International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 7, Issue 1, 244-250, 2026
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V7I1P137

Original Article

Bridging Operational Technology (OT) and Enterprise Analytics: A Framework for Integrating AVEVA PI with Cloud-Scale ELT Pipelines

Ekta Sojitra
Independent Researcher, Peoria, Illinois.

Received On: 19/01/2026 ***Revised On: 20/02/2026*** ***Accepted On: 26/02/2026*** ***Published On: 11/03/2026***

Abstract: *Industrial organizations struggle to bridge the gap between high frequency operational telemetry and cloud-scale enterprise analytics without sacrificing semantic integrity and incurring prohibitive computational costs. While the AVEVA PI System remains a mission-critical staple for contextualizing time-series data, traditional integration methods often trigger significant semantic loss during the transition to the cloud. This research introduces a reference framework utilizing Informatica IDMC and a layered Extract, Load, Transform (ELT) architecture to synchronize the PI System with Snowflake. Unlike standard migrations that flatten data into disconnected tags, this approach embeds operational semantics directly through the preservation of asset models and event context. By utilizing warehouse native Change Data Capture (CDC), specifically Snowflake Streams and Tasks, the framework replaces inefficient full refresh cycles with high performance incremental processing. To address the rigorous requirements of critical infrastructure, the architecture aligns with NIST SP 800-82 Revision 3 security standards and introduces a quantitative observability model based on Service Level Objectives (SLOs) for data freshness and reconciliation. I validated the framework through a utility. y scale implementation, which yielded a 95th percentile latency of 7.4 minutes and approximately 38% reduction in compute consumption compared to legacy methods. Ultimately, the aim of this research is to provide a practical, reproducible blueprint for organizations seeking to modernize industrial analytics while maintaining operational trust, performance, and continuity.*

Keywords: *AVEVA PI System, Informatica IDMC, Snowflake Streams/Tasks, Industrial Iot (Iiot), OT/IT Convergence, Data Observability, NIST SP 800-82, Cloud Data Warehousing.*

1. Introduction

The industrial sector is navigating a fundamental paradigm shift as it moves from isolated Supervisory Control and Data Acquisition (SCADA) systems toward integrated cloud native analytics. This IT/OT convergence is not merely a trend; it is a prerequisite for unlocking predictive maintenance, real-time demand forecasting, and heightened operational safety. However, the migration from plant floor historians like AVEVA PI System to modern cloud data warehouses like Snowflake remains technically fraught [1], [5]. Despite the promise of the cloud, the transition often breaks the very data relationships that make industrial telemetry valuable.

Current integration strategies typically stumble over three persistent hurdles:

- Semantic Loss: Traditional ETL processes frequently flatten sophisticated asset hierarchies into disconnected raw tags. This data dumping strips away the operational context essential for high level analysis [2], [3].
- Security Fragility: Bridging secure Operational Technology (OT) zones with public cloud environments often introduces vulnerabilities that clash with critical infrastructure standards, such as NIST SP 800-82r3 [6].
- Economic Inefficiency: High frequency telemetry generates massive data volumes, attempting to process these via full-table refreshes results in prohibitive compute costs and latencies that render real time monitoring impossible [8].

In this paper, I propose a comprehensive architectural framework designed to bypass these bottlenecks using a metadata driven, CDC aware pipeline by employing Informatica IDMC as a resilient orchestration layer [7], this research demonstrates a method for preserving source semantics while simultaneously optimizing cloud performance. I validate the framework's efficacy through a case study at a North American utility company, where the architecture modernized transformer health monitoring. The implementation not only met stringent regulatory requirements but also significantly reduced the Mean Time to Recovery (MTTR) [9].

The remainder of this paper is structured as follows: Section II evaluates existing literature; Section III establishes core system requirements; Sections IV and V detail the proposed architecture; and the concluding sections provide a

rigorous evaluation through the lenses of security, observability, and empirical validation.

