Redefining Industrial Networks Through Network Heuristics: A Conceptual Framework for EtherNet/IP Architectures

Robert Lodesky
Sales Director
HMS Networks
Presented at the ODVA
2026 Industry Conference & 24th Annual Meeting
October 21, 2026
Vilamoura, Portugal

Abstract

Industrial networks are commonly defined through topology, protocol conformance, and connectivity. While these descriptors remain essential, they are no longer sufficient to explain the behavior of contemporary automation systems operating across distributed control, diagnostics, engineering access, segmentation boundaries, and cross-domain integration. This paper argues that new networks require a renewed examination of current heuristics in networking, particularly when industrial architectures are treated only as static communication infrastructures rather than as dynamic systems of interaction.

The paper develops a conceptual framework that reinterprets eight networking heuristics through an industrial lens, focusing on interaction density, connection paths, communication-function fit, concentration points, and coordination load, specifically, through EtherNet/IP. Drawing on network analysis and communication theory, it distinguishes which heuristics can be translated productively into industrial-network thinking and which must be qualified before they can support technical reasoning. The central finding is that industrial-network behavior cannot be inferred from protocol and topology alone. Architectural reach, resilience, and scalability are also shaped by how traffic and control dependencies are distributed, where bridges connect segmented domains, and where hubs emerge as either useful coordinators or failure concentrators. This paper proposes a more precise analytical approach to the interpretation of industrial systems already well understood by practitioners. By redefining industrial networks as interaction systems rather than static infrastructures, the paper offers a stronger conceptual basis for architecture review, segmentation strategy, concentration analysis, interoperability planning, and resilience-oriented design.

Keywords
EtherNet/IP; Common Industrial Protocol (CIP); industrial network architecture; network heuristics; network science; segmentation; boundary governance; multicast; resilience; Monte Carlo simulation

Definition of Terms
Boundary quality: The effectiveness, transparency, observability, and recoverability of communication crossings between segments, zones, or functional domains.
Communication-function fit: The degree to which a communication pattern matches its operational purpose, including control, diagnostics, commissioning, configuration, or maintenance.
Concentration behavior: The extent to which communication or coordination depends on a limited number of nodes or paths.
Interaction density: The rate and intensity of exchanges among nodes and subsystems.

Vector richness: The number and diversity of meaningful communication paths available to a node or subsystem.

1. Introduction

Industrial networking has historically been described through physical topology, logical segmentation, protocol conformance, determinism, and interoperability. Those descriptors remain foundational. For engineers working with industrial communication systems such as EtherNet/IP, they are the primary language through which correctness, compatibility, and performance are judged. Yet the growing complexity of modern automation systems has exposed an analytical gap. Two industrial architectures may be equally compliant, similarly segmented, and comparable at the protocol layer, while still behaving differently in resilience, diagnosability, coordination load, and failure propagation.[2] [3]

This paper addresses that gap by proposing that industrial networks should be redefined as **interaction systems** rather than interpreted solely as static communication infrastructures. The claim is not that protocol rigor should be displaced by metaphor, nor that social-network concepts can be imported into industrial engineering without translation. The claim is narrower and more useful; recurring heuristics about activity, vectors, brokerage, hubs, and boundaries can help explain architectural behavior that topology diagrams alone often conceal, provided those heuristics are treated as conditional and analytically disciplined rather than as universal laws.

The paper makes four contributions. First, it reformulates eight proposed networking heuristics into a more rigorous conceptual framework suitable for industrial analysis. Second, it connects those heuristics to established findings from social network analysis, media-choice research, and coordination theory, while also incorporating literature that disputes overgeneralized claims. Third, it aligns the strongest heuristics with more formal network-science constructs such as preferential attachment, structural holes, modularity, and cascading effects.[9] Fourth, it reports a structured method and results from a previously built set of EtherNet/IP-oriented conceptual simulations, thereby giving the paper a stronger research-paper character rather than leaving it at the level of conceptual essay alone.[10]

2. Problem Statement and Motivation

The central problem is that conventional industrial-network definitions are strong at describing **what is connected** and **which protocols are present**, but weaker at explaining **how coordination load accumulates**, **where path dependency emerges**, **how boundaries shape operational behavior**, and **why formally correct systems can still become fragile or difficult to reason about**. Industrial-network engineering has long been excellent at identifying packet structure, timing requirements, deterministic behavior, and segmentation logic. It has been less explicit about the higher-order dynamics that arise when traffic, diagnostics, orchestration, and cross-domain dependencies become concentrated in particular nodes or interfaces.[2] [3]

