

Seamlessly integrating an acquisition: Building a complementary and competitive CRUD system on the core product

Kaarthikshankar Palraj *

University of Florida, USA.

World Journal of Advanced Research and Reviews, 2025, 26(01), 3448-3459

Publication history: Received on 18 March 2025; revised on 23 April 2025; accepted on 26 April 2025

Article DOI: <https://doi.org/10.30574/wjarr.2025.26.1.1452>

Abstract

Merging an acquired product into a core platform presents unique challenges when both systems share overlapping CRUD functionalities. This article details a synchronized dual-system integration strategy that maintains operational integrity across both environments while facilitating user transition to the core platform. Nowadays, there is a federated database architecture with bidirectional synchronization that preserves each system's native structure while ensuring data consistency. The implementation includes sophisticated conflict resolution mechanisms, a unified search infrastructure with de-duplication capabilities, and a security framework that harmonizes disparate permission models. The user experience engineering creates seamless interactions across interfaces through unified entry points and context-aware navigation, respecting platform-specific paradigms while enabling cross-system discovery. The resulting integration preserves valuable capabilities of both systems, minimizes user disruption, and creates natural pathways for cross-product adoption, offering insights for organizations facing similar acquisition integration scenarios.

Keywords: Dual-System Integration; CRUD Synchronization; Conflict Resolution; Federated Search; Cross-Platform User Experience

1. Introduction

Corporate acquisitions have emerged as a strategic imperative for organizations seeking competitive advantages in increasingly complex technological landscapes. These strategic maneuvers often aim to consolidate market positions, access innovative technologies, and leverage complementary capabilities to enhance product offerings. The integration process following technology acquisitions presents substantial challenges that extend beyond simple technical implementations, encompassing organizational structures, business processes, and knowledge management systems. Recent scholarly analyses have demonstrated that technological mergers and acquisitions significantly impact enterprise innovation through knowledge transfer mechanisms, organizational learning, and the reconfiguration of research and development activities, often resulting in enhanced innovative capabilities when properly executed [1]. The integration complexity increases substantially when both the acquiring and acquired entities possess overlapping functionalities—particularly in fundamental systems such as CRUD (Create, Read, Update, Delete) operations—requiring sophisticated strategies to harmonize disparate technological ecosystems.

The case study addresses this precise challenge: integrating an acquired product with parallel CRUD capabilities into a core platform while maintaining operational integrity across both systems. The conventional wisdom in post-acquisition integration often advocates for rapid system consolidation and migration; however, this approach frequently creates disruption for users, compromises data integrity, and introduces functionality gaps during transition periods. Research examining enterprise resource planning (ERP) system integrations during mergers and acquisitions has identified that maintaining parallel systems during transition periods while establishing robust data synchronization mechanisms often results in higher user satisfaction and reduced operational disruptions compared to

* Corresponding author: Kaarthikshankar Palraj.

immediate system consolidation approaches [2]. This insight informed the strategy to develop a more nuanced integration approach that preserves both systems' functionality while establishing seamless interoperability between them.

The research objectives focused on addressing three interconnected challenges inherent in dual-system integration scenarios. First, we designed a sophisticated database architecture capable of maintaining operational cohesion across both native and foreign systems without compromising the integrity of either environment. Second, we developed bidirectional synchronization mechanisms that ensure data consistency across platforms while optimizing performance and minimizing latency—a critical consideration for supporting real-time CRUD operations. Third, we engineered a user experience that facilitates natural migration toward the core platform while honoring users' established workflows within the acquired system. The literature on technology integration emphasizes that successful post-merger technology consolidation depends not only on technical compatibility but also on preserving workflow continuity for end-users, which directly influenced the approach to maintaining parallel operational capabilities [1].

The significance of the integration methodology extends well beyond immediate technical implementation concerns. By enabling users to interact fluidly with both systems, we created natural pathways for cross-product discovery and functionality adoption—a critical factor in realizing the anticipated synergies from the acquisition. This approach addresses a fundamental tension in technology acquisitions: the need to consolidate technical infrastructure while preserving the value propositions that made the acquired product successful. Studies of enterprise system integrations during mergers have established that the preservation of familiar interfaces and workflows during transition periods correlates strongly with user retention and satisfaction metrics, ultimately affecting the overall success of the acquisition integration process [2]. Furthermore, the architectural patterns developed during this integration effort provide valuable precedents for future acquisition scenarios involving similar functionality overlaps, potentially offering acceleration opportunities for subsequent integration initiatives while minimizing technical debt accumulation.

Table 1 Comparison of Integration Approaches for Acquisition CRUD Systems. [1, 2]

Integration Approach	Benefits	Challenges	Suitability
Complete System Migration	<ul style="list-style-type: none"> • Single unified system • Long-term maintenance simplicity • Unified user experience 	<ul style="list-style-type: none"> • High initial disruption • Risk of functionality loss • Extended transition period 	Best for systems with minimal unique functionality in the acquired system
Parallel Systems with No Integration	<ul style="list-style-type: none"> • No technical complexity • Preservation of all original functionality • Minimal initial disruption 	<ul style="list-style-type: none"> • Fragmented user experience • Duplicated data management • No cross-system functionality 	Temporary solution only; not sustainable.
API-Based Integration	<ul style="list-style-type: none"> • Moderate technical complexity • Preserves system independence • Flexible implementation timeline 	<ul style="list-style-type: none"> • Potential performance bottlenecks • Complex error handling • Limited real-time consistency 	Suitable for loosely coupled systems with infrequent data exchange
Synchronized Dual-System	<ul style="list-style-type: none"> • Seamless user experience • Data consistency across systems • Gradual transition path 	<ul style="list-style-type: none"> • Complex implementation • Sophisticated conflict resolution required • Ongoing synchronization maintenance 	Optimal for systems with overlapping but distinct value propositions

2. System architecture design

Designing a system architecture that accommodates dual-system integration presents unique challenges that extend beyond conventional database design principles. The integration of an acquired CRUD system with an established core

platform requires careful consideration of data models, synchronization mechanisms, conflict management, and overall system reliability. This section details the approach to these architectural challenges and the solutions implemented to achieve seamless integration.

2.1. Database Structure for Dual-System Integration

The architectural approach began with a comprehensive analysis of both the core and acquired systems' data models to identify commonalities, discrepancies, and integration opportunities. Rather than forcing a complete structural alignment—which often leads to compromises in both systems—we implemented a federation strategy that maintains the native schema integrity of each system while establishing robust mapping layers between them. This approach aligns with distributed database architectural principles, where data is stored across multiple locations while providing a unified access mechanism, enabling both systems to maintain their original structure while functioning as an integrated whole [3]. The core database was extended with dedicated tables to represent the acquired product's objects, incorporating all necessary fields while implementing additional supporting structures for access controls and organizational hierarchies.

The mapping layer employed bidirectional schema transformations to translate between the two systems' data representations. This included handling differences in data types, field names, relationship cardinality, and constraint models. We implemented a metadata repository to maintain these mappings and transformation rules, allowing for centralized management and evolution of the integration logic as both systems continued to develop independently. The metadata-driven approach provided flexibility to accommodate future changes without requiring extensive recoding of integration components, a critical consideration given the ongoing development in both systems. This strategy aligns with post-merger integration best practices that emphasize the importance of building adaptable systems that can evolve alongside changing business requirements during the consolidation process [4].

2.2. Data Synchronization Pipeline Implementation

The synchronization pipeline represents the operational core of the dual-system integration. We implemented an event-driven architecture that captures change events from both systems and propagates them bidirectionally in near real-time. Each system was equipped with change data capture (CDC) mechanisms that publish events to a centralized message broker when relevant data modifications occur. These events follow a standardized format containing the operation type, entity identifier, modified fields, timestamp, and originating system identifier—essential context for proper synchronization processing.

The synchronization service consumes these events, applies the appropriate schema transformations using the mapping layer, and executes corresponding operations on the target system. This process is transaction-aware, ensuring that complex operations spanning multiple entities maintain their atomicity when synchronized across systems. To optimize performance and network utilization, we implemented batching for high-volume change scenarios while maintaining near real-time propagation for typical operational loads. This approach leverages core principles of distributed database systems, which emphasize the importance of efficient consensus protocols and transaction management to maintain consistency across geographically dispersed data stores [3]. The implementation incorporated these principles through configurable synchronization parameters that balance performance requirements with consistency guarantees.

2.3. Conflict Resolution Mechanisms and Priority Hierarchies

Conflict resolution represents one of the most challenging aspects of bidirectional synchronization systems. When concurrent modifications occur to the same entity across systems, sophisticated resolution strategies are required to maintain data consistency while preserving user intent. The architecture implements a multi-tiered conflict resolution approach that combines automated resolution rules with selective manual intervention for complex scenarios.

The primary resolution strategy employs a field-level timestamp comparison that identifies the most recent update for each modified field rather than applying a simplistic "last writer wins" policy at the entity level. This granular approach preserves non-conflicting changes while intelligently resolving conflicts for specific fields. For scenarios where timestamp-based resolution is insufficient, we implemented a priority hierarchy that considers the system of origin, user roles, and operation context. The core system generally received higher priority for conflict resolution, but this could be overridden based on contextual factors such as the user's primary system affiliation or specific workflow states. This nuanced approach to conflict management reflects recommendations from post-merger integration frameworks that emphasize the importance of establishing clear data governance policies that account for operational realities while moving toward unified systems [4].

For complex conflicts that cannot be automatically resolved, the architecture includes an intervention mechanism that flags the conflict for manual resolution while maintaining system operation. Affected entities enter a specialized state that allows continued read access while restricting further modifications until the conflict is resolved. This approach balances system availability with data integrity requirements, a critical consideration in enterprise environments where system downtime must be minimized.

2.4. Technical Considerations for Maintaining System Reliability

Ensuring system reliability in a dual-database architecture requires comprehensive monitoring, fault tolerance, and recovery mechanisms. The implementation includes a dedicated monitoring subsystem that tracks synchronization health metrics, including event propagation latency, queue depths, error rates, and conflict occurrences. These metrics feed into both real-time alerting systems and historical analysis tools that help identify patterns and optimization opportunities.

The synchronization pipeline incorporates multiple reliability features, including idempotent operation handlers that prevent duplicate applications of changes, checkpoint mechanisms that enable recovery from interruptions, and circuit breakers that prevent cascading failures when subsystems experience issues. We implemented a comprehensive logging system that records all synchronization activities with sufficient detail to support both troubleshooting and audit requirements. These resilience patterns draw from established distributed database design principles, which emphasize that system reliability depends on the ability to handle partial failures gracefully while maintaining data consistency guarantees—a significant challenge in geographically distributed systems with independent failure domains [3].

Database consistency checking utilities run at scheduled intervals to verify synchronization effectiveness and detect any discrepancies that might have occurred despite the real-time synchronization mechanisms. When discrepancies are identified, the system can initiate corrective synchronization operations or flag entities for manual review, depending on the nature and severity of the inconsistency. This defense-in-depth approach creates multiple safeguards against data drift between systems, addressing one of the primary technical challenges in post-merger integrations. Research on system integration following acquisitions highlights that data inconsistencies represent one of the most common and disruptive issues during transition periods, often leading to operational disruptions and diminished user trust if not properly managed [4].

3. User experience engineering

The successful integration of an acquired product extends beyond technical database synchronization to encompass thoughtful user experience design that supports seamless interaction across both systems. This section explores the approach to engineering a cohesive user experience while acknowledging the distinct interface paradigms of each platform. The user experience strategy focused on minimizing cognitive load for users while enabling a gradual, organic transition toward the core platform.

3.1. Dual Read/Write Capability Across Interfaces

A fundamental requirement for the integration was enabling bidirectional interaction with data regardless of which interface users chose to employ. This dual read/write capability presented significant engineering challenges, as each system had developed distinct interaction patterns, validation rules, and data entry workflows. Rather than forcing conformity to a single model, we implemented an adaptive approach that preserved each system's native interaction paradigms while ensuring data integrity across both environments.

The implementation required developing comprehensive API layers for both systems that could translate interface-specific actions into standardized operations on the underlying data. These APIs implemented validation logic appropriate to each interface while ensuring that the resulting data met the requirements of both systems. This approach aligns with established integration patterns, such as the API-led connectivity model, which emphasizes creating experience APIs tailored to specific user interfaces while maintaining process and system APIs that handle standardized business logic and data access [5]. User testing revealed that maintaining interface-specific validation and interaction patterns significantly reduced error rates compared to enforcing a unified interaction model across disparate interfaces.

Each interface was enhanced with subtle visual indicators that provided awareness of the cross-system nature of the data without disrupting the native experience. Records originating from the alternate system were marked with unobtrusive badges, and tooltips provided context about cross-system behaviors when relevant. These awareness

features helped users build accurate mental models of the integrated system while minimizing confusion about data origin and modification capabilities. Research on cross-platform application development highlights the importance of maintaining platform-specific UI components and interaction patterns while ensuring data consistency across environments, particularly when users might switch between platforms depending on their context and needs [6].

3.2. Unified Home View Development Strategy

The home view represents the primary entry point for users and plays a critical role in shaping their perception of the integrated system. We developed a unified home view in the core platform that presented both native and foreign records in a cohesive manner, enabling users to browse, search, and interact with content regardless of origin. This unified view served as a bridge between systems, encouraging cross-system discovery while providing a consistent starting point for all workflows.

The design of this unified view required careful consideration of information architecture to accommodate the distinct organizational principles of both systems. We implemented a flexible taxonomy system that mapped the organizational structures of each platform, enabling records to appear in appropriate locations regardless of their origin. This approach allowed us to preserve the hierarchical organization familiar to users of each system while creating a coherent overall structure. The implementation leveraged the content-based router integration pattern, which directs content to appropriate destinations based on message content rather than predetermined routes, enabling the dynamic organization of heterogeneous data sources into a coherent presentation layer [5].

The implementation included intelligent default sorting and filtering options that prioritized relevance across the combined dataset while offering system-specific views for users who preferred to focus on a particular subset of content. These options were personalized based on usage patterns, gradually adapting to each user's preferences to optimize their experience. The unified view also incorporated contextual actions that respected the capabilities of each record's native system, ensuring that users were only presented with operations that could be reliably executed on the selected content.

3.3. Interface Consistency with System-Specific Adaptations

Maintaining consistency across integrated interfaces while honoring system-specific design patterns required a nuanced approach to user interface engineering. We developed a shared design system that established common patterns for core interactions while allowing for system-specific implementations that respected each platform's established visual language and interaction paradigms. This approach created a sense of cohesion without forcing artificial uniformity that would disrupt user expectations.

The shared design system defined consistent patterns for primary actions, navigation structures, and data presentation, creating a foundation of familiarity that users could transfer between systems. These consistent elements were implemented with visual treatments appropriate to each platform, maintaining their native aesthetic while creating recognizable behavioral patterns. System-specific adaptations were primarily focused on specialized functionality unique to each platform, allowing these features to maintain their original interaction patterns without compromise. This balanced approach mirrors best practices in cross-platform application development, which emphasize maintaining platform-native UI components and interaction paradigms while sharing business logic and data models across implementations to ensure functional consistency without compromising platform-specific user experiences [6].

To support this strategy, we implemented a comprehensive component library that encapsulated the shared interaction patterns while offering system-specific rendering options. This infrastructure enabled development teams to implement features with consistent behavior while respecting visual and interaction differences between platforms. The component architecture included robust context awareness that allowed components to adapt their behavior based on the hosting system, user permissions, and data origin, ensuring appropriate functionality regardless of context.

3.4. User Journey Mapping Through Integrated Systems

Understanding and optimizing user journeys across the integrated systems was essential for delivering a coherent experience. We conducted extensive journey mapping exercises to identify common workflows that spanned both platforms and designed integration points that facilitated smooth transitions when users needed to move between systems. These optimized pathways reduced friction in cross-system workflows and helped users develop accurate mental models of the integrated environment.

The journey mapping process involved analyzing usage data from both systems to identify common patterns and pain points in existing workflows. This analysis revealed several key journeys that frequently spanned both systems, particularly in collaborative workflows where different user roles had established preferences for specific interfaces. For each identified journey, we designed and implemented integration points that created smooth handoffs between systems, including deep linking capabilities, context preservation, and appropriate data transformations. This approach leverages the process orchestration integration pattern, which defines the sequence of steps needed to complete a business process that spans multiple systems while maintaining transaction integrity and context throughout the process [5].

To support users in navigating these integrated journeys, we developed contextual guidance features that offered assistance at key decision points. These features included subtle suggestions for alternate workflows, explanations of cross-system implications, and dynamic help content that was adapted based on the user's current context and history. Rather than forcing users into prescribed paths, these guidance elements supported informed decision-making while allowing flexibility in how users approached their tasks. This philosophy aligns with research on cross-platform user experience design, which suggests that effective cross-platform experiences should prioritize usability and context-appropriate interactions over strict visual consistency, recognizing that different platforms may serve different aspects of the user's overall journey with the product [6].

4. Search and Indexing Infrastructure

The integration of an acquired product into a core platform necessitated the development of a sophisticated search and indexing infrastructure to provide users with seamless discovery capabilities across both systems. Search functionality represents a critical entry point for users accessing content, making it essential that results appear unified and relevant regardless of which system houses the underlying data. This section details the approach to unified search, de-duplication algorithms, performance optimization, and scalability considerations.

4.1. Unified Search System Architecture

Creating a unified search experience across disparate data sources required an architecture that could harmonize different data models, access patterns, and relevance signals into a cohesive index and query system. Rather than attempting to force conformity between the two systems' native search capabilities, we implemented a federated search architecture with a central orchestration layer that coordinates queries across both platforms while normalizing and integrating results.

The architecture consists of three primary components: index adapters that transform system-specific data into a standardized search representation, a central search index that stores and updates these representations, and a query processor that distributes search requests and consolidates results. Each system maintains its specialized index structures optimized for its particular data model, while the central orchestration layer handles the translation between these specialized indexes and the unified search interface. This approach aligns with research on modern information integration systems, which emphasizes federated query processing over disparate data sources as an effective strategy when complete data warehouse consolidation is impractical or undesirable [7]. The index adapters implement incremental synchronization mechanisms that ensure the central search index remains current with both underlying systems, leveraging the change data capture mechanisms developed for the core data synchronization pipeline.

The query processor employs a sophisticated routing algorithm that analyzes each query to determine which underlying indexes should be consulted. This analysis considers query context, user permissions, and historical performance data to optimize the distribution of search requests. The processor normalizes relevance scores across different sources using a calibrated scoring model that accounts for system-specific biases, ensuring that results are ordered by true relevance rather than source system. This approach reflects established practices in distributed query processing for heterogeneous data stream systems, which emphasize the importance of adaptive query routing based on data characteristics and processing capabilities of constituent systems [8].

4.2. De-duplication Algorithm Development

A significant challenge in maintaining dual systems is the potential for duplicate or near-duplicate records to appear in search results when similar content exists on both platforms. We developed a multi-stage de-duplication algorithm that identifies and consolidates duplicate records at both indexing and query time, ensuring that users receive a clean, non-redundant result set regardless of underlying data duplication.

The algorithm employs a hybrid approach combining signature-based and similarity-based techniques. During indexing, each record is assigned a content signature derived from its critical fields, enabling quick identification of exact duplicates. For near-duplicates, we implemented a similarity detection process that uses term frequency-inverse document frequency (TF-IDF) vectors to calculate content similarity across potentially matching records. Records exceeding a configurable similarity threshold are flagged as potential duplicates and linked in the index. This approach allows for efficient de-duplication without requiring exact field-level matches, accommodating the reality that equivalent records might have minor variations across systems. The hybrid approach draws inspiration from research on entity resolution in federated information systems, which demonstrates that combining multiple detection techniques produces more robust results across heterogeneous data models than any single technique in isolation [7].

At query time, the system performs a final de-duplication pass on results before presentation to the user. When duplicates are identified, the system selects a canonical record to display based on a prioritization algorithm that considers the system of origin, record completeness, update recency, and user context. This dynamic selection ensures that users see the most relevant version of a record based on their specific context and needs rather than applying a static prioritization rule. Importantly, the system maintains awareness of all duplicate instances, allowing users to access alternate versions when necessary while preventing redundancy in the primary result presentation.

4.3. Performance Optimization for Cross-Database Queries

Maintaining responsive search performance across dual databases presented significant technical challenges, particularly given the need to coordinate queries, normalize results, and perform de-duplication within interactive timeframes. We implemented multiple optimization strategies to ensure that the unified search-maintained performance characteristics comparable to the native search capabilities of each individual system.

Query optimization began with the development of a predictive query classifier that analyzes incoming search requests to identify the most efficient execution strategy. This classifier considers query complexity, predicted result size, and historical performance data to determine whether to execute parallel or sequential searches across the underlying systems. For queries likely to return large result sets, the system implements an incremental fetching strategy that retrieves, and processes results in batches, allowing for rapid initial response while continuing to refine results as the user interacts with the search interface. These adaptive execution strategies align with current research on distributed query optimization, which demonstrates that selective query decomposition based on predicted execution costs outperforms static distribution strategies, particularly in environments with heterogeneous data systems that may have significantly different performance characteristics for similar queries [8].

We implemented a multi-level caching architecture to reduce redundant processing for common search patterns. This architecture includes result caches that store processed search results for frequent queries, fragment caches that preserve partial processing results for query components, and metadata caches that accelerate permission checks and de-duplication operations. Cache invalidation is tied to the change data capture mechanisms of both systems, ensuring that cached results remain consistent with the underlying data. The caching strategy is adaptive, automatically adjusting cache parameters based on observed usage patterns and system load to optimize memory utilization while maximizing cache effectiveness.

At the database level, we implemented specialized indexes optimized for the most common search access patterns identified through usage analysis. These indexes include composite structures that support efficient filtering across multiple dimensions, covering indexes that eliminate the need for table lookups on common result fields and partial indexes that focus on frequently queried subsets of the data. The index design strategy balanced query performance against index maintenance costs and storage requirements, guided by a comprehensive analysis of actual search behavior rather than theoretical access patterns. This approach reflects findings from research on federated query processing, which emphasizes the importance of adaptive indexing strategies that consider both the structural characteristics of the data and the observed query workload [7].

4.4. Scalability Testing and Implementation

Ensuring that the search infrastructure could scale to accommodate growing data volumes and user loads was critical for long-term sustainability. We developed a comprehensive testing methodology to evaluate scalability across multiple dimensions and implemented architectural patterns that support horizontal scaling as demands increase.

The testing methodology included synthetic load testing that simulated projected user growth, data volume testing that assessed performance impact as content repositories expanded, and complexity testing that evaluated system behavior with increasingly sophisticated search queries. These tests were automated and integrated into the continuous

integration pipeline, providing ongoing visibility into scalability characteristics as the system evolved. This methodology draws from established practices in distributed system evaluation, which emphasize the importance of multi-dimensional performance assessment under varying workloads to identify potential bottlenecks and scaling limitations before it impacts production environments [8].

Based on testing results, we implemented several architectural enhancements to improve scalability. The index structure was partitioned both horizontally (sharding data across multiple nodes) and vertically (separating different types of index information), allowing for targeted scaling of specific components as needed. The query processing layer was designed to scale horizontally, with stateless processors that can be added or removed dynamically based on demand. Load-balancing algorithms distribute requests across processors based on current utilization and query characteristics, ensuring efficient resource use under varying load conditions.

To support scaling beyond initial projections, we implemented a monitoring and alerting system that tracks key performance indicators and identifies potential bottlenecks before it impacts user experience. This system collects detailed telemetry data on query patterns, system resource utilization, and response time distributions, enabling both reactive scaling in response to current demands and proactive capacity planning based on observed trends. Research on distributed data stream systems highlights the importance of continuous monitoring and adaptation in maintaining performance as scale increases, particularly when workloads may change unpredictably over time [8]. The monitoring subsystem provides not only operational alerts but also valuable insights into usage patterns that inform ongoing optimization and feature development.

5. Security and Data Integrity Framework

Integrating an acquired product with an existing platform introduces significant security and data integrity challenges that must be addressed comprehensively. The dual-system nature of the integration creates potential vulnerabilities at intersection points while requiring consistent security controls across environments with different underlying architectures. This section explores the approach to developing a unified security and data integrity framework that maintains robust protection across both systems while enabling seamless user interactions.

5.1. Cross-System Permissions Architecture

Designing a permissions architecture that works effectively across integrated systems required harmonizing disparate security models without compromising the protection mechanisms of either platform. Rather than attempting to merge these models into a single unified structure—which would have required fundamental changes to both systems—we developed a permissions translation layer that maintains the native security model of each system while ensuring consistent access control across the integrated environment.

The permission architecture consists of three primary components: a unified permission model that defines the superset of access controls across both systems, system-specific adapters that translate between the unified model and native security structures, and a centralized policy enforcement point that ensures consistent application of security rules. The unified model establishes core permission concepts such as ownership, access levels, and inheritance rules, mapping these to the corresponding constructs in each system's native security model. This approach draws from established principles of access control in distributed systems, which emphasize the importance of creating interoperable security domains that can maintain consistent protection mechanisms while accommodating heterogeneous implementation details [9]. The translation layer maintains bidirectional mappings between permission structures, ensuring that changes in either system propagate appropriately to maintain consistent access controls.

A central challenge in implementing this architecture was handling permissions for entities that exist in both systems, particularly when user access might differ between platforms. We addressed this through a principle of most restrictive reconciliation, where the effective permissions for a user accessing an entity are determined by the most restrictive set of permissions across both systems. This conservative approach ensures that security is never inadvertently weakened through the integration process, though it necessitates careful attention to permission synchronization to prevent unnecessary access restrictions. This approach aligns with enterprise data governance frameworks that emphasize the principle of least privilege as a cornerstone of effective access control, particularly in environments where multiple security domains must interact while maintaining their individual integrity [10].

The architecture includes specialized handling for administrative operations that might have system-wide implications, implementing additional verification steps, and audit logging for these high-impact actions. This multi-layered approach

to permission management enables seamless user interactions while maintaining robust security controls appropriate to each system's native environment.

Table 2 Security Framework Components Across Integration Boundaries. [9, 10]

Component	Primary Function	Technical Implementation	Integration Consideration
Index Adapters	Transform system-specific data for a unified index	System-specific ETL pipelines with standardized output format	Must handle differing data models and field semantics
Central Search Index	Stores unified representation of all content	Distributed inverted index with entity-relationship awareness	Optimized for query patterns spanning both systems
Query Processor	Distributes and consolidates search requests	Adaptive routing engine with query classification capabilities	Balances completeness against performance.
Relevance Normalizer	Ensures consistent result ordering	Calibrated scoring model with system-specific adjustments	Prevents bias toward either system in results
De-duplication Engine	Eliminates redundant results	Hybrid signature and similarity-based detection with prioritization rules	Maintains awareness of all instances while presenting the canonical version

6. Foldering Structure for Integrated Content Management

Developing an integrated content management approach that supports intuitive organization across both systems presented significant challenges, particularly given the different organizational paradigms employed by each platform. We implemented a hierarchical folding system with bidirectional synchronization that enables consistent content organization while respecting the native structures of each system.

The folding architecture implements a unified organizational model that defines core concepts such as containers, hierarchies, and membership rules. This model is mapped to the native organizational structures of each system through dedicated adapters that translate between the unified model and system-specific implementations. For the core system, this mapping was relatively straightforward, as it already employed a hierarchical folder structure. The acquired system used a combination of tags and collections, requiring more complex transformations to maintain conceptual integrity across the integration boundary. This approach to organizational structure integration draws from research on access control models for distributed file systems, which demonstrates that capability-based models can effectively bridge disparate organizational structures while maintaining consistent security properties [9].

To support cross-system content organization, we implemented bidirectional synchronization mechanisms that propagate organizational changes between systems. When content is moved within the folder hierarchy in either system, corresponding updates are applied in the companion system to maintain consistent organization. This synchronization includes special handling for system-specific organizational features that may not have direct equivalents, such as mapping tag collections to virtual folders that preserve the semantic grouping without forcing an exact structural match. The synchronization approach follows the principle of semantic equivalence rather than structural identity, ensuring that the organizational intent is preserved across systems even when the exact implementation differs.

The folding system includes comprehensive conflict resolution mechanisms to handle scenarios where organizational changes conflict across systems. These mechanisms employ a combination of rule-based automated resolution for common conflicts and administrative intervention capabilities for complex scenarios, ensuring that organizational integrity is maintained even in edge cases. This conflict resolution approach aligns with enterprise data governance frameworks that emphasize the importance of establishing clear stewardship responsibilities and escalation procedures for resolving structural conflicts, particularly in environments where multiple organizational paradigms must coexist [10].

6.1. Real-Time Monitoring and Alerting Systems

Maintaining visibility into the health and security status of an integrated dual-system environment requires comprehensive monitoring capabilities that span both platforms while providing unified insights. We implemented a multi-layered monitoring architecture that collects telemetry from all components of the integrated system, performs correlation and analysis to identify potential issues, and generates appropriate alerts when intervention is needed.

The monitoring infrastructure collects data from multiple sources across both systems, including application logs, database metrics, synchronization events, security audit trails, and user interaction patterns. This diverse telemetry is normalized into a consistent format and stored in a centralized monitoring database that supports both real-time analysis and historical trend evaluation. The collection architecture implements adaptive sampling that increases data granularity when anomalies are detected, enabling detailed investigation without generating excessive data volume during normal operations. This approach is consistent with capability-based security models for distributed systems, which emphasize the importance of comprehensive audit trails that can track capability propagation and usage across security domains while maintaining manageable operational overhead [9].

Real-time analysis components process the collected telemetry to identify potential issues through a combination of threshold-based rules, anomaly detection algorithms, and pattern recognition techniques. These components are organized in a hierarchical structure, with low-level components focusing on specific subsystems while higher-level components perform cross-system correlation to identify complex issues that span integration boundaries. The analysis architecture includes specialized detectors for security-relevant events, synchronization anomalies, performance degradation, and data integrity issues, providing comprehensive coverage of potential concerns.

When potential issues are identified, the alerting subsystem generates appropriate notifications based on the nature, severity, and context of the detected condition. Alerts are routed to appropriate responders based on configurable rules that consider the affected systems, issue category, business impact, and time of day. High-priority alerts trigger immediate notification through multiple channels, while lower-priority issues are grouped into digests to prevent alert fatigue. This context-aware approach to alert management reflects enterprise data governance principles that emphasize the importance of establishing clear incident response procedures with defined ownership and escalation paths, ensuring that the appropriate stakeholders are engaged based on the nature and severity of detected issues [10].

The monitoring system includes a comprehensive dashboard that provides visibility into system health, security status, and operational metrics across both platforms. This dashboard supports both operational monitoring and executive reporting, with different views tailored to the needs of various stakeholders. The integrated visibility enables proactive management of the dual-system environment, supporting both day-to-day operations and strategic planning.

6.2. Data Synchronization Verification Protocols

Ensuring data integrity across synchronized systems requires robust verification mechanisms that can detect and resolve inconsistencies before they impact users. We implemented a comprehensive verification framework that combines continuous monitoring, scheduled validation, and on-demand reconciliation to maintain data consistency across both platforms.

The verification architecture implements multiple validation strategies operating at different levels of the system. At the transaction level, each synchronization operation includes integrity checks that verify successful completion and data consistency immediately after execution. These checks employ cryptographic hashes and record counts to provide efficient verification without requiring full data comparison. At the entity level, periodic verification processes compare key fields across corresponding records in both systems to identify any inconsistencies that might have developed despite transaction-level validation. At the system level, comprehensive integrity scans periodically validate the overall consistency of the databases, identifying any orphaned or misaligned records that require attention. This multi-layered approach to integrity verification draws from capability-based protection principles for distributed systems, which emphasize the importance of establishing verifiable security properties at multiple levels of abstraction to create defense-in-depth against both accidental corruption and malicious manipulation [9].

When verification processes identify inconsistencies, the system employs a graduated response strategy based on the nature and severity of the detected issues. Minor inconsistencies that don't affect user functionality are logged and scheduled for background reconciliation during low-usage periods. More significant issues that might impact users generate alerts for administrative review, with suggested resolution actions based on the specific inconsistency pattern. Critical inconsistencies that could lead to data loss or security vulnerabilities trigger immediate reconciliation with appropriate safeguards to prevent cascading issues. This tiered response strategy aligns with enterprise data

governance frameworks that recommend implementing graduated remediation processes that balance the importance of data integrity against operational continuity requirements, particularly in environments where immediate reconciliation might disrupt critical business processes [10].

The verification framework includes comprehensive logging and auditing capabilities that maintain detailed records of all verification activities, detected inconsistencies, and resolution actions. These records support both troubleshooting of specific issues and trend analysis to identify recurring patterns that might indicate underlying systemic problems. The audit trail maintains cryptographic integrity protections to ensure that verification records themselves remain trustworthy, supporting both operational needs and compliance requirements.

To minimize the performance impact of verification activities on production systems, the architecture implements resource-aware scheduling that adjusts verification intensity based on system load and business priority. Verification processes automatically reduce resource consumption during peak usage periods, while comprehensive verifications are scheduled during maintenance windows to minimize user impact. This adaptive approach balances the need for thorough verification with the operational requirement to maintain system performance and availability, reflecting governance principles that emphasize the need to balance security controls against business functionality requirements to achieve appropriate risk management outcomes [10].

7. Conclusion

The successful integration of an acquired product with overlapping CRUD functionality demonstrates that preserving dual systems can yield superior outcomes compared to forced consolidation. By implementing a federation strategy with robust synchronization, the system-maintained integrity while enabling seamless user interaction across platforms. The architecture's resilience against partial failures, combined with comprehensive monitoring and verification protocols, ensured data consistency without compromising performance. The user-centric design approach preserved established workflows while creating natural migration paths, resulting in high retention rates and increasing adoption of core platform features. The security and governance frameworks established clear protocols that maintained protection across system boundaries without introducing unnecessary restrictions. Beyond immediate technical benefits, this integration created opportunities for feature cross-pollination between platforms, enriching both environments through shared innovation. The patterns developed during this effort provide a blueprint for future acquisition scenarios, potentially reducing integration timelines and minimizing technical debt. As organizations continue to pursue strategic acquisitions, the ability to integrate disparate systems while preserving their unique value propositions will remain crucial to maximizing return on investment and delivering cohesive user experiences.

References

- [1] Lingling Suo et al., "The Impact of Technological Mergers and Acquisitions on Enterprise Innovation: A Review," Sustainability, 2023. [Online]. Available: https://www.researchgate.net/publication/373406671_The_Impact_of_Technological_Mergers_and_Acquisitions_on_Enterprise_Innovation_A_Review
- [2] Michael David Myers, von Vangerow, Andreas, "Integration of Different ERP Systems: The Case of Mergers and Acquisitions,," DBLP, 2011. [Online]. Available: https://www.researchgate.net/publication/221229380_Integration_of_Different_ERP_Systems_The_Case_of_Mergers_and_Acquisitions
- [3] Charlie Custer, "What is a distributed database, and how do they work?" Cockroach Labs, 2023. [Online]. Available: <https://www.cockroachlabs.com/blog/what-is-a-distributed-database/>
- [4] Shelley Bounague, "Navigating Post Merger Integration Challenges," Cloudficient, 2024. [Online]. Available: <https://www.cloudficient.com/blog/navigating-post-merger-integration-challenges>
- [5] Abhishek Gaurav, "Top 10 integration patterns for enterprise use cases," MuleSoft Blog, 2021. [Online]. Available: <https://blogs.mulesoft.com/api-integration/patterns/top-10-integration-patterns/>
- [6] Szymon Nitecki, "Cross Platform Mobile App Development – Pros and Cons," Netguru Blog, 2024. [Online]. Available: <https://www.netguru.com/blog/cross-platform-mobile-apps-development>
- [7] Alexandros Labrinidis et al., "Challenges and opportunities with big data," Proceedings of the VLDB Endowment, 2012. [Online]. Available: <https://dl.acm.org/doi/10.14778/2367502.2367572>

- [8] Niko Pollner, "Query Optimization in Distributed Heterogeneous Data Stream Systems," ResearchGate, 2021. [Online]. Available: https://www.researchgate.net/publication/355119069_Query_Optimization_in_Distributed_Heterogeneous_Data_Stream_Systems
- [9] Anita K. Jones, M. Satyanarayanan "Integrating security in a large distributed system," ACM Transactions on Computer Systems (TOCS), 1989. [Online]. Available: <https://dl.acm.org/doi/10.1145/65000.65002>
- [10] Deepa Madhavan, "Enterprise Data Governance: A Comprehensive Framework for Ensuring Data Integrity, Security, and Compliance in Modern Organizations," International Journal of Scientific Research in Computer Science Engineering and Information Technology, 2024. [Online]. Available: https://www.researchgate.net/publication/385146465_Enterprise_Data_Governance_A_Comprehensive_Framework_for_Ensuring_Data_Integrity_Security_and_Compliance_in_Modern_Organizations