loss?

RPO = “0” means, NULL data loss acceptable
RPO = “5” means, data loss in last 5 min acceptable

TREND: RPO = 0

Network **R**ecovery **O**bjective (**NRO**)

Time requirements for network availability.

Fundamentals of HA

Redundancy, Redundancy, Redundancy
Duplicate to eliminate single points of failure.

Early detection
To keep offline time as short as possible
Reduce risk of wrong interpretation and unnecessary failover
Keep offline time as short as possible (mean-time-to-repair MTTR)

Protect Data Consistency – Provide ability for data and file systems to return to a point of consistency after a crash.
Journaling databases
Journaling file systems
Mirroring
Routine database backups

Automate Detection and Failover - Let the system do the work in order to minimize outage windows.
Multipath
VIPA –Virtual IP Addresses
Monitoring and heart-beating
Clustered middleware
Clustered operating systems

Differences between HA and DR

High Availability - HA:

- Failover is typically realized via duplication and clustering
- Failover times measured in seconds and minutes
- Reliable synchronous inter-node communication

Disaster Recovery - DR:

- Failover is typically realized with 2 or more sites in case of disasters
- Failover times often measured in minutes and hours
- Unreliable inter-node communication assumed

How it works: Clustering

Computer Cluster

A computer cluster consists of a set of loosely connected computers that work together so that in many respects they can be viewed as a single system. (Wikipedia definition: Computer Cluster)

High Availability Cluster

A computer cluster where each cluster operates as workload node. When one node fails another node takes over the entire workload: IP address, data access, services, etc.

The key of High Availability is avoiding single points of failure

High Availability adds costs because of added complexity due to redundant resources in the environment

How it works: Split Brain

If the heartbeat between nodes fails, all nodes can still be active and:

- detect the other as failing
- the status of an unreachable node is unknown

Communication/heartbeat failures between cluster nodes can lead to isolated actions in separated partitions of the cluster

If those partitions each try and take control of the cluster, then it's called a split-brain condition

This can lead to data corruption or inconsistency; therefore, split brain must be inhibited

How it works: Quorum

Quorum is an attempt to avoid split brain for most kinds of failures

Your Cluster Management Software tries to make sure only one partition can be active

Quorum is the term for methods for enforcing which part of the cluster is active:

- A quorum server – as additional node-can decide more reliably
- Quorum server is in a quorum daemon

Most common kind of quorum is voting – and only one partition can run the cluster

How it works: Fencing

Fencing puts a fence around an errant node or inhibit a failed node from accessing cluster resources

This way one doesn't have to rely on correct behavior or timing of the errant node

This is often implemented via STONITH

– STONITH: Shoot The Other Node In The Head

Other techniques also work

– use of hardware or software watchdog timers
– self shutdown / restart
– Shared device for notification instead of heartbeat

Agenda

Why is Resiliency?
(At a business level, what is this all about?)

What is Resiliency?
(High Availability, Business Continuity, RTO… Huh?)

What can LinuxONE do for me?
(How does resiliency function, from hardware to containers?)

Sample Resiliency Framework
(Multi-site Build for Maximal Uptime)

IBM LinuxONE Resiliency Stack

Faster recovery when failures occur
- ✓ **System Recovery Boost (z/VM, VSE, TPF)**
- ✓ Pacemaker and Corosync
- ✓ HyperPAV
- ✓ Linux Auto-IPL
- ✓ z/VM Snapdump

React faster to workload fluctuations
- ✓ Virtual Flash Memory
- ✓ HyperPAV
- ✓ FICON 16s+
- ✓ Dynamic Partition Manager
- ✓ HiperDispatch
- ✓ Dynamic memory upgrade / downgrade

Improved workload scaling
- ✓ HiperDispatch
- ✓ HyperPAV
- ✓ NVMe support
- ✓ 32 GB FCP
- ✓ JVMS pause-less garbage collection
- ✓ FICON 16S+

Failure Avoidance
- ✓ IBM Z Business Resilience Stress Test
- ✓ Safeguarded Copy and Cyber Vault

Concurrent maintenance
- ✓ z/VM Single System Image
- ✓ Guest relocation and mobility
- ✓ Dynamic I/O Reconfiguration
- ✓ HyperSwap for z/VM
- ✓ GDPS Support

Problem Determination and Data Capture
- ✓ Auditd
- ✓ Logging and tracing of device drivers
- ✓ AutoIPL w/Standalone Dump

Preventative / Real-Time Insight
- ✓ Pacemaker and Corosync
- ✓ OpenShift Container Platform (OCP)

