Up Your Resiliency Game With Feedback from WSC

Share Orlando 24 February 2026

Meral Temel
z/OS Technical Engagement Leader
z/OS, Sysplex, IBM Z Resiliency & Performance

WSC IBM Z Resiliency Solutions CC Track leader
Member of z/OS Development Resiliency core team, IBM ZStack Resiliency team
WSC IBM Z Resiliency, IBM Z HW and z/OS Performance, z/OS and sysplex SME

IBM Z Worldwide System Center (Former Washington Systems Center)
Poughkeepsie NYUSA
Meral.temel@ibm.com
SHARE roles: SHARE z/OS System Programmer IBM Co-Track Leader ,SHARE Performance Track Leader

WSC IBM Z Resiliency Workshop

Around 60 Resiliency focus client engagements in 2.5 years in different service formats

- ❑ IBM Z Resiliency Workshops
- ❑ Standalone Infrastructure, z/OS and sysplex Health Check (combined with IBM Z Resiliency workshop)
- ❑ IBM Z and z/OS performance analysis (sometimes combined with above all)
- ❑ WSC-wide HC – Infrastructure,z/OS and sysplex part
- ❑ Architectural guidance

and other specific topic-oriented engagements which are not counted in are:

IBM Z HW implementation guidance from z/OS perspective, z/OS and sysplex updates, z/OS upgrade,z/OS components, sysplex new and old features' guidance, GDPS updates, Performance and Capacity management guidance...

WSC IBM Z Resiliency Workshop

WSC IBM Z Services
Quick Reference Guide

IBM Z Resiliency workshop

Overview

- **If you need to improve resiliency for your critical IBM Z business services, want to have an architectural review of your current environment and learn about existing IBM z Resiliency Solutions, WSC SMEs will help you design the architecture you need to meet your current and future business requirements and create a roadmap to move from your current environment to the future architecture.**

Contacts

Meral Temel
meral.temel@ibm.com

Miao Zhang-Cohen
miaozha@us.ibm.com

Worldwide Systems Center
IBM Z

Target Audience

- IBM Z clients interested in leveraging IBM Z resiliency solutions, best practices to enhance resiliency.
- IBM Z clients seeking an analysis of their platform's resiliency across architecture, processes, and other critical areas.

Why Use This Service?

- Worried about the financial and operational impact of service outages on your business?
- Need your critical business services running 24/7 without interruption?
- Want to eliminate planned outages and keep your systems continuously available?
- Unsure if you could fully recover from a data center disaster or unexpected outage?
- Looking to align your infrastructure policies and practices with industry-leading standards?
- Ready to unlock the full power of IBM Z Stack capabilities to achieve your availability goals?
- Interested in discovering IBM Z resiliency solutions—including GDPS, Cyber Vault, architectural best practices, z/OS, Sysplex, hardware, software, storage, and other resiliency features—that can help you maximize uptime and protect your business?

Service Provided

This is typically a 2–3 day onsite workshop. Our SMEs will review the client's architecture, ask questions, and share recommendations to improve platform resiliency. They will explain why these recommendations impact resiliency, using real-world examples.

Review of IBM Z platform architecture: hardware, LPARs, storage, replication, and tape environment. Includes workload distribution of LPAR architecture by business channels.

IT process best practices: recovery, operations, monitoring, capacity planning, performance management, risk management, and how to achieve/report business SLAs using IT data.

Functions, capabilities, and processes that help clients avoid outages and recover quickly.

Information on resiliency solutions such as GDPS, Cyber Vault, storage functions, IBM Z hardware, z/OS and sysplex capabilities, and AIOps solutions tailored to the client's environment.

Collaborate with the client to define current and future organizational and regulatory requirements, along with IT objectives that support them.

Develop a roadmap or target architecture based on industry and technology best practices to support current and future resiliency objectives.

Note: above tasks are representative and may be customized based on specific workshop focus

Deliverables

- **Comprehensive analysis** of the current platform environment to identify gaps, vulnerabilities, and resiliency risks.
- **Strategic recommendations** to enhance platform resiliency, tailored to the client's environment.
- **Detailed report** including IBM z Resiliency solution capabilities, supplemented by insights shared during the workshop.
- **High-level vision and target-state architecture** designed to meet both current and future resiliency objectives.
- **Roadmap documentation** outlining the implementation strategy for the target architecture and associated resiliency configurations.

Benefits

This workshop helps clients gain a clear understanding of IBM Z Resiliency Solutions and provides tailored recommendations, supported by a list of findings, to address any gaps and ensure improved platform resiliency.

Understanding what ‘Resilience’ mean

Case 1) **‘I want to improve my IBM Z Resilience, from where I can start ?’**
Case 2) ‘**I had an outage**, I want to improve my IBM Z Resilience, from where I can start ?’

Our goal is to do **Case 1** ,not second!.

Definition of Resilience

The ability of an information system to continue to:
1) Operate under adverse conditions or stress, even if in a degraded or debilitated state, while maintaining essential operational capabilities
2) Recover to an effective operational posture in a time frame consistent with mission needs. **(NIST SP 800-39)**

Oxford dictionary
Reliability : The quality of being trustworthy or of performing consistently well
Resilience is "**the capacity to recover quickly from difficulties**"

Merriam-webster dictionary
Reliability: the quality or state of being reliable (suitable or fit to be relied on : dependable, giving the same result on successive trials)
the extent to which an experiment, test, or measuring procedure yields the same results on repeated trials
Resilience:

the ability of something to return to its original size and shape after being compressed or deformed
an ability to recover from or adjust easily to adversity or change

Definition of Resilience

WSJ (How to create IT Resilience article)
Simply defined, IT resilience is an organization's ability to maintain acceptable service levels, no matter what challenges arise. Because so many things can cause service disruptions —hardware failures, natural disasters, malicious attacks, and operator error, among others—creating effective disaster recovery (DR) and business continuity (BC) plans to address vulnerabilities can be daunting

Redundant: Serving as a duplicate for preventing failure of an entire system (such as a spacecraft) upon failure of a single component
(Removing SPOFs! - Single point of failures)

Definition of Business Continuity and Business Resilience

Business resilience is a company's ability to overcome an unexpected business disruption and recover to **acceptable continuing operations**.

Business continuity is a set of processes and procedures that enable companies to keep running their business during a crisis.

Example:
Continuity ensures that your business can function if a natural disaster strikes your area. However, business resilience plans for continued success **even after acute events**.

*Business continuity would enable a company to function in the event of an earthquake. Business resilience takes **continuity to the next level**. Instead of a set number of processes, it is a framework integrated into the all company culture.*

Achieving Resiliency : Mainframe System Programmers' Guide as Approach

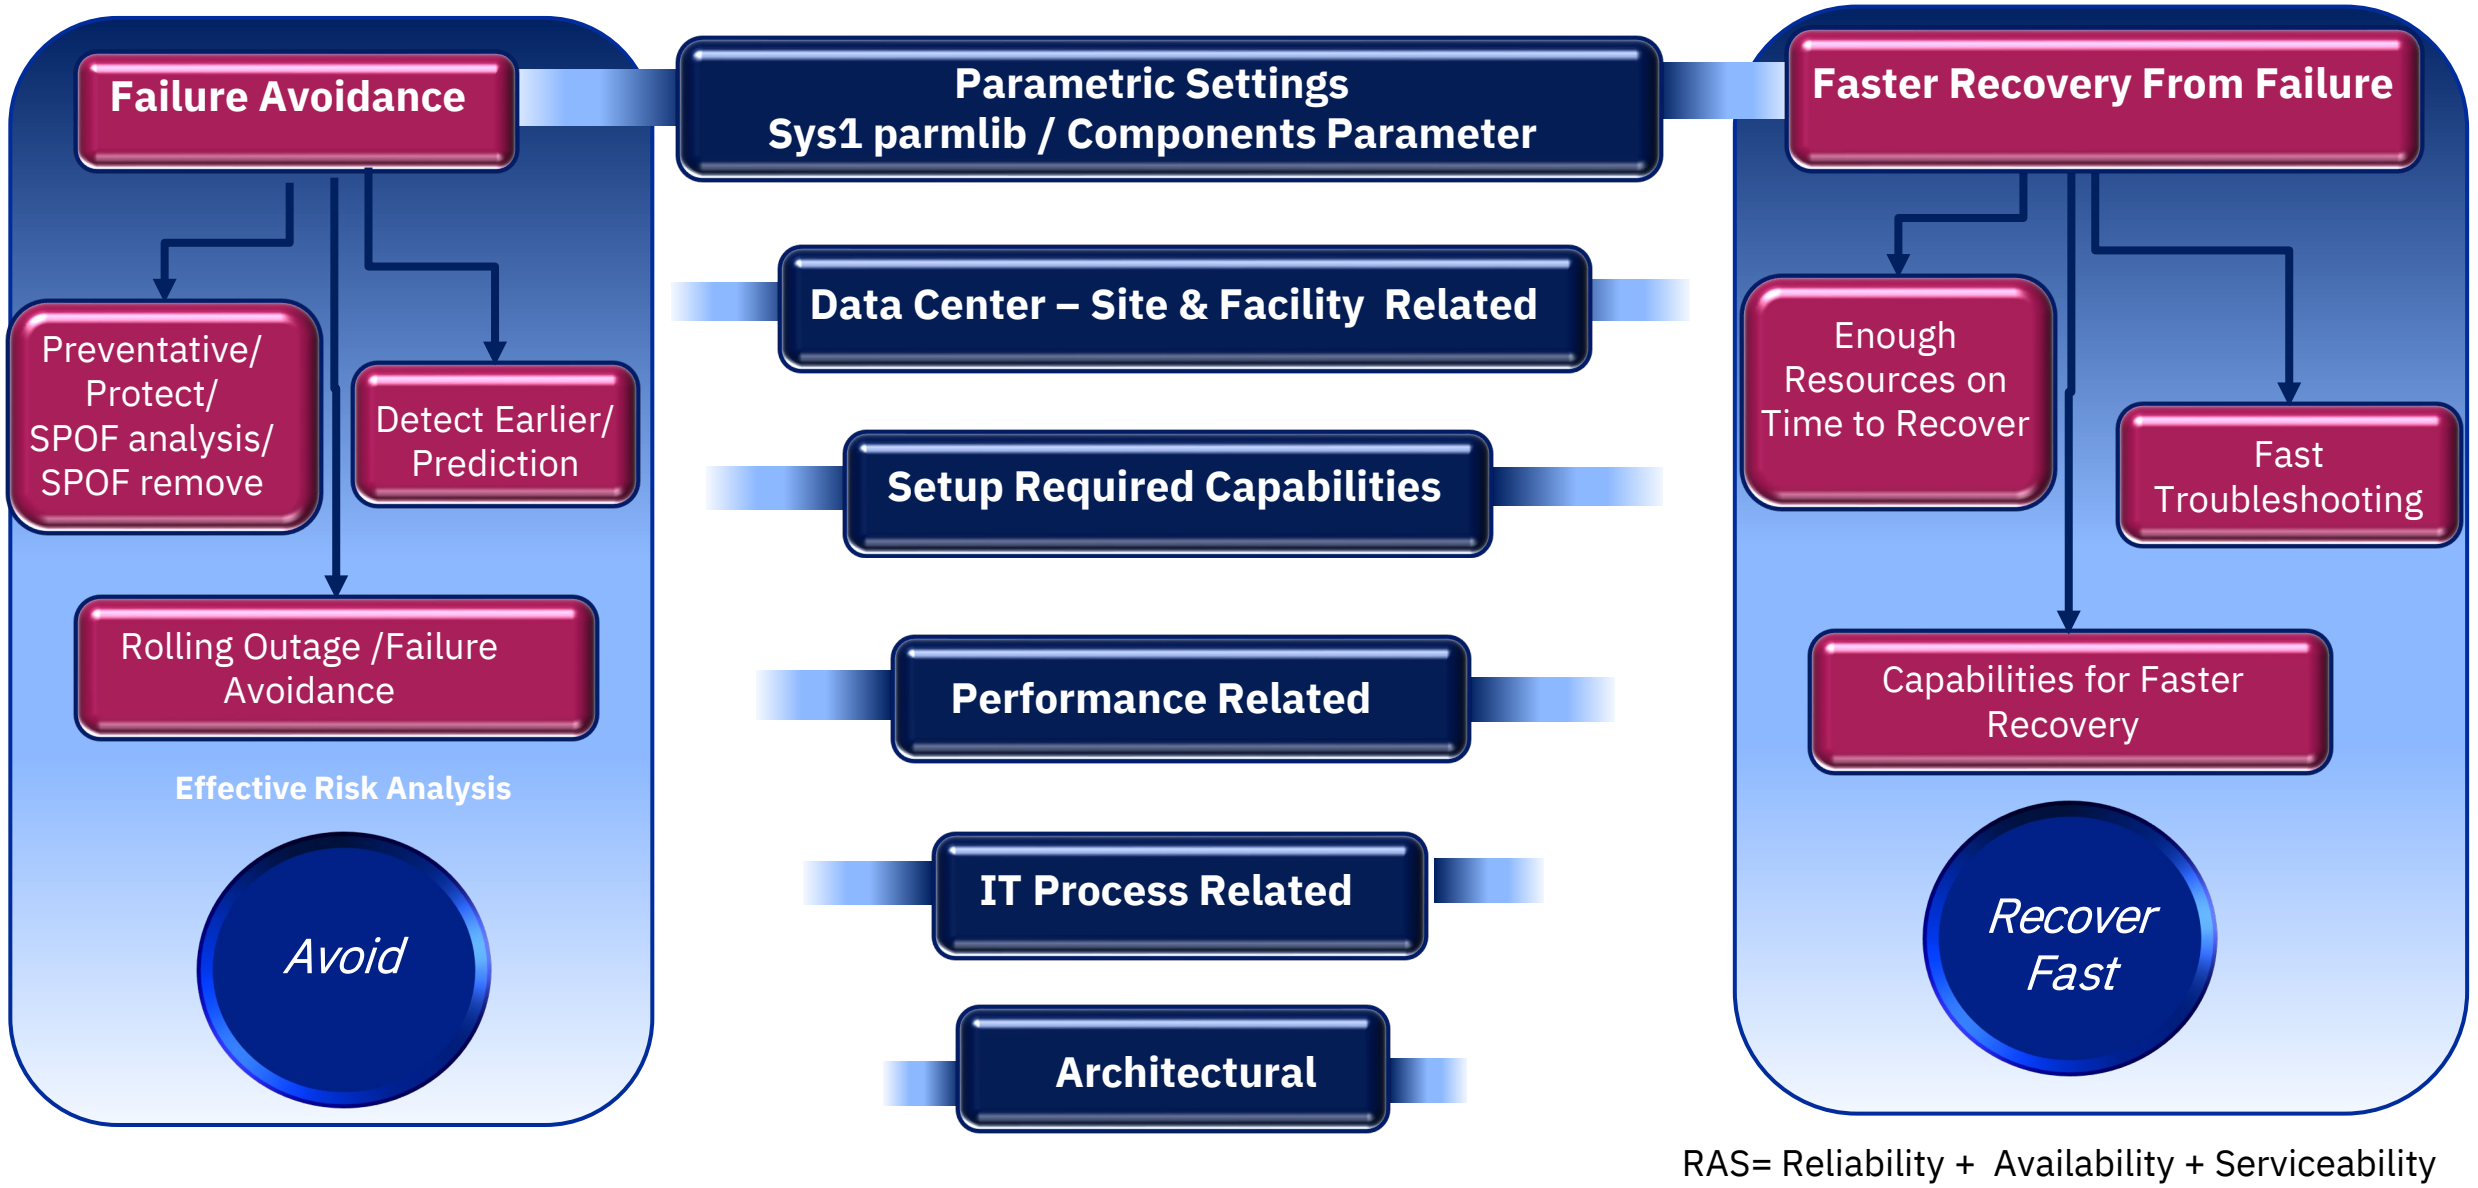

RAS= Reliability + Availability + Serviceability

5 Pillars of Resilience

The 5 pillars of resilience are:

- ✓ Prepare/identify,
- ✓ Protect,
- ✓ Detect,
- ✓ Respond
- ✓ Recover

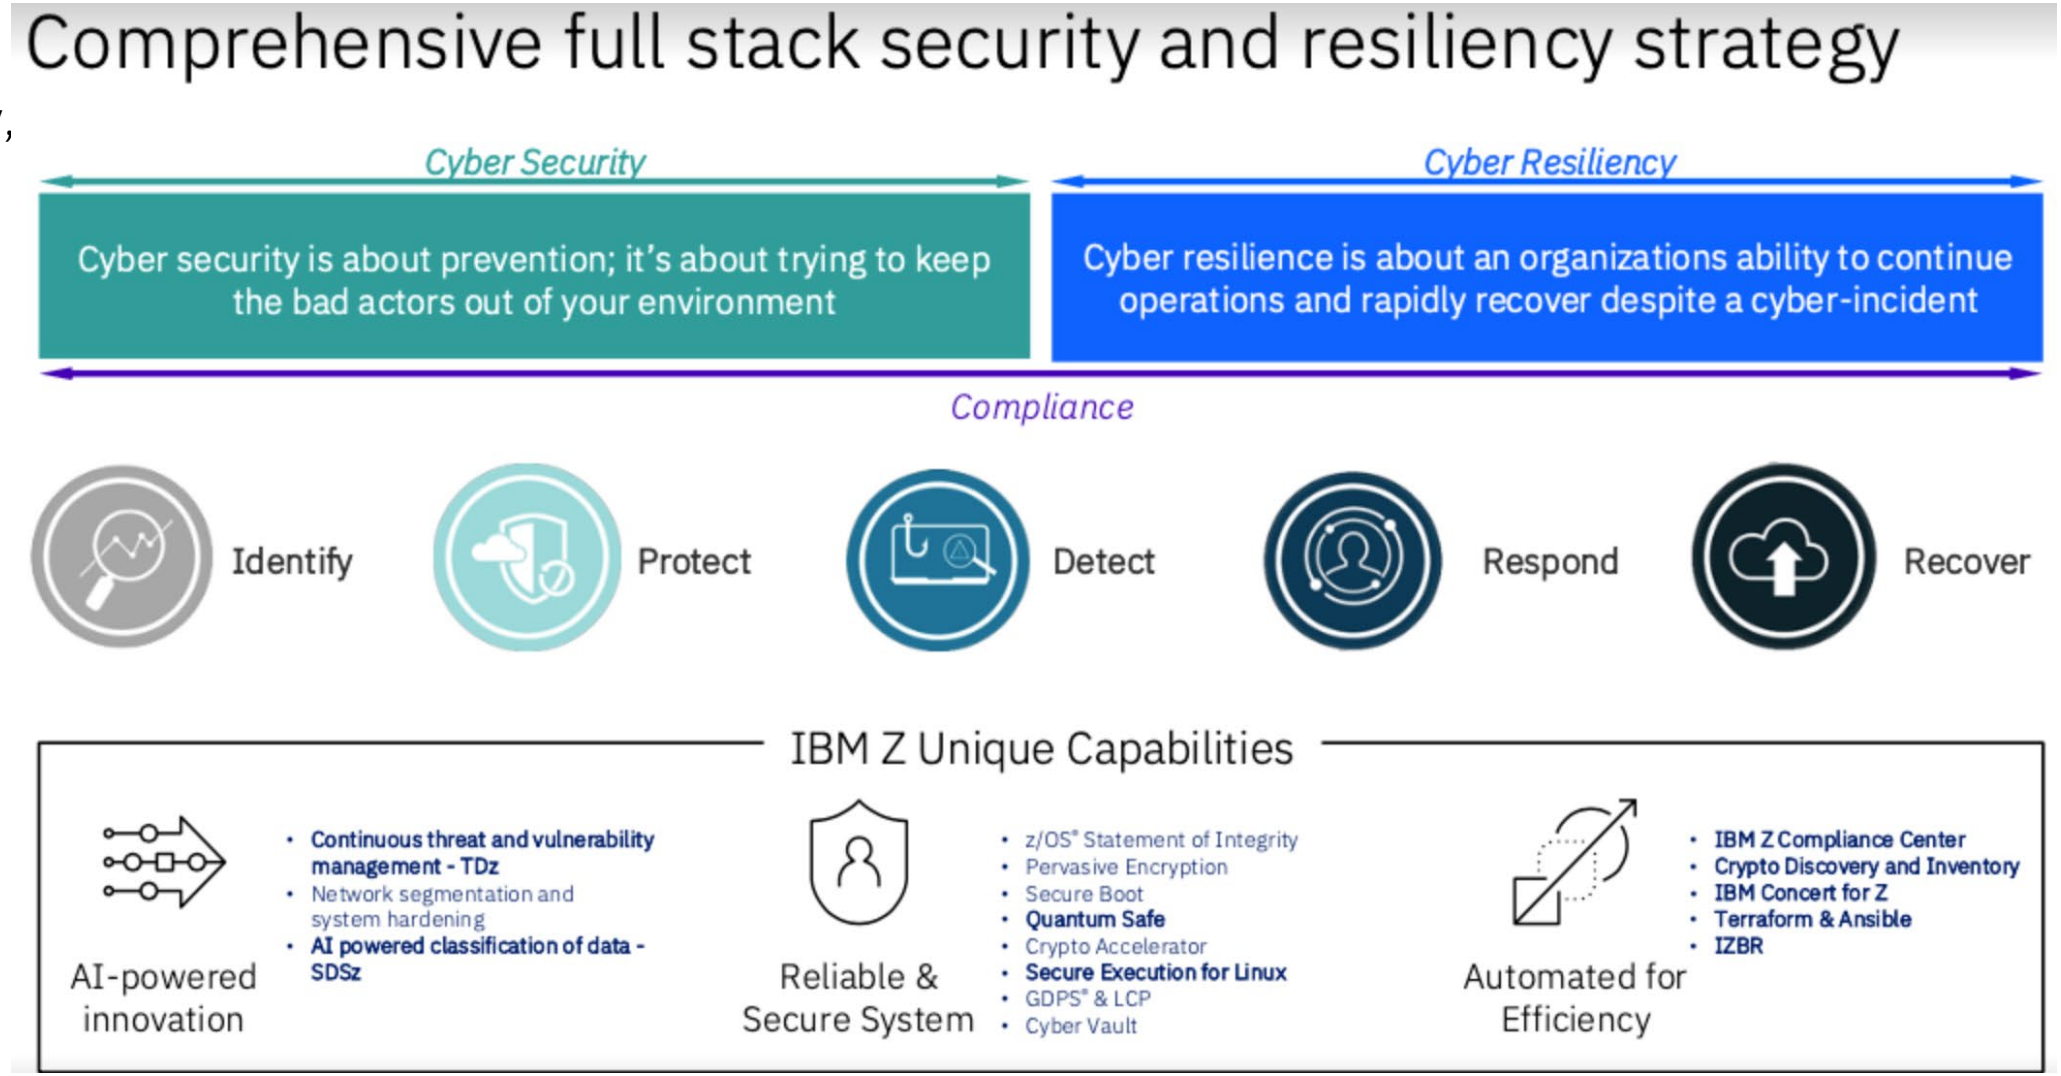

IBM Z – Leading the industry in IT resiliency and integration

Avoiding the cost of downtime

Ensuring access to critical apps / data

Maintaining productivity of users

Open to clients 24/7

IBM Z

Where mean time between failure is measured in decades

Parallel Sysplex

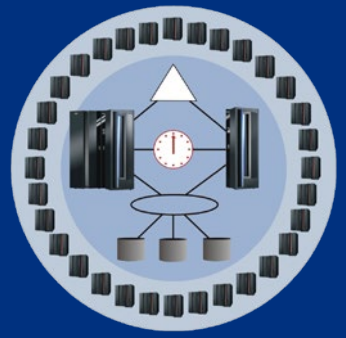

Designed for *application* availability of 99.9999999%

GDPS

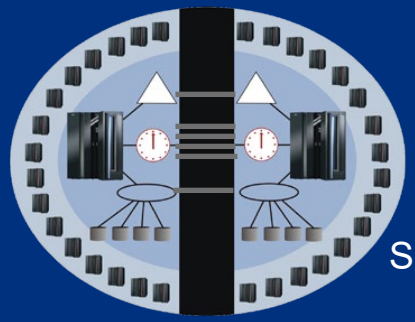

Industry leading solution for continuous availability / disaster recovery

Storage Synergy

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

IBM Z Cyber Resiliency Solution (Cyber Vault)