This matters because industrial architectures increasingly include not only controllers and devices, but also gateways, edge collectors, engineering workstations, historians, diagnostics

services, asset-management tools, and enterprise-facing data paths. In such environments, network behavior is affected not only by whether devices can communicate, but also by whether those communication patterns create useful optionality, hidden coupling, manageable concentration, or brittle boundary crossings. An architecture may be protocol-correct and still be operationally opaque.[2] [3] [8]

Problem dimension	Conventional strength	Analytical gap addressed in this paper
Topology	Describes physical and logical arrangement accurately	Does not fully explain why similar topologies behave differently under load or failure
Protocol conformance	Establishes correctness and interoperability	Does not identify emergent chokepoints, bottlenecks, or coordination asymmetries
Segmentation	Defines isolation and zoning logic	Does not guarantee that boundary crossings are well governed or resilient
Traffic analysis	Measures volume and message classes	Does not by itself reveal whether communication is functionally appropriate
Architecture review	Evaluates design intent and standards fit	Often under-describes interaction structure, hidden bridges, and concentration risk

The motivation for this paper is therefore methodological. Engineers need a language that supplements protocol and topology with a more behaviorally sensitive framework, without sacrificing technical precision. The present work argues that a qualified heuristic framework can serve that role.

3. Literature and Conceptual Foundation

The conceptual basis of this paper comes from three adjacent bodies of literature: **social network analysis**, **organizational communication**, and **coordination / centralization research**. These bodies of work are not identical to industrial networking, but they offer disciplined constructs for thinking about repeated interaction, brokerage, concentration, and communication effectiveness.

Social network analysis provides the first foundation. Standard network concepts such as **degree**, **betweenness**, **closeness**, and **eigenvector centrality** describe different forms of importance within a network.[1] Tie strength is commonly associated with the frequency and intensity of interaction, while weak ties and bridging positions are often valuable because they connect otherwise separated regions of a network.[1] [5] These ideas are directly relevant to industrial-network reasoning because they help distinguish simple busyness from strategic structural position.

A second foundation comes from research on communication and proximity. Work associated with Tomas Allen showed that communication frequency declines sharply with distance and that the decay applies not only to face-to-face interaction but to the use of all

communication media.[3] Subsequent research using sociometric badges found that face-to-face cohesion is positively associated with productivity, especially for complex tasks, suggesting that richer forms of interaction can matter for coordination under ambiguity.[2] At the same time, this line of research does not support a fixed universal ranking of channels. El-Shinnawy and Markus explicitly criticized media richness theory as an inadequate explanation of media choice in their comparison of electronic mail and voice mail, indicating that task and context remain decisive.[7]

A third foundation comes from work on centralization and coordination. Hossain and Wu argue that network centrality correlates with coordination ability, which supports the intuition that hubs can aid orchestration.[4] Yet recent evidence complicates any universal endorsement of centralization. Ding, Shi, and Xiao report that centralized communication structures outperform decentralized ones under weak-tie conditions, while decentralized structures perform better under strong-tie conditions.[8] The implication is that central hubs can be helpful, but only under certain relational and task conditions.

Literature stream	Core contribution	Relevance to industrial-network interpretation
Social network analysis	Centrality, tie strength, weak ties, brokerage, reachability[1] [5]	Supports analysis of hubs, bridges, path significance, and boundary crossings
Communication / proximity research	Distance effects, face-to-face cohesion, media fit[2] [3] [7]	Supports analysis of interaction density and communication-function fit
Centralization / coordination research	Coordination gains and centralization contingencies[4] [8]	Supports qualified analysis of hubs and concentration behavior
Critiques of one-dimensional tie models	Separation of capacity, frequency, redundancy[6]	Prevents overreliance on simplistic scalar notions of connectivity

The literature therefore supports the use of heuristics, but only in qualified form. It suggests that activity matters, but activity is not equivalent to strategic reach. It suggests that richer interaction can matter, but not through a universal multiplier. It suggests that hubs can coordinate, but not without risk of overload, bottlenecks, or context-dependent underperformance.[1] [4] [6] [7] [8]

4. Translating Heuristics into Industrial-Network Terms

The original heuristic set was framed in the language of **nodes**, **vectors**, **hubs**, and **interaction quality**. To make those claims useful for industrial protocol specialists, they must be translated into engineering terms. That translation is the conceptual bridge of the paper.