LinuxONE – Redundancy in Hardware Support

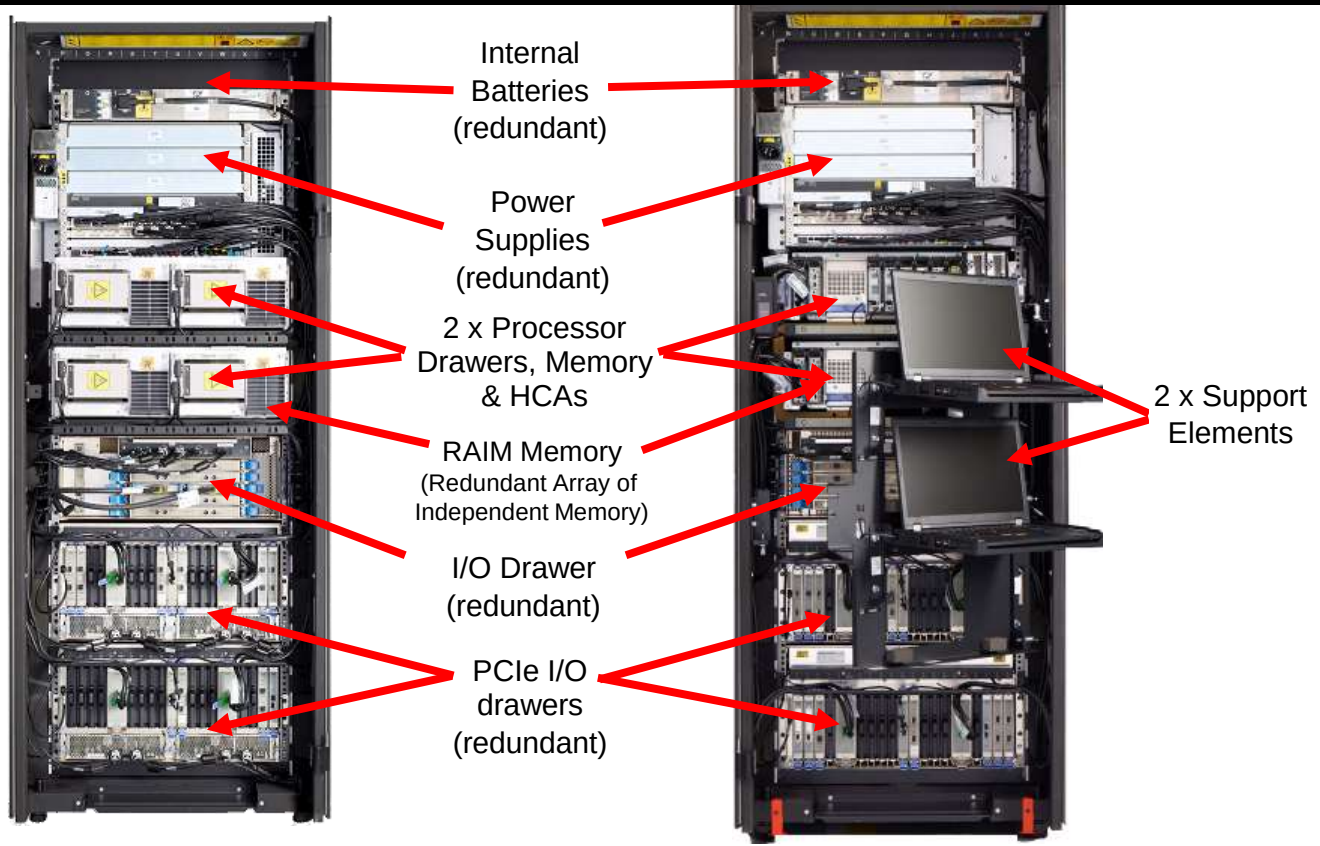

LinuxONE – Redundancy of Cache and Memory Protection

Built-in redundancy protects data against attacks on hardware

– ECC

– RAIM

– Defense against rowhammer etc.

Cache separation and virtualization tech prevents side-channel attacks

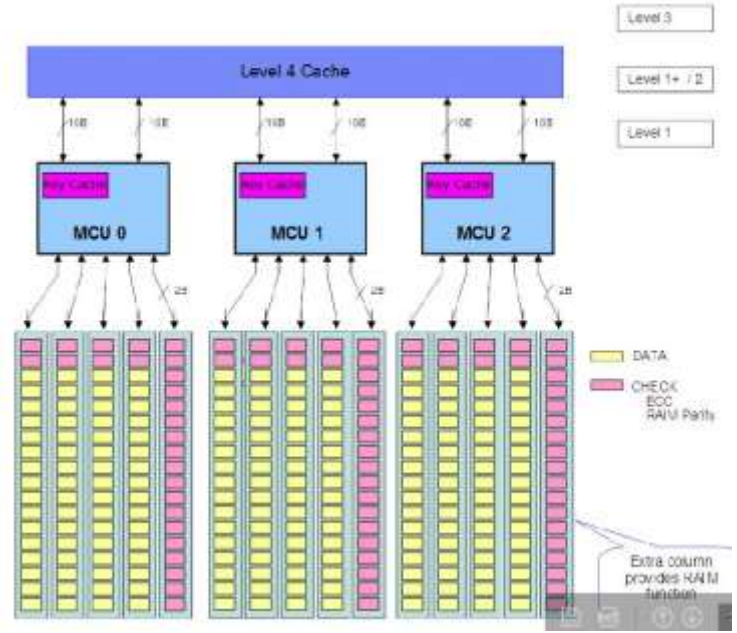

LinuxONE -- Core Sparing

- 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 node or drawer
- Any core (general processing core or I/O core) can failover to spare

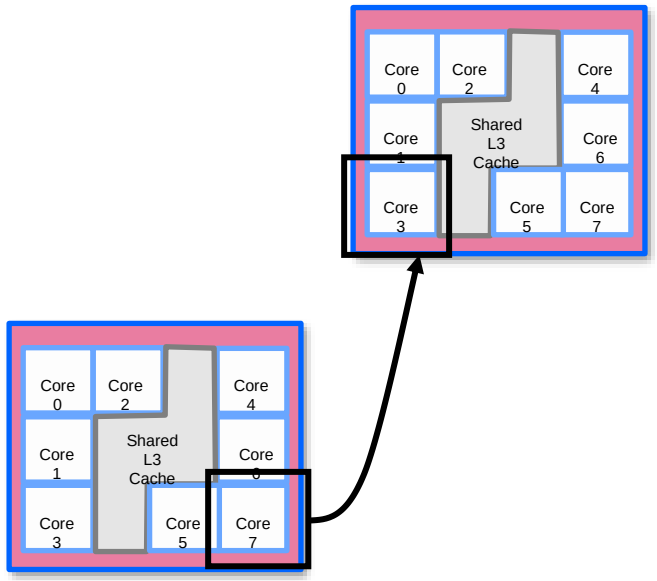

LinuxONE – Separation of Partitions and System Integrity

EAL5 – better than an air gap when it comes to data confidentiality and compute management

- Isolation of a logical partition at the architectural level (more on this in a moment)
- Controls on direct access to devices
- **Elimination** of covert channels
- Role-based access controls to a partition (or partitions), or hardware

With a few added bonuses:

- Controlled in-memory communication paths (**HiperSockets**)

IBM System Recovery Boost

Unleash your capacity to maximize your availability

Restore service and recover workloads substantially faster than on previous generations, built into IBM LinuxONE III with ***zero increase in IBM software licensing costs*.**

Faster shutdown and startup

Accelerate the planned shutdown, restart and recovery of images, middleware environments and client workloads to accelerate return to pre-shutdown SLAs.

Faster partition recovery

Accelerate Parallel Sysplex recovery processes to minimize disruption and expedite return to steady-state operations.

Faster GDPS automation actions

Drive faster and more efficient GDPS automation actions to rapidly reconfigure and recover your environment

Faster elimination of backlog

Utilize additional capacity for a fixed period during recovery, so you can process backlog faster after planned or unplanned downtime.

IBM z/VM

- Platform-native virtualization technology
- Architectural isolation and VM separation (EAL 4+)
- System Recovery Boost for host and guests
- Link aggregation and failover of Virtual Networking
- HyperPAV aliasing for I/O configuration
- HyperSwap support for rapid disk failover
- GDPS xDR for mirroring and data replication
- Single System Image clustering for planned outages and guest mobility

