What's New in DFSMS Orlando 2026

Monday, February 23, 2026

Barbara McDonald, IBM z/OS DFSMS Product Manager (bawhite@us.ibm.com)
Alexis Kapica, IBM z/OS DFSMS Product Manager (alexis.kapica@ibm.com)

Storage Management Track

SHARE
EDUCATE • NETWORK • INFLUENCE

- The **Storage Management Track** (formerly called the MVSS Project) focuses on the **education**, **professional networking**, **& industry influence** in the area of IBM Z Storage Management and the facilities and products necessary to support these functions.
- In z/OS environments, we discuss all of the components of DFSMS, and other vendor developed (OEM) packages used to provide supporting services. Among the many subjects covered are strategies used for storage management, maintaining data integrity, cloud storage, data backup/recovery, and data archival solutions.
- The track works closely with the Hardware Track (formerly the SAN/DISK/TAPE project) to address topics related to storage hardware.
- The project is always open, and new members are encouraged to provide continued and active participation. See one of the project officers or IBM representatives for more information.

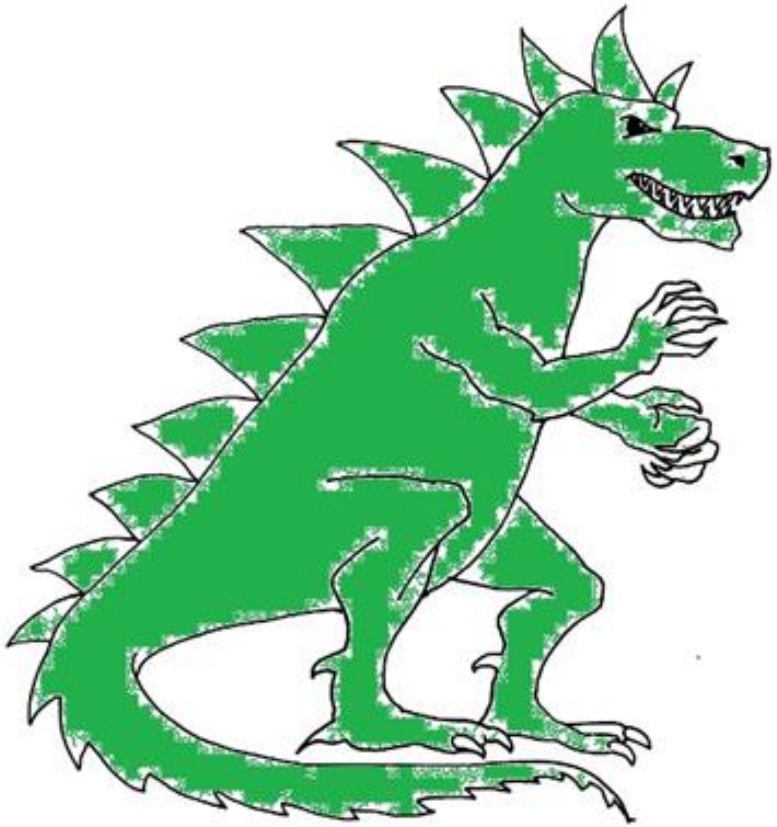

Storage Management Track
-- *SHARE Project Officers*

SHARE
EDUCATE • NETWORK • INFLUENCE

Project Officers	Title
Russell Witt Broadcom	Project Manager
Janet Sun Sun Coast	Deputy Project Manager / Scheduling
Tom Conley BMC	Deputy Project Manager / Scheduling
Edward Petka TIAA	Deputy Project Manager / Requirements
Marty Hasegawa Rocket Software	Deputy Project Officer / Volunteer Coordinator
John Hogan 21CS	Project Officer
John Crossno Phoenix SW	Project Officer
Chris Taylor	Project Officer

IBM DFSMS and Storage Team at SHARE this week!

SHARE
EDUCATE • NETWORK • INFLUENCE

Alexis Kapica
z/OS DFSMS Product Manager

Barbara McDonald
z/OS DFSMS Product Manager

Glenn Wilcock
CPO, STSM DFSMS

Robert Gensler
STSM, DFSMShsm
Product Owner

Alex Hong
DFSMS VSAM RLS Development

Frank McCune
Catalog / IDCAMS L2
Technical Support Lead

Trevor Geisler
DFSMS LS PDSE

Vanessa Sobik
TS7700 Development Manager

Shreya Nair
Infrastructure Product
Management

Beth Peterson
STSM - DS8000 Storage
CKD I/O Architecture / Dev

How do you envision the future of your z/OS storage management environment?

Share your ideas and experience with us! Come talk with us here at the event or send an email to the contacts if you're interested in being a part of the next evolution of storage management!

Storage Management Track

- **SHARE is a Volunteer organization**
- **Volunteer to CHAIR a session**
  - You get the best seat in the house
    - Talk to one of the project leaders if you're interested
- **Submit a topic for a presentation**
  - Call for Presentations for **Pittsburgh, PA** (August 16-20) opens soon. **Stay tuned to SHARE for more details!**
  - If you don't feel you can talk for an hour, submit the session anyway.
    - We may be able to find someone (other user, possibly a Vendor) that can also add some discussion on the same topic.
- User presentations / experiences are **ALWAYS** welcome and receive high schedule priority!
  - Contact the **project officers** for additional information and **Barb McDonald** (bawhite@us.ibm.com) if you might be interested in a joint user experience presentation with IBM.

Storage Management Collaborate and Socialize

- **Storage Management Free-For-All**
  - **Wednesday at 5:00PM in Salon 15.**
  - Everyone is welcome!
  - This session is a great opportunity to ask any question you might have about DFSMS and/or Storage Management or even share something you've done recently that might be a time saver or help to others.
  - IBM Development reps and SMEs, vendors, and clients typically attend to help answer questions.

Calling New to Z attendees…

z/OSMF Storage Management Plugin - SMS Like Never Before, Wed 8AM, Salon 15

Storage Management Collaborate and Socialize

Storage Management Project Dinner Wednesday 6:30PM

- Following the Free for All, meet in the Rosen Centre lobby at 6:15PM

The Pub

- 9101 International Drive – approx. 0.7 miles (15-20 min walk)
  - *The Pub features classic British pub fare along with some American favorites.*

The Pub menu:

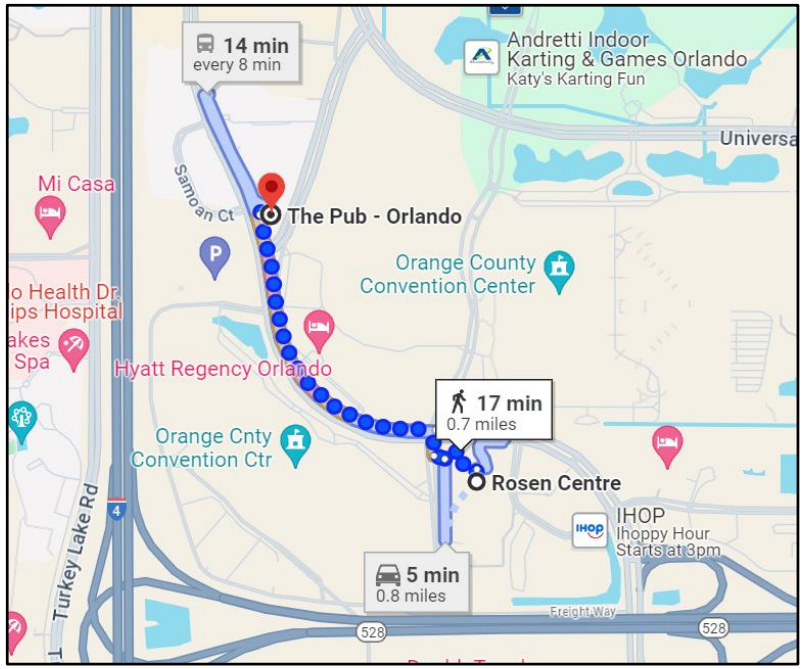

DFSMS: Data Facility Storage Management Subsystem

Storing, managing, protecting, and serving data on IBM Z®

SHARE
EDUCATE • NETWORK • INFLUENCE

DFSMS is the Data Subsystem for z/OS

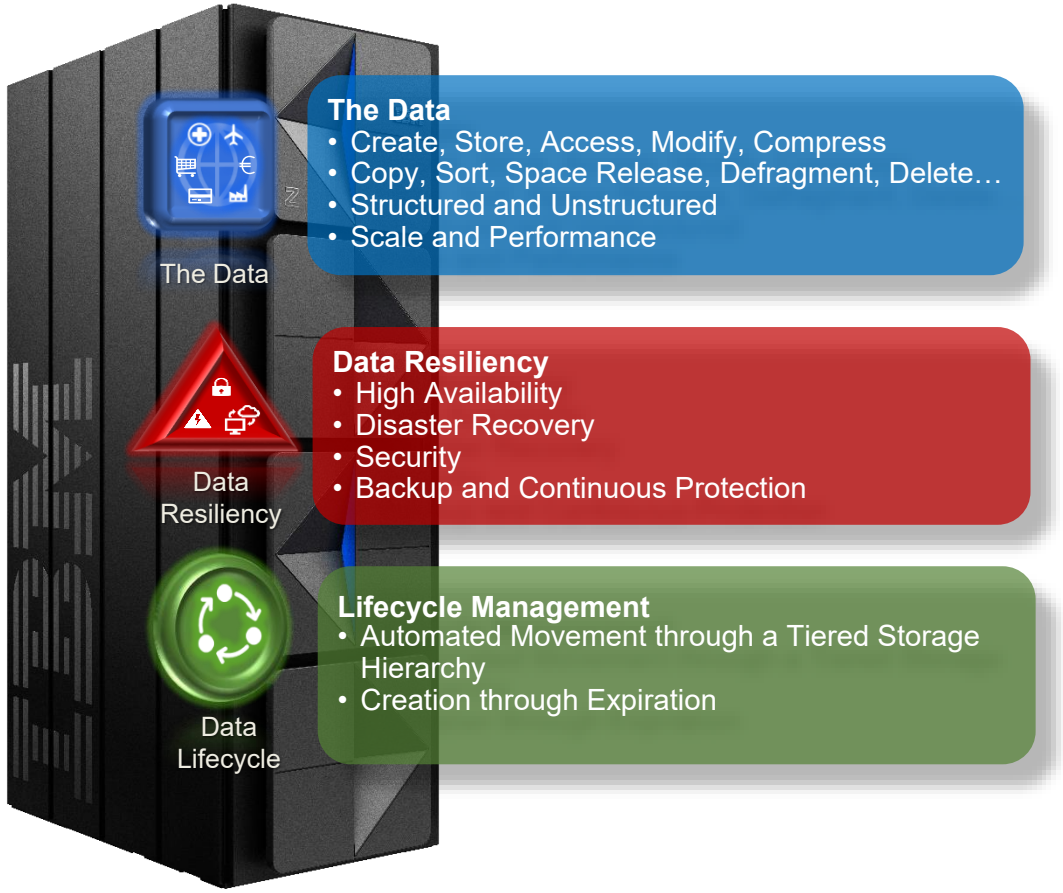

- **THE DATA supports business critical applications and modernization efforts by providing constraint relief, modernization, scalability, and performance of storage and data**
  - Manage exponential growth of structured and unstructured data
  - Leverage industry standard data stores, like NoSQL, and APIs to store, access and protect data within and across platforms
  - Create a resilient, modern infrastructure that integrates mainframe data, operations, and applications with hybrid cloud environments
- **DATA RESILIENCY strengthens business continuity by exploiting new opportunities and advancements in cyber resiliency**
  - Address increased security and compliance requirements
  - Enable 24x7 availability of systems and data to applications
  - Provide simpler, faster, and more reliable recovery operations
  - Provide trusted logical corruption protection
- **DATA LIFECYCLE enables utilization and management of seamlessly integrated online and offline tiers with data movement between the tiers, including Cloud**
  - Provide automated, policy-driven Lifecycle Data Management
  - Enable cloud object storage as an additional tier

Session Objectives

Cloud Storage

- CDA Compression
- Cloud Compatible Data Sets
- CDA Immutable Objects
- DSS and HSM Direct to Cloud

Simplify and Modernize

- z/OSMF Storage Management Plugin
- DFSMSrmm (RMM) z/OSMF Plugin Enhancements
- DFSMSrmm customer requirement enhancements
- EzNoSQL COBOL support and Python APIs
- Object Access Method (OAM) REST APIs
- Out-of-space avoidance enhancements

Security

- Dataset Encryption
- Granular control for data set encryption enablement
- Optimized encrypted data set copy
- VSAM GSR support for encrypted data sets

Blue = continuous delivery

Announcements

- **Next phase of IBM Storage DS8000 (10th Generation) family for the future, announced** August 26, 2025
  - *The IBM Storage DS8000 (10th Generation) family announced in September 2024 is recognized for innovations in performance, scale, security and synergy. With this announcement, the IBM Storage DS8000 (10th Generation) family grows to include two new models, aiming to continuously set the standard for modern, mission-critical enterprise storage solutions. We offer the latest-generation DS8000 technology to meet the needs of any enterprise storage clients – from rack mount and racked single frame systems to the highly scalable multi frame systems – while leading enterprise data storage's transition to the new era of cyber resilience and AI.*
  - ***DS8000 10th Generation R10.1 release*** blog post (Nick Clayton)

Find Out More…

- **What's New On the IBM DS8000 G10 Storage Family,** Tues 8AM

13

Announcements

- **IBM TS7700 R6.0**, announced October 22, 2024
  - *IBM TS7700 Release 6.0 builds upon and expands the capabilities of earlier TS7700 offerings while extending enterprise storage system leadership for IBM Z virtual tape systems. The TS7780 introduces the next generation hardware platform along with TS7780 R6 microcode that will be the foundation for the next several years. It is designed to provide better capacity efficiencies, higher performance for active data while reducing the expenses associated with archive, data retention, and backup operations behind IBM Z servers. Built with the most advanced POWER 9+ processors and 32GB FICON, it benefits from years of trailblazing research and deep collaboration between the IBM Storage and IBM Z teams to deliver unique business value for mainframe deployments.*
  - **The next generation VTS from IBM: The TS7780**
  - **IBM TS7700 R6.1**, announced August 12, 2025

Find Out More…

The Brand New TS7700-Taming the TBs to ExBs of IBM zSystems Data, Wed 3:45PM

14

Announcements

- ***IBM z/OS 3.2 unlocks the value of IBM z17***, dated July 22, 2025
  - Object Access Method (OAM) REST APIs
  - EzNoSQL COBOL support and Python APIs
  - z/OSMF Storage Management Plugin
  - DFSMSrmm (RMM) z/OSMF Plugin Enhancements
  - CDA Cloud Compatible Data Sets
  - CDA Compression
  - DFSMSdss (DSS) and DFSMShsm (HSM) Direct to Cloud
  - Out-of-space avoidance enhancements
  - Granular control for data set encryption enablement
  - VSAM GSR support for encrypted data sets
  - Optimized encrypted data set copy
  - DFSMSrmm customer requirement enhancements
  - Updated zHyperLink write statistics
  - Updated DFSMSdfp Checkpoint/Restart SOD

*"**The introduction of Python EzNoSQL APIs and cloud data access through DFSMSdfp and DFSMS Object Access Method is a shrewd move**. It acknowledges that mainframe data does not exist in a vacuum; it needs to be accessible to a broader ecosystem of developers and modern applications. By making mainframe data available via standard cloud APIs and supporting popular languages like Python for NoSQL database access, IBM is working to lower the barrier to entry for a new generation of developers… **This is an essential step towards modernizing how enterprises interact with their most valuable data assets**."*

*"Efficiency and automation enhancements in z/OS 3.2 are also significant… REST APIs for configuration and software updates, along with **simplified storage management UIs**, indicate a strong push towards reducing operational complexity. Mainframe administration has historically required deep specialized skills, and these updates **suggest IBM is trying to make the platform more approachable for a wider range of IT professionals**. "*

-- ***z/OS 3.2: Is Mainframe AI the Enterprise Game Changer?***
Steven Dickens, CEO HyperFRAME Research

Blue = continuous delivery

Announcements

Look for more z/OS continuous delivery enhancements every quarter on the z/OS Community page. To stay up to date, please 'Join Community' to ensure you don't miss an announcement!

★ **Statements of direction: z/OS,** published February 17, 2026

Strategic initiative to help simplify the z/OS user experience - *IBM intends to deliver a simplified z/OS user experience based on the needs of next-generation mainframe professionals and adapted to how they learn, work, and develop new skills.*

By integrating AI-driven insights and enabling the use of industry-aligned tooling, the new z/OS user experience is designed to make z/OS more intuitive and accessible, while preserving the unique value and capabilities of z/OS. ***IBM Project Polaris*** *will consist of a stream of early technical preview shipments that will be made available for experimental access to new and enhanced functions.*

- ***IBM z/OS 3.2 4Q25 enhancements,*** *posted November 20, 2025*
  - Security checking for VSAM data sets
  - DFSMSdss (DSS) Direct to Cloud enhancements
  - EzNoSQL Enhanced ResultSet Java APIs
- **IBM z/OS 3.1 1Q25 enhancements,** posted February 18, 2025
  - DFSMS Cloud Object Storage Enhancements
  - DFSMS Optimization Mode for Data Set Copy
- **IBM z/OS 3.1 4Q24 enhancements,** dated November 19, 2024
  - EzNoSQL Python API support
  - TLS 1.3 Support for Transparent Cloud Tiering (TCT)
  - Cloud Data Access (CDA) enhancements

Coming soon - z/OS 3.2 1Q26 enhancements!

Blue = continuous delivery
Hyperlink = backup

17

Connect with DFSMS on Social Media

Join the IBM DFSMS Community and Mainframe Data Management LinkedIn Group!

