

Virtual IMS User Group
October 10, 2023



EDGE

Z/OS Enterprise Information Server Overview

Build a dedicated z/OS EIS LPAR for Information-as-a-Service to support Analytics, Machine Learning, and Artificial Intelligence

Stan Muse

smuse@edgeconsulting.com

October 2023

Agenda

- Business and technical challenges
- The AI imperative
- Solution overview
- Business justification
- A few use cases
- Technical architecture
- Implementation steps
- Summary + Q&A

This presentation will be provided as .ppt, .pdf and .wmv files for downloading

A Few Business and Technical Challenges

In today's rapidly changing banking environment, financial institutions face numerous business challenges, including:




- Reduced deposits, loan losses, Rising interest rates due to inflation
- Increased risk from customers
- Heightened political and government regulatory scrutiny
- People costs, wages and benefits, employee retention and equity

The AI revolution is in full swing and financial institutions are racing to implement solution to cut costs and add value, presenting many new technical challenges:

- Cloud managed services: PaaS, SaaS, IaaS
- Increased complexity: Implementation of distributed computing models, from monolithic programs to micro services
- Security: Pervasive encryption, Disaster Recovery, Cyber resiliency, Logical Copy Protection (LCP)
- Core systems modernization, legacy systems retirement
- Distributed data management: many new unstructured database types other than VSAM, IMS, Db2
- Implementation of new revolutionary technologies like Analytics, ML & AI are **not optional**

Larry Summers: AI will be the most important technological development since fire and the wheel.
Harvard Macro Economist, Former US Treasury Secretary

Dr. Jordan Peterson: We are all already androids
Harvard psychologist



The AI Imperative

- AI/ML will be implemented and used heavily by most successful companies and is therefore unavoidable
- AI/ML will be very disruptive for the job market, just as “Data Processing “was.
 - The first jobs that were eliminated by DP were many thousands of accountants, then Telco Operators, Secretaries, Gas Attendants, Bank Tellers, etc.
 - The modern digital age would not have been possible otherwise.
- AI/ML and Robotics will eliminate and enhance many jobs at all levels.
- The combination of AI/ML and Robotics will be even more disruptive and life enhancing and is unstoppable
- AI/ML have the potential to vastly improve our lives while also reducing waste
 - Google Maps (Generative AI), Tax accounting software
- AI/ML is only as good as the data it has access to and can use for producing results
- AI/ML can and will be used to do harm by some bad actors
 - Expect many new regulations from politicians who know nothing about the technology
 - How will techies force AI to be legally and socially compliant without sacrificing valid results?



AI and Robotics
Now on Mars



AI/ML expectations:

- **Implementation and widespread use will be disruptive but unavoidable in order to be survive in the new digital era**
- **Will vastly improve our lives while reducing wasted time, money, energy, pollution, etc., and with better than human –only results**
- **Will require as much or more compute power as transaction processing does now, with much more data access**
- **Any technology can be used to do harm, we must not fear and avoid AI**



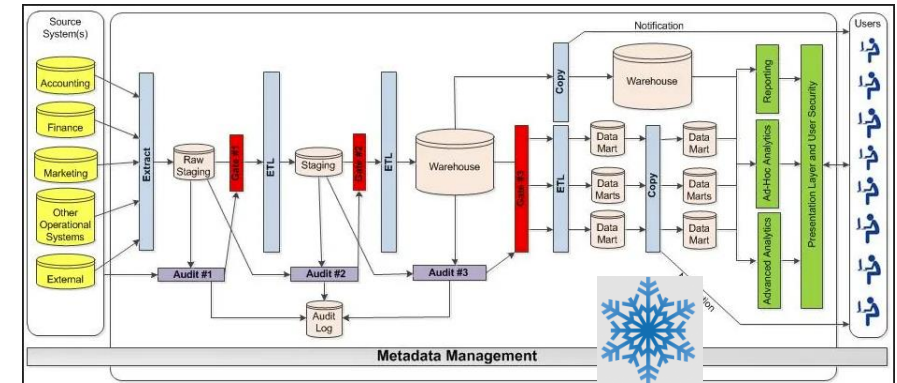


Enterprise Data Management Challenges Today

For decades clients have been proliferating data from the mainframe for analysis and providing services

- **ETL processes are complex, consume vast amounts of resources, and add no value to the data**
Daily ETL volumes have grown From gigabytes to terabytes to petabytes
ETL actually increases mainframe workload and costs
Data is often aggregated into a new schema, losing some meaning
- **Data duplication is very costly: Storage, Staff, Software Licenses**
COOs and CFOs can not quantify these costs in terms of TCO,
They are just waking up to the new Digital Age demands
- **Costly platforms, like Teradata, Exadata, and other boutiques consume IT growth budget**
The cost for cloud platforms like Snowflake and MongoDB are open ended
- **Data latency is a huge problem for many applications, like fraud detection & stock trading**
Day or week old data is not useful, losses continue to grow, many false positives
- **Data Security is compromised with thousands of ETL data feeds to who knows where**
Who needs to hack a mainframe? The lack of distributed data security results in expensive data breaches
Hackers may have many access points for data breaches and ransomware attacks
Data Encryption is difficult to manage over many copies on several platforms
- **Enterprise data architecture has become very complex or unmanageable**
Unsustainable server sprawl is a common problem, with runaway unquantifiable costs
- **Clients are struggling with how to provide mainframe Information-as-a-Service (IaaS)**
This is key to the success of Customer Analytics, Machine Learning, Artificial Intelligence, and many other new strategic solutions

ETL Multi-Process Steps



Enterprise Data Management and data access have become IT's most costly, and most visible, problems in the new Digital Age

Solution: Use the System of Record data in place with System z Enterprise Information Server

Goal: Provide seamless (primarily read-only) near real-time access to core systems data in place

- Allocate a separate z/OS LPAR as an Enterprise-wide Information Server
- Direct new AI/ML read only data access workload traffic to this new LPAR
- Stop expensive data proliferation to specialty servers
- Data Gravity – move the analysis to the data
- Single point of entry for data , especially for external users

Purpose: Establish System z as **THE** Corporate Data Server for:

- Cognitive Computing, AI, Machine Learning, Analytics and Reporting
- Mobile and distributed banking (90% may be read-only transactions)
- Data access for distributed or cloud applications
- Fraud Detection & Analysis

How: Provide templates for Operational Model, Project Plan, and Quick Start Services:

- Operational Model: minimal configuration for hardware, software, tools, connectivity
- Project implementation plan, Staffing and Education plan
- Z/OS and Db2, IMS, VSAM, Etc. services for Z/OS Sysplex data sharing implementation
- Cost/Benefit Analysis
- Negotiation points for future IBM purchases for read-only, non-revenue-producing workload pricing





Business Justification for zEIS

Reduce complexity, lower costs, improve security, reduce risk:

- Reduce Complexity:
 - Simplify data architecture, make the mainframe data easier to access & understand
 - One copy of data in place with one set of metadata
- Lower Costs:
 - Eliminate thousands of distributed servers and their software and support costs
 - Eliminate duplicate data costs, data privacy costs, data breaches
 - Stop paying for ETL processing, buying expensive ETL Tools and training
 - IBM is considering special hardware and software pricing for contained read-only workloads
- Improve Data Security:
 - Single entry point, pervasive encryption
 - Ability to Monitor Data Usage: internal and external, including ransom ware attacks
 - Ability to create a chargeback system for data usage for the digital bank or IaaS Clients
 - Better Data Governance: Data is growing exponentially, and most clients are not ready
- Reduce Risk: A place to try new solutions without impacting production systems
 - Avoid long mainframe production systems change control cycle for new data-centric solutions
 - Become more agile for delivering new strategic solutions like machine learning
 - Enable agile analytic development

A Cost/Benefits business case will be required because this will not be inexpensive and will need buy-in from the business units and executive management. Do not try to skip this task.