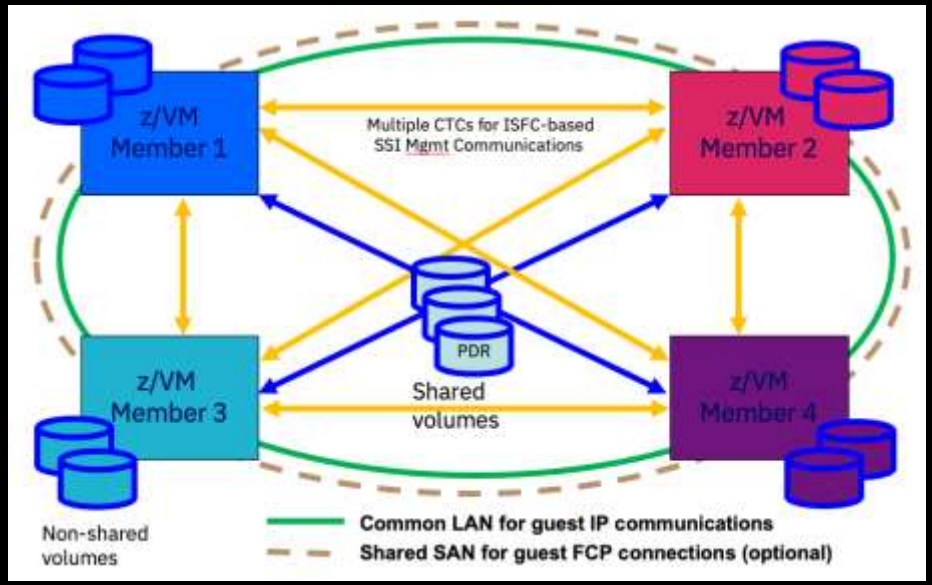

z/VM Single System Image – Cluster Configuration

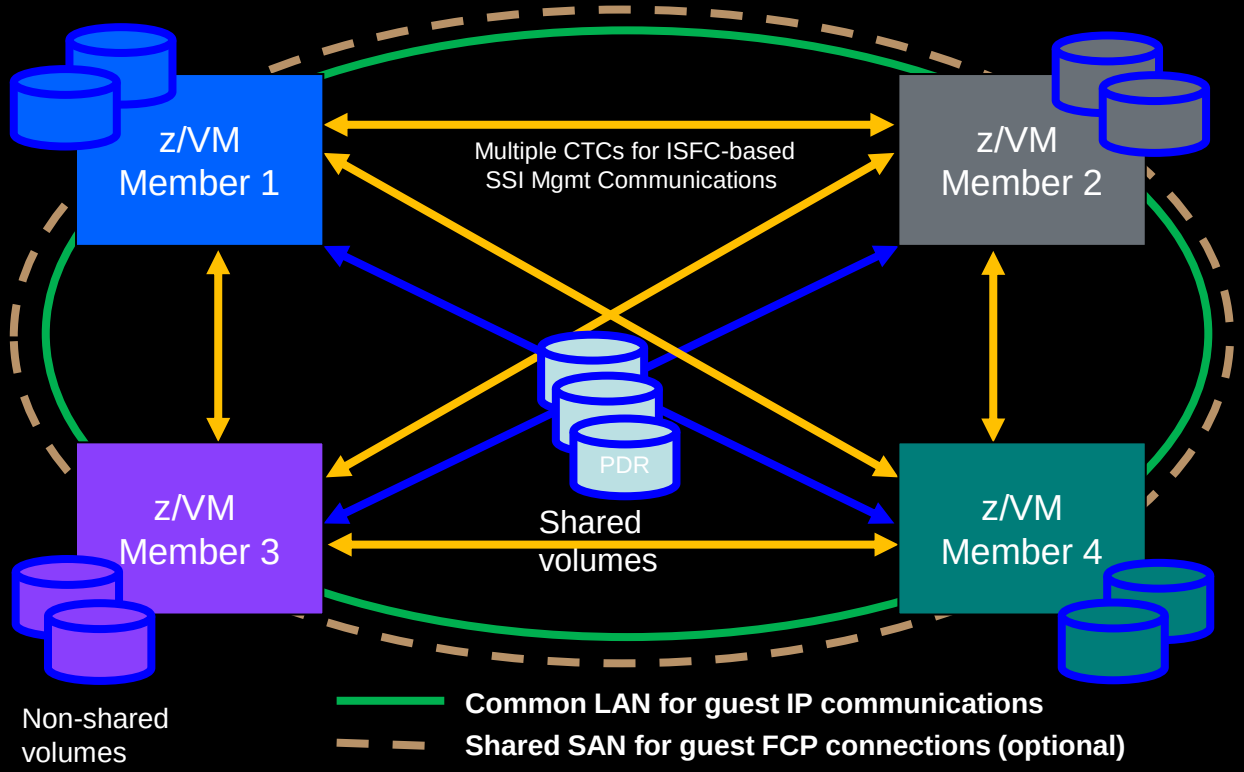

Availability guards:

- ✓ Linux guests can run anywhere in cluster and be relocated while running
- ✓ Multiple CTCs for redundancy, but even if all are lost, enters Safe Mode, allowing work to continue, but not starting new virtual servers
- ✓ Up to 4 + 4 Members (z/VM System) can be on separate CPCs for greater redundancy. Gold standard two CPCs with two members each.
- ✓ Planned outages addressed by moving virtual machines off member being serviced and then back again.
- ✓ Ability to move Volume with PDR (Persistent Data Record) while SSI Cluster is up.
- ✓ Typically exploit IBM Operations Manager for automation of restart and/or LGR

Risks:

- ➢ Unplanned out of one member disrupts 1/Nth of workload (see middleware clustering).
- ➢ All members fail at once or underlying CPCs.
- ➢ Unplanned outage of PDR volume (see GDPS).
- ➢ Insufficient resources for cascading failure.(Various IBM Z ability to add resources dynamically)

