

# MODERN SCADA ARCHITECTURES: UNLOCKING VALUE FROM INDUSTRIAL ASSETS

Trenton Franklin  
AUTOSOL  
2026

## EXECUTIVE SUMMARY

Supervisory Control and Data Acquisition (SCADA) systems underpin the safe and efficient operation of nearly every critical infrastructure sector, from wellheads to interstate natural gas pipelines. Over six decades, these systems have evolved through four distinct architectural generations, each driven by the limitations of its predecessor. Today, the industry stands at its most consequential inflection point: the convergence of cloud computing, edge intelligence, MQTT-based messaging, and artificial intelligence is reshaping not only how SCADA systems are built, but what they are capable of. The global SCADA market, valued at roughly 13 billion in 2025, is projected to reach \$20 billion by 2030 at a compound annual growth rate of 9.2%.<sup>1</sup>

This paper traces the SCADA evolution chronologically, examines the structural tensions created by modern oil and gas (O&G) operational demands, and evaluates the technologies emerging to meet them. It is written for engineers, operators, and technical leaders who interact with these systems daily and for the executives who fund them.

## 1. FOUR GENERATIONS OF SCADA ARCHITECTURE

SCADA systems emerged in the 1960s as centralized, mainframe-driven platforms designed to monitor and control geographically dispersed field equipment. The architecture of each generation reflects the computing hardware, networking standards, and operational priorities of its time.<sup>2</sup>

Figure 1 maps the four-generation chronology as a timeline. Table 1 provides a structured comparison of architectural attributes across generations, including a fifth emerging generation that represents a meaningful architectural discontinuity from the IIoT substrate of Gen 4.




Figure 1: Four Generations of SCADA Architecture: Key Technology Milestones

| Generation | Era         | Architecture                    | Key Technology                        | Primary Limitation               |
|------------|-------------|---------------------------------|---------------------------------------|----------------------------------|
| Gen 1      | 1960s-1970s | Centralized mainframe           | Proprietary RTUs, analog telemetry    | Cost, single point of failure    |
| Gen 2      | 1980s-1990s | Distributed - PDP/VAX computers | PC-based HMI, Modbus, DNP3            | Vendor lock-in, limited interop  |
| Gen 3      | 2000s-2010s | Networked SCADA, Purdue Model   | Ethernet, OPC-DA, historian databases | Air-gap erosion, IT/OT collision |

| Generation       | Era           | Architecture                      | Key Technology                                     | Primary Limitation                                             |
|------------------|---------------|-----------------------------------|----------------------------------------------------|----------------------------------------------------------------|
| Gen 4            | 2010s-2020s   | Cloud / Edge / IIoT hybrid        | MQTT, Sparkplug B, edge computing                  | Security surface, management complexity                        |
| Gen 5 (emerging) | 2020s-present | Autonomous / AI-native operations | LLMs, digital twins, self-optimizing control loops | Governance, liability, operator trust; industry debate ongoing |

**Table 1: SCADA Generations: Architecture, Technology, and Limitation**

The first generation relied on proprietary hardware and analog telemetry, with little interoperability. The second generation introduced distributed computers (such as PDP / VAX) and PC-based HMIs, reducing hardware costs but deepening vendor lock-in. The third generation, enabled by Ethernet and enterprise/control integration standards such as ISA-95<sup>3</sup>, produced the networked SCADA environments that most O&G operators still run today. The fourth-generation layers cloud connectivity and edge computing atop that installed base. What separates a potential fifth generation from Gen 4 is not additive capability but a change in the nature of the control relationship: from human-supervised automation to systems that reason, self-optimize, and increasingly act.

Figure 2 frames the Gen 5 hypothesis; its distinguishing characteristics, the technologies enabling it, and the governance questions the industry has not yet resolved.




**Figure 2: Gen 5 SCADA: Characteristics, Technologies, and Unresolved Questions**

The critical observation is that Gen 5 is not a clean break from Gen 4; it is built on top of the same edge, MQTT, and cloud platform. What changes is what the system does with the data once it arrives. That distinction matters for operators making architecture investments today: the infrastructure choices made in Gen 4 either enable or foreclose Gen 5 capabilities.