Original heuristic language	Industrial-network translation
Activity per node	Interaction density, polling intensity, produced/consumed traffic, diagnostic request frequency
Vectors	Path diversity, bridge links, alternate communication routes, cross-zone dependencies
Quality of interaction	Suitability of communication pattern for control, diagnostics, commissioning, configuration, and maintenance
Forced interaction	Required exchanges created by architecture, scheduled polling, mandatory status traffic, induced dependencies
High-capacity hubs	Controllers, gateways, aggregation switches, orchestration layers, protocol translators, visibility concentrators
Growth or stagnation	Scalability versus bottlenecking, graceful degradation versus failure concentration
Dual hubs and borders	Multi-domain coordination centers, semi-autonomous cells, bounded interfaces, inter-zone bridges
Artificial hubs	Imposed concentration points created by ownership structure, tooling, or governance rather than organic traffic demand

Once translated this way, the heuristic set becomes much less metaphorical. It can be interpreted as a set of questions about architecture. Where does activity actually concentrate? Which paths are meaningful alternatives and which are latent liabilities? Which communication patterns are aligned with their function, and which are merely accumulated over time? Where do hubs improve orchestration, and where do they create silent fragility?

5. An Integrated Conceptual Framework

The paper's core argument can be organized into six analytical constructs: **structural topology**, **interaction density**, **path structure or vector richness**, **concentration behavior**, **communication-function fit**, and **boundary quality**. Together, these constructs provide an explanatory model that is more behaviorally expressive than topology alone.[2] [3]

Framework component	Working definition	Expected effect on industrial-network behavior
Structural topology	The physical and logical arrangement of devices, controllers, switches, gateways, and segments	Determines what routes are possible and where failure exposure exists
Interaction density	Rate and intensity of exchanges among nodes and subsystems	Reveals local dependence, operational rhythm, and load concentration

Framework component	Working definition	Expected effect on industrial-network behavior
Path structure / vector richness	Number and diversity of meaningful paths available to a node or subsystem	Improves reach and optionality when paths are purposeful; increases coupling when unmanaged
Concentration behavior	Extent to which communication or coordination depends on a small number of nodes or paths	Can improve orchestration efficiency until it produces bottlenecks or fragility
Communication-function fit	Degree to which the communication pattern matches operational purpose	Improves diagnosability and performance when control, alarms, diagnostics, and engineering access are well matched to traffic class
Boundary quality	Effectiveness and transparency of crossings between segments, zones, or functional domains	Determines whether segmentation behaves as resilience, isolation, or concealed chokepoint

The causal logic of the framework is straightforward. **Topology** defines the space of potential relations. **Interaction density** determines which relations become behaviorally important. **Vector richness** and **boundary quality** shape whether the system preserves alternatives or concentrates dependency. **Concentration behavior** reveals where orchestration, translation, or visibility accumulate. **Communication-function fit** governs whether the resulting interaction pattern serves the actual purpose of the architecture. The outcome is a network that is more or less resilient, scalable, diagnosable, and operationally legible.[2] [3]

This framework supports the paper's central proposition.

> ***Industrial-network behavior is shaped by the interaction between topology, traffic intensity, path dependency, concentration points, boundary quality, and communication-function fit; therefore, systems with similar topologies may exhibit materially different resilience, scalability, and coordination behavior when their interaction patterns differ.***

A second proposition qualifies the role of heuristics.

> ***Network heuristics are most useful when treated as conditional descriptors of system behavior rather than as universal laws.***

6. What Holds and What Does Not

The value of the heuristic approach depends on separating what is well supported from what is overstated. This section evaluates the major claims one by one.

6.1 Interaction Density and Reach

The idea that more activity per node reduces distance is directionally plausible, but incomplete. Repeated interaction is a standard indicator of tie strength, and higher frequency often creates stronger local dependence.[1] [3] Yet reachability is not reducible to frequency. A locally busy controller can remain structurally peripheral if it interacts densely only within one bounded cell. Conversely, a gateway with lower total message volume may be strategically central because it bridges functional or zonal boundaries.