Top 10 Reasons for allocating a Dedicated z/OS LPAR as an Enterprise Information Server

1. Access Isolation

- Manage access to only one LPAR, instead of point-to-point solutions in every production LPAR
- Eases API Management, eliminates confusion about the best access solution
- Everyone using one copy of the data for one version of the truth
- Better Security: A single place from which to monitor data patterns and provide alerts
- Easier to shut down in case of detected denial of service, ransom ware intrusion, or data corruption

2. Workload Isolation

- The informational access workload is growing rapidly and will be as large or larger than the transaction and batch workloads today
- Isolating this workload will protect core transactional and batch systems running the business from runaway ad-hoc access
- Isolating this workload will allow for better hardware capacity planning, and performance tuning for all workloads, corruption detection
- Easier to understand where requests are coming from for planning: on-premises, cloud, or mobile

3. Better Hardware Resource Allocation and Usage

- Dedicated cores, (CPs, zIIPs, and specialty compression & encryption engines), without transactional interference are better for an informational workload
- Spare core capacity can also be allocated at a lower priority, allowing for higher overall core utilization for a CEC
- Dedicated memory for separate IMS, DB2 & VSAM buffers will provide better response times for both informational and transactional systems
- Communications traffic can be routed through dedicated OSA cards for no interference with current production applications
- Faster access for Linux on Z applications via Hypersockets or RoCE

4. Information Usage Chargeback

- Today there is no chargeback system for information access. This will be a requirement for TCO and ROI analysis
- Having a chargeback system for outside (B2B) access to core data could be a new revenue stream for the business
- Transform the Mainframe from a Cost Center to a Revenue Center

5. Infrastructure Simplification

- Having only one access point for informational requests greatly simplifies, and should reduce support costs, VS point-to-point solutions
- An easier to understand infrastructure aids with Agile Development, API Management, and new RESTful services for Business Units
- Provides a central place for access to an enterprise Metadata Manager for understanding the data
- Most FSS clients have thousands of distributed servers. Each distributed server is the equivalent of at least one LPAR.

Point to Point Architecture



Server Farm



Z16

Thousands of distributed servers could potentially be eliminated by just one or two Z LPARs

Top 10 Reasons for allocating a Dedicated z/OS LPAR as an Enterprise Information Server

6. z/OS Analytics Tools Availability

- Many analytics tools run well under z/OS or Linux on Z:
- New analytics and AI tools may not be allowed to run in the production transaction LPARs because of their data access patterns
- KPI dashboards and KYC analytics tools must be enabled to ALL enterprise and cloud data

7. Enterprise Data Virtualization & Federation for Real-time Analytics

- Near Real-Time access is required for many new applications, from mainframe and distributed databases for fraud detection or corruption prevention
- All data can be federated or joined, cleansed, summarized, and presented to the remote requestor, reducing the amount of data transmitted
 - [IBM Data Virtualization Manager](#)
 - [Infosphere Information Server](#)
 - [Infosphere Classic Federation Server for z/OS – DB2 & VSAM](#)
 - [Common Data Provider for Z Systems](#)
 - [z/OS Platform for Apache Spark, MDSS](#)
 - [Also OEM ETL Tools: Informatica, Syncsort, Veristorm, Open Source tools, RYO](#)
- Much ETL work may be eliminated with direct data access

8. Build an Active-Active Core Banking Applications Environment

- An Active-Active environment requires at least two swappable production LPARS,
- Plus a spare LPAR when servicing one of the two main production LPARS. This could be that.
- Also, provides the ability to test certain read-only functions before deployment to production.

9. Reduce Costs: Take Advantage of IBM Special Pricing for Hardware & Software

- IBM provides software Container Pricing for this solution & new workload which greatly reduces MLC costs
- Special pricing for new workload hardware via special bid
- Reduce the Number of Software Licenses: Point-to-Point solutions may require software licenses in every CEC/LPAR, increasing software costs

10. Ability to Respond to Regulators and Auditors

- GDPR is not just for the EU, it's Global, and US firms must be able to quickly respond to requests
ISO and other regulations will be just a stringent and complicated
- The ability to show data flows from System of Record source to final target database & user is required
and will be easier from one place Inability to respond comes with stiff penalties
- Provide consistent data obfuscation & masking for compliance from a single control point with [Optim Data Privacy](#)