Redundancy in Virtual Networking via Virtual Switches

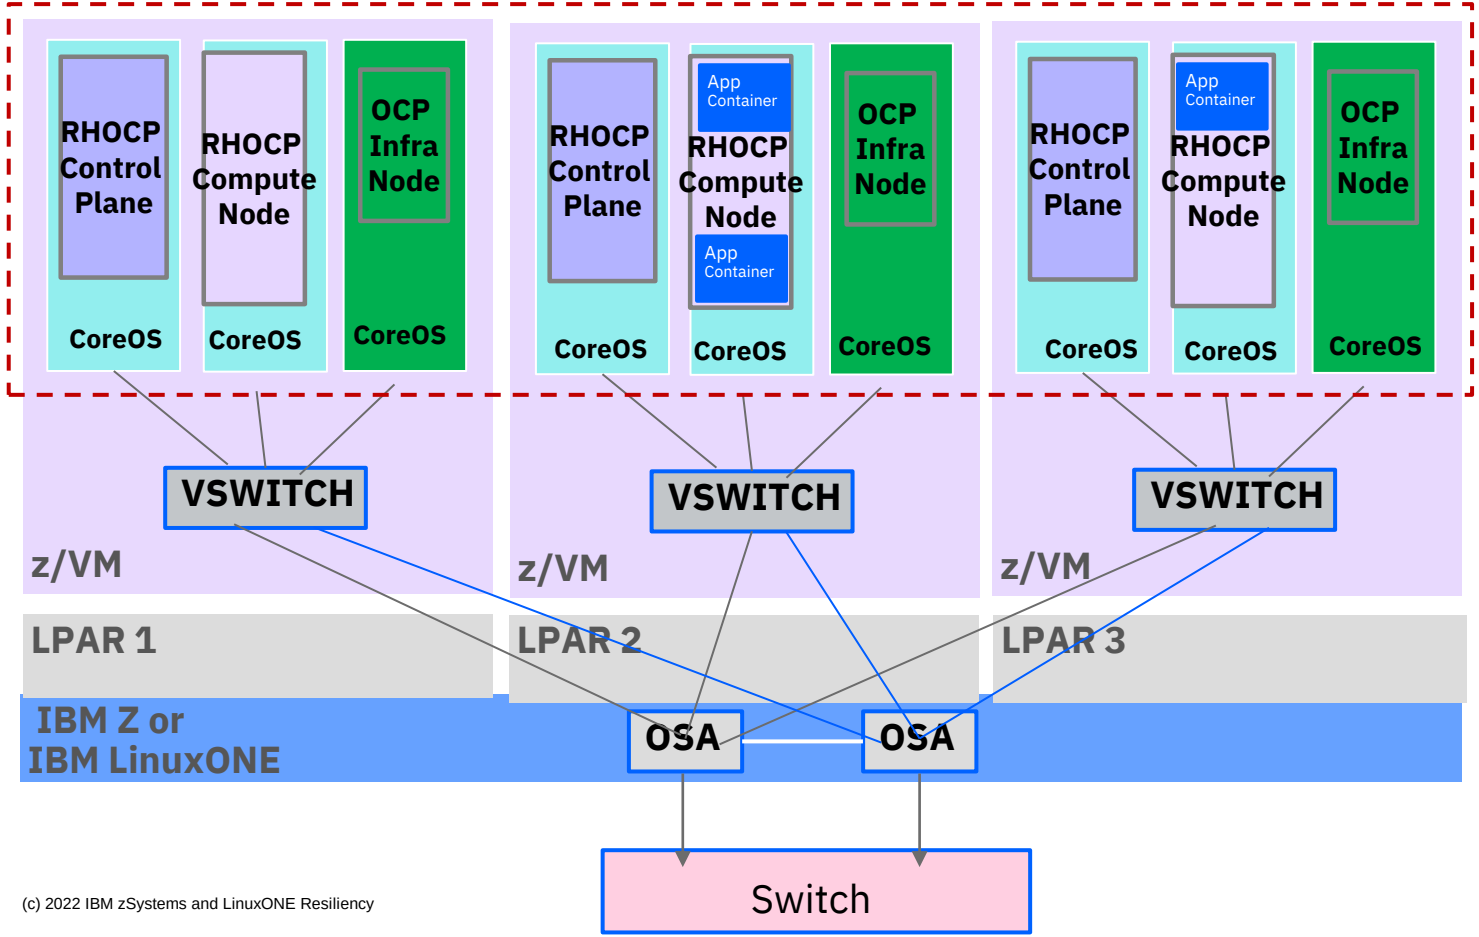

Multiple virtual devices at CoreOS host level (Layer 3)

Connect to z/VM Virtual Switch (Layer 2)

Link Aggregation to bind multiple OSA ports across multipe cards on a CEC

Intra-Vswitch Link to collate connections across physical hardware in an SSI cluster

Physical switches and cabling to bind sites together

200km range for metro mirroring (per GDPS)

Allows for network traffic separation at virtual, logical, or physical layers in accordance with security requirements

Guest Infrastructure

Once you've covered storage, hardware, and virtual infrastructure, we move on to workload availability

Consider:

- Failover / DR partitions – for when your entire partition needs to reappear in another location
  - Mind volume access, HW-critical settings, and crypto/key access
- VM duplication – having VM's or containers available to run workload in an emergency
  - Doesn't avoid an outage, but having the spare VM may shorten MTTR
- "Dead guest relocation" (shutdown and bring-up may be faster than guest mobility)

Linux-HA package - High Availability components

Heartbeat

- Messaging between nodes to make sure they are alive and available
- Action required if heartbeat stops after certain tries

Cluster-glue

- Everything that is not messaging layer and not resource manager

Resource-agents

- Agents running in clustered systems or remote
- Agents are able to start, restart or stop services

Pacemaker

- A Cluster Resource Manager (CRM)

OpenSAF checkpoint APIs

- SAF -> Service Availability Forum – created the Service Availability Specifications
- OpenHPI - The Hardware Platform Interface (HPI) abstracts the differences between hardware implementations, providing a uniform interface to hardware features.
- OpenAIS - The Application Interface Specification (AIS) specifies an interface that applications interchange information with the service availability middleware (i.e. CRM).

Pacemaker

Provides the policies, and starts/stops the virtual machines

Maintains a **Cluster Information Base (CIB)** in each host instance
- XML list of behaviors, directed by resource manager, informs policy engine
- This is where a client would define policy

Policy Engine takes that list of behaviors and maps them to cluster's current state

Cluster Resource Manager is the focal point for receiving cluster ops.
- One Designated Focal Point for the cluster
- The other hosts receive data from the **CRMd** via **corosync**
- The "Local Resource Manager" receives instructions from the CRMd and passes requests along to local resource agents (VirtualDomain, FileSystem, MailTo… general opertions)

STONITH is a fencing agent that detects if Pacemaker loses contact with one of the nodes in the cluster
- If Pacemaker thinks a node is down, STONITH will force it offline
  - "**S**hoot **T**he **O**ther **N**ode **I**n **T**he **H**ead"
- Fences the failed node to ensure data integrity

Corosync

- Both a component and daemon in a Linux HA cluster
- Handles **membership enrollment and member-communication** for any Linux instance enrolled in an HA cluster
- **Manages quorum rules and determination**
  - Quorum ensures that the majority of a given cluster agrees on what the world is
  - *Votequorum* as an interface for members to agree
  - Majority wins (special tiebreaker policy available for even-numbers of clusters)
- Also provides messaging for applications coordinating/operating across multiple members of a cluster

High Availability and Red Hat OpenShift Container Platform

Orchestration

- OpenShift implements a highly scalable infrastructure. To detect node problems or infrastructure service problems.
- There is always a set of the same (e.g. 3 API Server / etcd) for the control plane in case of a failure on in one instance the infrastructure keeps still running and operational
- HW is expected not reliable therefore Kubernetes implementation takes care it is running reliable.

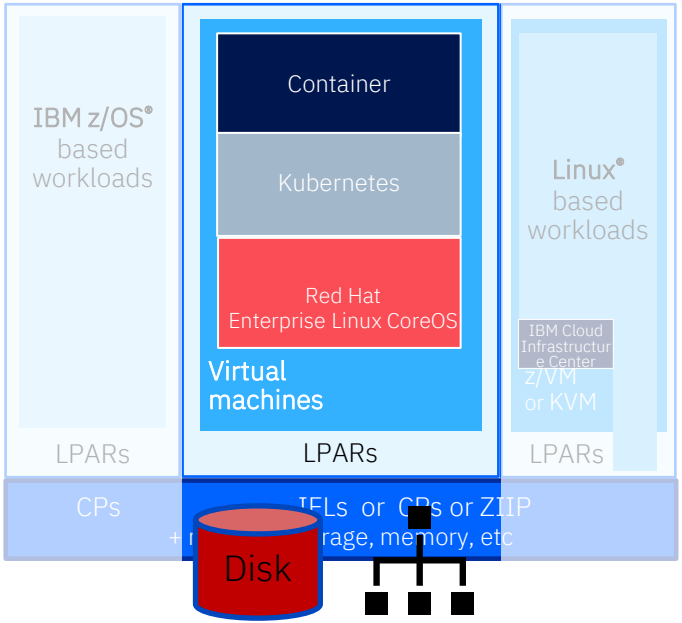

High Availability – Red Hat Infrastructure

Red Hat CoreOS takes care to keep single guests installed

- Integration for the platform, supporting devices, take care for recovery.
- Automated installation and container native OS and fully immutable
- Installation simplified through ignition
  - Case of problems with a single node – re-installation is possible and re-integration

