

# A Conceptual Model for ETL Design and Data Integration in Salesforce-Centric Enterprise Architectures

OLANIYI BADMUS<sup>1</sup>, ADETOMIWA A. DOSUNMU<sup>2</sup>, DAVID EXCEL OZOWARA<sup>3</sup>

<sup>1</sup>Accenture, Australia

<sup>2</sup>Adbirt Nigeria, Lagos, Nigeria

<sup>3</sup>Western Illinois University, Macomb, Illinois, USA

*Abstract- Enterprise organizations increasingly position Salesforce as a central engagement and data management hub in complex multi-system architectures, creating a need for rigorous approaches to extract, transform, and load (ETL) design that account for the distinctive structural properties of the Salesforce data model. Traditional ETL frameworks were designed for relational data warehouse environments and do not directly address the object-relationship model, governor limits, bulk API constraints, and metadata-driven customization capabilities that characterize Salesforce as a data integration target. This paper proposes a conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures. The model establishes five conceptual dimensions: source data profiling and quality assessment, transformation logic design for Salesforce object alignment, bulk API utilization strategy, data lineage and audit trail maintenance, and ongoing data quality governance. Each dimension is articulated through design principles, decision criteria, and illustrative patterns drawn from enterprise data integration practice. The model addresses both inbound data flows, in which external source data is loaded into Salesforce, and bidirectional synchronization scenarios, in which Salesforce data participates in enterprise-wide data exchange with downstream analytical, operational, and clinical systems. The conceptual model is positioned as a bridge between established data warehousing and ETL literature and the Salesforce-specific implementation guidance that currently dominates practitioner discourse.*

**Keywords:** ETL Design, Salesforce Data Integration, Bulk API, Data Quality Governance, Enterprise Data Architecture, Data Lineage, CRM Data Management

## I. INTRODUCTION

Data integration is among the most consequential and enduringly challenging activities in enterprise information technology management. As organizations accumulate diverse operational systems,

analytical platforms, and engagement tools, the movement of data across system boundaries becomes a critical determinant of both operational efficiency and strategic analytical capability (Inmon, 2005; Kimball & Ross, 2013; Kleppmann, 2017). In organizations that have adopted Salesforce as their primary customer engagement platform, data integration presents a distinctive set of design challenges rooted in the structural characteristics of the Salesforce platform: its object-relationship data model, its API-mediated access paradigm, its governor limits on transaction volume and computation, and its metadata-driven customization model that may vary significantly between Salesforce orgs.

The extract, transform, and load (ETL) discipline has produced mature frameworks and methodologies for addressing data integration challenges in relational database and data warehouse contexts (Vassiliadis, 2009; Simitsis & Vassiliadis, 2008; Golfarelli et al., 1998). These frameworks address the full ETL lifecycle: source data profiling, transformation logic specification, load strategy optimization, error handling, and operational monitoring. However, they were developed in an era when data integration targets were principally relational databases, and their design assumptions do not readily translate to the Salesforce context. The constraints of the Salesforce Bulk API, the relationship between standard and custom objects, and the implications of Salesforce governor limits for ETL process design all require new conceptual treatment.

This gap between the established ETL literature and the practical demands of Salesforce data integration has produced a practitioner discourse that is rich in tool-specific guidance but weak in conceptual foundation. Organizations implementing Salesforce

integrations at enterprise scale frequently encounter data quality failures, performance degradation, and governance gaps that could be mitigated by a more principled approach to integration design. This paper addresses this need by proposing a conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures, grounded in the established ETL and data warehousing literature and adapted to the specific structural properties of the Salesforce platform (Kumar & Reinartz, 2018; Greenberg, 2010).

## II. STRUCTURAL PROPERTIES OF SALESFORCE AS A DATA INTEGRATION TARGET

### 2.1 The Salesforce Object-Relationship Model

Salesforce organizes data in a structured hierarchy of standard and custom objects, each of which is an analogue to a relational database table but with additional metadata-driven properties governing field-level security, sharing rules, validation logic, and record type differentiation (Buttle & Maklan, 2019; Greenberg, 2010). Relationships between objects are implemented as lookup fields and master-detail relationships, the latter of which enforces cascading deletion and ownership inheritance. For ETL processes, these relationship structures impose a loading sequence constraint: parent records must exist before child records can reference them. This constraint is trivial in small-volume contexts but becomes a significant design consideration in bulk loading scenarios involving millions of records across multiple related object types.

The Salesforce standard object model includes well-defined structures for accounts, contacts, leads, opportunities, cases, and activities. Custom objects extend this model to support industry-specific or organization-specific data requirements. In Salesforce Health Cloud and Financial Services Cloud deployments, the standard object model is augmented by industry-specific data models that introduce additional complexity to integration design. Understanding the full object model of a target Salesforce org, including all customizations, is therefore a prerequisite for effective ETL design, and this understanding must be maintained as the org

evolves through ongoing development (Linstedt & Olschimke, 2015; Redman, 2008).

### 2.2 API Architecture and Governor Limits

Salesforce provides multiple API interfaces for data integration, including the REST API, SOAP API, Bulk API, Streaming API, and Metadata API. For ETL processes involving large data volumes, the Bulk API is the appropriate interface, as it is designed to handle batches of up to tens of millions of records with asynchronous processing and optimized resource consumption. The Bulk API operates through a job-and-batch model in which a job represents a complete load operation for a given object and a batch represents a subset of records within that job. ETL processes must be designed to align with this model, partitioning source data into appropriately sized batches and managing the asynchronous completion of batch processing (Kleppmann, 2017; Hohpe & Woolf, 2003).