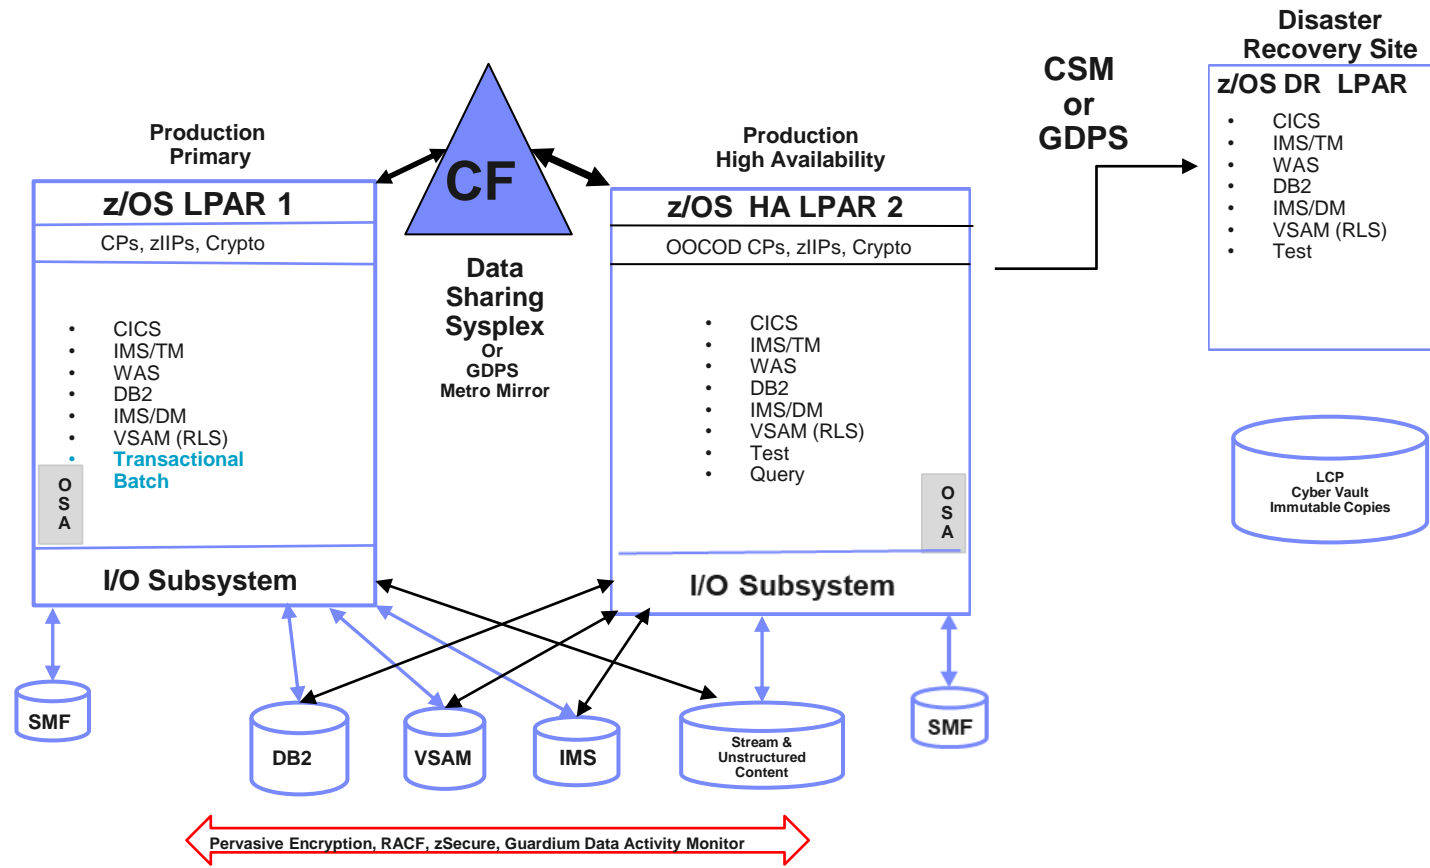
A Few zEIS Use Cases

zEIS can enable many new data uses:

- Near real-time access to core systems data as well as distributed data bases from a single view
- Separation of distributed ad-hoc query workload from production batch & transactional LPARs for resources, tuning, availability
- Everyone accessing and using the same data - Single version of the truth with one set of metadata
- Easier to grant/provide access than existing production Online & Batch LPARs, easier change control for new solutions
- Provide a Metadata catalog & defined interfaces, SQL and APIs to enterprise data and micro-services for easier access
- Ability to monitor usage patterns & stop access for intrusion detection, denial of service, data corruption, ransomware
- Employ data analysis & ad-hoc query tools that are not available on production Online & Batch LPARs
- Provide SQL access to IMS, VSAM, Flat File, SMF, Unstructured and Streaming data with data federation
- Stop purchasing & importing the same social media and industry data over and over
- Enable new technologies:
 - Hybrid Cloud applications access
 - Machine Learning, Cognitive Computing, Watson
 - Fin Tech Industry Solutions data access
- Better data security, eliminate the need for ETL & sending data everywhere
 - Keep data warehouse, data lakes on the mainframe
 - Provide a data obfuscation layer
 - z/OS DB2 Views using IIS or IDVM can expose (select & project) only data needed from all data sources

Typical Mainframe Production Data Sharing HA Conceptual Model:

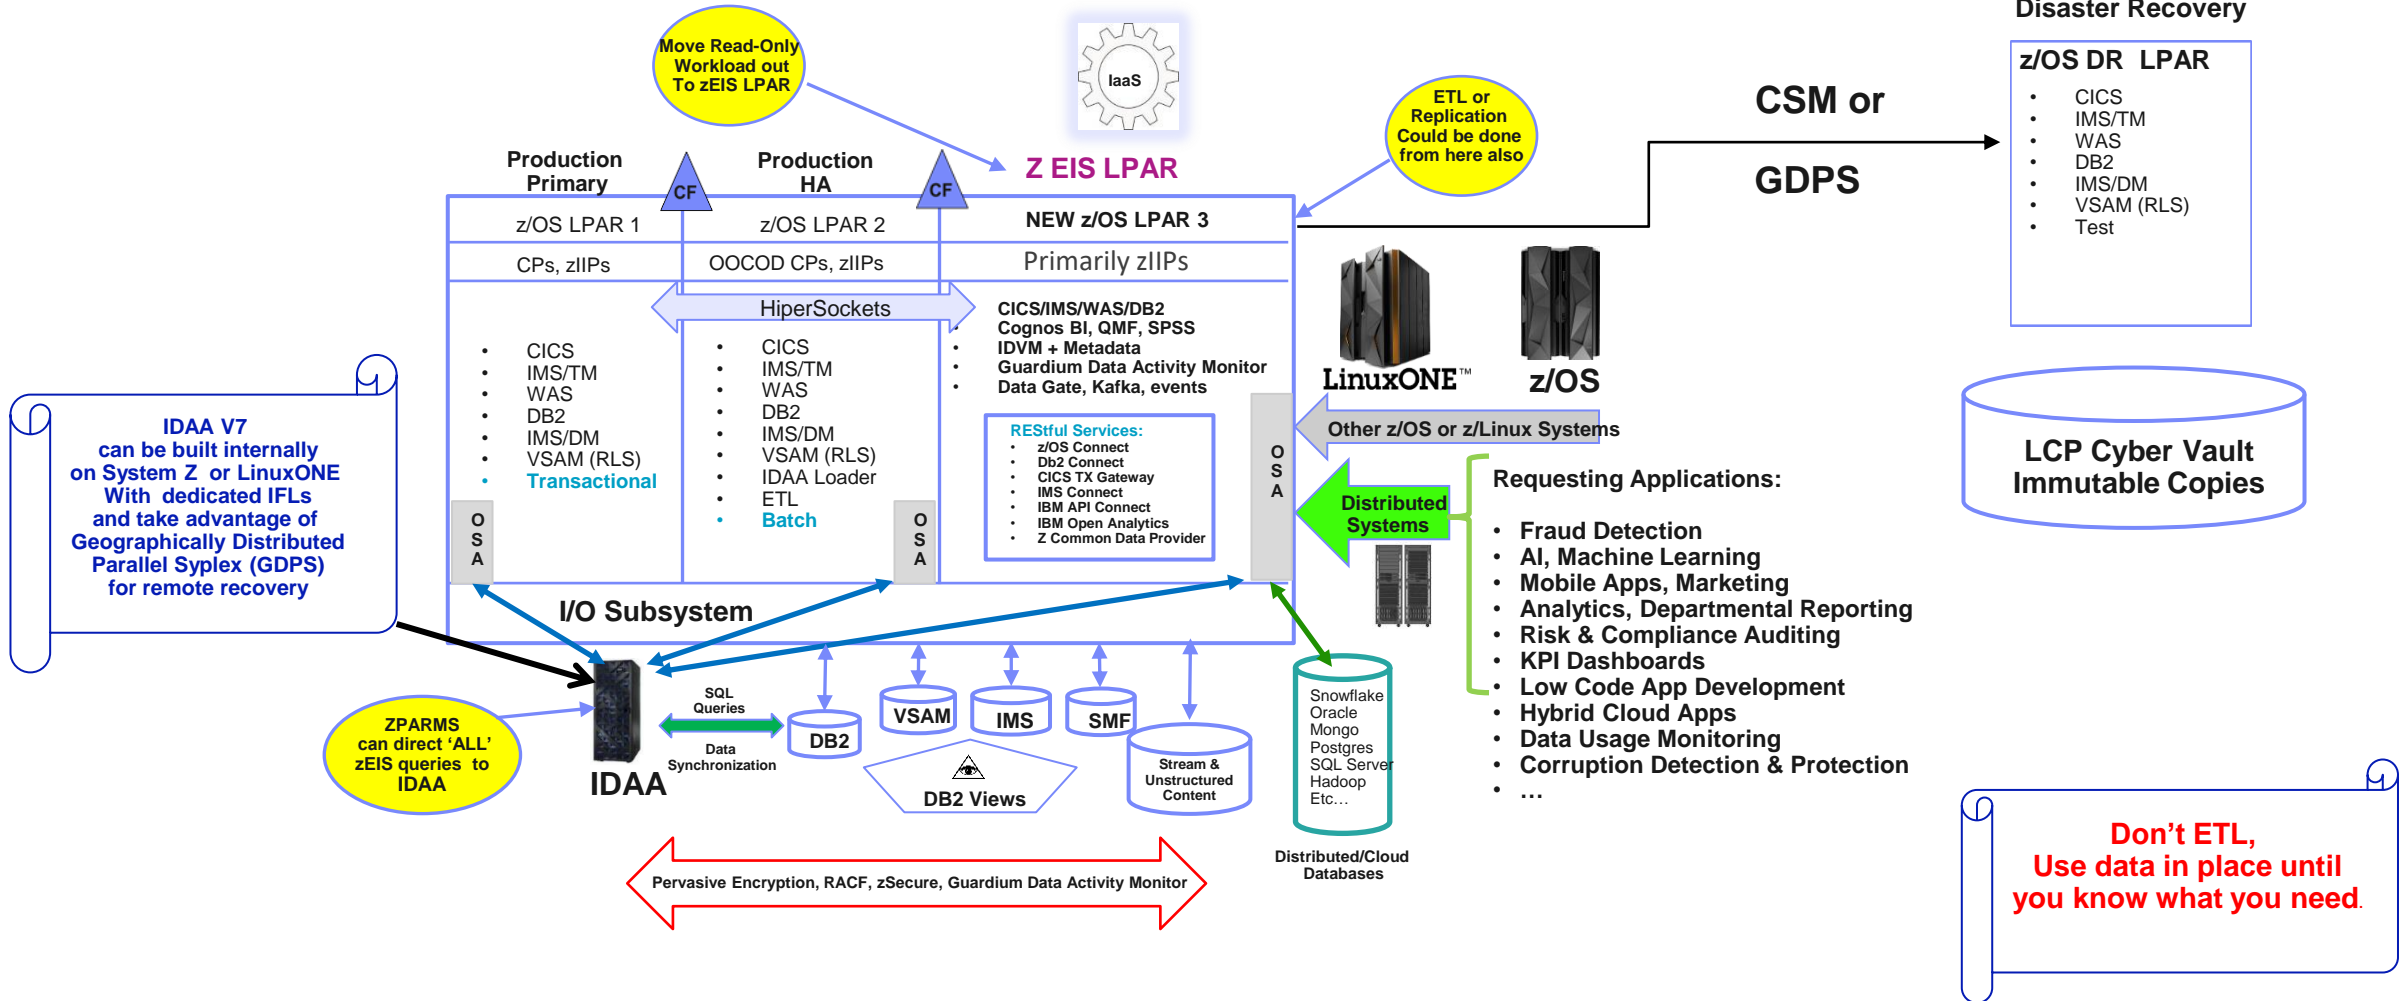
Most sophisticated clients already have at least TWO PRODUCTION LPARS in an active-active data sharing environment for high availability, plus remote site DR