Cyber Security & Cyber Resilience

Cyber Security

Cyber security is about prevention; it's about trying to keep the bad actors out of your environment

Cyber Resilience

Cyber resilience is about an organizations ability to continue operations and rapidly recover despite a cyber-incident

Organizations need to be both cyber secure and cyber resilient

Worldwide regulation

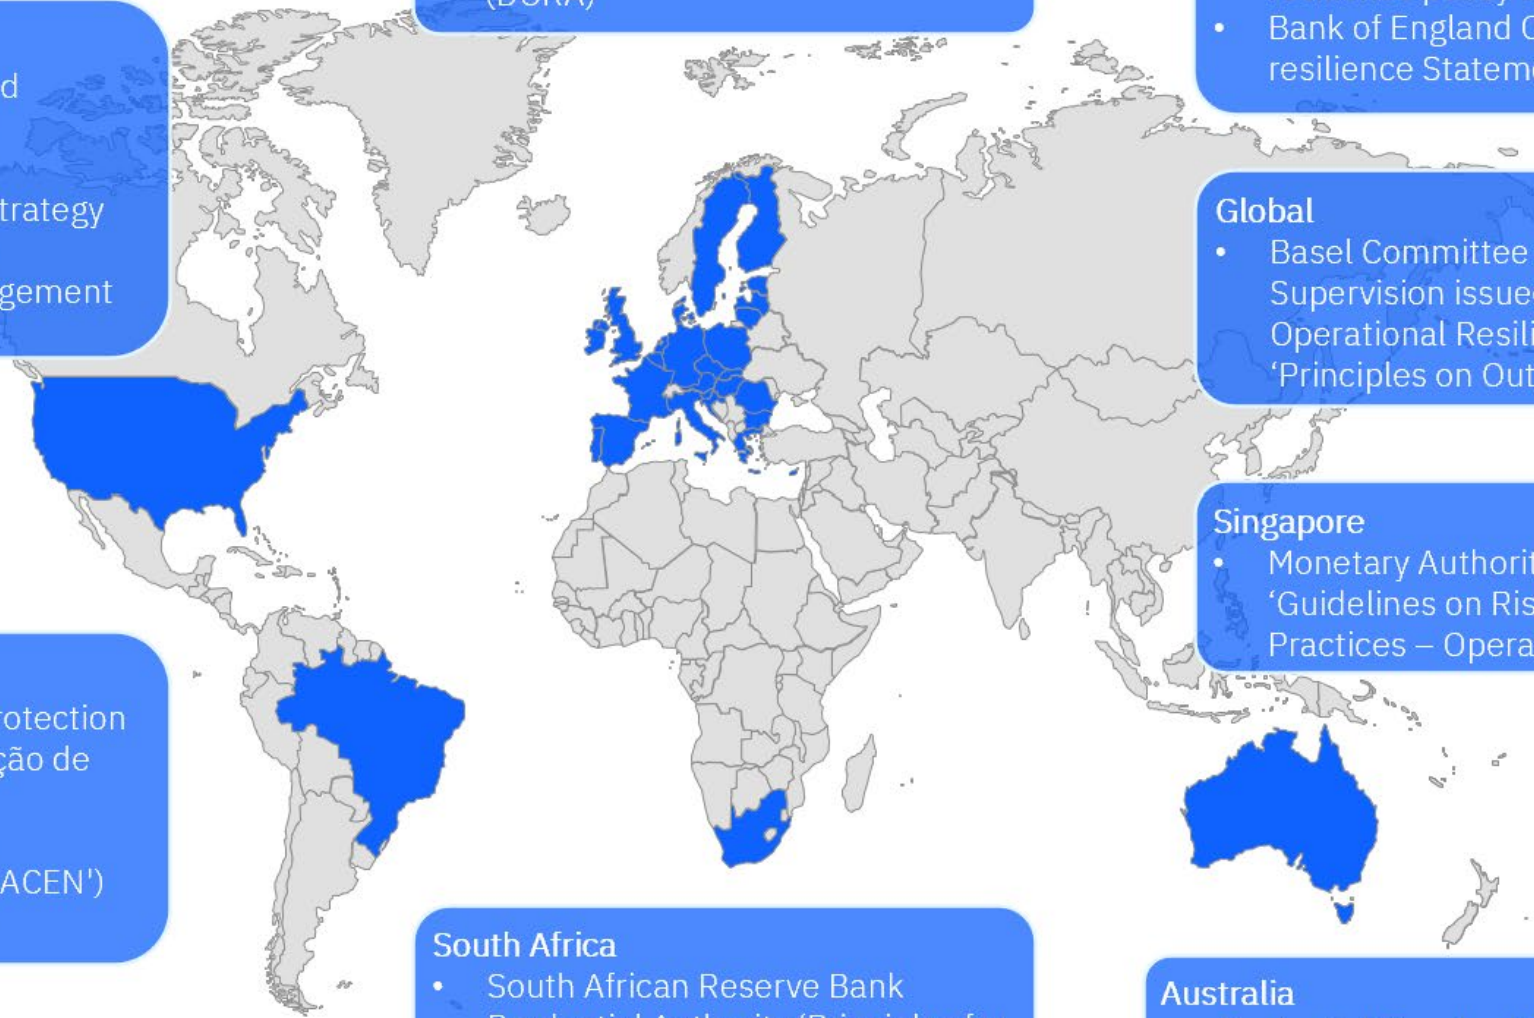

United States
- Interagency paper 'Sound Practices to Strengthen Operational Resilience'
- National Cybersecurity Strategy
- SEC Proposed Ruling for Cybersecurity Risk Management Rule 10

Europe
- Digital Operational Resiliency Act (DORA)

United Kingdom
- FCA PS21/3 Building operational resilience policy statement
- Bank of England Operational resilience Statement of policy

Global
- Basel Committee on Banking Supervision issued 'Principles for Operational Resilience' and 'Principles on Outsourcing'

Singapore
- Monetary Authority of Singapore 'Guidelines on Risk Management Practices – Operational Risk'

Brazil
- Brazilian General Data Protection Law ("Lei Geral de Proteção de Dados" or "LGPD")
- Resolution 4.502/2016
- Central Bank of Brazil ('BACEN') Resolution 4.893/2021

South Africa
- South African Reserve Bank Prudential Authority 'Principles for operational resilience'

Australia
- Prudential Standard CPS 230 - Operational Risk Management

Infrastructure Best Practices

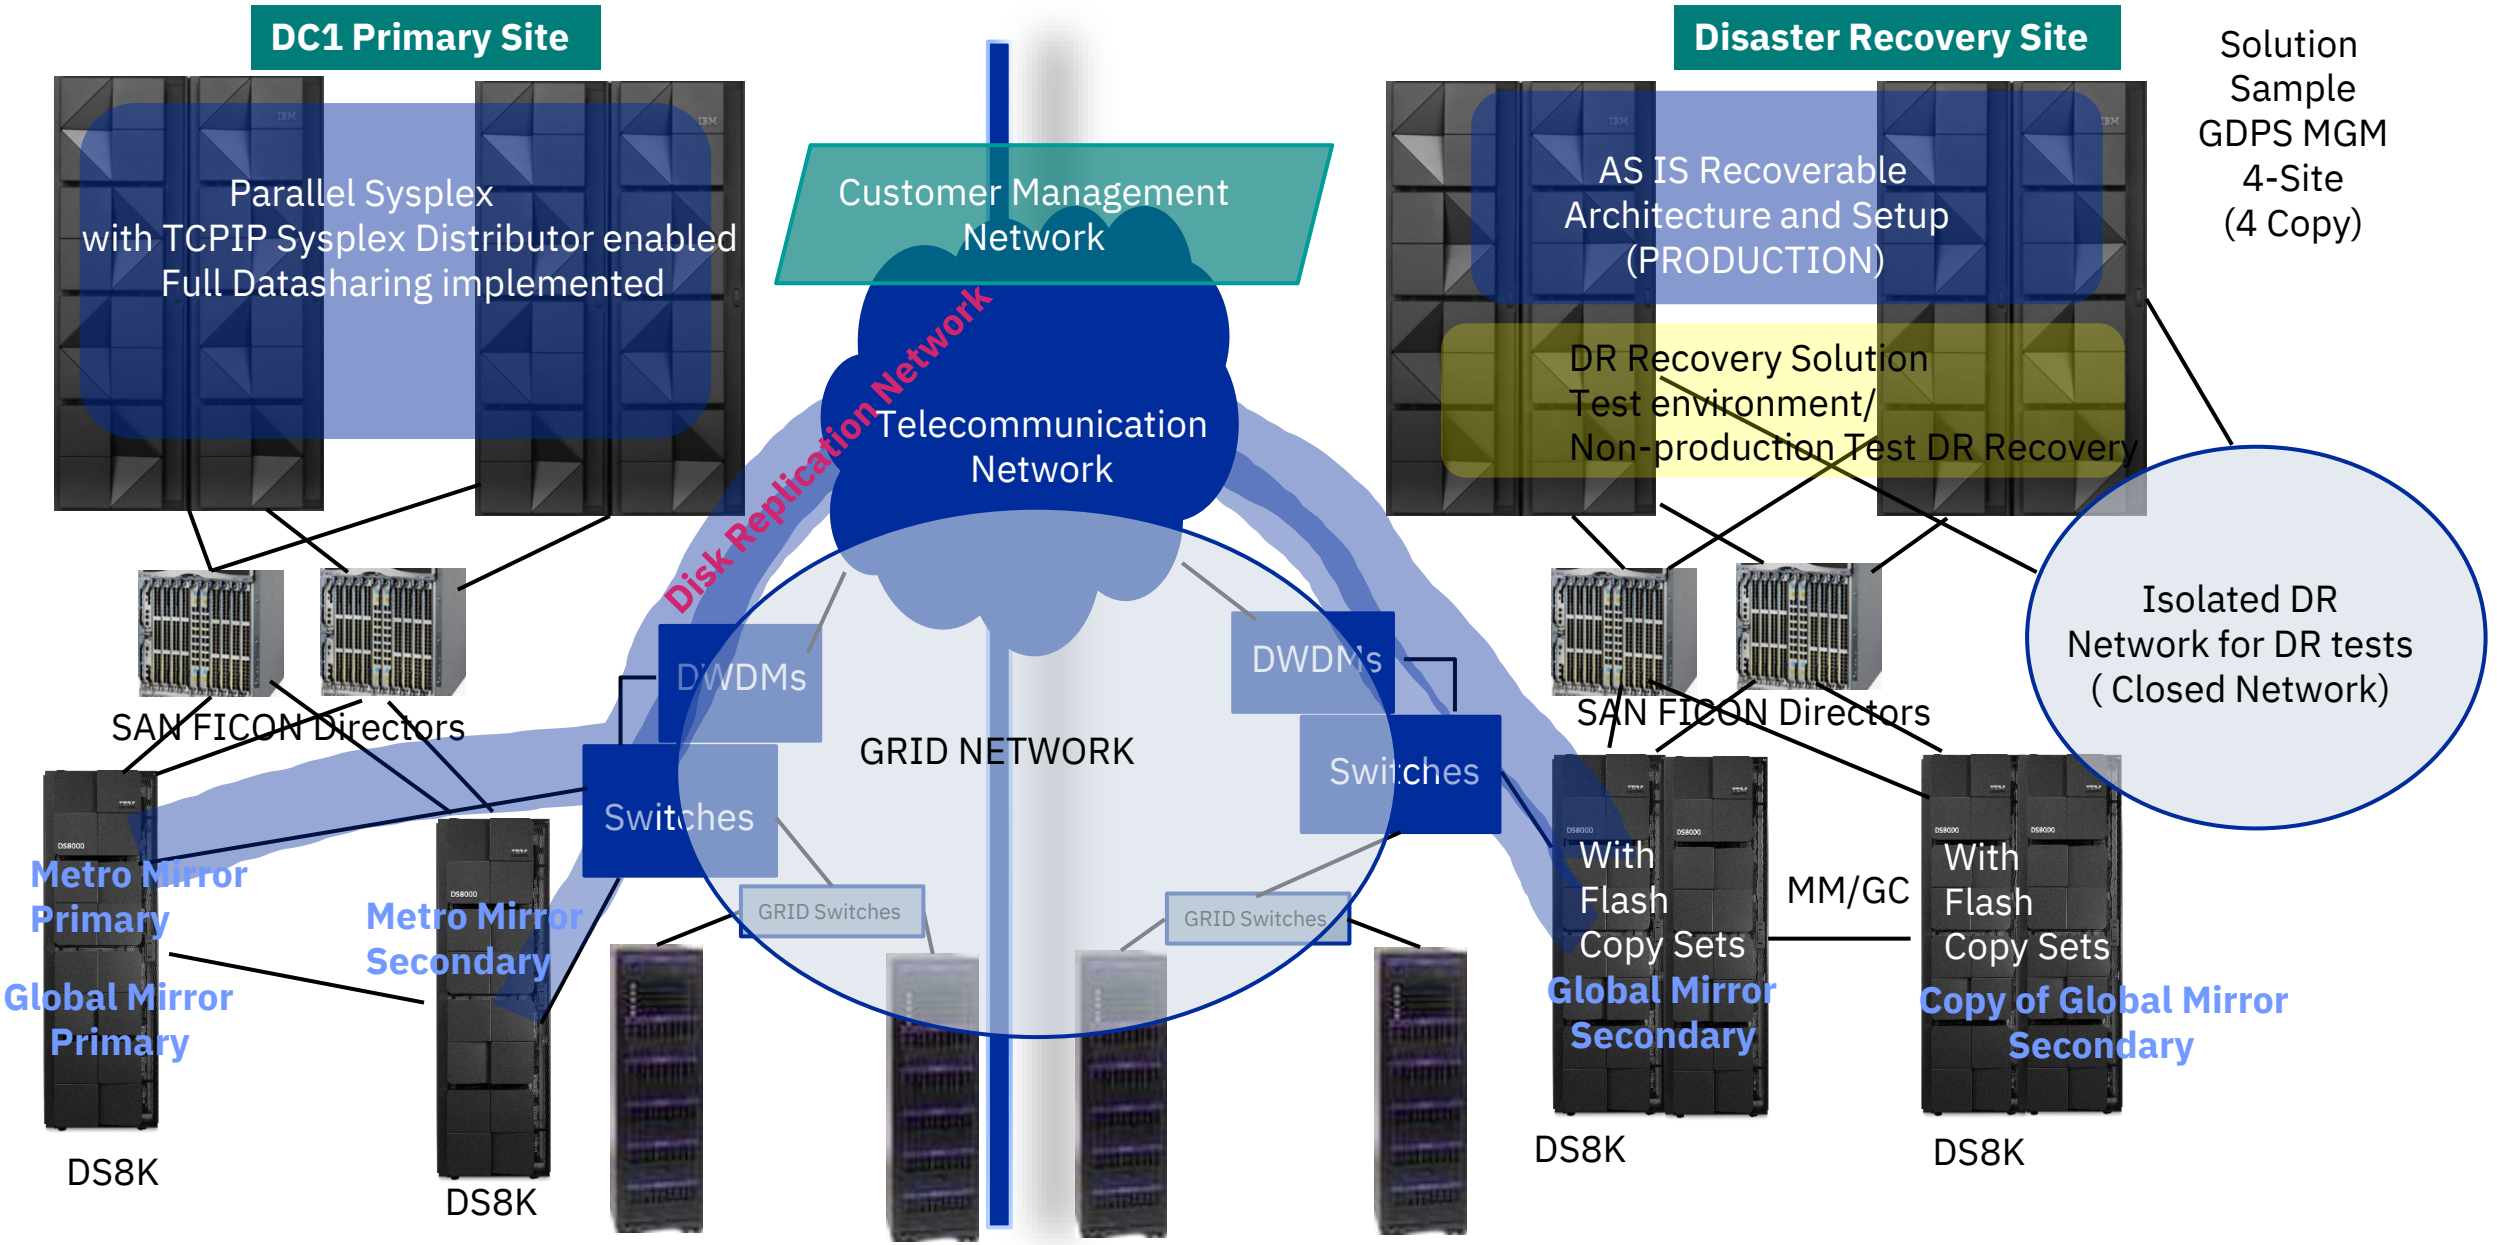

Infrastructure Items to consider – Best Practices

- **DC best practices** :
  Different rooms, tier level, operational maturity level, HW power connections,....
  Resilient DR Datacenter – GDPS 4-site, all lpars recovered
  Periodic tests, see next slide for specific to DR-site

- **HW management**
  No SPOF related to any HW box,
  Up-to-date HW models,
  Maintenance strategy
  Review HIPER MCLs biweekly

- **LPAR level (z/OS level)**
  Sysplex enabled , ALL applications support sysplex, all sysplex capabilities implemented,
  Spare LPARs in each CEC...
  Implementation of resiliency related functions that prevents z/OS /LPAR failures
  Leveraging sysplex distributor

Highlights of IBM Z Stack Innovations in Resiliency

Faster recovery when failures occur

- System Recovery Boost(All functions,including MiddleWare stop/restart support, Dump support)
- z/OS Automatic Restart Manager(ARM)
- Sysplex Failure Management,CFLRMGMT XCF Function
- System Status Detection Partitioning Protocol
- External CF, CF System/User Managed Duplexing
- zHyperLink (VSAM Verify) , DFHSM Dataset (z/OS & zFS File) Backup
- RLS Catalog , PDSESHARING(EXTENDED)
- z/OS AutoIPL
- IBM Z Backup Resiliency, DB2 advance recovery,IMS FDBR
- Flexible Capacity,CBU
- DVIPA,Sysplex Dynamic Workload Routing (Sysplex Distributor...)
- CICSPLEX,DB2,IMS Datasharing,MQ Queue Sharing
- z/OS IOS CU Recovery
- Z16 CF Improved IFCC Handling
- IBM Z Cyber Vault

Failure Avoidance

- z/OS Blocked Workload Support
- GRS Contention Analysis
- XCF Message Isolation, CFMONOPAVOID
- Message Flood Automation
- SMF Flood Automation
- Capacity Provisioning
- XCFCNFXX- LOGR Services Monitor
- Auto-WTOR Response
- IBM TST, IBM ZBURST
- Cyber Resiliency,
- IBM Thread detection for z/OS(TDz)
- IBM Concert For Z
- TFP-HW
- HMC LPAR Error Protection
- Z17 HMC Dual Control

React faster to workload fluctuations

- zHyperLink,zHyperwrite,DFSMS Striping
- FICON 32S
- Capacity Provisioning Manager
- TFP-HW, OOCoD
- Hiperdispatch
- More Memory (Larger Buffer Pools, GBPs…)

Improved workload scaling

- Hiperdispatch • Shared zFS • Z17 Performance improvements
- zHyperLink, zHyperwrite,Read from Secondary, zEDC
- JVMS pause-less garbage collection
- FICON 32S,z17 CL6 Throughput improvements
- Large Page Support, RLS Catalog
- CICPLEX,DB2,IMS Datasharing, MQ Queue Sharing
- XTCSIZE, zOS V1R10 VSM algorithm improvements,V1R8 RSM
- Sysplex Distributor, VSAM RLS 2 Lock Structures
- z16 CF Adapter Performance Improvements
- z16 CFCC Code Scalability Enhancements
- DFSMS Extended Addresability, DFSMS Striping

Concurrent maintenance

- Software coexistence and rolling IPLs
- z/OS concurrent maintenance, Fixcats
- Dynamic I/O Reconfiguration
- HyperSwap

Problem Determination and Data Capture

- Run Time Diagnostics(RTD)
- Dump cmd Parmlib members (IEADMCxx)
- AutoIPL w/Standalone Dump
- z/OSMF Incident Log
- SMF Logstreams (Performance & Scalability)
- zWorkload Information Correlator (zWIC)
- zWorkload Information Navigator (zWIN)
- New SMF Records & Fields

Preventative / Real-Time Insight

- z/OS Health Checker
- Predictive Failure Analysis (PFA)
- IBM Z Anomaly Analytics (IZAA now part of IBM Concert for Z)
- IBM Concert For Z
- IBM Z Operational Log and Data Analytics
- Heartbeat Monitoring – Omegamon, RMFMON, Intellimagic, OpenTelemetry

And many more…

Common Issues : Disaster Recovery Site gaps

❑ DR site gaps

- Importance and value of Disaster Recovery Site not well understood. (It is not there for just natural disasters or if DC fails)
- It is not a task that you have to do just to say ‘ Yes, I have DR, yes I am doing tests’. It is very valuable for team and company not only related to ‘natural disasters’ or DC site failures.
- There are clients who tried to move to DR site not related to natural disaster or Primary DC outages!.