2. Related Work

The literature surrounding industrial data integration spans several diverging domains, from specialized OT architectures to modern cloud native data engineering. This research sits at the intersection of these fields, synthesizing established historian practices with modern elastic compute capabilities.

Industrial Data Management: The foundational architecture of the AVEVA PI System relies on a strict separation of concerns, isolating the high throughput Data Archive from the contextual metadata within the Asset Framework (AF) and the transactional nature of Event Frames [2]-[4]. While these structures are robust within the OT perimeter, they are notoriously difficult to replicate in relational cloud environments. This paper builds upon these standards by ensuring that OT native hierarchies are preserved rather than flattened during the migration process.

Cloud Native Ingestion and CDC: Prior research from cloud providers has established resilient hosting patterns for PI Systems, emphasizing private connectivity and Multi-Availability Zone (Multi-AZ) redundancy to target five-nine availability [5]. Within the orchestration domain, Informatica IDMC provides a framework for parameterized, metadata driven taskflows [7]. I utilize these capabilities to manage the complex ingestion windows inherent in high frequency telemetry. Furthermore, the use of Snowflake Streams and Tasks for CDC is now recognized as a best practice for minimizing both compute latency and the cloud tax associated with full refresh pipelines [8].

Security and Observability Standards: Bridging the gap between the plant floor and the cloud necessitates a rigorous security posture. NIST SP 800-82 Revision 3 remains the definitive guide for OT security, prescribing a defense-in-depth strategy that mandates network segmentation and granular access controls [6]. However, security alone does not guarantee reliability. Recent advancements in system observability have moved beyond infrastructure level metrics (e.g., CPU/RAM) towards data centric telemetry. Following the logic in [9], I advocate for structured telemetry schemas and the application of SLOs to reduce the MTTR within these increasingly distributed and large-scale data systems.

3. Problem Statement and Requirements

The primary objective of this framework is to establish a reliable, secure, and context aware bridge between high frequency OT environments and cloud scale analytics. Standard ingestion methods fail in this domain because they treat industrial data as flat streams rather than structured operational assets. To address this, I have defined five core functional and non-functional requirements that govern the proposed architecture:

- R1 - Semantic Fidelity: The system must preserve the hierarchical integrity of PI AF templates and the episodic context of Event Frames because without this, downstream curated data marts and Key Performance Indicators (KPIs) lose their original operational meaning, rendering them useless for complex root cause analysis [2], [3].
- R2 - Temporal Heterogeneity: Industrial telemetry does not move at a single speed; the architecture must support mixed ingestion cadences ranging from 5-minute sub-cycles to daily aggregates. This necessitates idempotent pipeline logic and deterministic processing windows capable of gracefully handling late arriving telemetry [2].
- R3 - Incremental Efficiency (CDC): To mitigate the cloud tax and ensure data freshness, the framework mandates warehouse native CDC. It utilizes Snowflake Streams and Tasks to isolate delta changes, avoiding the high computational costs of full table refreshes [8].
- R4 - Hardened Security Posture: Bridging OT and IT zones introduces significant risk. All security controls must map directly to the NIST SP 800-82 Revision 3 standards, incorporating strict network segmentation, end-to-end encryption, and Role-Based Access Control (RBAC) to maintain a defense-in-depth strategy [6].
- R5 - Data Centric Observability: Modern industrial systems require more than just up/down infrastructure monitoring. This framework implements a telemetry model focused on data centric health, tracking specific metrics such as row-level reconciliation, latency percentiles, and MTTR [9].

4. Framework Overview

The proposed architecture adopts a layered, Source-to-Sovereign design aimed at maintaining industrial telemetry integrity while enabling cloud scale elasticity. A key design principle here is the decoupling of collection, transport, and transformation layers; this ensures that intensive IT analytical workloads never jeopardize OT operational stability. The framework is organized into five logical layers:

Layer 1: Collection and Contextualization (The OT Source)

At the edge, the AVEVA PI System functions as the primary data producer. Diverging from common sensor-to-cloud methods that strip away metadata, this framework extracts data specifically through the PI AF and Event Frames. This ensures that telemetry arrives in the cloud with pre-existing parent-child hierarchies and episodic context (e.g., specific downtime events) intact, directly satisfying the requirement for Semantic Fidelity (R1) [2]–[4].