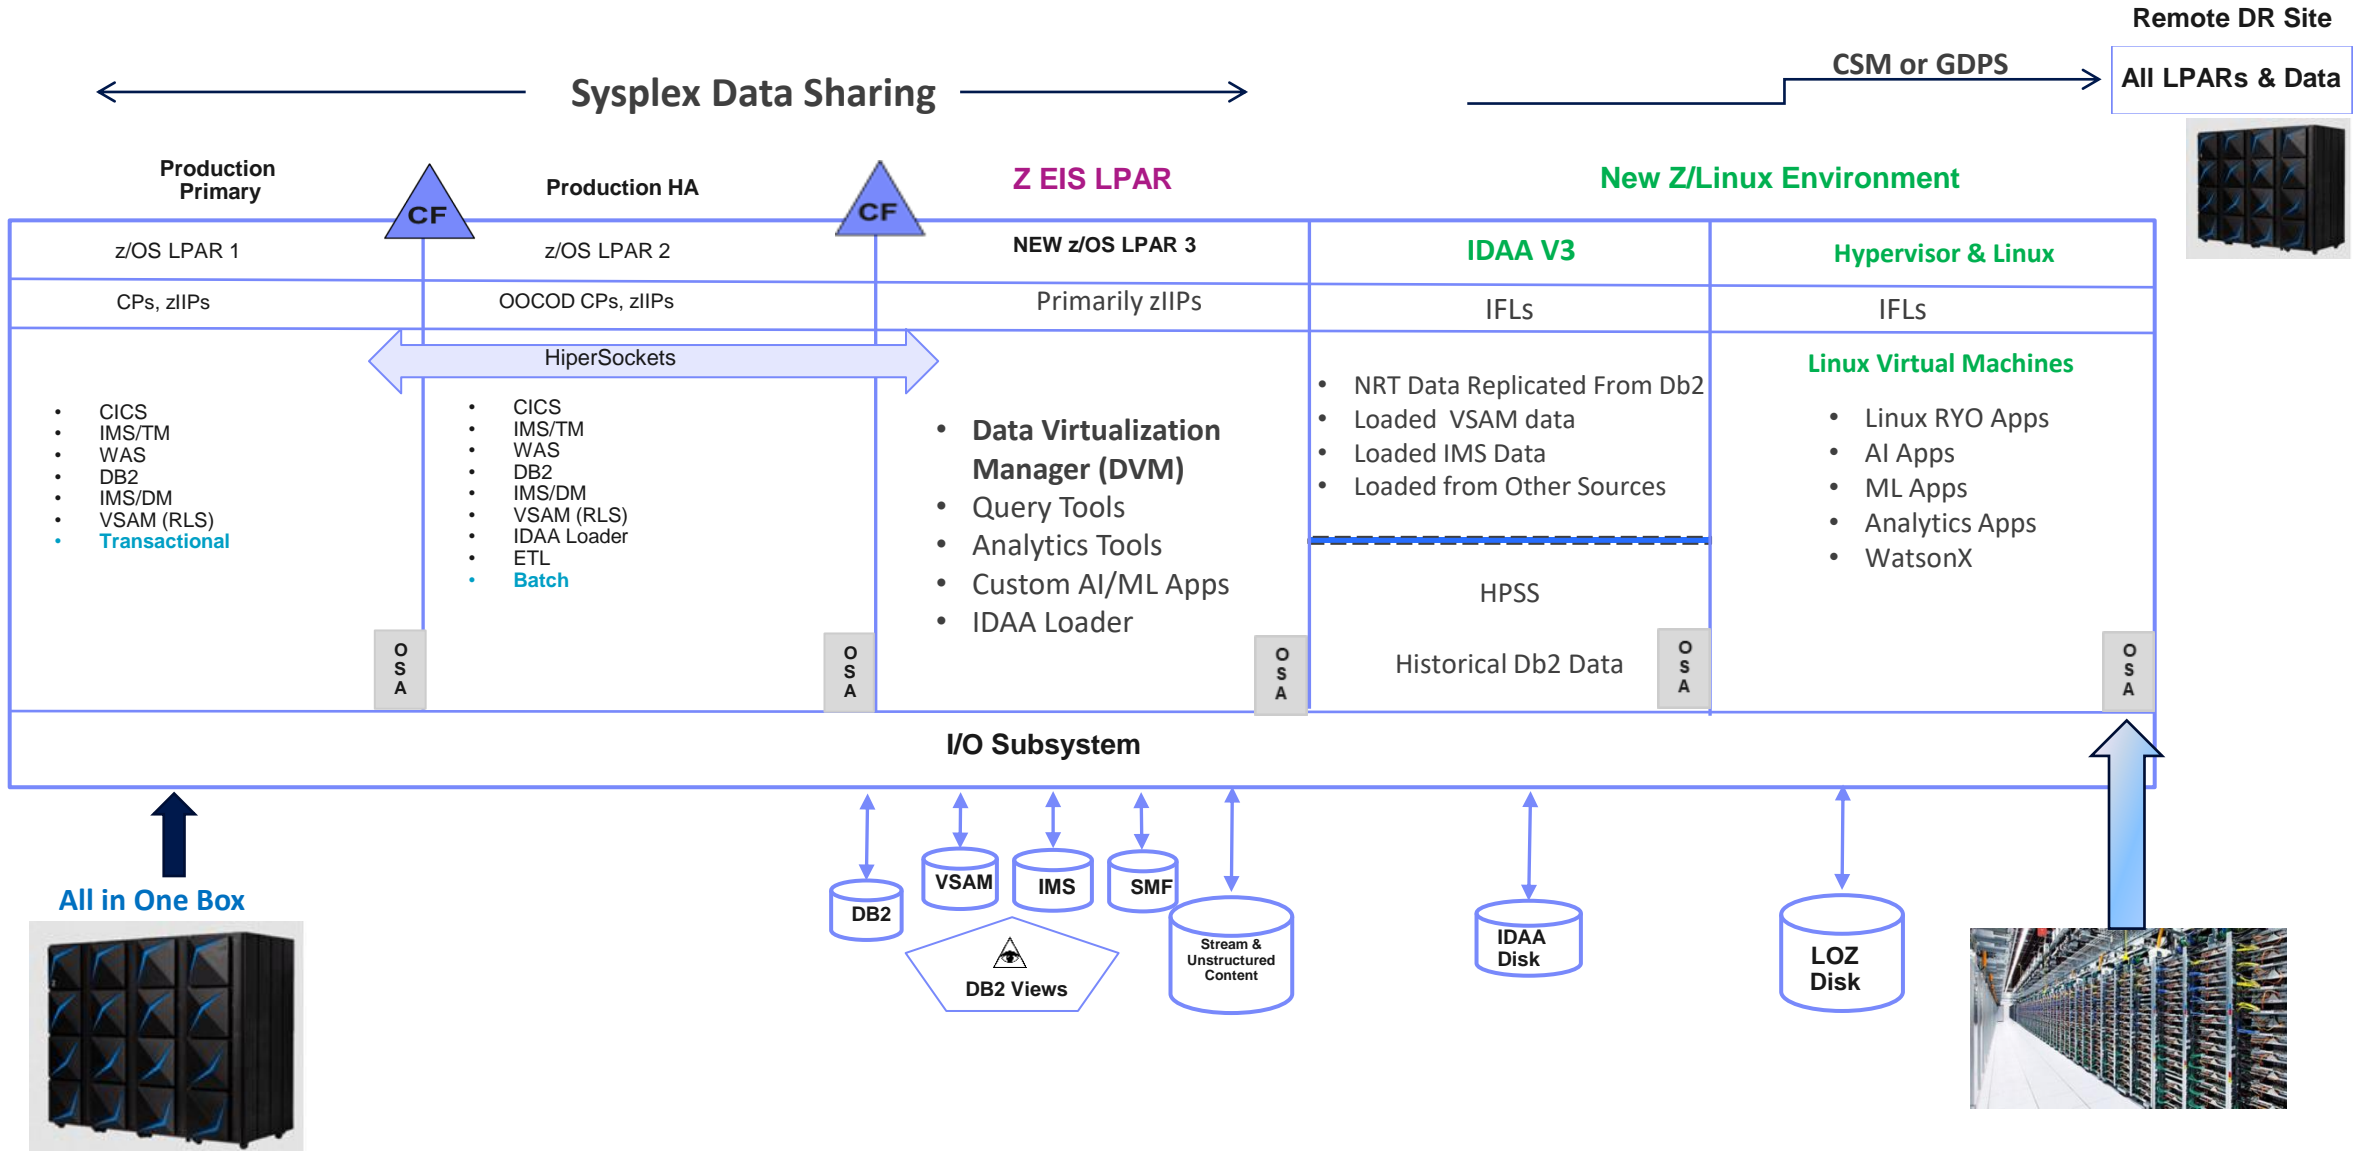
zEIS LPAR Conceptual Model:

A central access point or data bus for all data flowing from the mainframe and distributed servers

Next, Implement a separate, primarily read-only, z/OS LPAR as an Enterprise Information Server (zEIS) to provide Information-as-a-Service for the entire enterprise



Expanding the Architecture with IDAA and Linux on Z



CF, zEIS LPAR, IMS, Db2, and VSAM Data Sharing Implementation Steps

1. Allocate Coupling Facility LPAR on the HMC
2. Allocate the new zEIS LPAR on the Hardware Management Console (HMC)
3. Set up IMS Data Sharing Groups
4. Set up Db2 Data Sharing Groups
5. Set up VSAM Data Sharing Groups
6. Install Data Virtualization Manager (DVM) for accessing & joining all data via SQL
7. Connect query & AI tools access via DVM

**Note: IMS Block Level Data Sharing will probably require implementation of Geographically Distributed Parallel Sysplex (GDPS)
Edge Consulting has performed over 50 of these engagements for mainframe customers.**

Coupling Facility Implementation Steps

1. Use an existing CF or create a new one?
 - Use existing if possible
2. Review coupling facility sizing and requirements
 - Processing workload , Minimal locks but, Data Buffer considerations
 - Structure sizing
3. Format and define couple data sets, Primary and Alternate
 - Sysplex CDS
 - Function CDS
 - Couple Data Set sizing analysis
 - CFRM, ARM, SFM, WLM, etc. policies
4. Plan common timing reference - STP
5. Plan XCF signal paths
 - CTCs
 - Signal (List) structures
 - Define transport classes (COUPLExx)
6. Perform HMC actions
 - Define CF LPAR image
 - Verify CFCC levels and compatibility maintenance
 - Activate LPAR image
7. Systems must IPL to join Parallel Sysplex
 - Review relevant SYS1.PARMLIB members
 - Validate all systems have the same:
 - Sysplex name
 - CDS
 - signal connectivity
 - time reference



More Information LPAR Design Tool

<https://www.slideshare.net/abassani/lpar-design-tool-v8-t01-for-ibm-z-systems>

Relevant IBM Doc (hyperlinks)

[Planning for a coupling facility - IBM Documentation](#)

[Sharing a time source among LPARs - IBM Documentation](#)

[Planning signaling services in a sysplex - IBM Documentation](#)

[Planning the couple data sets - IBM Documentation](#)

[Defining a coupling facility - IBM Documentation](#)

IBM Z Server Time Protocol Guide – SG24-8480

Allocate the new zEIS LPAR Implementation Steps



1. Determine initial LPAR hardware requirements
 - Existing CEC or new CEC?
 - Network/Communication requirements
 - DASD requirements
 - Performance/processor requirements, at least one General Processor, several zIIPs
2. Plan and prepare system build requirements
 - Consider SYSRES, master catalog, DFSMS, WLM, VTAM/TCPIP, etc.
 - Review all members of new SYS1.PARMLIB
3. Define new system parameters in HCD, creating new IODF
4. Perform HMC actions
 - Define LPAR image, at least one GP, several zIIPs
 - Load new IODF to Support Element, creating new IOCDs
 - Activate LPAR image
 - IPL into new LPAR Install or Clone needed software
 - IMS
 - Db2
 - CICS
 - ISV products
 - Tools
 - **DVM**
 - Etc.