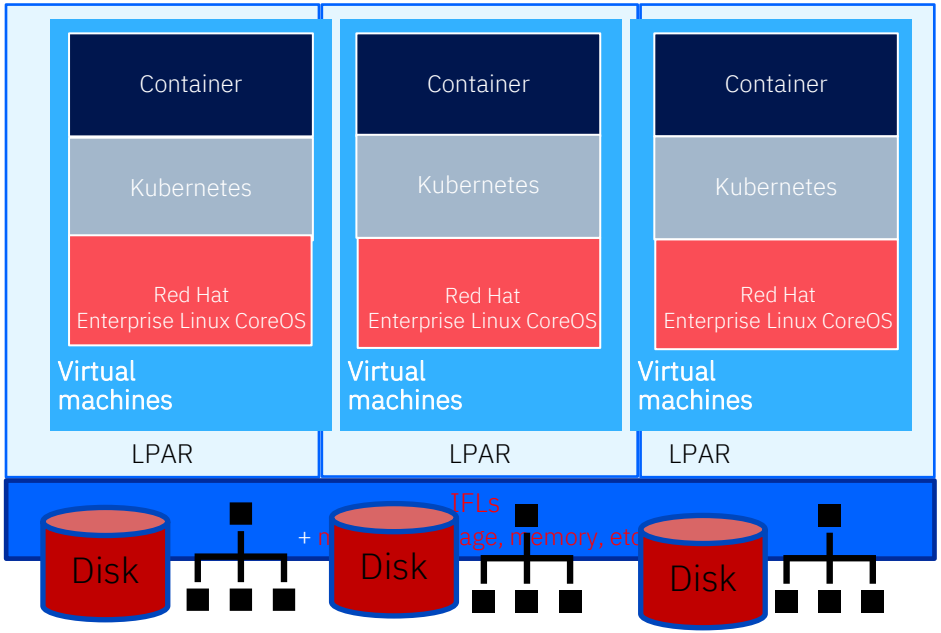

High Availability – IBM zSystems Infrastructure, Hypervisor & Storage

IBM Hardware and Hypervisor Components are highly reliable and keep the system running

- Additional facilities to raise avialablity are:
  - Split Cluster among CECs
  - Build a z/VM Single System Infrastructure to easy migrate/restart nodes on another cec
  - use GDPS for storage replication and manament to migrate guests.
- Depending on the container storage consider the right way to add Storage
- Think of how to mitigate mobility of the storage

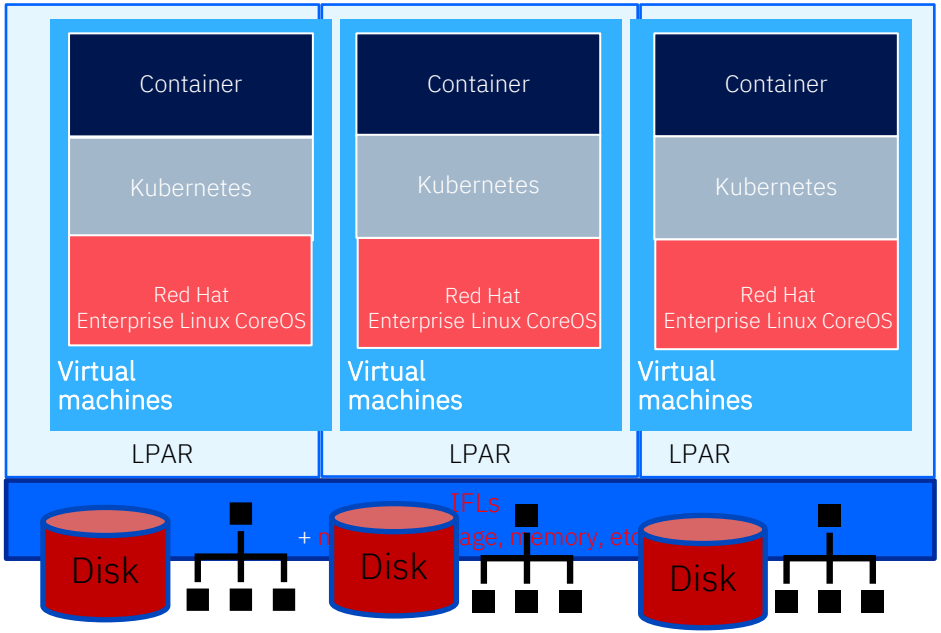

High Availability – IBM zSystems Infrastructure, Hypervisor & Storage

Multiple Sites – Stretch Cluster

- One Cluster stretched on multiple sides
- Increases availablility if:
  - Guests can migrated among to sides e.g. in side fail control plane can get back to 3 (minimum 2) have the cluster operational.
  - In case of a stretch cluster consider Round Trip Time (RTT) for etcd must not higher than 100ms better is less latency -> 50ms be on the save side

Multiple Sites – Two Clusters

- Requires Global Load Balanacer to switch among both sides
- Consider data replication for persistent workload
- Consider Orchestration via Red Hat Advanced Cluster Manager

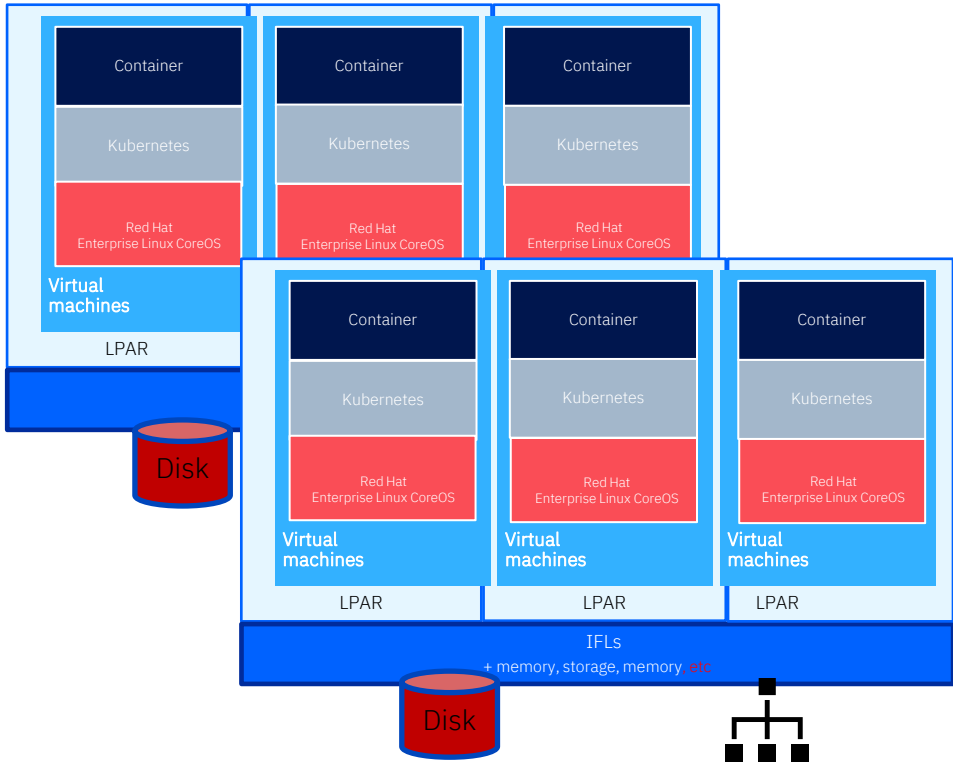

Recover fast with IBM GDPS with minimal data loss

IBM GDPS solutions are designed to:

- Deliver end-to-end application and data availability within a site or across multiple sites
- Automate recovery procedures for near-continuous availability and disaster recovery
- Provide extremely rapid RPO and RTO
- Provide multiple, secure, point-in-time copies of production data to help protect from logical corruption threats
- Monitor systems and storage that support open and IBM copy technology architectures, including IBM, Hitachi Vantara and Dell EMC (Metro Mirror only for non-IBM storage)
- Simplify management tasks with an easy-to-use interface and a central point of control