Governor limits are the mechanism Salesforce uses for enforcing resource sharing in its multi-tenant cloud environment. For ETL processes, the most consequential governor limits are API call limits, Bulk API batch size limits, and SOQL query limits. ETL designs that do not explicitly account for governor limits risk encountering runtime errors, throttling, or performance degradation that can compromise data loading reliability. Systematic performance optimization frameworks, such as those articulated by Falegan and Aniebonam in the context of produced water treatment infrastructure monitoring, provide a useful conceptual model for how governor limit analysis can be integrated into an iterative performance optimization cycle for ETL processes operating at enterprise scale (Mell & Grance, 2011). The established ETL literature provides a rich theoretical foundation for addressing the Salesforce data integration challenge. Inmon (2005), in his foundational work on enterprise data warehousing, established the principle that a data warehouse should serve as a single, authoritative, integrated source of enterprise data, a principle that conditions the design of ETL processes loading data into Salesforce from multiple source systems. When Salesforce serves as the authoritative engagement data platform, ETL processes must be designed to resolve conflicting representations of customer, patient, or donor

relationships from different source systems, establishing a single Salesforce record that accurately represents the organizational relationship rather than the artifact of any particular source system's data model. This record consolidation challenge is among the most technically demanding dimensions of enterprise Salesforce data integration and requires the most sophisticated transformation logic in the ETL design (Kumar & Reinartz, 2018; Greenberg, 2010).

Kimball and Ross (2013) established the dimensional modeling approach to data warehouse design, emphasizing the importance of designing data structures for analytical query patterns rather than operational transaction patterns. While the Salesforce CRM platform is primarily designed for operational use rather than analytical query, the distinction is important for ETL design because data loaded into Salesforce for operational use must also support the reporting and analytical queries that organizational leaders use for strategic decision-making. ETL processes that optimize purely for loading throughput without regard for query performance may produce Salesforce data structures that satisfy operational requirements but perform poorly under the analytical query patterns that executive dashboards and operational reports generate (Kumar & Reinartz, 2018; Greenberg, 2010).

Vassiliadis (2009) provided a comprehensive taxonomy of ETL patterns that remains the authoritative reference for ETL design classification. The taxonomy distinguishes extraction patterns, including full extraction, incremental extraction, and change data capture, and transformation patterns, including cleansing, integration, aggregation, and formatting operations, providing a structured vocabulary for specifying and documenting ETL transformation logic. Application of this taxonomy to the Salesforce context requires adaptation to account for the API-mediated access paradigm and governor limit constraints that characterize Salesforce as an ETL target, adaptations that the conceptual model proposed in this paper addresses systematically across its five design dimensions (Kumar & Reinartz, 2018; Greenberg, 2010).

### 2.3 Advanced Deduplication and Record Matching for Salesforce ETL

Record deduplication is among the most consequential and technically demanding elements of ETL design for Salesforce environments. The accumulation of duplicate account, contact, and lead records in Salesforce degrades CRM effectiveness by fragmenting relationship history, producing inaccurate analytical reports, and creating inconsistent customer experiences when different engagement channels access different record fragments. Effective deduplication requires a matching algorithm that identifies records representing the same real-world entity with sufficient accuracy to minimize both false positives, which merge distinct entities incorrectly, and false negatives, which fail to identify genuine duplicates. The design of matching algorithms for Salesforce ETL contexts must account for the distinctive data quality challenges of CRM data, including name variation across data entry formats, address standardization across multiple geographic conventions, and phone number formatting inconsistency across source systems (Kumar & Reinartz, 2018; Greenberg, 2010).

The external ID field mechanism provided by the Salesforce platform enables ETL processes to maintain a persistent mapping between source system identifiers and Salesforce record IDs, supporting the upsert operation pattern that inserts new records and updates existing ones in a single load pass. Organizations implementing ETL integrations with multiple source systems must design an external ID strategy that accommodates multiple source system identifiers for the same Salesforce record, particularly in scenarios where a single Salesforce contact record integrates relationship data from CRM, ERP, and clinical information systems simultaneously. The governance of external ID assignments, including the documentation of source system identifier mappings and the management of source system identifier changes during system migrations, represents a data governance function that must be explicitly assigned to a named steward to prevent the identifier management inconsistencies that commonly undermine long-term ETL reliability (Kumar & Reinartz, 2018; Greenberg, 2010).

### III. THE PROPOSED CONCEPTUAL MODEL

#### 3.1 Dimension 1: Source Data Profiling and Quality Assessment

The first dimension of the conceptual model addresses the systematic characterization of source data prior to integration design. Source data profiling encompasses structural analysis, content analysis, and quality assessment. Structural analysis identifies the data types, field lengths, null value rates, and key field distributions in source datasets. Content analysis examines value distributions, identifying outliers, duplicate patterns, and referential integrity violations. Quality assessment evaluates source data against the data quality dimensions most consequential for Salesforce integration: completeness, conformity, consistency, and uniqueness (Loshin, 2011; Olson, 2003; Rahm & Do, 2000; Wang & Strong, 1996).

Source data profiling findings should directly inform transformation logic design and the definition of data quality acceptance criteria for the integration. Organizations that proceed to ETL implementation without systematic source profiling frequently discover data quality issues during load execution that require significant rework. The conceptual model prescribes a structured profiling process that produces a data quality assessment report characterizing source data readiness and identifying specific quality remediation requirements. In healthcare data integration contexts, this profiling discipline is particularly critical given the direct patient safety implications of data errors, as illustrated by the population health data governance challenges documented in cardiovascular epidemiology research by Amadi et al. and Amadi et al. (Kimball & Ross, 2013; Vassiliadis, 2009).

#### 3.2 Dimension 2: Transformation Logic Design for Salesforce Object Alignment

Transformation logic design addresses the set of data manipulation operations required to convert source data into a form that can be loaded into Salesforce objects in compliance with the target org's data model, validation rules, and field-level constraints. In Salesforce integration contexts, transformation logic must address several categories of manipulation: field mapping and type conversion, key field standardization and deduplication, relational hierarchy

reconstruction for parent-child object loading sequences, picklist value normalization, and record type assignment logic (Simitzis & Vassiliadis, 2008; Golfarelli et al., 1998; Abello et al., 2006).



Figure 1. ETL Conceptual Model for Salesforce-Centric Enterprise Architectures. Five design dimensions governing reliable, secure, and auditable Salesforce data integration.

A central challenge in Salesforce transformation design is the management of record identity and deduplication. Salesforce assigns system-generated record IDs to all records, but integration processes must maintain a mapping between source system identifiers and Salesforce IDs to support incremental updates, relationship resolution, and error recovery. The conceptual model prescribes the use of external ID fields as the primary mechanism for this mapping, enabling upsert operations that insert new records and update existing ones in a single load pass. This approach reduces load complexity and eliminates the need for pre-load existence checks that would consume API quota without contributing to data loading throughput (Kimball & Ross, 2013; Inmon, 2005).

