Digital Transformation using the IBM Z Digital Integration Hub

Roy Duke Jr.
royduke@us.ibm.com
WW Sales Leader, IBM Z Digital Integration Hub

Agenda

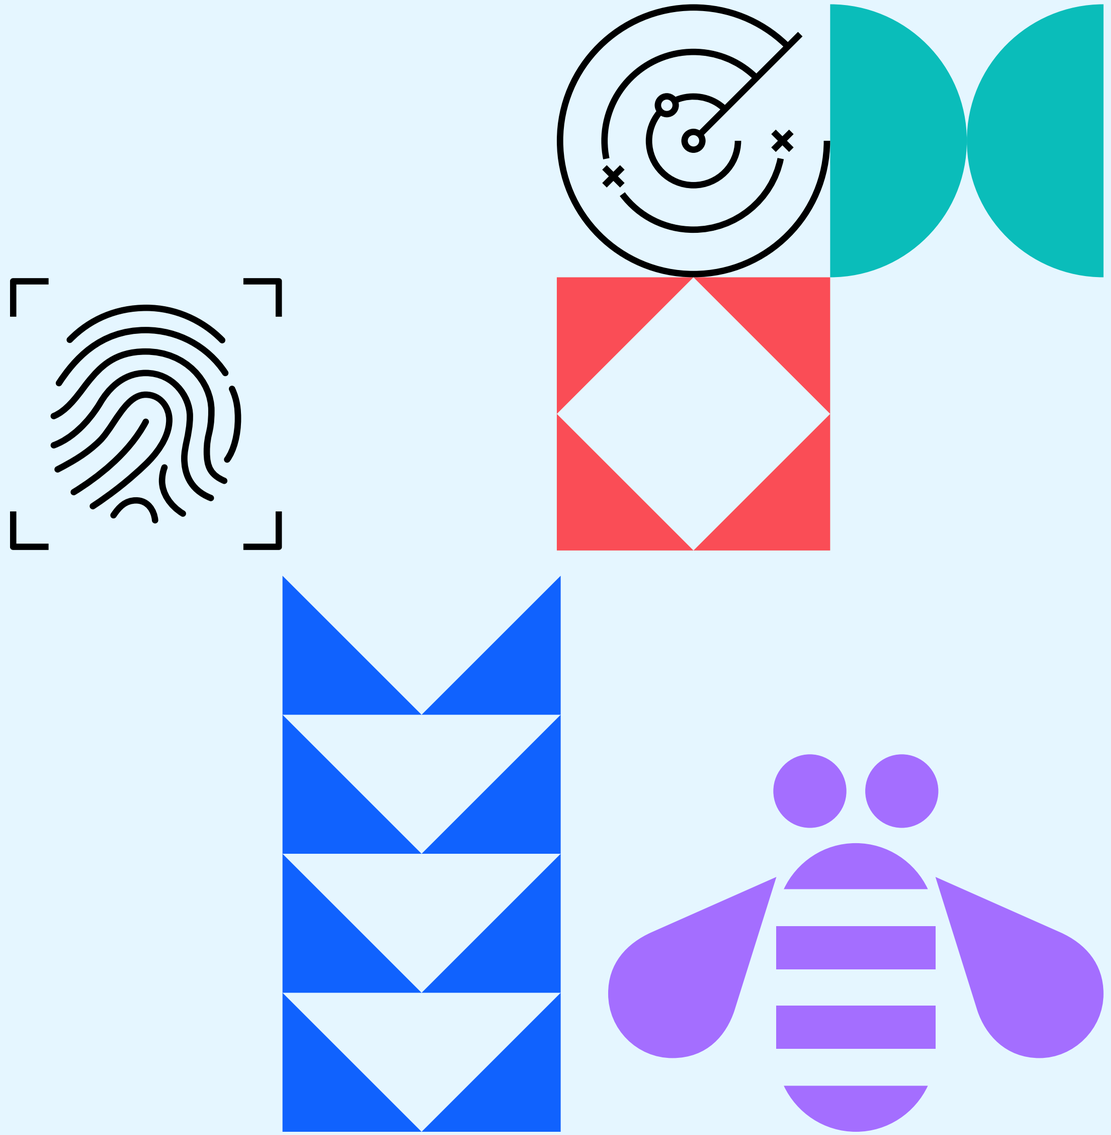

01 Introduction

02 zDIH Solution Overview

03 Digital Transformation at M&T

04 zDIH Use Cases

05 Getting Started

06 Q&A

Mainframe Barriers for Greater Agility and Integration

Real-time and event-based flows

How to share info at scale

Faster integration with cloud apps

Standards based & self-serve

Increased inquiry traffic & unpredictability

Optimization for query handling

Focus on digital and mobile channels

Ensure consistency of info

Digital transformation

Purpose-based modernization

Ecosystem expansion

Landscape for handling increased SOR interaction

Gaps:

- Stale information
- Spikey SOR impacts
- Limited eventing
- Inability to separate query & update traffic
- Cannot get needed info
- Ordered information (e.g. for transactions)

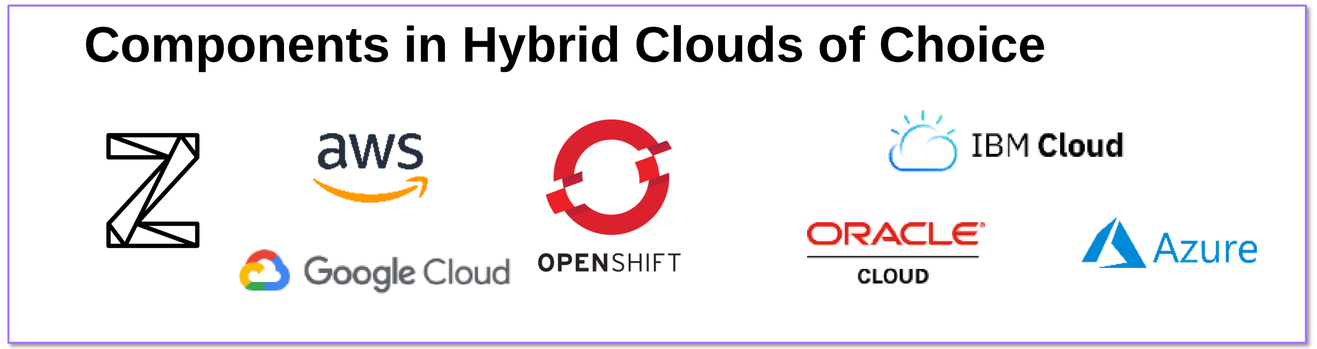

Access or move raw data

Direct APIs to individual core systems each time information is needed

z/OS

Core Transactional & Batch Applications

Raw Transaction

Core Systems Data

TXNS | ACH | Member | ... | Card

Existing ETL / ELT / CDC processes & tools

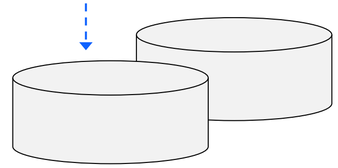

Warehouses, Data Lakes, Other Full Copies of Raw Data

Landscape for handling increased SOR interaction

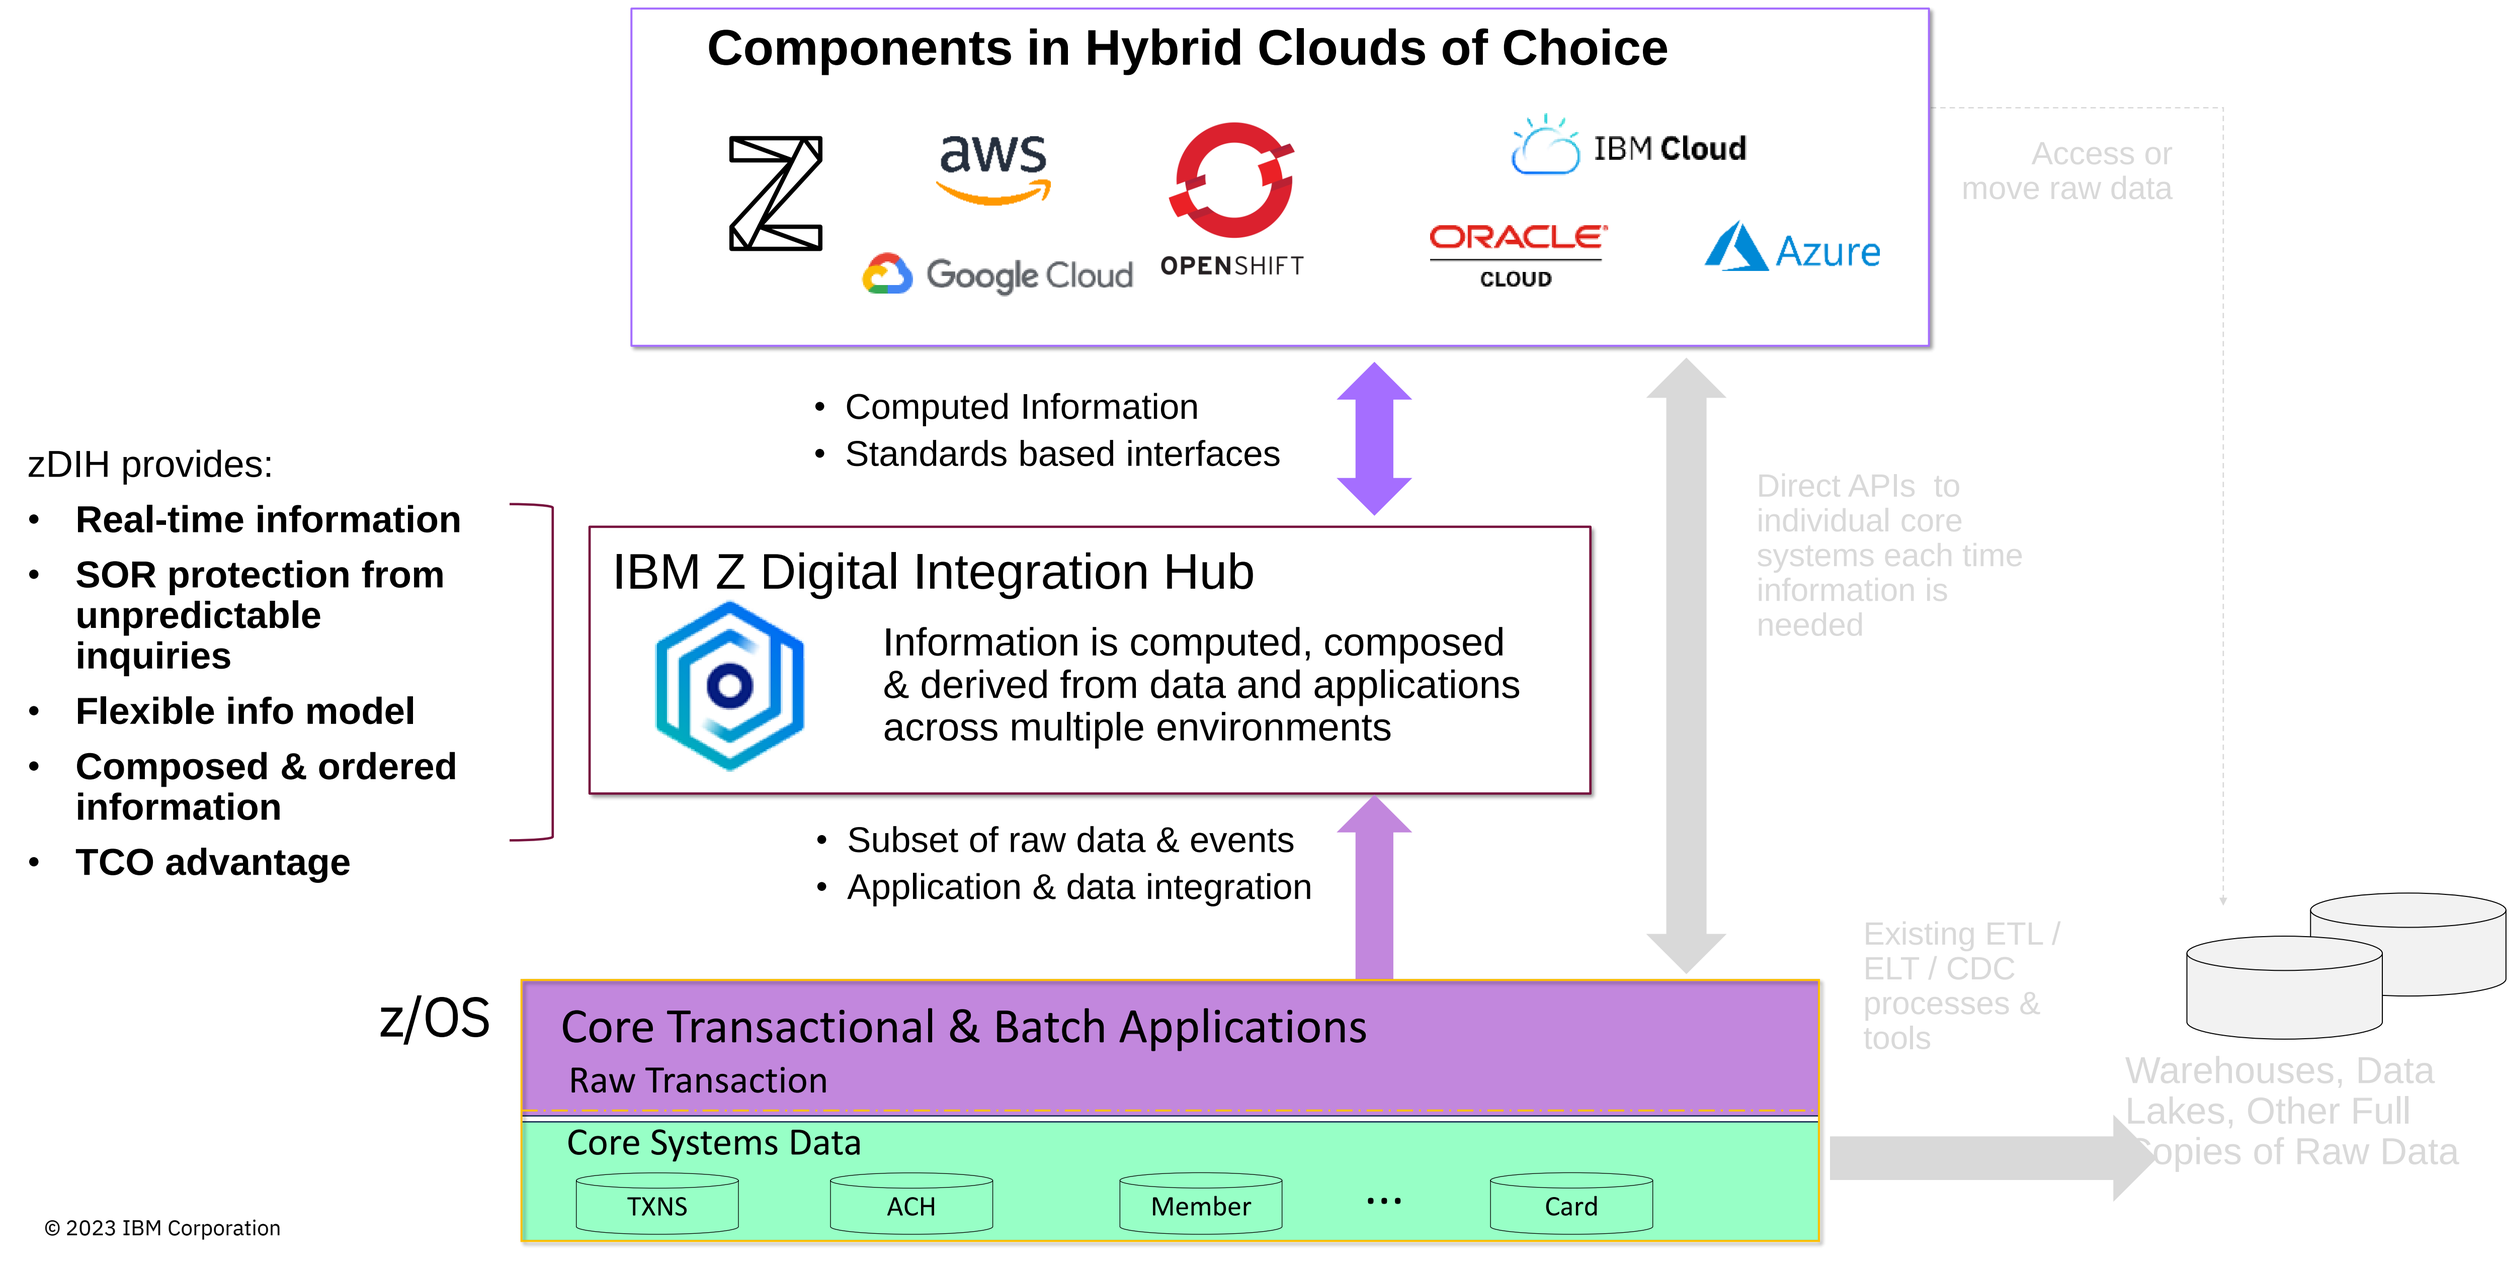

IBM Z Digital Integration Hub (zDIH) for Systems of Record

- **Real-time information flow** at scale between Systems of Record and hybrid cloud or end users
- **Faster development** of hybrid cloud applications due to decoupling with Systems of Record
- Accelerated **core systems integration** across the enterprise
- Incremental **application modernization** while avoiding disruption to core systems
- **Self-service** for business analysts without impacting core systems
- **Cost optimization** through separation of query processing from core transactions