Safeguarded Copy for logical corruption protection

- Provides up to 500 Safeguarded Backups
- Stores backups which is hidden and not accessible by any server (Air gap)
- Secures the data so it can only be accessed after a Safeguarded Backup is recovered to a separate recovery volume
- Performs data validation and forensic analysis to restore entire production data or individual volumes

Prevents sensitive point in time copies of data from being modified or deleted due to user errors, malicious destruction, or ransomware attacks

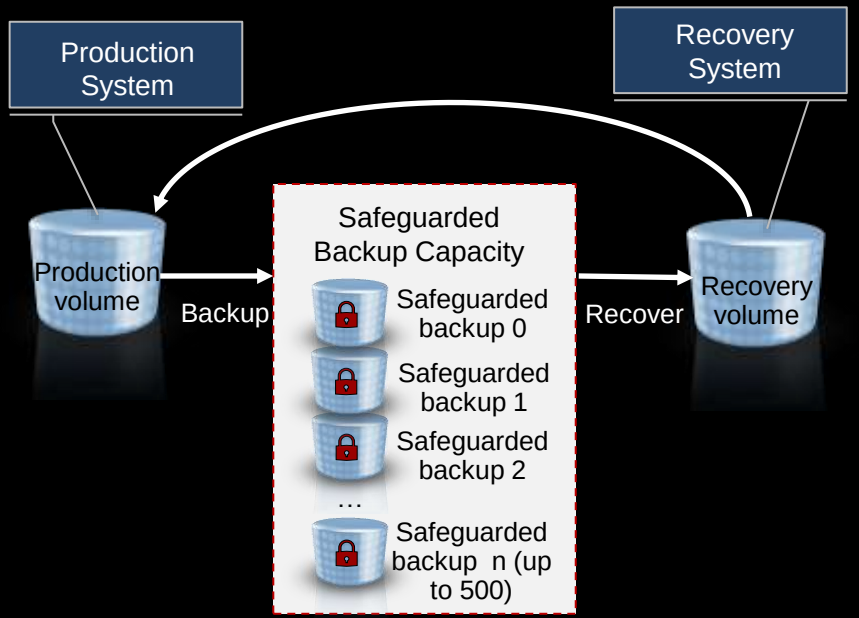

IBM DS8900F and IBM LinuxONE: Integration by Design

All-Flash enterprise storage for production environments with mission-critical requirements

Ultra low-latency
Fastest application response time

High availability
Better than seven 9's with HyperSwap technology

Advanced DR
3 and 4 site replication with 2 to 4 seconds RPO and less than 60 seconds RTO

Transparent Cloud Tiering
Up to 50% savings in IBM Z CPU utilization when transparently migrating data to the cloud

100% data encryption
Secure authentication and in-flight encryption to protect data wherever it resides

Cyber resiliency
immutable snapshorts to prevent data from being modified or deleted

IBM LinuxONE – Leading the industry in IT resiliency and integration

Avoiding the cost of downtime
Ensuring access to critical apps / data

Maintaining productivity of users
Open to clients 24/7

IBM LinuxONE

Where mean time between failure is measured in decades

z/VM Single System Image

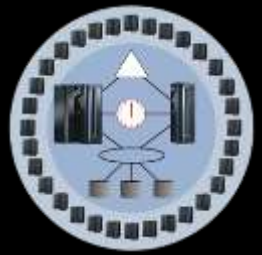

Designed for *guest mobility* and workload balancing

GDPS

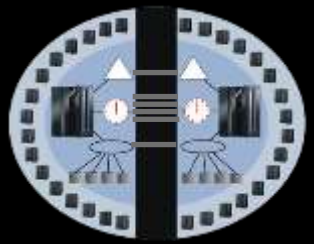

Industry leading solution for continuous availability / disaster recovery

Storage Synergy

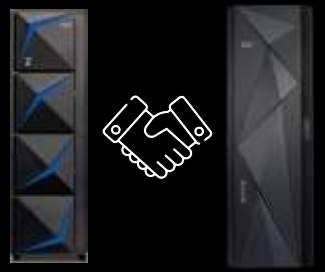

Integration by design to maximize technology advantages and
> seven 9's availability

Agenda

Why is Resiliency?
(At a business level, what is this all about?)

What is Resiliency?
(High Availability, Business Continuity, RTO… Huh?)

What can LinuxONE do for me?
(How does resiliency function, from hardware to containers?)

Sample Resiliency Framework
(Multi-site Build for Maximal Uptime)

Assumptions

- These claims are based on an optimum configuration for the calculation of availability for OpenShift Container Platform workloads running on IBM Z or LinuxONE hardware

  - *We assume proper configuration, whereupon all components (hardware/software) have appropriate back-up and redundancy to allow for outage mitigation and to ensure minimal performance impact, within client requirements*
    - *This includes "free space" on systems (CPU, storage, memory) in case of planned outages and workload redistribution*
    - *This includes configuring z/VM and RHOCP for High Availability and redundancy – virtual resources are resources*

  - *We assume that the greater datacenter ecosystem is configured for redundancy and high availability*
    - *This includes services such as DNS and load balancing for applications running in RHOCP*
    - *This includes physical cabling for networking*
    - *This includes power/electricity, fire suppression, etc.*

  - *We assume a competent and non-malicious set of system administrators*
    - *This includes testing fixes on a dev system, isolated from production workloads*

  - *We assume a properly configured z/OS system running GDPS and xDR exists somewhere pertinent in the enterprise.*
    - *Parallel Sysplex Best Practices: http://www.redbooks.ibm.com/abstracts/sg247817.html?Open*
    - *GDPS Configuration*

  - *We assume that all pertinent security guidelines have been followed with regards to configuration of the included Framework.*

Threats

- This resiliency model has been designed to withstand the following types of threats to service availability:

- Natural disasters
  - Multi-site configuration meant to support this, but flexibility depends on size/scope of disaster

- Planned outages to systems

- Power outages
  - Having backup generators and power failover should already be a physical availability consideration.

- Replication failures
  - Restoration of previous copies of workload, dependent upon forensic investigation
  - Data integrity check following (f.ex.) system upgrade

- While security was considered, cyber attack is outside the formal scope of our evaluation
  - Pending metrics on cyber resiliency and recovery when Cyber Vault solutions are in use within the managed environment