- When you can not recover on primary site for certain time for any reason, you may go for DR site to continue serving your business. It is not about ‘natural disasters

- DR tests are great opportunity for teams to experience on recovery actions , speeds up troubleshooting on primary site, decreases locally recoverable outages. Flexible capacity solution gives you chance to do more tests and stay longer!.

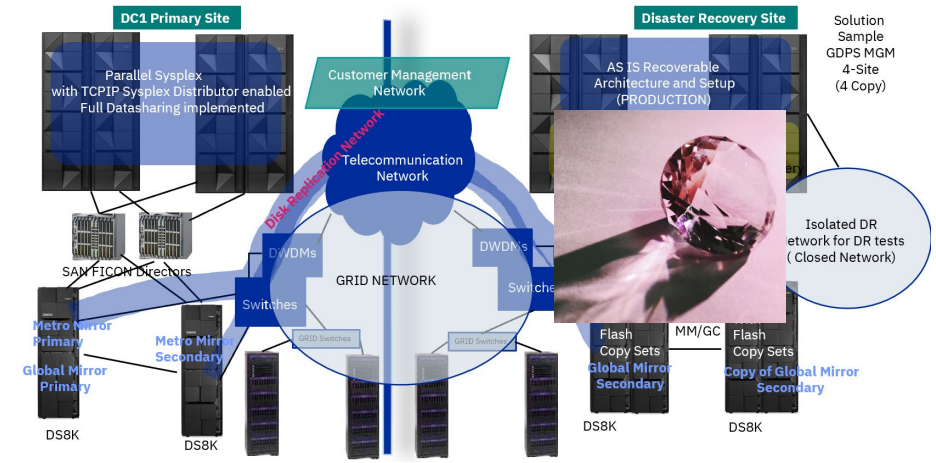

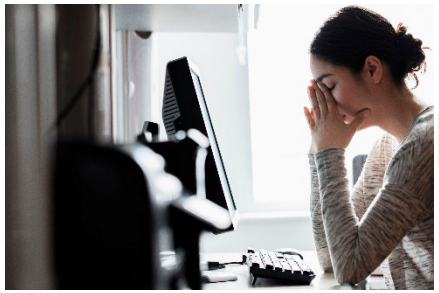

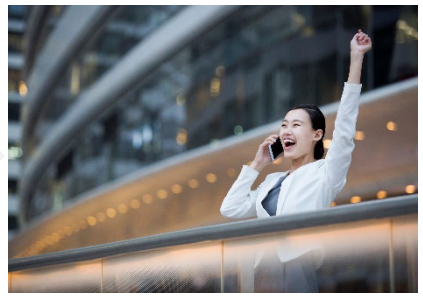

Flexible Capacity Enterprise and LiTe Edition vs. CBU

Today's Scenario:

- Ever-changing and complex regulatory environment
- More complex and longer periods of tests

Offering -> Available with z16 and newer models (Z17)
Protects investments made into CBU
Credit the remainder of a CBU contract period towards a Flexible Capacity solution
Creates an 'upgrade path' to Flex Capacity for IBM z16 Clients

More tests, longer duration

Description	Flexible Capacity (EE or LiTe)	CBU
Customer controlled	Yes	Yes
GDPS Enabled	Yes	Yes
IBM intervention required (On-Site)	No	No
Intersite and Intrasite activation	Yes (EE) / Intersite only (Lite)	Yes
Activated by	**One Flex Cap Record per serial#**	1 or more CBU Records / serial#
Replacement Record Limit	Up to to available machine capacity	Up to to available machine capacity
Total activated capacity	Up to owned capacity across all machines	Up to owned capacity across all machines
Activations/Deactivations per year	**12 (EE) / 4 (LiTe)**	DR and DR Test only
Maximum Stay out period	**12 months (EE) / 30/90 days (LiTe)**	10 days (Test) / 90 days (DR)
Use cases	Real Disaster Recovery / DR Test **Facility Maintenance** **Compliance Testing** **Proactive DR avoidance**	Disaster Recovery DR Test
Technology requirement	IBM z16 + TFP SW / LinuxOne 4	Systems Not WdFM
Tiered pricing	**Yes**	No
Carry forward	**Yes (N-1/N-2)**	When upgradeable

Common Issues : Disaster Recovery Site gaps

❑ DR site gaps -Continue

- No second sync copy (GDPS 3-site)
- Not all LPARs are recovered
- Not enough DR tests, CBU and flexible capacity tests are there waiting to be used!
- Closed network approach is not used
- Timeline of Disaster Recovery tests is either not monitored, or in high level. (Phases of actual recovery time)
- Projects to improve RTO has not been conducted or planned.
  HW, SW, z/OS, subsystem functions that will help decrease the RTO both in DR site and also help on any kind of local recoveries in primary site.
- DR availability not reported, monitored, tracked. Until you solve the issue, your RPO (Recovery Point Objective) continues to increase.

Common Issues : Disaster Recovery Site gaps – VTS Replication

❑ DR site gaps -Continue

- Selective data replication for VTS. Not all backups are archives are send to VTS in DR site

- It is important to send all backup and archive data to DR site. Client who use ‘selective data send to DR site’ mostly have issues related to either ‘ DFSMS dataset naming standards were not enforced ‘ or human error in definitions cause the important backup not to be send to DR site.

- This is important especially for clients where they did not setup yet their DFSMS ACS routines to enforce certain naming standards for Tape datasets. Ex: Is application developer able to create non standard tape datasets and if your ACS routine allows it , it may peak wrong management class which has defined in VTS as ‘do not copy’ to DR VTS. When you need those in DR site , you may not find your backup . Backup and archive data is as much important as primary data.
- So best approach is send all backups and archives – all VTS-tape data to DR site.

- There are clients who experienced loss of backup/archive data in DR site sadly. They could not reach to backup in a time where they desperately need it.

Importance of Data Center Properties & Capabilities & Tier Certifications

- ❑ Datacenter capabilities impacts the Resilience Level significantly
- ❑ Power outages/ wrong connection/SPOF UPS issues/cooling problems
- ❑ Cabling problems because of mature cabling architecture
- ❑ Tier level
- ❑ Operational Tier Level
- ❑ Same room , same corridor,....

uptime INSTITUTE

Tier I	Tier II	Tier III	Tier IV
Basic Capacity	Redundant Capacity Components	Concurrently Maintainable	Fault Tolerant
Site-wide shutdowns are required for maintenance or repair work. Capacity or distribution failures will impact the site.	Site-wide shutdowns for maintenance are still required. Capacity failures may impact the site. Distribution failures will impact the site.	Each and every capacity component and distribution path in a site can be removed on a planned basis for maintenance or replacement without impacting operations. The site is still exposed to an equipment failure or operator error.	An individual equipment failure or distribution path interruption will not impact operations. A Fault Tolerant site is also Concurrently Maintainable.

- Datacenter maturity level related outages are very common although they can be prevented
- Actions are done after an outage in several cases.
- Mainframe IT technical teams may not be in authorized persona to decide on big changes but they can definitely raise the risks and may trigger a change

Importance of Data Center Properties & Capabilities

Sample cases:

- UPS problem -> disk outage -> rolling outage
- UPS problem , no GDPS , long outage
- UPS problem, single CEC, outage
- Wrong wired power connections (HW CE rely on DC operation team’s wrong input)
- Cabling activity related to separate environment, causing IFCCs and connection issues in IBM Z platform because of not ‘good’ cabling framework /architecture in DC
- Wrong CF cables pulled out at the same time– sysplex outage
- Primary and secondary disks on same room in same DC.
- IBM Z HW in same room in same DC.
- FICON Directors in same rack, same room ,same DC
- HW maintenance – missing – explosure to known issue

Common Issues : Dataset naming standardization

- In most cases, application developers are responsible from backup/archive of datasets that use during batch processes. Make sure application developers do not overwrite management classes in their backup jobs. This is true for local recovery as well. There are cases where backup version of data is deleted/expired because of using wrong management classes by application developers in their JCLs. Review your dataset naming standards and force them in ACS routines.

- Dataset naming standards were not being visited for more than 10-15 years. The current design is making the environment hard to manage and open to errors:

- Make sure you have DFSMS dataset naming standardization implemented and your ACS routines enforce the standards. There are several cases where clients could not access the backup when needed in DR site just because tape dataset naming standards were not correct and data was not in DR . This is important especially for environments where application developers decide/implement which data to take backup.

- Lack of mature dataset naming standards also causes studies like ZEDC compression preparation/implementation / analyze / manage more time consuming. (Ex: 1.000.000 datasets without any clear dataset naming rules...)

- If you have in efficient ACS routines, allocation will take more time. It may cause z/OS uncaptured cputime to be higher. This can be also checked from SMF30ICU field.

Simple, absolute but many clients did not visit their dataset naming standards for many years and/or force rules in ACS routines

Common Issues : Disaster Recovery Site gaps

- DR site gaps –Continue

Make sure, you have setup that you can do DR recovery tests, without effecting availability of the DR replication and environment. Some clients use two flash copy sets on DR site for that reason.

- Follow the best practices also while managing maintenance of HW on DR site.

- Make sure you have enough bandwidth between two Datacenters. Having bandwidth issues will cause you to have delays in replication. Do not use share bandwidth with distributed platform replication (especially their backups) without defining ‘Quality of Service’ QoS.

Bandwidth sharing and quality setup optimizes network performance by prioritizing critical traffic (VoIP, video) using QoS/SQM settings on routers

- Conduct periodic RTO and RPO improvement projects to make sure you are using latest capabilities and solutions and that your RTO and RPO stays in high level of achievement

- Maintenance of HWs needs to be done based on best practices. HIPER MCLs needs to be tracked periodically.

- It is good practice to subscribe to ‘IBM My notification’ capability to be aware of about new HIPER MCLs

- Make sure, you do recover all production systems as is and all production LPARs in DR site.

- Make sure you get benefit from VTS GRID technology solution which is fastest way to have your migrated/archived and backup data in DR site. VTS have ability to do be able to recover from flash set in DR site.

- Make sure you have 2 different paths also used for Customer Management Network , which is needed for system programmers to access to DR site when needed /to do management activities for systems and HW management interfaces.

SPOF – Single Point of Failure Analysis

SPOF analysis for HW, SW, human resources, components in all layers....
If DC is impacted, if DC room impacted, if CEC fails, one lpar,one subsystem member fails ...

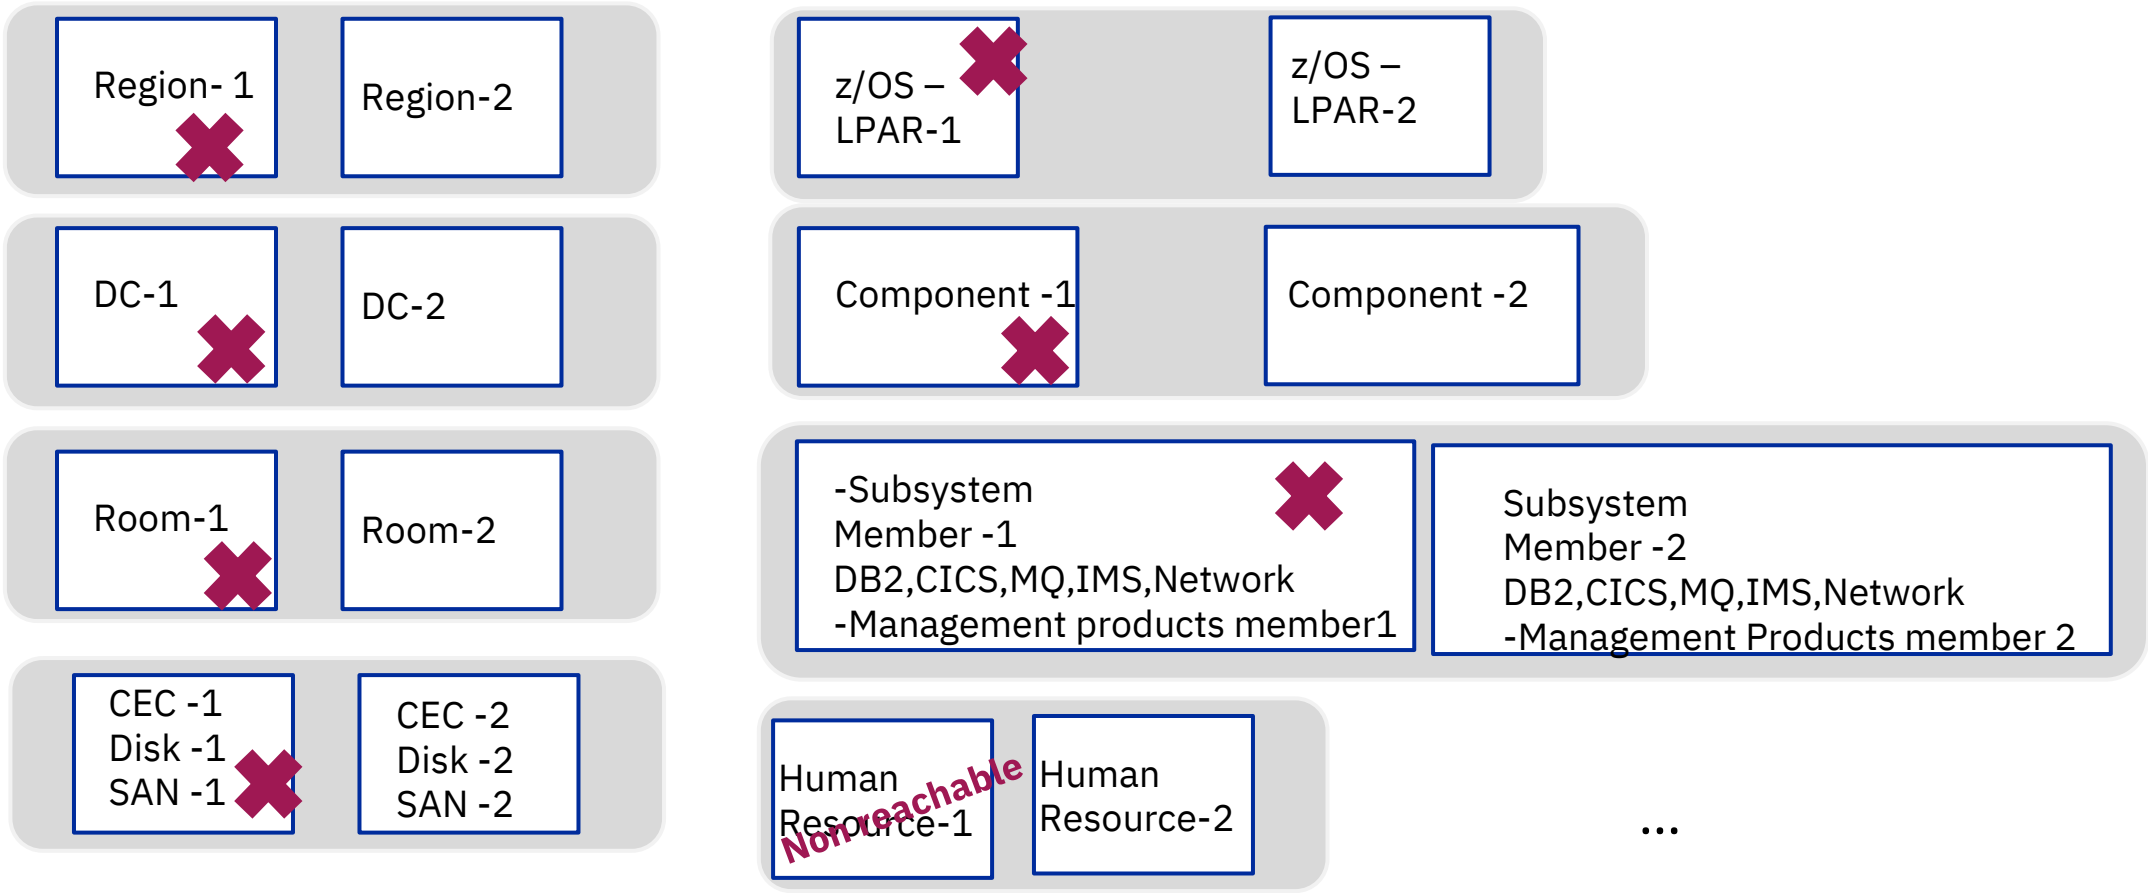

- Sysplex capabilities fully implemented for **every** application and products that supports HA.
- SPOF reviews integrated into 'change management' processes, CAB reviews
- Internal team meetings at least once in a year – internal SPOF analysis review studies
- SPOF review in design process, new architectural, infrastructure, application changes...

Common issues: Lack of It mainframe teams – business departments communication

- IT mainframe teams – business departments communication

- SLAs are not agreed with business departments .
- Business departments expecting 0 downtime, but IT setup is not aligning with that. In most cases, this is realized after an outage.
- It is important to have SLA agreements with business departments so that investment to resiliency projects can be supported as well. SLA agreements need to be agreed for each business channel.
- Capacity planning process does not include input from business departments growth projections.
  Example: Increase in number of ATMs, number of internet banking logins, ...

Common Issues : Cost of existing outages are not calculated for each minor/major service discontinuity

❑ Cost of existing outages are not calculated for each minor/major service discontinuity

- Example: Internet banking down for 5 minutes.
  Cost of the outage to business suppose to be calculated and reported.

Result: Investment to Resiliency is not approved (human resources , human resource time and HW/SW costs)
Cost of preventable outages is not visible to decision makers

Example of calculation method :

Between 10:00 – 10:10 100.000 transaction abends
Services were unresponsive between 10:00 – 10:20
Look for similar day completed business services
Earlier task to be completed: Profit from each service: profit/sec for similar day
Multiply profit/sec X 20 minutes X 60 seconds = XXXXXXX $
Or
Earlier task to be completed : Profit from each transaction completed = B
Look at similar day # of transactions completed = A
If outage did not happen A*B $ is the money that company gains
Cost of outage is A*B dollars

This study requires collaboration of IT and business departments

- Also reputation damage cost can exist and it should be added to above cost
- Also penalty cost can exist (Financial companies may have to pay penalty for their service unavailability : Credit Card Offices etc.). This should be added to above cost
- There can be also risk of losing business certifications , depending on duration of the outage and regulations rules in that country.