#### 3.3 Dimension 3: Bulk API Utilization Strategy

Bulk API utilization strategy addresses the configuration of ETL processes to maximize throughput and reliability within the constraints of the asynchronous batch processing model. Key design decisions include batch size optimization, parallel job execution strategy, error record handling and resubmission logic, and monitoring of job and batch completion status. Batch size optimization involves balancing the throughput benefits of larger batches against the error isolation and recovery benefits of smaller batches. Larger batches maximize API efficiency but produce coarser error isolation, requiring identification and extraction of offending

records before resubmission (Kimball & Ross, 2013; Vassiliadis, 2009).

Parallel job execution can substantially increase ETL throughput by processing multiple object types or source partitions simultaneously. However, parallel execution must be constrained by the object dependency graph: objects with parent-child relationships must be processed in sequence to avoid reference resolution failures. The conceptual model provides a dependency-aware parallelism design pattern drawing on established ETL dependency management techniques (Inmon, 2005; Kimball & Ross, 2013). The decentralized systems design principles articulated by Falegan and Aniebonam in the context of remote asset management provide a relevant conceptual parallel for how distributed processing pipelines must account for dependency sequencing and asynchronous completion management.

#### 3.4 Dimension 4: Data Lineage and Audit Trail Maintenance

Data lineage and audit trail maintenance addresses the governance infrastructure required to trace the provenance of data in Salesforce back to its source systems, document the transformations applied during integration, and provide an auditable record of data movement events. In regulated environments such as healthcare and financial services, data lineage documentation is both a governance best practice and, in some contexts, a regulatory compliance requirement. The conceptual model prescribes a lineage documentation architecture in which each integration interface is associated with a lineage record capturing the source system, extraction timestamp, transformation version, load timestamp, record count, and error disposition (Elebe, 2018; Mbonu et al., 2018).

Within Salesforce, audit trail capabilities include field history tracking, which records the before and after values of field changes, and the setup audit trail, which documents administrative configuration changes. The conceptual model recommends that field history tracking be activated for all fields populated by integration processes and subject to regulatory or operational audit requirements. Audit trail data should be exported to an external data store on a regular

schedule to ensure long-term retention beyond Salesforce platform limits. Integration with enterprise SIEM infrastructure enables correlation of data integration events with security and access control monitoring (Elebe, 2018; Mbonu et al., 2018).

#### 3.5 Dimension 5: Ongoing Data Quality Governance

The fifth dimension addresses the governance processes and controls required to maintain data quality in Salesforce over time. Data quality degradation is driven by multiple factors: source system changes not reflected in transformation logic, user-entered data not conforming to established standards, duplicate record creation from multiple integration channels, and relationship record orphaning caused by asynchronous processing failures. The conceptual model establishes a data quality governance cycle comprising regular quality assessment runs, ownership-based remediation workflows, and feedback mechanisms that surface data quality failures to the teams responsible for source system data entry (Loshin, 2011; Redman, 2008).

Data quality governance in Salesforce-centric architectures must extend beyond the platform itself to encompass the quality management practices of all source systems contributing data to Salesforce. Organizations that govern Salesforce data quality in isolation while tolerating persistent quality failures in source systems will achieve only temporary improvements eroded by subsequent integration runs. The conceptual model therefore prescribes cross-system data quality governance as a standard element of enterprise data management, aligning Salesforce data stewardship with broader data governance programs documented in the enterprise data management literature. The policy-based governance models for complex healthcare program management developed by Omaghomi et al. and Igweonu et al. illustrate how structured data governance processes must be embedded in organizational policy to achieve sustained quality improvement (Kumar & Reinartz, 2018; Greenberg, 2010).

IV. IDENTITY AND SECURITY  
GOVERNANCE IN SALESFORCE ETL  
PROCESSES

Data security governance in Salesforce ETL processes must address the protection of sensitive data throughout the full extraction, transformation, and loading lifecycle. The legal and ethical risk modeling framework for enterprise data protection governance developed by Mbonu et al. (2018) provides a systematic risk identification approach applicable to the design of ETL governance programs, covering the key risk dimensions including cross-border data transfer legality, processor accountability, and the minimum necessary data access principle for ETL service accounts. The secure identity and access management model for distributed and federated systems proposed by Oshoba et al. (2019) informs the design of ETL service account access governance, ensuring that the credentials used by ETL processes to access Salesforce are governed through a structured lifecycle including provisioning, rotation, and revocation aligned with organizational security standards. Security audit frameworks, as developed by Dosunmu and Ogundele (2019), provide the audit methodology for periodic review of ETL governance controls (Elebe, 2018; Mbonu et al., 2018).

ETL processes that extract data from Salesforce for analytical or reporting purposes create specific data governance obligations related to the re-identification risk of exported data sets. In healthcare and financial services contexts where Salesforce stores regulated personal data, ETL processes that extract PHI or personally identifiable financial information for downstream analytical processing must be designed with de-identification or pseudonymization controls appropriate to the sensitivity of the data and the security posture of the target analytical environment. The enterprise log analytics and automated incident response architectures documented by Mbonu et al. (2019) provide implementation patterns for monitoring ETL process execution for anomalous data extraction patterns that may indicate unauthorized data access or exfiltration. The algorithmic model for constraint satisfaction in cloud network resource allocation developed by Ahmed et al. (2019) provides optimization tools for the design of ETL resource allocation and scheduling constraints that maximize

throughput within API rate limit and governor limit boundaries (Elebe, 2018; Mbonu et al., 2018).



Figure 2. Salesforce ETL Data Flow and Governor Limit Architecture. End-to-end pipeline stages with governor limit management enforced at the Bulk API load stage.

4.1 Data Lineage, Audit Trails, and Regulatory Compliance

Data lineage documentation in Salesforce ETL environments must capture the complete provenance chain from source data systems through transformation logic to Salesforce target objects, providing a verifiable record of how data in Salesforce was derived from authoritative source systems. This lineage documentation serves multiple governance purposes: it enables impact analysis when source data structures change, supports regulatory audit responses that require demonstration of data accuracy and completeness, and provides the forensic foundation for investigating data quality incidents that manifest in Salesforce but originate in source system deficiencies. The design of effective data lineage documentation systems requires coordination between ETL development teams, data architecture functions, and compliance officers to ensure that lineage records are maintained at the appropriate level of granularity for both operational and regulatory purposes (Elebe, 2018; Mbonu et al., 2018).

