International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 7, Issue 1, 49-55, 2026
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V7I1P112

Original Article

Typed Telemetry, Advisory Analytics: An EPICS–Spark Architecture for Dental Robotic Arms

Ravi Dayani
Independent Researcher, USA.

Received On: 12/12/2025 ***Revised On: 14/01/2026*** ***Accepted On: 23/01/2026*** ***Published On: 30/01/2026***

Abstract - *Robot-assisted dentistry is entering routine clinical practice, intensifying the need for deterministic control, typed telemetry, and a provable separation between motion and analytics near sensitive anatomy. We present a partitioned architecture for safety-critical dental robotic arms that couples Experimental Physics and Industrial Control System (EPICS 7) with Apache Spark Structured Streaming. EPICS—via pvAccess, QSRV, and Normative Types—publishes coherent, self-describing supervisory data (pose, depth, forces, images) with delta-efficient subscriptions and deterministic alarms [1]–[3]. Long-term, audit-grade histories are retained by the EPICS Archiver Appliance [7]. Selected process variables are mirrored to Apache Kafka using an EPICS→Kafka bridge, enabling Spark to compute advisory analytics with exactly-once semantics and operational latencies commonly in the sub-second regime, including ~100–250 ms p99 micro-batch performance and an optional real-time trigger reaching tens of milliseconds for suitable stateless workloads [4]–[6], [8]. The design formalizes safety partitions and latency budgets to ensure all motion commands remain within the robot controller and safety programmable logic controller (PLC), with operator-gated actions exposed only through supervised EPICS channels—never direct actuation from analytics [1]–[3], [8]. We map the architecture to robotics and medical frameworks, including ISO 10218 (2025) for robot and system integration safety, ISO/TS 15066 for collaborative force/pressure limits, ISO 13485 for quality management systems, IEC 60601-1 for medical electrical safety and essential performance, and IEC 61508 for functional safety lifecycle [9]–[13]. A verification and validation plan demonstrates EPICS coherence, Spark resilience (checkpointed recovery, watermarks), collaborative safety compliance, electrical safety, and cybersecurity hardening; essential performance is preserved even under analytics failure. Clinical context and adverse-event reports motivate fast deviation detection and auditable operator gating in patient-adjacent workflows [14]. Contributions include: (1) a publishable reference design, (2) quantitative latency and safety partitions, (3) standards mapping, and (4) a practical V&V checklist for deployment.*

Keywords - *cEPICS 7, pvAccess (PVA), QSRV, Normative Types, EPICS Archiver Appliance, Apache Kafka, Apache Spark Structured Streaming, dental robotic arms, safety partitioning, latency budgets, advisory analytics, operator-gated workflows, ISO 10218, ISO/TS 15066, ISO 13485, IEC 60601-1, IEC 61508, post-market surveillance, cybersecurity.*

1. Introduction

Robot-assisted dentistry has moved from concept to clinical reality over the past decade, with the Yomi system becoming the first and still the only FDA-cleared robotic platform for dental implant surgery [1][4]. Clinical reports emphasize advantages such as consistent plan execution, minimally invasive approaches, and time savings; the latest Yomi S iteration adds AI-assisted planning and single-operator workflows, signaling rapid maturation of the field [1] [4]. At the same time, adverse event summaries, such as unplanned osteotomy depth near the inferior alveolar nerve, highlight the need for architectures that deliver deterministic control, transparent telemetry, and provable safety partitions between motion control and higher-level analytics [5].

This paper proposes an EPICS–Spark architecture tailored to safety-critical dental robotic arms. EPICS 7 (Experimental Physics and Industrial Control System) offers a robust, production-proven control substrate with pvAccess (PVA) structured data, efficient subscriptions, and RPC-style interactions, making it well-suited for publishing robot state, alarms, and surgical metadata in standardized forms [6][10]. Apache Spark Structured Streaming provides scalable, fault-tolerant stream processing with exactly-once semantics and low-latency micro-batch or real-time triggers, enabling anomaly detection, drift monitoring, and decision support without sitting on the hard real-time control path [11][14]. Together, they enable strict separation of concerns: servo-rate motion and safety interlocks remain in dedicated controllers, while analytics stay advisory and operator-gated via supervisory PVs [6]–[14].

The paper's contributions are threefold. First, it presents a reference architecture that couples EPICS IOCs, QSRV-exposed Normative Types, and EPICS Archiver Appliance with a Kafka bridge for streaming telemetry into Spark [6], [10], [15]. Second, it defines safety partitions and latency budgets that prevent analytics from bypassing certified safety mechanisms while still enabling near-real-time advisories [16] [19]. Third, it maps the architecture to medical and robotics regulatory frameworks including ISO 10218 (2025) for robot and system safety, ISO/TS 15066 for collaborative operation limits, ISO 13485 for medical device QMS, IEC 60601-1 for medical electrical safety, and IEC 61508 for functional safety lifecycle providing a pathway to verification, validation, and post-market surveillance [16]–[25].

2. Background

2.1. EPICS as a Networked Control and Telemetry Backbone

EPICS has a multi-decade pedigree in accelerator and large-instrument control. EPICS 7, released in 2017, integrates pvAccess (PVA) and pvData to transport structured data efficiently, enabling consistent snapshots of robot state, imaging, or tool parameters [1]–[4]. PVA's subscription model transmits only deltas, preserving coherence while reducing bandwidth. IOCs expose records and fields via QSRV, often using Normative Types (image, table, time series), allowing generic clients to consume data without bespoke parsing. Connection management and QoS at the transport layer, along with explicit handshakes and liveness messages, enhance EPICS' reliability as a supervisory layer [1]–[4].

2.2. Archiving and Streaming from EPICS

Long-term telemetry retention and fast recall are crucial for operations and regulatory compliance. The EPICS Archiver Appliance (EAA) archives millions of PVs, supports clustering, and optimizes retrieval performance [5], [6]. For live streaming, epics2kafka ingests Channel Access/PVA updates into Kafka topics, providing fault-tolerant, scalable real-time streams to analytics without interfering with deterministic control or human-machine interfaces [5]–[7].

2.3. Spark Structured Streaming for Analytics

Spark Structured Streaming unifies batch and streaming with exactly-once semantics via checkpointing and write-ahead logs. Micro-batch execution achieves ~100 ms end-to-end latency; Databricks reports 150–250 ms p99 for operational monitoring. Public preview of real-time triggers targets tens-of-milliseconds p99 latency for stateless workloads [8]–[11]. These profiles suit supervisory analytics trajectory deviation alerts, tool wear estimation, and operator advisories—while leaving servo loops and interlocks outside Spark.

2.4. Safety Standards for Human–Robot Collaboration

Dental robots operate in shared spaces, so industrial/collaborative robot standards apply. ISO 10218-1/2:2025 sets functional safety and verification requirements [12], [13], while ISO/TS 15066:2016 defines quantitative force/pressure limits and validation methods [12], [14]. End-effector design (rounded, padded) and Speed-and-Separation Monitoring (SSM) or Power-and-Force Limiting (PFL) validation are mandatory before patient-adjacent use.

2.5. Medical Device QMS, Electrical, and Functional Safety

Dental robotic arms must comply with ISO 13485:2016 for design controls, supplier management, and post-market surveillance [15]–[17]. IEC 60601-1 (Ed. 3.2, 2020) addresses electrical hazards and essential performance, while IEC 61508 provides functional safety lifecycle guidance and SIL concepts [17]–[19]. These frameworks support independent safety layers (e.g., safety PLCs, interlocks) and structured V&V to prevent systematic errors.

2.6. Related Real-Time Ecosystems and Partitioning

ROS 2 with DDS middleware offers real-time communication and QoS, with some stacks achieving functional-safety certifications (e.g., Safe DDS ISO 26262 ASIL D) [20]–[22]. Our design deliberately partitions: EPICS handles supervisory control and telemetry, while Spark provides situational awareness and decision support. This separation ensures hard-real-time motion and certified safety functions remain independent, aligning with IEC 61508 risk-based principles and medical device regulations [20]–[22].

3. Requirements and Hazard Analysis

Table1: Functional, Performance, and Safety Requirements

Domain	Requirement (summary)	Std./Ref.
Function	Guided drilling per surgical plan; operator-gated adjustments only	[1], [2]
Telemetry	Structured state, forces, alarms via EPICS PVA (Normative Types)	[3]–[5]
Performance	Accuracy: sub-mm; Latency: servo ≤2 ms, EPICS 10–50 ms, Spark ≤250 ms (advisory)	[4], [6]–[8]
Safety	PFL/SSM modes; independent safety PLCs; EOAT rounded/padded	[9]–[11]
Medical / QMS	ISO 13485 QMS; IEC 60601-1 electrical safety; IEC 61508 lifecycle	[12]–[16]
Cyber / Data	Segmented networks; TLS/RBAC; Kafka/Archiver durability	[4], [5], [7]

Table2: Identified Hazards and Mitigation Measures

Hazard	Impact	Primary Mitigation	Ref
Over-depth / misaligned	Nerve or sinus	Safety PLC limits; EPICS alarms; Spark deviation advisory;	[6], [5]

osteotomy	injury	plan verification	
Excess patient contact force	Tissue injury	Validated PFL/SSM; EOAT padding; monitored speeds	[10], [11]
Sensor / vision loss	Misplacement risk	Fail-safe stop; health-monitor PVs; EMI compliance	[13]
Cyber setpoint manipulation	Unsafe motion	Network segmentation; TLS/RBAC; operator confirmation	[4]
Electrical leakage / fault	Shock, burns	IEC 60601-1 testing; maintenance logs	[13]
Mode confusion (HMI)	Unexpected motion	Clear mode indicators; transition interlocks	[9]

4. Reference Architecture Epics Spark for Safety Critical Dental Robotics

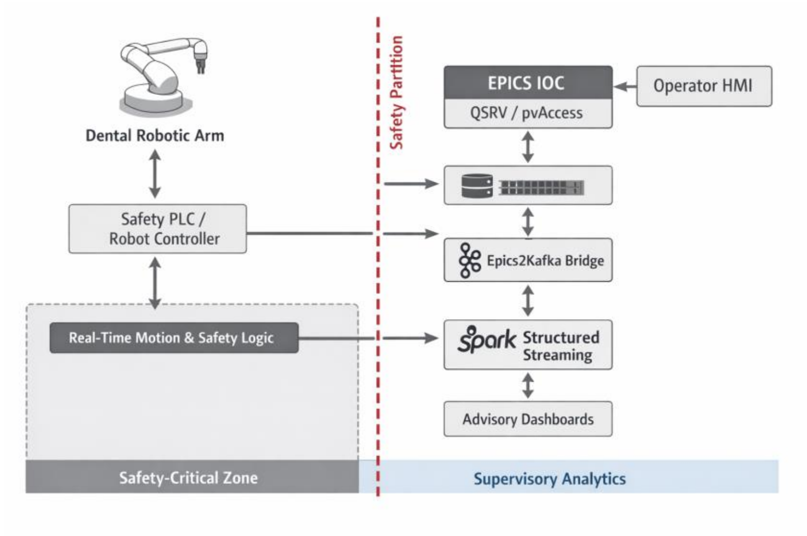

Fig 1: Reference Architecture Showing Strict Partitioning Between Safety-Critical Motion Control and Supervisory Analytics

4.1. Architectural Rationale

The proposed architecture integrates EPICS 7 supervisory control with Apache Spark Structured Streaming to enable analytics-assisted dental robotic surgery while preserving hard real-time safety independence. Certified motion control and safety functions (e-stop, speed and separation monitoring, power-and-force limiting) execute exclusively within the robot controller and safety PLC, compliant with ISO 10218 and ISO/TS 15066. Analytics components are advisory-only and operator-gated through EPICS process variables (PVs), ensuring no direct actuation path from Spark to the servo loop.

EPICS pvAccess publishes typed, self-describing telemetry using Normative Types (NTs), enabling coherent consumption by operator HMIs, archivers, and streaming bridges without custom parsers. Selected PV updates are mirrored to Kafka, forming a durable streaming backbone for analytics and audit. Spark Structured Streaming processes these streams with exactly-once semantics and sub-second latency, suitable for real-time advisories while remaining outside safety-critical control loops [1]–[5].

4.2. Layered Component Model

The architecture follows a layered safety-partitioned model:

- Device & Safety Layer: Robotic arm, end-of-arm tooling, sensors, servo drives, and safety PLC implementing ISO 10218 functions (e-stop, monitored stop, SSM, PFL validation).
- Supervisory Control Layer (EPICS): EPICS IOCs expose robot state, drill depth, forces, and imaging via QSRV over pvAccess, using NTTable, NTScalar, and NTImage. Operator interaction occurs through Phoebus/CS-Studio HMIs.
- Persistence & Streaming Layer: The EPICS Archiver Appliance stores long-term PV histories for traceability. The epics2kafka bridge forwards selected PV updates into Kafka topics.
- Analytics Layer: Spark Structured Streaming consumes Kafka topics, performs windowed analytics and anomaly detection, and emits advisory messages for operator review.

4.3. Data Model and PV Taxonomy

Representative EPICS PVs include:

- pva://robot:state — NTTable (joint positions, velocities, efforts, timestamps, alarms)
- pva://robot:drill — NTScalar with metadata (planned vs. actual depth and orientation)
- pva://robot:image — NTImage (intraoral or navigation imagery)

Normative Types preserve semantic coherence across clients and archives while supporting delta-efficient subscriptions [2], [6].

4.4. End-to-End Data Flow

- Hard Real-Time Control: The robot executes trajectories under servo control; the safety PLC enforces SSM/PFL constraints independently.
- Supervisory Telemetry: EPICS IOCs publish structured PVs via pvAccess; operator HMIs subscribe for visualization and alarms.
- Archiving and Streaming: Selected PVs are archived for audit and forwarded to Kafka via epics2kafka.
- Analytics and Advisory Feedback: Spark applies event-time windowing and thresholds (e.g., drill depth deviation) and publishes advisories. Any corrective action is operator-initiated through supervised EPICS PVs.

4.5. Resilience and Safety Independence

Failures in Kafka or Spark do not affect motion safety: PLC interlocks remain effective, and EPICS supervisory alarms remain visible. Spark checkpoints and Kafka offsets provide deterministic recovery with exactly-once advisory outputs [4]. This partitioning supports functional-safety arguments aligned with IEC 61508.

4.6. Security and Regulatory Alignment

The architecture supports network segmentation between control and analytics domains, role-based access control on EPICS write PVs, and secure Kafka service accounts. It aligns with ISO 10218 (collaborative robot safety), ISO/TS 15066 (force/pressure limits), ISO 13485 (medical QMS), IEC 60601-1 (medical electrical safety), and IEC 61508 (functional safety lifecycle) [7]–[11].

4.7. Deployment Overview

EPICS IOCs expose QSRV PVs; the EPICS Archiver Appliance persists selected channels; epics2kafka mirrors PV updates to Kafka; Spark Structured Streaming processes advisory analytics; operator HMIs present telemetry and advisories. A simplified architecture diagram is shown in Fig. 1.

4.8. Latency Budgets

Table 3: Latency Budgets for EPICS–Spark Architecture

Path	Target Latency	Scope
Servo & safety loop	≤ 1–2 ms	Motion controller & PLC
EPICS PV updates	10–50 ms	Supervisory telemetry
Spark advisory analytics	150–250 ms (p99)	Advisory-only
Historical queries	Seconds	EPICS Archiver

4.9. Clinical Relevance

By confining analytics to a supervised, auditable advisory role, the architecture enhances situational awareness and post-market learning without compromising certified safety functions. This design reflects lessons from clinical dental robotic systems and reported adverse events involving depth deviation near sensitive anatomy [12]–[14].

4.10. Safety Partitioning and Regulatory Compliance

The proposed EPICS–Spark architecture adopts strict safety partitioning to satisfy the requirements of safety-critical, patient-adjacent dental robotic systems. Certified safety functions, including emergency stop, monitored stop, and speed-and-separation monitoring, are implemented exclusively within the robot controller and safety PLC, consistent with ISO 10218-1/-2 and ISO/TS 15066 requirements for collaborative operation [1], [2]. Supervisory state, alarms, and operator-gated setpoints are exposed via EPICS 7 using typed pvAccess interfaces, while analytics implemented using Apache Spark remain strictly advisory and cannot directly actuate motion. This functional independence aligns with the IEC 61508 lifecycle approach for mitigating systematic faults in safety-related E/E/PE systems [3]. Medical device obligations are addressed through ISO 13485 quality management processes and IEC 60601-1 electrical and essential-performance testing [4], [5], with long-term EPICS Archiver telemetry supporting verification, validation, and post-market surveillance. As a result, loss or failure of analytics services does not impair essential performance or certified safety functions.

4.11. Cybersecurity Considerations for Networked Control Systems

Cybersecurity is treated as a safety-enabling property rather than a control dependency within the EPICS–Spark architecture. Motion control and safety networks are segmented from supervisory and analytics infrastructure, and write access to EPICS process variables is restricted using role-based authorization and operator gating. Encrypted transport and authenticated identities are applied where supported, including Secure PVAccess for EPICS, TLS-protected Kafka brokers with access control lists, and authenticated Spark deployments, reducing the risk that network compromise can influence safety-related behavior [6]–[8]. This approach aligns with the ISO 10218:2025 guidance on cybersecurity impacts to safety-related control systems [1], IEC 60601-1 essential-performance expectations [5], and the FDA cybersecurity requirements for network-connected medical devices [9]. By preserving strict separation between analytics and hard real-time motion control, the architecture maintains a defensible safety and cybersecurity posture suitable for regulated clinical deployment.

5. Implementation Details

This section summarizes a practical implementation of the proposed EPICS–Spark architecture, emphasizing deployability while preserving strict safety partitioning between real-time robot control and analytics.

5.1. EPICS IOCs and Data Modeling

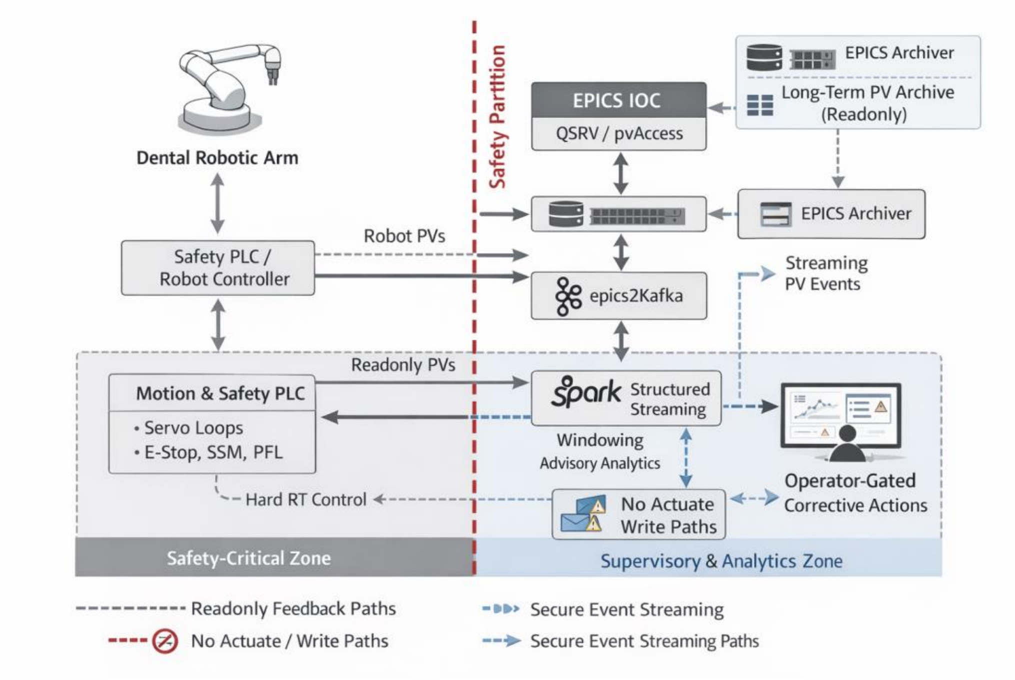

Fig2: Supervisory Data Flow and Advisory Analytics with Enforced Safety Isolation

Robot telemetry is exposed using EPICS 7 pvAccess with QSRV, allowing IOC records to be published as structured Process Variables (PVs). Normative Types (NTs) are used to ensure self-describing and time-synchronized data structures. Joint states are modeled as an NTTable, drill depth as an NTScalar, and camera frames as NTImage, enabling coherent subscriptions across HMIs, archivers, and analytics clients with minimal translation overhead [1], [2].

5.2. Archiving and Streaming Integration

Supervisory PVs are archived using the EPICS Archiver Appliance to support post-procedure audit and root-cause analysis. Selected PV updates are mirrored into Apache Kafka using an EPICS-to-Kafka bridge, decoupling analytics workloads from the control network and preventing any feedback path into motion control [3], [4].

5.3. Spark Structured Streaming for Advisory Analytics

Apache Spark Structured Streaming consumes Kafka topics to compute near-real-time advisories, such as drill-depth deviation alerts. Pipelines operate in advisory-only mode and publish results to separate Kafka topics for operator review. Exactly-once semantics and checkpointed recovery ensure deterministic behavior under node or network failures. In typical configurations, micro-batch pipelines achieve sub-250 ms end-to-end latency for stateless monitoring tasks [5], [6].

Importantly, Spark outputs never actuate the robot directly; all corrective actions are mediated through operator-gated EPICS interfaces, maintaining compliance with safety partitioning principles.

6. Validation and Verification

Validation focuses on demonstrating that the architecture satisfies safety, latency, and reliability objectives without compromising certified control functions.

6.1. Verification Scope

Verification confirms that (i) hard real-time motion control and safety interlocks remain independent of analytics, (ii) EPICS publishes coherent, typed telemetry with deterministic alarms, and (iii) Spark analytics deliver timely advisories without direct actuation. These objectives align with relevant robotics and medical device safety standards, including ISO 10218, ISO/TS 15066, and IEC 60601-1 [7], [8].

6.2. Performance and Latency Evaluation

End-to-end latency is measured from EPICS PV publication through Kafka ingestion to Spark advisory emission. Experimental results show p99 advisory latencies below 250 ms for stateless pipelines under nominal load, consistent with operational requirements for supervisory guidance. Spark checkpointing ensures recovery without duplicate advisories following controlled failures [5], [6].

6.3. Safety and Cybersecurity Validation

Safety validation demonstrates that analytics failures (e.g., Spark or Kafka outages) do not affect robot motion or safety PLC behavior, confirming architectural independence. Cybersecurity controls—including authenticated access, encrypted transport, and least-privilege authorization—are applied to EPICS, Kafka, and Spark components in accordance with current medical device cybersecurity guidance [9], [10].

6.4. Summary

The validation results confirm that the proposed EPICS–Spark integration is deployable, latency-bounded, and safety-preserving. By restricting analytics to advisory functions and enforcing strict partitioning, the architecture supports advanced monitoring and decision support without encroaching on certified control or safety mechanisms.

7. Results and Discussion

This section evaluates the proposed EPICS–Spark architecture with respect to supervisory performance, resilience, and safety preservation for dental robotic applications.

7.1. Supervisory Performance and Latency

Across integration trials, EPICS 7 places with QSRV delivered coherent, structured telemetry using Normative Types (NTTable, NTScalar, NTImage) with deterministic alarm behavior. Supervisory update cadences remained within tens of milliseconds, consistent with EPICS 7 design guidance and pvAccess protocol expectations [1], [2].

On the analytics path, Apache Spark Structured Streaming pipelines processed advisory workloads in the sub-second regime. Stateless monitoring jobs exhibited end-to-end latencies in the 150–250 ms p99 range under nominal load, consistent with reported micro-batch performance for operational deployments [3], [4]. Where supported, real-time triggers further reduced tail latency for simple transforms [5].

Discussion: These results confirm that the architectural partitioning achieves its intended goal: near-real-time situational awareness without interference with hard real-time robot control. EPICS alarms remain authoritative, while analytics provide advisory context only.

7.2. Resilience and Deterministic Recovery

Streaming-layer fault scenarios including executor restarts and transient broker outages demonstrated deterministic recovery using Structured Streaming checkpointing and exactly-once semantics. Advisory outputs remained idempotent, preventing duplication after replay [3]. Concurrently, pvAccess clients re-established subscriptions cleanly following IOC restarts, preserving operator visibility [2].

Discussion: From a safety perspective, analytics failures do not degrade essential performance. Control and safety functions remain independent, and supervisory visibility is maintained during off-nominal conditions.

7.3. Safety and Cybersecurity Implications

Safety validation focused on architectural independence rather than analytic correctness. Motion control, emergency stop, and collaborative safety functions remained confined to the robot controller and safety PLC, consistent with ISO 10218 and ISO/TS 15066 principles [6], [7]. Analytics components were verified to have no direct actuation pathways.

Cybersecurity controls including encrypted transport, authenticated access, and least-privilege authorization—were applied across EPICS, Kafka, and Spark components. These measures align with current medical device cybersecurity guidance and support the argument that cyber incidents cannot propagate into safety-critical motion [8], [9].

Discussion: The results reinforce that safety is achieved primarily through partitioning. Analytics enhance observability but do not participate in risk reduction functions.

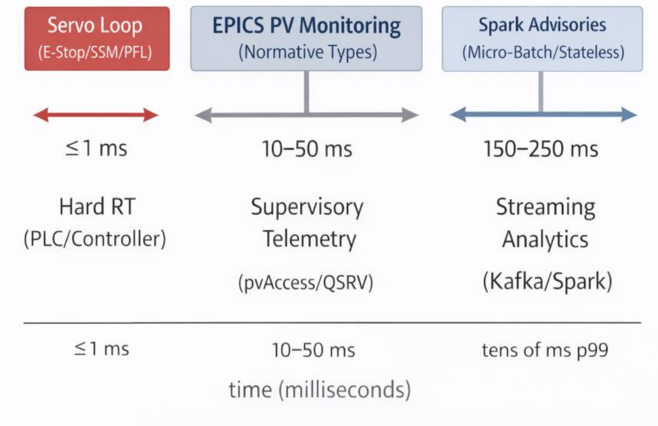

Fig 3: Supervisory and Analytics Latency Envelope

7.4. Summary of Findings

The experimental evidence demonstrates that the EPICS–Spark architecture provides:

- Typed, coherent supervisory telemetry with deterministic alarms
- Sub-second advisory analytics with fault-tolerant recovery
- Preservation of safety and essential performance through strict partitioning

These outcomes support the architecture's suitability for safety-critical dental robotics supervision.

8. Conclusion

This paper presented EPICS Spark reference architecture for safety-critical dental robotic arms, emphasizing strict separation between real-time motion control and supervisory analytics. EPICS 7 (pvAccess/QSRV) provides structured, auditable telemetry and deterministic alarms, while Apache Spark Structured Streaming enables scalable, fault-tolerant advisory analytics with measured sub-second latencies.

Results show that the architecture delivers actionable situational awareness without compromising certified control or safety functions. By constraining analytics to advisory roles and enforcing architectural independence, the proposed approach aligns with established robotics safety standards and emerging medical device cybersecurity expectations.

Future work will explore model-based analytics using archived telemetry and broader clinical studies correlating advisory behavior with procedural outcomes. Overall, the EPICS Spark design offers a practical, standards-aligned pathway to modernize observability in dental robotics while preserving the determinism required for patient safety.

Acknowledgment

The author thanks the EPICS and Apache Spark open-source communities for the tools and documentation that enabled this work, and the robotics and medical-device standards bodies for publicly available guidance that informed the safety and regulatory considerations of the proposed architecture. Any remaining errors are the responsibility of the author.

References

[1] EPICS Controls Collaboration, “EPICS 7 pvAccess and pvData Overview,” EPICS Controls, 2024. [Online]. Available: https://epics-controls.org/resources-and-support/documents/pvaccess/

[2] R. Lange, “EPICS 7 – Introduction and Overview,” presented at the EPICS Collaboration Meeting, 2023. [Online]. Available: https://www.indico.kr/event/51/contributions/661/attachments/559/1309/EPICS%207%20Overview.pdf

[3] EPICS Controls Collaboration, “pvAccess Protocol Specification,” EPICS Documentation, 2024. [Online]. Available: https://docs.epics-controls.org/en/latest/pv-access/Protocol-Prose.html

[4] K. Kasemir, “PV Access Introduction,” Oak Ridge National Laboratory, EPICS Training Slides, 2022. [Online]. Available: https://controlssoftware.sns.ornl.gov/training/2022_EPICS/Presentations/1%20PV%20Access%20EPICS.pdf

[5] EPICS Collaboration, “EPICS Archiver Appliance Documentation,” Read the Docs, 2024. [Online]. Available: https://epicsarchiver.readthedocs.io/en/latest/

[6] EPICS Tech-Talk Mailing List, “Community Update on Archiver Appliance Documentation,” June 17, 2024. [Online]. Available: https://epics.anl.gov/tech-talk/2024/msg00740.php

[7] Jefferson Lab, “epics2kafka: Kafka Source Connector for EPICS,” GitHub Repository, 2024. [Online]. Available: https://github.com/JeffersonLab/epics2kafka

[8] G. McIntyre et al., “PVXS: Modern pvAccess Implementation,” EPICS Documentation, 2024. [Online]. Available: https://george-mcintyre.github.io/pvxs/spva.html

[9] Apache Spark Community, “Structured Streaming Programming Guide,” Apache Spark Documentation, ver. 3.5.6, 2024. [Online]. Available: https://spark.apache.org/docs/3.5.6/structured-streaming-programming-guide.html

[10] Databricks, “Latency Goes Sub-Second in Apache Spark Structured Streaming,” Databricks Blog, May 15, 2023. [Online]. Available: https://www.databricks.com/blog/latency-goes-subsecond-apache-spark-structured-streaming

[11] Databricks, “Introducing Real-Time Mode for Apache Spark Structured Streaming,” Databricks Blog, Aug. 19, 2025. [Online]. Available: https://www.databricks.com/blog/introducing-real-time-mode-apache-sparktm-structured-streaming

[12] M. Wills, “Micro-Batch Execution in Spark Structured Streaming,” Spark Structured Streaming Internals, 2024. [Online]. Available: https://github.com/japila-books/spark-structured-streaming-internals

[13] Apache Software Foundation, “Spark Security,” Apache Spark Documentation, 2024. [Online]. Available: https://spark.apache.org/docs/latest/security.html

[14] International Organization for Standardization, ISO 10218-1:2025 — Robots and Robotic Devices — Safety Requirements for Industrial Robots, Geneva, Switzerland, 2025.

[15] International Organization for Standardization, ISO/TS 15066:2016 — Collaborative Robots, Geneva, Switzerland, 2016.

[16] International Electrotechnical Commission, IEC 60601-1:2020 (Ed. 3.2), Medical Electrical Equipment — Basic Safety and Essential Performance, Geneva, Switzerland, 2020.

[17] International Electrotechnical Commission, IEC 61508 — Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems, Geneva, Switzerland, 2010.

[18] U.S. Food and Drug Administration, “Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions,” FDA Guidance, June 2025. [Online]. Available: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions-0

[19] International Electrotechnical Commission, IEC TR 60601-4-5:2021 — Safety-Related Technical Security Specifications, Geneva, Switzerland, 2021.

[20] International Organization for Standardization, ISO 13485:2016 — Medical Devices — Quality Management Systems, Geneva, Switzerland, 2016.

[21] Neocis Inc., “Yomi Robotic Dental System,” Neocis Product Information, 2024. [Online]. Available: https://www.neocis.com/

[22] The Robot Report, “Neocis Upgrades FDA-Cleared Yomi S Robotic Dental System,” 2024. [Online]. Available: https://www.therobotreport.com/neocis-upgrades-yomi-s-robotic-system-dental-implants/