Balanced solutions designed to address different requirements

GDPS Metro	GDPS Global	GDPS Metro Global	GDPS Continuous Availability
Near-continuous availability and recovery at metro distances	**Disaster recovery at extended distance**	**Near-continuous availability regionally & recovery for 3-4 sites**	**Near-continuous availability, recovery & workload balancing**
Systems remain active Multisite workloads can withstand site and storage failures	Rapid systems DR with "seconds" of data loss	Metro near-continuous availability and out of region disaster recover	Continuous availability at unlimited distances
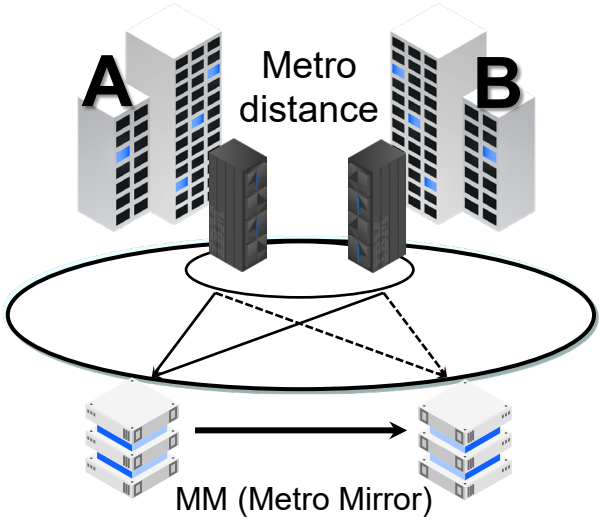	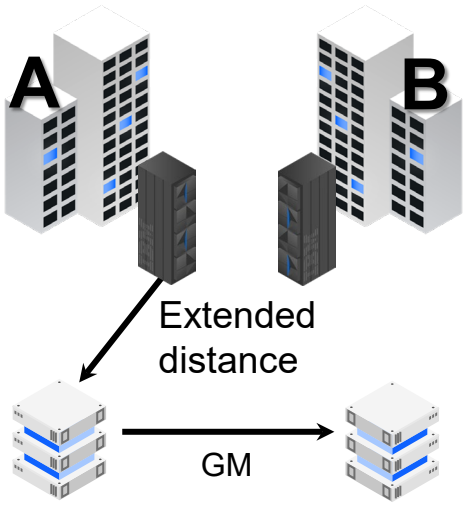	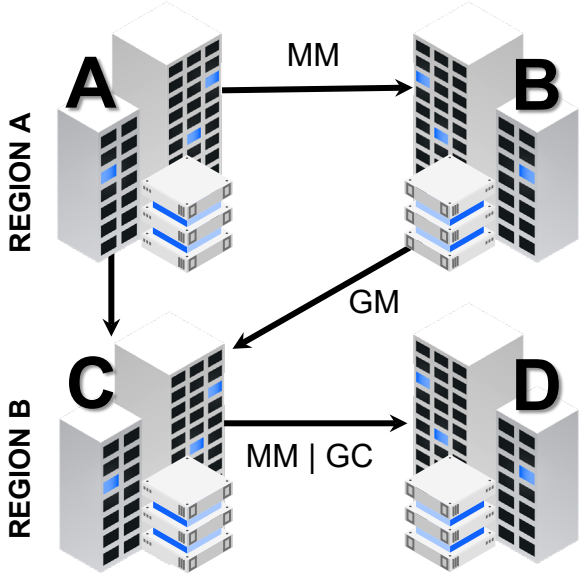	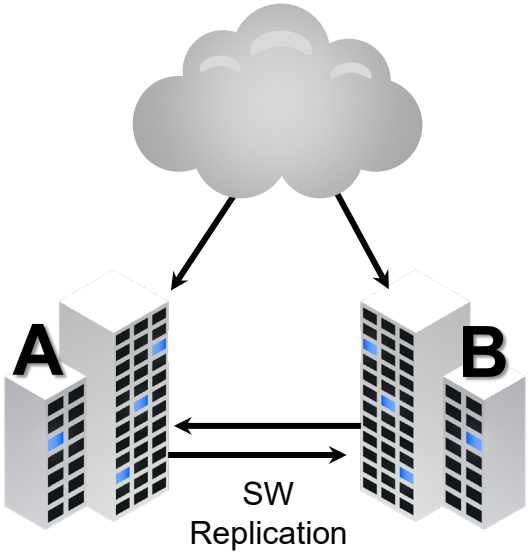
RPO 0 & RTO <60 min	***RPO 3-5 sec & RTO <60 min***	***RPO 3-5 sec & RTO <60 min***	***RPO 3-5 sec & RTO <60 sec***

IBM DS8000 Data Resilience

When bad things happen to your data - protection from failures, disasters and cyber attacks

Flash Copy

Point in time copy

Within the same system

Instantly accessible for fast recovery

Safeguarded Copy

Secure Immutable point in time copy

Within the same system

- Secure Immutable, not erasable
- Hidden to hosts and applications

Metro Mirror

Synchronous mirroring

Primary Site A — Metro distance Site B

- 303 Km; 500 Km with RPQ
- Performance improvements
- 50% less infrastructure

Global Mirror

Asynchronous mirroring

Primary Site A — Out of region Site B

Metro / Global Mirror

3/4/5/6 site cascaded or multi-target synchronous and asynchronous mirroring

Primary Site A — Metro distance Site B — Out of region Site C

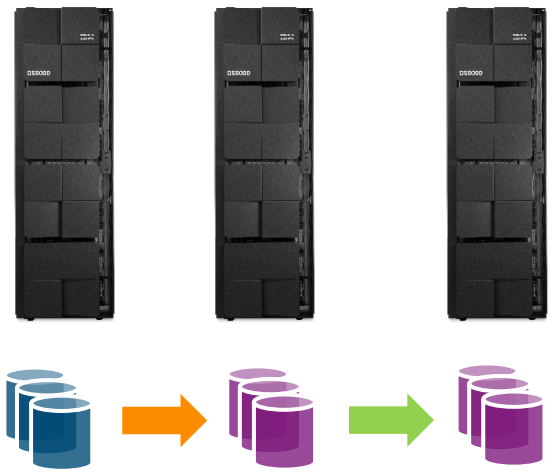

- Up to 50% bandwidth reduction over competitive solutions
- From 1 to 3 seconds RPO
- Superior application failover via CSM, GDPS or PowerHA

Replication architecture & health & performance

- Two DWDMs, Monitor DWDM, FCIP routers performance data. The availability of the telecommunication impacts your data replication and your ‘data consistence ‘ time difference, which is your RTO. It may also impact your production, depending on the architecture.
- Make sure DWDMs, FCIP routers, switches, SAN devices are configured with best performance parameters.
- Monitor and report replication performance, it impacts your RPO(Recovery Point Objective – how many seconds of data lost you can have in async replication).
- Make sure you have test environment for your DR replication solution .You need to use that environment , to test SWs and solutions/scripts before doing any changes directly in production replication environment.
- Several regulations require encryption being used while transferring data to other DC. FC/IP routers/ DWDMs have encryption capabilities.

GDPS 4.8 -Use current version – leverage enhancements & new functions

GDPS Metro (and HM) Highlights – Summary

- **GDPS Security enhancements**
  - Role-based security enhancements
    - New profiles for GDPS Scripts
  - GEOSEC tool updated for GDPS Scripts
  - Change Role Based Security setting SECURITY default to SAF
- **New Standard Actions panel to provide a LPAR (physical) view**
- **Change management**
  - Generate reports when testing or loading a new dasd configuration
  - Generate reports when testing or loading a new site table
- **Identification of active CC session(s) during a planned HyperSwap**
- **Support for GDPSIOST in Metro for post-HyperSwap host access cleanup and managing Utility device dynamic path grouping**

Note: xDR supported versions for z/VM proxy or production cluster in GDPS 4.8
- Linux: RHEL 8, RHEL9 and SLES15
- TSAMP: 4.1.1.1

GDPS GM and MGM Highlights – Summary

- **GDPS Security enhancements**
  - Role-based security enhancements
    - New profiles for GDPS Scripts
  - GEOSEC tool updated for GDPS Scripts
  - Change Role Based Security setting SECURITY default to SAF
- **New GDPS priced feature: GDPS Solutions Manager (GSM)**
  - Simplified systems management for clients with multiple sysplexes managed by GDPS
  - Consolidate up to 3 GDPS controlling system pairs into a single SYSPLEX(2 LPARs) to reduce complexity and lower system management overhead
  - Initial support for MGM 4-site topologies only
    - Some limitations at GA
- **New GDPS Topology: GDPS Metro Global Mirror 6-site (MGM 6-site)**
  - Enhanced resiliency by enabling 3 copies of data to be managed synchronously in the active region.

New GDPS Solutions Manager (GSM) feature

Example: 2 MGM4SITE environments moved to GSMPLEXes

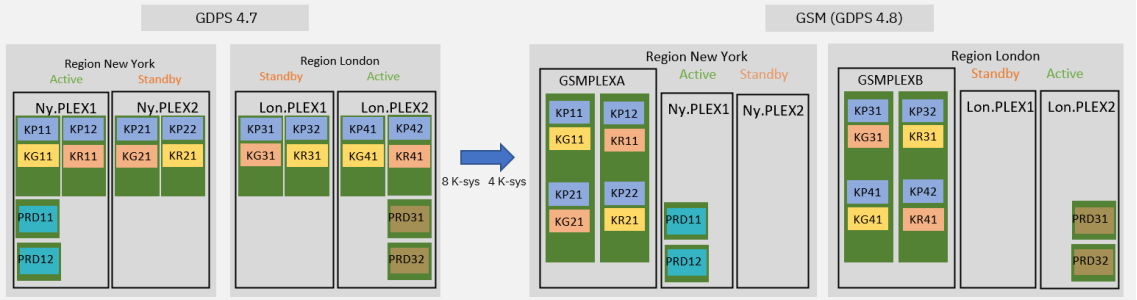

- **2 x MGM 4-site license**
  - 1 each for Plex1 & Plex2
- **No GDPS license restrictions on the number of LPARs in each sysplex**
  - It's two (e.g. PRD11, PRD12) in this case but it could be 1 or 5.... No change to the GDPS license required

- **2 x MGM 4-site license (No change to today)**
  - 1 each for Plex1 & Plex2
- **2 x NEW GSM features licenses**
  - 1 GSM feature per sysplex, up to 3 sysplex per GSM Plex
- **Allow use of z/OS Proxies within GDPS licensed sysplex (no charge)**
  - The client would have to deploy 2 z/OS proxies in both Plex1 & Plex2
  - There will be no additional license required for these proxies
- **Limitations at GA**
  - LCP Management functions not available
  - GDPS GUI not available for GSM agents
  - Health check restrictions
  - Limited protection when IPLing from HMC

Request for GDPS update briefing

MGM 6-site overview - requirements

Requirements are even further than 4-site -> 6-site

Background for MGM 6-Site

- Customer requirement:

> Production can only run if and only if 2 local synchronous copies and one remote consistent copy (DR copy) are available.

- Production cannot be maintained with MGM 4-site in case of a Metro primary failure (1 local copy left) or GM secondary failure (no DR copy). Failure of the RS2 DR copy prevents region switch (1 local copy left).

- IBM Statement Of Direction (SOD) to address the requirement:

> IBM intends to extend support for 6 replication copies where a Multi-Target Metro Mirror (MM3SITE) topology in one region is connected via Global Mirror to a second Multi-Target Metro Mirror environment in a second region.

MGM 6-site overview – New Topology

- New GDPS topology: **MGM6SITE**
- 2 server sites (SSn) per region, same as MGM4SITE and MM3SITE
- Expands MGM4SITE from 2 to 3 replication sites (RSn) in each region
  - Metro dual-leg instead of single-leg
- Provides two possible GM secondary sites
  - Incremental resynchronization used when changing GM secondary site (*new functionality provided in DS8K R9.4*)

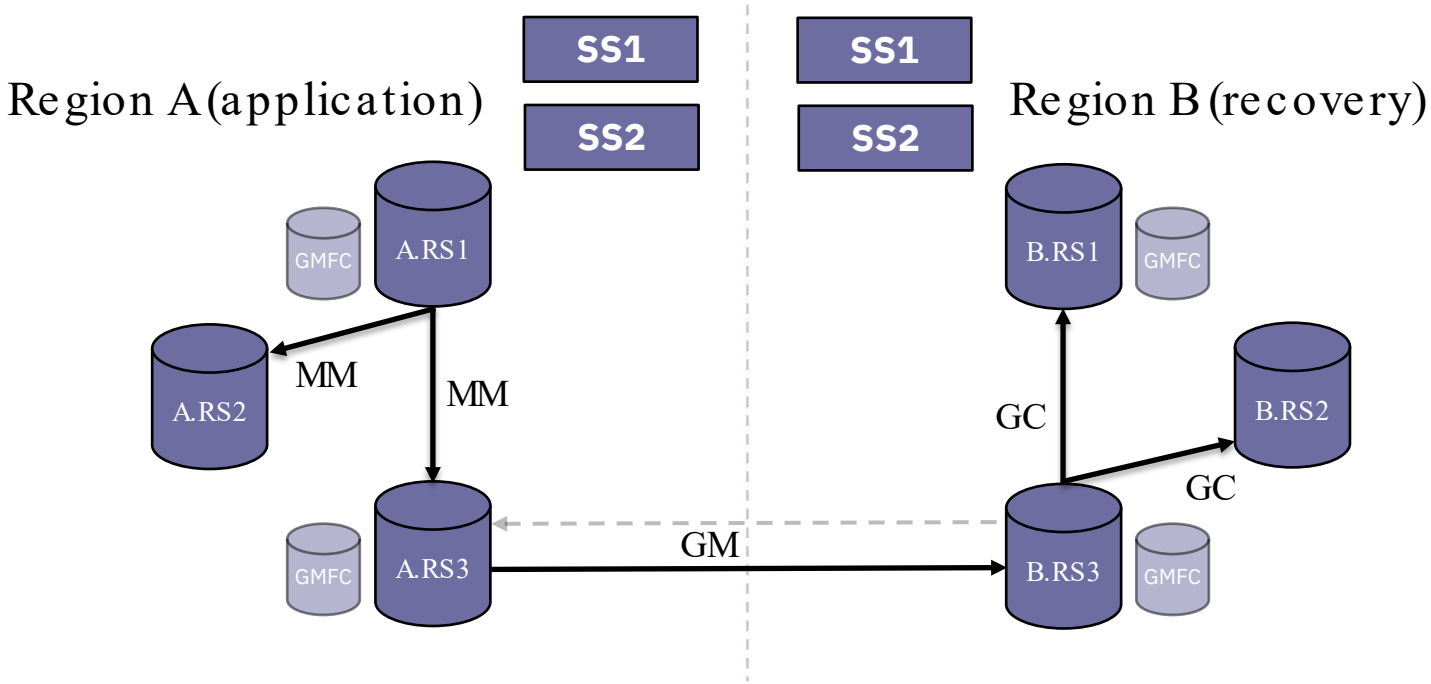

Application source code & development capabilities in DR site

- ❑ **Resiliency is recovery from failure /outage without an issue and continue operations in similar level of quality.**

- When there is an outage, it is likely that some UIs that create interfaces to end user definitely need to be updated.
- There might be other cases where you really need to update application related to the outage.
- This means you need latest version of your application source codes in DR site and also you may need to be prepared a development environment (at least compiler capability) in DR site / in production LPARs that are recovered in DR site.

- ❑ **Because of those reasons as well, it is also good practice to replicate non-production environments**. If you do not have that setup , at least make sure you have latest source code on production disk boxes so that incase of an outage, if production is recovered in DR site, you have the ability to compile/develop/update applications.

Common issues : Backup Plans – Be ready for a problem – Test Recovery Processes

Several Leadership Lessons To Learn From This Story…
Try to get ready for Possible Failures, What If Senarios

They had plan B,C,D,E.
They had place to test the plans!

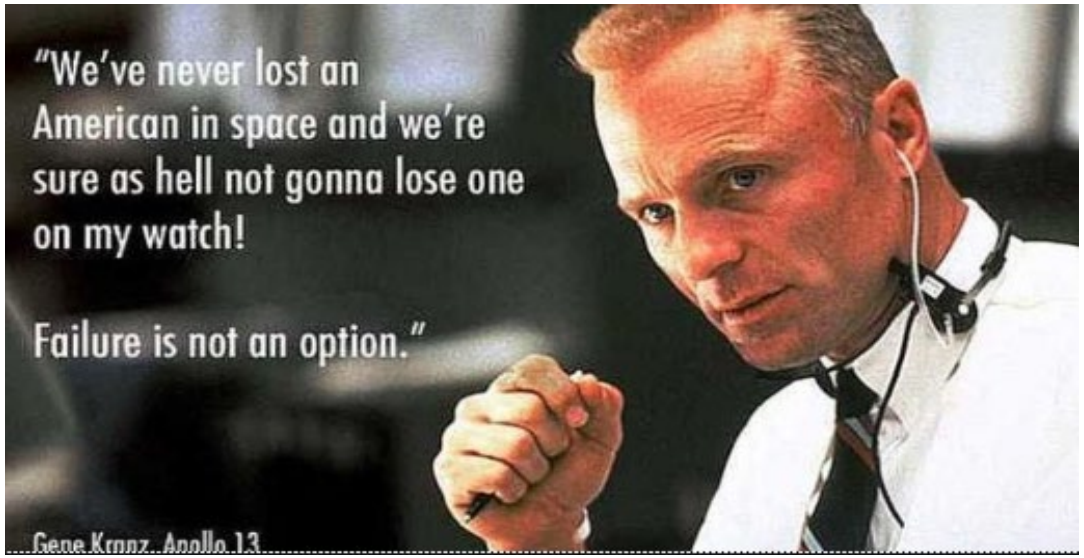

Always Have Plan B – Backup Plans!
There are series of methods To Let you a Success Failure!

https://www.youtube.com/watch?v=s_7PfocHTmc

https://www.forbes.com/sites/geoffloftus/2013/04/03/apollo-13-lessons-from-the-successful-failure/#1fceb25c5d0b


Focus | SPACE EXPLORATION

"Houston, we've had a problem."

Those are the exact words of Apollo 13 astronaut James A. Lovell Jr. after his spacecraft experienced an explosion on this day in 1970. With that sentence, the Apollo 13 mission changed from a glorious moon landing to a daring lifesaving rescue. Here's a look at the series of events.

BY KURT SNIBBE
STAFF

Apollo 13 was the seventh manned mission in the American Apollo space program and the third mission expected to land men on the Moon.

Not to scale

Command module
Service module
Lunar module

1. ROUGH START
Liftoff from Kennedy Space Center in Florida is at 2:13 p.m. April 11, on a Saturn V rocket. Five and a half minutes after liftoff one of the four engines shuts down two minutes early.

2. Saturn's third stage fires Apollo 13 on translunar course

3. Apollo docks with Aquarius (lunar module).

4. 8:54 p.m. April 12 Apollo changes trajectory to prepare for moon landing

5. EXPLOSION
10:08 p.m. April 13: 205,000 miles from Earth
The Accident Review Board concluded wires shorted that were damaged in preflight testing. The wires started a fire that eventually exploded oxygen tank No. 2 and damaged tank No. 1.

6. 3:43 a.m. April 14: Apollo changes back to free-return trajectory (164 miles above the moon instead of 69 miles high)

7. 9:41 p.m. April 14: The engine on the Aquarius lunar module is used for 4 1/2 minutes to put the spacecraft on course toward Earth.

8. 11:31 p.m. April 15: Midcourse correction refines course for Earth

9. 8:15 a.m. April 17: Apollo jettisons and photographs dead service module

10. 11:43 a.m. April 17: Astronauts transfer to command module then jettison their lifeboat, the Aquarius lunar module

11. SPLASHDOWN
1:08 p.m. April 17: The command module safely lands in the South Pacific. The astronauts board the USS Iwo Jima at 1:53 p.m.

Mission control celebrates

Human Errors

‘It is always a risk for everyone in every level to make mistakes/errors’
Highest portion of outages are because of human errors
There are functions that will prevent you from making mistakes
There are certain methodologies that will prevent you from making mistakes
Leverage all!

Common issue : HMC LPAR User error protection not used – Z15 and above

- Available since Z15.HMC user profile protection capabilities protects clients from outages caused by human errors while using HMC tasks .
- Clients can customize/enable added protection against user errors for LPAR targets
  - Client create Userid customization
    - Userid configured for 'view only'
    - Userid limited to CPCs / LPARs that are Non-Production
    - Userid limited to Non-disruptive tasks
  - Object Grouping
    - Organize objects (CPCs, LPARs) into Groups for clearer recognition of use
  - Object Locking
    - Lock all LPARs / CPCs from any disruptive task
    - Limit which Userids have the authorization to Lock/Unlock task
  - Task Confirmations
    - Basic Default confirmation pop-up for all tasks performed.
    - Customize individual Client Created Userid to enable **Disruptive Userid/Password** confirmation
    - Customize individual Client Created Userid userid/password confirmation for additional requirement to **add text input** (i.e., OS name to be affected)

- REASON of not using : ‘Not aware of it’

Z17 HMC Dual Control

Dual Control Design Highlights

Dual control adds an extra layer of security for critical tasks on the HMC.

Dual control enabled tasks require another level of verification from an approver before they can be run.