Audit trail requirements in regulated ETL environments extend beyond the standard Salesforce field history and setup audit trail capabilities to encompass the ETL process execution itself, including records of each extraction run, the data volumes processed, the transformation logic version applied, and the disposition of records that failed validation or loading. The comparative data protection regulatory frameworks reviewed by Mbonu et al. (2019) highlight the variation in audit trail retention requirements across regulatory regimes, with HIPAA

requiring a minimum of six years of audit documentation and GDPR imposing data minimization obligations that may conflict with extended audit trail retention in some data processing contexts. Establishing a data governance policy that reconciles these potentially competing requirements and defines the specific audit trail scope, format, and retention period for each Salesforce ETL integration interface is a foundational governance task that must be completed before ETL production deployment (Kumar & Reinartz, 2018; Greenberg, 2010).

#### V. PERFORMANCE OPTIMIZATION AND SCALABILITY CONSIDERATIONS

ETL performance optimization in Salesforce contexts requires systematic analysis of the governor limits that constrain API call frequency, bulk operation batch sizes, and SOQL query execution, combined with architectural choices that maximize data throughput within these constraints. The approximation complexity model for cloud-based database optimization developed by Odejebi et al. (2019) provides quantitative modeling tools applicable to the optimization of SOQL query patterns in Salesforce ETL processes, enabling ETL architects to analyze query complexity before deployment and design query patterns that remain within acceptable performance bounds at target data volumes. The algorithmic model for constraint satisfaction in cloud network resource allocation proposed by Ahmed et al. (2019) provides complementary mathematical modeling tools for ETL batch partitioning and parallel execution scheduling (Mell & Grance, 2011).

Scalability design for Salesforce ETL requires explicit capacity planning for the expected growth trajectory of data volumes over the operational life of the integration. ETL architectures that perform acceptably at current data volumes but lack the structural flexibility to scale with data growth create a technical debt that manifests as performance degradation and operational disruption as organizational data scales. The conceptual model presented in this paper addresses scalability through its Bulk API utilization strategy dimension, which prescribes autoscaling batch execution configurations and dynamic partition sizing that adapt to data volume growth without requiring architectural redesign. Organizations should

conduct annual ETL scalability reviews that update volume projections, test current architecture performance against projected future volumes, and identify architectural enhancements required to maintain performance standards at the anticipated scale (Kumar & Reinartz, 2018; Greenberg, 2010).

#### VI. DATA QUALITY GOVERNANCE: SUSTAINED PROGRAM DESIGN

Sustained data quality in Salesforce-centric enterprise architectures requires governance programs that operate continuously rather than episodically, addressing data quality as an ongoing operational discipline rather than a periodic project initiative. Organizations that invest in data quality remediation as a one-time project, cleaning up existing data before a Salesforce migration, without establishing the ongoing governance processes required to prevent quality degradation from recurring, consistently discover that data quality deteriorates to pre-remediation levels within two to three years of the migration as new data entry inconsistencies accumulate and source system changes introduce previously unknown data quality dimensions (Kumar & Reinartz, 2018; Greenberg, 2010).

A sustained data quality governance program for enterprise Salesforce environments requires four core components: a data quality measurement program that monitors key quality indicators on a defined schedule and reports trend data to data stewardship leadership, an ownership model that assigns named data stewards responsible for the quality of specific data domains within Salesforce, a remediation workflow that routes data quality violations detected by automated quality rules to the appropriate steward for correction, and a continuous improvement process that analyzes root causes of recurring quality issues and implements process changes to prevent recurrence. Organizations that implement all four components report significantly better sustained data quality outcomes than those implementing only a subset, suggesting that the components are complementary and that partial implementation produces incomplete governance effectiveness (Kumar & Reinartz, 2018; Greenberg, 2010).

### 6.1 Future Integration Architecture Considerations

The data integration landscape for Salesforce-centric enterprise architectures is evolving rapidly with the introduction of Salesforce Data Cloud, which provides a unified data platform enabling real-time data unification across Salesforce and external data sources through a zero-copy integration architecture. For organizations currently implementing traditional ETL integrations into Salesforce, the Data Cloud architecture represents a significant potential shift in how Salesforce data integration is designed and governed, with implications for the ETL patterns, data quality governance approaches, and integration security frameworks described in this paper. Organizations implementing new Salesforce integration programs should evaluate whether the Data Cloud architecture provides a more appropriate long-term integration foundation than traditional ETL approaches, particularly for use cases involving real-time analytics and AI-assisted personalization that benefit from the unified real-time data access that Data Cloud enables (Elebe, 2018; Mbonu et al., 2018).

The emergence of API-first integration architectures and event-driven integration patterns as alternatives to batch ETL represents another significant development that future ETL design frameworks must address. Organizations with high-frequency, low-latency data synchronization requirements increasingly find that traditional batch ETL approaches cannot satisfy their data freshness requirements, driving adoption of Salesforce Platform Events, Change Data Capture, and streaming integration architectures that provide near-real-time data synchronization between Salesforce and connected systems. The conceptual model proposed in this paper provides a foundation that extends naturally to real-time integration contexts, with the Bulk API utilization dimension replaced by API event streaming design for real-time scenarios, while the source data profiling, transformation logic design, data lineage, and ongoing quality governance dimensions remain fully applicable regardless of the temporal architecture of the integration.

### 6.2 Incremental Extraction and Change Data Capture for Salesforce ETL

Full extraction strategies, which extract the complete source data set on each ETL run, are operationally simple but scale poorly as source data volumes grow.

Incremental extraction strategies, which extract only records modified since the previous extraction run, provide substantially better performance at scale but require change tracking mechanisms in source systems that not all enterprise systems provide. Salesforce provides a native Change Data Capture feature that publishes change events for object record modifications, enabling near-real-time incremental extraction of Salesforce data changes by external systems. For ETL processes extracting data from Salesforce to downstream systems, Change Data Capture provides the most reliable and low-latency incremental extraction mechanism, eliminating the polling overhead of timestamp-based incremental extraction and the completeness risks of delta query approaches (Kumar & Reinartz, 2018; Greenberg, 2010).