Participate in the System Z Enthusiasts Discord

SHARE
EDUCATE • NETWORK • INFLUENCE

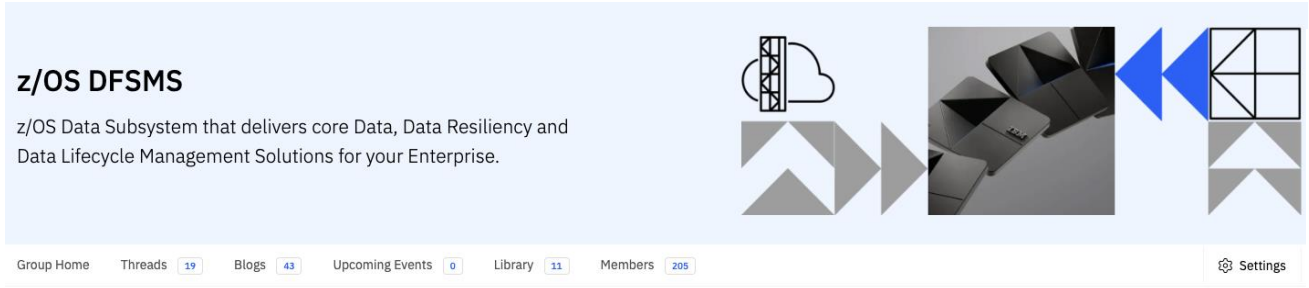

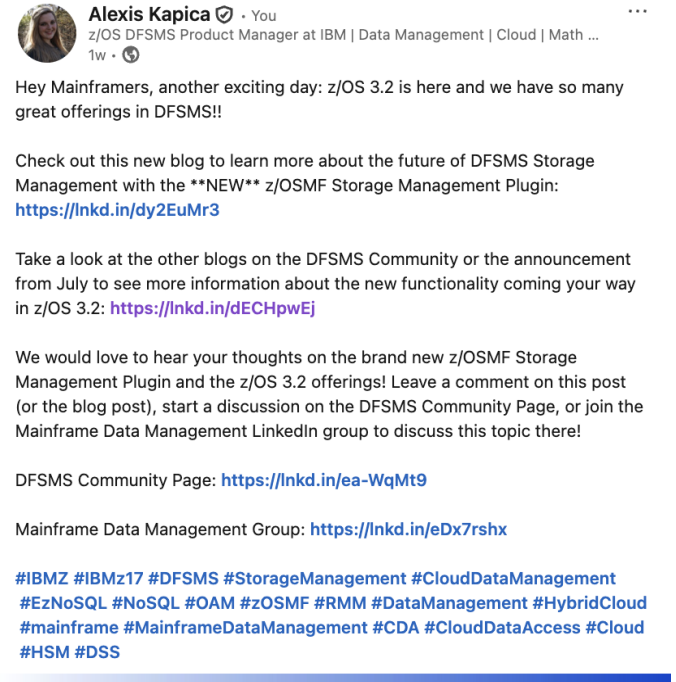

Announcing the brand-new z/OS DFSMS Guild!

SHARE
EDUCATE • NETWORK • INFLUENCE

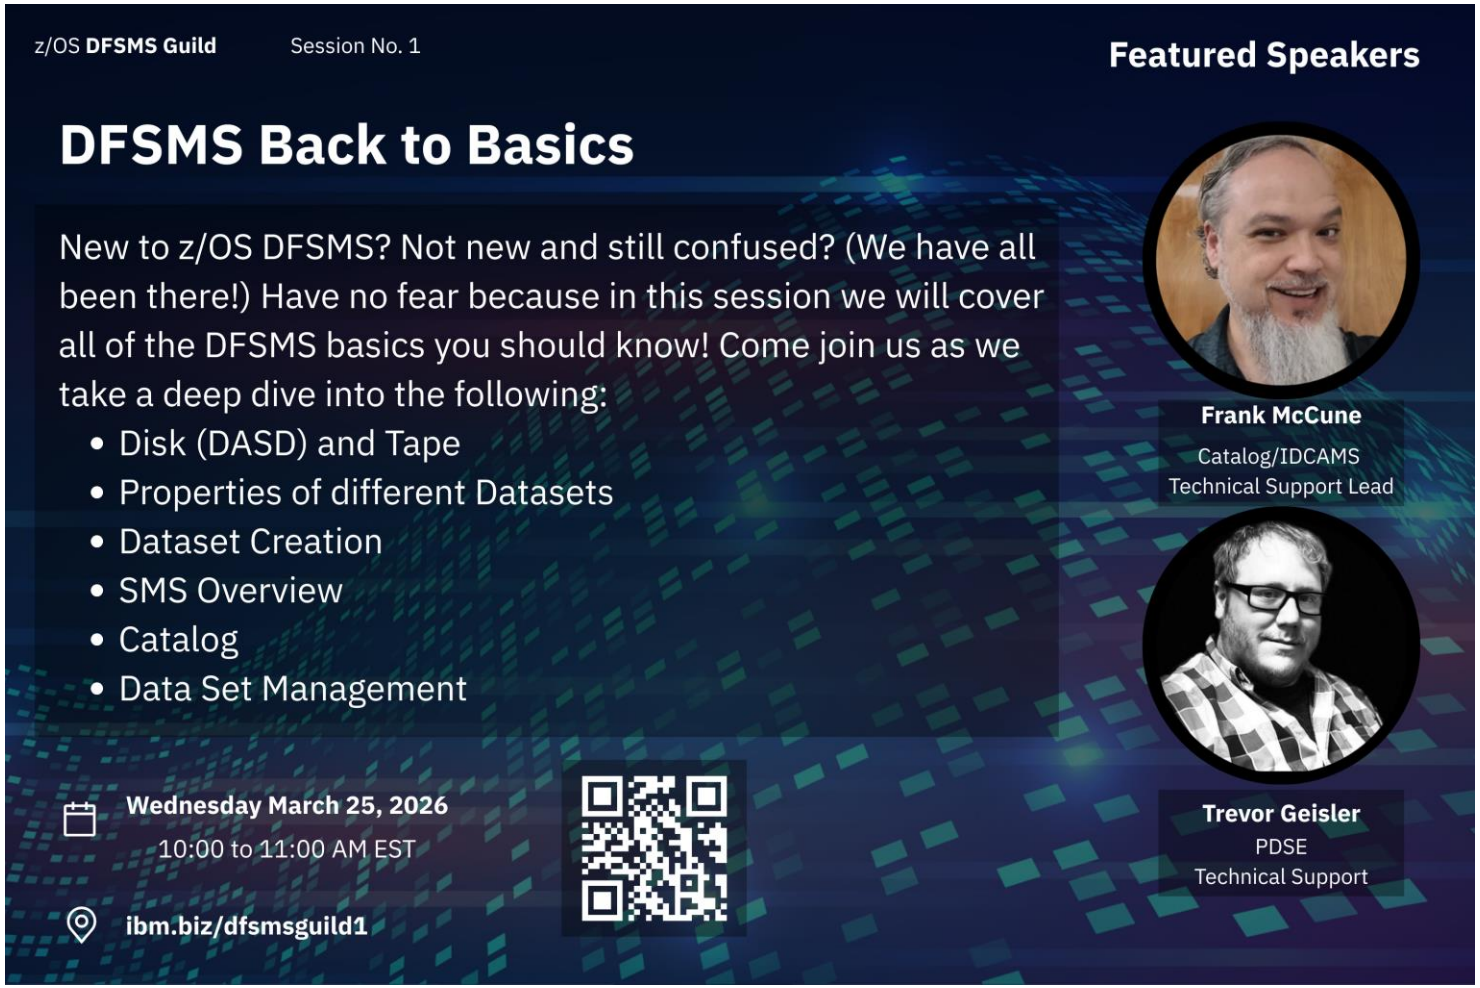

z/OS DFSMS Highlights

Simplify and Modernize

- z/OSMF Storage Management Plugin
- DFSMSrmm (RMM) z/OSMF Plugin Enhancements
- DFSMSrmm customer requirement enhancements
- EzNoSQL COBOL support and Python APIs
- Object Access Method (OAM) REST APIs
- Out-of-space avoidance enhancements

Blue = continuous delivery

Find Out More…

Wednesday
z/OSMF Storage Management Plugin - SMS Like Never Before, 8AM
Thursday
The Ins and Outs of VSAM and More!, 9:15AM **PDSE Best Practices: Getting the Most Out of Your PDSEs,** 10:30AM

z/OS DFSMS Personas

Karla
Experienced Storage Administrator

Remy
Mid-Tenure Storage Administrator

Elsa
Early-Tenure Storage Administrator

Simplify and Modernize – SMS Like Never Before

Karla
Senior Storage Administrator

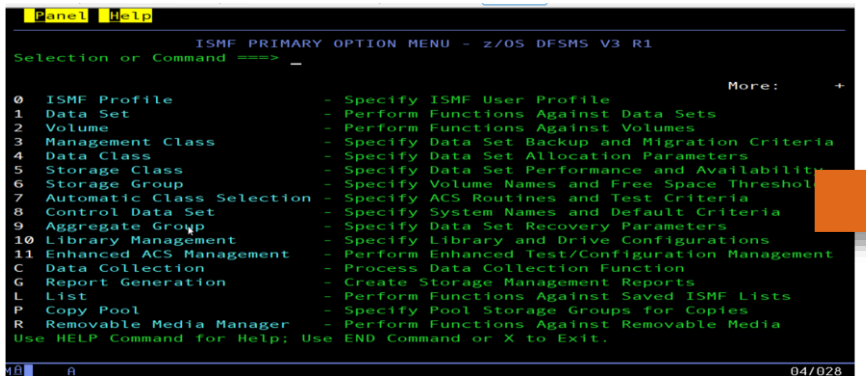

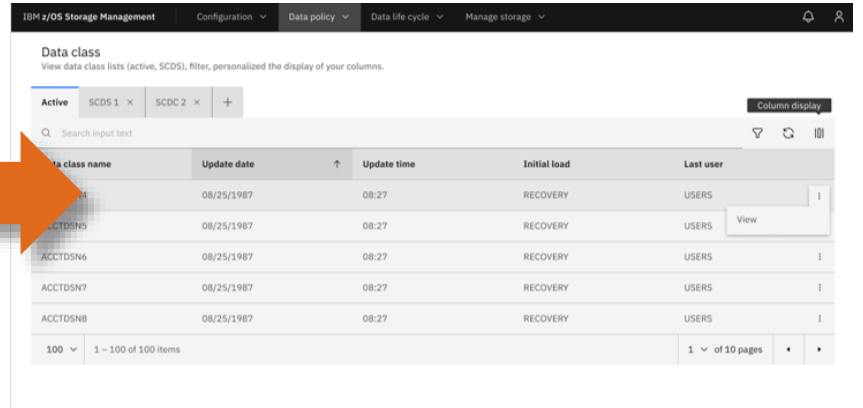

Elsa
Early Tenure Storage Administrator

z/OSMF Storage Management Plug-in

"The usability of the current ISMF solution is hard to use and involves a lot of hidden tricks that early tenure employees don't understand."
- Sponsor User

- ***z/OS 3.2:*** Provide enhanced storage management REST APIs and a new z/OSMF user interface for storage management.
- **View** the SMS configuration via z/OSMF
  - The beginning of a simplification journey and a foundation for automation.
- Storage Management **REST API updates**
  - Ability to specify a Source Control Data Set (SCDS)
  - New API to retrieve Automatic Class Selection (ACS) routine source

? **Why it Matters:** Provides a consistent and intuitive tool for early tenure storage administrators to perform z/OS storage management tasks, reducing manual steps, skills, and time needed for new users to be productive.

Simplify & Modernize

Dashboard & Navigation

- The starting view is the dashboard. Navigation is by top menu or by clicking on items.
- Multiple configurations can be viewed by adding them with the plus icon.

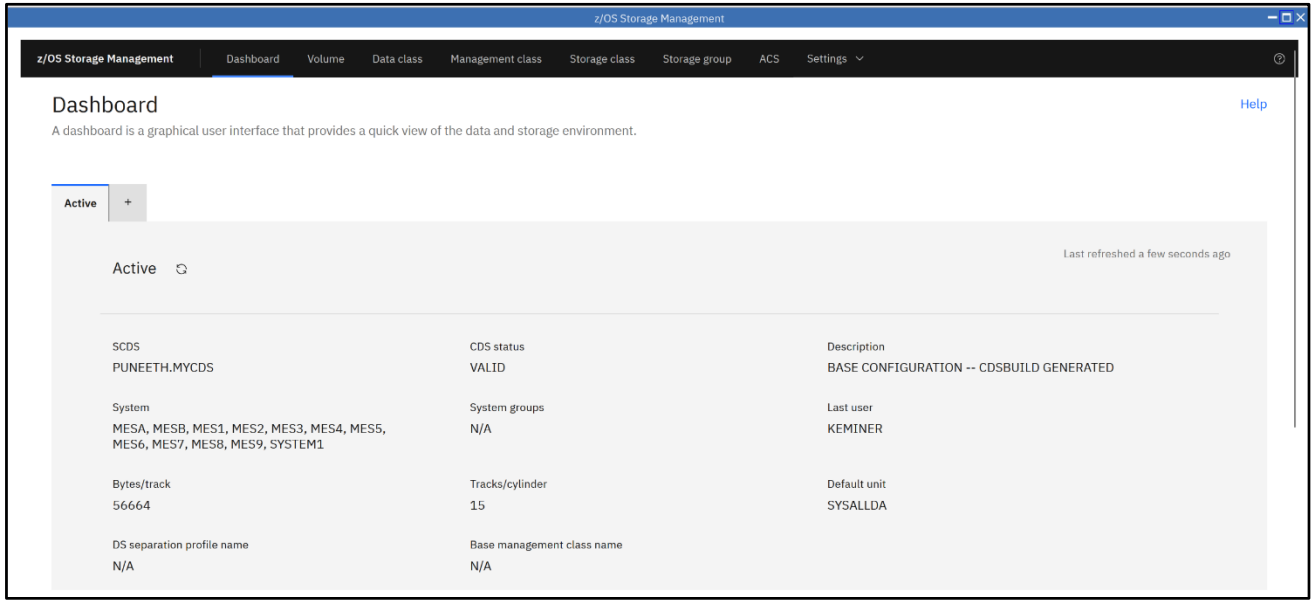

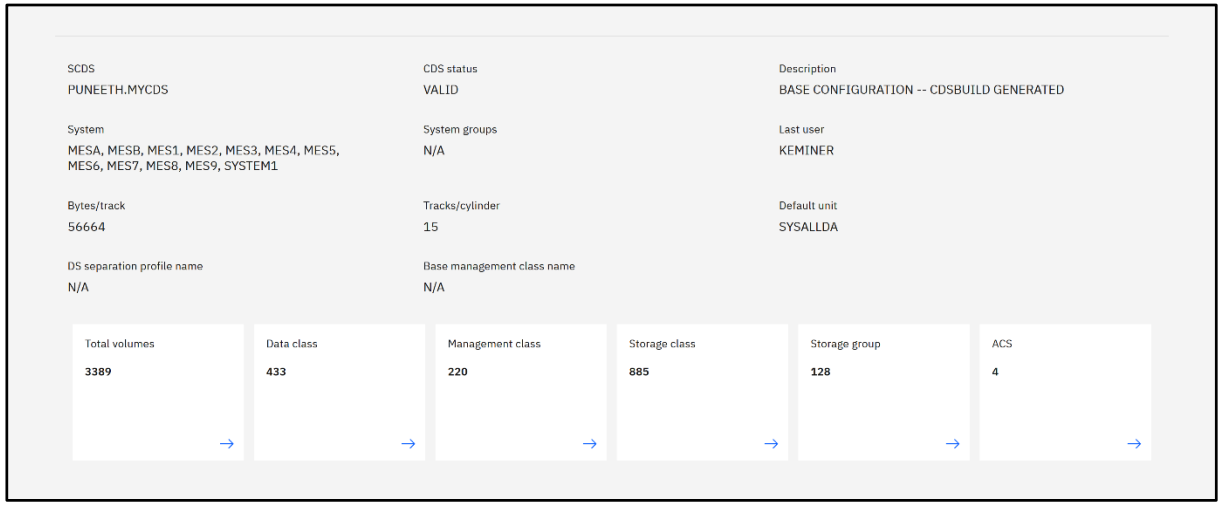

Simplify & Modernize

List Views

- The SMS constructs (Data class, Management class, Storage class, Storage group) in the configuration are displayed as tables.
- Within any table, filter and search is available to select the rows that are displayed.
  - Filter finds text or values in specific columns. Multiple filters are allowed with AND conditions.
  - Free form search finds text in any column.

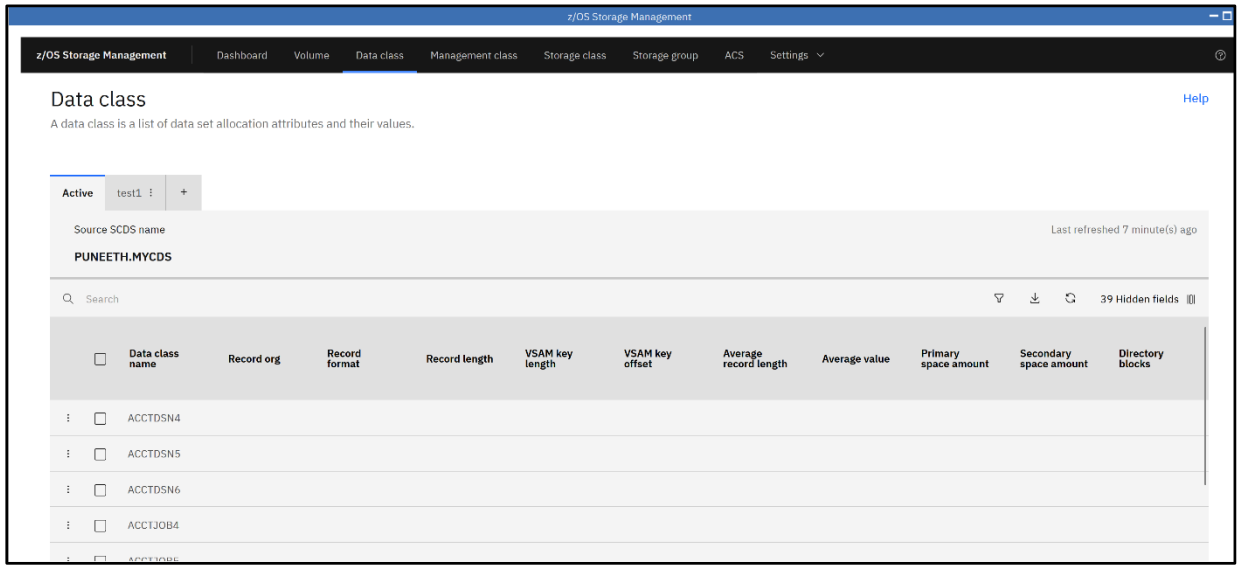

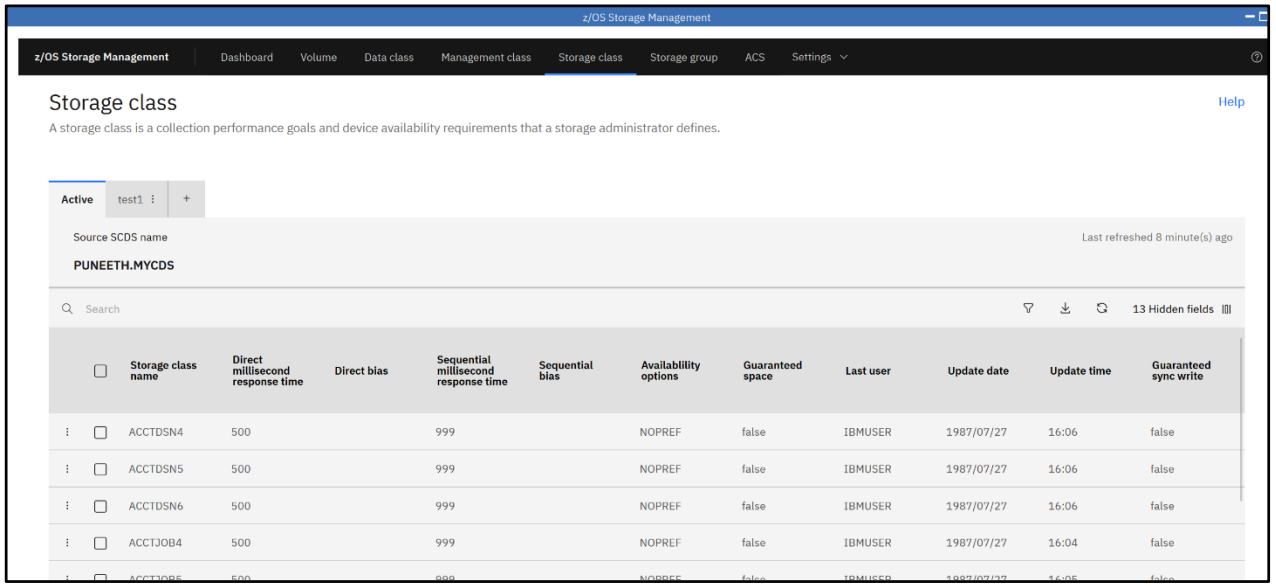

API Explorer (Swagger)