More Information

LPAR Design Tool

<https://www.slideshare.net/abassani/lpar-design-tool-v8-t01-for-ibm-z-systems>

Relevant IBM Doc (hyperlinks)

[Installation checklist - IBM Documentation](#)

[Making a copy of your system software \(cloning\) - IBM Documentation](#)

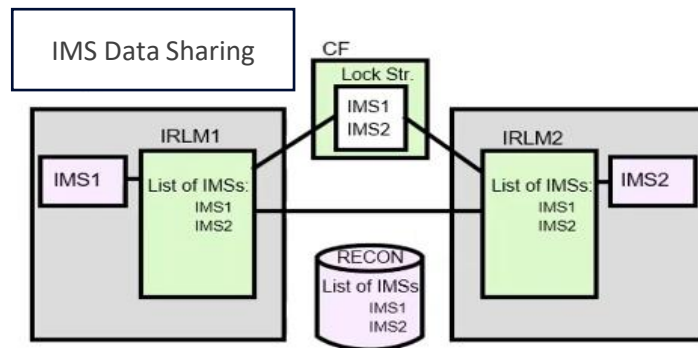
[Planning the couple data sets - IBM Documentation](#)

[Abstract for HCD Planning - IBM Documentation](#)

[Preparing the target system - IBM Documentation](#)

zEIS LPAR IMS Data Sharing Implementation Steps

1. Determine which IMS Subsystems/Databases will be Shared for zEIS LPAR (Read-Only)
 - Assuming first IMS Production Subsystem is already established
 - Use Block Level Data Sharing (BLDS), IRLM Required
 - HIDAM/PHIDAM, HDAM/PHDAM, HISAM, SHISAM, Secondary Indexes, DEDB, HSAM can be shared with BLDS
 - MSDB, GSAM cannot be shared
2. Register Databases SHARELVL(3) (Multiple LPARS & IRLMs)
 - Sharing is at the block level by multiple IMS subsystems on multiple Realms
3. Define ACCESS=RO or RD for read-only access to databases
 - PSB PROCOPT=GO (Dirty Read) or G (Read with Integrity) or GOT
 - IRLM with SCOPE=GLOBAL, understand IMS GLOBAL Commands
 - Several parameters to customize (eg. DEADLOK=, LTE=, TRACE=, LOCKTIME=,...)
 - **/Start DB dbname GLOBAL ACCESS=RO DBALLOC NOBACKOUT OPEN**
4. VSAM SHAREOPTION(3 3) and DISP=SHR
 - to allow DFSMSHsm to be started in a multiple DFSMSHsm host environment with GRS or equivalent function
5. Define structures in CFRM policy
6. CFNAMES statement: three coupling facility structure names must be passed to IMS: IRLM, OSAM, and VSAM
 - Each FF DB I/O requires CF access for OSAM/VSAM structures



More Information

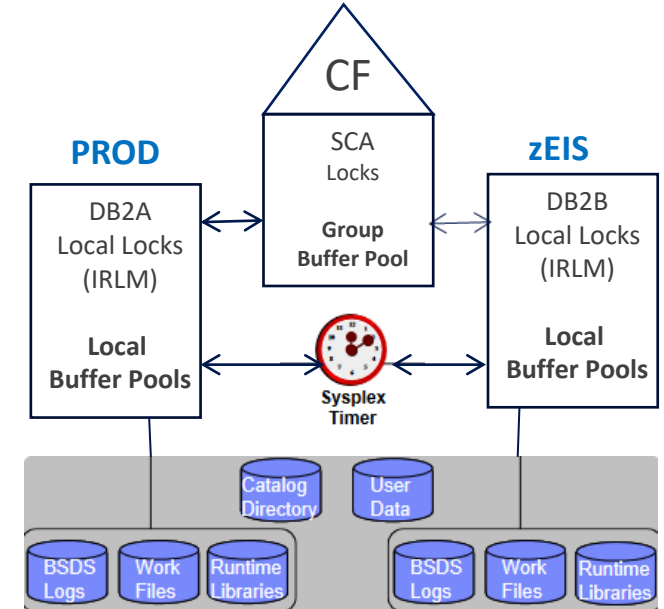
IMS in the Parallel Sysplex

- Volume I: Reviewing the IMSplex Technology – SG24-6908
- Volume II: Planning the IMSplex – SG24-6928
- Volume III: IMSplex Implementation and Operations – SG24-6929
- IMS Data Sharing Course with badge (CREDLY):

<https://community.ibm.com/community/user/ibmz-and-linuxone/blogs/michelle-carey1/2022/09/02/ims-data-sharing-course>

zEIS LPAR DB2 Data Sharing Implementation Steps

1. Determine which Db2 Subsystems/Databases will be Shared for zEIS LPAR (Mostly Read-Only)
 - Assuming first DB2 Production Subsystem & members & IRLM is already established
2. Run the CLIST for adding a new member (Tailors 63 DSN***** jobs to your environment)
3. Define the new Db2 data sharing member to IRLM and MVS: DSNTIJMV
4. Define address space startup procedures for the new Db2 data sharing member: DSNTIJMA
5. Alter the current active log data sets, DSNTIJIN uses SHAREOPTIONS (2 3)
6. Define Db2 initialization parameters
 - define initialization parameters in the subsystem parameter module (DSNZPxxx)
 - application defaults load module (DSNHDECP)
 - define the offline message generator CCSID (DSNHMCID),
 - update DDF related information in the BSDS
7. Initialize system data sets:
 - Initialize its BSDS and active log data sets after you define them and the member's Db2 initialization parameters
8. Define user authorization exit routines (optional)
9. Record Db2 data to SMF (optional)
10. Establish subsystem security (optional)
 - Who can access, also RACF
11. Connect IMS/TM and/or CICS to Db2 (optional)
12. Connect Db2 to TSO
13. Define Db2 to z/OS: After you run job DSNTIJMV, you must define Db2 in SYS1.PARMLIB
14. Start the Db2 subsystem
15. Create the default storage group, define temporary work files, and bind Db2 REXX Language Support
16. Create an image copy of the Db2 directory and catalog (optional)
17. Verify that enabling data sharing was successful (optional)
18. Ensure that the WLM address space is defined and available



Source: Mark Rader, IBM

More Information

Redbook: DB2 Data Sharing

- DB2 for z/OS: Data Sharing in a Nutshell– SG24-7322 (2007)

Installing or migrating to Db2 13

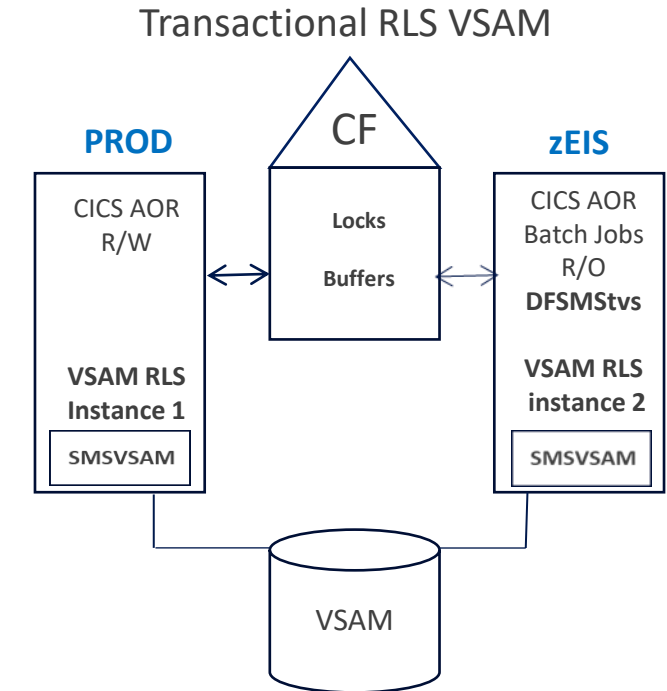
<https://www.ibm.com/docs/en/db2-for-zos/13?topic=installing-migrating-db2-13>

<https://www.ibm.com/docs/en/db2-for-zos/13?topic=sharing-adding-new-data-member>

zEIS LPAR VSAM Data Sharing Implementation Steps



1. Determine which VSAM Subsystems/Files will be Shared for zEIS LPAR (Read-Only)
 - Assuming first VSAM Production Subsystem is already established
 - Uses the same VSAM interfaces & data format
 - KSDS, ESDS, RRDS & VRRDS
 - **Must be SMS-managed**
 - One server per MVS image
 - Access mode specified on ACB or on JCL
 - Record level locking via CF, Serialization is at the record level rather than at the CI level
2. Define VSAM data set as **Recoverable**
 - READ/WRITE (GET/PUT) sharing across CICS AORs and
 - READ ONLY (GET) sharing with/across Batch Jobs
 - LOG(UNDO), LOG(ALL)
3. Specify RLS access in the ACB or the JCL
 - Open for input with read No Read Integrity option NRI
 - Or Open for input with Read Integrity (Consistent Read) CR - for completed updates only
 - Do NOT use “repeatable read” Consistent Read Explicit CRE due to locking
4. Interface with the RRS component of RRMS to define sync points (commit or back out)
5. Implement needed SMS SYS1.PARMLIB Changes



More Information

www.redbooks.ibm.com

- Transactional VSAM Presentation Guide SG24-6973
- Transactional VSAM Overview and Planning Guide SG24-6971
- Transactional VSAM Application Migration Guide SG24-6972
- VSAM Demystified SG24-6105

Transactional VSAM - Ruth Ferziger, IBM Retired

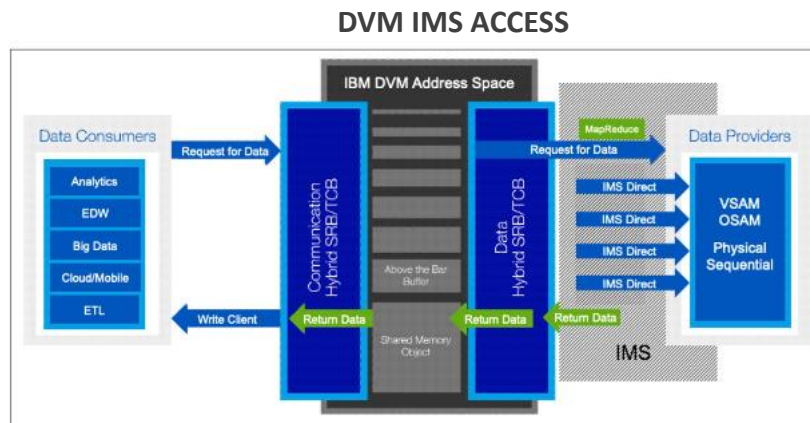
<https://www.vm.ibm.com/events/2005Z33a.PDF>

<https://www.vm.ibm.com/events/2005Z34a.PDF>

IBM Data Virtualization Manager Implementation Steps

IBM Data Virtualization Manager for z/OS® installs by using standard SMP methods common to other z/OS® software product installations and can be completed in ten steps.

1. Create the server data sets through the hlq.SAVZCNTL members AVZDFDIV, AVZGNMP1, and AVZEXSW1.
2. Set up the security application to use with the server by using one of the following hlq.SAVZCNTL members:
 - AVZRAVDB
 - AVZA2VDB
 - AVZTSVD
3. Configure Workload Manager (WLM) for optimum performance of the server.
4. Authorize the product LOAD library by using the Program Facility (APF).
5. Create a copy of the product libraries.
6. Configure the server to support DBCS to enable online use of databases.
 - For IMS, MapReduce must be enabled in the IN00 configuration file when IMS-Direct is turned on.
If you are performing an INSERT, UPDATE, or DELETE, IMS-Direct switches to DBCTL automatically, even when IMS-Direct is enabled
7. Customize the server to access your data sources in hlq.SAVZEXEC(AVZSIN00).
8. Configure the started task JCL that is in hlq.SAVZCNTL(AVZ1PROC) before you can start the server.
9. Configure the Command List (CLIST) that starts ISPF panels by using hlq.SAVZEXEC(AVZ).
10. Verify the installation by creating a virtual table and accessing its underlying VSAM file.



IMS Direct is fully zIIP-eligible and faster. When IMS Direct is unavailable, the DVM server uses DBCTL. Enable the IMSDIRECTENABLED parameter configuration member for IMS DIRECT access

IMS Open Database Access (ODBA) provides a callable interface that enables any z/OS recoverable, resource-managed z/OS address space, including the DVM server, to issue DL/I database calls to IMS.

More Information

Install Demo: <https://www.youtube.com/watch?v=h8PhKe8edGk>
IBM Data Virtualization Manager for z/OS SG24-8514

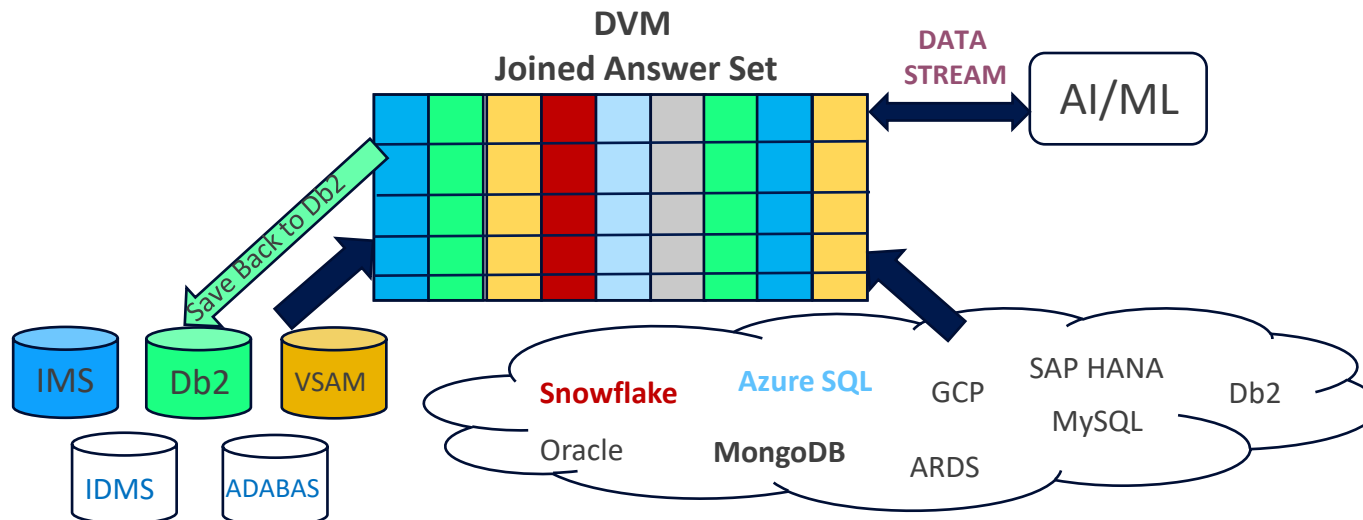
- <https://www.redbooks.ibm.com/redbooks/pdfs/sg248514.pdf>