The more rigorous formulation is therefore that **higher interaction density often increases local coupling and visibility, but effective reach depends on structural position, boundary crossings, and the significance of the paths involved**.[1] [5] [6]

6.2 Vectors, Bridges, and Useful Connectivity

The claim that more vectors are always better is also only partly defensible. Weak-tie and brokerage research shows that some ties matter disproportionately because they connect otherwise separate clusters.[5] Structural holes theory explains why bridges can create informational and coordinating advantage.[9] Yet the literature also warns against reducing network quality to tie count. Capacity, redundancy, and frequency are distinct dimensions, and more links can generate unmanaged coupling rather than resilience.[6]

The academic reformulation is that **vector richness improves performance when added paths create meaningful alternatives, not when they merely increase architectural ambiguity**. In industrial terms, a carefully added diagnostic route may improve observability, while a proliferation of ad hoc cross-links may create troubleshooting difficulty and hidden failure propagation.

6.3 Communication Quality and Function Fit

The claim that richer interaction is categorically more precise requires the strongest qualification. Face-to-face cohesion has been shown to support productivity in complex settings, and richer interaction can help with tacit coordination or ambiguity.[2] However, media richness theory does not support a universal, stable ladder of communication quality across all tasks.[7]

For industrial networks, the better formulation is not that one communication form is intrinsically superior in all cases, but that **communication patterns should be judged by functional fit**. Cyclic control traffic, scheduled diagnostics, alarms, event-driven visibility, and engineering access have different requirements. More communication does not necessarily mean better communication, and a richer channel is not automatically a better one when the task is routine, deterministic, or documentation-oriented.[2] [7]

6.4 Forced Interaction Versus Congruent Interaction

The earlier claim that forced interaction matters as much as congruent interaction is also weakly supported. Mandated exchanges can create exposure, consistency, and architectural

discipline. Yet they do not necessarily create trust, durability, or adaptive value. In industrial terms, required polling or mandatory status exchange can stabilize coordination, but it can also create burdensome traffic or artificial dependency if not aligned with function.

This heuristic is best preserved in softened form: **induced interaction can create structural dependency or regularity, but not necessarily durable relational or operational value**.[6] [7]

6.5 Hubs, Concentration, and Cascading Effects

The strongest part of the heuristic set concerns hubs. Centrality theory, preferential attachment, and betweenness all support the idea that some nodes become disproportionately important to flow and coordination.[1] [4] [9] Yet this same literature clarifies the danger: the most useful coordinators are often the most dangerous bottlenecks. Centralized structures can outperform under weak-tie conditions, while decentralized ones can outperform under strong-tie conditions.[8]

The defensible conclusion is that **hubs are analytically important because they can both improve orchestration and concentrate fragility**. Hub effects are real, but conditional. They should be understood as amplification mechanisms rather than presumed best practice.[4] [8]

7. Alignment With Formal Network-Science Heuristics

One reason the heuristic framework becomes more academically credible is that several of its strongest claims align with established graph-theoretic heuristics. A later ten-heuristic taxonomy of social-network concepts transferable to IT, logistics, and industrial networks shows that the present framework's strongest elements map onto **triadic closure**, **preferential attachment**, **betweenness centrality vulnerability**, **community detection via modularity**, **structural holes and brokerage**, **information cascades**, **k-core decomposition**, **interdependent cascades**, and **degree assortativity**.[9]

Formal heuristic	Best matching claim in this paper	Analytical gain
Triadic closure	Activity and local vector formation	Explains clustering and local thickening of interaction
Preferential attachment	Convergence toward hubs	Provides a mechanism for hub emergence
Betweenness centrality vulnerability	Hub concentration and chokepoint risk	Formalizes failure concentration and bridge criticality
Community detection / modularity	Borders, clusters, and bounded domains	Explains structural rather than accidental segmentation

Formal heuristic	Best matching claim in this paper	Analytical gain
Structural holes and brokerage	Bridges and vector richness	Clarifies why some paths carry disproportionate strategic value
Information cascades and thresholds	Growth, stagnation, and nonlinear propagation	Adds dynamic logic to hub and dependency effects
K-core decomposition	Resilience through dense support structures	Identifies stable architectural cores
Interdependent cascading failures	Layered industrial and digital dependence	Extends analysis beyond a single network layer
Degree assortativity	Hub-to-hub structure and dual-center architecture	Refines analysis of core interconnection patterns