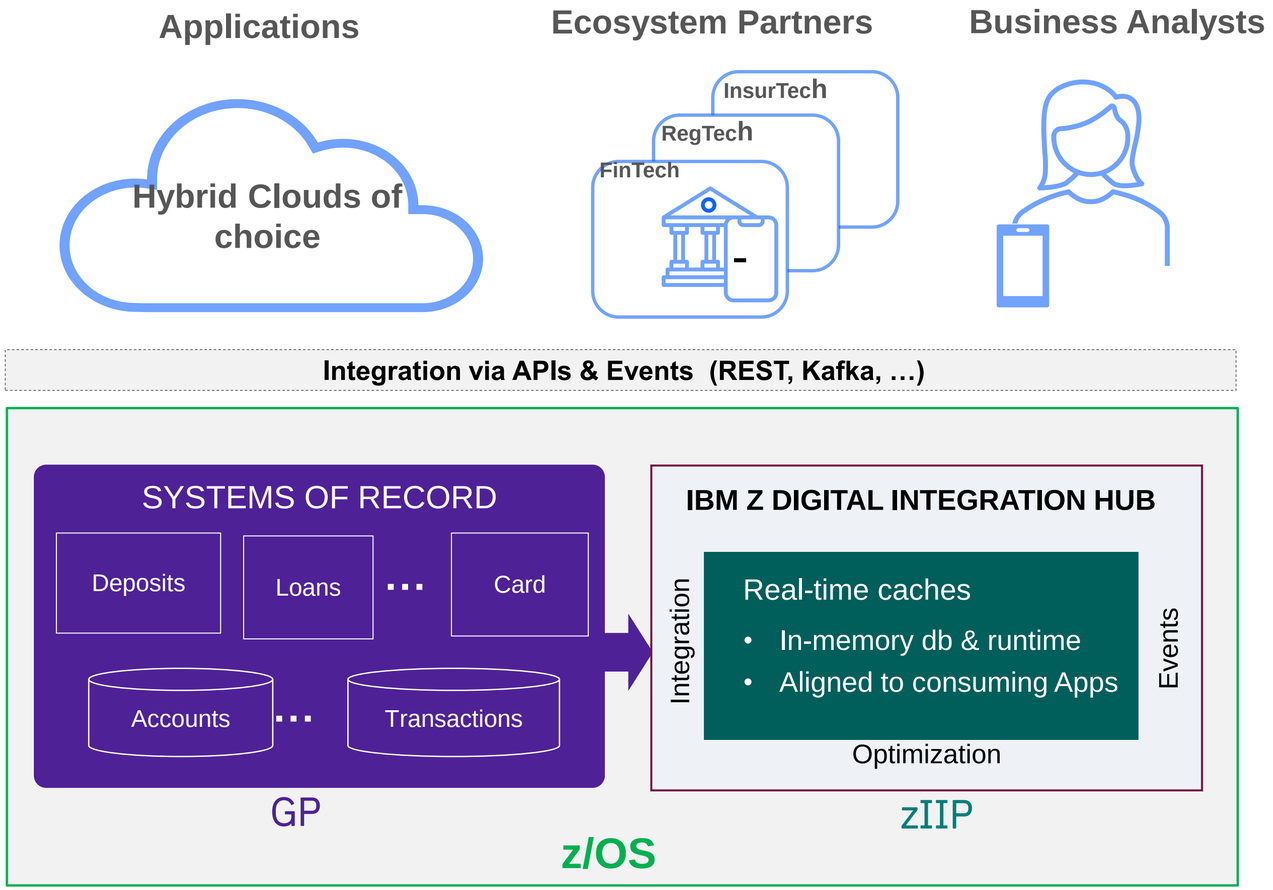

Gain faster ROI through hybrid cloud integration

Create new channels with expanded ecosystems

Leverage high-value investments

IBM zDIH technical overview

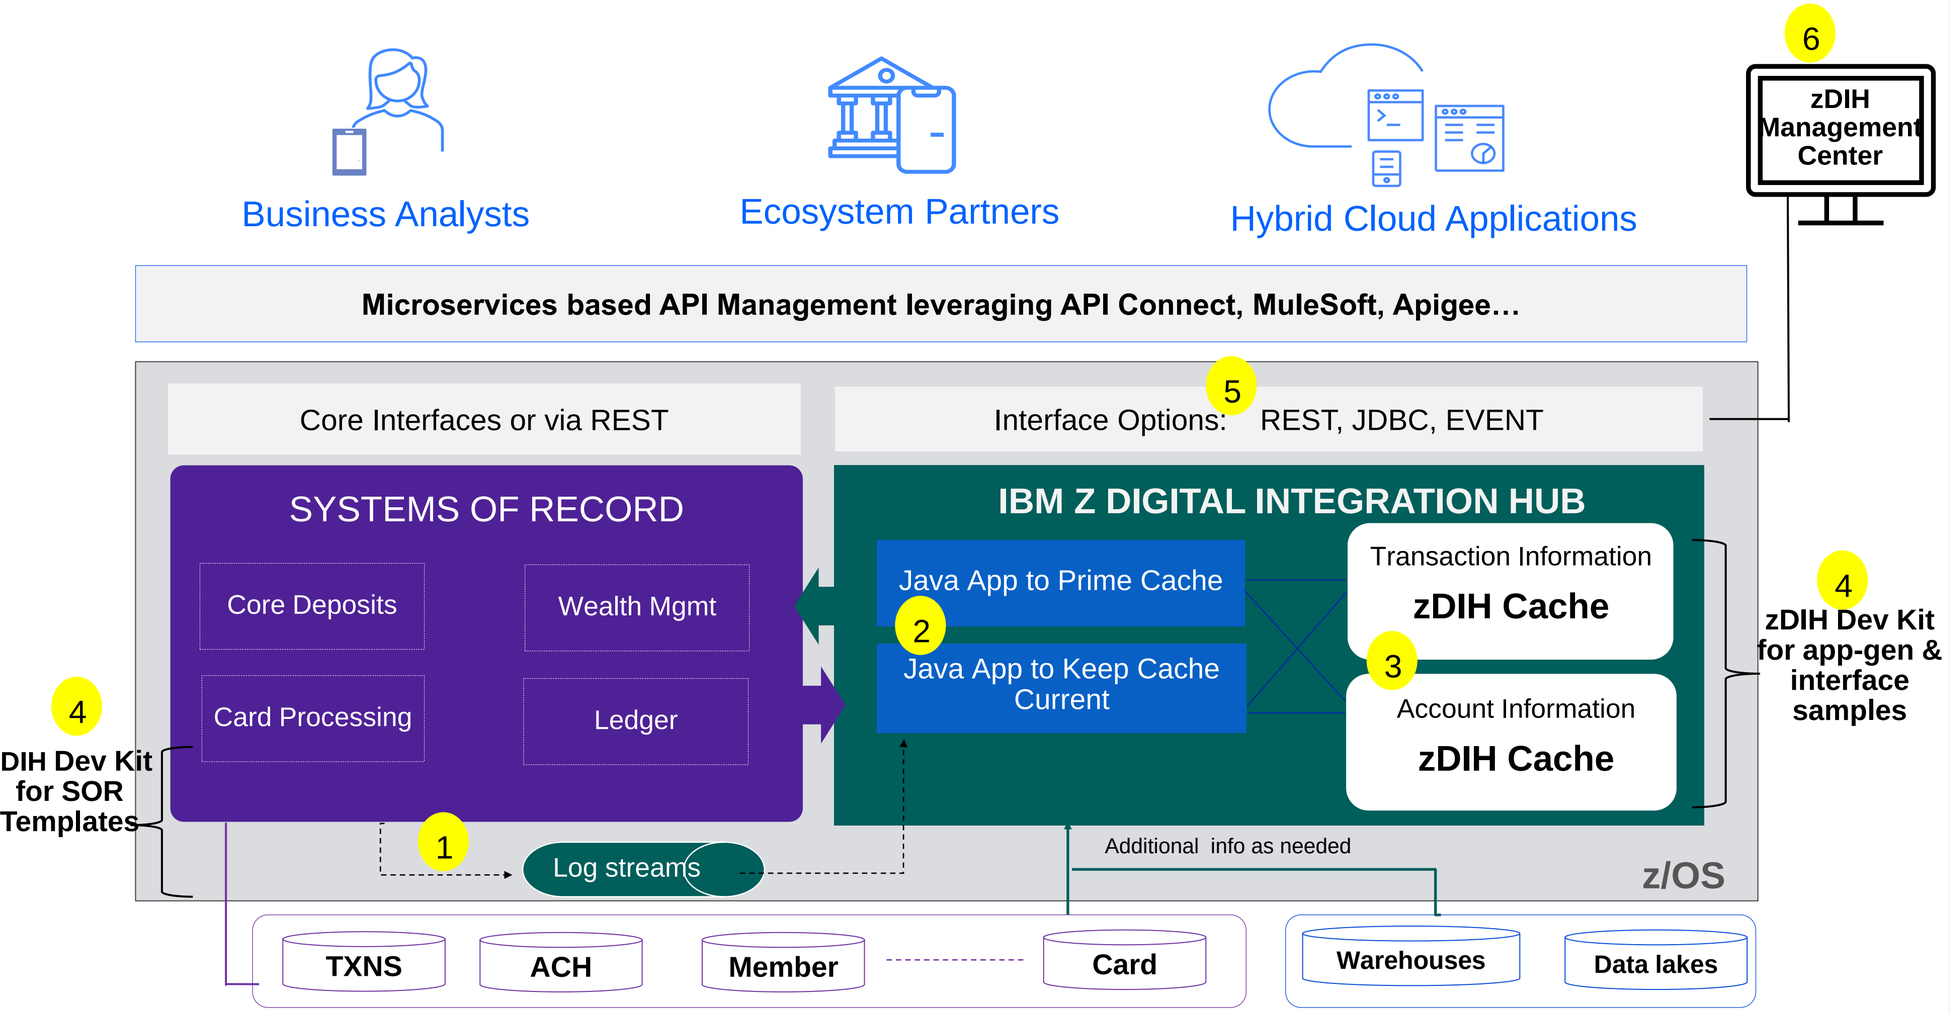

IBM zDIH Components

1. Efficient core systems integration
2. Java applications to leverage available skills
3. In-memory caches to accelerate processing
4. zDIH Developer Kit reduces code effort:
   - Auto App Generator
   - SOR integration templates
   - Interface samples (REST Kafka)
5. Standard interfaces
6. Management Center for monitoring zDIH

Various Integration Options:

- CICS App events + zDIH CICS Event Adapter + zDIH CICS TRUE
- zDIH CICS Exit & CICS TRUE
- Direct Write to zDIH log streams
- Other options

System of Record Integration using zDIH log streams

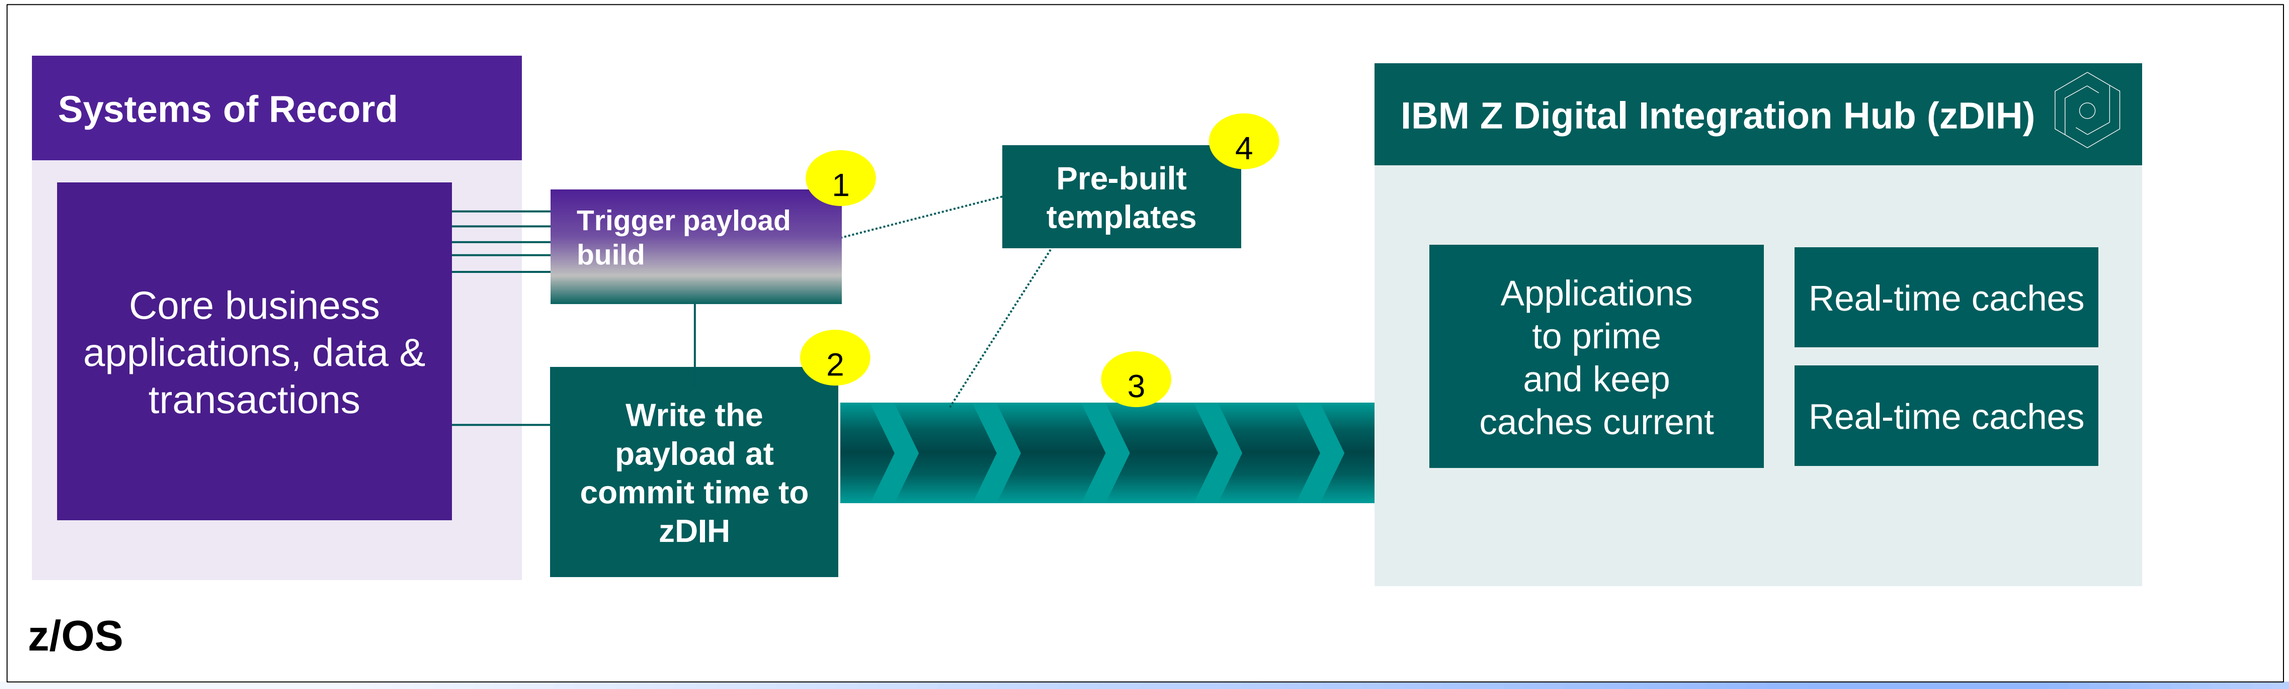

1. At <u>select</u> points from the application, build the complete payload to be shared with zDIH – use facilities such as CICS App events, CICS TS queues, CICS containers and/or zDIH exits
2. At syncpoint COMMIT of a logical unit of work (LUOW): leverage built-in capabilities such as CICS Task Related User Exit (TRUE) to write payload to zDIH -- for example, to zDIH log streams
3. The zDIH log streams are managed by the System Logger component of z/OS and offer high throughput, low latency, ordered communication
4. zDIH pre-built templates to accelerate system of record integration, build payloads and write to zDIH log streams with adaptability for IMS applications and batch.

CICS Application Integration with zDIH log streams

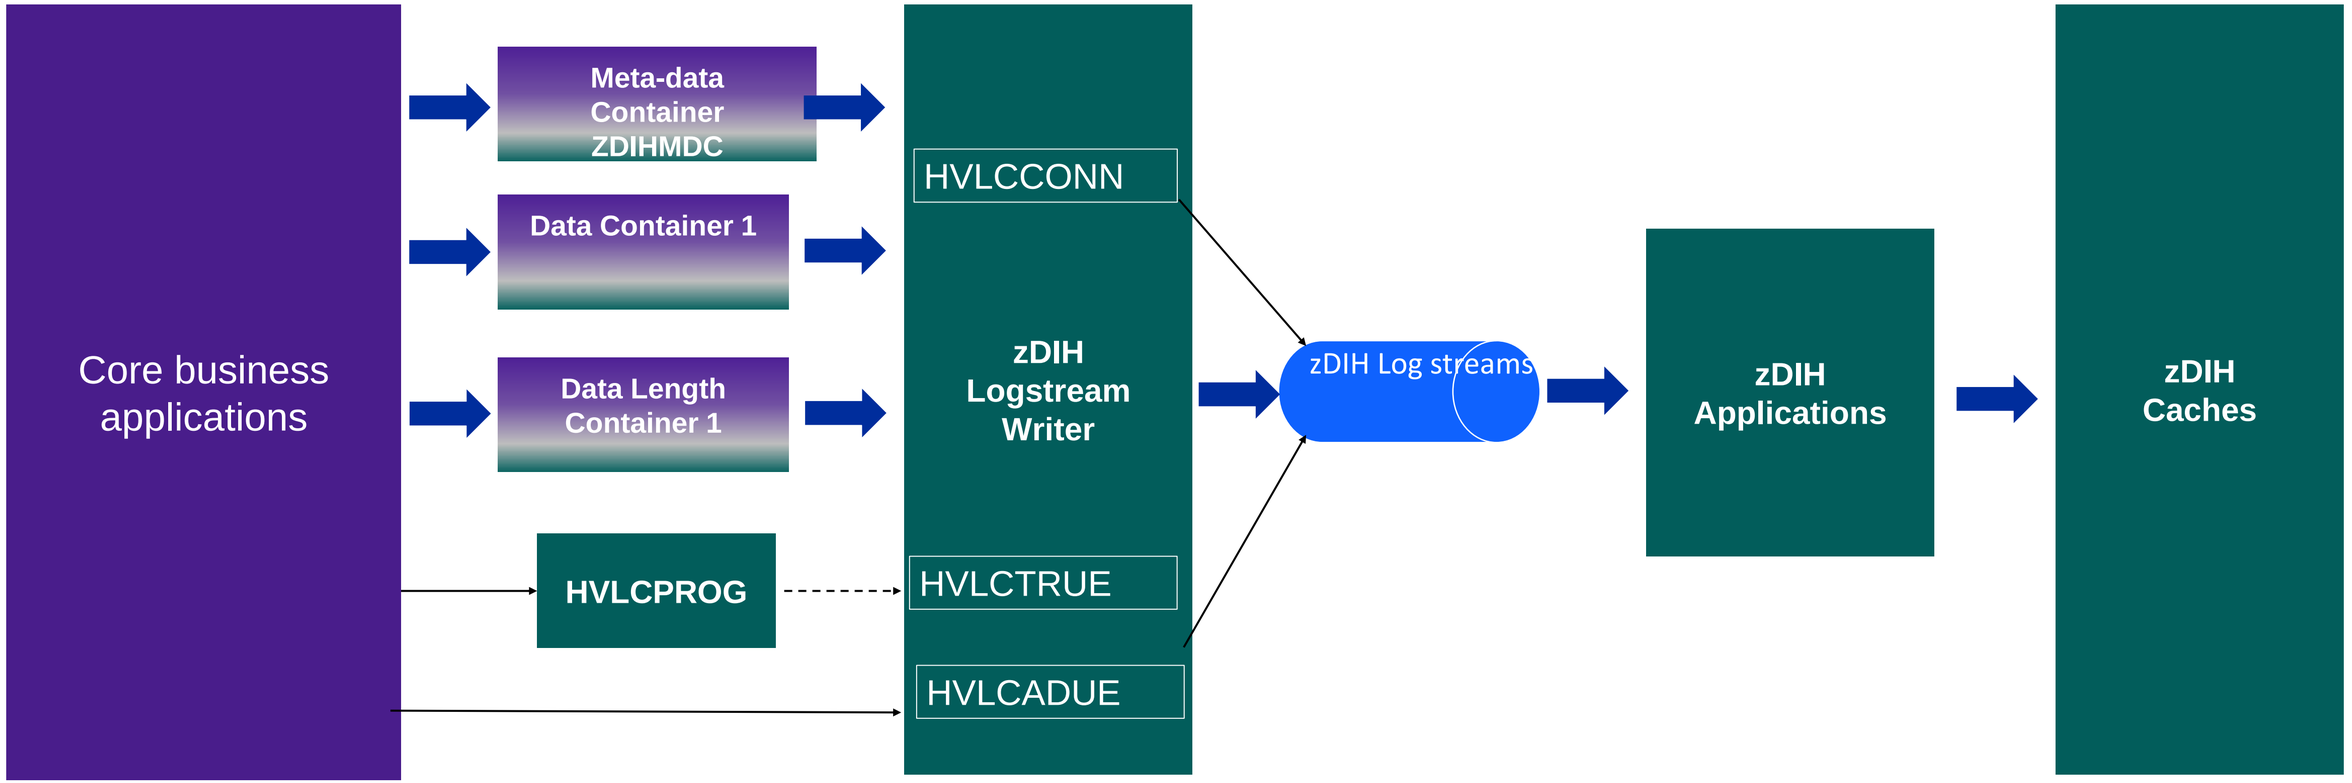

HVLCPROG - Registering HVLCTRUE to a CICS application
HVLCCONN - Connecting to a log stream
HVLCTRUE - Writing to a log stream from a CICS task-related user exit
HVLCADUE - Writing to a log stream directly from a CICS user exit

ZDIHMDC – Meta Data Container

- To share data container name, data length container and log stream name with zDIH log stream writers

Patterns of System of Record Integration with zDIH: CICS

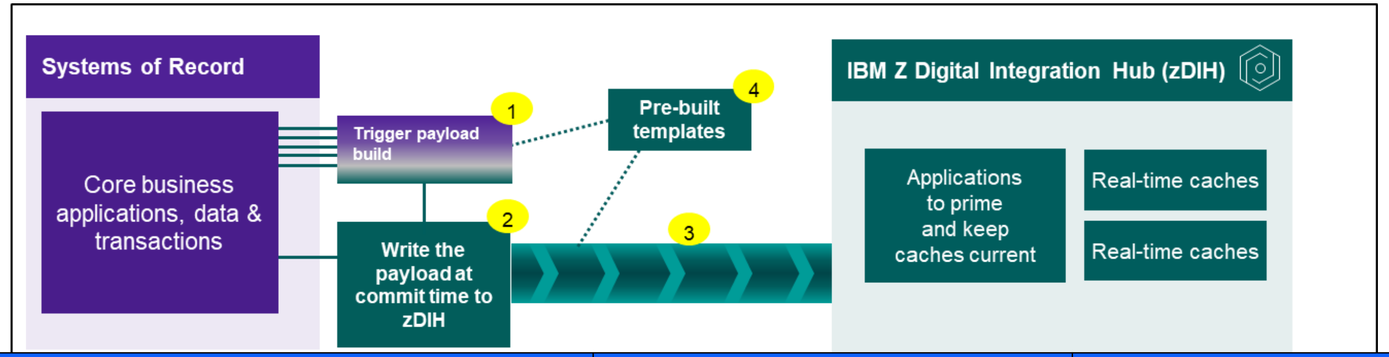

Options	Description	Use case fit	Characteristics
CICS App Events + zDIH templates for: • CICS Custom EP adaptor • zDIH CICS TRUE	• Leverage CICS Application Events capture points such as VSAM write/rewrite, PUT container, write to TSQ in order to build the zDIH payload • Leverage zDIH CICS TRUE to write payload to zDIH logstream	When the information needed for the zDIH caches is available at one or more of the predefined CICS application capture points	• No application code change • Ordering guaranteed • Multiple updates as a set guaranteed • Only committed information • Very low latency; low MIPS
zDIH templates for: • zDIH CICS App Exit • zDIH CICS TRUE	• Leverage a zDIH sample exit at custom points in the application to build the zDIH payload • Leverage zDIH CICS TRUE to write payload to logstream	When the information needed for zDIH caches is not available at any of the predefined CICS application capture points	• Minor application change to call the zDIH sample exit • Ordering guaranteed • Multiple updates as a set guaranteed • Only committed information • Very low latency; low MIPS
CICS to MQ to zDIH	• Leverage existing application writes to MQ for z/OS to communicate payload to zDIH	When the application already has defined points of integration with MQ that also align to the information needed in zDIH caches.	• No application code change • Ordering, multiple updates and committed information possible but depends on where and how the MQ Puts are initiated from the CICS application • Likely more latency & more MIPS

Low code approach

Create IBM Z Digital Integration Hub applications and caches in minutes with the IBM zDIH Developer Kit