Layer 2: Secure Landing and Transport

Consistent with NIST SP 800-82 Revision 3 (R4), the transport layer implements "Data Diode" logic via private networking, such as a VPN or AWS Direct Connect. Data is staged in a Secure Landing Zone, typically an S3/Blob storage bucket or a cloud hosted PI Data Archive. This buffer plays a critical role ensuring the enterprise analytics environment has no direct, unbuffered path into the plant

network, effectively isolating the control domain from external queries [5], [6].

Layer 3: Ingestion and Orchestration (IDMC)

Informatica IDMC serves as the orchestration engine, governing the flow between the landing zone and the data warehouse. Rather than using static scripts, I employ a metadata driven approach where ingestion windows are dynamically calculated through parameterized mapping tasks. This layer resolves the challenge of Mixed Frequencies (R2) by triggering specific taskflows based on temporal offsets or file-arrival signals, ensuring idempotent data loading [7].

Layer 4: Curation and CDC (Snowflake)

Core analytical processing occurs within Snowflake using a high-performance ELT pattern. This layer utilizes Snowflake Streams to track row-level deltas in the landing tables, while Tasks automate the MERGE logic into the final curated tables. By shifting the computational gravity to Snowflake native CDC, the framework ensures that only new or modified telemetry is processed. This satisfies the CDC and Performance requirement (R3) while drastically lowering the compute tax associated with large scale industrial datasets [8].

Layer 5: Governed Consumption and Observability

The final layer surfaces data for BI, machine learning, and executive reporting. Unlike traditional reporting layers, this environment is wrapped in a data centric observability model. By monitoring the Heartbeat of the pipeline, precisely tracking freshness and row level reconciliation in real-time; the system provides the transparency required for operational trust and rapid MTTR (R5) [9].

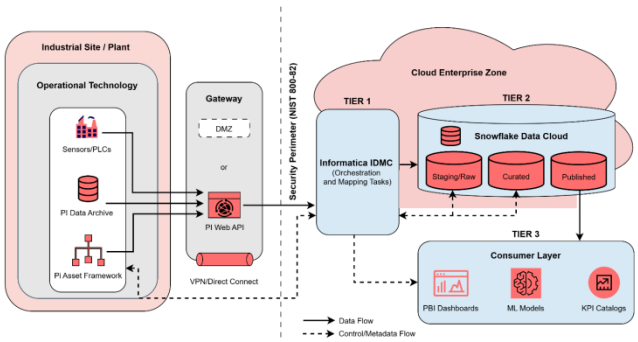

Fig 1: Reference Architecture

5. Architecture and Data Flow

The implementation of this framework follows a deterministic sequence of operations designed to transition high fidelity OT telemetry into an enterprise ready format. This process moves beyond simple data ingestion, focusing on a coordinated handshake between the extraction, transport, and transformation layers.

5.1. Access and Semantic Extraction

Data extraction originates at the PI Web API layer, serving as the programmatic gateway to the historian. Rather than pulling raw, disconnected tags, the framework targets PI AF attributes and Event Frames. This approach preserves the object-oriented structure of the industrial data from the point of origin. To optimize payload efficiency, requests are formatted as compressed CSV or JSON. Crucially, I enforce AF template governance upstream; by version controlling the asset model, the framework pre-emptively mitigates the risk of downstream schema drift [2]-[4].

5.2. Connectivity and Secure Egress

Adhering to the defense in depth requirements of NIST SP 800-82 Revision 3, the framework mandates dedicated private networking such as a Site-to-Site VPN or AWS Direct Connect for all data egress. By ensuring telemetry never traverses the public internet, the architecture minimizes the attack surface. For utility scale resilience, the PI System is ideally hosted within a cloud native Virtual Private Cloud (VPC) across multi-AZ, providing a hardware agnostic failover mechanism against localized infrastructure outages [5].