Dual Control Target per User Role
- Object and
- Task

Dual Control Approver per User Role
- Any User Role for task/object authorization control
- Can also create special User Role with list of specific users for DC approval

Dual Control Management Execution Requests for Approval
- Run by requester restricted to time window in the future
- Run by requester
  - No time window restriction
- Run immediately
  - Automatically without requester further involvement
- Run on a specific date and time

Dual Control CPCs supported (Requires z17 HMC 2.17.0)
- z17 CPC (no restrictions)
- z16 & z15 CPCs
  - Should remove Single Object Operations in any role applied to a User under Dual Control
  - *Perform Model Conversion (Capacity on Demand)* & *Change LPAR Cryptographic Controls* tasks not available for HMC 2.17.0 z16 & z15 targets

Dual Control external interfaces (UIs, WS APIs, BCPii v2, IBM HMC Mobile)

Summary of Common Issues

- Tools to speed up Predict, Detect, Prevent, Speed up recovery are needed to be reviewed and consider to leverage

- Free tools , functions, capabilities of HW, z/OS,SW, Middleware
- IBM Z Log & Data Analytics / syslog analyze mechanism
- IBM Concert for Z
- IBM Z Backup Resiliency
- Cyber Vault

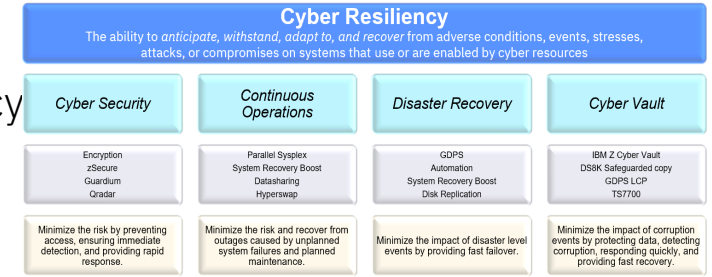

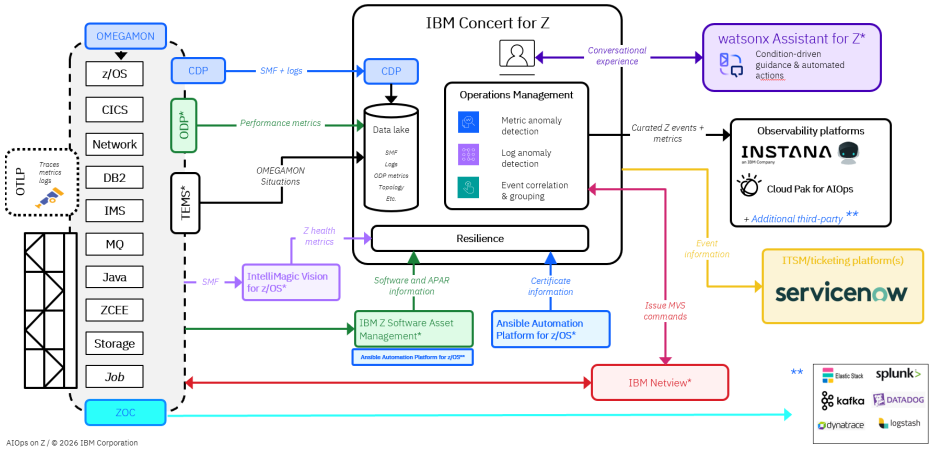

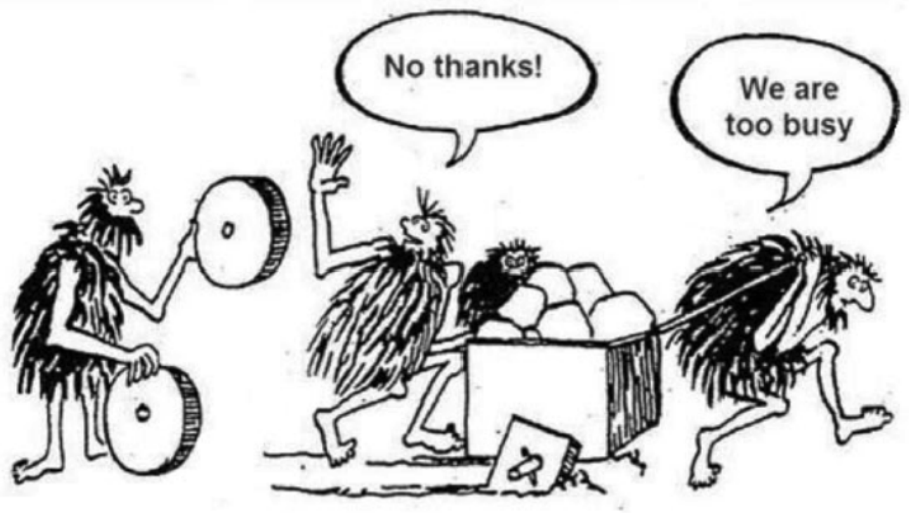

Tools & Solutions that addresses improved resilience

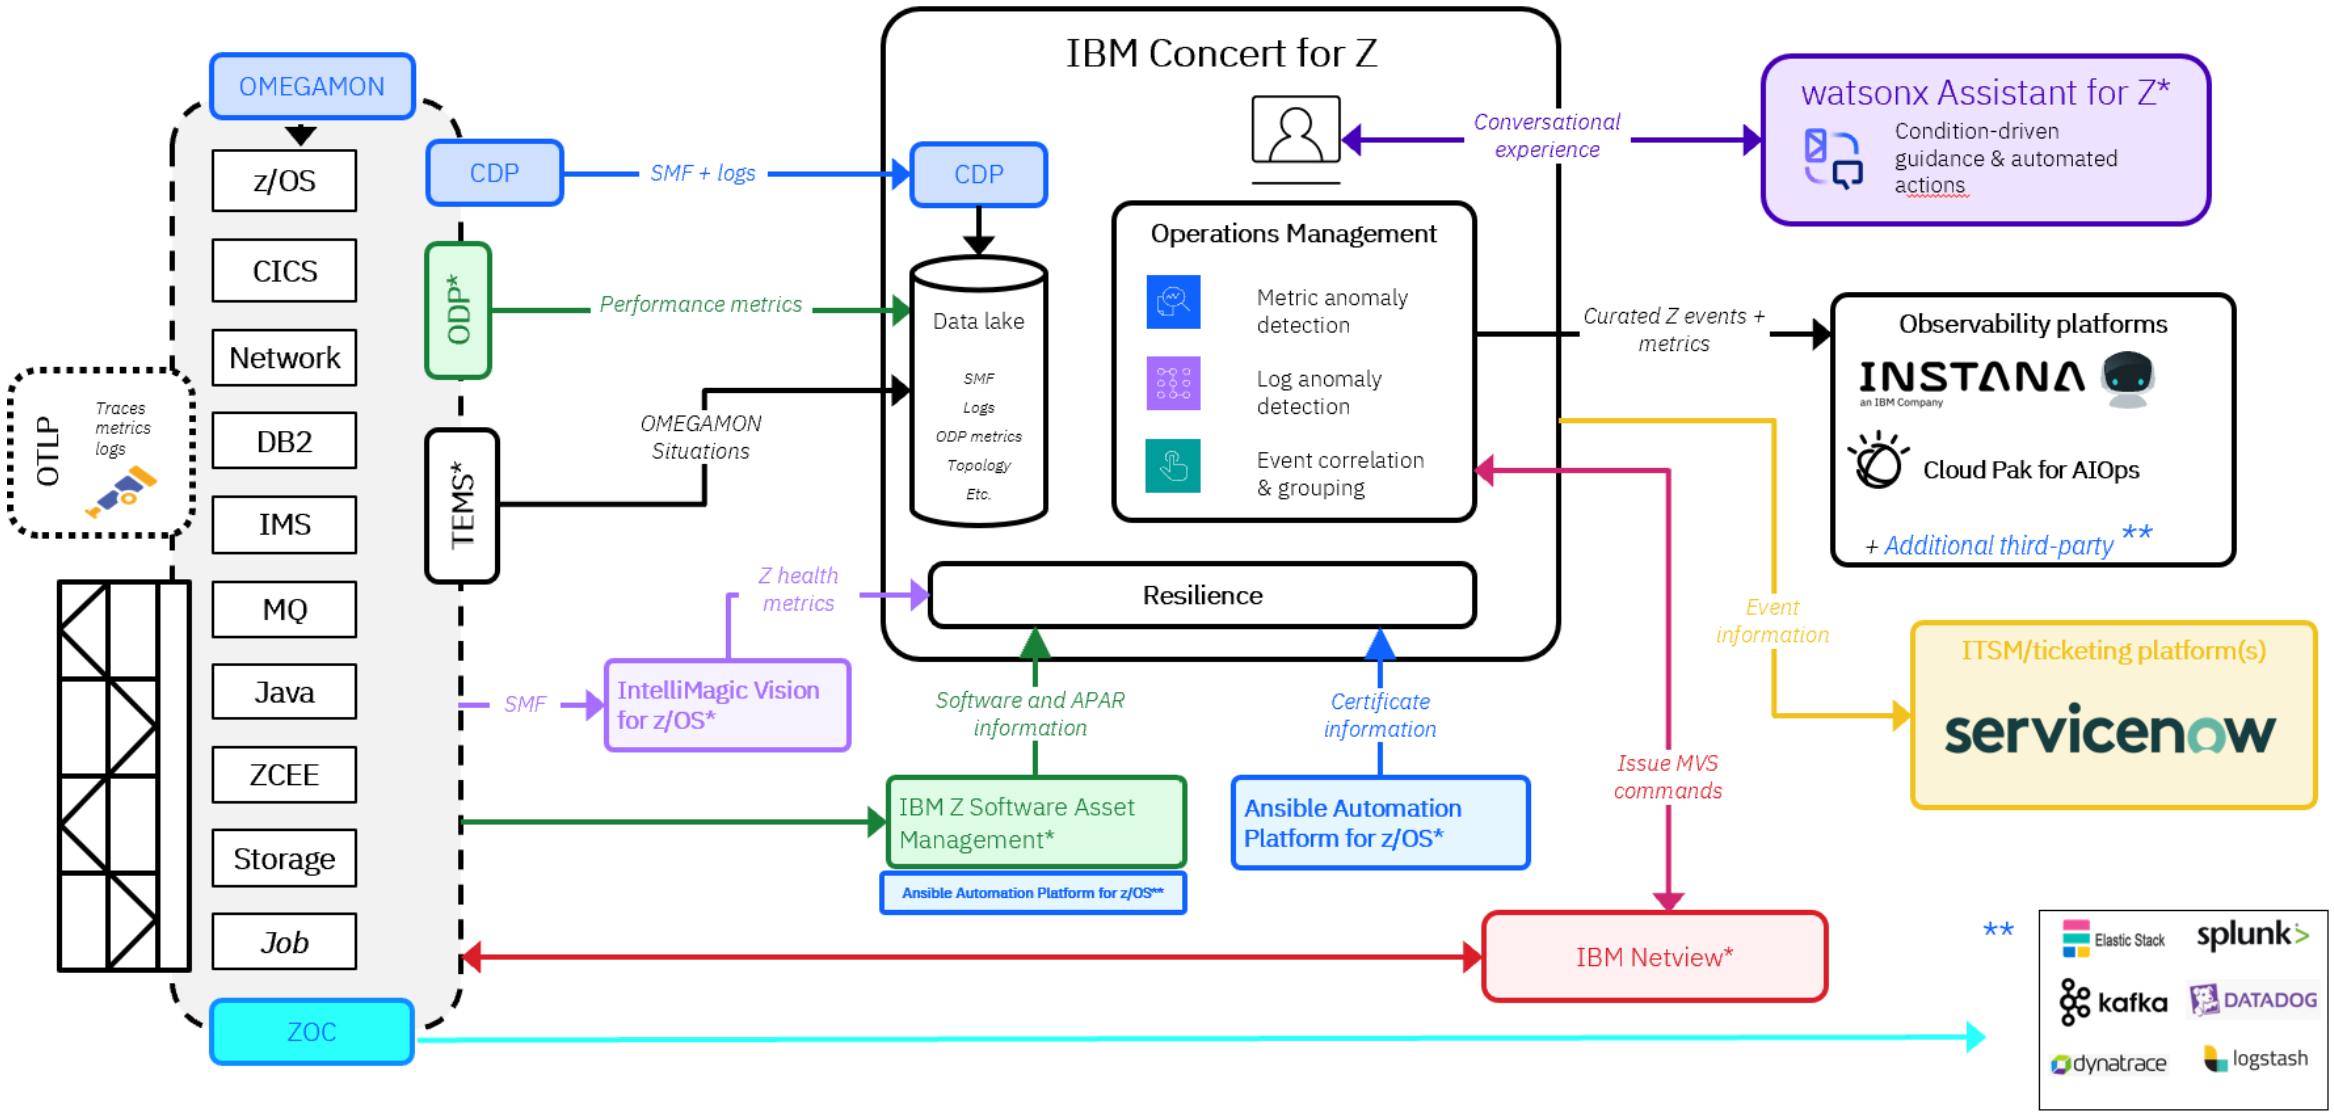

Tools & Solutions that addresses improved resilience

How IBM Concert for Z works

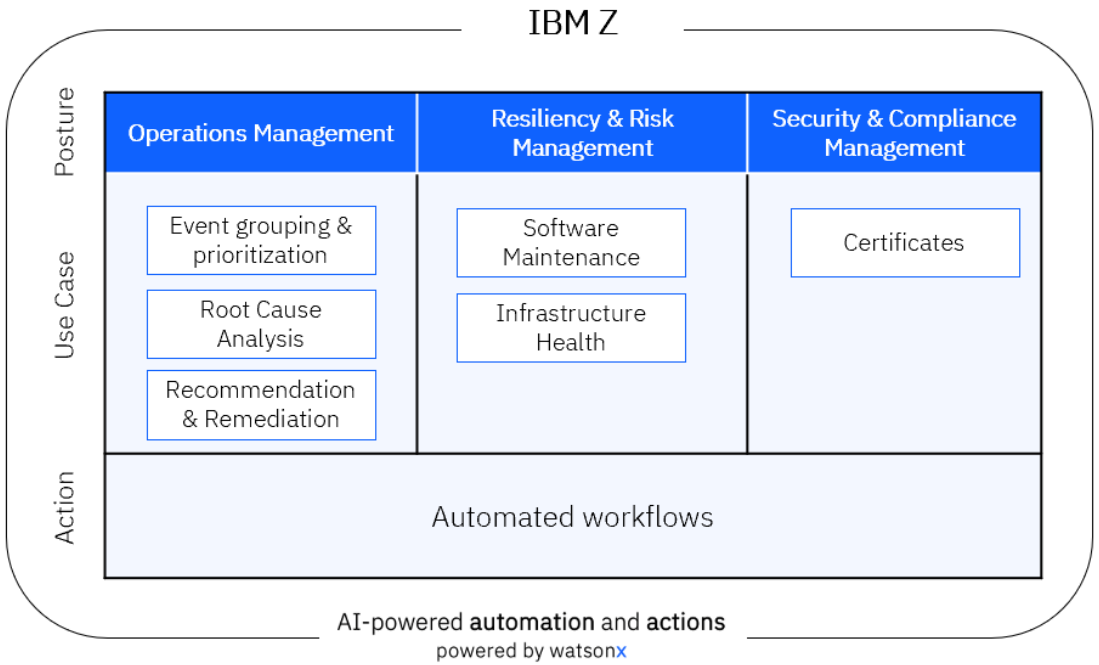

Resilience Dashboard

Discover

AI-driven grouping and prioritization of events for faster root cause analysis. Identification of potential risks.

Seamless integrations with your existing AIOps tools

Understand

Analyze and determine root cause by correlating data and providing easy access to relevant information

Visualize insights or explore them through a natural language chatbot.

Recommend

Context-aware recommendations tailored to each use case.

GenAI-powered insights and recommendations

Act

Take actions based on Gen-AI-powered recommendation tied to automation or incident management process.

Powered by watsonx

Key Benefits

- ✓ Reduce risk
- ✓ Improve resilience
- ✓ Reduce cost
- ✓ Save time
- ✓ Increase productivity

(Un)predictable high spiking business critical workloads

This cpu usage pattern is common for many system programmers : During issues / Workload increase

Changes in workload behaviors are caused by:

- Social, Mobile and Traditional Media
- Mobile payments
- FinTech companies, Apple Pay, iDeal, etc..
- Regulators (PSD2/Instant Payments)
- More card payments / less cash payments
- New Customer Demands
- Dynamic change between Batch and On-Line
- Scattered workplaces (work from home)
- **Side effects of errors in applications**
- **Issues in outside of z platform – Peak workload that was delayed because of bottleneck**

Disruptive events:

- Workforce strikes
- Elections and other political changes
- Geopolitical actions

Pandemic outbreaks can cause:

- Stock & Currency exchange volatility
- Cancellations of Hotel / Flight tickets / Rentals
- Logistical challenges
- D-DoS and other hacker attacks

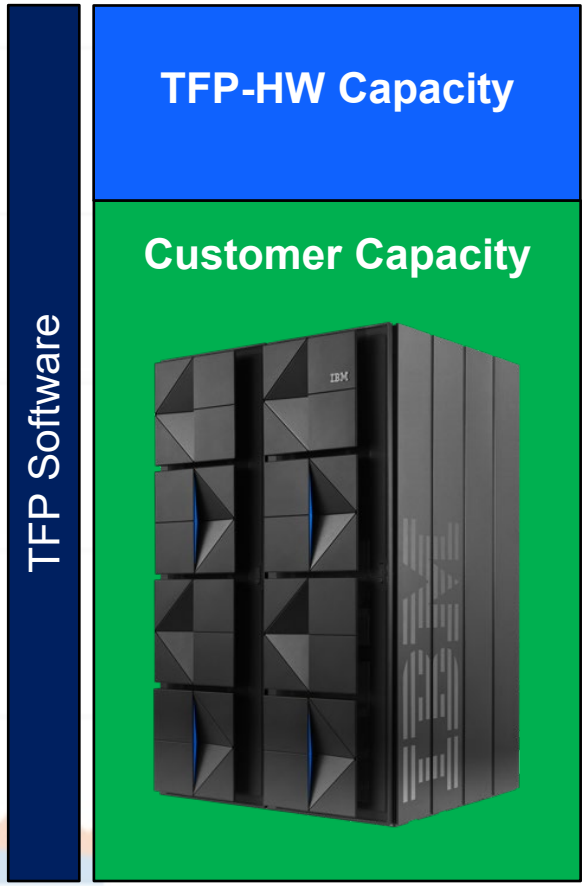

Protects system/subsystems from catastrophic impacts of peak cpu usage because of errors

Avoid Outage

Recover Fast

} TFP-HW effects Resiliency from both perspectives

Increase the toleration level of system to abnormal situations

Extra capacity covers up most of the time several issues

Issue outside IBM z , Solved, bulk workload comes

Issue in an application that causes high cpu usage

TFP-HW (Tailor Fit Pricing HW)

Time that everything is impacted
This can cause rolling issues

CICS – Max task reach
DB2 – Max Dbat reach, max thread reached
Delays in subsystems internal tasks if WLM policy is not efficient

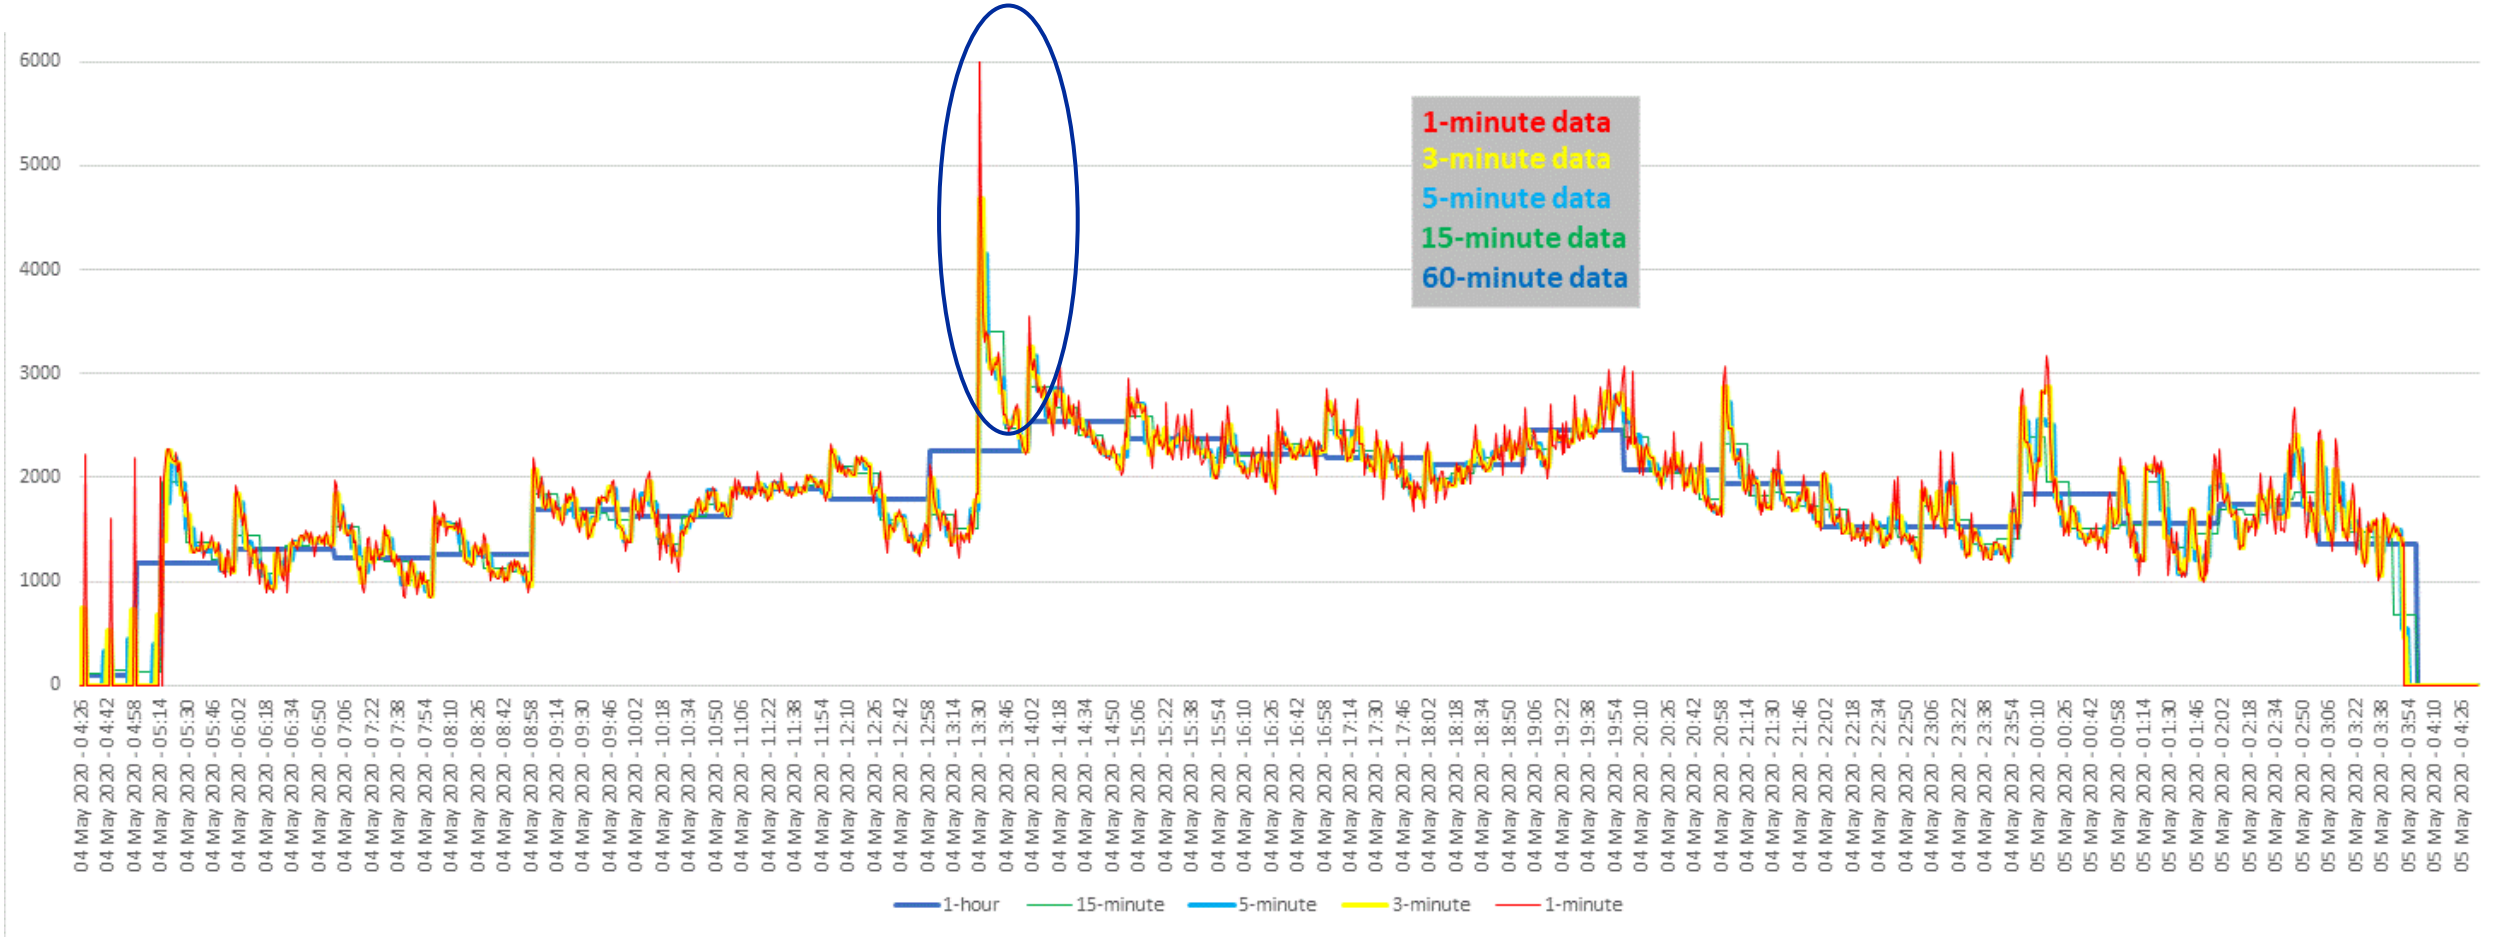

TFP-HW (Tailor Fit Pricing HW)