SHARE
EDUCATE • NETWORK • INFLUENCE

Storage Management APIs

Show/Hide | List Operations | Expand Operations

GET	/zosmf/storage/rest/v1/acsroutines	Get ACS Routine info of DC/SC/MC/SG
GET	/zosmf/storage/rest/v1/acsroutines/{con-type}	Get ACS Routine info based of DC/SC/MC/SG
GET	/zosmf/storage/rest/v1/baseconfig	Get base class properties
GET	/zosmf/storage/rest/v1/cds/{scds-name}	This operation retrieves the result from an SCDS activation request
POST	/zosmf/storage/rest/v1/cds/{scds-name}	This operation activate a specified SMS source control data set
PUT	/zosmf/storage/rest/v1/cds/{scds-name}	This operation validates a specified SMS source control data set
GET	/zosmf/storage/rest/v1/dataclasses	Get a list of data classes that match specified storage properties
GET	/zosmf/storage/rest/v1/dataclasses/{dc-name}	Get the data class that matches a data class name
GET	/zosmf/storage/rest/v1/managementclasses	Get a list of management classes that match specified storage properties
GET	/zosmf/storage/rest/v1/managementclasses/{mc-name}	Get the management class that matches a management class name
GET	/zosmf/storage/rest/v1/storageclasses	Get a list of storage classes that match specified storage properties
GET	/zosmf/storage/rest/v1/storageclasses/{sc-name}	Get the storage class that matches a storage class name
GET	/zosmf/storage/rest/v1/storagegroups	Get a list of storage groups that match specified storage properties
GET	/zosmf/storage/rest/v1/storagegroups/{stg-name}	Get the storage group that matches a storage group name
POST	/zosmf/storage/rest/v1/storagegroups/{stg-name}/volumes	Define a volume list to a specified storage group in the SMS source control data set
GET	/zosmf/storage/rest/v1/volumes	Get a list of volumes that match specified storage properties
GET	/zosmf/storage/rest/v1/volumes/{volume-ser}	Get the volume that matches a volume serial number

Recent maintenance: PH66694

- ***Add Volume POST API*** enhanced with validation checks to ensure the generated volume serial numbers adhere to proper naming conventions and to prevent performance-related issues, the API now enforces a limit of 100 volumes per request.

How to get more information

z/OSMF Guild from October 29, 2025!
Click here for the replay and details!

How to get more information

To learn more, visit the **Storage Management** content solution page!!

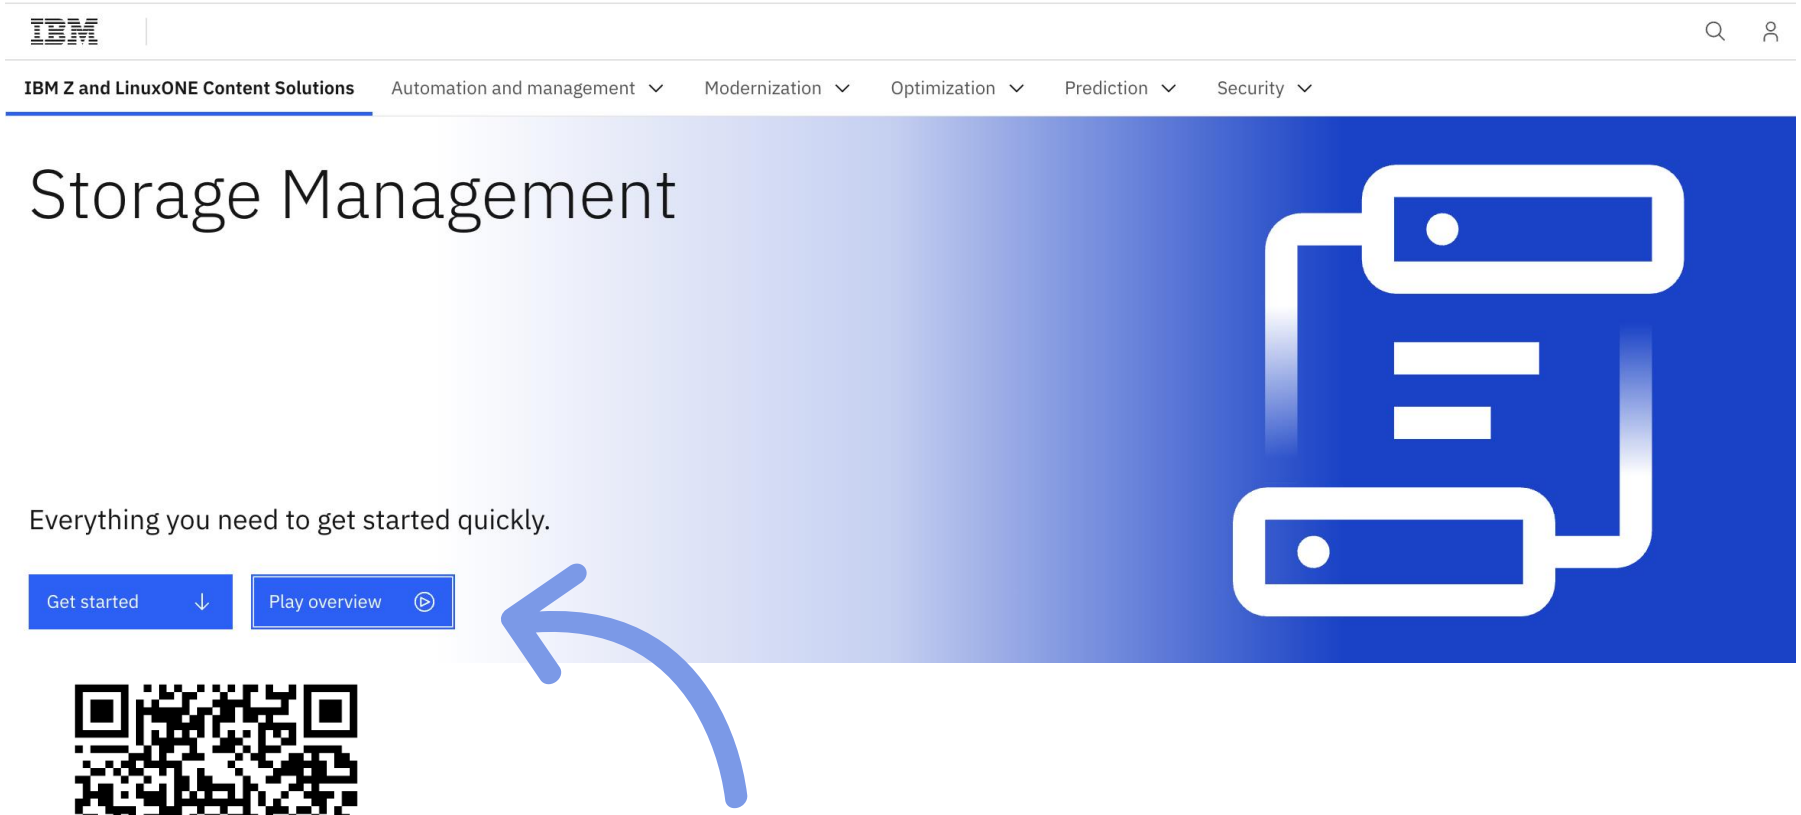

Check out the Overview video for a detailed look at what's available with z/OS 3.2!

Find Out More…

z/OSMF Storage Management Plugin - SMS Like Never Before, Wed 8AM

Simplify and Modernize

z/OSMF DFSMSrmm (RMM) Plug-in

- Support for a modern graphical user interface via a z/OSMF plugin that augments the existing TSO and ISPF dialog support.
- ***z/OS 3.2:*** **Customizable, at-a-glance, Dashboard view.**

? ***Why it Matters:*** Provides immediate visibility into the status of the RMM Address Space, recent housekeeping activity, and volume metrics—making system monitoring more intuitive and actionable.

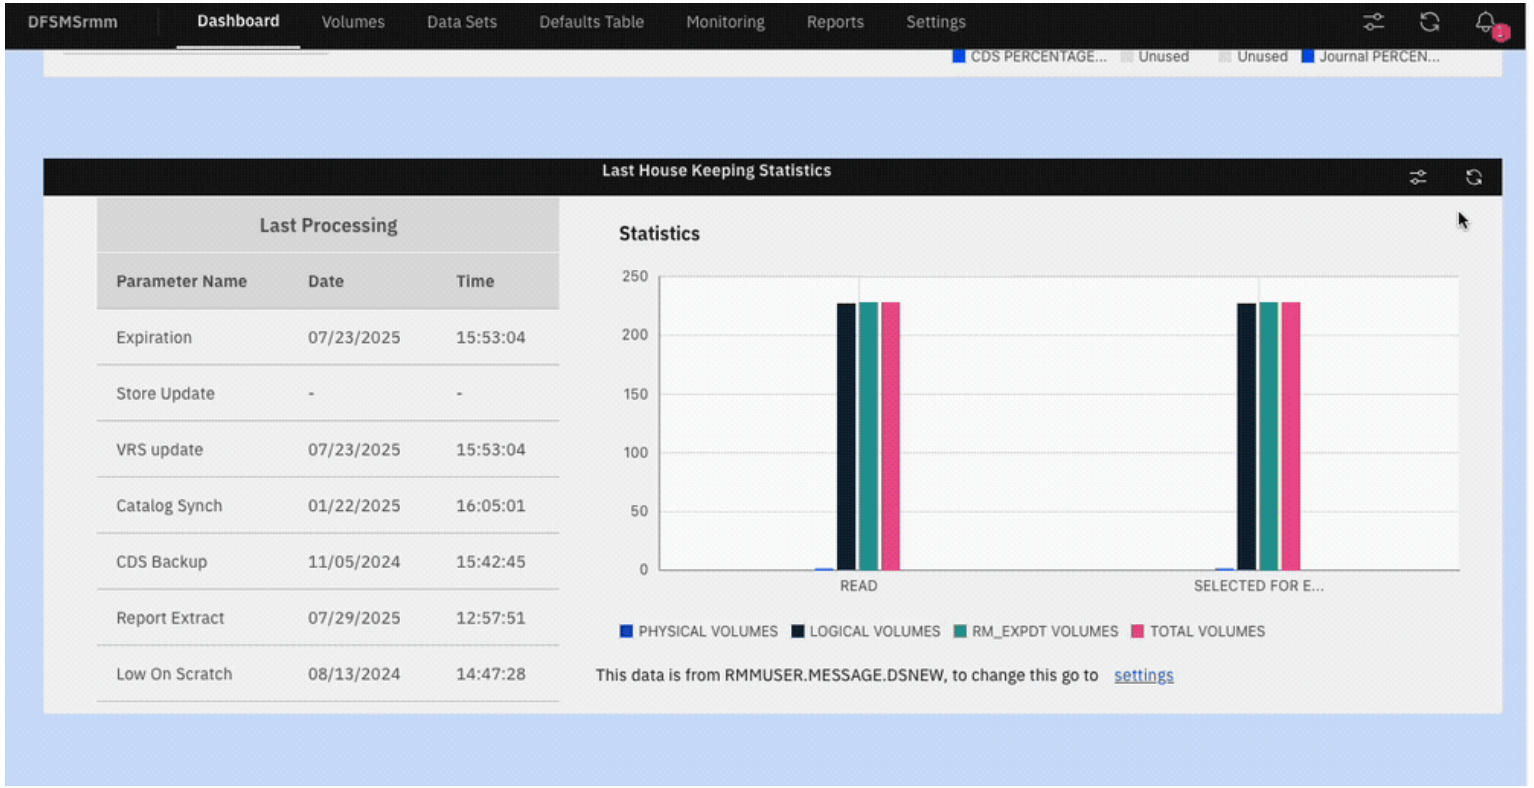

Simplify and Modernize

Customizable, at-a-glance, Dashboard view

- The Dashboard tab is the primary splash screen that displays customizable widgets that relay information on the RMM CDS, Journal, parameters, and status.
  - To use the Dashboard tab, you need READ access to the STGADMIN.EDG.MASTER resource profile.
- Widgets on the screen are customizable and able to be reloaded without affecting the other widgets.
  - Screen Options view to show available widgets:
    - RMM Properties
    - Last Housekeeping Statistics
    - Percentage Retention Policies
    - Percentage Status
  - Each of the described widgets are interactable, allowing the user to gain deeper insights.

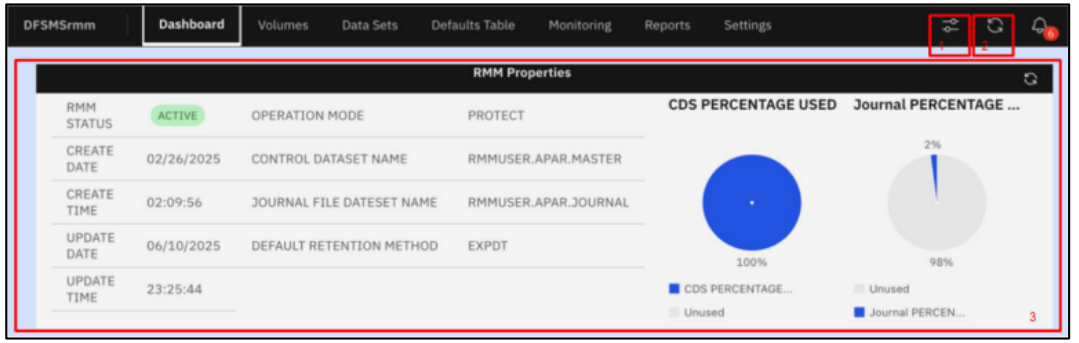

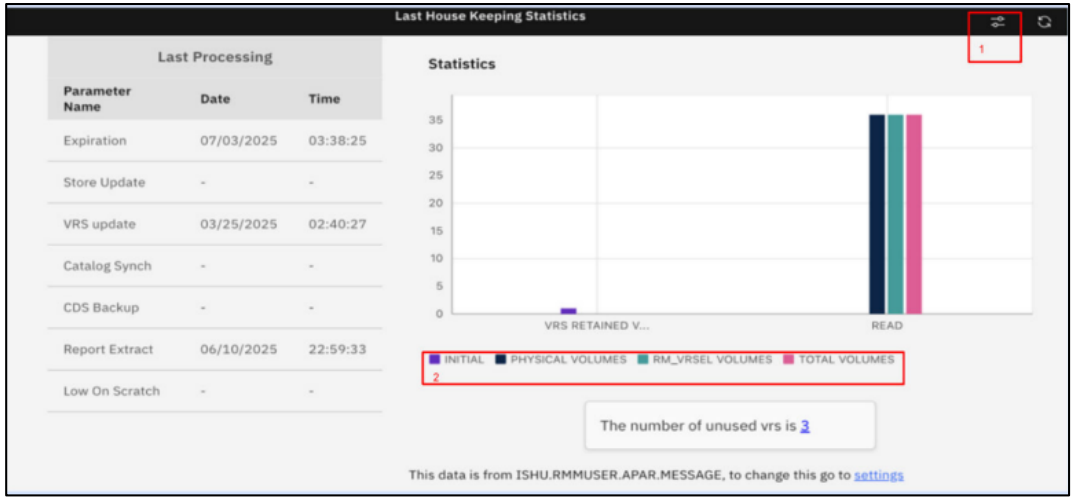

Simplify and Modernize

z/OSMF DFSMSrmm (RMM) Plug-in

- ***z/OS 3.2:*** Ability to generate, view, and/or download RMM extended reports based on shared templates with a flexible interface for sorting and filtering.

? ***Why it Matters:*** Streamlines report creation, allowing users to generate default and custom reports with fewer keystrokes, reducing complexity and saving time.