This alignment matters because it shows that the framework is not merely rhetorical. Its strongest claims are legible within recognized network-science language. At the same time, the alignment also reveals the framework's limits. The graph-theoretic taxonomy does not strongly support claims about media richness or voluntary versus forced interaction, reinforcing the conclusion that those heuristics must remain more provisional.[7] [9]

8. Organizational Structure as a Precursor to Network Structure

The framework becomes stronger when extended sociotechnically through **Conway's Law**. Conway argued that organizations design systems that mirror their own communication structures.[10] Applied to industrial networking, this suggests that architecture is shaped not only by technical requirements, but also by how engineering, controls, operations, diagnostics, and IT/OT teams coordinate their responsibilities.

> *"Any organization that designs a system (defined broadly) will produce a design whose structure is a copy of the organization's communication structure." — Melvin Conway*[10]

In practical terms, recurring chokepoint gateways, over-centralized visibility layers, isolated diagnostics paths, or weakly governed inter-zone interfaces may reflect organizational ownership boundaries as much as protocol necessity. This observation does not replace the technical framework; it adds an upstream explanatory layer. Some concentration patterns and segmentation behaviors may be organizational artifacts before they are engineering artifacts.[10]

Framework construct	Conway-informed reading	Industrial implication
Interaction density	Sparse collaboration during design may yield sparse or brittle integration paths	Review team interaction patterns alongside traffic patterns

Framework construct	Conway-informed reading	Industrial implication
Path structure	Reporting lines and ownership can shape which systems connect directly and which require mediation	Some indirect paths may be organizationally rather than technically induced
Concentration behavior	Centralized decision-making often produces centralized coordination nodes	Technical hub risk may mirror managerial hub risk
Boundary quality	Organizational silos can become network boundaries	Some segmentation protects; some merely encodes separation
Communication-function fit	Ownership-driven design can privilege responsibility boundaries over operational fit	Traffic design should be evaluated against plant function, not only governance

This yields an additional proposition.

Proposition	Statement
P6	Industrial-network architectures will tend to reproduce the communication structures, ownership boundaries, and coordination asymmetries of the organizations that design and maintain them; therefore, some chokepoints and segmentation patterns should be interpreted as organizational artifacts as well as engineering decisions.

9. Experimental Method

To supplement the conceptual framework with a more explicit evaluation layer, the following EtherNet/IP experiments were implemented as a **Monte Carlo comparative model** in spreadsheet form and then reproduced programmatically for reporting consistency. The objective of the experiment set was not to validate EtherNet/IP protocol mechanics directly, but to test whether the framework's propositions generate stable comparative differences when competing architectural concepts are scored under uncertainty.

Each experiment compared two or three architecture variants representing distinct design choices. Every variant was defined by **five weighted factors**, with weights summing to 1.0. For each factor, a **mean** and **standard deviation** were assigned on a bounded 1-5 conceptual performance scale, where higher values represented stronger architectural quality under that factor. The simulation then drew random values for each factor from a normal distribution using the assigned mean and standard deviation, clipped the results to the 1-5 scale, and computed a weighted composite score for each architecture variant in each trial.

The workbook design used **1,000 trials per experiment** and reported six summary statistics for every variant: **mean**, **10th percentile**, **median**, **90th percentile**, **standard deviation**, and **win probability**. Win probability was defined as the share of trials in which a variant achieved the highest composite score in its experiment. To provide a stable reported version

for this paper, the experiment set was re-run with a fixed random seed of **42**, preserving the original assumptions while making the outputs reproducible.

Method element	Implementation in the reported experiments
Experiment type	Monte Carlo comparative scoring model
Domain framing	EtherNet/IP-oriented conceptual architecture scenarios
Trials per experiment	1,000
Scoring scale	1-5 bounded conceptual quality scale
Distribution form	Normal draws with truncation to the 1-5 bounds
Composite score	Weighted average of five factors per variant
Reproducibility approach	Fixed-seed rerun of the previously built workbook assumptions
Primary outputs	Mean, P10, median, P90, standard deviation, win probability

The experiments therefore function as **conceptual structured comparisons** rather than empirical plant studies. Their methodological value lies in making the framework operational, transparent, and falsifiable at the level of architecture alternatives. Their limitation is equally important: the inputs are theory-informed priors, not plant measurements.