Averaging effect: Comparison of 1-minute measured data to recalculated 3-, 5-, 15- and 60-minute MSU usage

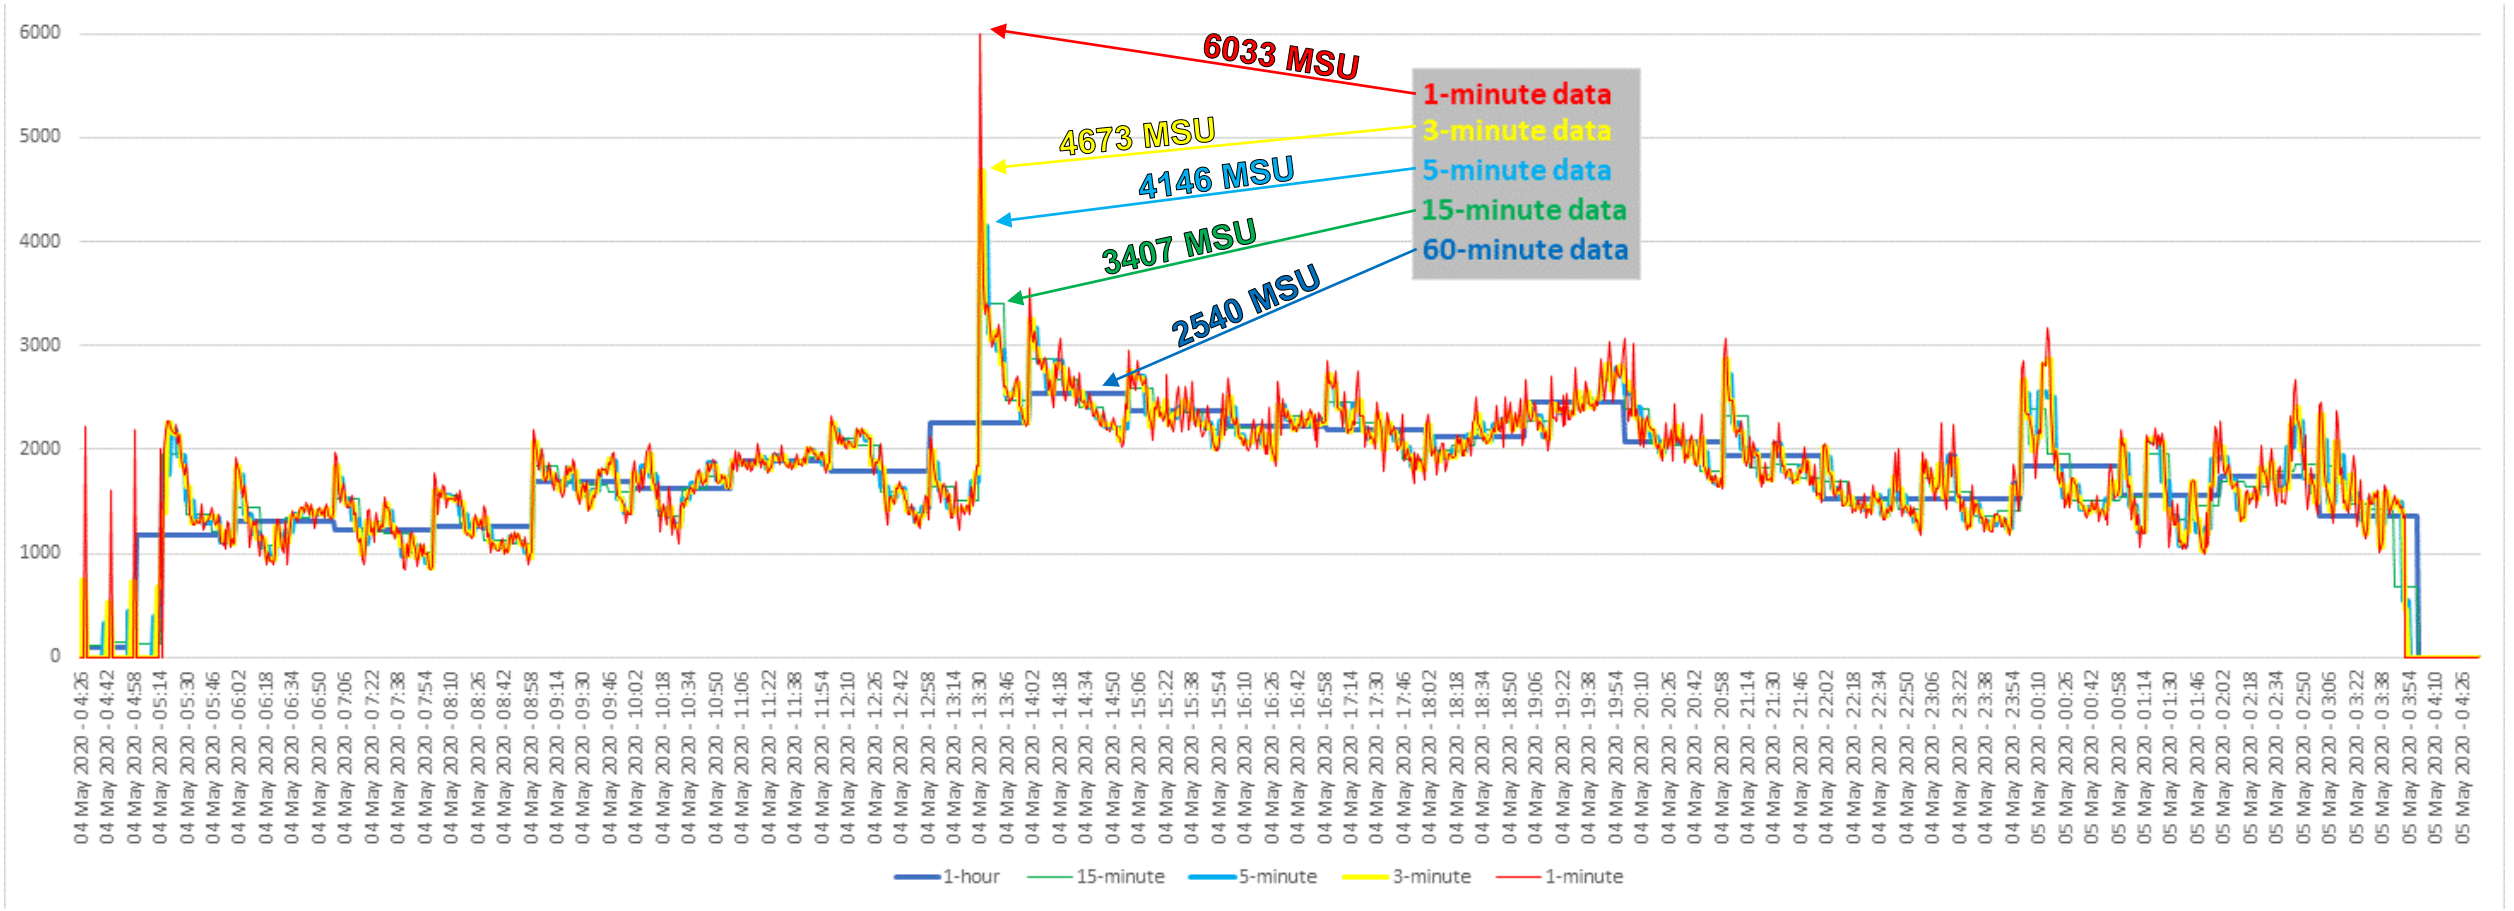

Hardware Models : N vs N-1 – Benefits of using current model

z17

One hardware model, Five Features, 1-4 19” Frame System, Maximum 8 cores/chip, 2 chips/DCM
1 integrated I/O accelerator/chip.
5.4Ghz.

Up to 85 user partitions,
32 TB per partition,
208 CPUs/zIIPs/IFLs per partition, up from 200.

Up to 16 TB per z/OS LPAR as of z/OS V2.5

Channel Subsystem scalability
•Up to six (6) Channel Sub Systems (CSSs)
• 4 Subchannel Sets per LCSS

IBM Z Integrated Accelerator for AI

HiperDispatch Enhancements

Hardware Instrumentation Services (CPUMF)

New machine instructions

IBM Flexible Capacity for Cyber Resilience

IBM System Recovery Boost

Integrated I/O Architecture (DPU)

IBM Spyre AI Accelerator (4Q2026)

Network Express

BCPii support for z/OS Identity Mapping

Power Consumption Reporting

TFP for Hardware
- Workload Classification
- Replacement Capacity

z17	z16
•2 CP chips on a Dual Chip Module (DCM), 5.5 GHz •L1 Private 128K instruction & 128K data •L2 Shared 36 MB / core, 270 MB effective shared •10 x Private/Shared 36 MB L3 caches	•2 CP chips on a Dual Chip Module (DCM), 5.2 GHz •L1 Private 128K i & 128K d •L3 Shared 32 MB / core, 192 MB effective shared
684 GB HSA, 64 TB maximum, 16 TB per drawer	256 GB HSA, 40 TB maximum, 10 TB per drawer
Max coupling CHPIDs per CEC of all types: 384 8 subchannels for IC CHPID	Max coupling CHPIDs per CEC – 384 7 subchannel for IC CHPID
ICA-SR 2.0 for short-reach coupling • 4 CHPIDs/port for CS5 • Same limits as z15/z16 with 48 adapters/96 ports.	ICA-SR 1.1 Max ICA SR per CEC 48 adapters/96ports (same as z15)
Coupling Express3 LR 10Gb/25Gb optics for long-reach • 4 CHPIDs/port for CL5/CL6 • 32 adapters, 64 ports per CEC	Coupling Express2 LR 10Gb (CX6-DX) PCIe adapter
CF Level 26 •Parallel Sysplex scalability, virtualization, consolidation, and density enhancements. •Removal of support for CF Flash Memory and CF images using dedicated GP processors.	CF Level 25 •Retry buffers for cache and lock commands •Cache residency time metrics •Scalability improvements •CF Request latency/performance improvements
10 GbE and 25 GbE RoCE Express 3 SR and LR (CX6-DX)	10 GbE and 25 GbE RoCE Express 3 SR and LR (CX6-DX)
zHyperLink® Express 2.0 • Maximum 16 Adapters /32 ports	zHyperLink® Express1.1 • Maximum 16 Adapters /32 ports
FICON Express 32 -4P	FICON Express 32S
OSA Express7S 1.2 1/10/25 Gb	OSA Express7S 1.2 1/10/25 Gb
Crypto Express8S	Crypto Express8S

z16

One hardware model,
Five Features, 1-4
19” Frame System

Up to 85 user partitions, 32 TB per partition,
200 CPUs/zIIPs/IFLs per partition, up to 224 Pus

Up to 16 TB per z/OS LPAR with z/OS V2.5

Channel Subsystem scalability
•Up to six (6) Channel Sub Systems (CSSs)
•4 Subchannel Sets per CSS

IBM Z Integrated Accelerator for AI

HiperDispatch Enhancements

Hardware Instrumentation Services (CPUMF)

New machine instructions

IBM Flexible Capacity for Cyber Resilience

IBM System Recovery Boost

z/OS Validated Boot

Z15
LICWithdrawn dates passed
Not leveraging Resiliency, Performance
Scalability Enhancements

z/OS Continuous Delivery – Be in current version

How can I continue to get the latest enhancements ? ? What is the 'NEW' benefit of being in current release!!!

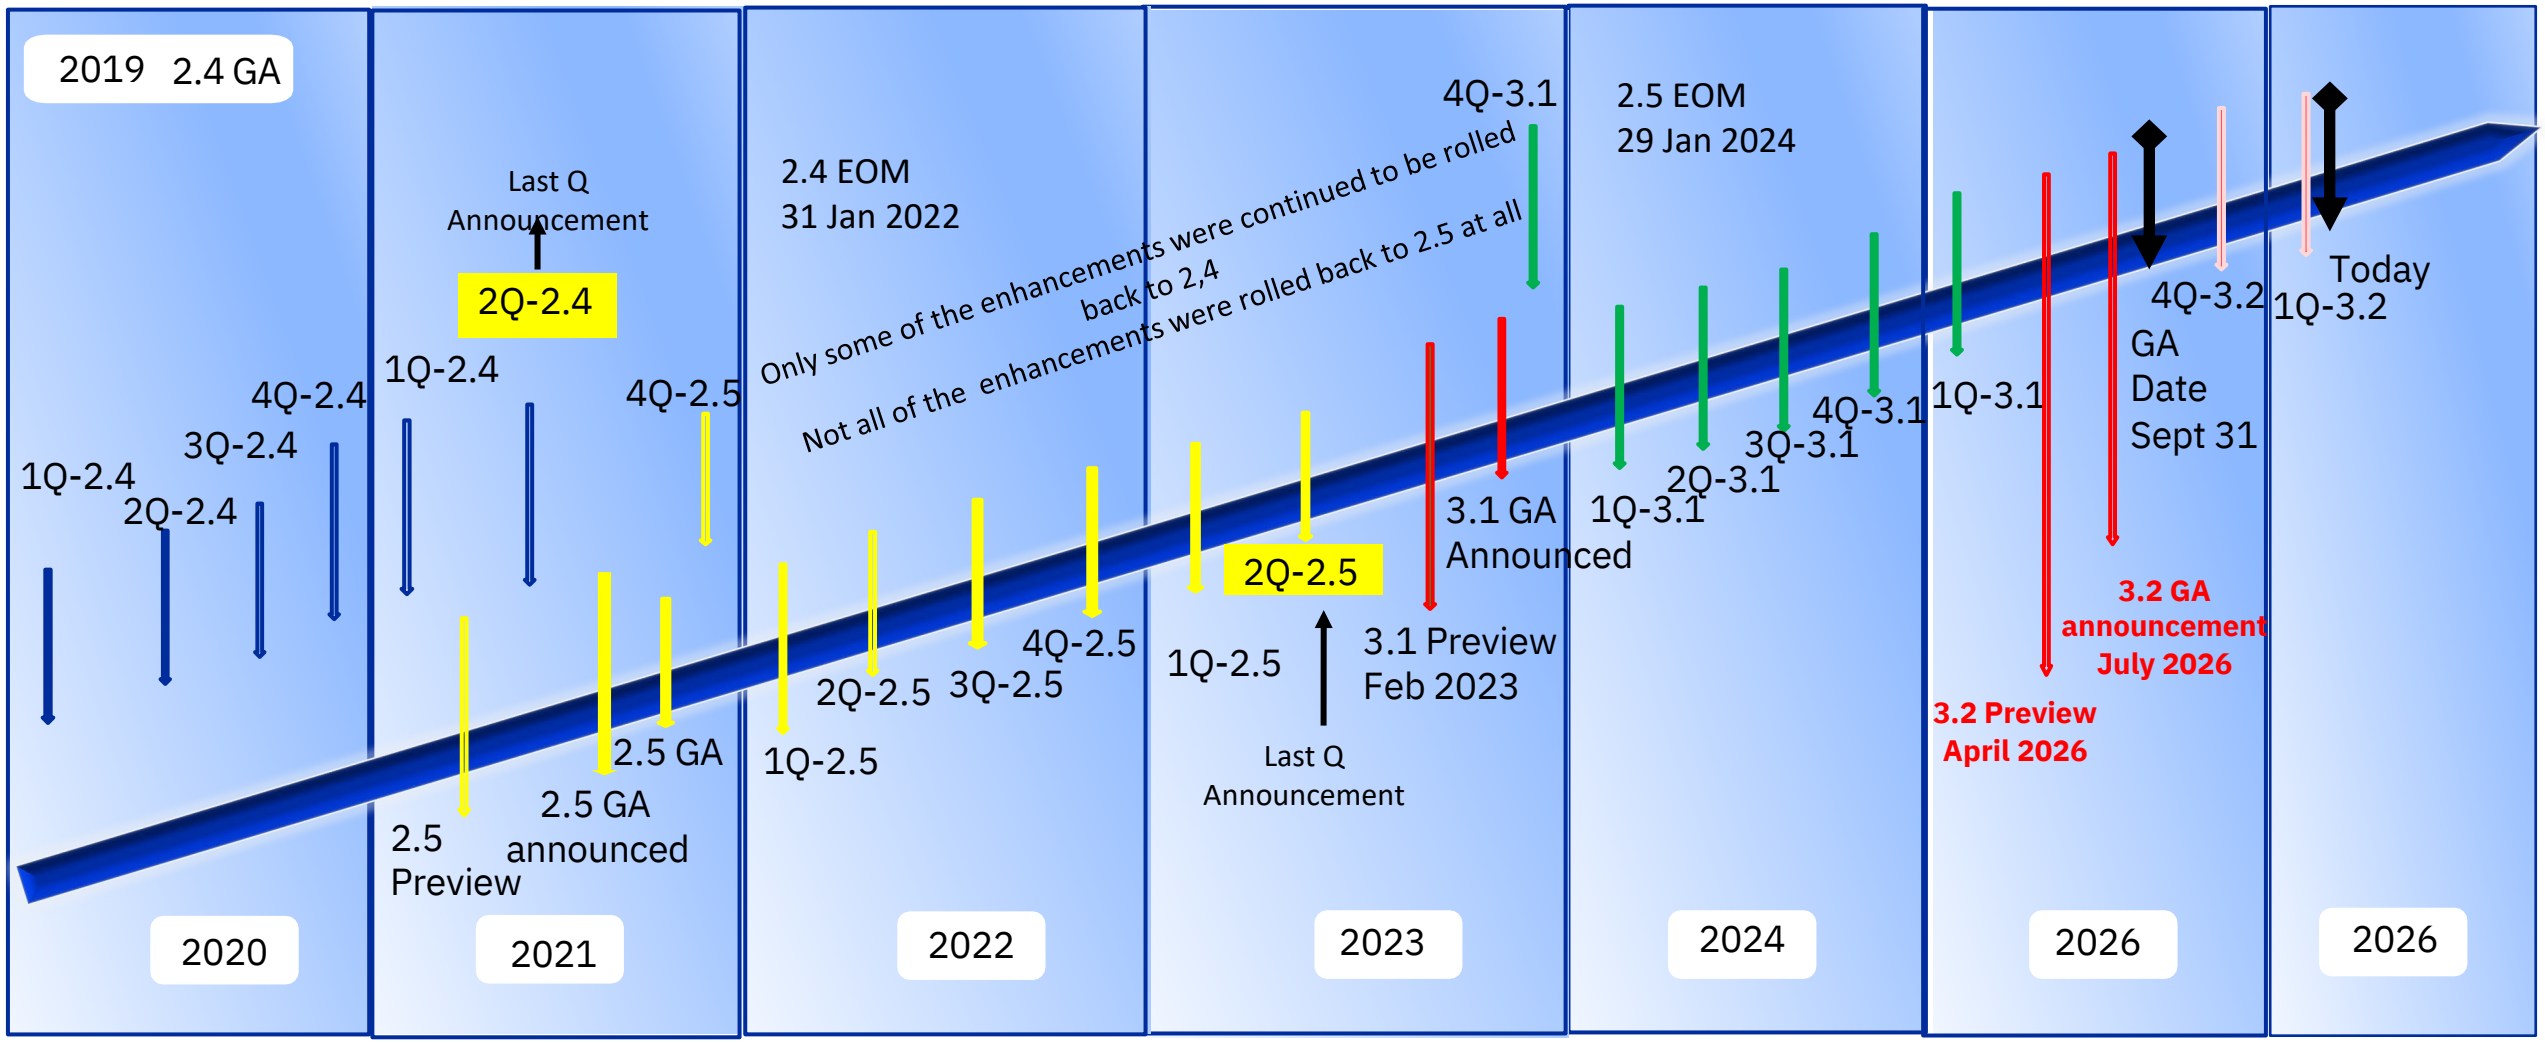

Summary of Common Issues

- z/OS, Sysplex and subsystem best practices / functions that impacts directly resiliency have not been implemented
  Check Appendix for list of majority of the ‘z/OS and sysplex best practices’
  Slide 13 shows some highlighted feature

More information will be address in today’s session: Come to session to hear more about these…

‘WSC z/OS System Programmer Hot Topics From Orlando’ Tuesday 02:30 – 03:30 EST Salon 13

Several best practices are not implemented, Common issues with % of frequency in all studies

- %95 z/OS Health checker exceptions not integrated into monitoring process. HC severities are lowered.
- %90 XCF functions : CFLRMGMT and several others not used.
- %70 SRB middleware stop/restart not used (WLM policy update), SRB dump support not implemented
- %80 SSDP not implemented
- %80 ARM usage is rare.
- %90 XCFCNFXX Logr monitoring not active
- %20 Traces are left active
- %80 VSMBESTFITCSA(YES) not coded
- USEV1R19Rules(YES) coded
- SWBSTORAGE(ATB) not used
- %95 VTOC_USE_ZHPF not used
- %10 , Long distance CF Adapter used for short distance
- %70 CF Performance not monitored, reported, tracked
- %50 Memory ‘minimum available frame queue’ not monitored

Common issues: z/OS Health Checker

- Z/OS Health Checker exceptions were not integrated into monitoring processes
- z/OS Health Checker severity levels lowered from high sev to low sev by clients.
- Critical z/OS HC exceptions were not recognized and actually they were the ones that warned client to prevent the outage they had experienced.
- One client had health checks(*,*) inactive in HC policy -> This means they are not using it.