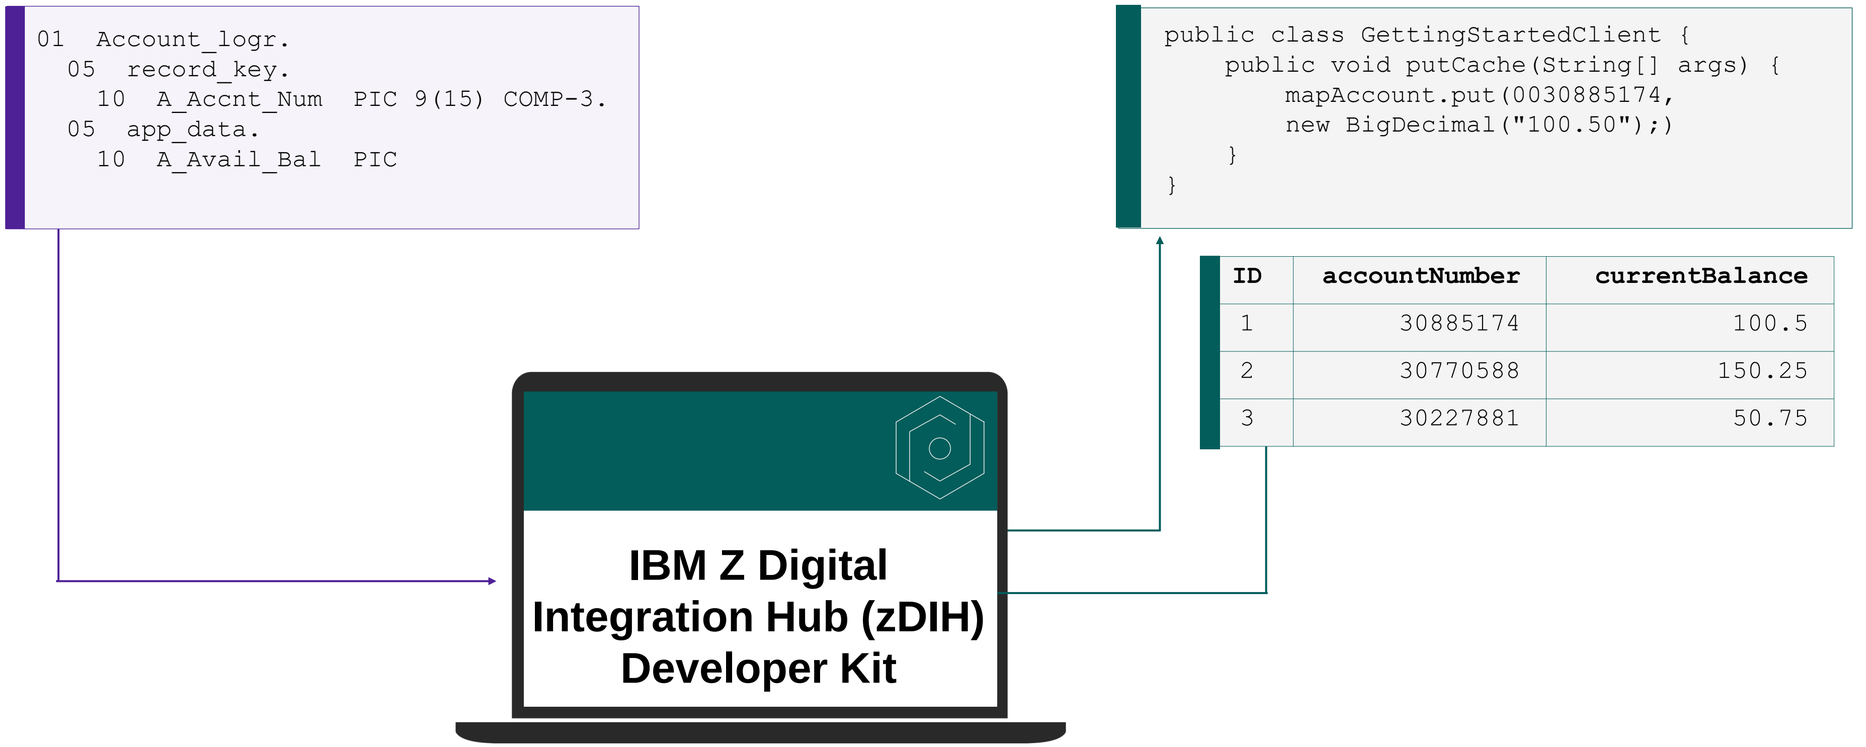

- Create zDIH applications and caches based on COBOL copybooks of the systems of record information to be shared with zDIH (e.g. zDIH log stream copybook formats)
- Robust customization parameters for flexibility and ease of use
- Resulting Maven Java project can be imported into IDE of choice for integration with DevOps pipelines

IBM zDIH application DevOps

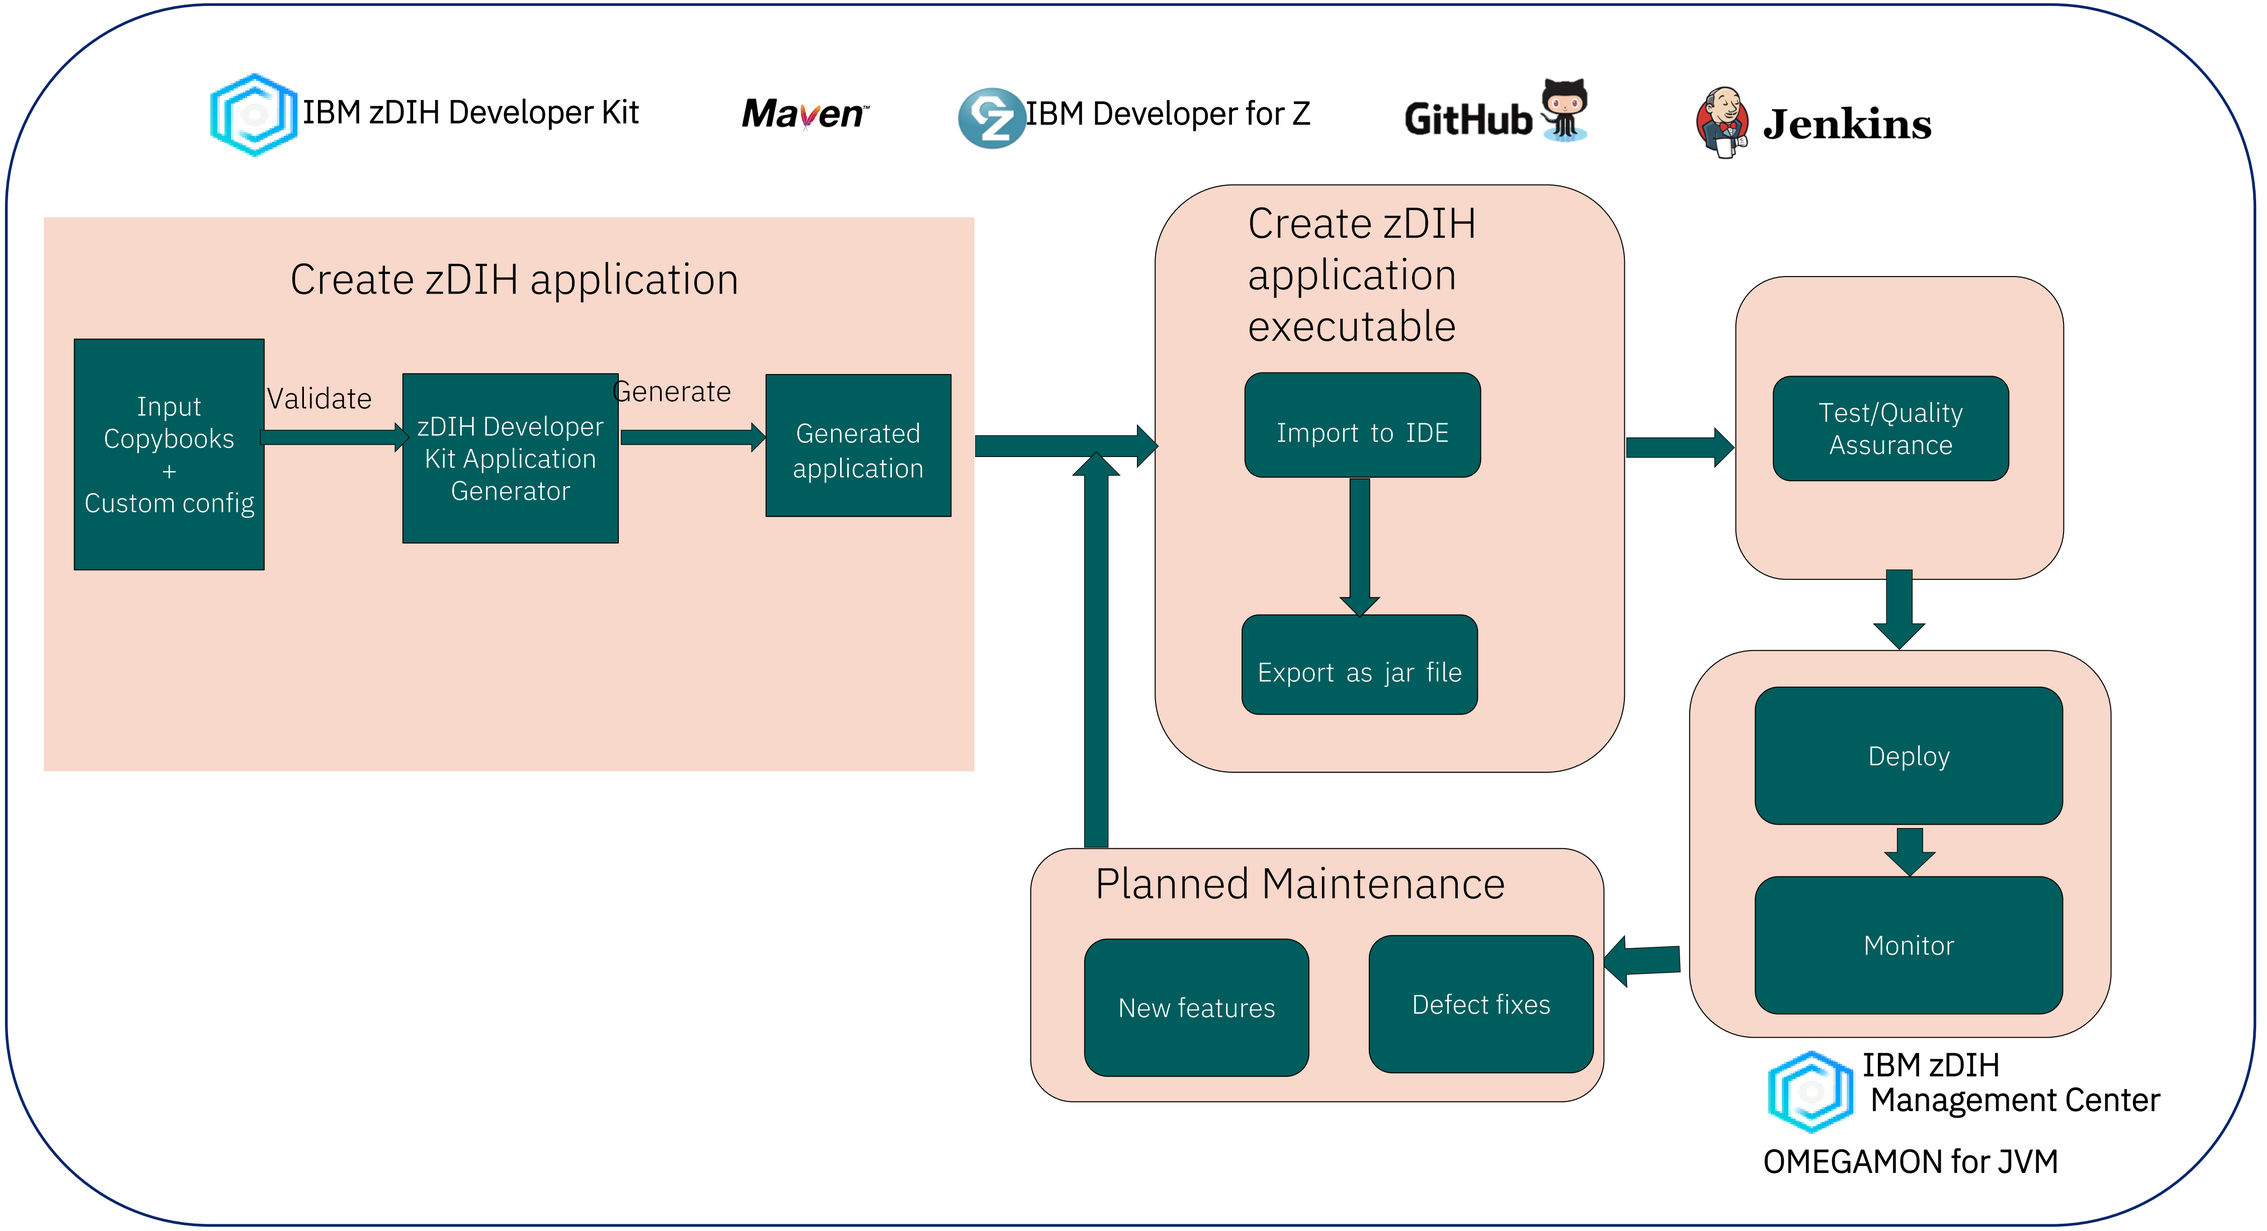

Monitoring

Monitor IBM Z Digital Integration Hub clusters, Java environment, and z/OS resources for optimal performance

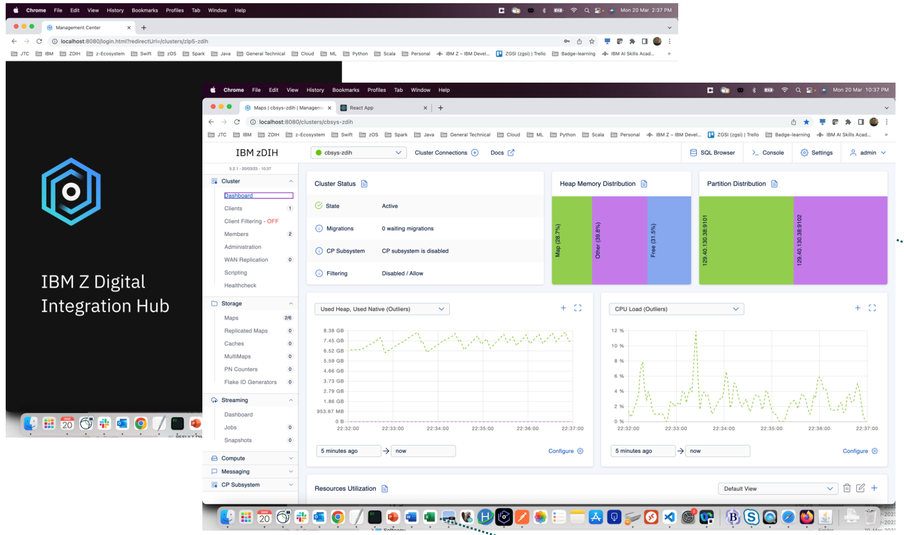

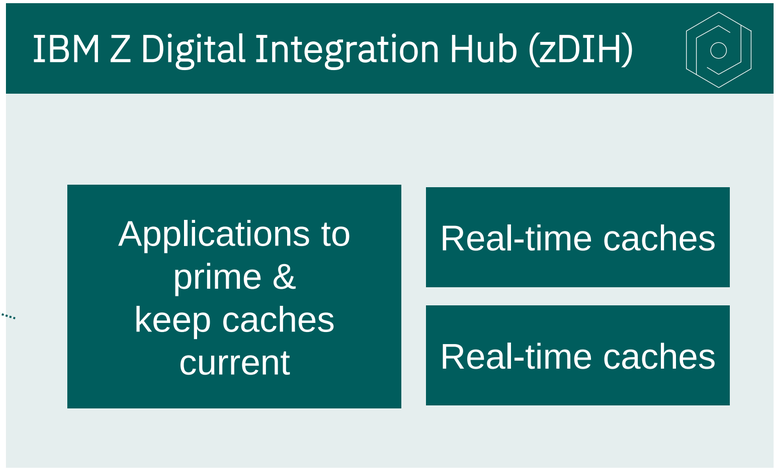

- IBM Z Digital Integration Hub Management Center shows cache entries, memory and heap utilization, node configurations and status
- Use tools such as Java Garbage Collection Memory Visualizer (GCMV), Java Health Center and IBM OMEGAMON for JVM on z/OS to monitor performance and function of the JVMs used by IBM zDIH
- Monitor z/OS resources used by IBM zDIH with standard IBM z/OS tooling (SMF records, RMF, SYSLOG, etc.)
- For more information see: IBM zDIH Product Documentation: Monitoring zDIH

IBM zDIH System Structure

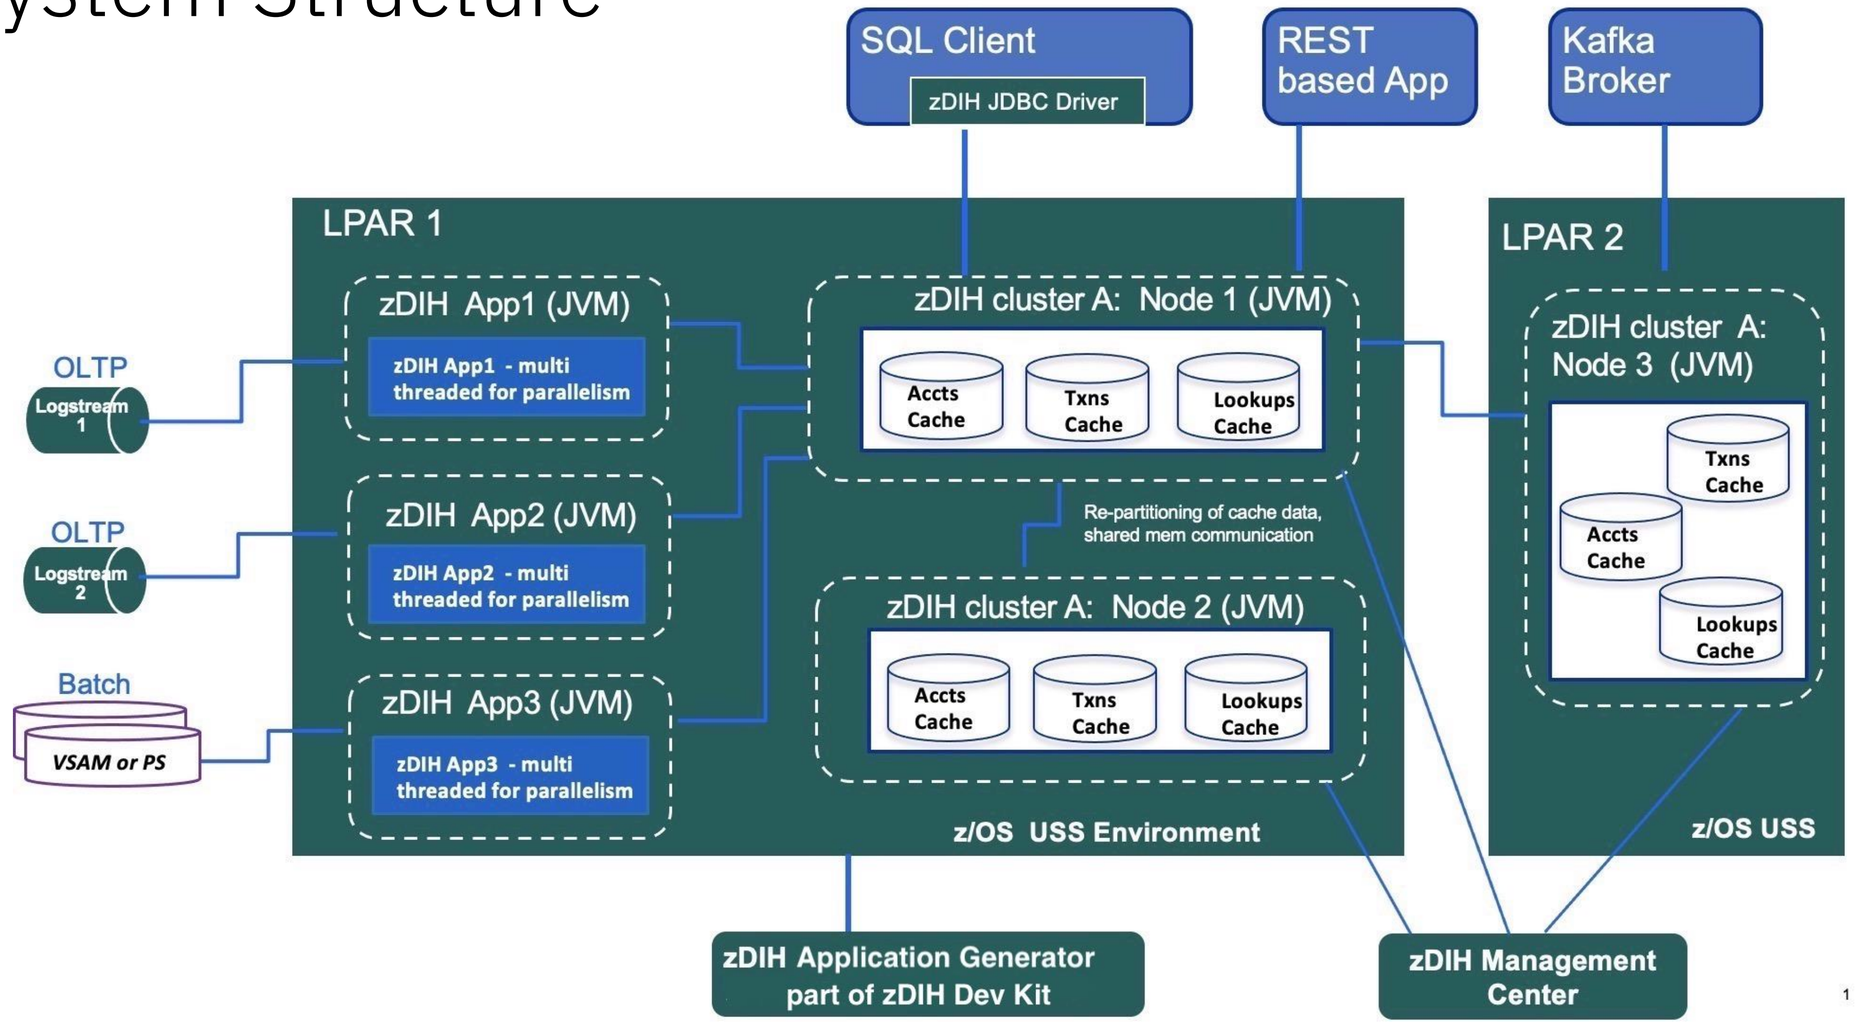

Key Points

1) The zDIH caches can run in a single node or multi-node configuration, each node is in a JVM under USS running with multiple threads
2) Each zDIH app runs in its own JVM & either reads from a given log stream and populates zDIH caches –or- reads from MVS datasets such as DB2, VSAM, PS or GDGs to read select data for priming zDIH caches
3) In a single LPAR, all communication over IP protocol is always optimized by z/OS to use shared memory to avoid network latency
4) In multiple LPAR approach, SMC-D can be configured if on same CEC or SMC-R if between different CECs to optimize any network delays; SMC-D and SMC-R provide significant performance advantages

IBM zDIH HA/DR Example

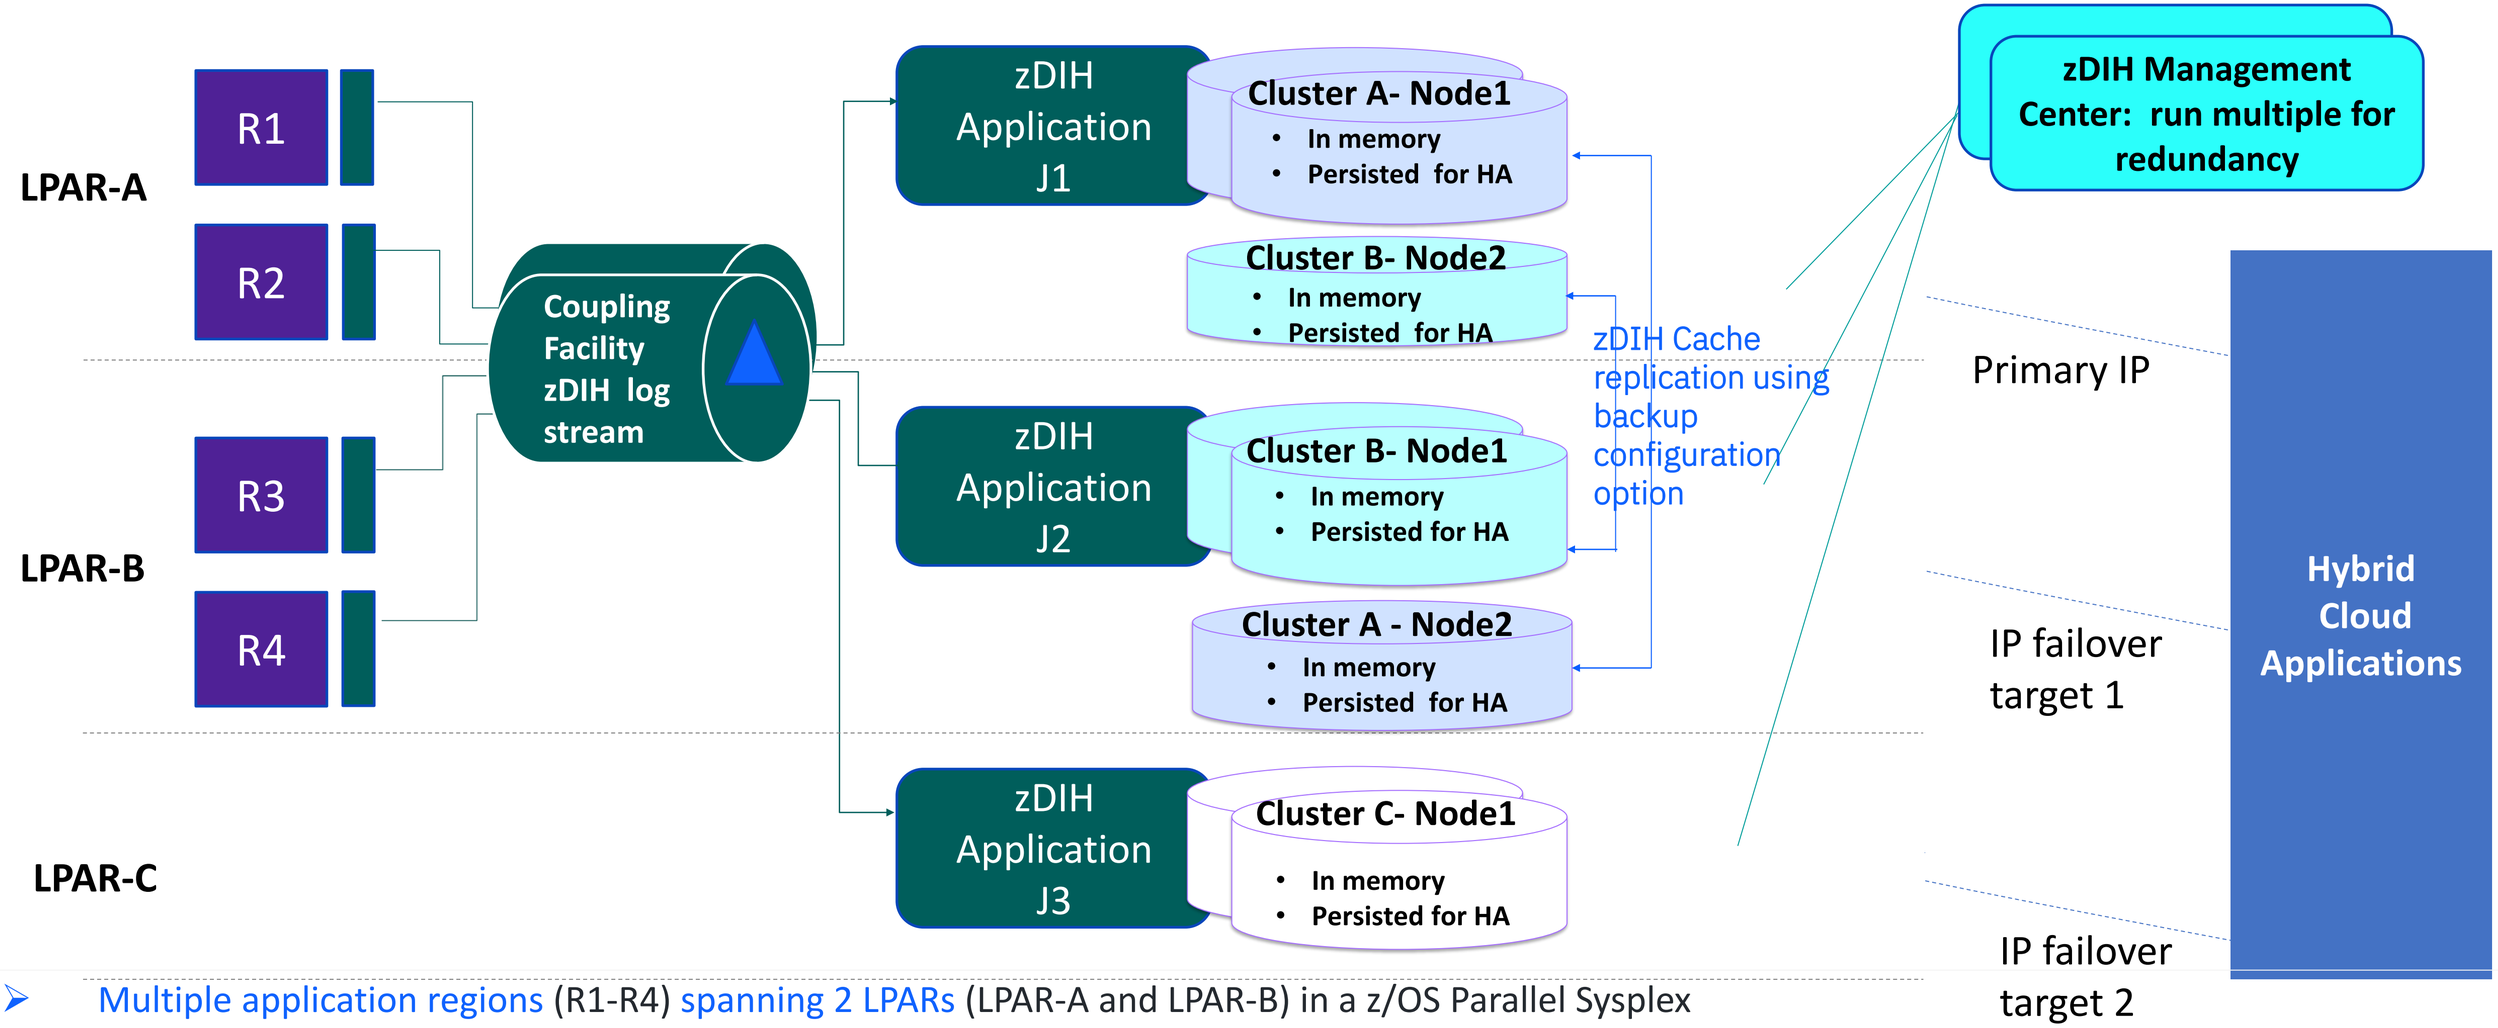

- Multiple application regions (R1-R4) spanning 2 LPARs (LPAR-A and LPAR-B) in a z/OS Parallel Sysplex
- IBM zDIH log streams defined in the Coupling Facility that are also defined with duplexing mode
- Three concurrent zDIH clusters across the 3 LPARs in a Parallel Sysplex (Cluster A, Cluster B and Cluster C)
- Each zDIH cluster reads from the same log stream and keeps its zDIH caches current; zDIH Clusters A, B, and C also use persistence.
- zDIH Clusters A and B also use a multi-node approach for zDIH clusters, replicating their partitions between zDIH nodes that span LPAR-A and LPAR-B
- Dynamic VIPA (DVIPA) with IP failover is used for automatic failover between zDIH clusters for hybrid cloud applications that are querying the zDIH
- For more information, see: IBM zDIH Product Documentation: HA/DR Considerations

M&T Bank

How do banks make accurate business decisions in real-time?

M&T Bank was searching for a faster, more efficient way to share core banking information with hybrid cloud applications—without impacting production systems.

M&T collaborated with IBM on a Z Digital Integration Hub (zDIH) engagement to modernize and better integrate their z/OS® applications with downstream consumers.

With zDIH, the bank mitigates risk and potential fraud by providing hybrid cloud applications with sub-second current information from systems of record. M&T also improves time to value up to 40% for data-driven applications and enables business analysts to respond customer issues.

Agile integration and real-time information flow at scale between Systems of Record and hybrid cloud applications

“Without this pilot and collaborative approach, our collective organizations would not have achieved the great outcomes we did.”

Russell Plew

Technology Senior Manager, M&T Bank

Solution Components

- IBM z15
- IBM Z Digital Integration Hub

Read the full story

Example banking customer challenges

Business challenge:

- Limited event processing capabilities
- Raw data not readily consumable
- Moving all raw data not viable due to transaction volumes
- Applications & business analysts need real-time information
- Increased query interaction interferes with OLTP performance
- Time to value delays for new apps due to specific data extracts

Solution objectives:

- Event Centric interaction without disruption
- Flexible data delivery & presentation for multiple consumers
- Performant & cost-effective solution with low latency

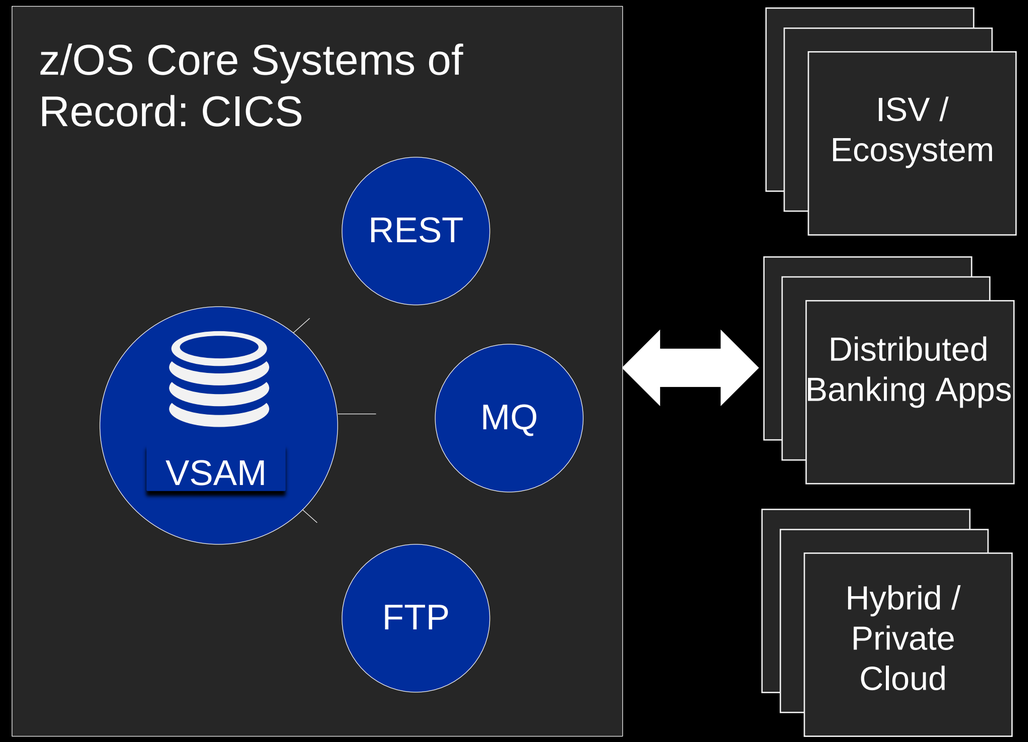

IBM Z Digital Integration Hub for modernizing core information flow

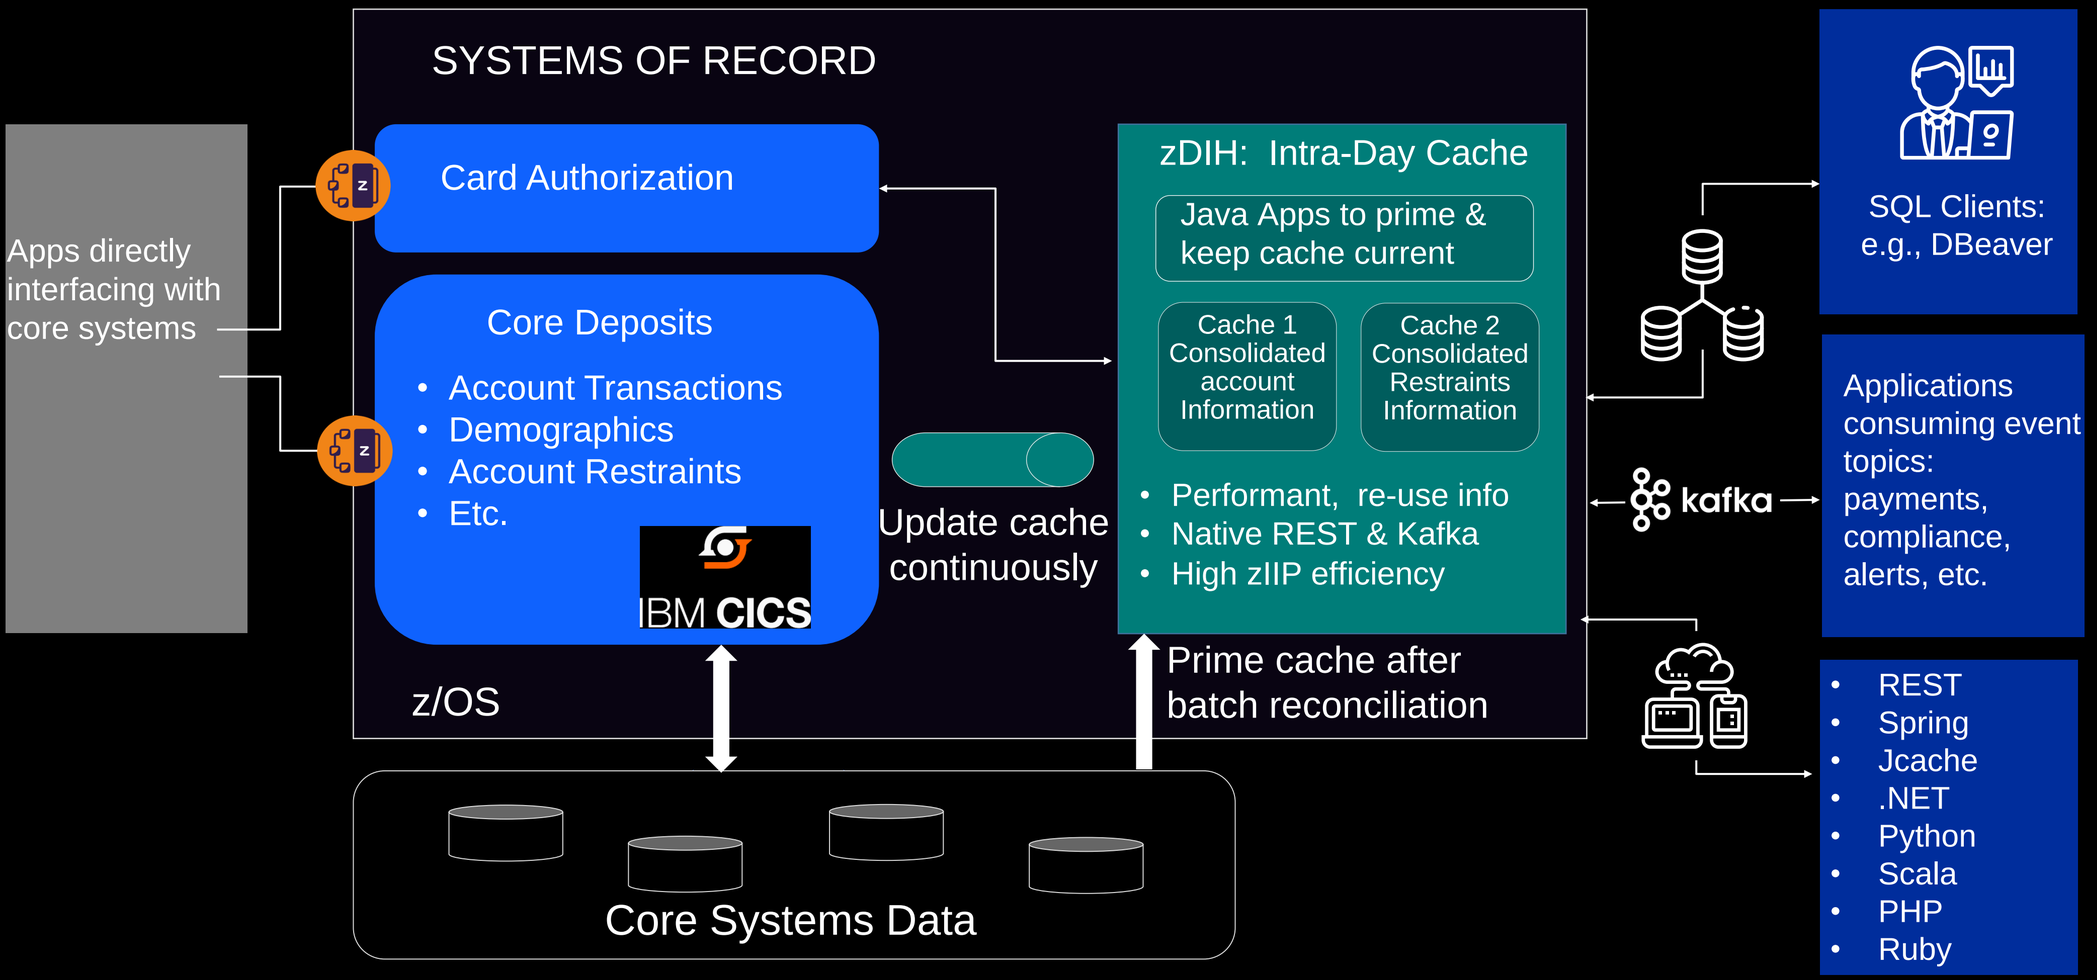

Business analysts can **self-serve** inquiries and detect anomalies without impact to core systems leading to ***better agility***

Apps get **real-time consumable information** leading to more ***optimal business decisioning***

Numerous, **standards interfaces** to intra-day cache information leading to ***faster app development***

Example customer results with zDIH

Previously	With zDIH
z/OS Applications – Currency: 3+hr old – Increased compliance risk window	z/OS Applications + Currency: sub-second + Significant mitigation of risk
Hybrid Cloud Development – Specific data extracts, leads to elongated new app development cycles – No self-serve for BAs, leads to longer time to service customer issues	Hybrid Cloud Development + 40% faster time-to-value for data-driven hybrid cloud apps + Full self-serve for BAs, reduces customer issue servicing time
Price performance – Minimal use of specialty cores (cost disadvantage)	Price performance + 95+% use of specialty cores (cost advantage)

“With Z DIH, M&T is able to reimagine the ways our core banking platforms can share information and events with our consumers—whether for data analysis or application consumption—and become a full-fledged member of our digital ecosystem transformation.”

“Self-service by our users and hybrid cloud applications to intelligible/consumable data and events through modern methods (JDBC, ODBC, Kafka, REST) with little to no modification to our core applications is a paramount need.”

Russell Plew
Technology Senior Manager
M&T Bank

Optimized Command-Query Responsibility Segregation (CQRS)

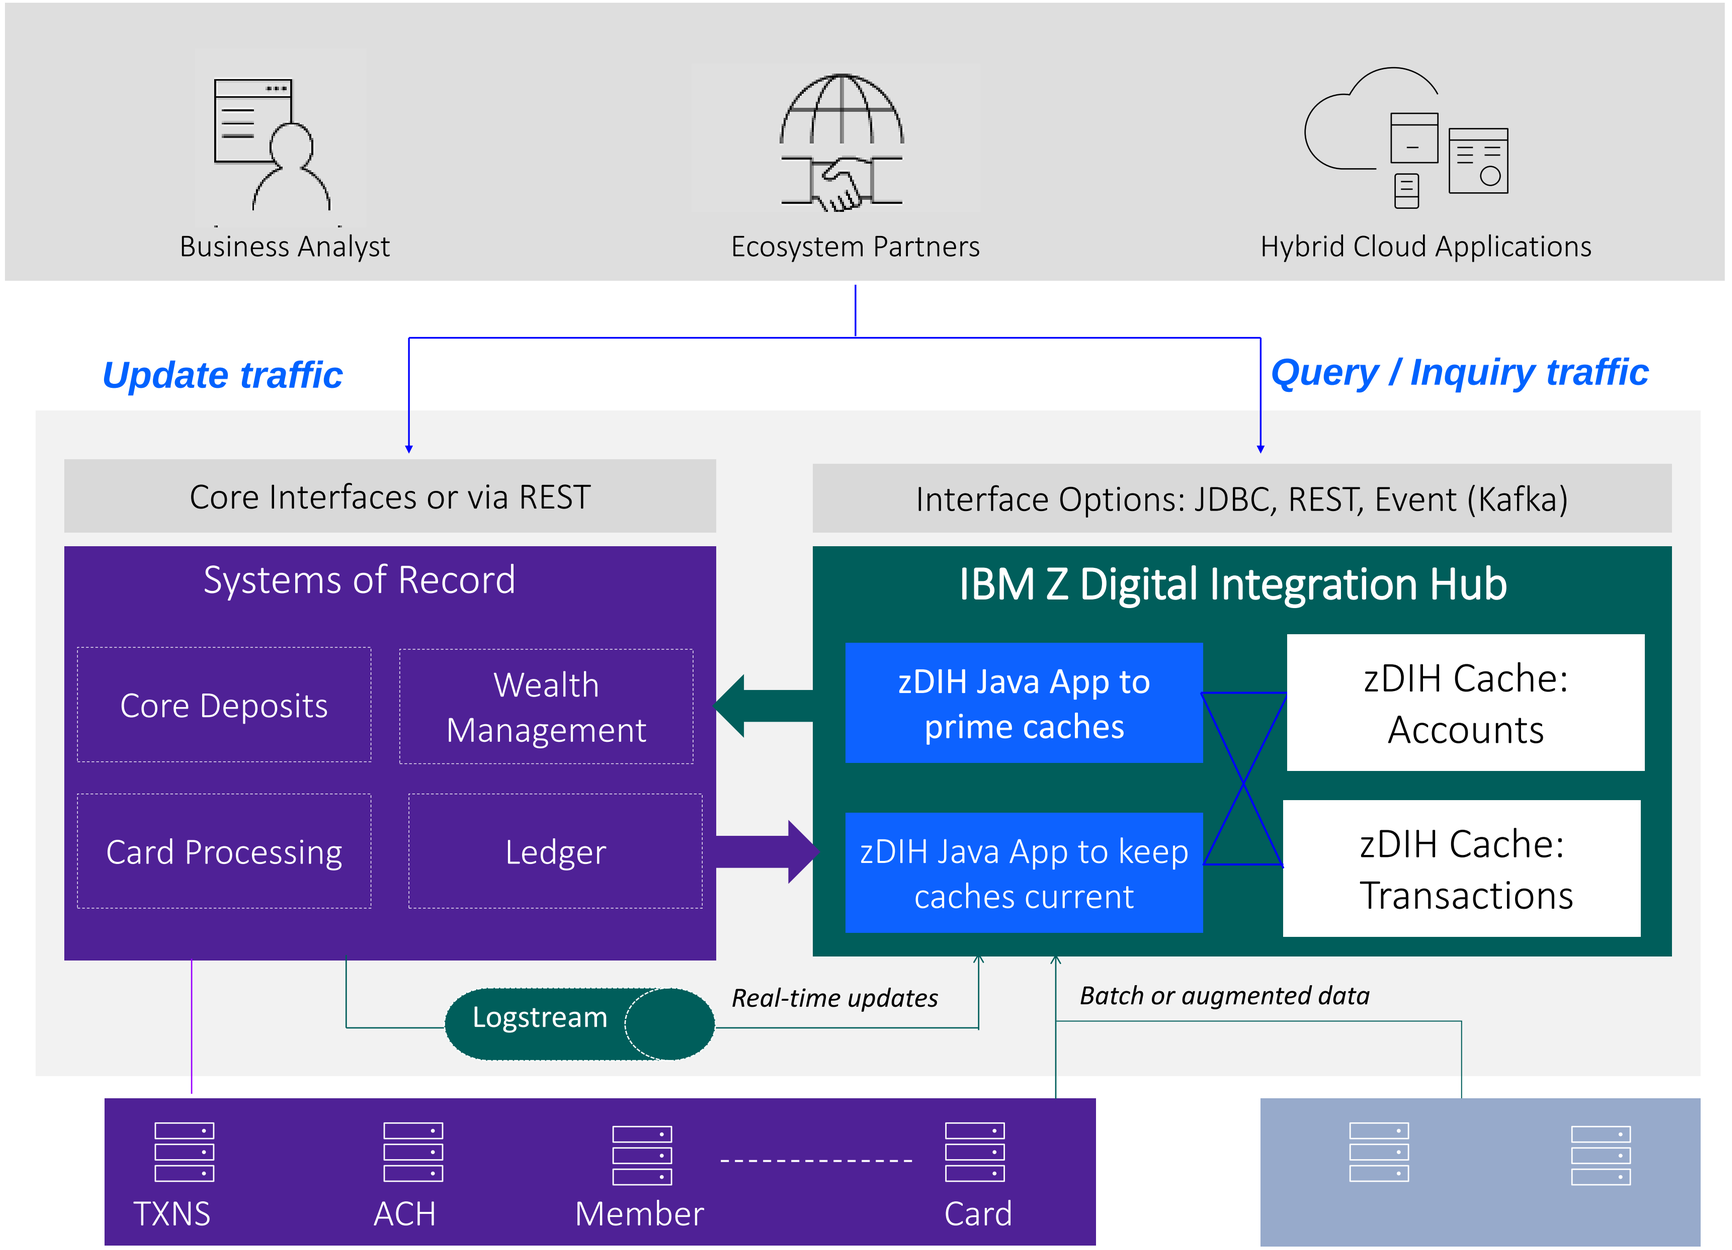

Use case examples:

- Optimized inquiry of intraday balances for current accounts and savings
- Real-time information about wholesale or retail payment entities
- Credit/debit card activity

Benefits:

- SoR not impacted
- Real-time at scale
- Events: proactive updates
- Consumable information
- Standards based interaction
- Composed info (e.g. balances)
- TCO advantage
- Selectivity about information shared

Bank Servicing after CQRS

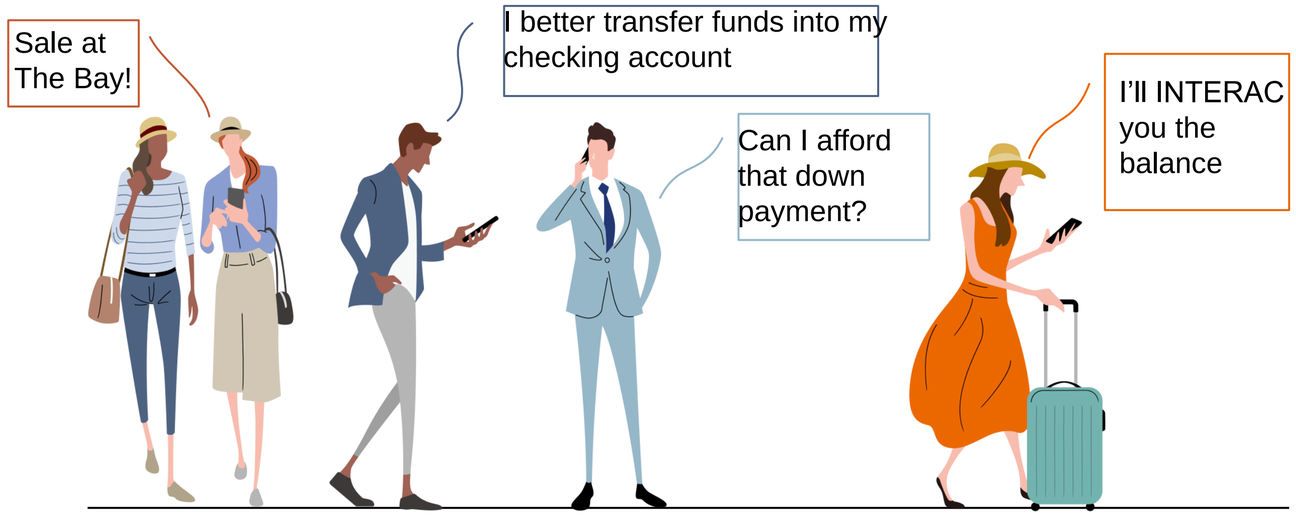

Command Traffic <30% (e.g., move $ from one account to another)

Inquiry Traffic >70% (e.g., what is my account balance?)

Systems of Record

Core business applications
- Deposits
- Loans
- Credit card
- Customer information
- Payments

IBM Z Digital Integration Hub

Applications to prime & keep caches current

Real-time caches

Real-time caches

Checking Accounts

Daily Transactions

Savings Accounts

Credit Card Accounts

Loan Accounts

z/OS

- Real-time information shared across the enterprise
- Cost optimization through zIIP and separation of Command and Inquiry Traffic (CQRS) reduces processing load on SoR
- Increased developer agility resulting in greater TTV
- High throughput with low latency for maximum scalability
- Flexible Information model
- Support for Open Standard Interfaces

IBM zDIH use case: share intra-day running balances

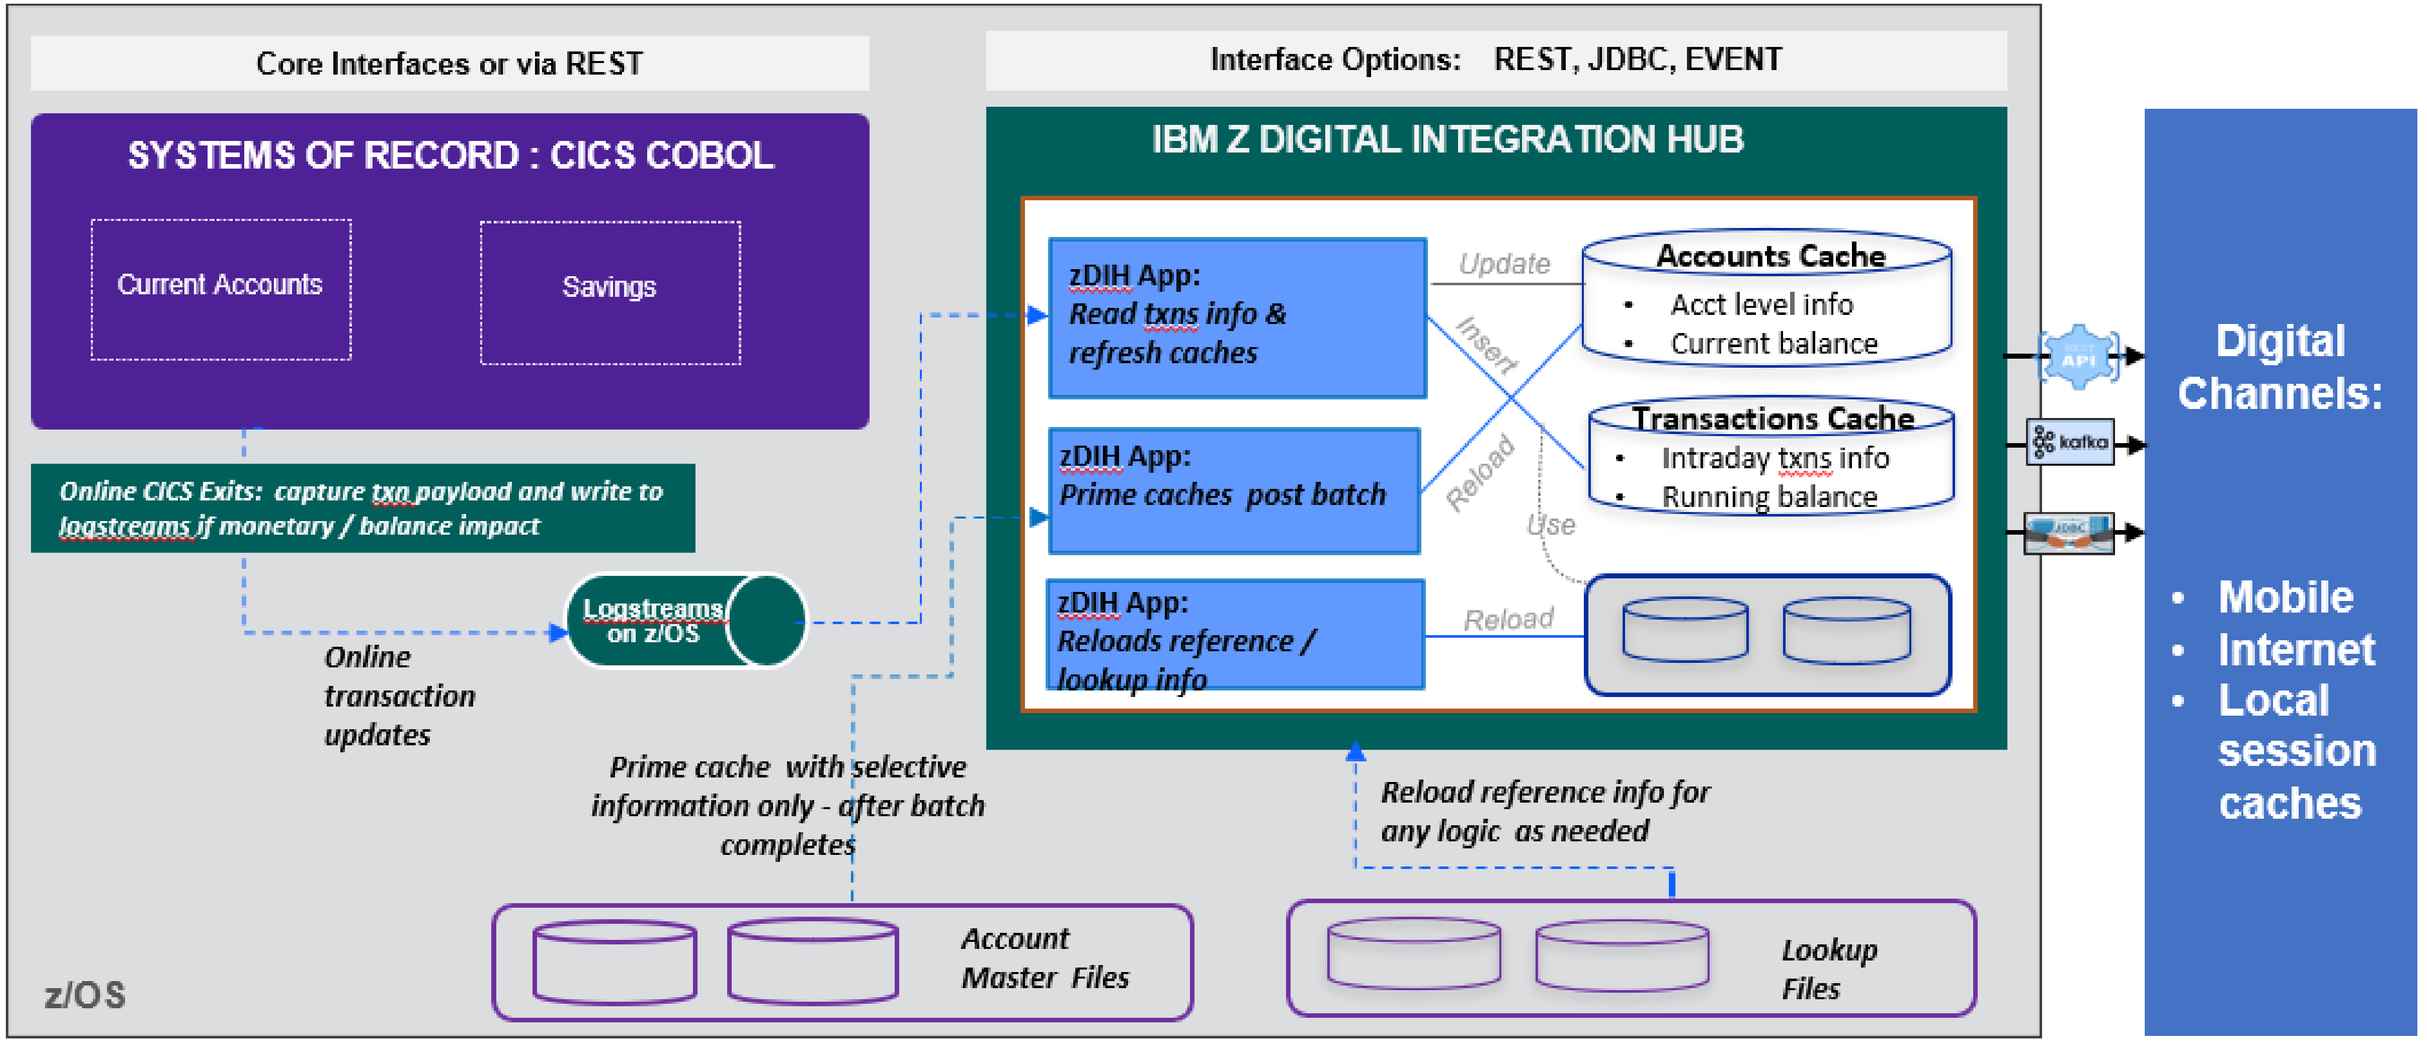

- The exits identify when composed information such as real-time balances should be re-computed and added to the payload
- Avoids duplication of business logic since existing computations derive the composed information (e.g. compute balances)
- Use this approach for information that is frequently queried or highly valuable

IBM zDIH use case: reduce need for data copies

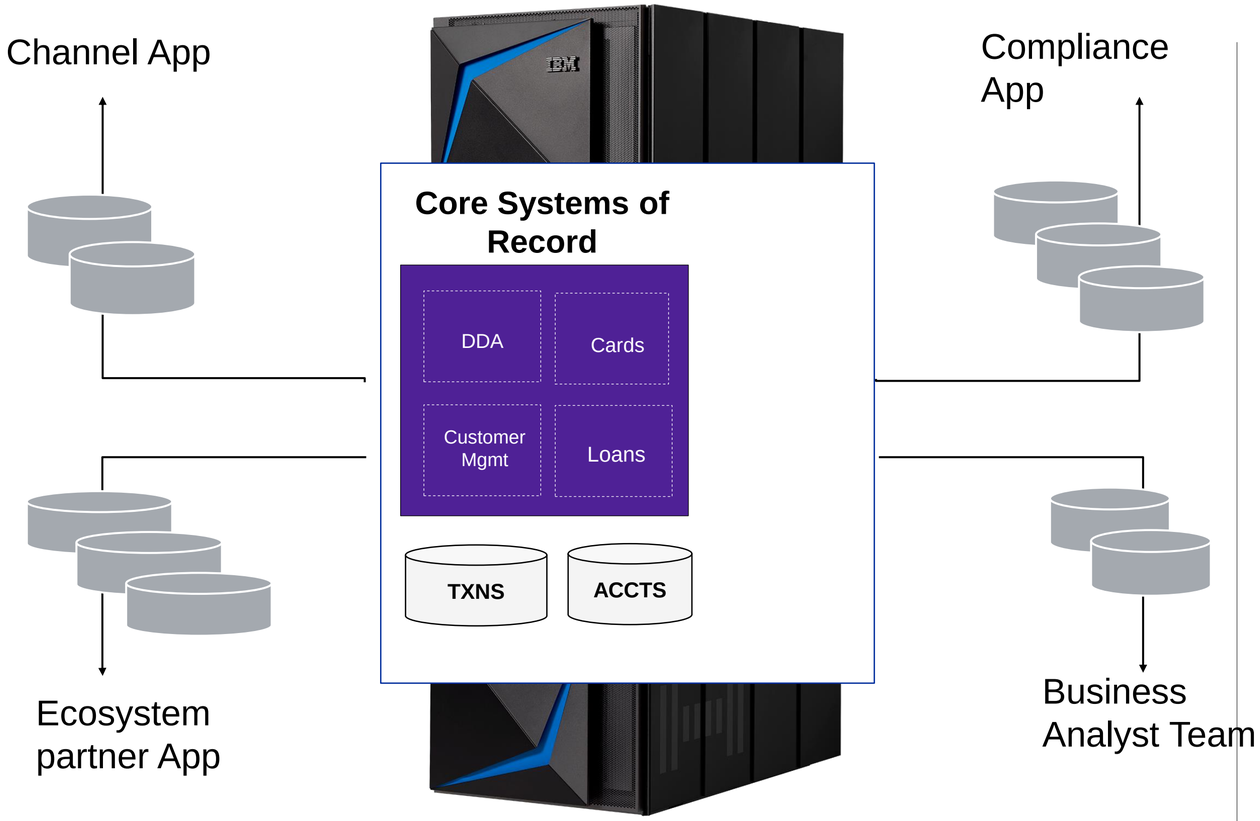

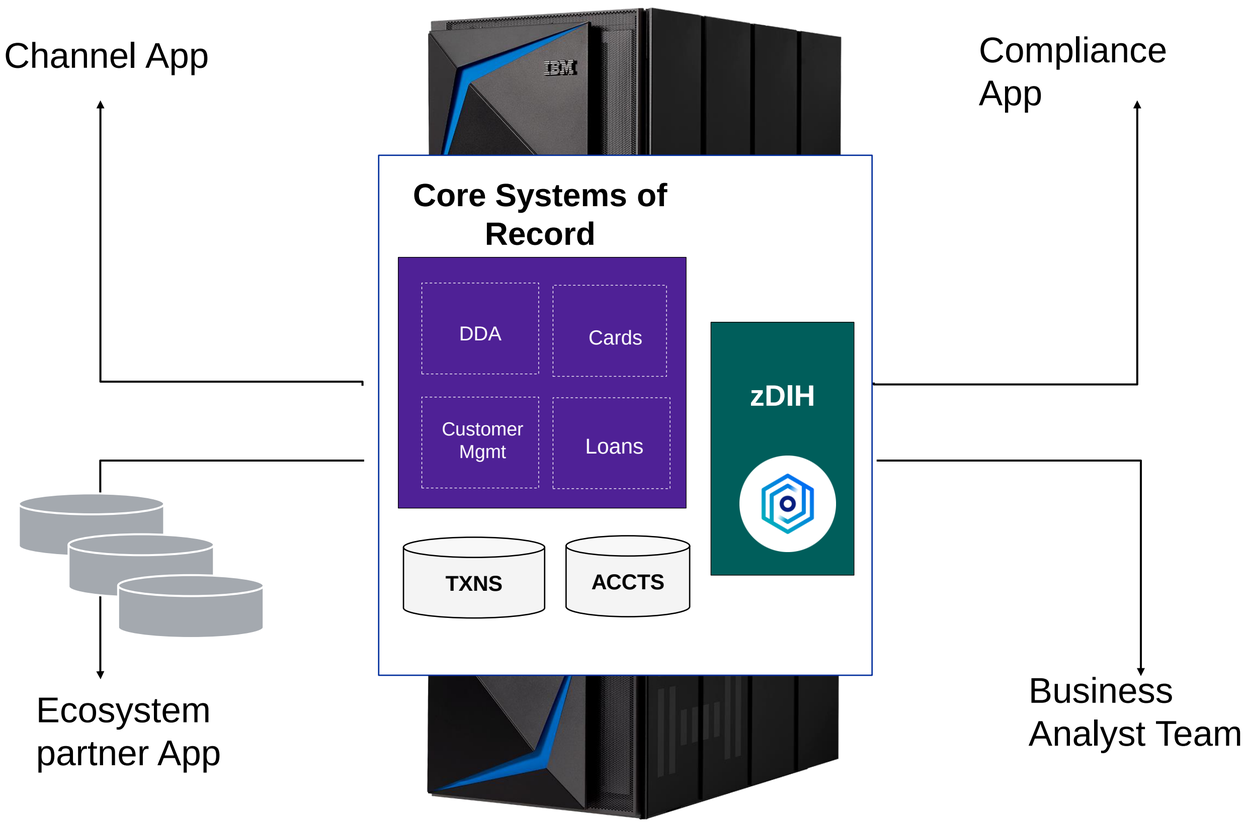

Today's Landscape: typically Extract Transform & Load

- Proliferation of custom data extracts – for different constituents & Apps
- Increasing costs of extraction, governance, security and maintenance
- Latency in information presented
- Significant volumes of unused data moved frequently