Our Selected Resiliency Framework

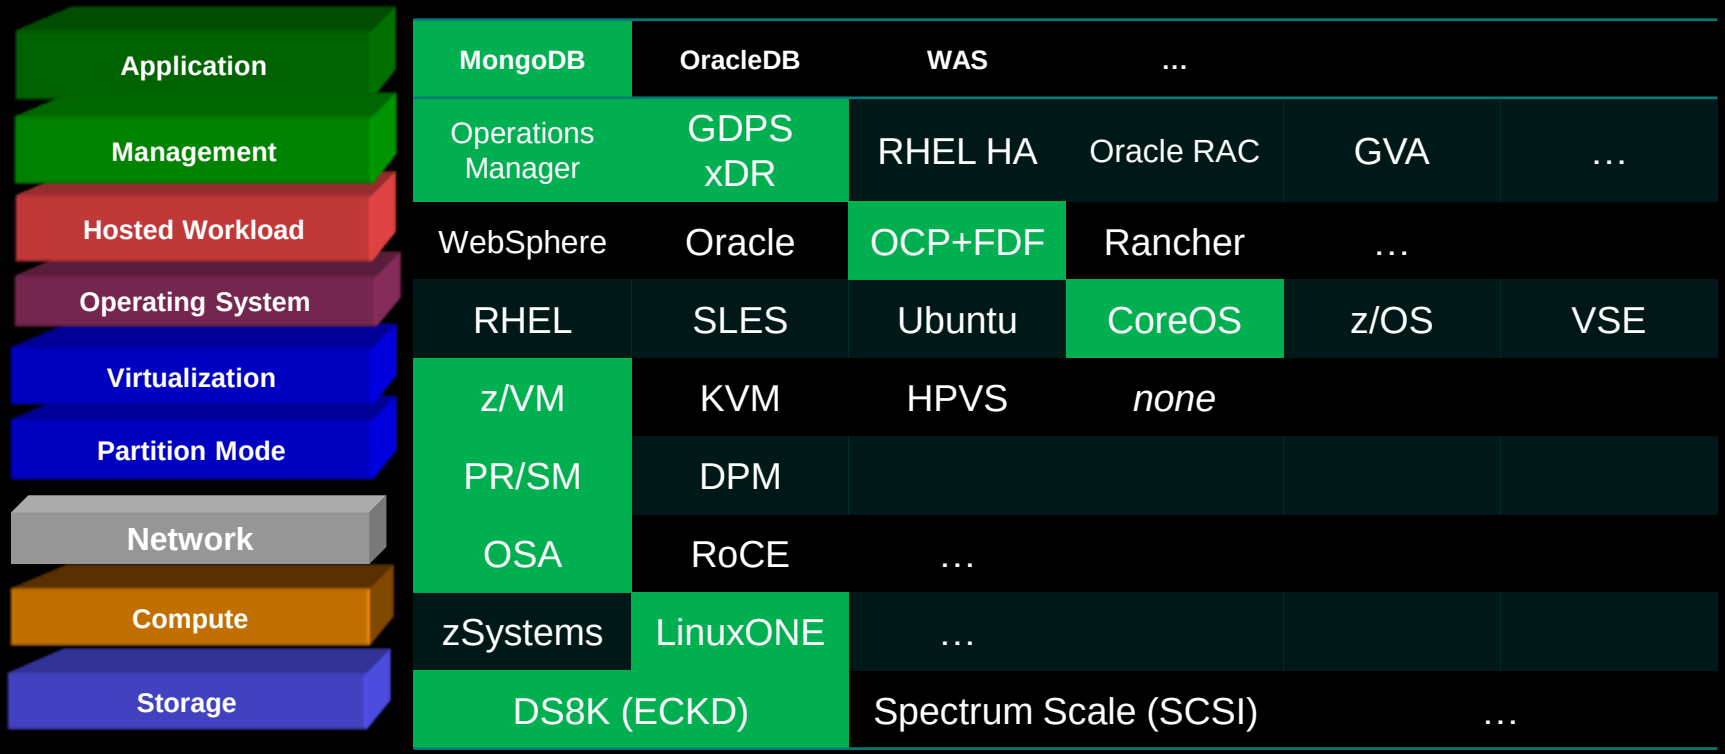

Components of the Sample Resiliency Framework

- Multiple physical sites at regional distance
- At least two (2) Storage units (DS8K-series configured for ECKD)
- One (1) IBM z16 machine per site for running z/OS 2.5, GDPS 4.5 (K-sys and xDR)
- LinuxONE Emperor 4 machines, 2+ CPCs (at least one per site), each configured PR/SM mode (not DPM)
- z/VM V7.x with virtual networking, HyperSwap, and Single System Image
  - Minimum of three (3) z/VM partitions (out of 4 maximum) in an SSI cluster
  - z/VM V7.3, GA 3Q2022, increases allowable cluster to 8 members
  - IBM Operations Manager for z/VM, V1.6
  - GDPS 4.5 xDR Proxies for z/VM
- RHOCP V4.10 (CoreOS) configured for containerization resiliency
  - *Red Hat Openshift Data Foundation 4.10 (ODF) for local storage*
  - *SAMP for RHOCP Control Node and Compute Node*
  - *Control and Compute nodes should not be converged in a single guest*