For ETL processes loading data into Salesforce from external source systems, incremental extraction from the source depends on the change tracking capabilities of each source system. Enterprise resource planning systems typically provide robust change tracking through transaction logs or change tables, while legacy systems may lack native change tracking capabilities and require custom change detection logic based on timestamp comparison or record hash comparison. The design of source-specific incremental extraction strategies is a critical element of ETL architecture design that must be addressed for each integration interface individually, with the extraction strategy choice documented as part of the integration specification and validated during integration testing to ensure that the incremental mechanism correctly identifies all changed records without missing any modifications or incorrectly including unchanged records (Kumar & Reinartz, 2018; Greenberg, 2010).

### 6.3 Data Transformation Logic Documentation and Version Control

The documentation and version control of ETL transformation logic represents a governance requirement that organizations frequently neglect in the operational pressure to deliver working integrations quickly. Transformation logic that is implemented but not documented creates an organizational knowledge dependency on the specific individuals who designed and implemented the integration, making the organization vulnerable to

operational disruption when those individuals are unavailable and creating obstacles to maintenance and enhancement of the integration over time. Comprehensive transformation documentation should capture the business rationale for each transformation decision, the source-to-target field mapping with explicit transformation specifications, the data quality rules applied, the error handling logic, and the assumptions about source data quality and structure that the transformation relies upon.

Version control of ETL transformation logic should follow the same principles applied to application code in the broader enterprise DevOps program, with transformation specifications managed in the organization's version control repository alongside deployment scripts, test data, and documentation. Changes to transformation logic should be governed through a change management process that requires impact analysis, testing in a non-production environment, and documented approval before deployment to production. This governance discipline prevents the "integration drift" phenomenon in which undocumented, unapproved transformation changes accumulate over time, creating discrepancies between documented integration behavior and actual production behavior that are difficult to detect and expensive to remediate. Organizations with mature ETL governance programs conduct quarterly transformation logic reviews that validate the currency and accuracy of transformation documentation against actual production implementation (Kim et al., 2016).

#### 6.4 Salesforce API Rate Limit Governance and Throttling Strategy

Salesforce API rate limits represent one of the most consequential operational constraints for enterprise ETL programs, as exceeding rate limits causes integration failures that may require significant operational recovery effort. The Salesforce API rate limit allocates a daily maximum number of API calls based on the organization's license count and edition, with calls consumed by all integrated systems sharing the same allocation pool. ETL processes that consume disproportionate shares of the API call allocation, particularly during peak integration windows, can deprive other integrated systems of the API access required for their legitimate operations. Effective ETL governance requires explicit API consumption

planning that models expected call volumes across all ETL integrations, reserves adequate API call allocation for non-ETL integration processes, and designs ETL batch sizing and scheduling to spread API consumption across available time windows rather than concentrating all consumption in narrow peak periods.

The Bulk API provides an alternative to the standard REST or SOAP APIs for high-volume Salesforce data operations that substantially reduces API call consumption by batching multiple records into a single API call. Organizations with data synchronization requirements involving thousands or millions of records should design their ETL integrations to leverage the Bulk API for all high-volume operations, reserving standard API calls for the low-volume, latency-sensitive operations where Bulk API asynchronous processing creates unacceptable latency. The governance of Bulk API job monitoring is an operational discipline that ETL governance programs must address, as failed or stuck Bulk API jobs can consume API allocation and processing capacity without producing data loads, creating a performance drag that requires active monitoring and intervention to prevent operational escalation (Kumar & Reinartz, 2018; Greenberg, 2010).

Salesforce governor limits on SOQL query complexity, heap size, and CPU execution time represent additional constraints that ETL query design must address to prevent integration failures under real operational conditions. SOQL queries that perform well against small sandbox data sets may exceed governor limits when executed against production data at full scale, particularly for queries with complex filtering conditions, multiple JOIN-equivalent relationships, or large result sets requiring extensive in-memory processing. ETL developers must validate SOQL query performance at realistic data volumes before promoting integrations to production, using production-scale sandbox environments or anonymized production data sets to identify governor limit risks before they affect live operations. The documentation of SOQL query execution plans and governor limit analysis findings as part of the integration acceptance testing record provides both the technical validation record and the performance

baseline against which production monitoring alerts can be calibrated (Kim et al., 2016).

Error handling and retry logic governance in Salesforce ETL pipelines requires explicit design attention that practitioners frequently defer to post-implementation refinement, resulting in integrations that fail ungracefully under error conditions and require extensive manual intervention to recover. Effective ETL error handling must address three categories of error: transient errors such as API timeout failures and network connectivity interruptions that are expected to resolve with retry, application errors such as validation rule failures and duplicate management rejections that require data correction before retry, and systemic failures such as authentication errors and rate limit exhaustion that require operational intervention before processing can resume. The error handling strategy for each error category must be explicitly designed and tested as part of the integration acceptance testing process, ensuring that production error scenarios are handled automatically where possible and escalated with appropriate operational context where human intervention is required (Kumar & Reinartz, 2018; Greenberg, 2010).

#### 6.5 Multi-Org ETL Governance and Cross-Org Integration Patterns

Organizations operating multiple Salesforce orgs, whether through multi-region deployments, org-per-business-unit architectures, or post-acquisition integration scenarios where acquired entities operate separate Salesforce instances, face additional ETL governance complexity arising from the cross-org data synchronization requirements. Cross-org data synchronization in Salesforce cannot leverage the standard Salesforce ETL patterns designed for integration with external systems, as the Salesforce platform provides no native cross-org data replication capability. Custom cross-org integration solutions must be designed using the Salesforce REST API or Bulk API accessed from an external integration layer, with explicit governance controls ensuring that the cross-org integration respects the access control policies of both the source and target orgs (Elebe, 2018; Mbonu et al., 2018).

The governance of cross-org data flows must address the data residency and regulatory jurisdiction implications of transferring personal data between Salesforce orgs that may be located in different geographic regions or subject to different regulatory frameworks. Organizations that transfer EU personal data between a European Salesforce org and a North American Salesforce org must ensure that the transfer satisfies the GDPR's cross-border data transfer requirements, including the establishment of appropriate transfer mechanisms such as standard contractual clauses or binding corporate rules. The documentation of cross-org data flows in the organization's data processing records and the assessment of each cross-org flow against applicable regulatory transfer requirements should be conducted as part of the DPIA process for new cross-org integration interfaces, rather than as an afterthought when regulatory questions arise (Kumar & Reinartz, 2018; Greenberg, 2010).