IBM Data Virtualization Manager Benefits for AI/ML



IBM Data Virtualization Manager provides ease of access and several other benefits:

1. Access relational and non-relational IBM Z data through modern APIs including
 - HTTP
 - SOAP,
 - SQL and no SQL (REST when combined with IBM z/OS Connect)
2. Direct real-time read/write access to relational and traditional non-relational IBM Z data sources
 - Ability to save answer set back to Db2 for AI accountability
3. Mainframe databases include: Db2 IMS VSAM IDMS ADABAS, Oracle
4. Security: Exploits IBM Z security features such as pervasive encryption and security protocols
 - RACF, ACF2 and Top Secret
 - Provides native database-level security within Db2
5. Cost & Performance: Runs on IBM z Systems® Integrated Information Processors (zIIP)
 - Efficient and scalable for getting and combining data from many structured and unstructured sources



More Information

DVM Users Guide

- https://www.ibm.com/docs/en/dvm/1.1.0?topic=SS4NKG_1.1.0/havuga10/topics/kc_welcome_user.html

IBM DVM for z/OS SG24-8514

- <https://www.redbooks.ibm.com/redbooks/pdfs/sg248514.pdf>

First zEIS Project POC Steps



- Understand your pain points, get an Executive Sponsor (CIO) and Business Unit sponsors
- Present the zEIS LPAR concept, socialize, gain commitment to proceed, free Edge zWorkshop
- Identify a BU sponsored POC, document criteria for success, commitment to move forward
- Use test system LPAR for POC with IBM Data Virtualization Manager or query tool
- Allocate a small z/OS zEIS LPAR : (1) CP, (2) zIIPs, (1) TB memory
- zEIS LPAR POC's:
 - Build the entire z/OS EIS and z/Linux environment on ZD&T (Parallel Sysplex Edition) – Runs under RedHat Linux on X86
 - Demonstrate z/OS Connect RESTful access to DB2 & VSAM data via CICS or IDVM
 - Demonstrate access from z/OS to distributed databases via IDVM
 - Move a small ETL workload to zEIS, convert to direct data access
 - Provide direct data access to a BU who has requested ETL
 - Use DVM to access SMF data & develop a utilization report
 - Install & Implement IBM Data Activity Monitor
 - Use DVM to join DB2 & VSAM data for a BU, analyze cost VS ETL
 - Develop a metadata catalog with DVM
 - Develop some models for a BU with SPSS Modeler, IDAA
 - Address new reporting or analytics BU backlog & build it
 - Demonstrate data access from, z/Linux apps
 - Build a z/OS CICS micro service with API

Edge Consulting can provide Project Management, Architecture, Z/OS Sysplex Systems Programming, Database services for success.

More Information

IBM zD&T 12.0.0

- <https://www.ibm.com/docs/en/zdt/12.0.0> IBM DVM for z/OS

IBM zD&T Demos:

- https://mediacenter.ibm.com/media/Quick+LookA+Provisioning+a+ZD%26T+Test+Environment/1_j9mflk0y

Summary

- All industries must embrace AI and quickly move into a new digital data business model
- 90%+ of Financial Services customer data originates on System z
- 95%+ of Financial Services core applications workload runs on System z
- 90% - 95% of mainframe workload is Read-Only
- Many Financial Services clients run all core production applications in a single z/OS LPAR, (Mainframe Middle Ages) while they have thousands of distributed servers, with many ETLing data daily from z/OS
- The Mainframe's strength is data management and its shared everything architecture: Processors, I/O subsystem, huge memories, connectivity
- The new digital data distribution workload will be as big, or bigger than, current batch and transactional applications work, Core Operational Systems data must be enabled to serve the new digital workload, while being isolated and protected
- Firms may have no choice but to build the zEIS to comply with new regulations for protecting data and providing AI equity
 - Must be able to protect and recreate all data used to produce AI Results
- Establishing and employing a new z/OS Enterprise Information Server LPAR to handle the digital information distribution workload may be the **key to the survival** of financial services, and other industry, firms

Edge Consulting can assist with all aspects of building the zEIS server

Starting with a short scoping engagement for the project

<https://www.edgesolutionsandconsulting.com/>

AI on YouTube:

https://www.youtube.com/watch?v=lkxspUt0_gY
<https://www.youtube.com/watch?v=YzcaKGWEB1Q>
<https://www.youtube.com/watch?v=qxpbTyeDZp0>



- z/OS MVS Setting Up a Sysplex - SA22-7625
- z/OS MVS Programming: Resource Recovery - SA22-7616
- z/OS MVS Diagnosis: Tools and Service Aids - GA22-7589
- z/OS MVS Programming: MVS Assembler Services Guide - SA22-7605
- z/OS MVS Programming: Authorized Assembler Services Reference, Volume 2 - SA22-7610
- Z/ OS Parallel Sysplex Configuration Cook Book, Volumes 1-3 - SG24-2075, SG24-2076, SG24-2077
- CICS Recovery and Restart Guide, SC33-1698
- CICS Transaction Server for z/OS 6.1
 - <https://www.ibm.com/docs/en/cics-ts/6.1>
 - <https://www.ibm.com/docs/en/cics-ts/5.2?topic=tuning-using-vsam-record-level-sharing>
- Transactional VSAM Presentation Guide SG24-6973
- Transactional VSAM Overview and Planning Guide SG24-6971
- Transactional VSAM Application Migration Guide SG24-6972
- Transactional VSAM – Ruth Ferziger, IBM Retired
 - <https://www.vm.ibm.com/events/2005Z33a.PDF>
 - <https://www.vm.ibm.com/events/2005Z34a.PDF>
- VSAM Demystified SG24-6105
- DFSMSStvs Application Migration Guide SG24-6972
- IMS in the Parallel Sysplex
 - Volume I: Reviewing the IMSplex Technology – SG24-6908
 - Volume II: Planning the IMSplex – SG24-6928
 - Volume III: IMSplex Implementation and Operations – SG24-6929
- Free IMS Data Sharing Course with badge (CREDLY):
<https://community.ibm.com/community/user/ibmz-and-linuxone/blogs/michelle-carey1/2022/09/02/ims-data-sharing-course>
- DB2 for z/OS: Data Sharing in a Nutshell– SG24-7322 (2007)
- Installing or migrating to Db2 13
 - <https://www.ibm.com/docs/en/db2-for-zos/13?topic=installing-migrating-db2-13>
 - <https://www.ibm.com/docs/en/db2-for-zos/13?topic=sharing-adding-new-data-member>
- LPAR Design Tool
<https://www.slideshare.net/abassani/lpar-design-tool-v8-t01-for-ibm-z-systems>
- White Paper: Coupling Facility Configuration Options
https://www.researchgate.net/publication/342570694_Coupling_Facility_Configuration_Options_-_Updated_2020
- IBM Redbook IBM Data Virtualization Manager for z/OS SG24-8514
<https://www.redbooks.ibm.com/redbooks/pdfs/sg248514.pdf>
- IBM Z Development and Test Environment: <https://www.ibm.com/docs/en/zdt/12.0.0?topic=overview>



Trademarks

IBM®, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Red Hat® and Ansible® are trademarks or registered trademarks of Red Hat, Inc. 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.

Other product and service names might be trademarks of IBM or other companies.



Thank You for Attending!

Please email me for any follow up questions or help with the project, don't be shy

Stan Muse

smuse@edgeconsulting.com

www.linkedin.com/in/stan-muse

<https://www.edgesolutionsandconsulting.com/>