Sample Framework in Action

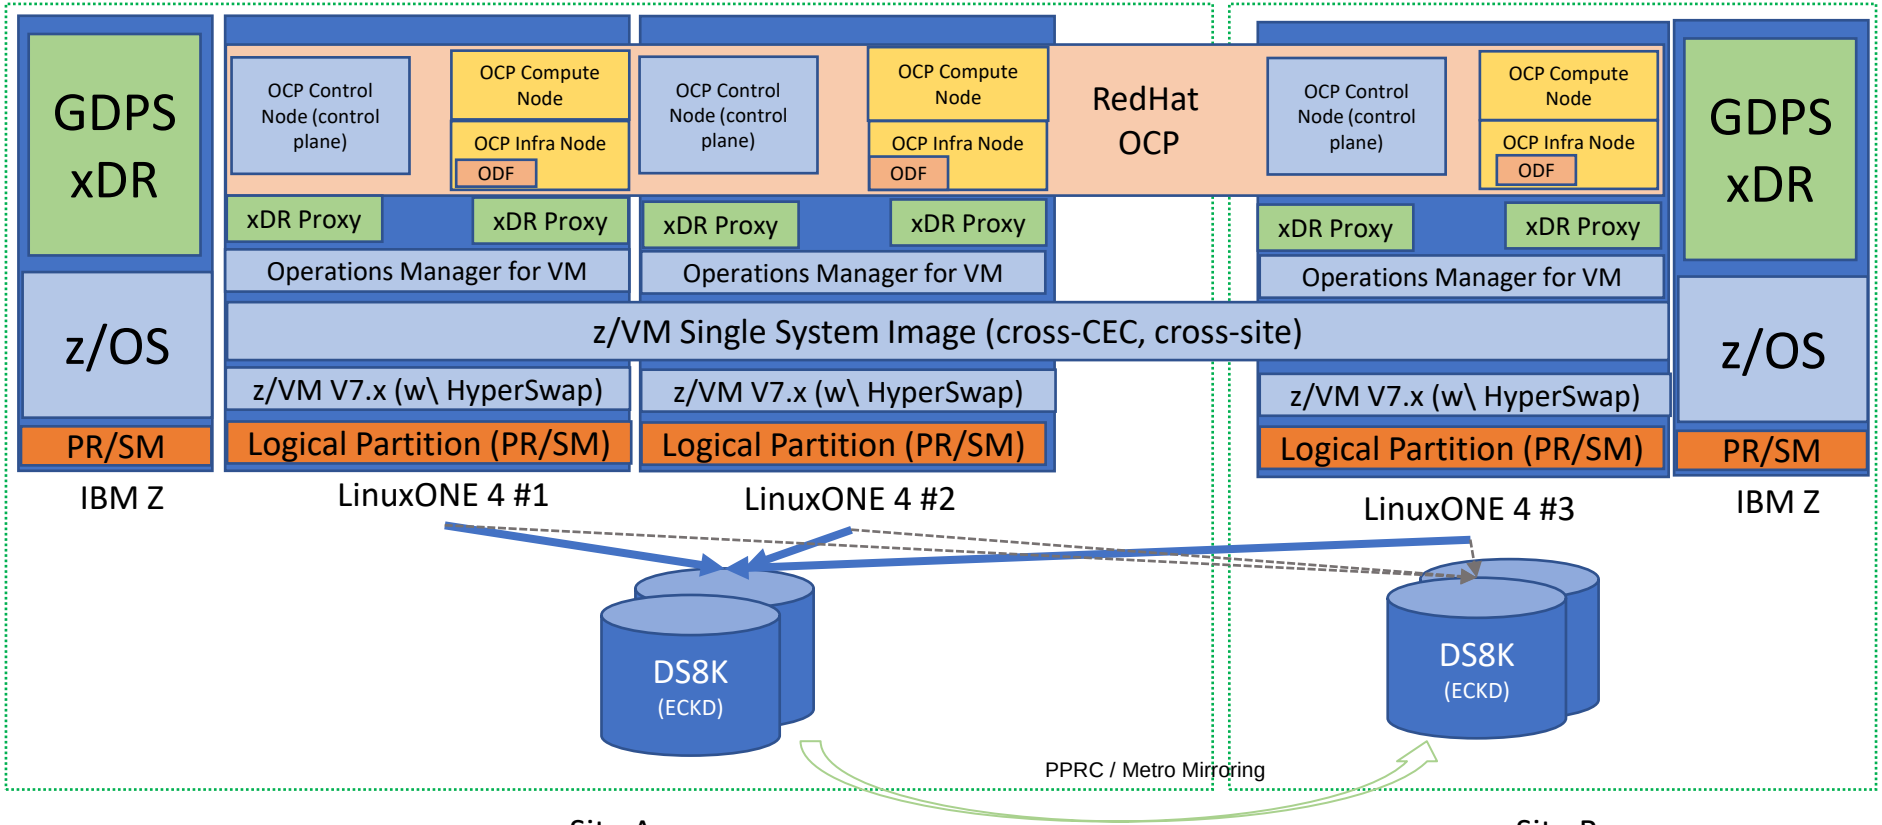

Redundancy at All Levels

A physical data center provides five 9's of availability

LinuxONE hardware is backed up at both component, feature, and CPC levels

Multiple z/VM partitions allow for guest mobility and planned outages

Operations Manager for z/VM keeps z/VM healthy and acts on problems as they happen

Red Hat OpenShift Container Platform balances workload and restarts applications as necessary

GDPS xDR proxies running on z/VM can relocate guests, restart virtual machines, restart containers, and allow for HyperSwap operations

And DS8000 series storage is itself rated to six 9's of availability

+ **LCP (via DS8K SafeGuarded Copy or GDPS)**

		1 Resilient/Reliable	2 Failover	3 Fault Tolerant	4 Continuous Availability
System Availability	Service Characteristics	Z Monitoring Suite IBM Service Management Suite IBM Storage Management Suite	IBM Infrastructure Suite	Z App Perf Management Connect Z Performance Capacity Analytics Load balancer/workload router	Cloud Pak for Watson AIOps GDPS Continuous Availability for Db2
	Compute	LinuxONE II/III with at least one CPC	LinuxONE II/III with at least 2 CPCs 2 - LinuxONE II/III systems z/VM Single System Image with redundant links	2 - LinuxONE II/III systems z/VM SSI with nodes on multiple CPCs across close data centers Optional 3rd system and data center	2 - LinuxONE II/III systems z/VM SSI with nodes on multiple CPCs across close data centers LPAR/system/data center for DB Optional 3rd/4th system/data center
	Data	RAID/mirroring within attached storage (SAN256/SAN512)	RAID/Mirroring + FlashCopy within data center (DS8K)	HyperSwap GDPS Virtual Appliance, GDPS Metro data center mirroring and FlashCopy (DS8K)	HyperSwap GDPS Metro, Metro Global, Global Cross data center mirroring and FlashCopy (DS8K)
Foun-dation		CBU, On/Off CoD, SRB, zBuRST, IZTA, IZBR			
Data recovery		TS4K Tape Storage iCoS (Cloud Object Storage)	FlashCopy (DS8K) TS7700 Tape Farm iCoS	Auto recovery from GDPS mirror FlashCopy from same or other data center (DS8K) TS7700 Tape Farm iCoS	Auto switch to running Database Auto recovery from GDPS mirror FlashCopy from same or other data center (DS8K) TS7700 Tape Farm iCoS

zBuRST

Accelerate adoption & deployment of key technologies while delivering quality levels of service

It’s time to take cost, complexity, and risk out of the testing equation.

With zBuRST, you can perform load and stress testing at scale to ensure maximum quality of service and business resiliency while minimizing cost and complexity.

Minimize the cost

Leverage a low-cost, end-to-end integrated testing and validation environment that mirrors your production systems.

Avoid business disruptions

Easily validate changes and upgrades to ensure robustness and minimize disruption in production environments.

Accelerate deployment

Increase testing efficiency to reduce change risk and deploy new technologies faster.

Maximize quality, ensure compliance

Simplify testing at scale to improve quality of application changes and avoid regulatory penalties.

How did you do?

- **What level of Business Continuity do you require, and do you fulfill that requirement end-to-end?**
- Can you remove people as Single Points of Failure?
- Can you reduce recovery times using accurate and automated processes?
- Can you stress test innovations at production scale?
- Can you optimize recovery processes?
- How do you verify accuracy and viability of your recovery?
- What does it cost you to fail?

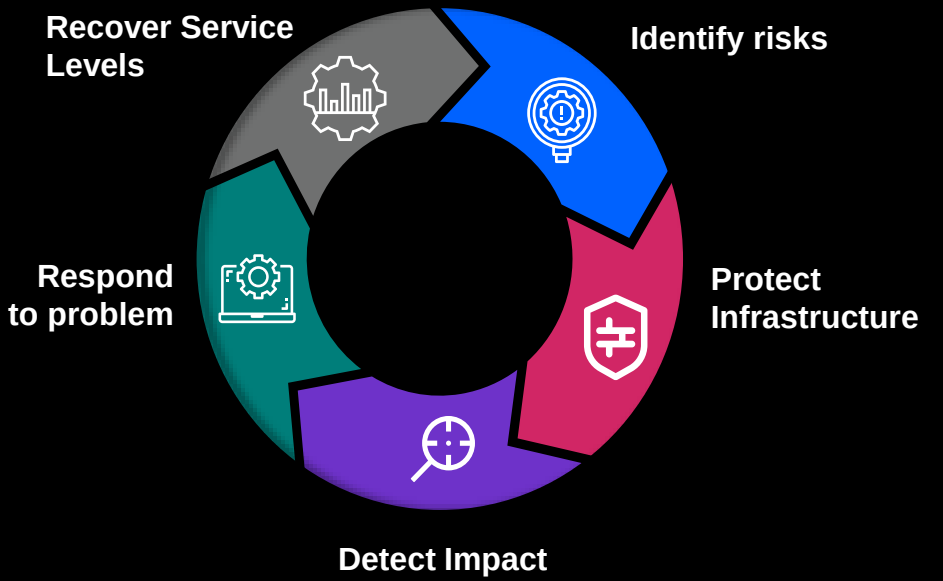

Thank You!

Thanks for this content to:

Brian W. Hugenbruch, CISSP
IBM LinuxONE Resiliency Lead
bwhugen@us.ibm.com