10. Experimental Results

The results were directionally strong across all five experiments. In every case, the winning architecture variant aligned with the framework's more qualified view of interaction, vectors, concentration, and boundaries. The simulations consistently favored designs that improved **cross-boundary reach**, **disciplined path diversity**, **function-aligned communication**, **bounded concentration**, and **governed segmentation** over alternatives defined by local busyness, indiscriminate connectivity growth, generalized status traffic, centralized chokepoints, or opaque bridges.

10.1 Cross-Experiment Summary

Experiments 1–5 examine one recurring factory-floor architecture rather than five unrelated networks: local EtherNet/IP Class 1 cyclic I/O remains inside production cells, while supervisory coordination, engineering, diagnostics, lifecycle access, and data collection cross governed cell boundaries. Each experiment holds the underlying control role and required operational functions constant, then varies one interaction property—reach, path structure, communication-function fit, concentration, or boundary quality.

Shared factory-floor reference architecture. The multi-cell scenarios use four cells, each with one controller and sixteen adapters (64 adapters total). Local controllers retain machine-level control. A supervisory controller, engineering workstation, and historian/edge service require bounded cross-cell access. The fixed 69–71-device scenarios and nominal 100 Mb/s switched-Ethernet assumption are transparent reference conditions, not claims that every plant has this scale or topology.

Why this is a credible aggregate. It represents modular plants in which machines must control locally while line- or plant-level services must remain visible, diagnosable, maintainable, and recoverable. It is structurally representative—not a statistical average—because the same tensions occur at larger scale: bounded cell control, cross-zone access, change management, shared-service dependence, and recovery after ordinary failures.

Experiment	Architectural comparison	Value to architecture review
1	High local I/O activity versus lower-volume supervisory reach.	Separates local traffic intensity from system-wide coordination leverage; a busy controller is not necessarily the most consequential dependency.
2	Minimal segmentation, purposeful alternate paths, and cross-coupled access.	Distinguishes recoverable path diversity from coupling debt. Resilience is added by bounded, documented alternatives—not by open-ended relationships.
3	One frequent pattern or continuous broad status versus function-aligned traffic.	Tests whether control, routine status, detailed diagnostics, and lifecycle activity each use communication appropriate to their purpose.
4	Distributed services, one centralized hub, and bounded role-specific hubs.	Makes concentration inspectable: coordination can simplify normal operation, but a shared intermediary expands the blast radius of failure or change.
5	Disciplined, observable boundaries versus opaque gateway dependency.	Shows why zone diagrams alone are insufficient: permitted crossings, local evidence, recovery behavior, and chokepoint control determine whether segmentation protects.

Aggregate architectural conclusion. The five experiments replace incomplete proxies for quality—busyness, connection count, information volume, centralization, and segmentation alone—with a combined assessment of purpose, reach, path governance, concentration, and boundary behavior. The resulting design principle is to keep time-sensitive control close to the machine and make cross-cell access deliberate, role-bounded, observable, and recoverable.

Interpretation boundary. *The reported Monte Carlo results are theory-informed conceptual architecture comparisons. They do not measure actual EtherNet/IP throughput, latency, jitter, packet loss, controller load, DLR recovery, security enforcement, or vendor-specific behavior.*

Experiment 1: High Activity Does Not Equal High Reach

Experiment 1 compares two roles within a fixed four-cell EtherNet/IP reference architecture containing one supervisory controller, four cell controllers, and sixty-four adapter devices. In

the busy-local condition, one cell controller exchanges 1 ms cyclic I/O with sixteen local adapters using nominal 96-byte produced input and 32-byte consumed output payloads. This yields an application-payload estimate of 16.384 Mb/s, excluding Ethernet, IP, transport, CIP, encapsulation, and retransmission overhead. In the broader-reach condition, a supervisory controller exchanges a 128-byte received summary and a 64-byte sent command with each of four cell controllers every 10 ms, yielding a nominal application-payload estimate of 0.6144 Mb/s. The supervisory controller has four direct peers but architectural visibility across all four cells and their sixty-four adapter devices. The comparison is intended to isolate local traffic intensity from cross-cell coordination reach; it is not a conformance, timing, or packet-level performance test.

The broader-reach design achieved a mean score of **3.762**, versus **2.188** for the locally busy alternative, and won **100.0%** of trials. This result supports the proposition that interaction density and strategic reach should not be conflated. In industrial terms, high cyclic traffic inside one cell may indicate local importance, but it does not imply system-wide reach or architectural leverage.