Common issues: z/OS Health Checker

Of Health Checks Group By Owner

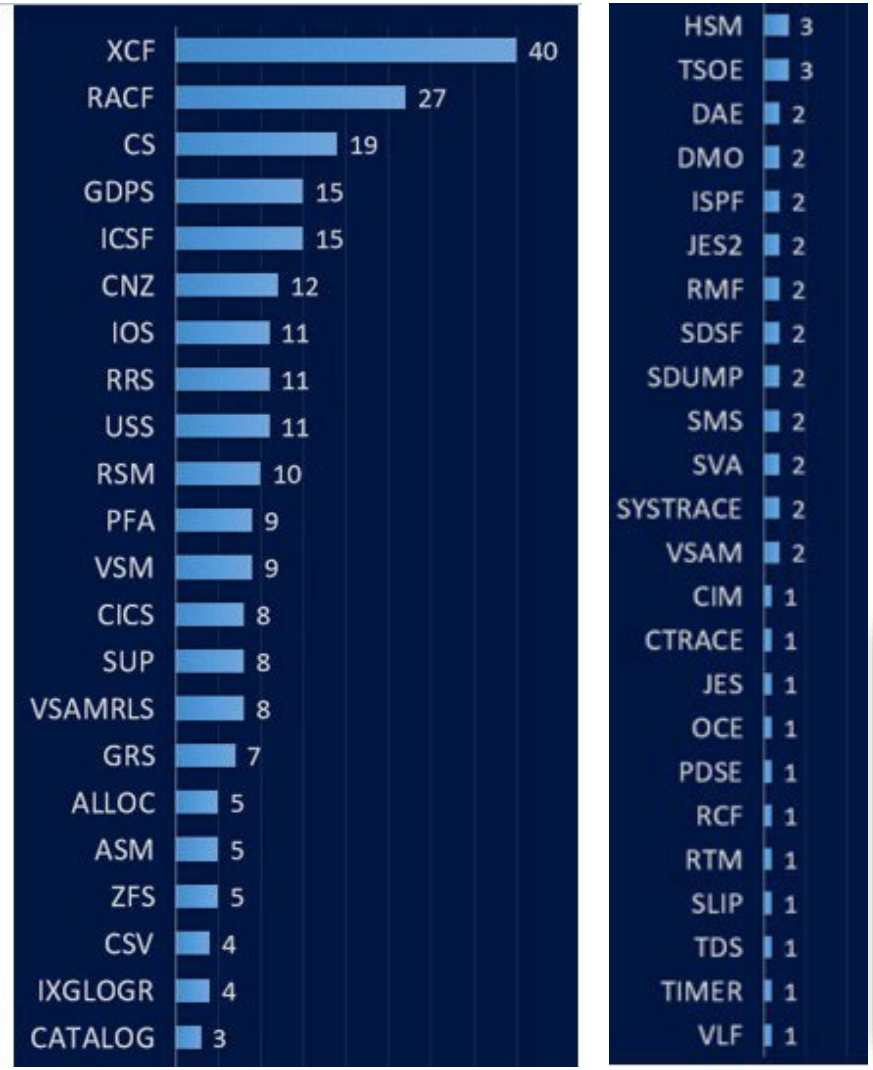

Close to 300 Base Health Checks exist that will help you improve your resiliency.

How to Print Health Checks Information in Batch Mode

SYS1.SAMPLIB(Hzsprint)

```
//HZSPRINT EXEC PGM=HZSPRNT,TIME=1440,REGION=0M,PARMDD=SYSIN
//SYSIN DD *
CHECK(*,*) ,EXCEPTIONS
//SYSOUT   DD SYSOUT=A,DCB=(LRECL=256)
```

CHECK(*,*)	Shows all HCs
CHECK(*,*), EXCEPTIONS	Shows HCs with exceptions

Common Issues : Questions to ask

1. Are you tracking IBM.Function.Health Checker fixcat ?
2. Are you monitoring HC alerts ? Do you have automation for HC alerts ?
3. When did you last check your policy to make sure you activated back HCs temporary deactivated before?
4. Did you deactivate or decrease the priority of critical HC – I know let me check that later -?
5. What is the default setting for new HCs according to your policy ?

New Health Checks

New in z/OS 3.2:

USS_ENVIRONMENT **WLM_SCLASS_SYSSTC**
WLM_OPT_PARM_CPENABLE

New in z/OS 3.1 Base and with CDs:
WLM_SCLASS_SYSSTC (OA62316) CD 1Q-2026 (z/OS V2R5 and z/OS 3.1)
ZFS_CACHE_PERFORMANCE(APAR OA63911, applies to z/OS 3.1 and V2R5) **CD 1Q-2026**
ZFS_EXCEPTIONS (APAR OA63911, applies to z/OS 3.1 and V2R5) **CD 1Q-2026**
ICSF_CLEAR_KEYS
ICSF_STATUS
IXGLOGR_LOCALBUFFERUSAGE (APAR OA64676, applies to z/OS V2R5 and z/OS 3.1)
RSM_FREEMAINEDFRAMES
SUP_ASVT_ABOVE_16M

Changed in z/OS 3.1:

RACF_PASSWORD_CONTROLS (added password phrase interval) VSAMRLS_QUIESCE_STATUS (APAR OA64048)

Changed in z/OS V2R5:

RACF_SENSITIVE_RESOURCES
XCF_TCLASS_CLASSLEN

- If you have statements that suppresses these new HCs, you can not see exceptions. Do not code statements to suppress by default.
- Update your check customization for modified IBM Health Checker for z/OS checks.
- Changes that IBM makes to the checks provided by IBM Health Checker for z/OS can affect any updates you might have made.

New in z/OS V2R5:
VSM_CheckRegionLoss
RACF_ADDRESS_SPACE
RACF_ERASE_ON_SCRATCH
RACF_PROTECTALL_FAIL
RACF_PTKTDATA_CLASS
RACF_SYSPLEX_COMMUNICATION
IOS_ENDPOINT_SECURITY_LCUPATHS
ZOSMIGV2R5_NEXT_CS_OSADLH
ZOSMIGV2R5_NEXT_CS_LSA
ZOS31MIG_SSH_CONFIG(APAR OA65071 applies to z/OS V2R4 and V2R5)
ZOS31MIG_SSHD_CONFIG(APAR OA65071 applies to z/OS V2R4 and V2R5)

Common issues : SRB middleware stop/restart is not used

z/OS System Recovery Boost Summary

Stage	Boost Class[2]	Description	Duration	Usage	Trigger
1	IPL Boost and Shutdown Boost z15, z16	IPL / Startup ShutDown GDPS® Enhancements[3] Standalone Dump	60 minutes At most 30 mins N/A Dump time or max 60 mins	Once per LPAR Once per LPAR N/A Speed boost only	IPL PROC IEASDBS GDPS Script IPL SADMP
2	Recovery Process z15, z16	Sysplex Partitioning Recovery CF Structure Recovery CF DataSharing Member Recovery Hyperswap Recovery	2 mins 1 min per structure 1 min per lock structure 2 mins	30 mins in 24 hours per eligible LPAR Shared Among Invocations Only 2 Reserved zIIPs brought online	Automatic Automatic Automatic Automatic
3	Recovery Process z16	SVC DUMP Middleware Start/Stop/Recycle Hyperswap load boost	2 mins[1] 5 mins 2 mins		CHNGDUMP RPBMINSZ= WLM Policy Automatic

[1] In order to see a benefit from zIIP Boost, you will need to turn on dump optimization, via the CHNGDUMP SET,SDUMP,OPTIMIZE=YES command.
[2] WLM will implicitly set all single-period importance 1 or 2 work as CPU Critical for all boost classes for duration of boost
[3] GDPS provides configuration and orchestration parallelization, no SRB related activities

New in Z17 : Dynamic I/O Activation SRB

SRB Enhancements

Boosts are: 1. Speed: subcap can run fullcap 2. zIIP: allowing workload onto zIIPs On by default in IEASYSxx BOOST=SYSTEM	SMP/E FIXCAT IBM.Function.SystemRecoveryBoost — Updated 10 April 2026										
	z16 and z17										
	z15							z16 and z17			
	System Recovery Boost			Recovery Process Boost* at MCL P46602.005 for IBM z15 Driver 41C (Bundle S29)				Recovery Process Boost*			
	IPL Startup	Standalone Dump (no zIIP boost)	Shutdown	Sysplex Partitioning – planned or unplanned removal	CF Structure Recovery – rebuild or duplex	CF Datasharing Member Recovery-disconnect or failure from locking resources	Planned/Unplanned HyperSwap	SVC Dump	STC start/restart	HyperSwap configuration load	Dynamic I/O Activate
Intended Duration	60 min	60 min	30 min	2 min	1 min	1 min	2 min	2 min	5 min	2 min	2 min
Basis of use	Auto	Auto	S IEASDBS	Auto	Auto	Auto	Auto	CHNGDUMP SET,SDUMP, RPBMINSZ threshold	WLM service definition BOOST attribute	Auto	Auto
z/OS 2.3	PTF	PTF	PTF	PTF	PTF	PTF	PTF				
z/OS 2.4	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	
z/OS 2.5	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	PTF	
z/OS 3.1	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	PTF
z/OS 3.2	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES

*Recovery Process Boost limited to 30 min per day per LPAR in aggregate. In V2.4 and higher, can be enabled or disabled with S IEASRB,CLASS=RP,REQ=DISABLE|ENABLE

GDPS provides configuration and orchestration parallelization in GDPS V4R2 and higher.

IBM Z Backup Resiliency

IBM Z Backup Resiliency provides a **smarter, better, faster** recovery process that can **prove your compliance** because it provides a **three-pronged data management approach** to know:

- How it's used
- Who's using it
- Where the data is located

IBM Z Backup Resiliency is also part of Cyber Resiliency Solution (Cyber Vault)

When implementing a robust cyber resiliency strategy to protect your enterprise against the increasing threat of ransomware and provide disaster recovery, IBM Z Backup Resiliency forms part of the solution to bring your production workloads back online by working together with technologies, such as IBM DS8000 Safeguarded Copy.

IBM Z Data Resiliency Manager

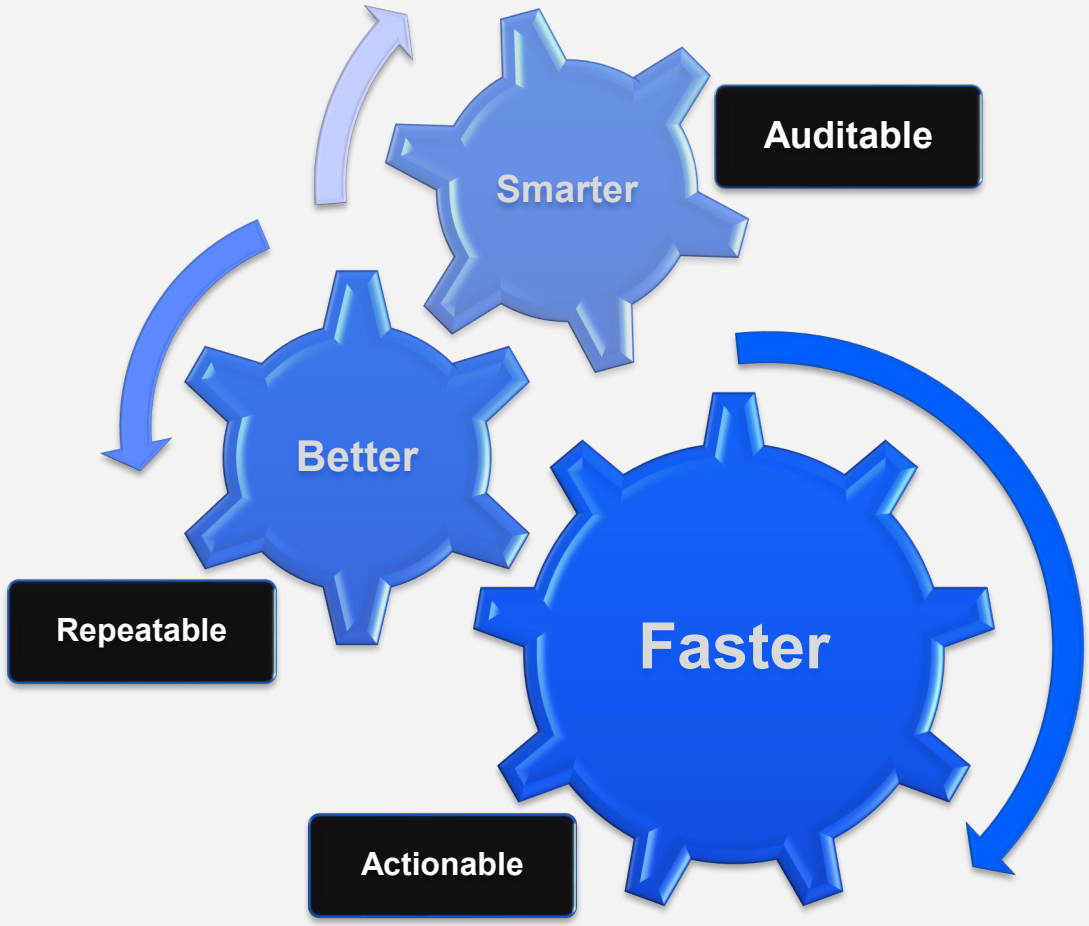

IZBR – z/OS "Data Aware" Data Resiliency Manager

Simplified, Centralized, Persona-Driven Data Resiliency

Robust Data Resiliency

- ✓ Identify critical data
- ✓ Identify missing backups
- ✓ Identify all downstream jobs that may need to rerun after a restore to complete recovery
- ✓ Forensics aids

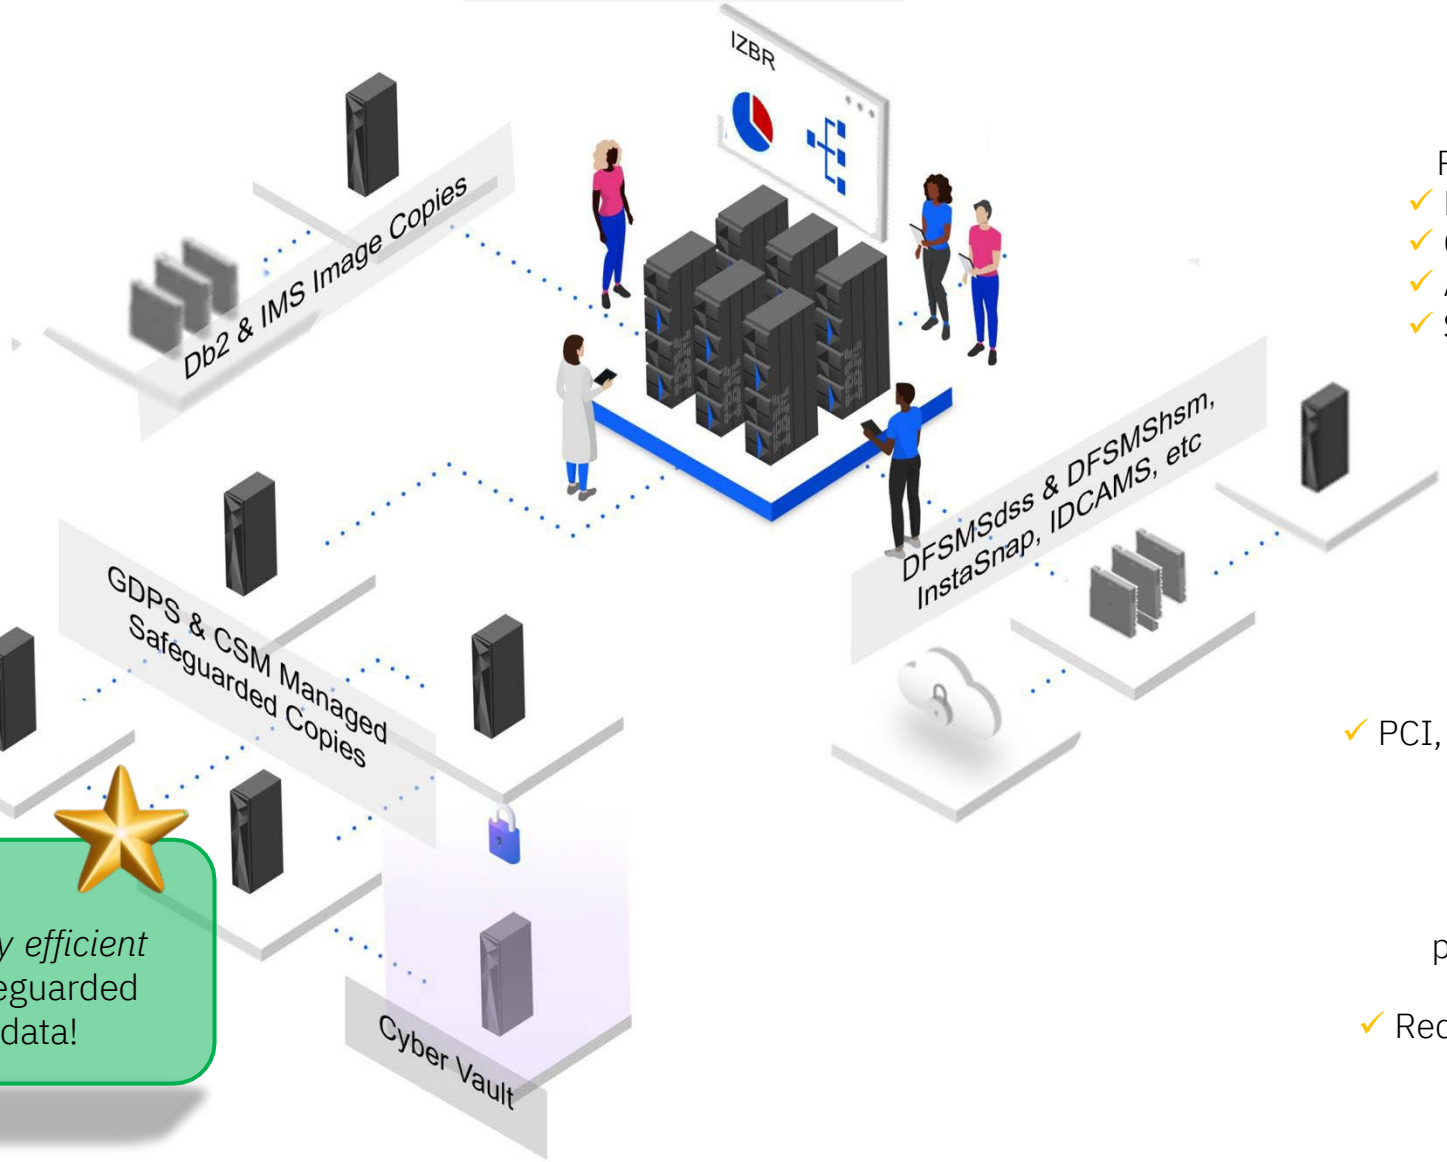

Zero Trust:
Today, IZBR provides *a very efficient* surgical recovery from Safeguarded Copy for *nonDatabase* data!

Single source of truth leveraged by *All Personas*

Persona based dashboards:

- ✓ LOB Owner
- ✓ Compliance Team
- ✓ Application Owners
- ✓ Storage Admins

Comply with regulations and audits

- ✓ PCI, DORA, Bank of England , etc.

Simplification Reduce

processing cost

- ✓ Reduce excessive backups, etc.