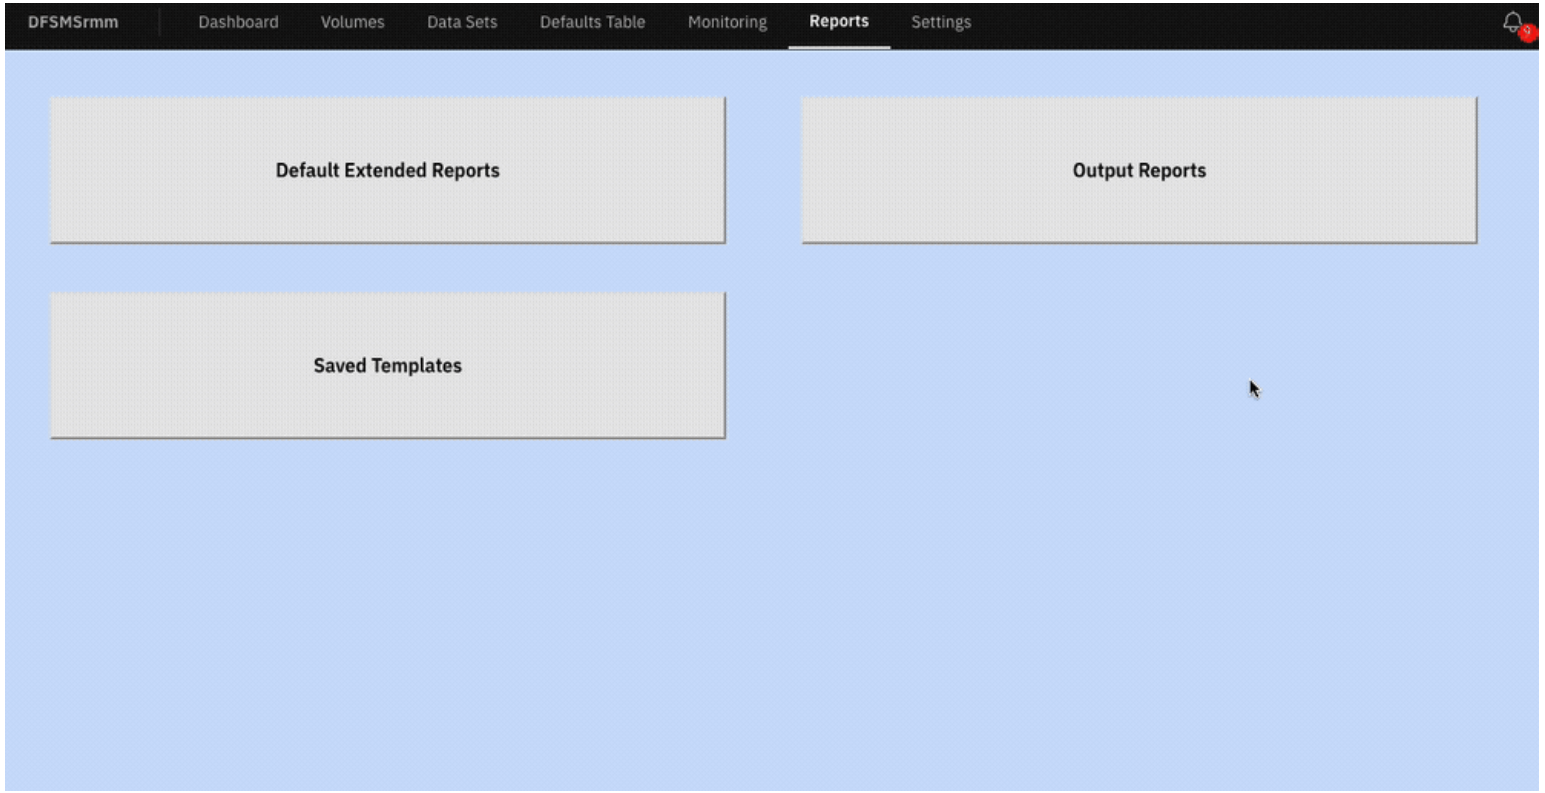

Simplify and Modernize

Streamlined, customizable report generation

- The **Reports Tab** can be used to generate RMM extended reports, view/download generated reports, and create report templates using a standard base report.
  - New section within the plugin, bringing an existing RMM feature ISPF panel to the z/OSMF plugin.
    - Default reports - RMM supplies 15 default reports
    - Output (previously generated) reports
    - Saved templates
  - ★ The “Saved Templates” feature provides one-click report submission without the need to locate JCL manually.
  - ★ Report output can be retrieved by the browser for viewing or to download, eliminating the need to be familiar with ISPF dialogues or complicated Sort JCL.

? ***Why it Matters:*** Streamlines report creation, allowing users to generate default and custom reports with fewer keystrokes, reducing complexity and saving time.

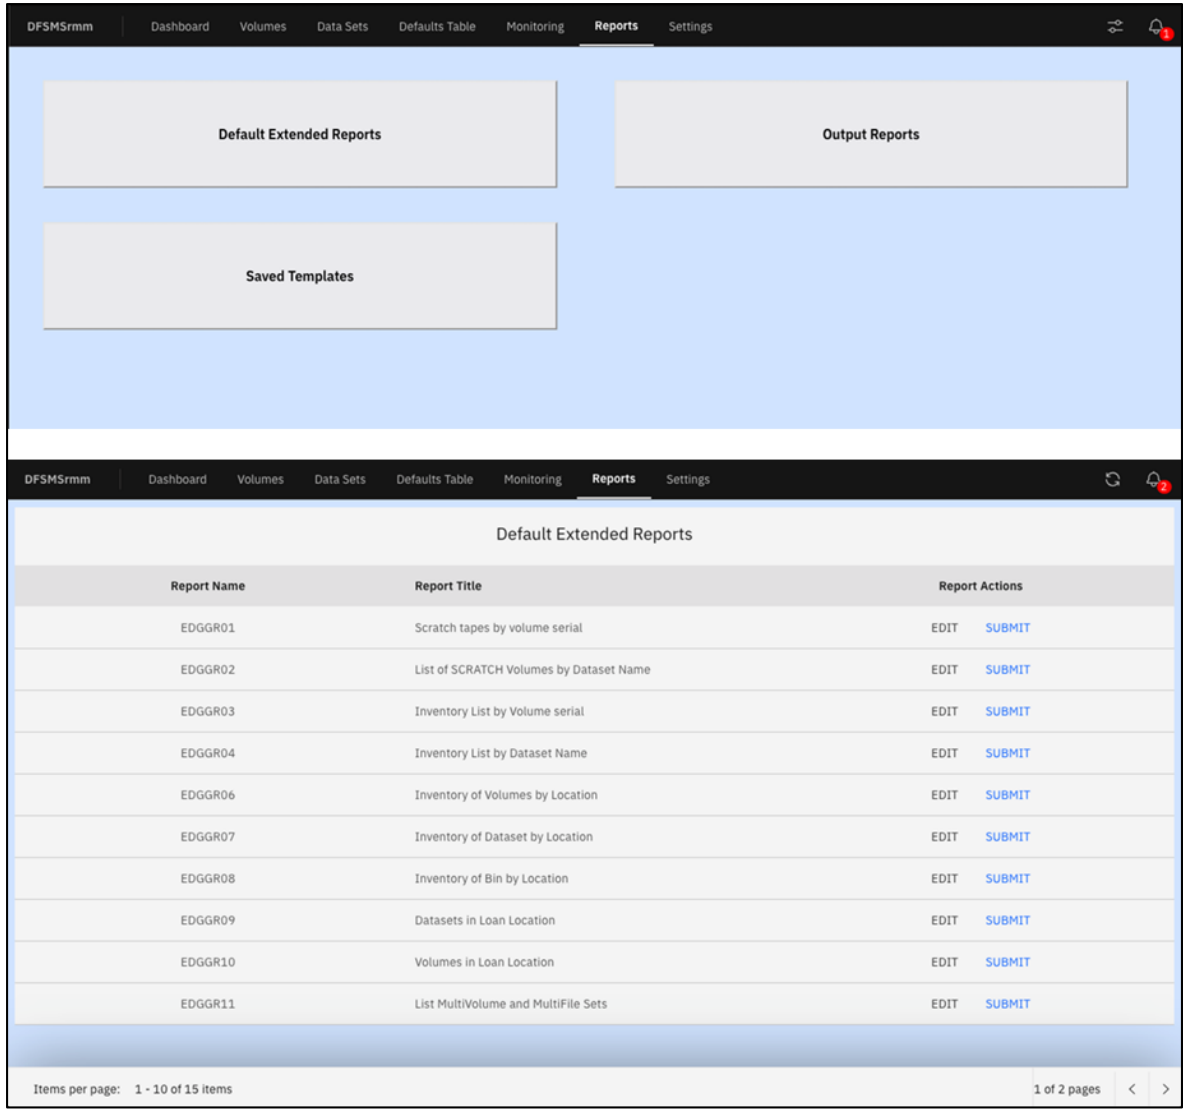

Simplify and Modernize

z/OSMF DFSMSrmm (RMM) Plug-in

- ***z/OS 3.2:*** Improved search capabilities, allowing for scalable searching and providing the ability to do queries and reporting in a more efficient way.
  - Server-side pagination on volume and dataset requests enables users to view all contents of their CDS as single pages at a time and results in performance gains for data queries.
  - Improved usability with Export to CSV.
    - Additional options on what the data export will contain and a better user experience to show work is being done.
  - Exposed more fields from the CDS into the RMM Search capabilities within the RMM plug-in and the TSO interfaces.
    - Fields for searching on dates now include time as well as more dates to search on.
  - ★ For more information see Simplifying your RMM Environment – A Dashing New Look blog post in the DFSMS Community.
- ? **Why it Matters:** Provides performance and usability improvements - data browsing is now faster, request wait times are shorter, and full CDS access is supported, enabling a more complete and responsive user experience.

Simplify and Modernize

DFSMSrmm (RMM) Customer Requirements

- Today, within VRSEL, data sets may be managed entirely by its catalog status by combining WHILECATALOG(YES) and MCATTR(VRSELXDI), but the equivalent doesn't exist for EXPDT.
- ***z/OS 3.2:*** For EXPDT retention method, provide an **'ONLY'** option for **WHILECATALOG**, which will expire an uncatalogued dataset without considering its expiry date.
  - Specified through SMS management class, the Default Table, or RMM PARMLIB member, EDGRMMxx,.
  - See **APAR OA67189** for toleration support on z/OS V2.5 and above.
    - New feature can only be specified on z/OS 3.2, and with toleration support, lower-level systems in a plex will comprehend WC(ONLY) for EXPDT datasets and will expire/not expire as applicable.
  - ***Requirement(s) met:***
    - ZOS-I-3586 EXPDT retention method needs an option for WHILECATALOG control that is independent of expiration date and relies on catalog status only
    - ZOS-I-2233 Z/OS RMM RM(EXPDT) Catalog Control function for expdt=99000

? **Why it Matters:** Simplifies device management retention policies, allowing more tape retention management decisions into SMS constructs and ACS routines and enables usage of the EXPDT retention method as a viable replacement for the existing VRSEL retention method.

Simplify and Modernize

EzNoSQL

- Use NoSQL for z/OS to enable applications to create NoSQL key-value databases, based on open standard JSON file format, ***directly*** on z/OS, which can then be accessed in real-time, at scale, and with transactional consistency.
  - Integrated ***base*** z/OS solution that provides a comprehensive set of APIs designed to simplify application development and leverage the scalability, security, resiliency, and performance provided by z/OS.
- ***1Q25 new enhancement:*** New **EzNoSQL** APIs to utilize Python on z/OS, in addition to C or Java, to build modern, scalable NoSQL database applications.
  - See **APAR OA66418** on z/OS 3.1. PTFs are available!
    - Requires Python 3.10 or greater to be installed on z/OS.
  - ★ For more information see ***NoSQL for z/OS? Even EZer!*** blog post. Content Solution website has everything needed to get started!
- ? **Why it matters:** Python support with EzNoSQL not only expands z/OS database management versatility but also significantly empowers AI and ML.

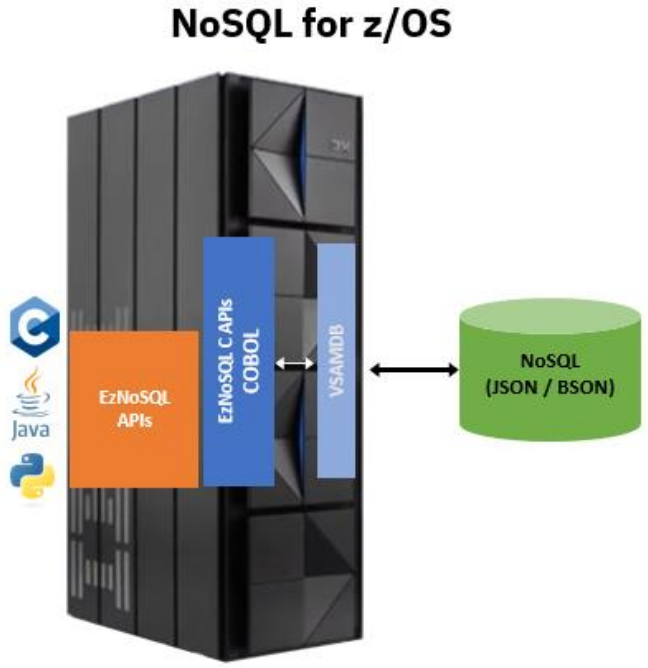

"By making mainframe data available via standard cloud APIs and supporting popular languages like Python for NoSQL database access, IBM is working to lower the barrier to entry for a new generation of developers… This is an essential step towards modernizing how enterprises interact with their most valuable data assets." -- z/OS 3.2: Is Mainframe AI the Enterprise Game Changer?
Steven Dickens, CEO HyperFRAME Research

COBOL Exploitation of EzNoSQL

SHARE
EDUCATE • NETWORK • INFLUENCE

IBM Enterprise COBOL for z/OS 6.5 unleashes the power of IBM z17 and modernizes COBOL applications with user-defined types and VSAMDB support, dated April 8, 2025

- **VSAMDB (NoSQL database) support**
  - Supports reading, writing, updating, and deleting of JSON documents in a VSAMDB file using COBOL's file processing statements. This function is compatible with IBM EzNoSQL for z/OS data.
  - Connect to an EzNoSQL database (VSAMDB) directly from COBOL and allow applications to share a NoSQL database (VSAMDB) between Python, Java, C, assembler, and COBOL programs. These programs can all read, write, update, and delete the same JSON document in a VSAMDB file!
  - ★ Use as a convenient method for converting traditional VSAM records into JSON (or vice versa) and stored into an EzNoSQL database - see sample program IGYVDB1J shipped with the compiler.
  - **Availability date:** 13 June 2025

? **Why it Matters:** The VSAMDB feature helps you modernize your mission-critical COBOL applications and provides application architects with a wider set of options for choosing a programming language they are accustomed to.

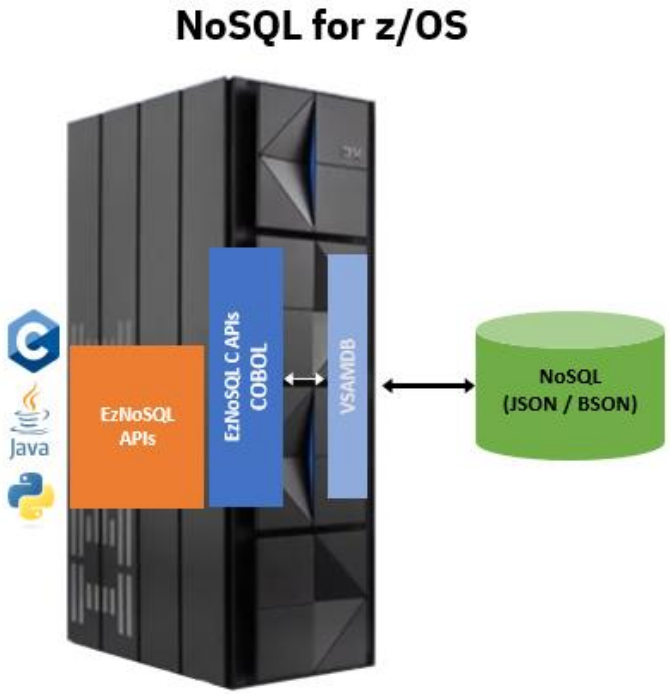

Follow Terri Menendez on LinkedIn to keep up with the latest on EzNoSQL. Check out her latest post on COBOL support for EzNoSQL.

Find Out More…

The Ins and Outs of VSAM and More!, Thur 9:15AM

Simplify and Modernize

OAM REST APIs

- OAM is the access method for storing unstructured (object) data on z/OS, and today OAM's existing interface uses high-level assembler (HLASM).
  - Most voted on OAM requirement, ZOS-I-2086, Enabling of a rest service/https support to use OAM functions
- ***z/OS 3.2:*** OAM provides a REST API for storing, retrieving, deleting, querying, and changing management characteristics of unstructured (object) data.
  - Leverage industry standard interfaces, similar to Amazon S3, allowing REST applications to access existing OAM managed object data on z/OS.
    - DFSMS Cloud Data Access (CDA) used to enter S3 credentials.
  - Designed to easily add new applications and/or transition existing OAM applications to REST.
    - Regardless of how the objects are stored, either through REST or through the native HLASM API, they're fully accessible by either API.
  - New REST front end interfaces with OAM's existing support and data managed by OAM is kept on z/OS.
    - All existing OAM function continues to work as-is
  - See **APAR OA64282** for support on z/OS 3.1. PTFs are available!
  - ★ For more information see ***z/OS S3 Object Store with OAM's Simple Twist -- Take a REST*** blog post in the DFSMS Community.

? **Why it matters:** With DFSMS OAM, users can modernize their access and management of unstructured data on z/OS, while allowing distributed environments simple access to core business data through industry standard REST APIs.

Real Examples!

SHARE
EDUCATE • NETWORK • INFLUENCE

Users > pssobik > dev > demo > demo.asm

```
1  ********************************************************************
2  *
3  * DESCRIPTIVE NAME: Object Storage Request Sample Interface #3 for
4  *                   Multiple OAM Address Spaces
5  *
6  * (Modeled after CBROSR2 V2R2)
7  *
8  * Licensed Materials - Property of IBM
9  * 5650-ZOS
10 * COPYRIGHT IBM CORP. 2017
11 *
12 * FUNCTION: Provides a generalized interface for the Object Storage
13 *           Request (OSREQ) macro in a Multiple OAM environment.
14 *
15 * OPERATION: This routine is called with a parameter area that
16 *            defines the function and pointers necessary to invoke
17 *            the OSREQ macro.
18 *
19 *            If it is determined that an OSREQ function is requested,
20 *            then the OSREQ parameter list is filled in with an
21 *            MF=M form of the macro. The function is executed via an
22 *            MF=E form.
23 *
24 *            1. Validity check the DATAAREA Header.  Exit if error.
25 *            2. Fill in the OSREQ PARM_LIST with all of the optional
26 *               keywords using MF=M form of the macro.
27 *            3. If FUNCTION_REQUEST = "ACCESS  "
28 *               a. IF CBROSR3 was compiled with IADD option, then
29 *                  set IADDRESS OSREQ macro keyword to the address of
30 *                  the DB2 library entry point DSNHLI using the MF=M
31 *                  form of the macro.
32 *               b. ELSE set IADDRESS OSREQ macro keyword to
33 *                  IADDRESS_PTR and set DB2ID OSREQ macro keyword
34 *                  to DB2ID_PTR using  the MF=M form of the macro
```

Users > pssobik > dev > demo > demo.py > ...

```
1  import os
2  import boto3
3  from botocore.exceptions import ClientError
4  import json
5  import pprint
6  import xml.dom.minidom
7
8
9  # Create EZU27 S3 client
10 client = boto3.client('s3',
11                         endpoint_url='http://ezu27.tuc.stglabs.ibm.com:9080/',
12                         aws_access_key_id='s3accesskey',
13                         aws_secret_access_key='s3secretkey',
14                         region_name='production')
15
16
17 # Store one OAM object
18 response = client.put_object(Body=open('demo.txt', 'rb'),
19                              Bucket='group02',
20                              Key='demo.txt')
21 print('')
22 print('Put')
23 print(json.dumps(response, indent=2))
24
25 # Retrieve one OAM object
26 response = client.get_object(Bucket='group02',
27                              Key='demo.txt')
28 print('')
29 print('Get')
30 print(response['Body'].read().decode('utf-8'))
31 print('')
```

Simplify and Modernize

OAM REST APIs

- **Software Dependencies**
  - This support uses existing capabilities in the IBM WebSphere Liberty Application Server.
- **For z/OS 3.1, refer to the following APARs:**
  - OAM **APAR OA64282**
  - Alternate CDA panels for key entry (**OA65068**)
  - DFPMKDIR support (**OA65069**) for code installation (/usr/lpp/dfsms/oam).
  - Program Property Table (PPT) support for CBRWOLA (**OA66996**)
- **Code Levels**
  - Java 11 version 11.0.17.0 or greater
  - z/OS Liberty version 24.0.0.9 or greater
- **Additional Information**
  - Existing OSREQ parameters and API, see OAM Application Programmer's Reference
  - AWS S3 interface, see Amazon Simple Storage Service API Reference

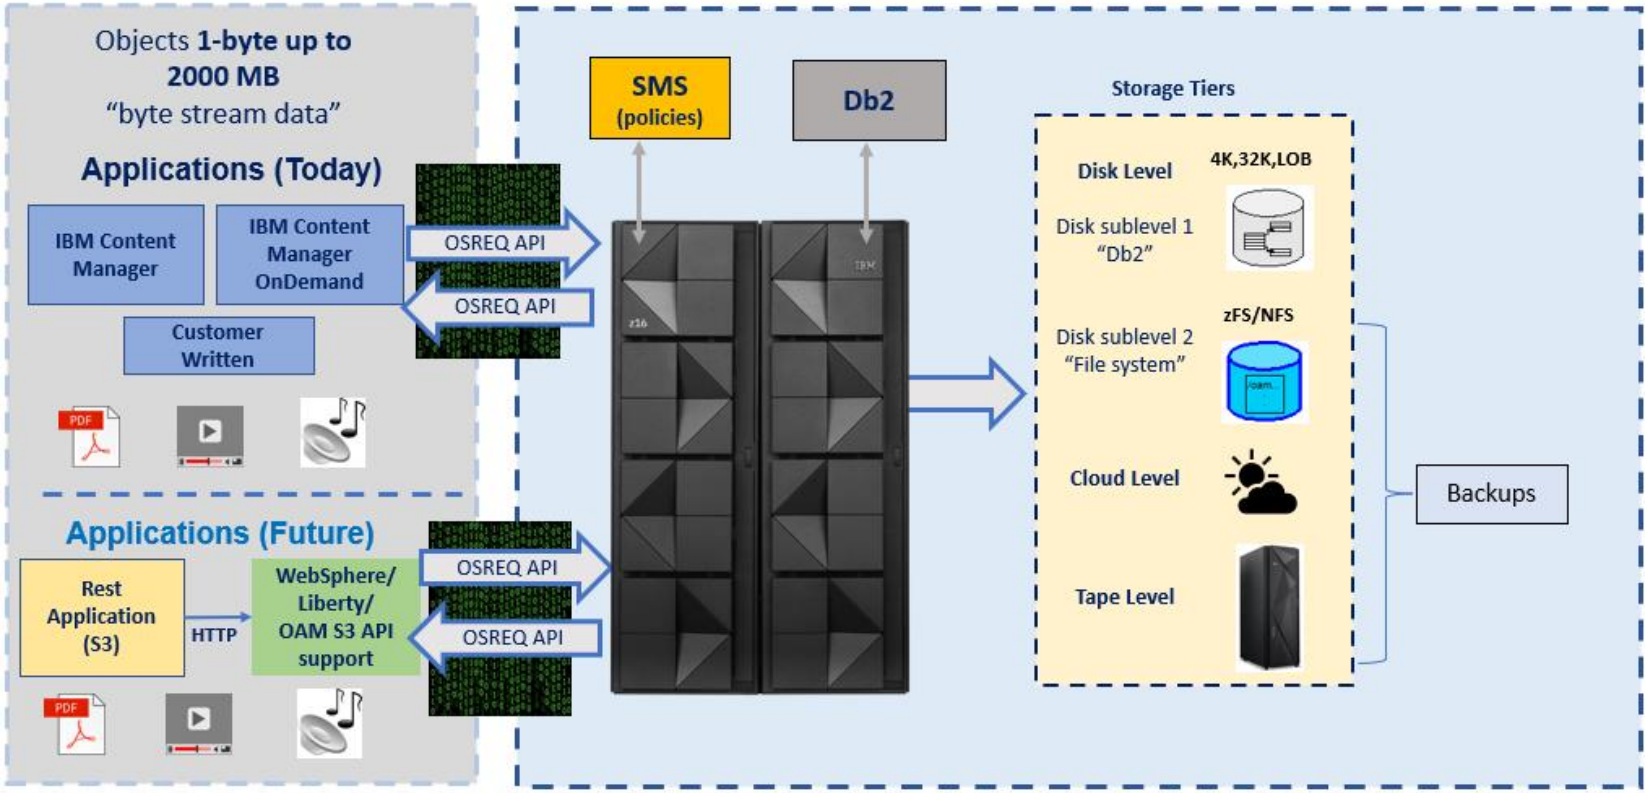

Simplify and Modernize

Out-of-space avoidance enhancements

- Many users may monitor available space on volumes, but not how full the VTOC or INDEX structures are becoming.
  - Possible to have plenty of free space on the volume, but not enough space in the VTOC or INDEX to accommodate the request - could render volumes unusable for new allocations until intervention is done to give the VTOC or INDEX more space.
- ***z/OS 3.2:*** Provide information needed to identify volumes whose VTOCs and INDEXes are running out of space.
  - DADSM LSPACE macro, SMF 19 record, ISMF volume list panel for DASD volumes, Naviquest ACBQVAR1 exec, z/OSMF Storage plug-in volume queries, IDCAMS DCOLLECT output record type VL.
  - ★ **SMS will use this information to improve volume selection and reduce instances where a volume is chosen as it meets the criteria for available space, but the allocation fails as not enough room exists in the VTOC or INDEX.**
    - **IGDSMSxx - AUTO_VOLSTATUS_VTOC(YES|NO)** plus SMS message changes
    - **IGD002I (updated)**
    - **IGD418I (new) VTOC SPACE *capacity* ON VOLUME *volser.* CURRENT VOLUME STATUS IS *status*.**
    - **IGD419I (new) VTOC INDEX SPACE *capacity* ON VOLUME *volser.***

? **Why it Matters:** Externalizing when VTOCs or INDEXes are becoming full helps to proactively avoid out-of-space conditions and enables more intelligent SMS volume selection.

Simplify and Modernize

Out-of-space avoidance enhancements

- Process to automatically create a VSAM Volume Data Set (VVDS) typically happens when the first catalog or SMS-managed data set is defined on the volume. When a VVDS is allocated incorrectly, it requires much manual effort to replace with a new VVDS.
- ***z/OS 3.2:*** Provide the capability to ***create a VVDS automatically and immediately upon initialization of a volume***.
  - New optional parameter for ICKDSF INIT
    - **VVDS(primary,secondary)** - Define and create VVDS dataset on device initialization.
      - Restrictions: Device must be able to vary online, and there must be no devices with the same VOLSER online. PRIMARY and SECONDARY maximum allocation size is 65535.
    - **ONLINE** - Offline device remains online after VVDS creation during initialization. Only valid with the VVDS parameter.
  - See **APAR PH55904** (all ICKDSF changes ship via APARs).
    - The Device Manager address space must be restarted to include this fix: MODIFY DEVMAN,RESTART
  - ***Requirement met:*** ZOS-I-3378 ICKDSF utility to include VVDS creation as part of INIT

? **Why it Matters:** More proactively manage z/OS storage and avoid disruptive outages.

Simplify and Modernize

Out-of-space avoidance enhancements

- Today, a sufficiently full PDSE can fail a delete operation with an “out-of-space” error (e.g. ABEND D37-04).
  - Often occurs when a fix is removed from a SMPPTS (defined as a PDSE).
  - Workarounds often involved creating a new, larger PDSE and copying old data, often requiring a reconfiguration of the SMP/e data set list.
- ***z/OS 3.2:*** PDSEs allow deletion of members from a full data set without encountering out of space conditions.
  - ***Requirement(s) met:*** ZOS-I-1375 Unable to delete member in PDSE 100% utilized (31 votes)

? **Why it Matters:** Avoids manual intervention to recover from such conditions and facilitate PDSE member data availability.

Find Out More…

PDSE Best Practices: Getting the Most Out of Your PDSEs, Thur 10:30AM

z/OS DFSMS Highlights

Data Lifecycle

Cloud Storage

- CDA Compression
- Cloud Compatible Data Sets
- CDA Immutable Objects
- Cloud Data Access
- DSS and HSM Direct to Cloud
- DFSMS Cloud Data Manager (CDM)

Blue = continuous delivery

Find Out More…

Tuesday
Integrating Cloud Objects in Batch Processing, 9:15AM Unlocking a New Era of z/OS Storage with Cloud Object Integration, 2:30PM
Wednesday
Are You Leveraging the Latest DFSMShsm Enhancements?, 1:15PM **Understanding the IBM z/OS Cloud Storage Solutions**, 3:45PM The Brand New TS7700-Taming the TBs to ExBs of IBM zSystems Data, 3:45PM

Cloud Storage

Cloud Storage

z/OS clients are integrating cloud object storage into traditional disk and tape environments to create a hybrid storage architecture that leverages the strengths of on-premise disk and tape storage while adding the intrinsic strengths of cloud solutions for backup, archive, and unstructured data.

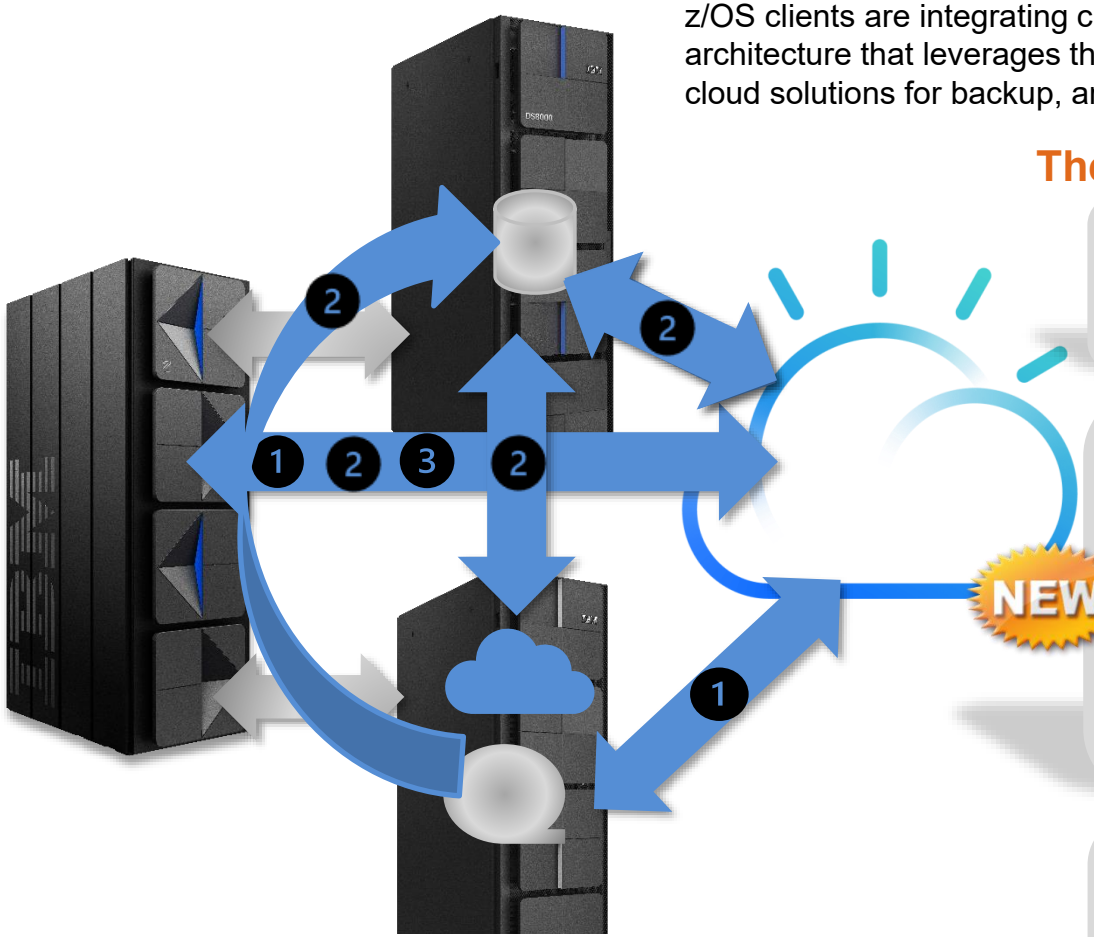

There are existing solutions that leverage Cloud Object Storage …

Solutions transparent to z/OS – no application changes

1. TS7700 Cloud Storage Tier seamless, storage-based host transparent solution
   IBM Cloud Tape Connector for z/OS (CTC) Software-Only Virtual Tape solution

Object aware solutions – no / minimal application changes

2. DS8900 Transparent Cloud Tiering (TCT) and TS7700 DS8000 Object Store enable DFSMS backup and archive with *none* of the data passing through z/OS, *minimizing MIPS*
   DSS / HSM software only direct-to-cloud solution
   OAM provides native support for leveraging object storage directly for z/OS unstructured data (PDFs, audio, video, etc)
   DFSMS Cloud Data Manager (CDM) manages Object Storage data ensuring z/OS governance of z/OS data in the cloud

Store Application Data Directly

3. Cloud Data Access (CDA) provides simple APIs for interacting with the various object storage providers and managing credentials
   GDKUTIL enables batch access to cloud objects (IEBGENER-like utility)
   DFSMS Cloud Data Manager (CDM) manages Object Storage data ensuring z/OS Governance of z/OS data in the cloud

Find Out More…

Understanding the IBM z/OS Cloud Storage Solutions, Wed 3:45PM

Cloud Storage

Cloud Data Access / Cloud Object Utility

- ***1Q25 new enhancement:*** CDA is enhanced to be able to compress data before storing it in cloud object storage and to automatically decompress the data after retrieving it.
  - New optional compression parameter for **GDKWRITE** and **GDKGET** APIs to request either **zEDC or gzip compressio**n
    - **GDKWRITE / GDKGET**: **"compression"** - allowed values are "*zEDC*" and "*gzip*"
    - **GDKGET**: **"decompress"** – allowed values are "true", "false", "attempt-inflate"
  - **GDKUTIL** utility is updated to support a new keywords on the **UPLOAD** and **DOWNLOAD** commands.
    - **UPLOAD**: **COMPRESSION(zEDC|gzip)**
    - **DOWNLOAD**: **DECOMPRESS(zEDC|gzip|zlib|NONE**
  - Available on z/OS 3.1 with **APAR OA66536.** PTFs are available!
  - ★ For more information read the ***Game Changing Series blog post*** in the DFSMS Community.
- ? **Why it Matters:** Compressing data before sending it to the cloud can potentially provide overall performance benefits as less data is transferred.

Cloud Storage

Cloud Data Access (CDA) / Cloud Data Utility

- ***z/OS 3.2:*** **Cloud Compatible Data Sets**
  - ***New enhancement:*** CDA preserves data set record boundaries, allowing data sets to be restored upon download exactly as they were during upload.
    - If the local data set does not exist in the z/OS Catalog, CDA will create the local data set based on the metadata tags associated with the cloud object, captured when the object was uploaded.
  - In addition to already supported sequential data (PS, PDS members, and z/OS USS files), **VSAM KSDS** and **ESDS**, and record format U data sets are now supported.
  - ★ Old DFSMSdss backups on disk or tape can now be moved to object storage and downloaded to other systems for recovery.
    - ***Requirement met:*** ZOS-I-3499 - Download RECFM U.
  - New optional parameter for **GDKWRITE** API:
    - **GDKWRITE** and GDK_PATH_DATALOCATION mode: **"cloudformat"** - allowed value is "*record*"
  - **GDKUTIL** utility is updated to support this enhancement with a new keyword on the **UPLOAD** command:
    - **UPLOAD**: **FORMAT(RECORD)**
  - Available on z/OS 3.1 with **APAR OA66579** and on z/OS 3.2 with **APAR OA67785**. PTFs are available!
  - ★ For more information read the ***Game Changing Series blog post*** in the DFSMS Community.

? **Why it Matters:** Improved compatibility between z/OS datasets and cloud objects, allowing z/OS datasets to be shared between sysplexes, including DFSMSdss backup datasets.

Cloud Storage

Cloud Data Access / Cloud Object Utility

- ***4Q25 new enhancement:*** CDA is enhanced to utilize object store versioning and locking functionality to create ***immutable objects***.
  - Leverage **S3 Object Lock** feature to protect objects from accidental or malicious deletion by preventing modifications or deletions during a defined retention period.
    - **S3 Versioning** is automatically enabled and works together with Object Lock to prevent locked object versions from being permanently deleted or overwritten using a write-once-read-many (WORM) model.
  - New support for both **GDKUTIL** and **CDA APIs** to work with objects in buckets with object lock and/or versioning enabled:
    - Version ID and Entity Tag (Etag) are returned. New **VERSIONID()** keyword.
    - Enhanced **LIST** support to report more information about an object, including Version ID.
    - Ability to delete ALL versions for a specified object.
    - CDA Provider File enhancements to support versioning and locking bucket functionality.
  - Available on z/OS 3.1 with **APAR OA67881.** PTFs are available!
  - ★ For more information read the ***Game Changing Series blog post*** in the DFSMS Community.

? **Why it Matters:** Immutable storage protects data against threats like ransomware, accidental deletion, or internal threats and is indispensable for secure backups, regulatory compliance, and building true cyber resilience.

Cloud Data Access (CDA) Resources

Content Solution Page

Find Out More…

Integrating Cloud Objects in Batch Processing,
Tue 9:15AM

CDA Blog Series:

- Introduction to CDA
- CDA Enhancements: New and Updated Keywords (OA65224)
- CDA Enhancements: New and Updated Keywords Deep Dive
- CDA Microsoft Azure Support (OA65925)
- CDA Signed Payload Support (OA65989)
- CDA Enhancements: Long Object Name, Text Conversions, Credential Management (OA65990)
- CDA API Enhancements (OA65990)
- DFSMSdfp CDA Compression (OA66536)
- DFSMSdfp CDA Cloud Compatible Data Sets
- Cloud Data Access and Object Immutability
- Getting SMF data to Object Storage

Related CDA Blogs:

- Sharing z/OS data in the cloud using base z/OS utility GDKUTIL
- Using GDKUTIL to copy backup tape data to an S3 compatible cloud storage
- Files & Sequential Data Sets & Objects, oh my!

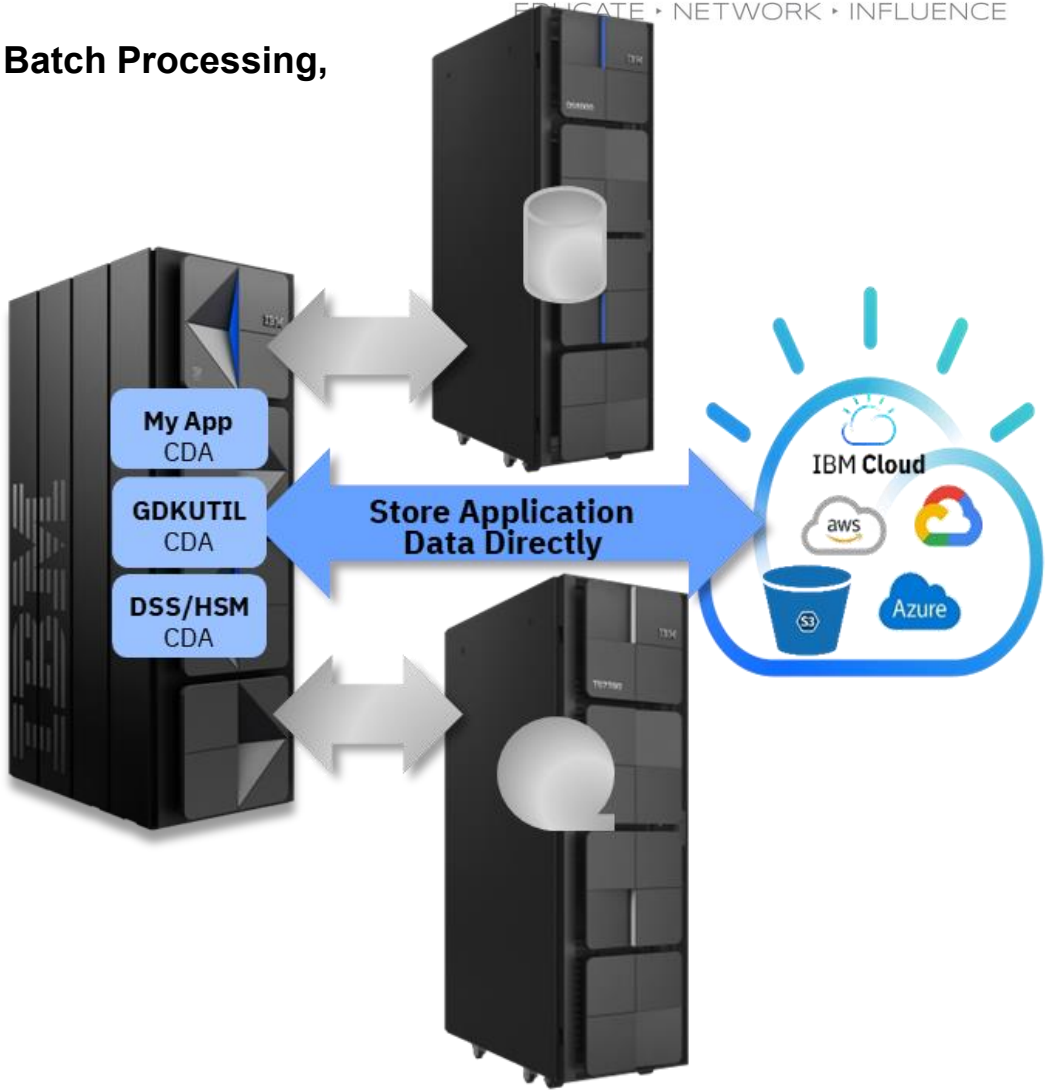

Cloud Storage

DFSMSdss (DSS) Direct to Cloud

- ***1Q25 new enhancement:*** DSS **DUMP** and **RESTORE** commands are enhanced to use CDA APIs to store and retrieve objects directly from cloud object storage.
  - DSS will utilize the new CDA compression support, and zEDC can be specified for these new cloud object storage backups.
    - Physical full volume and logical data set are both supported.
  - Requires a CDA cloud provider file with appropriate configuration details for each cloud intended to be used for DSS.
    - New keyword, **CDAPROVIDERFILE**, allows DSS to first check for a CDA provider file before defaulting to an SMS network connection.
    - Sample provider files are shipped in '/usr/lpp/dfsms/dss/samples/'
  - Available on z/OS 3.1 with **APAR OA66450 and OA67394.** PTFs are available.
  - ★ For more information see ***In the clouds with DFSMSdss*** blog post in the DFSMS Community. More to come!!

? **Why it Matters:** This new, software only direct-to-cloud, vendor agnostic solution provides another option to incorporate cloud object storage into z/OS environments, providing users with more flexibility, more choices, and more ways to simplify their storage strategy.

Cloud Storage

DFSMSdss (DSS) Direct to Cloud

- ***4Q25 new enhancement:*** DSS Direct to Cloud solution provides support for Cloud Object Versioning.
  - DSS will utilize the new CDA immutable storage support to store multiple versions of the same backup. Repeated backups don't overwrite existing data; rather, they create new versions.
    - Specify new DSS keyword **VERSION** during DUMP.
    - Save the unique **DUMPVERSIONID** generated by DSS to specify during RESTORE.
    - Use **CLOUDUTILS LIST VERSIONS** to display all current versions in the cloud, along with their DUMPID and last modified timestamp.
  - Requires CDA cloud provider file updates to specify versioning.
  - Available on z/OS 3.1 with **APAR OA67510 and OA67781.** PTFs are available.
  - ★ For more information see ***In the clouds with DFSMSdss – Part 2*** blog post in the DFSMS Community.

? **Why it Matters:** Safeguard backups from accidental or malicious changes and avoid using system variables to create unique dump names for manual versioning. No need to manually manage deleting older versions if the object storage provider supports maintaining a specific number of versions.

Cloud Storage

DFSMShsm (HSM) Direct to Cloud - HSM CDA DIRECT

- ***2Q25 new enhancement:*** HSM **MIGRATE, RECALL, DUMP,** and **RESTORE** functions are enhanced to leverage CDA to store and retrieve objects directly from cloud object storage.
  - **Primary supported functions:** MIGRATE, RECALL, Full Volume DUMP, RECOVER, FRDUMP, FRRECOVER
    - **Other supported functions:** LIST, AUDIT MEDIACONTROLS, Secondary Space Management (SSM), HDELETE, and DDELETE
  - HSM leverages all CDA supported cloud providers except for Azure and new CDA compression support to compress data in-flight during MIGRATE and Full Volume DUMP operations.
  - Requires CDA setup via key file, config file and provider file with minor modifications for use with HSM.
    - Configure CDA for the user ID the DFSMShsm started task is running under.
    - Sample provider files in /usr/lpp/dfsms/gdk/samples/providers directory may be used.
  - Available on z/OS 3.1 with **APAR OA66204 (toleration) and OA66198.** PTFs are available. Be aware of prerequisites! Refer to APAR documentation for OA66198
    - To activate the new function, use new parameters to the existing SETSYS CLOUD command: **SETSYS CLOUD(NAME(name) CDAPROVIDER)**
  - ★ For more information see ***Going Above and Beyond the Clouds with HSM*** blog post in the DFSMS Community.

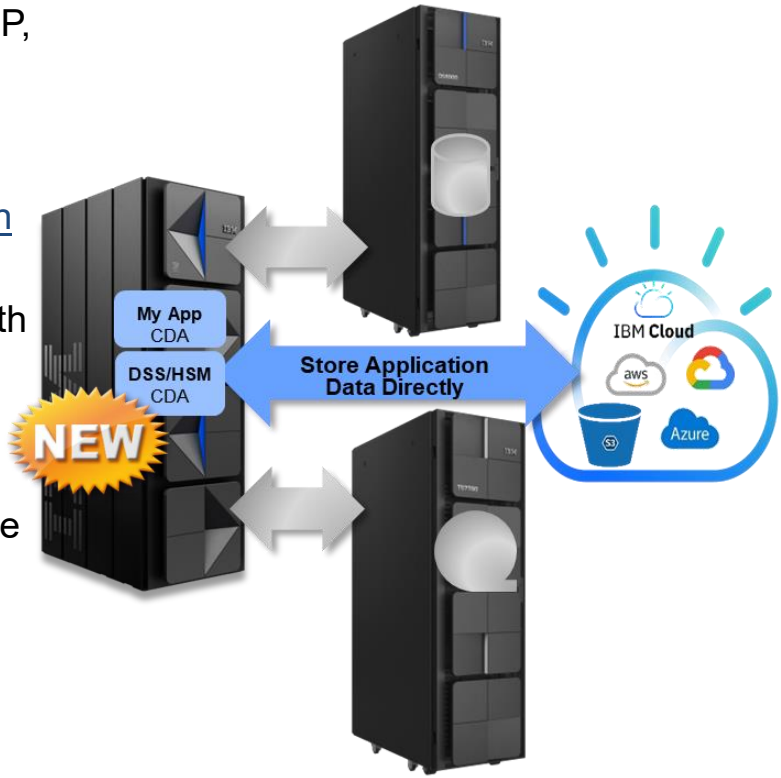

? **Why it Matters:** This new, software only direct-to-cloud, vendor agnostic solution provides another option to incorporate cloud object storage into z/OS environments, providing users with more flexibility, more choices, and more ways to simplify their storage strategy.

Find Out More…

Are You Leveraging the Latest DFSMShsm Enhancements?, Wed 1:15PM

Cloud Storage

DFSMS Cloud Data Manager (CDM)

- Separately priced, stand-alone z/OS software to provide automated, ***policy-based*** and command-driven object storage data management at scale.
- **CDM V1.1:** Enables HSM to move large amounts of candidate data from ML2 to MLC, *while retaining important migration and expiration dates*.
  - Utilize DS8K Transparent Cloud Tiering (TCT), TS7700 Object Store features, and/or HSM CDA DIRECT to reduce MSU consumption.
- ***New*** **CDM V1.2**: Seamlessly move data between disk/tape and cloud object storage at ***scale***
  - Managed, parallel movement of all the data sets specified by a single partially-qualified name
    - Alternative to FTP, NAS, etc
    - Share data sets across sysplexes & enterprises
    - Immutability

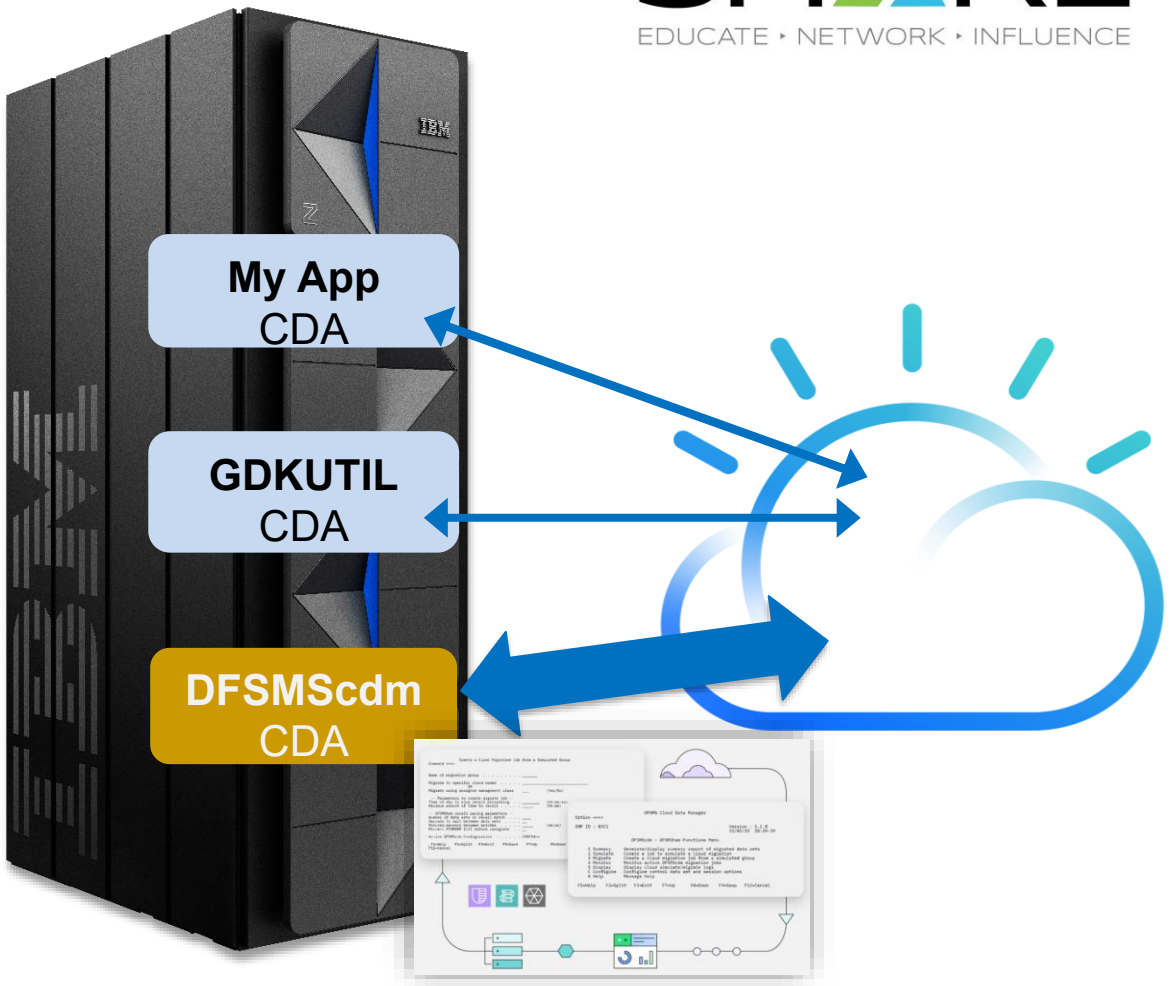

DFSMS Cloud Data Manger (CDM)

Find Out More…

Unlocking a New Era of z/OS Storage with Cloud Object Integration, Tues 2:30PM

z/OS DFSMS Highlights

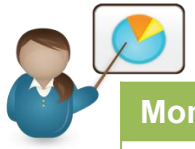

Find Out More…

Monday
Elevate Your Data Resiliency Framework, 10:30AM
Tuesday
Up Your Resiliency Game With Feedback from WSC, 8AM Stronger, Smarter Operations: How a Large Italian Bank Reinvented Data Resiliency Management with IZBR, 10:30AM Unlocking the Power of z/OS HyperSwap for High Availability of Storage, 2:30PM **Unlocking Resilience: How a Wall Street Bank Created a Secure and Automated 4-Site Replication Solution,** 3:45PM
Wednesday
ICF Catalog Resiliency - Understanding How to Protect and Recover Your Catalogs, 9:15AM **From Corruption to Restoration: The Power of HSM's Enhanced CDS Recovery,** 10:15AM
Thursday
Improve Your RTO and RPO With a Multi-Site Resiliency Solutions, 9:15AM World-Class High Availability and Disaster Recovery With IBM GDPS 4.8, 1:15PM

Security

- Dataset Encryption
- Granular control for data set encryption enablement
- Optimized encrypted data set copy
- VSAM GSR support for encrypted data sets

Blue = continuous delivery

Security

- **Dataset Encryption**
  - Various factors are driving the need for clients to adopt extensive use of encryption across their enterprises, including compliance mandates and the threat of data breaches.
  - ***Ongoing enhancements:*** Provide users the ability to encrypt their data sets, using either SAF or SMS policies, ***without changing their application programs***
    - Data set level granularity
    - Supports separation of access control for data set and encryption key label
    - Enabled through RACF and/or SMS policy
      - Security policy: RACF data set profile DFP segment
      - Explicitly: JCL, Dynamic Allocation, TSO Allocate, IDCAMS DEFINE
      - SMS policy: Data class
    - Provides mechanisms for audit readiness

Designed to take advantage of the processing power of the z15 and beyond

Security

Dataset Encryption

- ***SOD Preview:*** IBM intends to enhance pervasive encryption to perform encryption within the access methods for ***tape data sets***. It is expected to be ***transparent to the application program*** unless it uses EXCP.
  - This support extends to both SMS managed tape libraries as well as non-SMS managed tape, cataloged and non-cataloged data sets, and physical and virtual cartridge tape are supported.
  - A tape volume can have a mix of encrypted and non-encrypted data sets.
  - This support requires IBM standard label tapes.
  - Full function support will be available on z/OS V2.5 and above.
- **Support will look very similar to basic and large format DASD encryption:**
  - Ability for the security administrator to indicate to the system that tape data sets are eligible for data set encryption.
    - New resource in the FACILITY class – ***STGADMIN.SMS.ALLOW.DATASET.TAPE.ENCRYPT***
  - Ability to supply a key label through the similar sources as used for DASD encrypted data sets.
    - RACF DFP segment, JCL statement, or data class
    - To open an encrypted tape data set, the user must have SAF authority to both the data set and the key label.

Security

Dataset Encryption

- **Support will look very similar to basic and large format DASD encryption**
  - Applications using standard BSAM and QSAM APIs require no, or minimal changes*, to access encrypted tape data sets.
  - Applications using EXCP require changes to access encrypted tape data sets by performing encryption/decryption of data.
    - The **IGGENC** access method encryption macro can be used to simplify encryption/decryption of data while ensuring compatibility with the access methods.
  - The encrypted attribute cannot be added to an existing data set, but IDCAMS REPRO, IEBGENER, ICEGENER or a similar copying program can be used to copy the data set into a new encrypted data set.
  - ***Requirement met:*** **ZOS-I-2192** RACF Granular Control for encryption of basic and large format datasets
    - New keyword on the RACF DATASET profile DFP segment to indicate that a key label should be honored for the tape data set.
      - **ENCRYPTTYPES(INTAPE | EXTAPE | NOTAPE)**

    ★ **NOTE:** In addition to the Tape options, the ENCRYPTTYPES keyword will also support options for PDSEs and basic and large format datasets. See **APAR OA68029** on z/OS V2.5 and above.

? **Why it matters:** Allows applications to encrypt tape data sets securely using **BSAM and QSAM APIs with *no/minimal changes and using EXCP with provided interfaces.***

Security

DFSMS Optimization Mode for Data Set Copy

- *z/OS 3.2:* Enhancements provide the ability to move data within the access methods without decrypting/uncompressing.
  - New interface for **VSAM encrypted data sets and sequential encrypted extended format data sets** to optimize copying from a source to target encrypted data set, with similar characteristics, where both have the same key label.
  - Makes the copy operation more secure since the data remains encrypted throughout the copy.
  - Also provides the capability to optimize copying **sequential extended format compressed data sets** with similar characteristics using BSAM so that the data does not have to be decompressed and recompressed.
  - **IDCAMS REPRO** provides the initial exploitation of this new capability.
  - See **APAR OA63434** and **OA68062** on z/OS V2.5 and above.

? **Why it matters:** Intended performance benefit to not having to decrypt and/or decompress the data on input and then re-encrypting/recompress the data on output.

Security

VSAM GSR support for encrypted data sets

- ***z/OS 3.2:*** Enhancements to VSAM now provide the ability to use GSR buffering with an encrypted data set.
  - Prior to this enhancement, VSAM GSR buffering could not be used with an encrypted data set. VSAM OPEN failed with message ***IEC161I 121-163.***
  - ***Requirement met:*** ZOS-I-2166 VSAM GSR support for encryption KEYLABEL
  - See **APAR OA65855** on z/OS V2.5 and above. PTFs are available.

? **Why it matters:** Expand the use case for VSAM data set encryption to include GSR buffering.

z/OS DFSMS and Storage Highlights "Cheat Sheet"

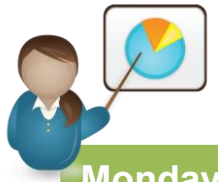

Find Out More…

Monday

- Elevate Your Data Resiliency Framework, 10:30AM
- **What's New in DFSMS**, 2:30PM

Tuesday

- ☺**Up Your Resiliency Game With Feedback from WSC,** 8AM
- What's New On the IBM DS8000 G10 Storage Family, 8AM
- **Integrating Cloud Objects in Batch Processing**, 9:15AM
- Stronger, Smarter Operations: How a Large Italian Bank Reinvented Data Resiliency Management with IZBR, 10:30AM
- Unlocking a New Era of z/OS Storage with Cloud Object Integration, 2:30PM
- Unlocking the Power of z/OS HyperSwap for High Availability of Storage, 2:30PM
- ☺**Unlocking Resilience: How a Wall Street Bank Created a Secure and Automated 4-Site Replication Solution,** 3:45PM

Wednesday

- **z/OSMF Storage Management Plugin - SMS Like Never Before**, 8AM
- **ICF Catalog Resiliency - Understanding How to Protect and Recover Your Catalogs**, 9:15AM
- **From Corruption to Restoration: The Power of HSM's Enhanced CDS Recovery,** 10:15AM

Wednesday (con't)

- **Are You Leveraging the Latest DFSMShsm Enhancements?**, 1:15PM
- IBM Z Cyber Vault Explained - Soup to Nuts and Nose to Tail, 2:30PM
- The Brand New TS7700-Taming the TBs to ExBs of IBM zSystems Data, 3:45PM
- **Understanding the IBM z/OS Cloud Storage Solutions**, 3:45PM

Thursday

- Improve Your RTO and RPO With a Multi-Site Resiliency Solutions, 9:15AM
- **The Ins and Outs of VSAM and More!**, 9:15AM
- **PDSE Best Practices: Getting the Most Out of Your PDSEs,** 10:30AM
- World-Class High Availability and Disaster Recovery With IBM GDPS 4.8, 1:15PM

Blue = Storage HW sessions

Do you want to influence the future of DFSMS?

- Want to be a Sponsor User and have a direct impact on the future of DFSMS?
- Talk to us right here at the event, or
- Scan the QR Code below to fill out our **DFSMS sponsor user interest form**:

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.

Any Questions?

SHARE
EDUCATE • NETWORK • INFLUENCE

Information Overload

Ask Questions!!

Storage Management Free-For-All, Wed 5PM

Want to attend an in-person

IBM z/OS Academy?

Learn, Interact and Network with IBMers and peers

May 5th- 7th, 2026	**Fall 2026**
IBM Tech Campus	**IBM US**
Ehningen, Germany	**New York, USA**

These **free** events are designed for early tenure z/OS system programmers (2-10 years), but all are welcome!

Training and presentations include topics on new z/OS capabilities, best practices, career tips, and **much more!**
Subscribe to the community page today to stay informed about future events!

Register now for Ehningen/ Germany:

Join our IBM Community: https://ibm.biz/zOSAcademy
Questions? Contact us at zOS.Academy.USA@us.ibm.com **or** zOS.Academy.Europe@de.ibm.com

Your feedback is important!

SHARE
EDUCATE • NETWORK • INFLUENCE

Submit a session evaluation for each session you attend:

www.share.org/evaluation

Thank You for Attending!

BACKUP

IBM Z Configuration Collection Program: Sign up form

SHARE
EDUCATE • NETWORK • INFLUENCE

Overview

Sign up to participate in the ***IBM Z Configuration Collection Program*** — the first step in IBM Z's initiative to simplify configuration across the Z stack. Our goal is to reduce complexity and optimize performance and value.

Due to the challenges with current configuration settings, we are partnering with clients to take a data-driven approach. By analyzing existing configurations, we aim to identify patterns that will inform decisions on what should be deprecated, internalized, or optimized to default settings.

Process

01	02	03	04
Sign up to receive an invitation to the Early Release Program site. (*click to agree to access*)	***Download and run*** the provided scripts on each system.	***Review and scrub*** the output data to remove any information you do not wish to share with IBM.	***Upload*** the files to IBM.

For questions, contact zos@ibm.com with the email subject line: Z Configuration Collection Program

IBM

Announcements

- **IBM Storage DS8000 10th Generation,** announced September 10, 2024
  - *The next-generation of IBM Storage DS8000, is the latest innovation in enterprise-class storage for IBM Z mainframe architectures. Designed to harness the full power of IBM Z, this 10th generation of IBM Storage DS8000 is designed to provide organizations access to critical workloads, consistent and optimized data performance, and a modular architecture to adopt the latest IBM research-backed technologies to fuel business growth while monetizing data.*
  - Accelerating Innovation in Mainframe Storage for Next-Generation Workloads with the new IBM DS8000
  - The new DS8000 announced at IBM Storage Strategy Days 2024

78

Announcements

- **Look for more z/OS continuous delivery enhancements every quarter on the z/OS Community page.** To stay up to date, please 'Join Community' to ensure you don't miss an announcement!
- **IBM z/OS 3.1 4Q24 enhancements,** dated November 19, 2024
  - EzNoSQL Python API support
  - TLS 1.3 Support for Transparent Cloud Tiering (TCT)
  - Cloud Data Access (CDA) enhancements

Blue = continuous delivery
Hyperlink = backup

79

DADSM LSPACE Macro

New fields added to the EXPDATA area:

```
LSPDDCET 68(X'44')   4 DCE size of the volume in tracks
LSPDVTSZ 72(X'48')   4 Size of the VTOC in tracks
LSPDIXSZ 76(X'4C')   4 Size of the INDEX in tracks
LSPDPCUV 80(X'50')   1 Percentage of used VTOC records
LSPDPCUI 81(X'51')   1 Percentage of used INDEX records
```

LSPACE SMF TYPE=19 record- additions in yellow

62	3E	SMF19FL1	1	binary	SMF flags 1. **Bit** **Meaning when set** 0 SMF19CYM. Returned data is for a volume that has CYL-managed space. 1 SMF19DNT SMF19DCT not equal to SMF19TRK 2 SMF19IXD Volume has a disabled INDEX. 3 SMF19NIX Volume has no INDEX.
80	50	SMF19PUV	1	binary	Percentage of used VTOC records
84	54	SMF19PUI	1	binary	Percentage of used INDEX records
124	7C	SMF19VLI	12	EBCDIC	Volume size information

LSPACE SMF TYPE=19 record- additions in yellow

132	84	SMF19DCT	4	binary	DCE size of the VTOC in tracks
VTOC/INDEX Information					
136	88	SMF19VXI	8	EBCDIC	VTOC/INDEX information
136	88	SMF19VTS	4	binary	Size of the VTOC in tracks
140	8C	SMF19ISX	4	binary	Size of the VTOC INDEX in tracks

ISMF – new columns added

2.1 (VOLUME SELECTION ENTRY PANEL) and View 9 10 11 44 45

```
                                                                    VOLUME LIST
Command ===> _

Enter Line Operators below:

     LINE        VOLUME  INDEX      FREE      FREE      VTOC %   INDEX %
     OPERATOR    SERIAL  STATUS     DSCBS     VIRS      USED     USED
    ---(1)----   -(2)--  --(9)---   -(10)--   -(11)--   -(44)-   -(45)--
                 EV9LIA  ENABLED       4763       518        21        18
                 EV9LIB  ENABLED       4650       532        23        16
                 EV9LIC  ENABLED       5416       535        10        15
                 EV9LID  ENABLED       5738       551         4        13
                 EV9LIE  ENABLED       5835       552         3        12
    ----------   ------ ------------------------------------------------- BOTTOM  OF  DATA
```

```
                                                                    VOLUME LIST
Command ===> _

Enter Line Operators below:

     LINE        VOLUME  INDEX      FREE      FREE      VTOC %   INDEX %
     OPERATOR    SERIAL  STATUS     DSCBS     VIRS      USED     USED
    ---(1)----   -(2)--  --(9)---   -(10)--   -(11)--   -(44)-   -(45)--
                 1P0401  ENABLED       1430       311         1         1
                 1P0403  DISABLED     65499   -------         1         0
                 1P0404  ENABLED       1196       585        18         7
    ----------   ------ ------------------------------------------------- BOTTOM  OF  DATA
```

Naviquest Support and Example

Two new parameters can be used on the SYSIN DDNAME statement when running the ACBQVAR1 EXEC in batch:

- **VTOC%USED**
- **INDEX%USED**

JCL invoking ACBQVAR1

```
//*******************************************************************
//GENREP  EXEC ACBJBAOB,
//       PLIB1=SYS1.DGTPLIB,
//       TABL2=IBMUSER.TEST.ISPTABL
//ISPFILE DD  DSN=IBMUSER.DASDVOL.REPORT,DISP=OLD
//SYSTSIN  DD *
PROFILE PREFIX(IBMUSER)
ISPSTART CMD(ACBQVAR1 DASDLST SORT(VOLSER)) +
NEWAPPL(DGT) BATSCRW(132) BATSCRD(27) BREDIMAX(3) BDISPMAX(99999999)
/*
//SYSIN DD  *
VOLSER
INDXSTAT
DEVICETYPE
FRAG
USEATTR
LGEXTRK
VTOC%USED
INDEX%USED
INIASRES
OWNERID

/*
```

z/OSMF Volume query example

Request to get a list of volumes:
GET zosmf/storage/rest/v1/volumes?filter=DOMUSB
Example output:

```
[
              {
                "storageGroupName": "DOMUS",
                "lastUser": "KNMORT",
                "updateDate": "2003/08/07",
                "totalCapacityForTMS": 0,
                "volumeSerial": "DOMUSB",
                "updateLevel": 0,
                "freeSpace": 0,
                "storageGroupStatus": "ENABLED",
                "largestFreeExtentForTMS": 0,
                "updateTime": "11:22",
                "freeSpaceForTMS": 0,
                "fullVolumeLastUsed": 0,
                "trackRegionLastUsed": 0,
                "totalCapacity": 0,
                "largestFreeExtent": 0,
                "usedVTOCRecordsPercent": 10,
                "usedVTOCIndexRecordsPercent": 15,
                "status": [
```

Example output (continued):

```
                    {
                      "confirmedSmsStatus":
"DISABLED/ALL",
                      "sysType": "1",
                      "requestedSystemStatus":
"NONE",
                      "sysName": "S1I",
                      "mvsSystemStatus": "OFFLINE"
                    },
                    ...
                    {
                      "confirmedSmsStatus":
"NONE",
                      "sysType": "2",
                      "requestedSystemStatus":
"NONE",
                      "sysName": "XCFLOCAL",
                      "mvsSystemStatus": "OFFLINE"
                    }
                  ]
                }
              ]
```

DCOLLECT – new fields in blue in the record type 'VL'

- VOLUME RECORD FLAG INFORMATION
- 84(X'54') CHARACTER 1 DVLFLAGS FLAGS AND RESERVED
- 1... .... DVLCONV 1 = VOL IS IN CONVERSION
- .1.. .... * RESERVED
- ..1. .... DVLVTCIX 1 = VTOX IX, 0 = OSVTOC
- ...1 .... DVLIXDIS 1 = VTOC Index Disabled
- .... 1... DVLVTCOV 1 = VTOC Alert Override, when
- Warning ON, VLDSTSMS modified
- .... .111 * RESERVED

- 438(X'1B6') UNSIGNED 1 DVLPVTOC % USED VTOC RECORDS
- 439(X'1B7') UNSIGNED 1 DVLPIDX % USED VTOC INDEX RECORDS

SMS IGD002I message

Update existing IGD002I variants legend section with two (2) new indicator and text description for “/” and “=”.

```
VOLUME UNIT MVS SYSTEM= 1 2 3 4 5 6 7 8 STORGRP NAME
1P0201 0E55 ONRW / . . . . . . . S1P02
****************************** LEGEND ******************************
. THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
+ THE STORAGE GROUP OR VOLUME IS ENABLED
- THE STORAGE GROUP OR VOLUME IS DISABLED
* THE STORAGE GROUP OR VOLUME IS QUIESCED
D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
> THE VOLSER IN UCB IS DIFFERENT FROM THE VOLSER IN CONFIGURATION
/ THE VOLUME IS QUIESCED FOR NEW ALLOCATIONS BY THE SYSTEM
= THE VOLUME IS DISABLED FOR NEW ALLOCATIONS BY THE SYSTEM
SYSTEM 1 = SYSTEM1 SYSTEM 2 = SYSTEM2 SYSTEM 3 = SYSTEM3
SYSTEM 4 = SYSTEM4 SYSTEM 5 = SYSTEM5
```

z/OS DFSMS Highlights

Simplify and Modernize

- EzNoSQL

Blue = continuous delivery

Simplify and Modernize

"Best of distributed matched with the best of Z"

- New **NoSQL key:value document store** provides the ability to store open standard BSON/JSON (UTF-8) objects directly on z/OS
  - **CRUD** (Create, Read, Update, and Delete) properties

  PLUS

  - **ACID** (Atomicity, Consistency, Isolation, Durability) transaction properties
    - ***Transaction consistency is the biggest obstacle to adoption of NoSQL***

Modernize solutions and workloads directly on z/OS

- Use the latest tools/APIs (ie NoSQL) to more quickly develop applications with dynamic, flexible data schemas
- Native APIs leverage the flexible data schema, allowing for faster application development

Application data is platform independent

- Use the same development tools and practices across the entire organization (Z and other platforms)
- Write application programs using a modern API, and don't worry about where the data is located

Leverage Z qualities of service - scalability, security, resiliency, performance

- 24x7 SYSPLEX data sharing through VSAM RLS
- Built-in data set level encryption and integration with z/OS pervasive encryption
- Integrated disaster recovery
- Data lifecycle management with HSM and DSS
- Dynamic, high-performance indexing to alternate keys for faster queries/analytics
- Provide transactional support through Transactional VSAM (TVS)

Less work for the application programmer and less work for systems administration

z/OS DFSMS Highlights

Security

- Dataset Encryption

Blue = continuous delivery

Why not use sequential extended format for encryption?

BEST PRACTICE: Use encrypted sequential extended format data sets

- SMS-managed DASD data sets
- Benefits
  - Supports 123 extents per volume
  - Can be multi-striped for performance
  - Can be compressed format for space savings
    - Unless opened for UPDATE
  - Can be encrypted for data protection
- Consideration
  - Each physical block is fixed length (short blocks are padded to create full blocks) and has a 32-byte suffix.
- Allocated with DSNTYPE keyword
  - JCL DSNTYPE=EXTREQ or EXTPREF
  - SMS Data class DSNTYPE=EXTR or EXTP

- Restrictions
  - System data sets (such as Catalogs, SHCDS, HSM data sets) should not be created as extended format, unless otherwise specified.
  - Cannot be opened for EXCP processing
  - Cannot be temporary data sets
  - Cannot be VIO data sets
  - Not supported by Checkpoint/Restart
  - Not supported by hiperbatch

Identify data sets that are eligible. Use zBNA to evaluate candidates for extended format.

Security

Convert to Non-encryption on First Open

- Currently OPEN issues an abend for an encrypted data set when it encounters one of the following data set encryption restrictions: BLKSIZE or LRECL minimums not met, or hardware keys not supported (DCBKEYLE not zero).
- ***2Q24 CD:*** Enhancements provide the ability to allow an empty encrypted sequential basic format or large format data set to be converted to a non-encrypted data set when open processing detects certain data set encryption requirements have not been met, instead of failing the OPEN request.
  - A new FACILITY class resource, **STGADMIN.SMS.DATASET.ENCRYPT.PREFERRED**, with at least **READ** authority is required to enable this support.
  - See **APAR OA66122** on z/OS V2.5 and above. PTFs are available!!
    - Fix Categories: IBM.Function.DataSetEncryption
    - Keyword: DSENCRYPT/K

? **Why it matters:** Potentially avoid costly application outages due to OPEN abends, especially with large data set encryption conversion efforts.

z/OS Data Set Encryption – Evaluate impact

Estimating CPU Cost of Data Protection

Version 2.6.1
Available
12/17/2024

IBM Z Batch Network Analyzer (zBNA) Tool

- A no charge, “as is” tool to analyze batch windows
- Available to Clients, Business Partners, and IBMers
- Windows PC based, and provides graphical and text reports
- Available on Techdocs:
  - IBM Z Batch Network Analyzer

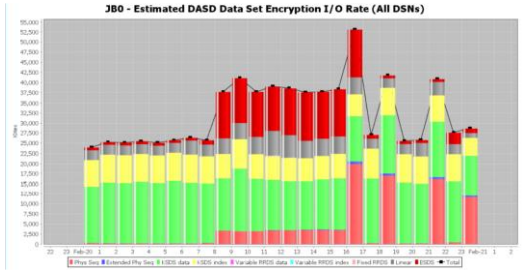

Enhanced to help clients estimate encryption CPU overhead based on actual client workload SMF data

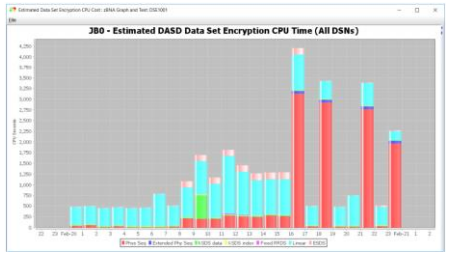

- Top Candidate List for data set Encryption selectable by Service Class and Job Name filters
- Ability to
  - filter the data by time
  - estimate the CPU cost of encrypt / decrypt activities for candidate data sets
  - estimate zEDC Compression and then encrypt / decrypt candidate data sets

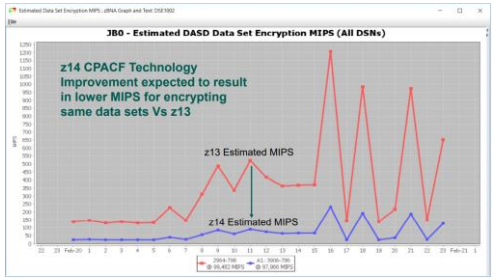

z/OS DFSMS Highlights

Data Lifecycle

Cloud Storage

- Cloud Data Utility / Cloud Data Access
  - Long Object Name
  - Text Conversion Overrides
  - Credential Management
  - Skip Key File and Provider File Disk Read
  - &GDK_SECONDARY symbol

Blue = continuous delivery

Cloud Storage

Cloud Data Access / Cloud Object Utility

- ***4Q24 new enhancement:*** Multiple enhancements intended to simplify management and use of cloud objects from z/OS.
  - **Long Object Name**
    - Requirement (ZOS-I-4119) to allow continuation characters or some way to specify an object name longer than 72-80 characters long.
    - New **GDKUTIL** DD name called **OBJNAMEX**
      - Read each line of the DD to build the full object name, with leading and trailing whitespace from each line discarded, and the remaining characters are put together to create the long object name.
    - ***Makes specifying object names easier***
  - **Text Conversion Overrides**
    - Requirement (ZOS-I-4123) to provide more options for text conversions than the original IBM-1047 EBCDIC to ISO8859-1 ASCII.
    - **GDKUTIL** allows specification of the codeset name for the local text and the codeset name for the remote text.
      - New sub-parameters: **CONVERT(<local_codeset>,<remote_codeset)**
    - New optional parameters for **GDKGET/GDKWRITE**: "**localCodePage**" and "**remoteCodePage**"
    - Default (ie IBM-1047 and ISO8859-1) can be overridden by key/value pairs in the provider file.
    - ***Allows for more options when converting text.***
  - Available on z/OS V2.5 and 3.1 with **APAR OA65990.** PTFs are available!

Cloud Storage

Cloud Data Access / Cloud Object Utility

- ***4Q24 new enhancement:*** Multiple enhancements intended to simplify management and use of cloud objects from z/OS.
  - **Credential Management**
    - Requirement (ZOS-I-4228) with regards to credential management and the need from many security departments to change cloud credentials every 30 days. Can be overwhelming!!
    - Provide **GDKUTIL** users the ability to perform credential management via batch JCL.
      - **CREDENTIALS** command with **ADD**, **DELETE**, and **LIST**.
      - Keyword **ZUSERID(<userid>)** is recognized for the **CREDENTIALS** command, allowing the credentials of that z/OS UserID to be managed.
      - **CREDENTIALS** command will perform the credentials operation for the named **PROVIDER** (and optional **BUCKET/OBJNAME**).
      - Expect DevOps pipelines can be used to submit batch JCL that will add new credentials via this method.
    - New **JSON APIs** will take an input JSON string and break it into the equivalent parameters to pass to the original equivalent CDA credential management API: **gdkkeyad(), gdkkeydl(), gdkkeysr(), and gdkkeygr()**. Results are returned as an output JSON string.
    - ***Makes managing user credentials faster and easier.***
  - Available on z/OS V2.5 and 3.1 with **APAR OA65990.** PTFs are available!
  - For more information see the z/OS MVS Programming: Callable Services for High-Level Languages and a set of Game Changing Series blog posts in the DFSMS Community.

? **Why it Matters:** Multiple enhancements delivered via continuous delivery that simplify management of cloud objects from z/OS and enable more use cases leveraging object storage.

Cloud Storage

Cloud Data Access / Cloud Object Utility

- ***4Q24 new enhancement:*** Multiple enhancements intended to simplify management and use of cloud objects from z/OS.
  - **Skip Key File and Provider File Disk Read**
    - Requirement (ZOS-I-4080) to not read the key file and provider file from disk on every API call. Especially a concern for those applications where performance is critical and use of CDA session support with GDKINIT/GDKTERM (key file and provider file read one time) is not feasible.
    - Optional parameters for CDA APIs allow passing in of the key file contents as well as the provider file contents.
      - **providerFile** – CDA will bypass reading the provider file from the disk when this optional parameter is passed. CDA will make a copy of the passed data and use the JSON parser on that for the CDA operations.
      - **keyFile** – CDA will bypass reading the user's gdkkeyf.json file, instead making a copy of the value and using the JSON parser on that for the CDA operations.
    - ***Applications may read the data in a non-performance sensitive flow and pass data as needed to the CDA API.***
  - Available on z/OS V2.5 and 3.1 with **APAR OA65990.** PTFs are available!
  - For more information see the z/OS MVS Programming: Callable Services for High-Level Languages and a set of Game Changing Series blog posts in the DFSMS Community.

? **Why it Matters:** Multiple enhancements delivered via continuous delivery that simplify management of cloud objects from z/OS and enable more use cases leveraging object storage.

Cloud Storage

Cloud Object Utility

- ***4Q24 UPDATE:*** **&GDK_SECONDARY** symbol is now recognized and resolved for z/OS data sets. This value may also be passed as part of the metadata string passed on the **"metadata"** optional parameter.
  - ***4Q23 new enhancement:*** Requirement (ZOS-I-3616) to allow specification of metadata to be associated with the Cloud Object on an UPLOAD command
    - New keyword, **METADATA(<dd_name>)**, that specifies the DD name containing the key:value pairs that should be sent to the Cloud Object server for association with the object.
      - The provider file for the cloud object storage server will also need to be updated
    - The GDKUTIL LIST command can be used to display the metadata associated with a specified object.
    - Available on z/OS V2.5 or later with **APAR OA64874.** PTFs are available!
  - Available on z/OS V2.5 and 3.1 with **APAR OA65990.** PTFs are available!
  - For more information see the z/OS MVS Programming: Callable Services for High-Level Languages and a set of Game Changing Series blog posts in the DFSMS Community

? **Why it Matters:** Ability for cloud objects created by GDKUTIL to have configurable metadata included with the object.

Simplify and Modernize

RMM Plug-in for z/OSMF

- Support for a modern graphical user interface via a z/OSMF plugin
  - The plug-in augments the existing TSO and ISPF dialog support and consists of 4 tabs:
    - Volumes tab – displays filtered list of volumes and volume details
    - Monitoring tab – monitors current status of RMM (similar to LISTCNTL)
    - Defaults table tab – lists defaults table policies of RMM
    - Settings tab - settings for the WUI, including communication settings
  - Use the plug-in to view data sets defined to RMM and the related data set information – export this data to a CSV format file
  - See **APARs OA59499 and OA59727** on z/OS V2R4.

? **Why it Matters:** Intended to help users manage RMM in a modern way using latest UI and UX technologies available.

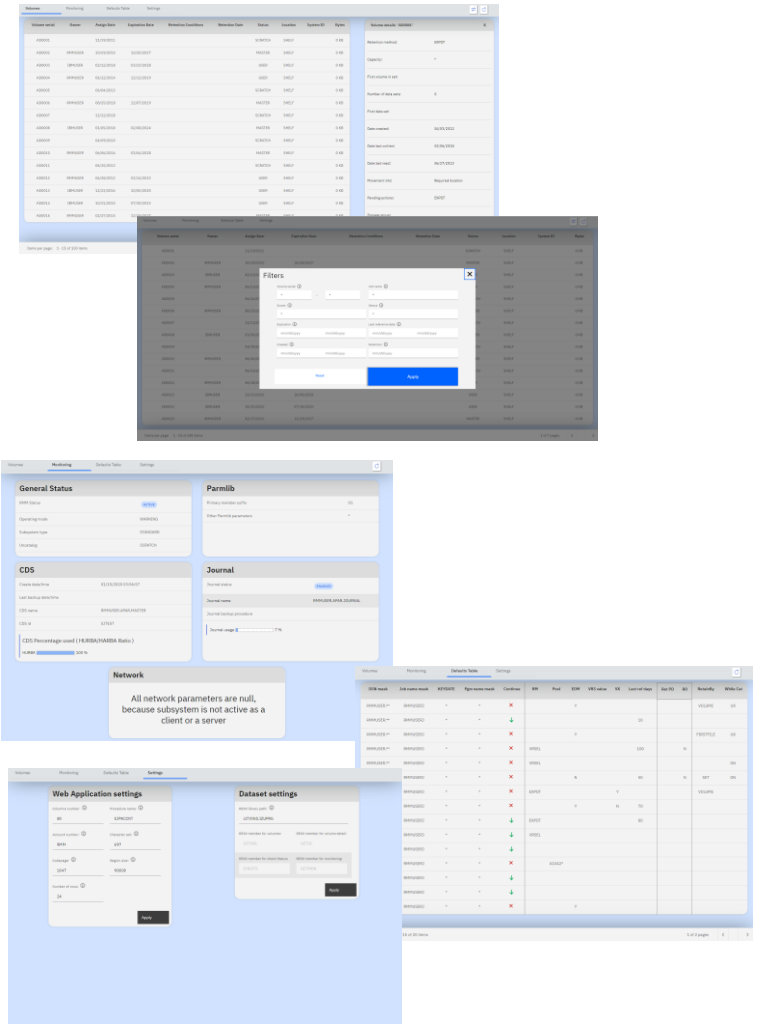

BACKUP

z/OS DFSMS Highlights

Scale and Performance

- zHyperLink Enhancements

Blue = continuous delivery

Scale and Performance

zHyperLink write support for multi-volume data sets

- With the DFSMS Media Manager support for zHyperLink writes, Db2 12 for z/OS can improve log throughput by exploiting zHyperLink for active log writes.
  - zHyperLink write processing can speed up commit by making the log write faster.
  - Designed to improve transaction response time, reduce the latency for log writes, and potentially speed up commits of the transactions.
- ***4Q22 CD :*** Enhancements to Media Manager provide support for multivolume data sets.
  - See **APARs OA59611** on z/OS V2R3 and above. PTFs are available!

? **Why it matters:** Expands the zHyperLink use cases for Db2 log data sets and IMS write ahead data sets.

102

IBM Z Batch Network Analyzer (zBNA)

IBM Z Batch Network Analyzer (zBNA)

- A no charge, “as is” tool to analyze batch windows
- Available to Clients, Business Partners, and IBMers
- Windows PC based, and provides graphical and text reports
  - Including Gantt charts and support for Alternate Processors
- Available on Techdocs
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132

Version 2.5.2
Available
July 18, 2023

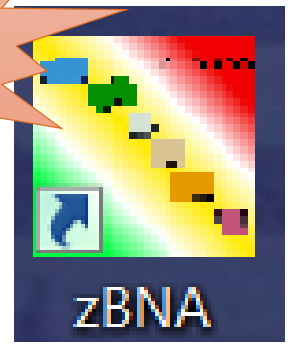

Enhanced for zHyperLink Support

- ***New with V2.4.2:*** Added zHyperLink write estimates for Db2 log data sets
- Top Data Set Candidate List for zHyperLink
- Data filtering capability
- Ability to estimate the benefit of zHyperLink I/O activity
- Generate customer reports with text and graphs to show zHyperLink benefit
- Provide support to aggregate zBNA LPAR results into CEC level views

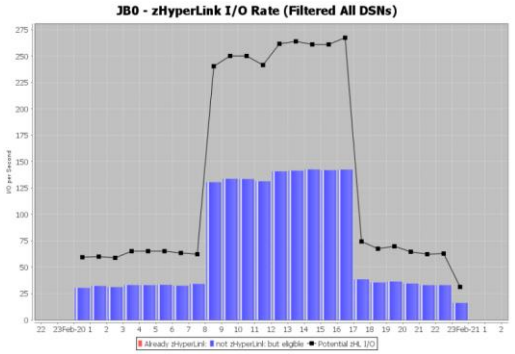

103

Storage Synergy

- **zHyperLink™**
  - New zHyperLink Technology accelerates transaction processing for z/OS
    - Lowering database transactional latency is critical to accommodate the use of new data sources and an increase in transaction volumes while enabling traditional and new mobile workloads to meet their SLA's.
  - **Critical components:**
    - **DS8880**
      - zHyperWrite protocols built into zHyperLink protocols for acceleration of database logging with continuous availability
      - Investment protection for clients that already purchased the DS8880
      - New zHyperLinks compliment, but do not replace, FICON channels
    - **zNext**
      - New I/O paradigm transparent to client applications for extreme low latency I/O processing
    - **z/OS, DB2 and VSAM**
      - New approach to I/O processing. Allows reduction of I/O interrupts, context switching, L1/L2 cache disruption and reduced lock hold times typical in transaction processing work loads.

? **Why it Matters:** Provide a simple, transparent and consumable approach to enable extreme low latency I/O for the acceleration of transaction processing for DB2 on z/OS and VSAM applications, batch window reduction, and improved scalability while leveraging the reuse of existing storage assets.

Storage Synergy - zHyperLink

- zHyperlink is fast enough so that the CPU can just wait for the data
- Point to point connection – 150 meters max distance from processor
- zHyperlink does not replace existing FICON channel paths
  - Traditional I/O used when request is not eligible for zHyperlink or if zHyperlink request fails (e.g., read cache miss) and for device initialization

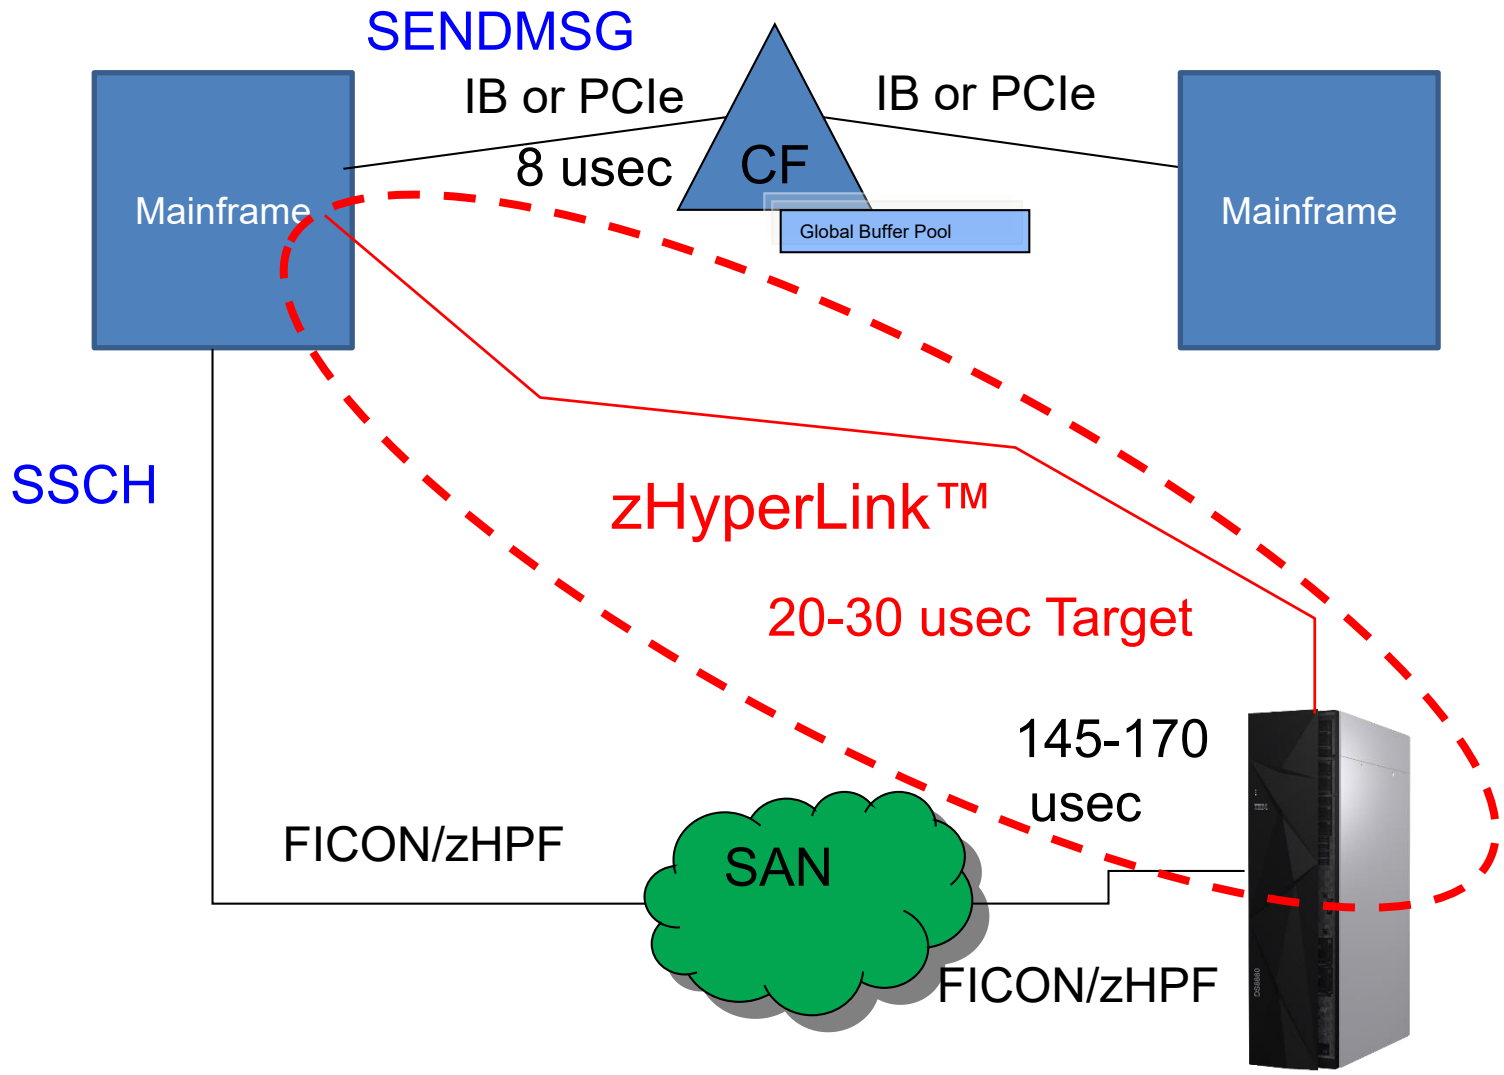

Trademarks and Disclaimers

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

CICS*
DB2*
DFSMS
DFSMSdfp
DFSMSdss
DFSMShsm
DFSMSrmm
DFSORT
DFSMS
DS4000
DS6000
DS8000
Enterprise Storage Server*
ESCON*
FICON*
FlashCopy*
GDDM*
GDPS*
geoManager*
HiperSockets
HyperSwap
IBM*
IBM logo*
ImagePlus*
IMS
Intelligent Miner
Language Environment*
Lotus*
MQSeries*
Multiprise*
OMEGAMON*
OS/390*
Parallel Sysplex*
PR/SM
QMF
RACF*
Rational*
RMF
System i
System z
System z9
System Storage
Tivoli*
TotalStorage*
Virtualization Engine
VisualAge*
VM/ESA*
VSE/ESA
VTAM*
WebSphere*
z/Architecture*
z/OS*
z/VM*
z/VSE
zSeries*
zSeries Entry License Charge

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both

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

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

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

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

Other company, product, or service names may be trademarks or service marks of others.

NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

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

Information is provided “AS IS” without warranty of any kind.

Trademarks and Disclaimers (continued)

NOTES:

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

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

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

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

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.