5.3. Orchestration and Metadata Driven Ingestion

Informatica IDMC functions as the control plane, managing data movement from the secure gateway to the Snowflake staging environment. To avoid the rigidity of hard coded pipelines, I utilize parameterized mapping tasks. In this model, processing windows (e.g., start/end timestamps) are dynamically calculated based on the metadata of the last successful execution. This metadata comprising row counts, high water marks, and execution durations is captured in a centralized operational log, providing the foundation for both lineage tracking and automated failure recovery [7].

5.4. CDC Aware Transformations in the Warehouse

Once data reaches the Snowflake staging layer, the architecture shifts to a warehouse native CDC pattern. I defined Snowflake Streams on the staging tables to capture row level deltas (inserts and updates) in real-time. Snowflake Tasks then automate the MERGE logic, upserting only the delta records into the curated and published layers. These tasks are structured as Directed Acyclic Graphs (DAGs) using AFTER dependencies. This ensures a strict execution order: transformations only trigger once the preceding ingestion task has successfully been committed, preventing partial data loads and maintaining referential integrity [8].

6. Security and Governance Framework

The integration of OT telemetry with cloud analytics introduces a significant attack surface that necessitates a structured, multi-layered security posture. This framework aligns its control objectives with the NIST SP 800-82 Revision 3 guidelines for Industrial Control System (ICS) security, centering on four defensive pillars.

6.1. Network Segmentation and Least Privilege

To mitigate the risk of lateral movement between the plant floor and the enterprise cloud, the architecture enforces strict network segmentation. The OT environment remains isolated, with data egress restricted to a hardened Demilitarized Zone (DMZ) via the PI Web API. I apply the

principle of least privilege across all service layers: Informatica IDMC service accounts are limited to read-only access for specific AF templates, while Snowflake Role-Based Access Control (RBAC) ensures that only authorized analysts can query the published data layer. This isolation ensures that a compromise in the analytics cloud cannot pivot back into the control domain [6].

6.2. Encryption and Data Integrity

Protection of telemetry is mandatory at every stage of the data lifecycle. Encryption in transit is strictly enforced using TLS 1.2 or higher for all API calls and tunnel traffic (VPN/Direct Connect). At rest, data is protected natively within the Snowflake environment using AES-256 multi-key encryption. To verify data integrity, the framework incorporates digital signatures and checksums during the IDMC ingestion phase. This ensures that telemetry remains untampered from the moment it leaves the PI System until it is committed to the cloud warehouse.

6.3. Centralized Monitoring and Incident Response

Observability in this framework extends into the security domain through centralized log aggregation. Access logs from the PI Web API, IDMC task execution metadata, and Snowflake login histories are fed into a centralized Security Information and Event Management (SIEM) system. This centralized pane of glass enables real-time detection of anomalous access patterns or unauthorized egress attempts, facilitating a response cadence that aligns with industrial safety protocols rather than just IT standard timelines [6].

6.4. Governance and Change Control

To maintain a Single Source of Truth, the framework implements a rigorous change control process for data definitions. All modifications to the PI AF templates or Snowflake curated schemas follow a promotion-based lifecycle (Development, Testing, and Production). This governed approach prevents the emergence of shadow data structures, unauthorized or unversioned tables, and ensures that operational semantics remain consistent as the enterprise scales.