Slide :Glenn Wilcock

IZBR Persona-Based Dashboards

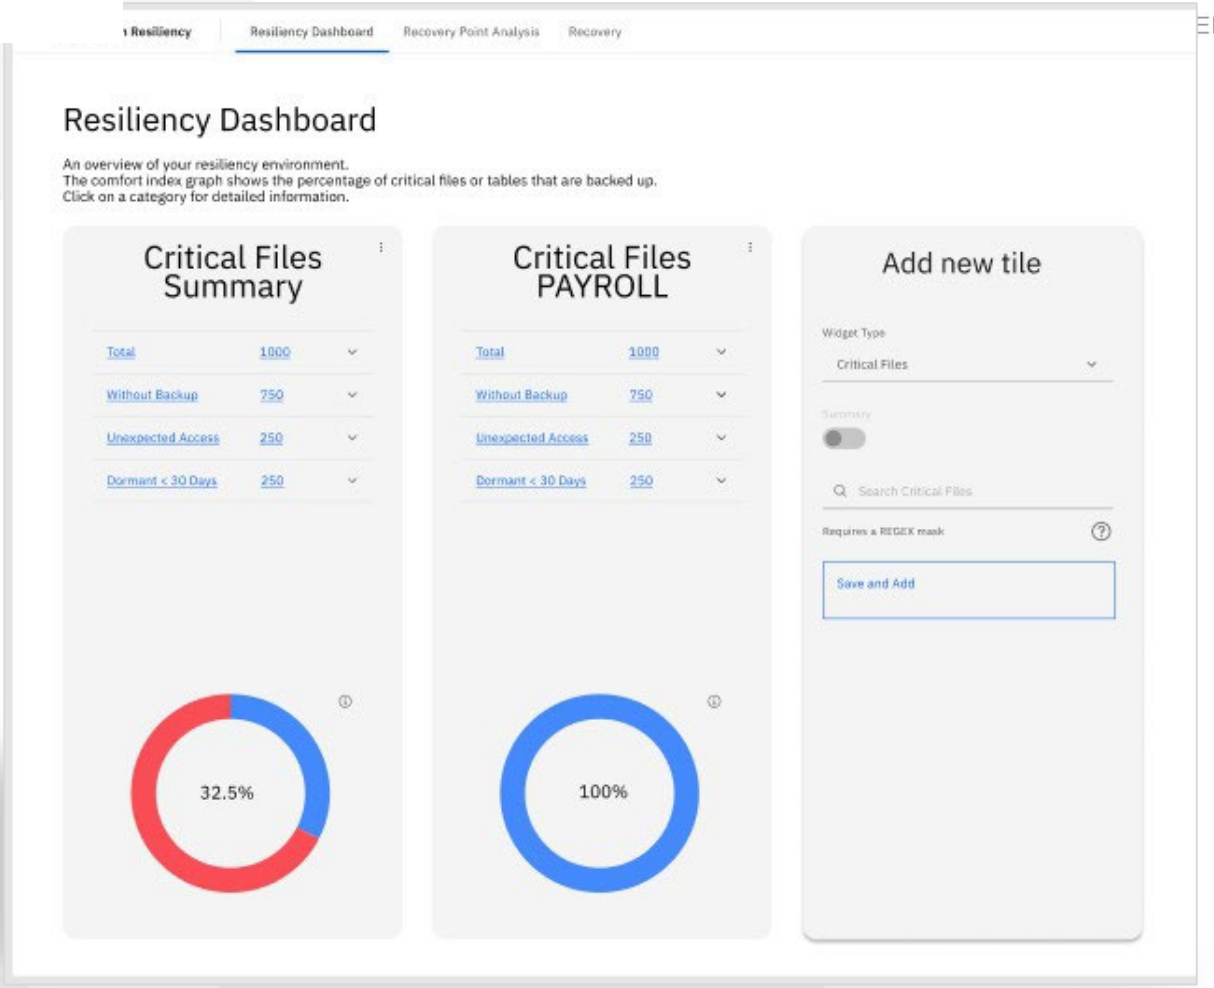

Single source of truth leveraged by *All Personas*

Persona based dashboards:

- ✓ LOB Owner
- ✓ Compliance Team
- ✓ Application Owners
- ✓ Storage Admins

IBM Z Backup Resiliency

IBM Z Backup Resiliency delivers high-value resiliency management of non-database managed data and applications by using detailed analytic reporting to provide insights to reduce manual approaches required to manage data outside of database control. By reducing dependency on domain expertise and time-consuming, error-prone analysis to determine the impact of data corruption incidents, application data inter-dependencies and vulnerabilities are determined quickly to reduce enterprise business risk.

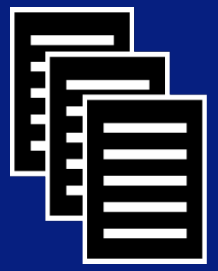

Auditable

Understanding of application data **interdependencies** and critical **recovery points**

Repeatable

Point-in-time, **granular Level** of application recovery of non-database systems

Actionable

Real-time comprehensive **inventory of datasets and associated backups** – pseudo log of batch applications

IBM Z Backup Resiliency

IBM Z® Backup Resiliency automates the analysis, backup and restoration of batch application data to provide operational resiliency and reduction of business risk.

It provides an inventory of data usage and backups and has the automation capabilities to recover batch data quickly for operational or disaster recovery events.

It can also generate reports that detail the usage and relationships of batch applications and data.

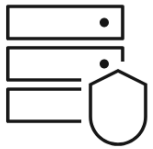

Impact of data corruption minimized
Address your business risks by recovering data from any corruption down to a granular level. Point-in-time recovery capability to non-database systems ensures near continuous data protection.

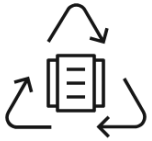

Data visibility and recovery
Get real-time visibility into jobs that are running and open data sets that might be at risk. Point-in-time inventory of data sets and backups enables faster, page-driven recovery.

Workload interdependency insight
Increase your understanding of application interdependencies and recovery points with resiliency and audit gap reports.

IBM Z Backup Resiliency

IZBR integrates near real-time analytics into the backup and recovery process to improve the speed and accuracy of your recovery decisions and processes for improved application availability.

Through analytics, IZBR can identify application-critical data sets (that is, data sets that must be recovered to restart an application from a point in time). IZBR understands the interdependencies between data sets, jobs, and applications, providing automated backup and recovery of non-database application data.

Once critical data sets are identified, IZBR monitors them using real-time SMF analysis of both the batch jobs and the backup utilities. Batch job analysis is used to track which jobs have used the critical data sets, while backup utility analysis creates and maintains a list of current backups for use in the event of a recovery. IZBR centralizes tracking of backup data, integrating all information into one repository.

Balanced solutions designed to address different requirements

GDPS Metro	GDPS Global	GDPS Metro Global	GDPS Continuous Availability
Near-continuous availability and recovery at metro distances	**Disaster recovery at extended distance**	**Near-continuous availability regionally & recovery for 3-4 sites**	**Near-continuous availability, recovery & workload balancing**
Systems remain active Multisite workloads can withstand site and storage failures	Rapid systems DR with "seconds" of data loss	Metro near-continuous availability and out of region disaster recover	Continuous availability at unlimited distances
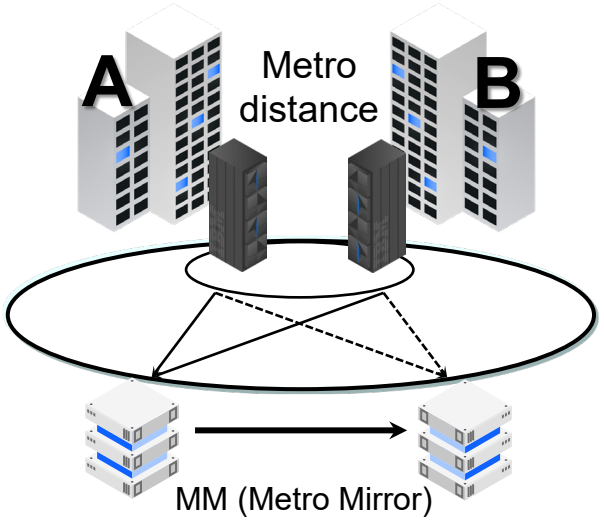	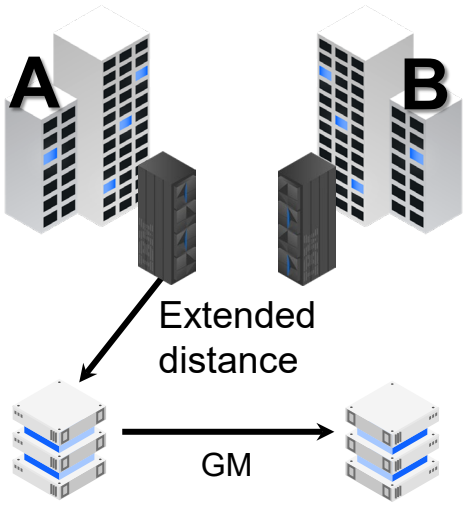	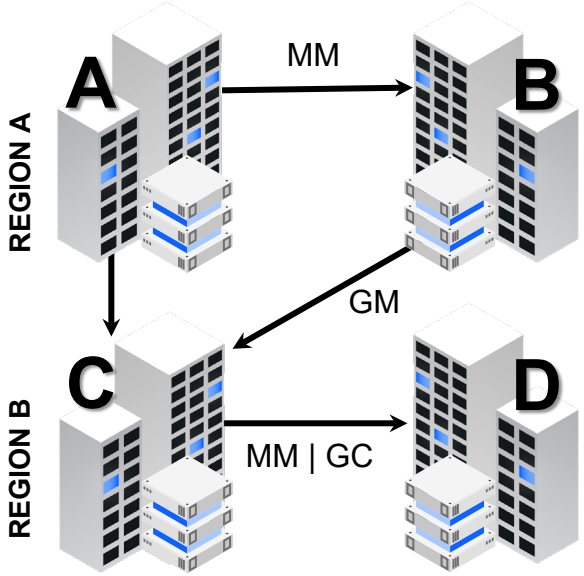	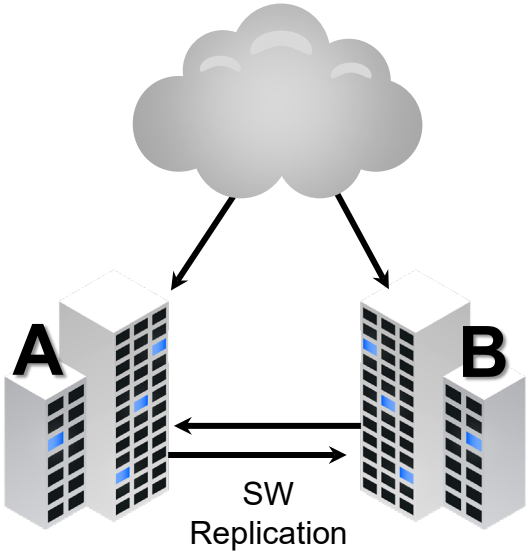
RPO 0 & RTO <60 min	***RPO 3-5 sec & RTO <60 min***	***RPO 3-5 sec & RTO <60 min***	***RPO 3-5 sec & RTO <60 sec***

IBM DS8000 Data Resilience

When bad things happen to your data - protection from failures, disasters and cyber attacks

Flash Copy

Point in time copy

Within the same system

Instantly accessible for fast recovery

Safeguarded Copy

Secure Immutable point in time copy

Within the same system

- Secure Immutable, not erasable
- Hidden to hosts and applications

Metro Mirror

Synchronous mirroring

Primary Site A — Metro distance Site B

- 303 Km; 500 Km with RPQ
- Performance improvements
- 50% less infrastructure

Global Mirror

Asynchronous mirroring

Primary Site A — Out of region Site B

Metro / Global Mirror

3/4/5/6 site cascaded or multi-target synchronous and asynchronous mirroring

Primary Site A — Metro distance Site B — Out of region Site C

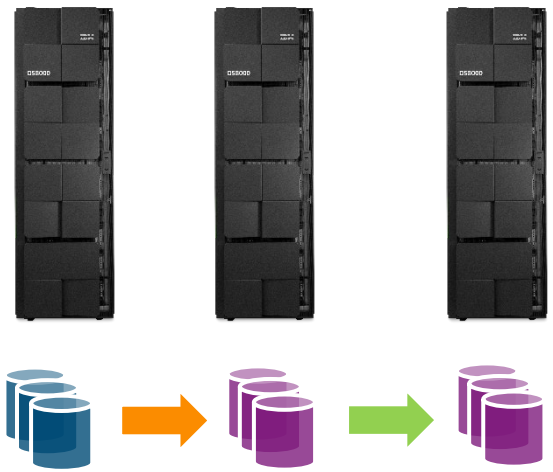

- Up to 50% bandwidth reduction over competitive solutions
- From 1 to 3 seconds RPO
- Superior application failover via CSM, GDPS or PowerHA

Additional Information

Web sites:

- GDPS https://www.ibm.com/products/gdps
- IBM zSystems https://www.ibm.com/z
- IBM zSystems Resiliency https://www.ibm.com/z/resiliency
- Storage https://www.ibm.com/storage
- Redbook – GDPS Family: An Introduction to Concepts and Capabilities http://www.redbooks.ibm.com/abstracts/sg246374.html?Open

GDPS Web site resources

- GDPS: The Enterprise Continuous Availability / Disaster Recovery Solution white paper
- GDPS pre-requisite information
- GDPS training schedule links
- GDPS hardware qualification letters
- E-mail: gdps@us.ibm.com

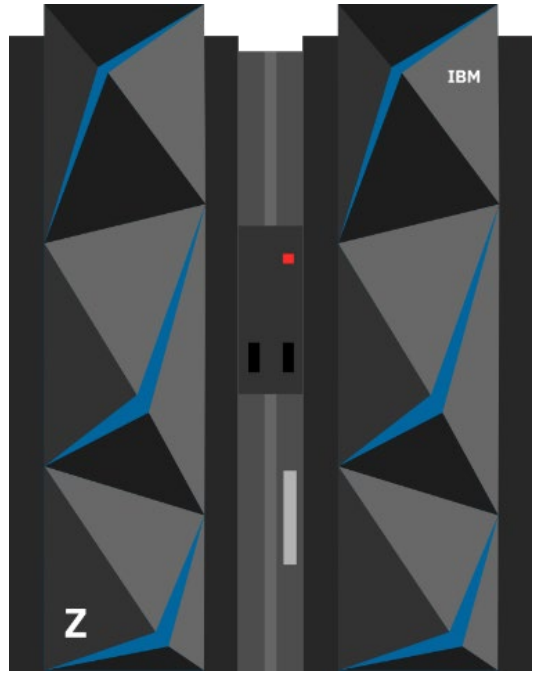

AIOps Offerings

AIOps Software	Target
IBM Z Anomaly Analytics	Enable problem identification, isolation and resolution on IBM Z through analysis of structured and unstructured operational data.
IBM Z Operational Log and Data Analytics	Gain insights from SMF records, RMF, z/OS System log, syslogD, Joblog output, application logs (CICS and Websphere) and generic files
IBM Z Application Performance Management Connect	Track transaction information from z/OS subsystems to APM solutions.
IBM Operational Decision Manager	Discover, capture, analyze, automate and govern rules-based business decisions on premises or on the cloud.
IBM DB2 Analytics Accelerator for z/OS	Leverage real-time insight from data at the point of origin.
IBM DB2 AI for z/OS	Enhance database performance with machine learning.
IBM Z Data Privacy for Diagnostics	Leverage machine learning to detect and redact PII from diagnostic dumps.
IBM Cloudpak for AIOps	Deploy advanced, explainable AI across the ITOps toolchain.

IBM z Content Solutions Center

Single Engaging place for everything you need to understand and use functions and products!

https://www.ibm.com/support/z-content-solutions/

Automation & Management

- ServerPac Installation Using z/OSMF
- IBM z/OS Change Tracker
- Red Hat Ansible Certified Content for IBM Z
- Software update with z/OSMF
- z/OS Management Services Catalog

Optimization

- Cloud Infrastructure Center
- Cloud Provisioning and Management for z/OS
- Integrated Accelerator for zEDC
- Journey to LinuxONE
- Journey to sustainability with IBM LinuxONE
- **Resilience**
- System Recovery Boost
- Tailored Fit Pricing for IBM Z

Modernization

- Automating and shift-left testing for z/OS hybrid applications
- Continuous delivery and deployment
- Continuous integration for the hybrid cloud developer experience
- Discover and plan for z/OS hybrid applications
- EzNoSQL for z/OS
- IBM Z and Cloud Modernization Stack
- IBM Z Distribution for Zowe
- Z Digital Integration Hub
- z/OS Cloud Data Access
- z/OS Connect
- z/OS Container Platform
- z/OS Container Extensions (zCX)
- zCX Foundation for Red Hat OpenShift

Security

- ServerPac Installation Using z/OSMF
- IBM z/OS Change Tracker
- Red Hat Ansible Certified Content for IBM Z
- Software update with z/OSMF
- z/OS Management Services Catalog

Prediction

- AI Infusion into z/OS
- IBM Open Data Analytics for z/OS
- Journey to AI on IBM Z and LinuxONE
- Journey to open data analytics
- Machine Learning for IBM z/OS
- Spyre Accelerator for IBM Z and LinuxONE

Currently in IBM z Content Solutions

SFM (Sysplex Failure Management) Policy Settings: Best/Proven Practices

CFSTRHANGTIME (SFM Policy)

Specifies the time interval, in seconds, that a coupling facility structure connector can remain unresponsive before the system takes action to relieve a hang in a structure-related process
Default: NO Value : 900

MEMSTALLTIME(SFM Policy)

Number of seconds after which XCF is to take action to resolve XCF signaling stall conditions. MEMSTALLTIME(n), N=n seconds. Default is NO, then N=MAX(FDI, 120 seconds). Value: 300

Confail(SFM Policy)

It indicates whether SFM is to handle signaling connectivity failures for the sysplex. Default: Yes
Value: Yes

Weight (SFM Policy)

Value based on client Request

PROMPT(SFM Policy)

Do not specify, use Isolatetime instead.

SSUMLIMIT(SFM Policy)

Indicates the length of time a system can remain in the state of not updating the heartbeat and sending signals(the amount of time a system will remain in a “semi-sick” state) Once the SSUMLIMIT has been reached the specified action will be initiated against the system (Isolatetime(0))
Default: None Value: 900

ISOLATETIME(SFM Policy)

Use Isolatetime(0) . Default is proven practice.
Tells that the system is to immediately partition an unresponsive system

Implement System Status Detection Partition Protocol (SFM with BCPii)

XCF exploits BCPii services to Detect failed systems, Reset systems
Improved availability by reducing duration of sympathy sickness,Eliminate manual intervention in more cases, Potentially prevent human error that can cause data corruption

Sysplex Failure related items CoupleXX : Best/Proven Practices

Cleanup (CoupleXX member)

Cleanup interval is the maximum number of seconds allowed for members of a group to clean up their processing before the system is put into a wait state. Value: 15 (Also default)

MSGISO (Default Enabled)
XTCSIZE (Default enabled)

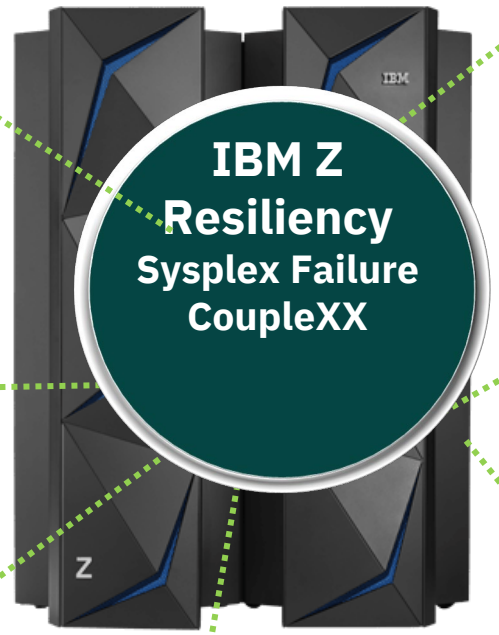

FDI (CoupleXX member)

User Interval Value : 165
Spin Interval Value : 165
User OpNotify Value : 3

CFLCRMGMT(CoupleXX)

LossConn Recovery management is enhancement of MSGBASED protocol and it enhances average CF LOSSCONN recovery time by processing CF structures serially, rather than in parallel during loss of CF connectivity.
Default: Disabled Value: Enable

Critical Paging (CoupleXX member)

It protects critical pages from being paged out when system is experiencing a critical real frame shortage(31-bit CS, Critical for paging defined ASes including all Dses,PLPA) Important requirement for GDPS clients
Default: Disabled Value: Enable

CFSTRQMON (CoupleXX)

CF requests that are delayed longer than expected trigger messages that warn of a potential problem. (IXL053E, IXL054I, IXL055I, and IXL056I)
Default: Disabled Value: Enable

CFMONOPAVOID (CoupleXX)

CF informs to z/OS when the structure is consuming a disproportionate share of CF resources. That z/OS instance responds to the CF notification by limiting CFrequests for that structure until the monopolization condition clears. Default: Disabled Value: Enable

DUPLEXCFDIAG(CoupleXX member)

When the completion of a duplexed coupling facility request is delayed, software and hardware problem determination data is to be collected.
(Cases where coupling facility might break duplexing for the structure that is associated with the delayed command)
Default: Disabled Value: Enable

Experience more with IBM

Visit us at the IBM Booth #113

After a full day of technical sessions, take a break with us!

Connect with our experts, snap a photo with the z17 Plexi or the latest Telum II, and get an up-close look at our Spyre Accelerator.

Come back each day for fresh topics and demos at our expert stations.

Think 2026

Join 5000+ senior business and technology leaders who are seizing the AI revolution to unlock unprecedented growth and productivity at Think 2026.

Find out more information using the QR code below.

IBM Digital Asset Haven

IBM Digital Asset Haven is the operational backbone for financial institutions and regulated enterprises entering the digital asset economy.

Find out more information using the QR code below.

Your feedback is important!

SHARE
EDUCATE ▸ NETWORK ▸ INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation