Journal of Advances in Developmental Research (IJAIDR)



E-ISSN: 0976-4844 • Website: <u>www.ijaidr.com</u> • Email: editor@ijaidr.com

Implementing Zero-Copy Integration in Hybrid Cloud Environments for Seamless Salesforce Data Exchange

Maneesh Gupta

Salesforce CRM Architect/ Evangelist Zionsville, USA <u>Maneesh_83@yahoo.co.in</u>

Abstract

In the IT world, hybrid cloud architectures have become integral to digital transformation strategies. According to Flexera's 2024 State of the Cloud Report, 94% of enterprises have adopted cloud services, with 39% implementing hybrid cloud models to balance on-premises infrastructure with public cloud platforms¹. This shift introduces complex data integration challenges, particularly when integrating mission-critical systems like Salesforce.



Salesforce maintains a dominant position in the Customer Relationship Management (CRM) market, holding a 21.7% market share as of 2024. Its role as a central hub for customer data and business process automation underscores the necessity for efficient and secure data integration across hybrid environments. Traditional data integration methods, such as Extract-Transform-Load (ETL) processes, often result in latency, increased operational costs, and data duplication risks. These issues are exacerbated when real-time data access is required across disparate systems².



Zero-copy integration emerges as a compelling solution to these limitations. By enabling real-time, read-only access to source systems without data duplication, zero-copy architectures facilitate seamless interaction between Salesforce and on-premises systems. Techniques such as data virtualization, federated queries, and external object modeling allow for efficient data exchange, reducing latency and maintaining control over sensitive information. Salesforce's Data Cloud exemplifies this approach, offering zero-copy capabilities that harmonize data across sources via metadata, eliminating the need for data duplication or migration³.

This whitepaper explores the implementation of zero-copy integration strategies to modernize hybrid cloud data exchange for Salesforce. It addresses security, scalability, and performance considerations, providing architectural best practices and technology options to support enterprise adoption of zero-copy methods in hybrid environments.

1. Introduction

Hybrid cloud architectures have emerged as a strategic approach to balance the benefits of both onpremises infrastructure and public cloud services. Hybrid cloud combines and unifies public cloud, private cloud, and on-premises infrastructure, providing businesses with the flexibility to optimize for cost, performance, and compliance.

Organizations adopt hybrid cloud models for various reasons. Compliance requirements often necessitate that certain sensitive data remain on-premises, while less sensitive workloads can be moved to the public cloud to leverage scalability and cost-efficiency. Performance considerations also play a role; latency-sensitive applications may perform better when hosted closer to end-users or integrated with existing on-premises systems. Furthermore, many enterprises continue to rely on legacy systems that are not easily migrated to the cloud, making a hybrid approach more practical.

Salesforce, as a leading customer relationship management platform, is often at the center of these hybrid environments. While Salesforce operates in the cloud, it frequently needs to integrate with onpremises systems such as enterprise resource planning software, databases, and other legacy applications. This integration is crucial for maintaining data consistency and providing a unified view of customer information across the organization⁴.

However, integrating Salesforce within a hybrid cloud environment presents several challenges:





- Latency: Data exchange between cloud-based Salesforce and on-premises systems can introduce latency, affecting real-time data access and application performance.
- **Redundancy:** Traditional integration methods may involve duplicating data across systems, leading to inconsistencies and increased storage costs.
- Security: Transferring sensitive data between on-premises and cloud environments raises security concerns, including data breaches and compliance violations.

Addressing these challenges requires innovative integration strategies that ensure seamless, secure, and efficient data exchange between Salesforce and on-premises systems within a hybrid cloud architecture. This whitepaper explores the concept of zero-copy integration as a solution to these challenges, aiming to enhance data exchange processes and overall system performance.

2. The Problem with Traditional Data Integration Approaches

As organizations increasingly adopt hybrid cloud architectures to balance on-premises infrastructure with public cloud services, integrating data across these environments presents significant challenges. Traditional data integration methods, such as Extract, Transform, Load processes, often fall short in addressing the complexities of hybrid environments, particularly when integrating with platforms like Salesforce.

2.1 ETL Inefficiencies

ETL processes involve extracting data from various sources, transforming it to fit operational needs, and loading it into target systems. While conceptually straightforward, ETL can be technically complex and challenging to implement, leading to slowness and inefficiencies. For instance, two-thirds of ETL projects exceed their initial time estimates, indicating the complexity and unpredictability of these processes⁵.

2.2 Data Duplication and Syncing Delays

Traditional integration methods often involve duplicating data across systems, leading to inconsistencies and increased storage costs. In hybrid environments, maintaining data consistency becomes more challenging due to the asynchronous nature of batch processing and the potential for data to become outdated between sync cycles. This can result in decision-making based on stale or inconsistent data.

2.3 Latency Issues in Batch Pipelines

Batch processing introduces latency, as data is collected, processed, and loaded at scheduled intervals rather than in real-time. This delay can hinder timely access to critical information, affecting business operations that rely on up-to-date data. In hybrid cloud environments, the latency is further exacerbated by the need to transfer data across different infrastructures, leading to potential bottlenecks and reduced performance.



2.4 Security and Compliance Risks in Data Movement

Transferring data between on-premises systems and the cloud introduces security vulnerabilities, including the risk of data breaches during transit. Ensuring compliance with regulations such as GDPR and HIPAA becomes more complex when data is moved across different jurisdictions and storage environments. Implementing robust encryption and access controls is essential, but traditional integration methods may not provide the necessary granularity or real-time monitoring capabilities.

2.5 Operational Overhead and Scalability Limits

Managing traditional data integration processes requires significant operational effort, including monitoring ETL jobs, handling failures, and scaling infrastructure to accommodate growing data volumes. This overhead can divert resources from strategic initiatives and limit the organization's ability to scale efficiently. As data sources and business requirements evolve, the rigidity of traditional integration approaches can hinder agility and responsiveness.

Traditional data integration methods present several limitations in hybrid cloud environments, including inefficiencies in ETL processes, data duplication, latency issues, security risks, and operational overhead. These challenges necessitate the exploration of more modern, efficient, and secure integration strategies to meet the demands of today's dynamic business landscape.

3. What is Zero-Copy Data Integration?

Zero-copy data integration is an architectural approach that enables systems to access and utilize data directly from its source location, eliminating the need for physical data movement or duplication. This method contrasts with traditional data integration techniques, such as Extract, Transform, Load, which involve copying data into centralized repositories for processing and analysis.

By allowing real-time access to data across disparate systems, zero-copy integration reduces latency, minimizes storage costs, and mitigates the risks associated with data redundancy and synchronization errors. This approach is particularly beneficial in hybrid cloud environments, where data resides across both on-premises and cloud-based systems.

3.1 Technical Foundations

Implementing zero-copy data integration relies on several key technologies:

Technical Foundations of Zero-Copy Data Integration:





- **Data Virtualization:** This technique creates a virtual data layer that provides a unified view of data from multiple sources without moving the data. It enables applications to query and manipulate data in real-time, regardless of its physical location.
- Application Programming Interfaces (APIs): APIs facilitate direct communication between applications and data sources, allowing for seamless data access and integration without the need for data replication⁶.
- **Shared Memory:** In certain scenarios, shared memory techniques allow multiple processes to access the same memory space, enabling efficient data sharing without copying.
- **Metadata Federation:** This involves creating a centralized metadata repository that provides information about data across various sources, enabling consistent data interpretation and access without physical data consolidation⁷.

3.2 Benefits

Adopting zero-copy data integration offers several advantages:

- **Real-Time Access:** By accessing data directly from its source, organizations can obtain up-todate information instantly, enhancing decision-making processes⁸.
- **Cost Efficiency:** Eliminating data duplication reduces storage requirements and associated costs.
- **Reduced Risk:** Minimizing data movement decreases the likelihood of data breaches and synchronization errors, enhancing data security and integrity.

3.3 Applicable Architectures

Zero-copy data integration aligns well with several modern architectural paradigms:

- **Data Mesh:** This decentralized approach treats data as a product, with each domain owning its data pipelines. Zero-copy integration supports data mesh by enabling seamless data access across domains without duplication.
- Lakehouse: Combining elements of data lakes and data warehouses, lakehouse architectures benefit from zero-copy integration by allowing real-time data access without the need for data movement.
- Service-Oriented Architectures (SOA): In SOA, services communicate over a network to provide functionality. Zero-copy integration enhances SOA by facilitating efficient data sharing between services without redundant data transfers.

Zero-copy data integration offers a modern solution to the challenges of data access and management in complex, hybrid environments. By leveraging technologies like data virtualization, APIs, shared



memory, and metadata federation, organizations can achieve efficient, secure, and real-time data integration without the drawbacks of traditional methods.

4. Applying Zero-Copy to Salesforce in Hybrid Clouds

In hybrid cloud environments, where data resides across both on-premises systems and cloud platforms, efficient and secure data integration becomes paramount. Salesforce, as a central hub for customer relationship management, often needs to interact with various external data sources without compromising performance or data integrity. Zero-copy integration offers a solution by enabling real-time access to external data without the need for physical data movement or duplication.

4.1 Salesforce as Both Data Producer and Consumer

Salesforce functions dually as a data producer, generating customer interactions, sales data, and service records, and as a data consumer, requiring access to external data such as inventory levels, financial records, or supply chain information. This bidirectional data flow necessitates a seamless integration mechanism that maintains data consistency and accessibility across disparate systems.

4.2 Utilizing Salesforce Connect and External Objects

Salesforce Connect is a powerful tool that facilitates zero-copy integration by allowing Salesforce to access external data sources in real-time. It achieves this through the use of External Objects, which are representations of data stored outside of Salesforce. These objects enable users to interact with external data as if it were native to Salesforce, supporting operations like search, reporting, and workflow automation.

The integration leverages the Open Data Protocol (OData), a standardized protocol for building and consuming RESTful APIs. By connecting to OData-compliant endpoints, Salesforce can retrieve metadata and data from external systems, allowing for dynamic interaction without data replication. This approach is particularly beneficial for integrating with on-premises databases or legacy systems that support OData.

4.3 Connecting On-Premises Databases Without Replication

To connect on-premises databases to Salesforce without data replication, organizations can expose their databases as OData endpoints using middleware solutions or data virtualization platforms. Once the OData service is established, Salesforce Connect can be configured to access these endpoints, creating External Objects that map to the external data structures. This setup allows Salesforce users to view and interact with on-premises data in real-time, eliminating the need for data duplication and ensuring data consistency.

4.4 Handling Schema Mapping, Transformations, and Metadata

Effective zero-copy integration requires careful handling of schema mapping and data transformations. Salesforce Connect retrieves metadata from the external data source, enabling the creation of External Objects that mirror the structure of the source data. Administrators can define relationships between External Objects and standard Salesforce objects, facilitating seamless data interactions. Additionally,



data transformations can be applied at the middleware layer or within the external system to ensure compatibility with Salesforce's data model.

4.5 Querying ERP Data from Salesforce UI

Consider an organization that uses an on-premises Enterprise Resource Planning system to manage inventory and financial data. By exposing the ERP system's data as OData endpoints and configuring Salesforce Connect, the organization can create External Objects representing the ERP data within Salesforce. Sales representatives can then access real-time inventory levels and financial information directly from the Salesforce interface, enabling informed decision-making without the need for data replication or manual synchronization.

This zero-copy integration approach enhances operational efficiency, reduces data redundancy, and ensures that users have access to the most current information across systems.



Benefits of Zero-Copy Integration

5. Addressing Latency and Performance Challenges

In hybrid cloud environments, where data traverses both on-premises systems and cloud platforms, latency and performance issues are prevalent. Zero-copy integration strategies aim to mitigate these challenges by enabling real-time data access without duplication.

5.1 Challenges with API Response Times Across Environments

API response times can be adversely affected in hybrid settings due to factors such as network latency, service call delays, and database access times. These delays can hinder real-time data access and degrade user experience. For instance, network latency, often the most significant contributor, results from the time taken for data to travel across an internet connection⁹.

5.2 Edge Computing and Caching Layers

Implementing edge computing and caching layers can significantly reduce latency. Edge computing brings computation closer to the data source, minimizing the distance data must travel and thereby reducing latency. Caching frequently accessed data at the edge ensures faster data retrieval and improved application responsiveness¹⁰.



5.3 Federated Queries and Pushdown Optimization

Federated queries allow systems to access and query data across multiple sources without data movement. However, performance can be enhanced through predicate pushdown, where query filters are applied as close to the data source as possible, reducing the volume of data transferred and processed. For example, Amazon Athena supports predicate pushdown to improve federated query performance¹¹.

5.4 Best Practices for High-Throughput, Low-Latency Integrations

To optimize performance in hybrid cloud environments:

Integration Tips for Peak Performance:



- **Optimize Network Configurations:** Utilize high-bandwidth, low-latency connections such as AWS Direct Connect or Azure ExpressRoute to ensure efficient data transfer between on-premises and cloud systems .
- **Implement Data Virtualization:** Use data virtualization techniques to access data in real-time without physical movement, reducing latency and improving scalability.
- Monitor and Analyze Performance Metrics: Regularly monitor API response times, network latency, and system throughput to identify and address performance bottlenecks promptly.

By adopting these strategies, organizations can enhance the performance and responsiveness of their hybrid cloud integrations, ensuring efficient and real-time data access across environments.

6. Security and Compliance in Zero-Copy Integrations

Zero-copy data integration offers significant advantages in hybrid cloud environments by enabling realtime data access without duplication. However, this approach necessitates robust security measures and strict compliance adherence to protect sensitive information and meet regulatory requirements.

6.1 Role-Based Access Control and Data Masking

Implementing Role-Based Access Control (RBAC) is essential in zero-copy architectures to ensure that users access only the data necessary for their roles. RBAC restricts data access based on user roles within an organization, aligning with the principle of least privilege and supporting compliance with regulations like GDPR¹².



Data masking further enhances security by obscuring sensitive information, allowing users to interact with data without exposing confidential details. This technique is particularly useful in environments where data is accessed across various systems and by multiple stakeholders.

6.2 Auditing and Data Residency Considerations

Maintaining comprehensive audit logs is critical for tracking data access and modifications in zero-copy integrations. Audit trails provide visibility into who accessed data, when, and what actions were performed, facilitating compliance with standards like SOC 2 and HIPAA¹³.

Data residency requirements mandate that data remains within specific geographic boundaries. Zerocopy integrations must ensure that data access complies with these requirements, often necessitating the use of localized data centers or cloud regions¹⁴.

6.3 Compliance with GDPR, HIPAA, and SOC 2

Zero-copy integrations must adhere to various regulatory frameworks:

- **GDPR:** Requires data minimization and protection of personal data, which zero-copy supports by eliminating unnecessary data duplication and enabling precise access controls.
- **HIPAA:** Mandates the safeguarding of protected health information (PHI). Zero-copy architectures can enhance HIPAA compliance by restricting data movement and ensuring secure, controlled access to PHI.
- SOC 2: Focuses on the security, availability, and confidentiality of data. Zero-copy integrations support SOC 2 compliance through stringent access controls, audit logging, and data protection measures.

6.4 Advantages Over Traditional Approaches

Traditional data integration methods often involve data replication, increasing the risk of data sprawl and potential breaches. Zero-copy integration mitigates these risks by:

- Reducing Data Duplication: Eliminates redundant data copies, minimizing exposure points.
- Enhancing Data Control: Maintains data within its original secure environment, allowing for consistent security policies.
- **Improving Compliance:** Facilitates adherence to data protection regulations by simplifying data governance and auditability .

Zero-copy data integration, when implemented with robust security measures and compliance considerations, provides a secure and efficient framework for data access in hybrid cloud environments.



7. Implementation Strategy & Architecture Blueprint

Implementing zero-copy integration within a hybrid cloud environment for Salesforce necessitates a structured approach that encompasses architectural design, tool selection, and phased execution. This section outlines a high-level architecture and a step-by-step roadmap to guide organizations through the implementation process.

7.1 Reference Architecture for Salesforce Zero-Copy in Hybrid Cloud

A robust zero-copy integration architecture typically comprises the following layers:

- **Data Sources:** These include on-premises databases, cloud-based data warehouses (e.g., Snowflake, Databricks), and other enterprise systems ¹⁵.
- **Data Virtualization Layer:** This layer employs platforms like Denodo or TIBCO to create a unified data access layer, enabling real-time data integration without physical data movement.
- Access Layer: Salesforce Connect serves as the primary interface, allowing Salesforce to access external data sources through OData-compliant endpoints, facilitating seamless integration.
- Security and Governance: Incorporates role-based access controls, data masking, and auditing mechanisms to ensure data security and compliance with regulations such as GDPR and HIPAA.
- Monitoring and Optimization Tools: Utilizes monitoring solutions to track performance metrics, identify bottlenecks, and optimize query execution across the integrated systems.

This architecture enables Salesforce to interact with external data sources in real-time, eliminating the need for data duplication and ensuring data consistency across the enterprise.

7.2 Phased Implementation Roadmap

1. Assessment:

- Data Inventory: Catalog existing data sources, identifying data critical to Salesforce operations.
- **Integration Requirements:** Determine the specific integration needs, including data access patterns, latency requirements, and compliance considerations.
- **Infrastructure Evaluation:** Assess current infrastructure capabilities to support data virtualization and real-time data access.

2. Tool Selection:

• **Data Virtualization Platform:** Choose a platform like Denodo or TIBCO that supports OData protocols and integrates seamlessly with Salesforce Connect.



- Salesforce Connect: Leverage Salesforce Connect to establish real-time connections with external data sources via OData endpoints.
- Security Tools: Implement security solutions that provide robust access controls, data masking, and auditing capabilities.

3. Pilot Implementation:

- **Prototype Development:** Develop a pilot project focusing on a specific use case, such as integrating customer data from an on-premises ERP system into Salesforce.
- **Testing and Validation:** Conduct thorough testing to validate data access, performance, and security measures.
- **Stakeholder Feedback:** Gather feedback from end-users and stakeholders to refine the integration approach.

4. Full Rollout:

- Scalability Planning: Plan for scaling the integration to additional data sources and business units.
- **Training and Documentation:** Provide comprehensive training and documentation to ensure smooth adoption across the organization.
- **Deployment:** Execute the full-scale deployment, ensuring minimal disruption to existing operations.

5. Monitoring and Optimization:

- **Performance Monitoring:** Continuously monitor system performance, focusing on query response times and data access patterns.
- **Optimization:** Implement optimizations such as query caching and load balancing to enhance performance.
- **Compliance Auditing:** Regularly audit data access and usage to ensure ongoing compliance with relevant regulations.

7.3 Team Roles and Responsibilities

• Data Architects: Design the integration architecture and oversee tool selection.



- Integration Specialists: Develop and implement the integration solutions, ensuring seamless connectivity between Salesforce and external data sources.
- Security Officers: Ensure that data security and compliance requirements are met throughout the integration process.
- **Project Managers:** Coordinate the implementation phases, manage timelines, and facilitate communication among stakeholders.
- **Business Analysts:** Gather requirements, define use cases, and validate that the integration meets business objectives.

By following this structured approach, organizations can effectively implement zero-copy integrations in hybrid cloud environments, enhancing data accessibility, reducing redundancy, and maintaining compliance with regulatory standards.

8. Technology Deep Dive: Tools and Protocols for Zero-Copy Integration

Implementing zero-copy integration within hybrid cloud environments necessitates a comprehensive understanding of the tools and protocols that facilitate seamless data access without duplication. This section delves into the key technologies and standards that underpin zero-copy integration, particularly in the context of Salesforce.

8.1 Salesforce-Specific Capabilities

Salesforce Connect with External Objects Using OData 2.0/4.0: Salesforce Connect enables real-time access to external data sources by representing them as External Objects within Salesforce. Utilizing OData 2.0 or 4.0 protocols, it allows Salesforce to interact with external systems without data replication.

Named Credentials and External Services: Named Credentials simplify the authentication process for external systems by storing endpoint and authentication details securely within Salesforce. External Services, in conjunction with Named Credentials, allow for the seamless invocation of external APIs, facilitating integration without custom code.

Custom APIs and Apex Integrations: For scenarios requiring bespoke integration logic, Salesforce provides the capability to develop custom APIs and Apex code. This approach offers flexibility in handling complex integration requirements while maintaining the zero-copy paradigm.

8.2 Data Virtualization and Federation Platforms

Denodo: Denodo offers a data virtualization platform that creates a unified data access layer, allowing real-time integration across disparate data sources. It supports OData protocols, making it compatible with Salesforce Connect for seamless data federation.



TIBCO Data Virtualization: TIBCO provides a data virtualization solution that enables the abstraction of data from multiple sources, presenting it as a single, unified view. This facilitates real-time data access and integration with Salesforce without the need for data replication¹⁶.

Red Hat JBoss Data Virtualization: Red Hat's solution offers a lean data integration platform that provides unified data access across various sources, including Salesforce, supporting real-time data consumption without duplication.

8.3 Protocols and Standards

REST, GraphQL, and OData: These protocols enable live querying of data across systems. OData, in particular, is integral to Salesforce Connect, facilitating real-time access to external data sources¹⁷.

JDBC/ODBC: These drivers are essential for connecting legacy on-premises systems to modern platforms. They allow for the integration of traditional databases into the zero-copy architecture, ensuring seamless data access.

Data Mesh Architecture: This approach promotes decentralized data ownership and governance, aligning with the principles of zero-copy integration by enabling domain-specific data management and access.

8.3 Security Protocols

OAuth: OAuth provides a secure, token-based authentication mechanism, allowing for delegated access to resources without sharing credentials, essential for secure zero-copy integrations.

TLS Encryption: Transport Layer Security (TLS) ensures that data in transit is encrypted, protecting against eavesdropping and tampering during data exchanges between systems.

SAML: Security Assertion Markup Language (SAML) enables single sign-on (SSO) and federated identity management across hybrid cloud domains, facilitating secure authentication and authorization.

8.4 Performance Optimization Tools

Smart Caching Strategies: Implementing in-memory caching for frequently accessed data can significantly reduce latency and improve performance in zero-copy integrations.

Pushdown Computation: This technique involves executing data processing operations as close to the data source as possible, reducing data movement and enhancing performance.

By leveraging these tools and protocols, organizations can effectively implement zero-copy data integration strategies within hybrid cloud environments, ensuring real-time data access, enhanced security, and optimized performance.



9. Future Outlook

Zero-copy integration has emerged as a transformative approach for organizations aiming to streamline data access and enhance operational efficiency within hybrid cloud environments. By eliminating the need for data duplication, zero-copy methodologies facilitate real-time data access, reduce storage costs, and mitigate the risks associated with data sprawl.

Salesforce's implementation of zero-copy integration, particularly through its Data Cloud, exemplifies the practical benefits of this approach. Organizations can now access and query data across multiple systems without the overhead of traditional Extract, Transform, Load processes, leading to more agile and responsive business operations.

Looking ahead, the convergence of zero-copy integration with emerging technologies such as edge computing and artificial intelligence is poised to further revolutionize data management strategies. Edge computing enables data processing closer to the data source, reducing latency and bandwidth usage, which is particularly beneficial for applications requiring immediate data analysis and response.

Moreover, the integration of AI at the edge (known as Edge AI) allows for real-time decision-making and analytics, enhancing the capabilities of zero-copy architectures. This synergy supports the development of intelligent systems capable of processing and acting upon data instantaneously, a critical requirement in sectors such as healthcare, manufacturing, and autonomous vehicles¹⁸.

In terms of compliance, zero-copy integration offers significant advantages by maintaining data within its original secure environment, thereby simplifying adherence to regulations like GDPR, HIPAA, and SOC 2. This approach minimizes the risks associated with data movement and duplication, ensuring that organizations can more effectively meet stringent compliance requirements.

In conclusion, zero-copy integration is an important advancement in data management, offering organizations a pathway to more efficient, secure, and responsive operations. As technology continues to evolve, embracing zero-copy methodologies will be essential for organizations seeking to maintain a competitive edge in an increasingly data-driven world.

References:

1. Luxner, T. (2024, October 7). Cloud computing trends: Flexera 2024 State of the Cloud Report. Flexera Blog. https://www.flexera.com/blog/finops/cloud-computing-trends-flexera-2024-state-of-the-cloud-report/

2. Deepesh, & Deepesh. (2025, January 29). Salesforce Statistics 2025: Revenue, Market Share, Customers, Employees, Acquistions& More. Cyntexa. https://cyntexa.com/blog/salesforce-statistics/

3. Data Cloud - Zero copy connectivity. (n.d.). Salesforce. https://www.salesforce.com/data/connectivity/zero-copy/

4. Banks, S. (2015, July 24). Cloud & On-Premise integration for Salesforce. Perficient Blogs. https://blogs.perficient.com/2015/07/24/cloud-on-premise-integration-for-salesforce/

5.Mponzi, P. (2022, April 13). 5 Challenges of Data Integration (ETL) and how to fix them. Datavail. https://www.datavail.com/blog/5-challenges-of-data-integration-etl-and-how-to-fix-them/



Journal of Advances in Developmental Research (IJAIDR)

E-ISSN: 0976-4844 • Website: www.ijaidr.com • Email: editor@ijaidr.com

6. Cognixia, &Cognixia. (2023, March 15). What is Zero-Copy Integration for Enterprise APIs? Cognixia. https://www.cognixia.com/blog/what-is-zero-copy-integration-for-enterprise-apis/

7. Perera, T. (2024, August 1). Zero-copy Data Integration (ZCI): towards a new era of data management. https://www.tyrell.co/2024/08/zero-copy-data-integration-zcdi-towards.html

8. Kihn, M. (2024, September 9). What is Zero copy? (and how it works). Salesforce. https://www.salesforce.com/blog/zero-copy/

9. Tackling latency issues in hybrid cloud microservices management. (2024, December 30). Java Tech Blog. https://javanexus.com/blog/tackling-latency-hybrid-cloud-microservices

10. UMATechnology. (2025, May 2). Caching Layer Optimizations in edge cloud networks monitored using Prometheus - UMA Technology. UMA Technology. https://umatechnology.org/caching-layer-optimizations-in-edge-cloud-networks-monitored-using-prometheus/

11. Best practice 15.6 – Prevent unnecessary data movement between systems and applications - Data Analytics Lens. (n.d.). https://docs.aws.amazon.com/wellarchitected/latest/analytics-lens/best-practice-15.6-prevent-unnecessary-data-movement-between-systems-and-applications..html

12. Healey, R. (2024, November 11). GDPR access Controls: Choosing the right model for your organisation. Formiti. https://formiti.com/gdpr-access-controls-which-one-is-right-for-you/

13. Convergence, I. (2023, October 25). Improving security and compliance through data integration. IT Convergence. https://www.itconvergence.com/blog/improving-security-and-compliance-through-data-integration/

14. Amazon Web Services, Inc. (2025). AWS Well-Architected Framework Data Residency and Hybrid Cloud Lens. https://docs.aws.amazon.com/pdfs/wellarchitected/latest/data-residency-hybrid-cloud-services-lens/data-residency-hybrid-cloud-services-lens.pdf

15. Orun, M., & M, T. (2025, March 21). Salesforce Data Cloud Zero copy: When (and when not) to use it. Salesforce Ben. https://www.salesforceben.com/salesforce-data-cloud-zero-copy-when-and-when-not-to-use-it/

16. TIBCO CloudTM Integration for Salesforce 1.0.0. (n.d.). https://docs.tibco.com/products/tibco-cloud-integration-for-salesforce-1-0-0

17. Keller, D. (2024, August 13). OData vs. GraphQL: When to Use Which? Five. https://five.co/blog/odata-vs-graphql/

18. The future of Cloud Computing in Edge AI - Data Center Catalog. (n.d.). Data Center Catalog - Catalog of World Data Centers. https://datacentercatalog.com/news/2025/the-future-of-cloud-computing-in-edge-ai