Transition to Read-Transform-Serve for Efficiency

- Leverage efficient in memory processing to create caches aligned to information model of consumers
- More self-serve options for constituent apps and developers, leading to reduced time to value
- Serve real-time / near-real time info
- Align with event-driven enterprise architectures

IBM zDIH use case: 2-Way communication to/from cloud apps

- Asynchronous communication back *from* cloud apps
- SoRs can retrieve the necessary information non-disruptively – either via JDBC or pulling from reverse zDIH logstream

SOR app retrieves enriched info from caches via JDBC

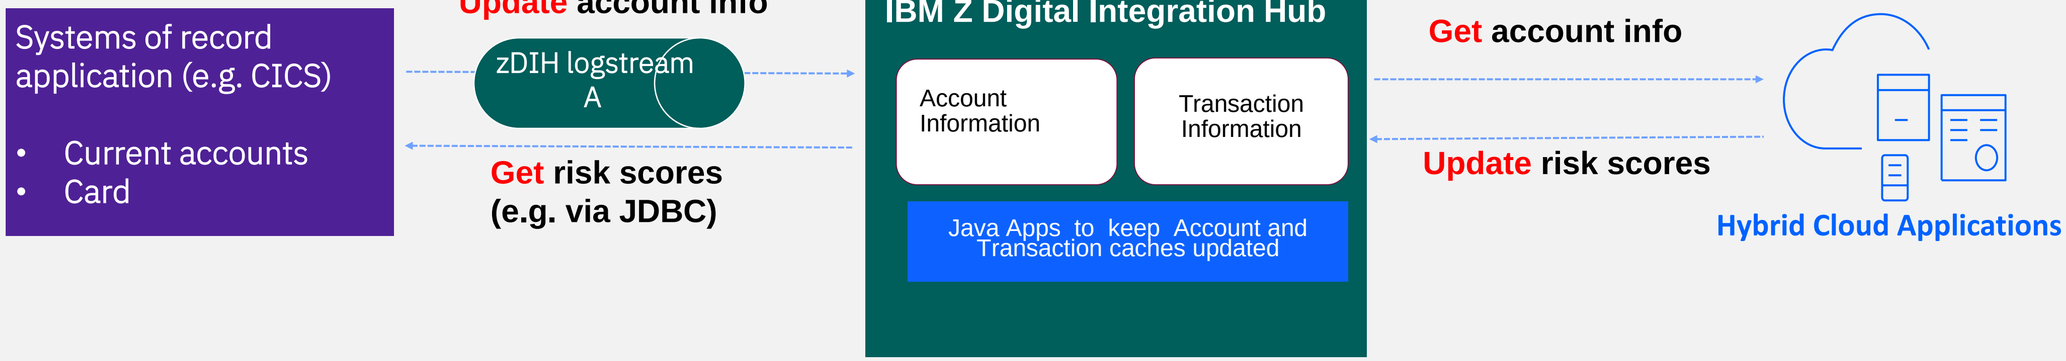

SOR app retrieves info continuously via zDIH reverse logstream

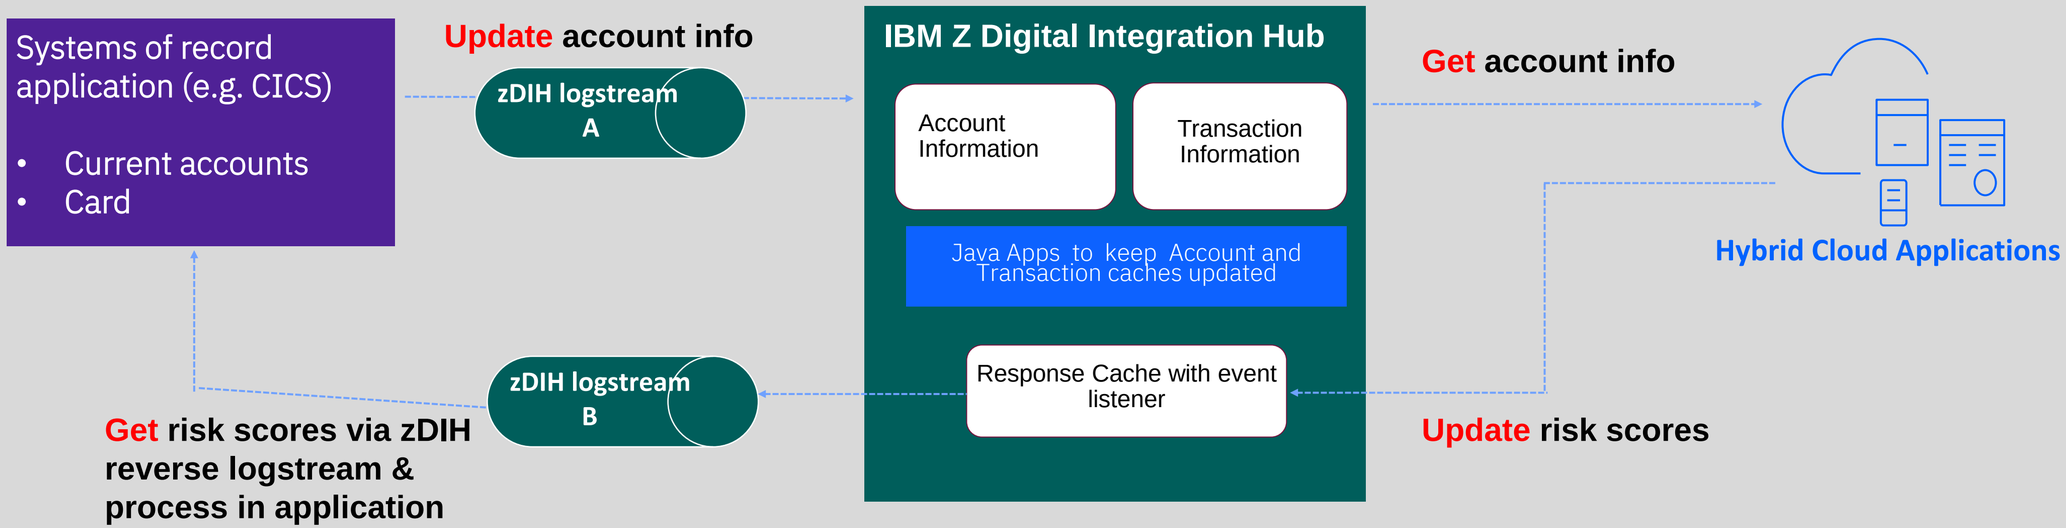

IBM zDIH use case: Enhancing transaction history

Without zDIH

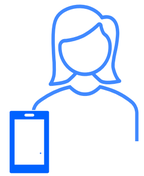

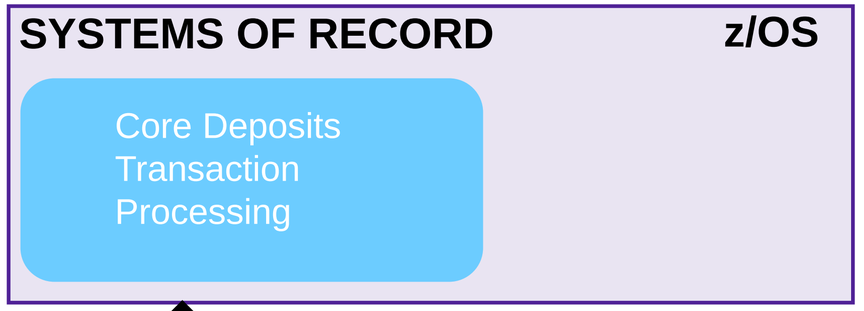

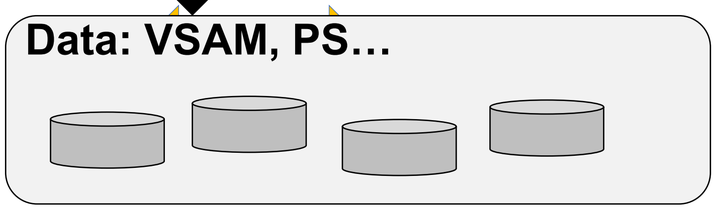

Transaction history searchable for limited time frame (couple months)

- Customers often request access to more history than available
- Results in custom work per request and longer time to satisfy client needs

With zDIH

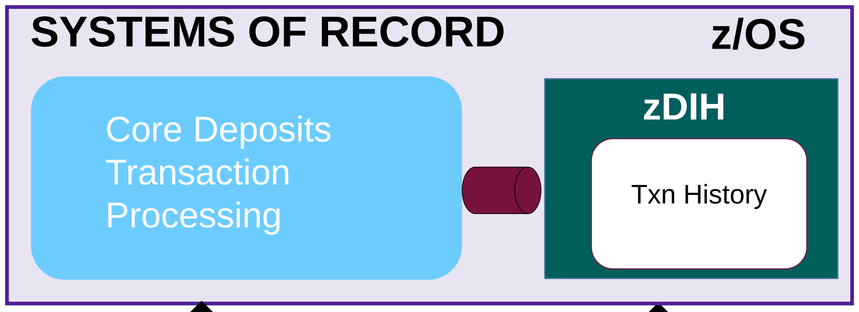

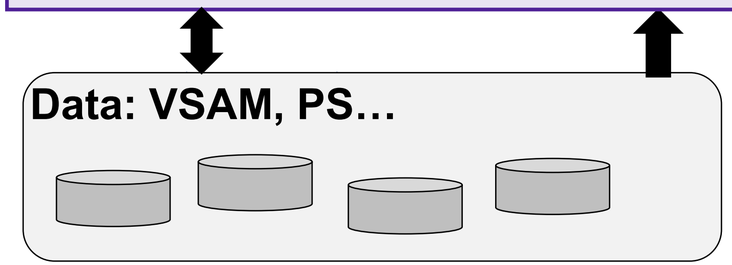

Transaction history over multiple years, fully searchable and available on demand

- Extended history can be made available, on-demand and fully searchable at speed
- Keep only relevant attributes in cache
- Both in-memory and native persistence leveraged, transparent to consumers

Real-time information sharing between z/OS core systems & Public Cloud

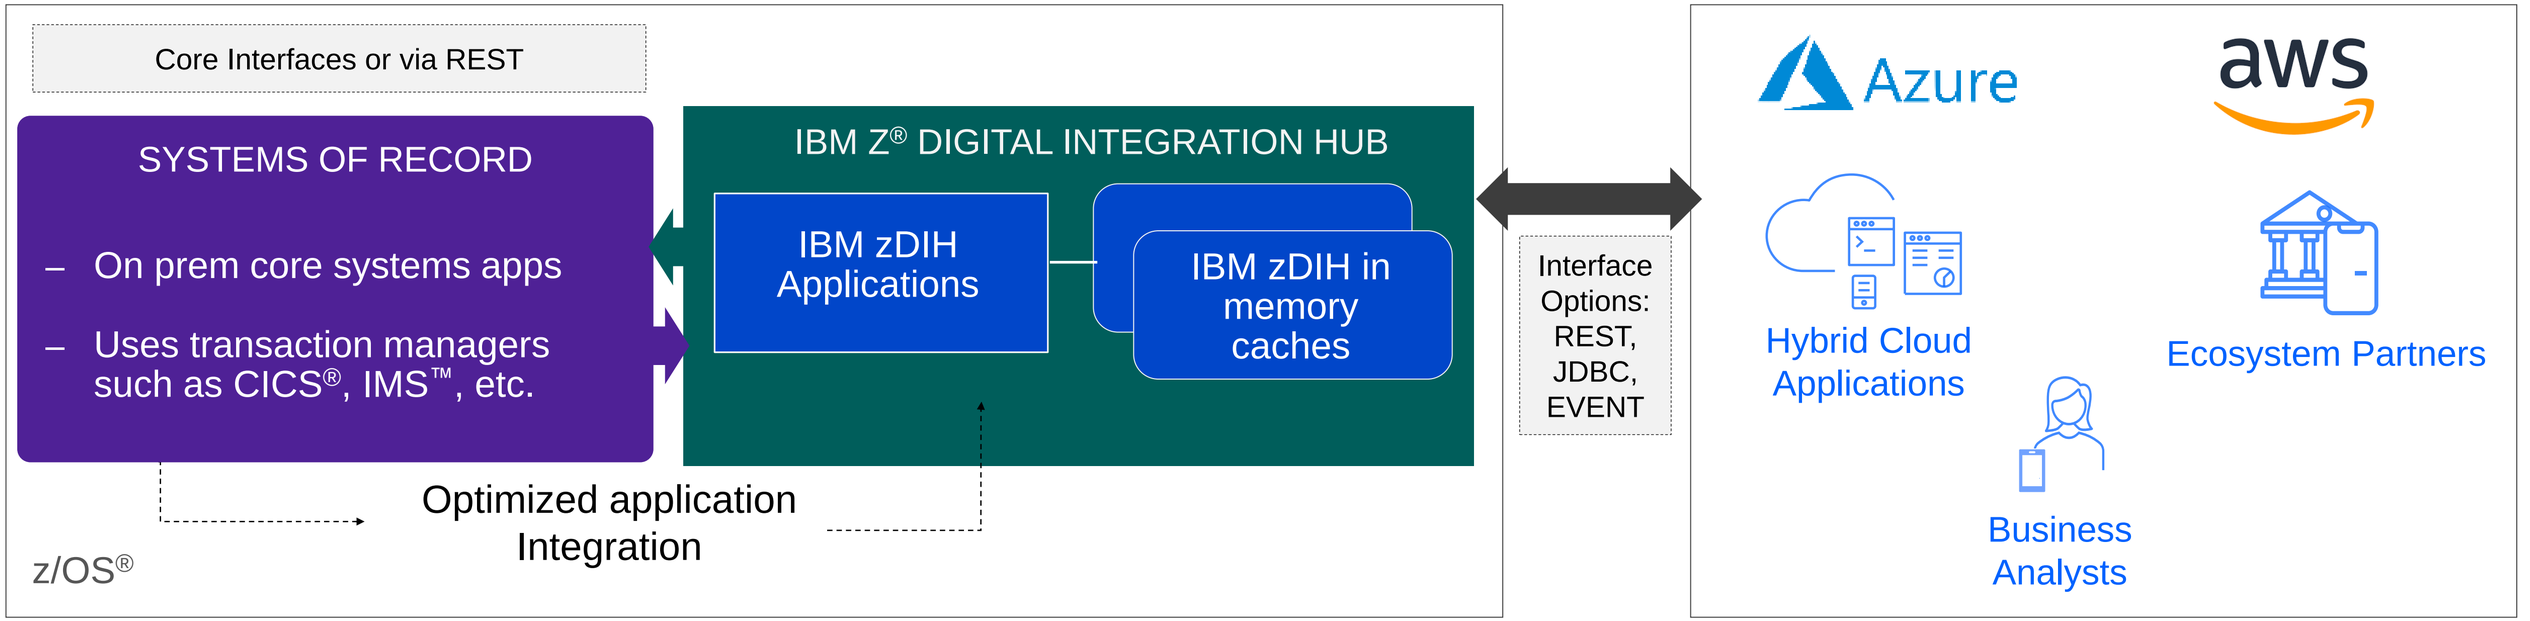

- Aggregated, curated / composed real-time information from core systems of record
- Efficient integration between core on premise z/OS systems applications and cloud native applications
- Minimal impact to mission critical application environments
- Standards based interaction enabling flexibility and decoupling from specific data contexts & data access formats

Azure (pattern 3): https://techcommunity.microsoft.com/t5/azure-migration-and/accelerate-mainframe-application-modernization-with-ibm-and/ba-p/3691322

AWS (pattern 2): https://aws.amazon.com/blogs/apn/modernize-mainframe-applications-for-hybrid-cloud-with-ibm-and-aws/

Combine information across multiple systems of record

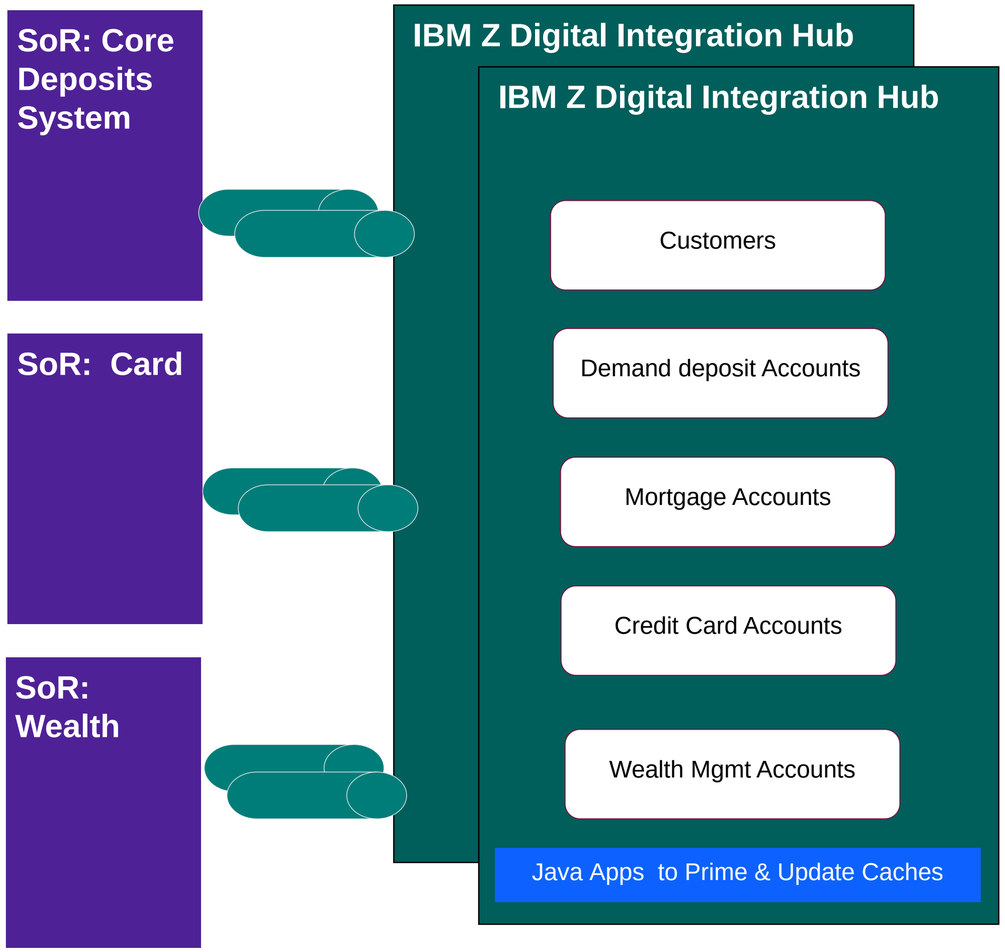

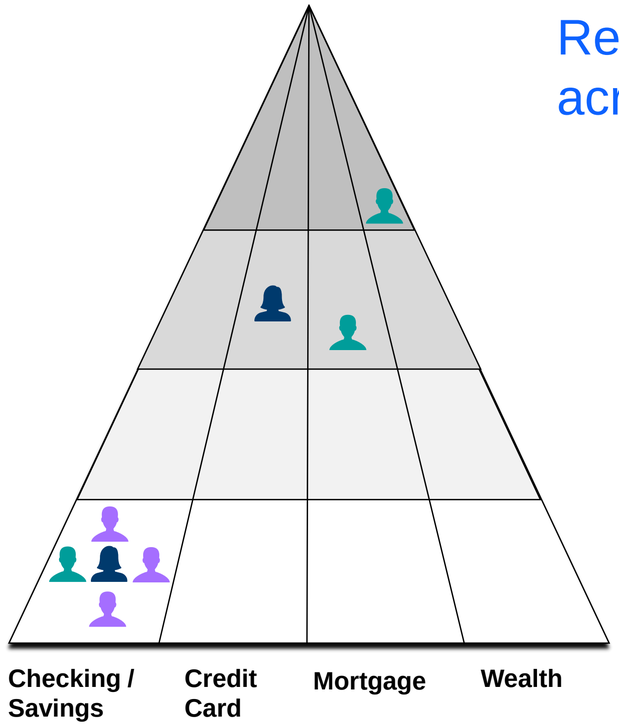

Real time view of customer positions across multiple products

- Find all customers > 5 years of business, with more than 2 accounts that have greater than 5 transactions this month, show real-time balances
- Leverage information for targeted cross product offers

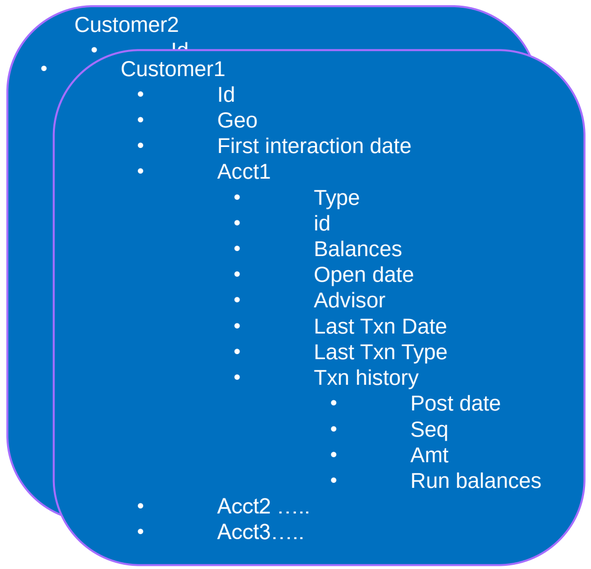

IBM zDIH use case: integrate with multi-cloud or edge caches

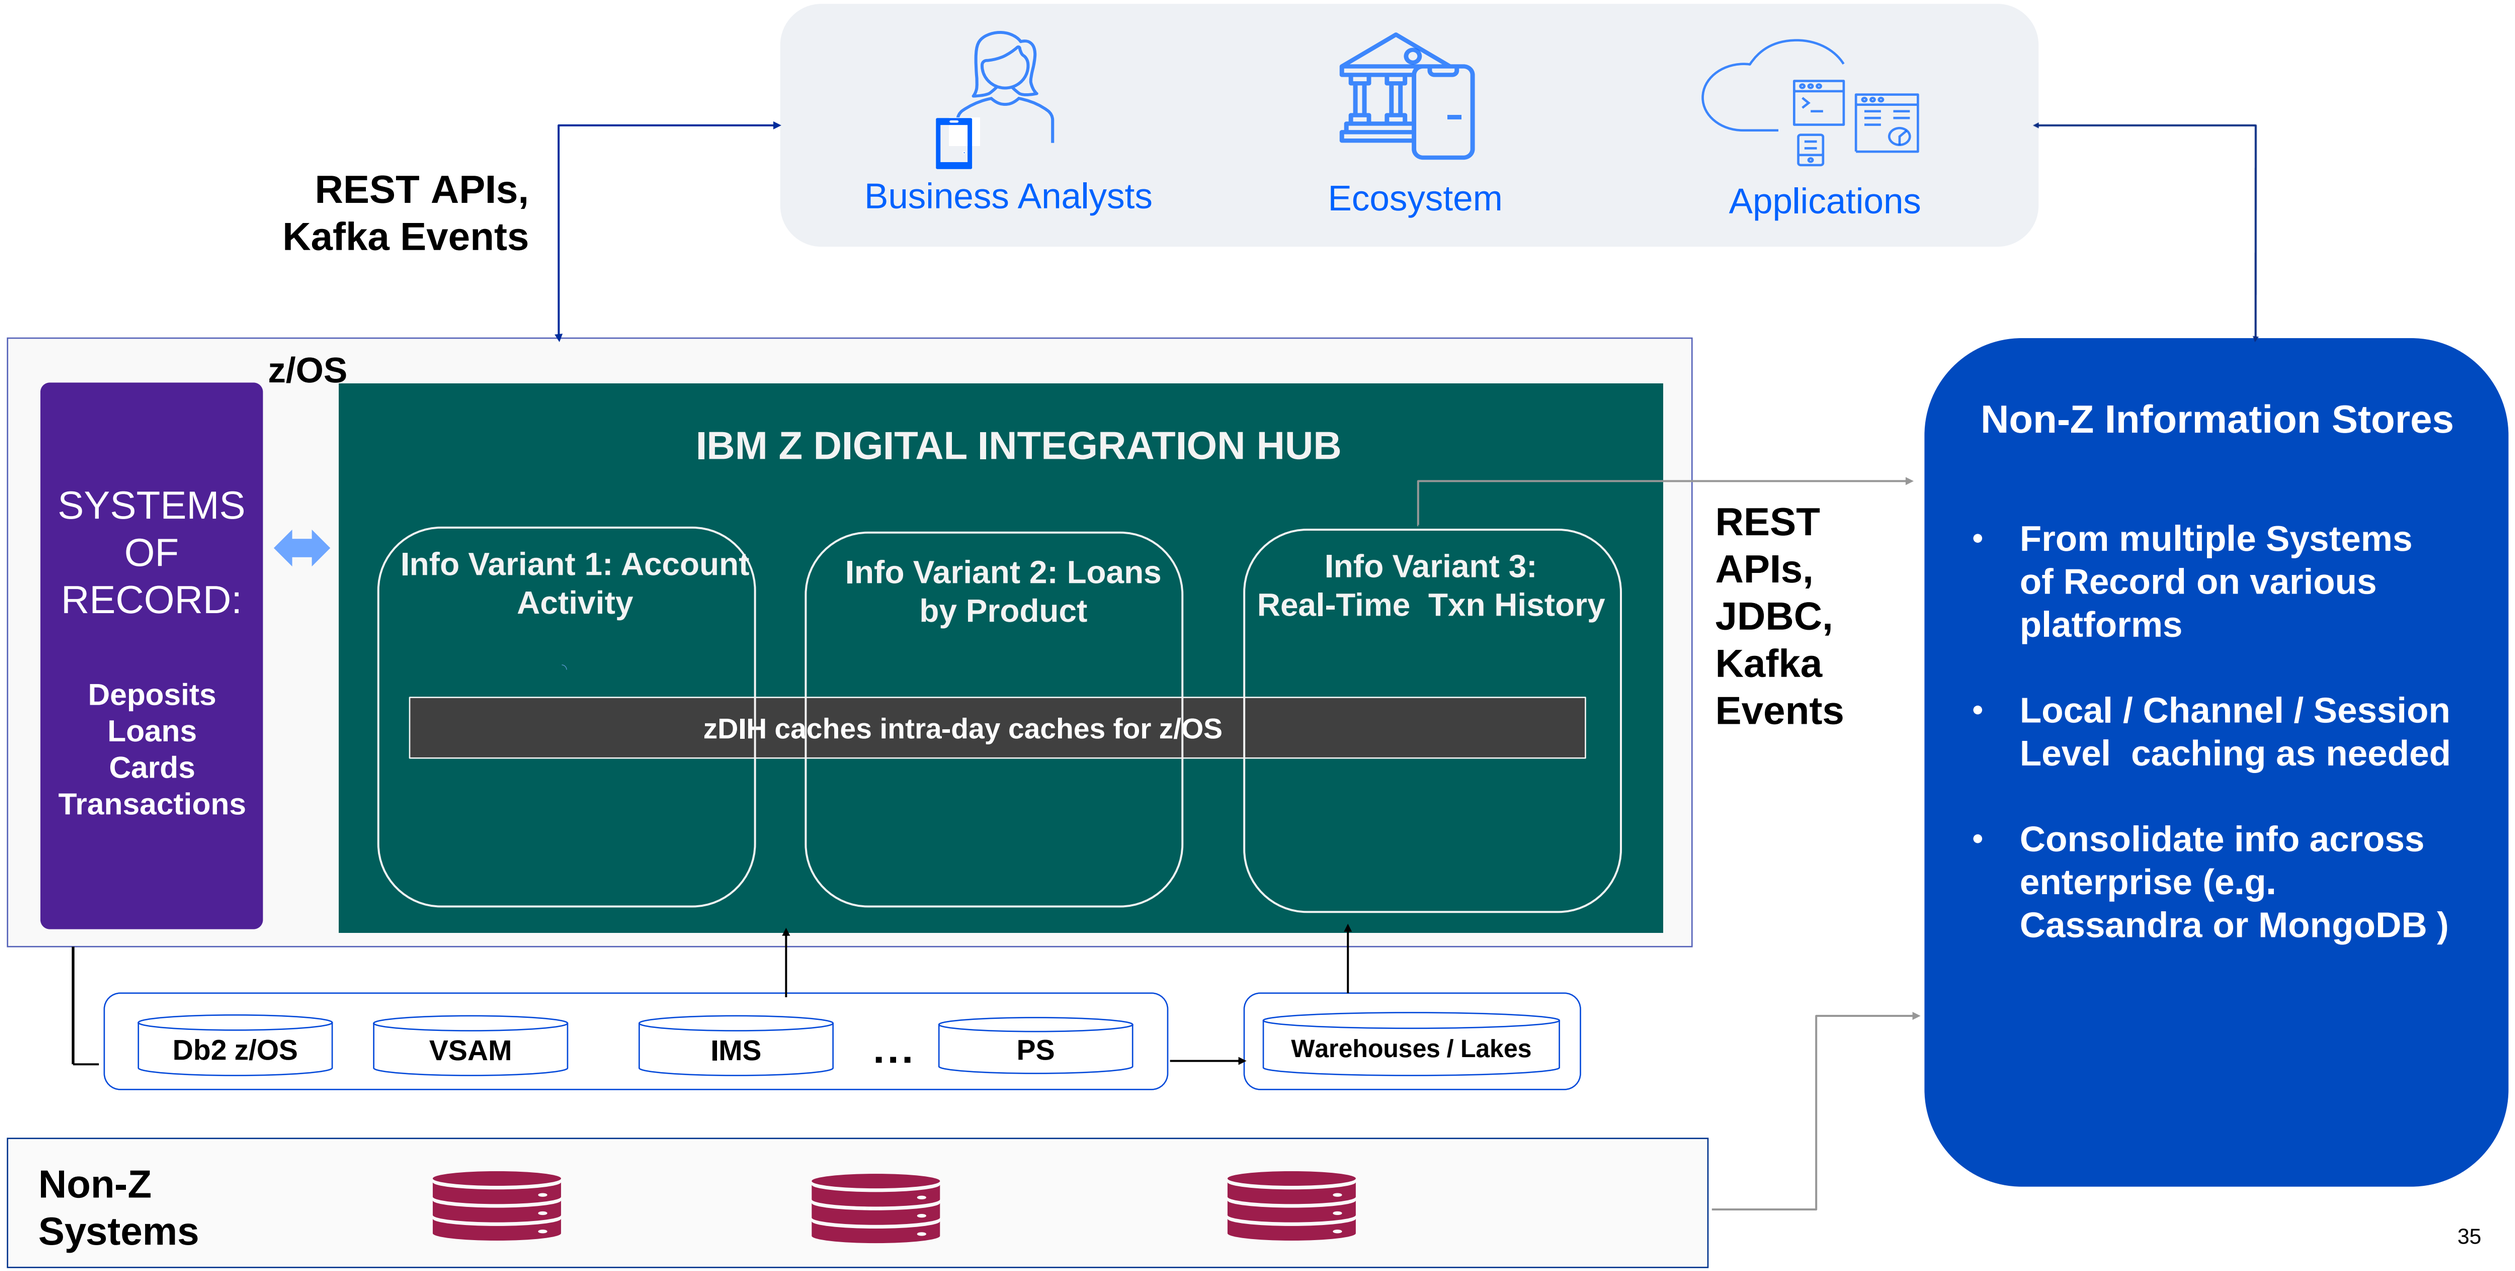

Examples **good fit** use cases:

- Surface information which is composed / aggregated as opposed to all raw data
- Implementing optimized CQRS (separating inquiry and update interactions) for downstream consumers
- Hybrid cloud application has latent information, and needs real-time or more current info
- Transition to more event-driven approach for information flow from systems of record
- Create desired information from combination of batch & online
- Efficient information sharing across multiple z/OS applications
- Re-use composed information by multiple cloud consuming applications

Examples **not good fit** use cases:

- Move all z/OS core systems data to the cloud or another environment
- Access to all core systems of record data for adhoc query interaction
- Stream all data off the platform
- Cache all data from a system of record in zDIH
- AI / ML model training that requires access to granular raw data
- As a replacement for core systems transaction manager or database (e.g. DB2 for z/OS, CICS, IMS, etc.)

For more information about selecting an initial use case for zDIH, see: IBM zDIH Product Documentation: Selecting Initial Use Case

IBM Z Digital Integration Hub engagement

- Can be conducted as a Proof of Concept
- Implemented in short time-boxed duration
- Precise deliverables based on scope & approach
  - ✓ Java apps to prime cache & maintain cache currency aligned to your use case
  - ✓ Foundation to implement additional use cases
  - ✓ Performance recommendations
  - ✓ Recommended configurations for rapid deployment
  - ✓ Knowledge transfer & design documentation

Time boxed POCs to quickly deliver value

Integrate teams & grow skills across IT and application owners

Accelerated use case implementation, deployment & expansion

IBM Z Digital Integration Hub Benefits

- Serve core system of record information to downstream application more efficiently with lower latency
- zDIH provides a source of truth, serving information that is current and ordered
- Consuming applications can access information using standard interfaces such as REST, JMS, JDBC, Kafka directly from the z platform.
- Curate data and information, enriched and tailored to your consuming applications
- Flexible information model allow applications to be developed faster with less effort
- Use of specialty engines (zIIPS) and the CQRS application modernization pattern offloads MIPS and protects the integrity of core systems of record
- Low code integration with your systems of record applications
- Utilize your existing high value z assets, securely, efficiently, and cost effectively, while leveraging all the qualities of the z platform

Additional IBM Z Digital Integration Hub resources

Z Digital Integration Hub landing page

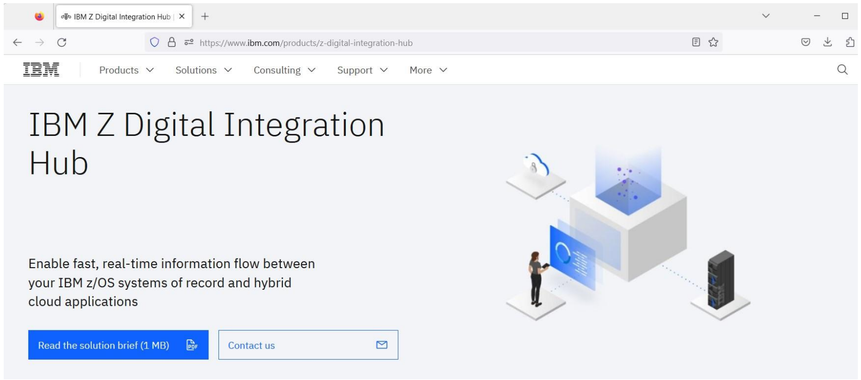

zDIH Product Page

Enterprise API management

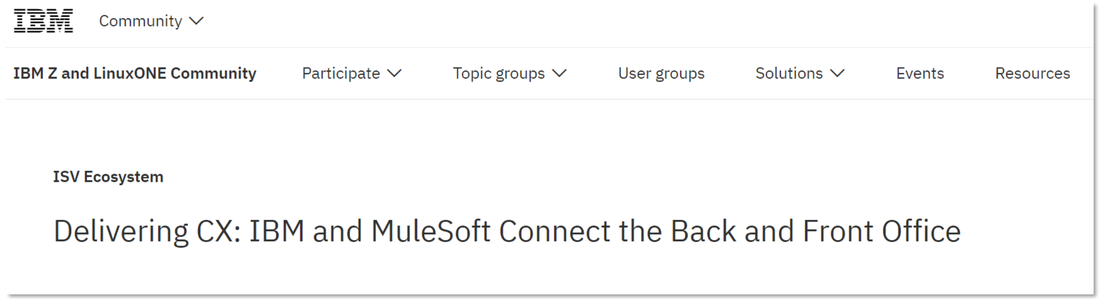

MuleSoft integration with zDIH

M&T Bank

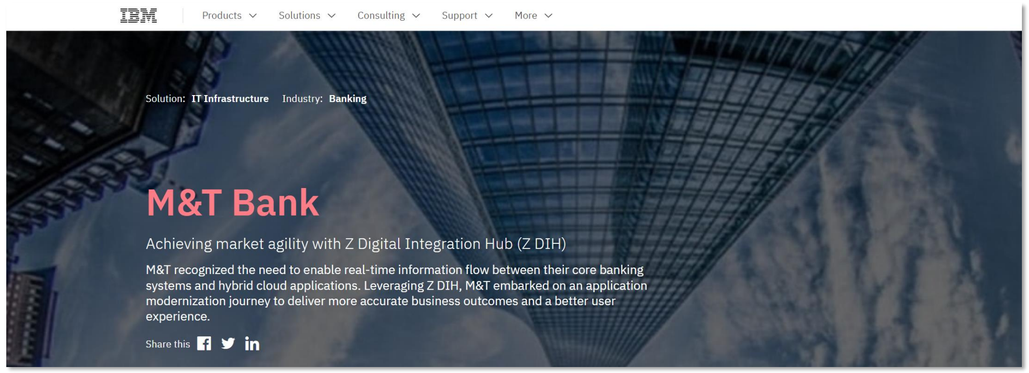

zDIH case study and reference

Modernize mainframe applications

Hybrid Cloud with IBM, AWS, and Azure

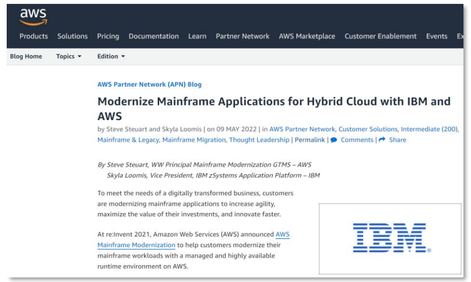

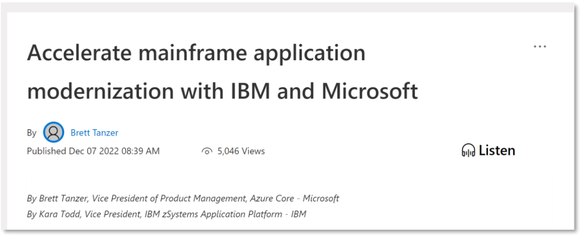

AWS: Pattern 2 Azure: Pattern 3

Q&A

Roy Duke Jr.

RoyDuke@us.ibm.com

How to get started with a zDIH POC

Initial workshop with IBM SMEs

- Identify & validate candidate use cases for technical suitability
- Create high-level architecture flow
- Select high value use case
- Identify business benefits

Follow-on workshops with IBM SMEs

- Define POC scope & success criteria
- Identify detailed architectural flows & data requirements

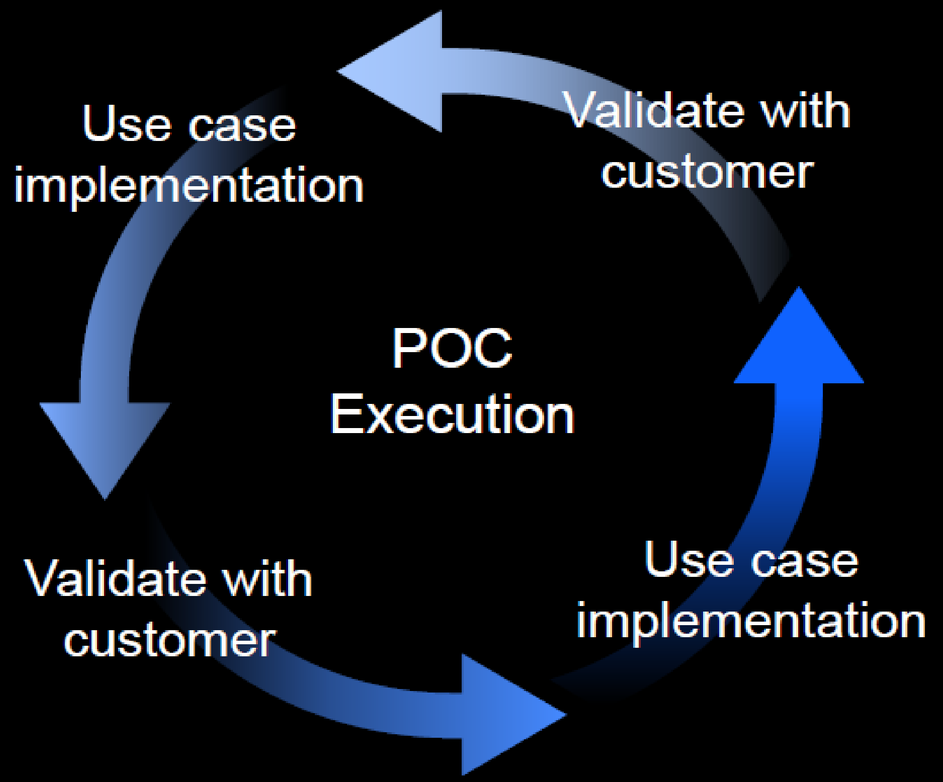

IBM zDIH PoC Approach Detail: Execute the PoCs

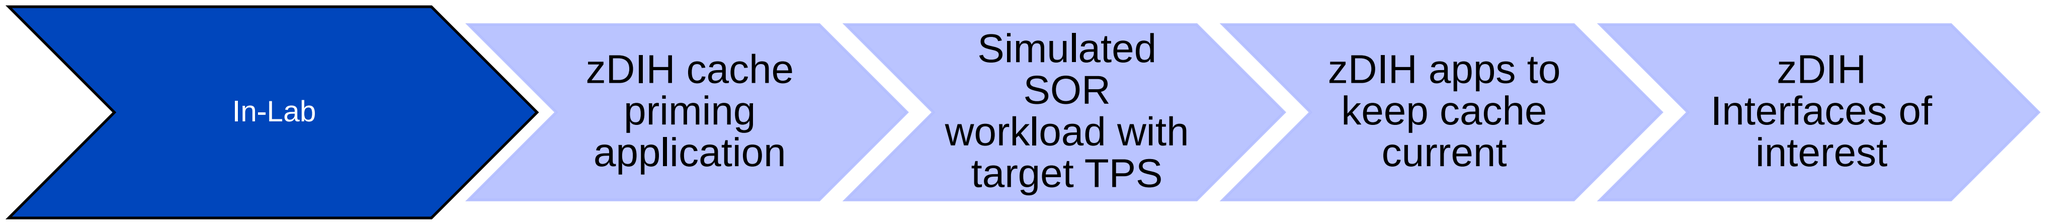

Uses IBM generated test data based on cache structure definitions aligned to use case scope
Client Skills: Application technical lead for systems of record; Periodic Q&A sessions or emails.

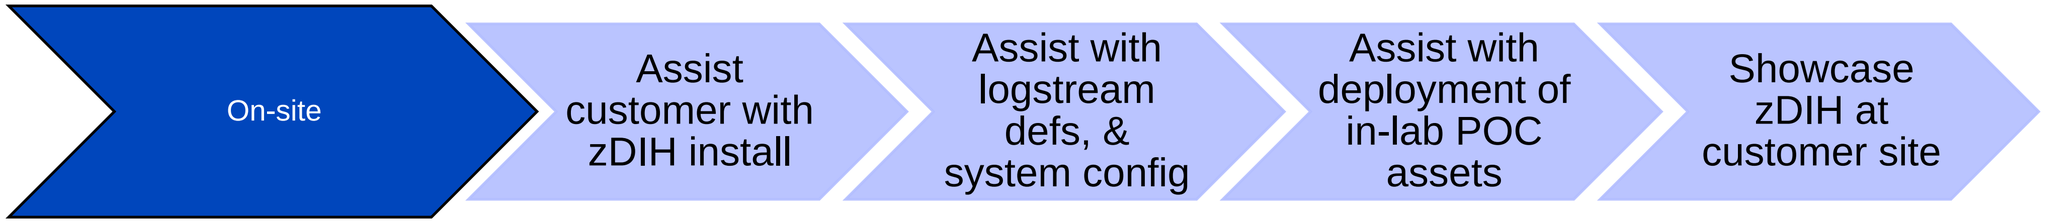

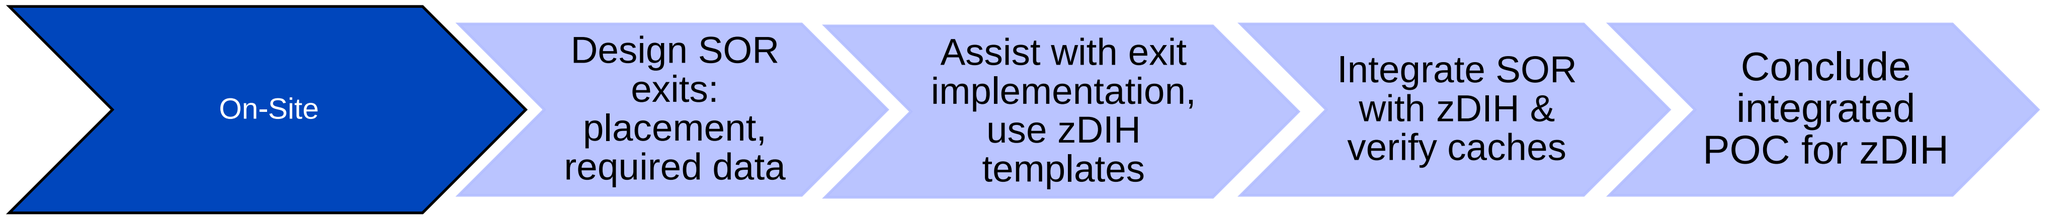

IBM zDIH Installation & Deployment Guidelines

Installation Pre-requisites

- IBM zDIH is SMP/E installable package, running in a JVM on Unix System Services in z/OS
- Software pre-reqs
  - z/OS 2.4 & above
  - Java for z/OS: IBM Semeru Runtime Certified Edition for z/OS, Version 11.0
  - Bash 4.3 & above
- Hardware pre-reqs:
  - z14 & above
  - Rule of thumb requirements to start a POC: 2-4 zIIPs & 50 gig memory
- Management Server: Any of Linux®, Unix, Windows environments

General Deployment Guidelines

- For HA/DR purposes, recommend installing zDIH in at least 2 LPARs
- For best performance, recommend installing zDIH in same LPAR as one of the core business applications which has the information to be shared (do not need to install zDIH in every LPAR)
- zIIPs and Memory needs can be staged over time
  - Requirements depend on use case, SLAs, and existing available installed capacity
  - Hardware sizing for initial deployment based on measurements captured during POC

IBM zDIH security mechanisms

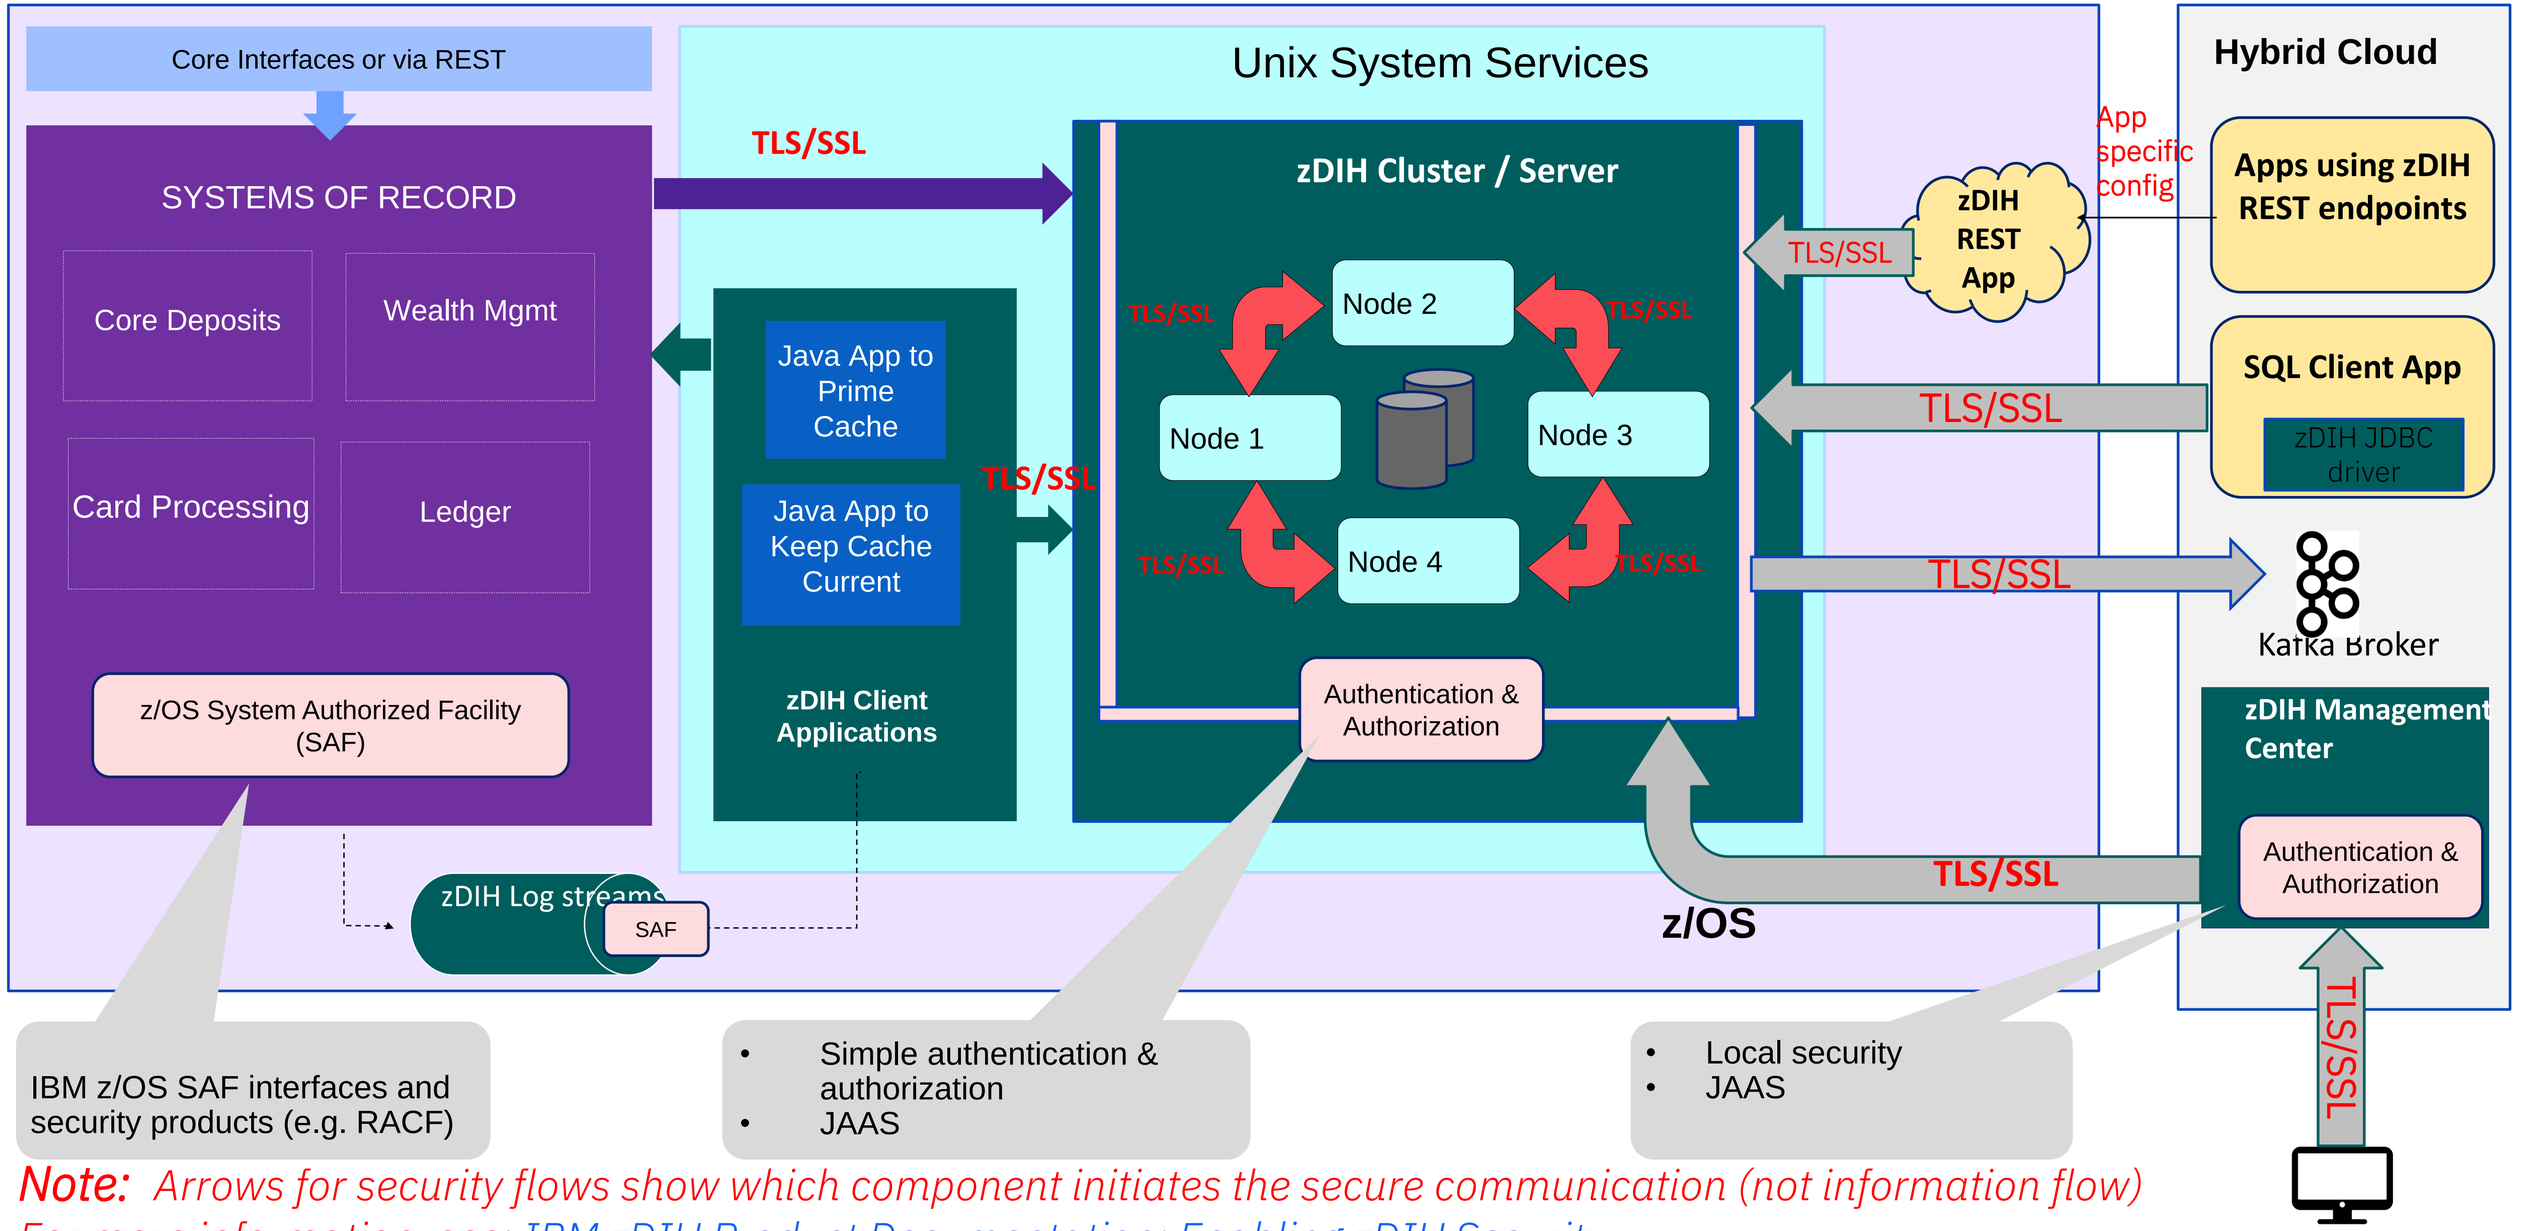

Note: *Arrows for security flows show which component initiates the secure communication (not information flow)*

For more information, see: IBM zDIH Product Documentation: Enabling zDIH Security

IBM zDIH use case: integrate enterprise API management

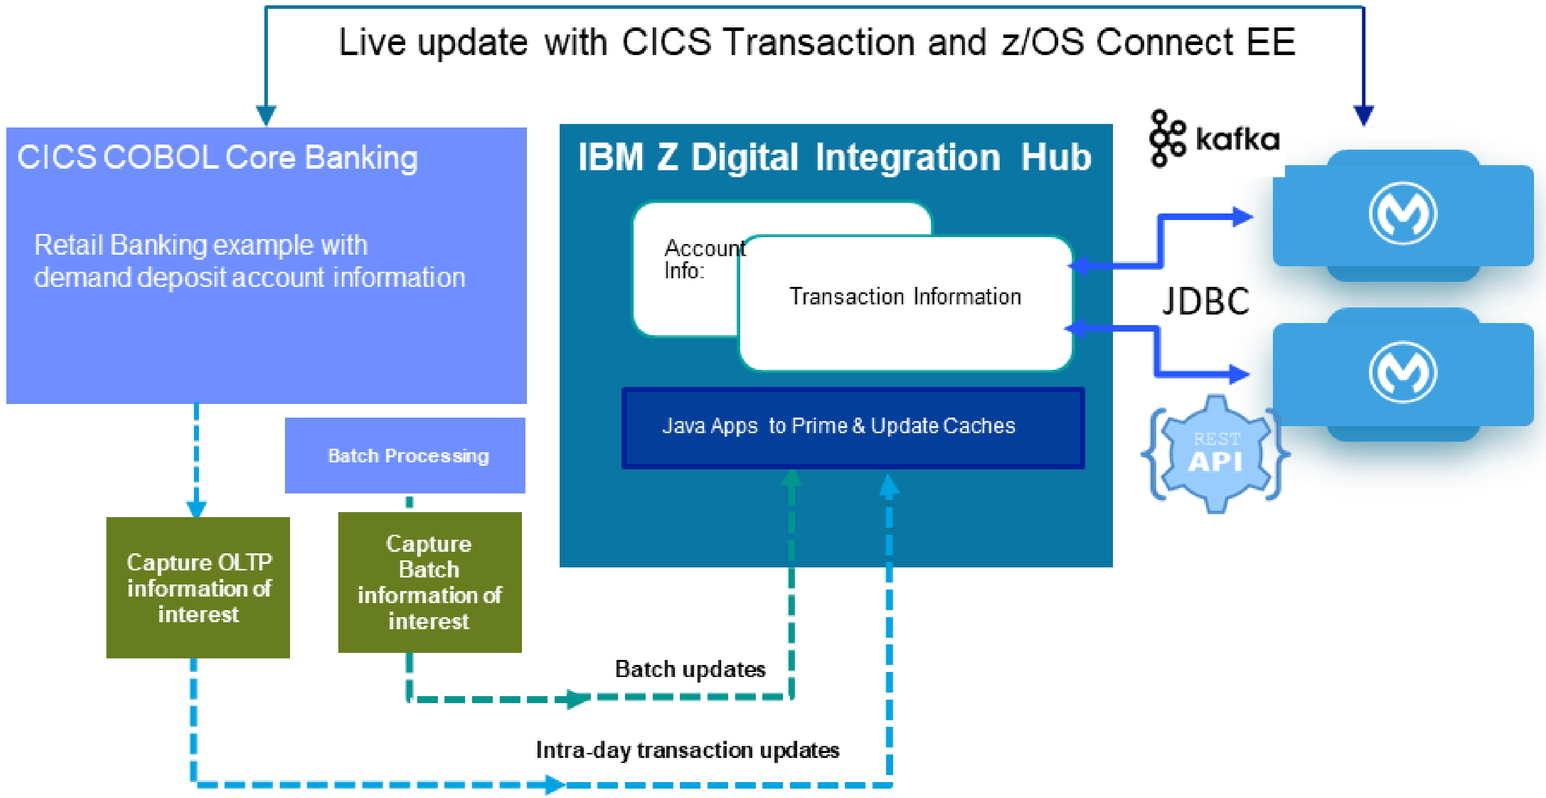

- Real-time information flow at scale between Systems of Record and MuleSoft
- Faster & optimized integration between Mulesoft and core applications running on IBM Z
- Cost optimization through separation of query processing from core transactions

Expand the ecosystem

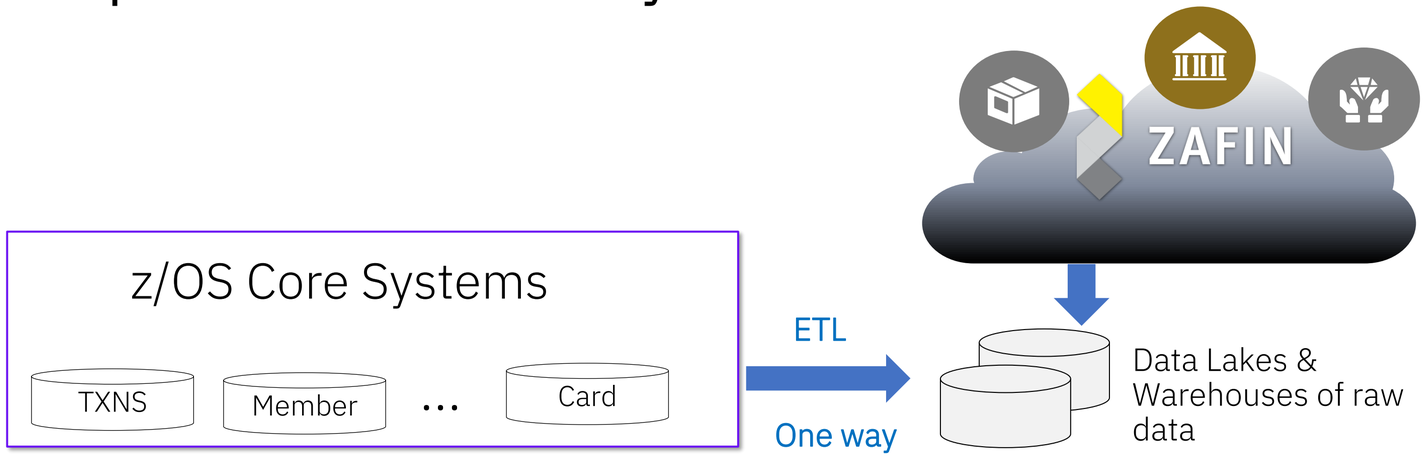

Without zDIH:

Zafin consumes stale ETL data for calculation of fees, rates, rewards, offers and product pricing

One way ETL causes delayed SOR use of product pricing recommendations

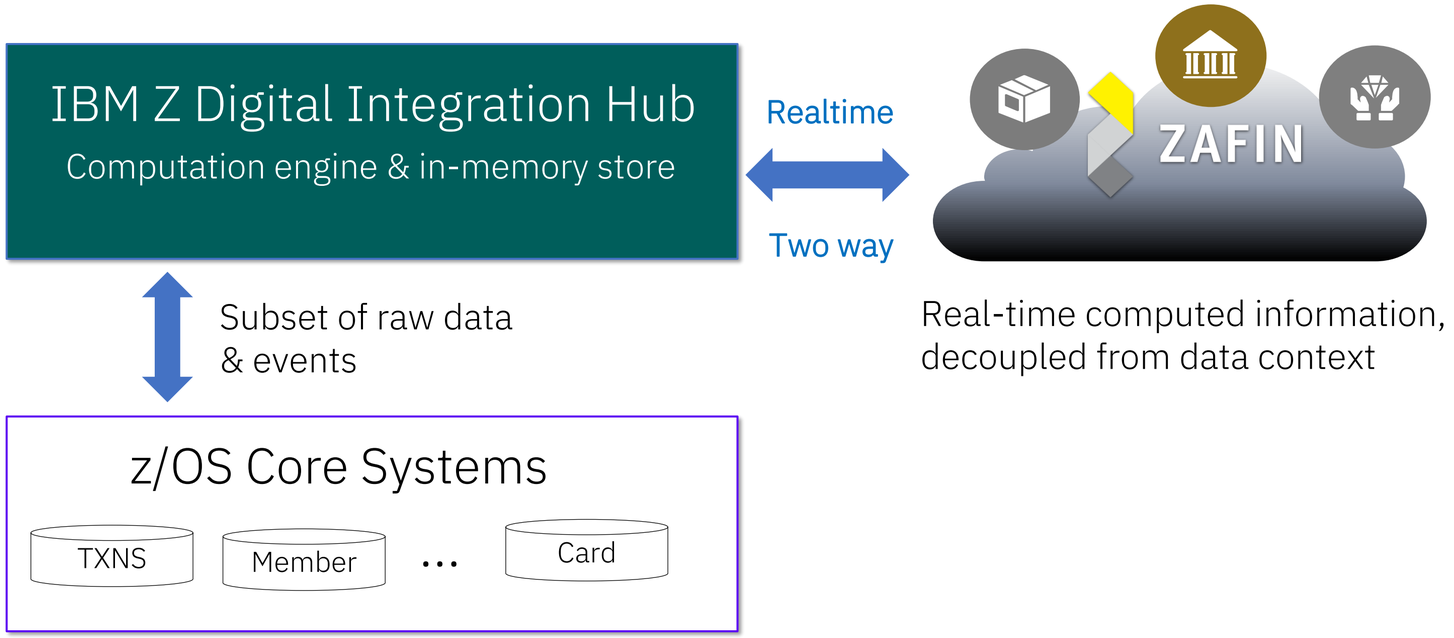

With zDIH:

Real-time information flow from core systems to Zafin enables better business outcome through more accurate product and pricing controls

Two-way communication between ecosystem & zDIH ensures SOR can process pricing recommendations in real time, delivering business agility