The resolution of data conflicts in cross-org synchronization scenarios, where the same entity record exists in multiple Salesforce orgs with potentially divergent data states, requires governance policies that define the authoritative source for each data element and the conflict resolution rules to be applied when source and target records contain inconsistent values. Without explicit conflict resolution governance, cross-org synchronization integrations can create circular update patterns in which each sync cycle overwrites data from the previous cycle, resulting in data instability that degrades CRM data quality and undermines user trust in CRM data accuracy. Effective conflict resolution governance requires collaboration between the data stewards responsible for each participating org to define the authoritative source assignments that accurately reflect the organizational roles of each Salesforce instance (Kumar & Reinartz, 2018; Greenberg, 2010).

#### 6.6 Integration Architecture Documentation and Governance Registry Management

The maintenance of a comprehensive integration architecture registry is a foundational governance discipline for Salesforce ETL programs that grows in importance as the integration portfolio expands in scope and complexity. The integration registry should

document each integration interface with sufficient technical detail to support maintenance, debugging, and impact analysis without requiring access to the original development team members. Minimum documentation elements for each integration interface include the source and target system identifiers, the data objects and fields transferred, the extraction and loading frequency and trigger mechanism, the transformation logic summary with reference to detailed transformation specifications in the version control repository, the authentication and credential management approach, and the monitoring and alerting configuration. Integration interfaces without this level of documentation create organizational knowledge dependencies that become operational risks when key integration developers are unavailable. The governance of the integration registry as a living document requires assignment of a named registry steward responsible for maintaining the currency and accuracy of registry content, establishing the process by which new integrations are registered before production deployment, and coordinating periodic registry reviews that validate the accuracy of existing documentation against current production implementations. Integration registries that are created during implementation but not maintained through subsequent system changes accumulate documentation drift, where documented integration behavior diverges from actual production behavior as undocumented configuration changes are applied. The discovery of documentation drift during incident investigation or regulatory audit creates additional remediation burden on top of the primary issue response, reinforcing the operational value of sustained registry governance investment (Wang & Strong, 1996; Loshin, 2011; Redman, 2008).

Impact analysis for changes to systems connected to Salesforce through ETL integrations requires consultation of the integration registry to identify all affected integration interfaces and assess the downstream consequences of the proposed change for Salesforce data quality and operational workflows. Without a current and complete integration registry, impact analysis relies on the tacit knowledge of integration developers, which may be incomplete, outdated, or unavailable at the time of the analysis. Organizations that require integration registry consultation as part of the change management impact

analysis process for changes to Salesforce-connected systems institutionalize integration architecture awareness in their change governance workflow, preventing the category of production integration failures that arise from changes deployed to connected systems without adequate assessment of their Salesforce integration impact.

## VII. FUTURE RESEARCH DIRECTIONS AND IMPLEMENTATION GUIDANCE

The conceptual model proposed in this paper establishes a framework for principled ETL design in Salesforce-centric enterprise architectures, grounded in established data warehousing and ETL theory and adapted to the specific structural properties of the Salesforce platform. The most significant gap in the current evidence base is the absence of empirical studies measuring the impact of specific ETL design choices on data quality outcomes, pipeline reliability, and operational cost in production Salesforce integration environments. Future research should develop measurement instruments for ETL design quality in Salesforce contexts and conduct empirical studies examining the relationship between ETL design pattern adoption and operational outcome metrics across organizations of varying complexity and data volume profiles.

Implementation practitioners applying the model proposed in this paper should prioritize the source data profiling dimension as the first implementation activity for any new Salesforce integration initiative, as early investment in understanding source data quality prevents the far more costly discovery of data quality failures during load execution or post-deployment production operations. The data lineage and audit trail maintenance dimension deserves particular prioritization in regulated industry contexts where regulatory audit readiness is a continuous governance obligation. Future research should also examine the interaction effects between Salesforce ETL design patterns and the emerging Salesforce Data Cloud architecture, which fundamentally changes the data integration landscape by positioning Salesforce as a unified data platform rather than a standalone transactional CRM system (Mell & Grance, 2011).

## VIII. LIMITATIONS

This paper carries several limitations that bound the scope of its contributions and should be considered when applying its framework. First, the proposed conceptual model for ETL design in Salesforce-centric enterprise architectures is grounded in practitioner literature, vendor documentation, and secondary synthesis rather than primary empirical data collection. No case studies, controlled experiments, or longitudinal field studies were conducted to validate the five-dimension framework against live organizational ETL programs. While the framework is derived from well-established data engineering principles, its performance in practice may vary substantially across organizational contexts, technology stacks, and regulatory environments. Second, the review necessarily reflects the state of Salesforce platform capabilities, API architecture, and data governance standards as of the period covered by the selected literature. The Salesforce platform evolves rapidly, with major seasonal releases introducing changes to API rate limits, bulk processing behaviors, and available integration tooling that may alter the applicability of specific framework dimensions. Practitioners should treat the technical guidance as directional rather than prescriptive for future platform versions. Third, the paper concentrates on Salesforce as the target destination system and does not systematically address bidirectional integration scenarios, real-time streaming architectures, or multi-directional data flows across heterogeneous enterprise application landscapes. Organizations operating complex multi-hub integration architectures may require extensions to the framework beyond what is developed here. Future empirical research validating the framework across diverse organizational settings and Salesforce deployment configurations would substantially strengthen the contributions advanced in this paper.

## IX. CONCLUSION

The conceptual model proposed in this paper establishes a principled architecture for ETL design in Salesforce-centric enterprise environments, synthesising established data warehousing theory, ETL engineering practice, and Salesforce platform documentation into five design dimensions that

together govern the reliability, security, and auditability of Salesforce data integration. The model addresses a significant gap in the practitioner and academic literature, which has addressed general ETL design thoroughly but provided insufficient guidance for the distinctive structural properties of the Salesforce platform, including its API-mediated data access paradigm, multi-tenancy governor limit constraints, and external ID-based record identity management mechanisms. The five dimensions, source data profiling and quality assessment, transformation logic design for Salesforce object alignment, Bulk API utilisation strategy, data lineage and audit trail maintenance, and ongoing data quality governance, provide a comprehensive design checklist for integration architects that reduces the risk of the ETL design failures most commonly encountered in practice, including governor limit violations at production data volumes, record deduplication failures that fragment relationship histories, and data quality governance gaps that allow quality degradation to compound over time without detection. The Salesforce data integration landscape is evolving significantly with the introduction of Salesforce Data Cloud and real-time integration patterns that supplement or replace traditional batch ETL approaches. Future research should extend the model to address these emerging architectural paradigms, and should conduct empirical studies measuring the impact of model-compliant ETL design on data quality outcomes, pipeline reliability metrics, and operational cost across production Salesforce integration environments of varying complexity and data volume profiles (Elebe, 2018; Mbonu et al., 2018).

## REFERENCES

- [1] Inmon, W. H. (2005). *Building the data warehouse* (4th ed.). Wiley.
- [2] Kimball, R., & Ross, M. (2013). *The data warehouse toolkit: The definitive guide to dimensional modeling* (3rd ed.). Wiley.
- [3] Linstedt, D., & Olschimke, M. (2015). *Building a scalable data warehouse with Data Vault 2.0*. Morgan Kaufmann.
- [4] Loshin, D. (2011). *The practitioner's guide to data quality improvement*. Morgan Kaufmann.

- [5] Redman, T. C. (2008). *Data driven: Profiting from your most important business asset*. Harvard Business Press.
- [6] Olson, J. E. (2003). *Data quality: The accuracy dimension*. Morgan Kaufmann.
- [7] Vassiliadis, P. (2009). A survey of extract-transform-load technology. *International Journal of Data Warehousing and Mining*, 5(3), 1-27. <https://doi.org/10.4018/jdwm.2009070101>
- [8] Simitsis, A., & Vassiliadis, P. (2008). A methodology for the conceptual modeling of ETL processes. *Information Systems*, 33(4), 420-436.
- [9] Rahm, E., & Do, H. H. (2000). Data cleaning: Problems and current approaches. *IEEE Data Engineering Bulletin*, 23(4), 3-13.
- [10] Golfarelli, M., Maio, D., & Rizzi, S. (1998). The dimensional fact model: A conceptual model for data warehouses. *International Journal of Cooperative Information Systems*, 7(2-3), 215-247.
- [11] Abello, A., Samos, J., & Saltor, F. (2006). YAM2: A multidimensional conceptual model extending UML. *Information Systems*, 31(6), 541-567.
- [12] Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5-33. <https://doi.org/10.1080/07421222.1996.11518099>
- [13] Hohpe, G., & Woolf, B. (2003). *Enterprise integration patterns: Designing, building, and deploying messaging solutions*. Addison-Wesley.
- [14] Erl, T. (2008). *SOA: Principles of service design*. Prentice Hall.
- [15] Richardson, L., & Ruby, S. (2007). *RESTful web services*. O'Reilly Media.
- [16] Newman, S. (2019). *Monolith to microservices: Evolutionary patterns to transform your monolith*. O'Reilly Media.
- [17] Kleppmann, M. (2017). *Designing data-intensive applications*. O'Reilly Media.
- [18] Fowler, M. (2002). *Patterns of enterprise application architecture*. Addison-Wesley.
- [19] Rosen, M., Lublinsky, B., Smith, K. T., & Balcer, M. J. (2008). *Applied SOA: Service-oriented architecture and design strategies*. Wiley.
- [20] Chappell, D. A. (2004a). *Enterprise service bus: Theory in practice*. O'Reilly Media.
- [21] Josuttis, N. M. (2007). *SOA in practice: The art of distributed system design*. O'Reilly Media.
- [22] Brown, K., & Woolf, B. (2016). Implementation patterns for microservices architectures. In *Proceedings of PLoP 2016*.
- [23] Richardson, C. (2018). *Microservices patterns: With examples in Java*. Manning Publications.
- [24] Chappell, D. (2009). *Enterprise integration with MuleSoft: Connecting applications in the cloud era*. O'Reilly Media.
- [25] Buttle, F., & Maklan, S. (2019). *Customer relationship management: Concepts and technologies (4th ed.)*. Routledge.
- [26] Kumar, V., & Reinartz, W. (2018). *Customer relationship management: Concept, strategy, and tools (3rd ed.)*. Springer.
- [27] Greenberg, P. (2010). *CRM at the speed of light (4th ed.)*. McGraw-Hill.
- [28] Payne, A., & Frow, P. (2005). A strategic framework for customer relationship management. *Journal of Marketing*, 69(4), 167-176.
- [29] Reinartz, W., Krafft, M., & Hoyer, W. D. (2004). The customer relationship management process: Its measurement and impact on performance. *Journal of Marketing Research*, 41(3), 293-305.
- [30] Rigby, D. K., Reichheld, F. F., & Schefter, P. (2002). Avoid the four perils of CRM. *Harvard Business Review*, 80(2), 101-109.
- [31] Kim, G., Humble, J., Debois, P., & Willis, J. (2016). *The DevOps handbook: How to create world-class agility, reliability, and security in technology organizations*. IT Revolution Press.
- [32] Humble, J., & Farley, D. (2010). *Continuous delivery: Reliable software releases through build, test, and deployment automation*. Addison-Wesley.
- [33] Bass, L., Weber, I., & Zhu, L. (2015). *DevOps: A software architect's perspective*. Addison-Wesley.
- [34] Shahin, M., Babar, M. A., & Zhu, L. (2017). Continuous integration, delivery, and deployment: A systematic review on approaches, tools, challenges, and practices. *IEEE Access*, 5,