Experiment 2: More Vectors Can Improve Resilience or Create Coupling Debt

Experiment 2: More Vectors Can Improve Resilience or Create Coupling Debt. Experiment 2 evaluates path structure within a fixed 69-end-device EtherNet/IP reference architecture comprising one supervisory controller, four cell controllers, and sixty-four adapter devices. Each cell controller maintains the same local Class 1 cyclic I/O relationships with sixteen local adapters in every alternative; the comparison therefore concerns only the structure and governance of cross-cell relationships and infrastructure paths.The minimal-segmented alternative permits four explicit-message supervisory relationships, one to each cell controller, across a single designated path per cell. The purposeful-path-diversity alternative preserves that same bounded set of approved CIP relationships but provides one documented alternate infrastructure path for each cell's supervisory boundary crossing. A DLR-capable local cell segment may be used where appropriate to support controlled physical-path resilience; the experiment does not assume standardized EtherNet/IP multipath CIP routing. The cross-coupled alternative retains the same endpoint population but permits 64 directed cross-cell logical relationships: four supervisory-to-controller relationships, twelve controller-to-controller relationships, and forty-eight controller-to-remote-adapter relationships. These additional accesses represent routine explicit-message diagnostic, commissioning, configuration, or coordination exchanges that expand across cell boundaries without an equivalent role-boundary or recovery definition.The original conceptual Monte Carlo model reports mean scores of 3.451 for minimal segmentation, 4.034 for purposeful path diversity, and 2.322 for the cross-coupled architecture. Purposeful path diversity achieved the highest score in 99.0% of the original model's trials, while the cross-coupled alternative did not lead in any trial. The result should be interpreted narrowly: additional paths add value when they improve recoverable, observable, and purpose-bounded alternatives. It does not establish that any particular EtherNet/IP device, topology, DLR implementation, or message rate will produce these scores in a deployed plant.

Experiment 3: Communication-Function Fit Outperforms Generic Richness

Experiment 3 compares three communication-pattern designs in a fixed EtherNet/IP packaging-cell reference system comprising one cell controller, sixteen adapter devices, one HMI/engineering workstation, and one historian or edge collector. All conditions preserve the

same local Class 1 cyclic control objective, diagnostic objective, lifecycle-access objective, visibility objective, endpoint population, topology, and nominal 100 Mb/s capacity assumption. The independent variable is the degree of communication-function fit. The high-frequency uniform-access condition uses one 10 ms explicit-message polling pattern for broad non-control information. The function-aligned condition retains local Class 1 I/O for time-sensitive control, uses a planned low-rate routine status set, initiates detailed diagnostics only when needed, and confines commissioning/configuration access to authorized lifecycle activity. The continuous-general-status condition distributes a broad status payload continuously to multiple consumers. The model evaluates control fit, diagnostics fit, load efficiency, signal clarity, and lifecycle usability with the stated weights, using 1,000 bounded-normal Monte Carlo trials with fixed seed 42.

Results: Under the stated conceptual assumptions, function-aligned patterning achieved a mean composite score of 4.188 and led in 100.0% of the 1,000 trials. High-frequency uniform access achieved a mean of 2.734, and continuous general status achieved a mean of 2.283; neither led in any trial. The result indicates that the decisive design variable is not raw communication volume, but whether the selected EtherNet/IP communication pattern fits the role of the information being exchanged. The function-aligned condition was assessed as stronger because it preserved local cyclic control for time-sensitive functions, limited routine visibility to a bounded status set, triggered detailed diagnostics when operationally needed, and kept lifecycle configuration activity distinct from normal production traffic. These results are conceptual architecture comparisons and do not constitute measured claims about packet-level performance, a specific RPI, or a universal EtherNet/IP traffic policy.

Experiment 4: Hubs Enable Coordination Until They Become Failure Concentrators

Experiment 4 compares three allocations of the same system-level EtherNet/IP functions within a fixed reference system of one supervisory controller, one engineering workstation, one historian or edge collector, four cell controllers, and sixty-four adapter devices. In every condition, each cell controller retains its local Class 1 cyclic I/O relationships with sixteen local adapters; supervisory coordination, engineering/maintenance access, and historian/edge collection are represented as twelve required cell-level system functions. The distributed-control condition assigns these functions directly to the three respective services. The centralized-hub condition makes one intermediary responsible for all twelve functions. The bounded-dual-hub condition assigns supervisory coordination to one intermediary and engineering plus historian/edge functions to a second intermediary, with a limited documented interface and hub-specific recovery behavior. The endpoint population, local control role, function set, nominal capacity assumption, and test events are held constant. The independent variable is the degree and form of hub dependency. The model scores coordination efficiency, failure containment, maintainability, bottleneck resistance, and fallback readiness using the stated weights and 1,000 bounded-normal Monte Carlo trials with fixed seed 42.