## 2. THE PURDUE MODEL: FOUNDATION, STRENGTHS, AND MODERN FRICTION

The Purdue Enterprise Reference Architecture (PERA), codified in ISA-95 and commonly referred to as the Purdue Model, provides the dominant framework for understanding separation between operational technology (OT) and information technology (IT) in industrial environments.<sup>3</sup> Its hierarchical zone structure, from physical process devices at Level 0 through enterprise business systems at Level 4, was designed to minimize cyber exposure and protect safety-critical systems from IT-side failures and threats.

Figure 3 illustrates the Purdue Model hierarchy with the DMZ correctly positioned at the IT/OT boundary. Table 2 maps each level to its components and the specific friction it faces in modern deployments.




**Figure 3: Purdue Model Hierarchy with DMZ Placement at IT/OT Boundary**

| Level | Name               | Components                                  | Modern Challenge                                                       |
|-------|--------------------|---------------------------------------------|------------------------------------------------------------------------|
| 0     | Physical Process   | Sensors, actuators, valves                  | Embedded edge intelligence blurs this boundary                         |
| 1     | Field Control      | PLCs, RTUs                                  | Edge nodes adding compute at this layer                                |
| 2     | Supervisory        | HMI, local SCADA                            | Often bypassed in cloud-direct architectures                           |
| 3     | Operations / Site  | Historian, MES, batch control, SCADA Server | DMZ (Level 3.5) sits above at IT/OT boundary                           |
| 3.5   | DMZ Boundary       | Firewalls, Gateways, Replication Servers    | Blurred by cloud-native architectures, bi-directional trust challenges |
| 4     | Business Planning  | ERP, EAM, corporate IT                      | Cloud integration increasingly common                                  |
| 5     | Enterprise Network | Corporate WAN, internet                     | Expanded attack surface via IT/OT convergence                          |

**Table 2: Purdue Model Levels, Components, and Modern Challenges**

The industrial demilitarized zone (DMZ), sometimes called Level 3.5, sits at the boundary between OT (Level 3) and IT (Level 4), not between the supervisory and operations layers. This distinction matters for correctly scoping firewall rules, historian

replication, and data diode deployments. The model remains foundational. Its logical separation of zones and network paths provide a durable framework for IEC 62443-compliant architectures.<sup>7</sup>

However, the assumption of clean layer separation has eroded. Cloud-direct sensor connections, remote access requirements, and demand for real-time production data have created paths that bypass intended boundaries. As AUTOSOL and Neeve co-presented at ENTELEC 2026, the Purdue Model is not dead; it must be reinterpreted through the lens of zero-trust principles and software-defined perimeters rather than treated as a static physical air-gap<sup>10</sup>. Regulators are moving the same direction. Frameworks like NIST CSF 2.0, the TSA Pipeline Security Directives, and CISA's Zero Trust Maturity Model expect segmentation based on identity and access, not just fixed network layers. As a result, real compliance is shifting from what diagrams say to what systems enforce in day-to-day operation.

3. THE POLL/RESPONSE BOTTLENECK AND REPORT-BY-EXCEPTION

For most of SCADA history, master stations communicated with remote terminal units (RTUs) and programmable logic controllers (PLCs) via poll/response: the master queries each device on a fixed schedule, and the device returns its current state. This model is deterministic and simple, but it does not scale gracefully. In many architectures a master polling 500 RTUs every 60 seconds is not actually collecting real-time data; it is collecting a 60-second stale snapshot from each device, sequentially.

Figure 4 illustrates the architectural difference between poll/response and MQTT report-by-exception data flows. Table 3 compares the two approaches across key attributes.




Figure 4: Poll/Response vs. MQTT Report-by-Exception Data Flow

| Attribute         | Poll/Response                        | MQTT / Report-by-Exception                   |
|-------------------|--------------------------------------|----------------------------------------------|
| Data flow trigger | Master polls each device on schedule | Device publishes only on change or threshold |
| Bandwidth usage   | High and predictable                 | Low, spiky on events                         |
| Latency           | Bounded by poll interval             | Near real-time on change                     |
| Scalability       | Limited by master poll capacity      | Scales horizontally via broker               |
| Protocol examples | Modbus, DNP3, ICCC                   | MQTT, MQTT Sparkplug B                       |

| Attribute        | Poll/Response                                 | MQTT / Report-by-Exception                              |
|------------------|-----------------------------------------------|---------------------------------------------------------|
| Security posture | Implicit trust within network zone            | mTLS, per-client ACLs, topic namespace governance       |
| Best fit         | Deterministic process control, small networks | Large distributed IIoT, bandwidth-constrained telemetry |

**Table 3: Poll/Response vs. MQTT Report-by-Exception**

MQTT (Message Queuing Telemetry Transport), developed originally for satellite pipeline monitoring and now an OASIS standard, inverts this model.<sup>9</sup> Field devices publish data only when a value changes beyond a configured threshold, a pattern called report-by-exception (RBE). A broker receives and routes published messages to any authorized subscriber. The result is dramatically lower average bandwidth consumption and near-real-time event propagation, at the cost of added architectural complexity and broker reliability requirements.

The Sparkplug B specification, maintained by the Eclipse Foundation, extends MQTT with a standardized payload schema and state management protocol specifically designed for industrial SCADA use cases.<sup>8</sup> Sparkplug B defines birth/death certificates for devices, enabling the broker to track which devices are online and what data they last reported; capabilities that raw MQTT does not provide natively.

#### 4. SENSOR-TO-CLOUD GOVERNANCE AND THE MQTT BROKER AS ENFORCEMENT POINT

The move to cloud-connected SCADA introduces a governance challenge that is often underestimated in early architecture discussions: who controls what data leaves the OT environment, when, and in what form? In traditional poll-based architectures, the master station implicitly controlled data collection scope. In MQTT-based cloud architectures, that control must be deliberately engineered into the broker layer.

Figure 5 illustrates the full edge-to-cloud architecture, with the MQTT broker positioned at the DMZ boundary as a governance enforcement point.




**Figure 5: Edge-to-Cloud Architecture with MQTT Broker as Governance Enforcement Point**

A MQTT broker coupled with proper certificate lifecycle management functions as a governance enforcement point, not merely a message bus. Mutual TLS (mTLS) authentication ensures that only provisioned devices with valid certificates can connect. Per-client access control lists (ACLs) define which devices may publish to which topic namespaces, preventing lateral data

leakage between production assets. A structured topic hierarchy provides an auditable record of data provenance. Combined with unidirectional trust per host, this architecture supports a zero-trust data plane without reliance on physical air gaps.

Nozomi Networks' 2024 OT/IoT Security Report found that 62% of OT environments had at least one internet-connected device in a zone not designed for external connectivity.<sup>15</sup> Broker-enforced governance directly addresses this exposure class.

5. O&G TELEMETRY REALITIES: FROM RADIO TO LEO SATELLITE

Oil and gas field operations present telemetry challenges that differ materially from manufacturing or utility environments. Production sites are often remote, temporary, and geographically dispersed across hundreds of miles. Table 4 compares the telemetry options available to O&G operators today, from legacy serial links to low-Earth orbit satellite.

| Medium                   | Typical Bandwidth | Latency      | Coverage                 | O&G Relevance                       |
|--------------------------|-------------------|--------------|--------------------------|-------------------------------------|
| Serial / Leased Line     | 9.6-56 Kbps       | Low          | Point-to-point           | Legacy infrastructure, still common |
| Licensed Radio (UHF/VHF) | 9.6-128 Kbps      | Low          | Line-of-sight            | Dominant in remote fields           |
| Cellular (4G/LTE)        | 10-100+ Mbps      | Low-Medium   | Wide (carrier-dependent) | Primary modern telemetry path       |
| LEO Satellite (Starlink) | 50-200+ Mbps      | Low (~20 ms) | Global                   | Disruptive for remote/offshore      |

Table 4: O&G Field Telemetry Options: Bandwidth, Latency, and Coverage

Licensed radio (UHF/VHF) remains common across mature producing basins. Cellular (4G/LTE) has become the dominant new deployment medium. Low-Earth orbit (LEO) satellite constellations, led commercially by SpaceX's Starlink, represent one of the more disruptive developments in O&G telemetry in decades.<sup>22</sup> Starlink's ability to deliver 50-200+ Mbps with sub-30ms latency at remote sites without ground infrastructure changes the economic calculus for edge computing, high-frequency data collection, and remote video-based inspection.

6. EDGE COMPUTING: THE RELIEF VALVE AND ITS TRADEOFFS

Edge computing, the execution of compute workloads at or near the data source rather than in a centralized data center, addresses several simultaneous constraints: bandwidth saturation, latency requirements, and cloud dependency. For O&G operations, edge nodes deployed at compressor stations, meter runs, or production pads can aggregate data from dozens of local RTUs, run local analytics, and publish only meaningful results upstream.<sup>14</sup>

LNS Research found that industrial companies deploying edge analytics reduced unplanned downtime by an average of 23% and cut data transmission costs by 40-60% compared to cloud-direct polling architectures.<sup>12</sup> However, edge computing introduces management complexity that many O&G operators are unprepared for. An edge fleet of 500 nodes across a producing basin represents 500 independent software update targets, 500 certificate lifecycle management obligations, and 500 potential security ingress points if misconfigured.

The architectural question is not whether edge computing is beneficial; it demonstrably is. The real question is whether the operator's organization has the operational maturity to manage it responsibly. Zero-trust edge architectures, as discussed at ENTELEC 2026, directly address this by treating every edge node as an untrusted endpoint that must authenticate continuously rather than as a trusted internal device.<sup>10</sup>

7. AI AND PREDICTIVE MAINTENANCE: THE ROI CASE

Artificial intelligence applications in O&G SCADA environments are maturing from proof-of-concept to production deployment. The most commercially validated use case remains predictive maintenance: using historical and real-time sensor data to identify equipment degradation signatures before they produce failures.<sup>16</sup> This capability sits at the leading edge of what may become Gen 5 SCADA; systems that do not merely report equipment state but actively reason about it.

McKinsey & Company estimated that AI-driven predictive maintenance can reduce maintenance costs by 10-25%, extend equipment life by 20-40%, and reduce unplanned downtime by up to 50% in capital-intensive industries.<sup>13</sup> For upstream O&G, where compressor failures or wellhead equipment outages translate directly to deferred production, these figures represent material financial outcomes.

To ground these percentages in dollars: the U.S. Department of Energy's Federal Energy Management Program found that a functional predictive maintenance program can yield savings of 30-40% over reactive maintenance strategies.<sup>24</sup> For a midsize upstream operator spending \$8-12M annually on field maintenance across 300 well sites, that range implies \$2.4-4.8M in recoverable cost. Deloitte Insights reported that a pilot deployment of predictive analytics on a single asset class at a large chemical manufacturer produced an 80% reduction in unplanned downtime and approximately \$300,000 in savings per asset.<sup>25</sup> At enterprise scale, Shell's deployment of AI-driven predictive maintenance, powered by C3 AI, now monitors more than 10,000 pieces of equipment across upstream, manufacturing, and integrated gas assets globally, ingesting 20 billion rows of sensor data per week and generating over 15 million failure predictions daily. Shell has described the ROI as compelling, with equipment and maintenance representing a significant percentage of operating costs that AI-based monitoring enables them to lower costs through reduced production interruptions, avoided unplanned downtime, and extended asset life.<sup>26</sup>

Saudi Aramco reported in its 2022 Integrated Report that digital production optimization programs contributed to measurable production uplift at scale across its upstream assets.<sup>18</sup> Aberdeen Group's 2023 research found that companies with mature predictive maintenance programs achieved 25% greater uptime and 10x better fault detection rates compared to reactive maintenance organizations.<sup>16</sup>

The prerequisite for effective AI-driven maintenance is high-quality, high-frequency sensor data; precisely the data that modern MQTT and edge architectures are designed to deliver. The feedback loop is intentional: better telemetry infrastructure enables better AI, and demonstrated AI ROI justifies further telemetry investment.

## **8. CYBERSECURITY IMPLICATIONS OF THE MODERN SCADA STACK**

The architectural evolution described in this paper expands the attack surface of industrial control systems in ways that traditional perimeter-based security models do not adequately address.<sup>5</sup> Dragos' 2023 Year in Review found that ICS-specific malware variants increased by 87% year-over-year, with energy sector targets accounting for the largest share of observed incidents.<sup>5</sup>

Claroty's 2024 Global State of CPS Security report found that 75% of OT environments experienced a cyberattack in the preceding 12 months, with an average business impact of \$1.5M per incident.<sup>6</sup> CISA's Cross-Sector Cybersecurity Performance Goals establish baseline controls for OT environments including network segmentation, secure remote access, and asset inventory management.<sup>21</sup>

IEC 62443 provides the most applicable security framework for industrial control systems.<sup>7</sup> Its risk-based, zones-and-conduits approach maps naturally to both the Purdue Model and the zero-trust edge architectures being deployed to replace or augment it. Security requirements should be treated as first-order architectural constraints, not post-deployment overlays.

## **9. WHY MOST SCADA MODERNIZATIONS STALL: EXECUTION, NOT ARCHITECTURE**

The preceding sections describe a target architecture: event-driven telemetry, broker-enforced governance, edge intelligence, and AI-driven decision support. In practice, the operators who reach that target are rarely the ones that chose the best technology stack; they are the ones that addressed the nontechnical problems first. SCADA modernization projects rarely stall on protocol selection or vendor evaluation. They stall because of data quality, the need to integrate equipment that cannot be replaced within the current budget cycle, workforce gaps, misaligned IT and OT incentives, and the operational reality that a 500-node edge fleet behaves nothing like a single SCADA server in a control room. Cybersecurity is one dimension of this readiness problem; it is not the whole of it.

Five frictions recur across the deployments I have seen. First, data governance and data quality: an MQTT broker without a unified namespace or metadata strategy creates a faster path to bad data, not better data. Sparkplug B's typed payloads help, but only when they are paired with deliberate naming conventions and clear ownership. Second, legacy integration: greenfield architectures rarely reflect brownfield realities. In practice, modernization must accommodate legacy assets that cannot be retired within the current operational or capital planning cycle. Phased coexistence through protocol gateways is the realistic path, not rip and replace. Third, workforce knowledge: the engineers who understand the existing system are often nearing retirement, while the engineers who are comfortable with modern technologies such as Kubernetes and certificate rotation often do not understand how the field operates. Fourth, IT and OT collaboration: the two organizations have spent decades optimizing for different objectives, uptime versus data integrity, and a converged architecture without converged governance produces predictable conflict at every change window. Fifth, scalability and operational management: a fleet of edge nodes is a software product the operator now owns, along with all the lifecycle obligations that implies, and most O&G operators are not staffed like software companies. For smaller companies, orchestration becomes a monumental task they are not equipped to handle.

The practical strategy is the same posture Section 8 recommends for cybersecurity: treat these issues as first-order architectural constraints, not deployment afterthoughts. That means budgeting for data governance work alongside the broker license, cross-training programs alongside the edge hardware refresh, and joint IT and OT operating models before the first cloud connection is provisioned. Frost & Sullivan's analysis of industrial IT/OT convergence concluded that the operators capturing real value from the transition are those that explicitly invest in bridging the workforce and skills gap, not those that treat it as a downstream HR problem.<sup>27</sup> The architecture described in this paper is achievable today with commercially available technology. Whether any given operator can deploy and sustain it is, far more often than not, an organizational question rather than a technical one.

## **10. CONCLUSION**

The SCADA systems that monitor and control oil and gas infrastructure today are the product of six decades of iterative architectural evolution. They are also the foundation upon which the next generation of intelligent, AI-assisted operations will be built. The transition from centralized polling architectures to distributed, event-driven, cloud-connected systems is not a discrete project with a completion date. It is an ongoing operational posture that demands continuous attention to telemetry quality, data governance, cybersecurity, and organizational capability.

The Purdue Model is not obsolete; as the ENTELEC 2026 discussion framed it, it is in need of reinterpretation.<sup>10</sup> Its logical separation of zones and conduits remains sound engineering. What has changed is the assumption that physical network boundaries can be maintained as a primary control. Zero-trust principles, MQTT broker-enforced governance, and edge intelligence are the mechanisms by which the Purdue Model's intent can be preserved in a world of cloud connectivity, remote access, and LEO satellite backhaul.

For students new to the gas measurement industry, the central lesson is this: every data point that informs a production decision, custody transfer, or regulatory report originates with a sensor, travels over a communication path, and is processed by software that someone designed, secured, and maintains. Understanding the architecture of that journey, and its failure modes, is as important as understanding the measurement itself.

## REFERENCES

1. MarketsandMarkets. (2025). SCADA market by component, communication technology, end-user industry and region - Global forecast to 2030. MarketsandMarkets Research.
2. Boyer, S. A. (2010). SCADA: Supervisory control and data acquisition (4th ed.). ISA.
3. International Society of Automation. (2010). ANSI/ISA-95: Enterprise-control system integration. ISA.
4. Zhu, Y. (2009). Multivariable process control: Efficient transfer of knowledge. Wiley.
5. Dragos. (2023). Year in review: ICS/OT cybersecurity 2023. Dragos, Inc.
6. Claroty. (2024). The global state of CPS security 2024: Business impact of disruptions. Claroty.
7. International Electrotechnical Commission. (2019). IEC 62443: Security for industrial automation and control systems. IEC.
8. Cirrus Link Solutions. (2022). MQTT Sparkplug specification v3.0. Eclipse Foundation.
9. Eclipse Foundation. (2023). MQTT version 5.0. Eclipse Foundation.
10. Neeve. (2026, April 7). The Purdue model is dead; long live the Purdue model: Building zero trust edge architectures for oil & gas [Conference presentation]. ENTELEC Conference & Expo, Galveston, TX.
11. Automation World. (2022, March). Edge computing's role in the evolution of industrial operations. Automation World.
12. LNS Research. (2023). Industrial transformation study: Operational technology modernization. LNS Research.
13. McKinsey & Company. (2022). Capturing value at scale in discrete manufacturing with Industry 4.0. McKinsey Global Institute.
14. ARC Advisory Group. (2024). Industrial IoT and edge computing: Market analysis. ARC Advisory Group.
15. Nozomi Networks. (2024). OT/IoT security report: Insights for critical infrastructure operators. Nozomi Networks.
16. Aberdeen Group. (2023). Predictive maintenance: The strategic approach to reliability. Aberdeen Group.
17. Deloitte Insights. (2023). The future of maintenance: AI-driven asset performance management. Deloitte.
18. Saudi Aramco. (2022). Integrated report 2022. Saudi Aramco.
19. Wood Mackenzie. (2024). Digital oilfield technology: Market outlook 2024-2030. Wood Mackenzie.
20. Gartner. (2024). Magic quadrant for industrial IoT platforms. Gartner.
21. CISA. (2023). Cross-sector cybersecurity performance goals. Cybersecurity and Infrastructure Security Agency.
22. SpaceX. (2024). Starlink for business: Enterprise connectivity solutions. SpaceX.
23. Yokogawa. (2022). OpreX asset performance management: Case study collection. Yokogawa Electric Corporation.
24. U.S. Department of Energy. (2010). Operations & maintenance best practices: A guide to achieving operational efficiency (Release 3.0). Federal Energy Management Program.
25. Deloitte Insights. (2017). Industry 4.0 and predictive technologies for asset maintenance. Deloitte.
26. C3 AI. (2022, March 8). Shell achieves major milestone: Scales artificial intelligence predictive maintenance to 10,000 pieces of equipment using C3 AI [Press release]. BusinessWire.
27. Frost & Sullivan. (2022). Edge computing drives industrial IT/OT convergence in oil & gas, utilities, manufacturing, and automotive verticals. Frost & Sullivan.