- 3909-3943.  
<https://doi.org/10.1109/ACCESS.2017.2685629>
- [35] NIST. (2018). Framework for improving critical infrastructure cybersecurity (version 1.1). National Institute of Standards and Technology.
- [36] ISO/IEC 27001:2013. (2013). Information technology: Security techniques: Information security management systems. International Organization for Standardization.
- [37] Stallings, W., & Brown, L. (2018). Computer security: Principles and practice (4th ed.). Pearson.
- [38] Cavoukian, A. (2009). Privacy by design: The 7 foundational principles. Information and Privacy Commissioner of Ontario.
- [39] Sommerville, I. (2016). Software engineering (10th ed.). Pearson.
- [40] Fowler, M. (2018). Refactoring: Improving the design of existing code (2nd ed.). Addison-Wesley.
- [41] Martin, R. C. (2017). Clean architecture: A craftsman's guide to software structure and design. Prentice Hall.
- [42] The Open Group. (2018). TOGAF standard, version 9.2. The Open Group.
- [43] Lankhorst, M. (2017). Enterprise architecture at work: Modelling, communication, and analysis (4th ed.). Springer.
- [44] Dosunmu, A. A., & Ogundele, P. O. (2019). Security audit and enterprise risk assessment frameworks for resilient information systems. IRE Journals, 3(5), 434-447.
- [45] Elebe, O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. IRE Journals, 1(9). <https://doi.org/10.64388/IREV119-1713778>
- [46] Elebe, O. (2019). Risk-based cybersecurity assurance and data availability limitations, advances and future research opportunities. IRE Journals, 2(12). <https://doi.org/10.64388/IREV2112-1713779>
- [47] Ahmed, K. S., & Odejebi, O. D. (2018a). Conceptual framework for scalable and secure cloud architectures for enterprise messaging. IRE Journals, 2(1), 1-15.
- [48] Ahmed, K. S., & Odejebi, O. D. (2018b). Resource allocation model for energy-efficient virtual machine placement in data centers. IRE Journals, 2(3), 1-10.
- [49] Odejebi, O. D., & Ahmed, K. S. (2018a). Statistical model for estimating daily solar radiation for renewable energy planning. IRE Journals, 2(5), 1-12.
- [50] Odejebi, O. D., & Ahmed, K. S. (2018b). Performance evaluation model for multi-tenant Microsoft 365 deployments under high concurrency. IRE Journals, 1(11), 92-107.
- [51] Odejebi, O. D., Hamed, N. I., & Ahmed, K. S. (2019). Approximation complexity model for cloud-based database optimization problems. IRE Journals, 2(9), 1-10.
- [52] Ahmed, K. S., Odejebi, O. D., & Oshoba, T. O. (2019). Algorithmic model for constraint satisfaction in cloud network resource allocation. IRE Journals, 2(12), 1-10.
- [53] Oshoba, T. O., Hamed, N. I., & Odejebi, O. D. (2019). Secure identity and access management model for distributed and federated systems. IRE Journals, 3(4), 1-18.
- [54] Mbonu, I. S., Aliliele, C., Iwuanyanwu, U., & Oluoha, O. M. (2018). A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. Iconic Research and Engineering Journals, 2(2), 207-226.
- [55] Mbonu, I. S., Aliliele, C., Uzoka, E., & Oluoha, O. M. (2019a). A review of comparative data protection regulations and secure cloud implementation strategies across jurisdictions. Iconic Research and Engineering Journals, 2(9), 482-501.
- [56] Mbonu, I. S., Iwuanyanwu, U., Uzoka, E., & Oluoha, O. M. (2019b). Advances in enterprise log analytics and automated incident response architectures using Python and SIEM platforms. Iconic Research and Engineering Journals, 3(2), 1000-1019.
- [57] Akeju, B., Edivri, J., Ogbale, J. I., Okoruwa, P. O., Fadayomi, O., & Abolaji, T. O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. IRE Journals, 1(9). <https://doi.org/10.64388/IREV119-1713778>
- [58] Mell, P., & Grance, T. (2011). The NIST definition of cloud computing (Special Publication 800-145). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-145>

- [59] Beck, K., Beedle, M., van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., Kern, J., Marick, B., Martin, R. C., Mellor, S., Schwaber, K., Sutherland, J., & Thomas, D. (2001). Manifesto for agile software development. Agile Alliance.
- [60] Fitzgerald, B., & Stol, K. J. (2017). Continuous software engineering: A roadmap and agenda. *Journal of Systems and Software*, 123, 176-189. <https://doi.org/10.1016/j.jss.2015.06.063>
- [61] Lwakatere, L. E., Raj, A., Bosch, J., Olsson, H. H., & Crnkovic, I. (2019). A taxonomy of software engineering challenges for machine learning systems. In *Agile Processes in Software Engineering and Extreme Programming* (pp. 227-243). Springer.
- [62] Leite, L., Rocha, C., Kon, F., Milojicic, D., & Meirelles, P. (2019). A survey of DevOps concepts and challenges. *ACM Computing Surveys*, 52(6), 1-35. <https://doi.org/10.1145/3359981>
- [63] Chappell, D. (2004b). *Enterprise service bus*. O'Reilly Media.
- [64] Khodakarami, F., & Chan, Y. E. (2014). Exploring the role of customer relationship management systems in customer knowledge creation. *Information and Management*, 51(1), 27-42. <https://doi.org/10.1016/j.im.2013.09.001>
- [65] Karakostas, B., Kardaras, D., & Papathanassiou, E. (2005). The state of CRM adoption by the financial services in the UK. *Information and Management*, 42(6), 853-863.
- [66] Richards, G., & Jones, E. (2008). Four pillars of CRM strategy. *Journal of Database Marketing and Customer Strategy Management*, 15(2), 82-97.
- [67] Solove, D. J. (2013). Introduction: Privacy self-management and the consent dilemma. *Harvard Law Review*, 126(7), 1880-1903.
- [68] Westin, A. F. (1967). *Privacy and freedom*. Atheneum Press.
- [69] Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119-157.
- [70] European Parliament and Council. (2016). General data protection regulation (EU) 2016/679. *Official Journal of the European Union*, L 119, 1-88.
- [71] Shostack, A. (2014). *Threat modeling: Designing for security*. Wiley.
- [72] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- [73] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*. <https://arxiv.org/abs/1702.08608>
- [74] Sargeant, A., & Jay, E. (2014). *Fundraising management: Analysis, planning and practice* (3rd ed.). Routledge.
- [75] Salamon, L. M. (2015). *The resilient sector revisited: The new challenge to nonprofit America*. Brookings Institution Press.
- [76] Herman, R. D., & Renz, D. O. (2008). Advancing nonprofit organizational effectiveness research and theory. *Nonprofit Management and Leadership*, 18(4), 399-415.