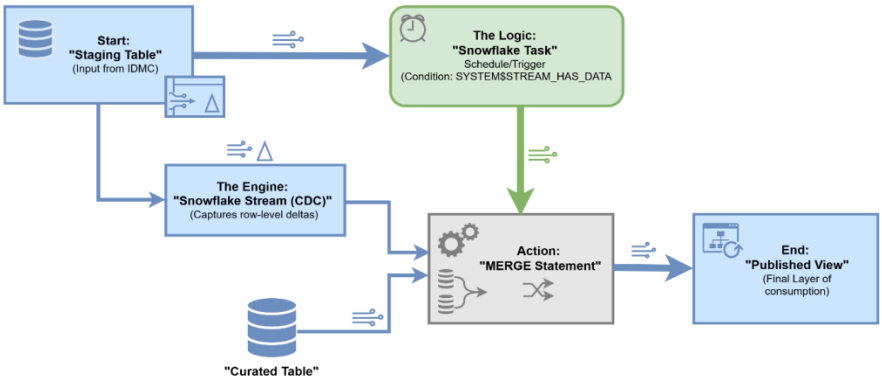

Fig 2: CDC Logic and DAG

7. Observability and Slos

Modern OT-to-Cloud pipelines require observability that transcends traditional infrastructure heartbeat checks. This framework implements a structured telemetry model that treats data integrity and delivery as first class citizens, utilizing specific SLOs to minimize the MTTR [9].

7.1. Data Centric Telemetry Model

The observability layer is anchored by a job registry and telemetry schema within Snowflake. This architecture captures four distinct pillars of pipeline health:

- Freshness (t_{lat}): I define freshness as the temporal delta between the source timestamp in the PI Data Archive (T_{src}) and the availability timestamp in the Snowflake published layer (T_{pub}): $t_{lat} = T_{pub} - T_{src}$
- Volume and Shape: The system monitors statistical anomalies in row counts and data distributions, such as a sudden drop in sensor events using automated detectors to identify potential edge-device failures.
- Reconciliation Rate: This involves a row count parity check between the PI Web API extraction count and the Snowflake staging record count, ensuring zero data loss during the transit phase.
- Failure Taxonomy: Rather than simple "success/failure" flags, I categorize task failures (e.g., authentication errors, timeouts, or schema mismatches) to accelerate root-cause analysis and automate recovery workflows.

7.2. Service Level Objectives (SLOs)

To validate the framework's effectiveness, I established the following SLOs as a baseline for operational readiness in a utility scale environment. These targets represent the minimum acceptable performance for mission critical industrial analytics.

Table 1: Framework Performance Targets (Slos)

Metric	Definition	Target (SLOs)	Measurement Point
Freshness	End to End ingestion latency	$P_{95} \leq 10$ mins	Snowflake published view
Reconciliation	Source to Target parity	≥ 99.9 %	IDMC audit logs
Anomaly Rate	Significant shape/volume shifts	< 1/day	Telemetry detectors
MTTR	Mean time to recover from failure	≤ 30 mins	Incident tracker

7.3. Instrumentation and Dashboards

Telemetry is surfaced through an automated SLA Heatmap, providing a color-coded visualization of pipeline health over time. This dashboard enables data engineers to identify failure hotspots and correlate them with upstream OT network events or downstream warehouse maintenance windows. By moving from reactive troubleshooting to proactive monitoring, the framework ensures that the data remains trustworthy for executive decision making.

8. Case Study: Utility Scale Implementation

To validate the proposed framework, I conducted a pilot implementation at a large scale North American utility provider. The organization required the integration of telemetry from over 40 electrical substations into a central Snowflake data lake to support predictive transformer maintenance which was previously hindered by data silos and high latency.

8.1. Implementation Environment

The utility's OT environment managed approximately 250,000 tags within an AVEVA PI System. I utilized the PI AF to model transformer assets, specifically capturing oil temperature, load, and ambient conditions. Event Frames were configured to trigger on excursion events, such as temperature spikes exceeding 90°C. Connectivity was established through a 1 Gbps AWS Direct Connect circuit, providing the secure, low latency foundation required for the cloud hosted Informatica IDMC and Snowflake environments.

8.2. Pipeline Configuration

The architecture was configured to handle Temporal Heterogeneity (R2) identified earlier:

- Critical Stream: 5-minutes intervals for transformer load data, utilizing Snowflake Streams to facilitate high frequency incremental merges.
- Strategic Batch: 24-hours aggregations for executive reliability reporting and long-term trend analysis.
- C. Performance Evaluation and Results

I evaluated the framework against the SLOs defined in Section VII over a 30-day window. The results, summarized in Table II, confirm that the CDC aware approach successfully met or exceeded all performance targets.

Table 2: Case Study Performance Results

Metric	Target (SLOs)	Measured Result (Mean)	Compliance Status
Freshness (P_{95})	≤ 10 mins	7.4 mins	Met
Reconciliation	≥ 99.9 %	99.98 %	Met
Anomaly Rate	< 1/day	0.4/day	Met
MTTR	≤ 30 mins	22 mins	Met

8.3. Operational Impact

The transition to this framework marked a significant departure from the utility's legacy methods. Previously, the organization relied on manual CSV exports and full table reloads, which incurred a freshness latency exceeding 24 hours and prohibitive compute costs. By implementing the CDC aware architecture (R3), the utility reduced Snowflake credit consumption by approximately 38% while gaining near real-time visibility into asset health.

The NIST aligned security controls (R4) provided the necessary documentation and technical boundaries to satisfy a North American Electric Reliability Corporation (NERC) CIP audit regarding OT/IT data segmentation. This demonstrates that the framework is not only technically performant but also regulatory compliant in highly scrutinized critical infrastructure environments.

9. Discussion

The pilot implementation at the North American utility confirms that a high fidelity, CDC aware bridge between OT and cloud analytics is technically viable and economically superior to legacy methods. However, the transition from theory to a utility scale environment revealed several architectural tensions that merit further analysis.

9.1. Semantic Fidelity and the Data Janitor Problem

The decision to utilize the PI AF as the source of truth rather than raw data tags is foundational to this framework's success. By extracting data through AF, the cloud data warehouse inherits a pre-validated physical model. This approach significantly reduces the data janitor burden typically found in enterprise analytics, where data scientists often struggle to reconcile raw sensor IDs with physical assets. As evidenced in the case study, this semantic preservation was the primary driver for the rapid deployment of transformer health models, as the data arrived in Snowflake pre-contextualized.

9.2. The Financial Imperative of Warehouse Native CDC

A distinct finding of this research is that Snowflake Streams and Tasks are not merely performance optimizations; they are requirements for financial sustainability. Traditional overwrite or full reload strategies scale poorly as tag counts grow, leading to exponential increases in Virtual Warehouse uptime. By restricting processing to MERGE only increments, I achieved approximately 38% reduction in credit consumption. This suggests that for industrial datasets, where most records are immutable time series entries, CDC is the only viable path for large scale ingestion.

9.3. Cloud Hosting vs. On-Premises PI: The Latency Trade-off

While the case study utilized a cloud hosted PI System to simplify egress, this approach introduces a dependency on edge-to-cloud bandwidth and network stability. For organizations managing strictly air gapped or bandwidth constrained remote sites, a cloud only model may be insufficient. In such scenarios, a hybrid architecture where a local PI Archive buffers data before relaying it to a cloud-based PI Collective would better satisfy the safety requirements of NIST SP 800-82 Revision 3 while ensuring operational continuity during network outages.

9.4. Limitations and Technical Constraints

Despite the success of the implementation, I identified several constraints that warrant attention:

- Schema Evolution Rigidities: While the framework handles data updates effectively, structural changes to an AF template (e.g., adding a new telemetry attribute) currently require a manual update to the Informatica IDMC mapping layer. Future iterations should explore dynamic metadata discovery to automate schema propagation.
- Stream Retention Sensitivity: Snowflake Streams are sensitive to pipeline inactivity. If a source system remains offline beyond the Snowflake data retention period (typically 14 days), the stream becomes stale and requires a manual reset, potentially leading to data gaps.
- Observability Overhead: Implementing data centric telemetry introduces its own computational cost. While this overhead was negligible in the case study compared to the reliability gains, smaller organizations must carefully balance the granularity of their SLOs with their available compute budget.

10. Conclusion and Future Work

This paper has detailed a reproducible, five-layer architectural framework designed to bridge the gap between high fidelity industrial telemetry and cloud native analytics. By integrating the semantic richness of the AVEVA PI System with the elastic scale of the Snowflake Data Cloud via Informatica IDMC, this research addresses the core technical and economic tensions of IT/OT convergence.

The research yielded three primary contributions to the field of industrial data engineering:

- Semantic Continuity: I demonstrated that extracting data through the PI AF and Event Frames ensures that operational meaning is preserved from the edge to the cloud, successfully fulfilling the requirement for Semantic Fidelity (R1).
- Architectural Efficiency: The application of warehouse native CDC specifically Snowflake Streams and Tasks proved both cost effective and performant. The implementation achieved approximately 38% reduction in compute overhead while maintaining a 95th percentile freshness of 7.4 minutes.
- Standardized Reliability: By aligning pipeline controls with NIST SP 800-82 Revision 3 and implementing a data centric observability model, this work establishes a blueprint for industrial grade reliability that satisfies both stringent security audits and operational SLOs.

The utility scale implementation confirms that this framework is a viable solution for organizations seeking to scale predictive maintenance and executive reporting without compromising data integrity or security.

Future Work: I intend to expand this research into three key areas. First, I will investigate Automated Schema Contracts to develop a mechanism for detecting and propagating upstream PI AF template changes to cloud schemas without manual intervention. Second, I plan to explore Edge-to-Cloud AI Integration, precisely the deployment of lightweight ML models at the PI edge that leverage the same semantic structures defined in this framework. Finally, a Comparative TCO Analysis will be conducted to evaluate this CDC based ELT approach against emerging real-time streaming architectures, such as Kafka-to-Snowflake, to determine the most cost-efficient path for varied industrial data volumes.

References

1. AVEVA, "AVEVA PI System Operations Data Management," 2026. [Online]. Available: https://www.aveva.com/en/products/aveva-pi-system/
2. AVEVA, "PI System Architecture, Planning and Implementation," Learning Manual (Version 2025, PI Server 2024). [Online]. Available: https://cdn.osisoft.com/learningcontent/pdfs/PISystemArchitecturePlanningAndImplementationWorkbook.pdf
3. AVEVA, "Understand Event Frames in PI AF," 2025. [Online]. Available: https://docs.aveva.com/bundle/pi-server-l-af-pse/page/1021923.html
4. AVEVA, "PI Web API Reference," Developer Documentation, 2025. [Online]. Available: https://docs.aveva.com/bundle/pi-web-api-reference/page/help.html
5. Amazon Web Services, "Guidance for Hosting AVEVA PI System on AWS," 2026. [Online]. Available: https://aws.amazon.com/solutions/guidance/hosting-aveva-pi-system-on-aws/

6. National Institute of Standards and Technology (NIST), "Guide to Operational Technology (OT) Security," NIST Special Publication 800-82, Revision 3, 2023. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.pdf
7. Informatica, "Informatica Intelligent Data Management Cloud – Cloud Data Integration Documentation," 2025. [Online]. Available: https://docs.informatica.com/integration-cloud/cloud-data-integration/current-version/introduction/informatica-resources/informatica-documentation.html
8. Snowflake, "Introduction to Streams and Tasks," 2026. [Online]. Available: https://docs.snowflake.com/en/user-guide/data-pipelines-intro
9. S. Karumuri, F. Solleza, S. Zdonik, and N. Tatbul, "Towards Observability Data Management at Scale," *SIGMOD* Record, 2020. [Online]. Available: https://people.csail.mit.edu/tatbul/publications/sigmod_record20.pdf.
10. Reddy, R. R. P. (2024). Enhancing endpoint security through collaborative zero-trust integration: a multi-agent approach. International Journal of Computer Trends and Technology, 72(8), 86-90.