Results: Under the stated conceptual assumptions, bounded dual hubs achieved a mean composite score of 3.914 and led in 74.0% of the 1,000 trials. Distributed control achieved a mean of 3.741 and led in 26.0% of trials, while the centralized hub achieved a mean of 2.569 and did not lead in any trial. The result indicates that a system-level hub may increase coordination efficiency while simultaneously reducing failure containment, bottleneck resistance, and fallback readiness when all functions depend on the same intermediary. The bounded dual-hub condition improved the balance of these factors by dividing the twelve

required system-level functions into defined role-specific dependency sets. The result should be interpreted as a conceptual architecture comparison rather than a measured claim about a specific EtherNet/IP device, topology, or redundancy implementation.

Experiment 5: Boundary Quality Determines Whether Segmentation Protects or Constrains the System

Experiment 5 compares two multi-zone EtherNet/IP reference architectures containing one supervisory controller, one engineering workstation, four cell controllers, and sixty-four adapter devices. Both conditions preserve the same four-cell segmentation, the same local Class 1 cyclic I/O relationships between each cell controller and its sixteen local adapters, the same authorized supervisory functions, the same nominal network-capacity assumption, and the same fault challenge: loss of one cell's designated boundary-forwarding function or primary upstream path. The independent variable is boundary quality, treated as a compound architecture condition. The disciplined-bridge condition bounds cross-zone functions to defined explicit-message relationships, exposes local forwarding state and counters, documents recovery behavior, and avoids undisclosed concentration. The opaque-boundary condition retains the same segmentation diagram but relies on incompletely documented gateway behavior, accumulated cross-zone access, limited local diagnostic disclosure, and hidden concentration. The model evaluates isolation quality, bridge transparency, interoperability control, recovery behavior, and chokepoint avoidance with equal weights, using 1,000 bounded-normal Monte Carlo trials and fixed seed 42.

Under the stated conceptual assumptions, disciplined bridge governance produced a mean composite score of 4.038, compared with 2.147 for opaque boundary dependency. The disciplined condition led in 100.0% of the 1,000 trials; its 10th–90th percentile interval was 3.798–4.281, compared with 1.853–2.448 for the opaque condition. The 1.891-point difference does not demonstrate a measured protocol performance gain. Rather, it indicates that, when two architectures have equivalent endpoint counts, top-level segmentation, local cyclic control, and intended cross-zone functions, a boundary designed for role-bounded access, local observability, recovery, and controlled concentration is assessed as materially more resilient and governable than an opaque gateway dependency. The result supports Proposition 5: segmentation becomes protective only when the behavior of the crossings is deliberately engineered.

11. Discussion

The experiments change the force of the paper in a useful way. Without them, the framework remains a persuasive conceptual reinterpretation. With them, it becomes a more structured analytical program. The reported results show that the framework can discriminate systematically among competing architecture concepts and that its discriminations are not trivial. The winning variants are not those with the highest raw activity, the largest number of connections, or the strongest centralization. They are the variants that embody **discipline in reach, discipline in connectivity, discipline in communication fit, discipline in concentration, and discipline in boundary governance**.

This is precisely why the framework is useful for industrial protocol specialists. It does not ask them to abandon protocol correctness, traffic engineering, or topology review. It asks them to extend those practices with a better language for interpreting architectural behavior. In practical terms, the results support five engineering judgments.

Engineering judgment

Local traffic intensity is not equivalent to architectural significance

Added paths create value only when they are purposeful and bounded

Communication design should follow function rather than volume maximization

Concentration should be managed, not maximized

Segmentation quality depends on bridge governance, not merely on zone labels

The results also sharpen the paper's theoretical asymmetry. The strongest empirical-style support is concentrated around **vector structure, concentration behavior, and boundary quality**. These are the domains where the conceptual framework, the literature, and the simulation outputs all point in the same direction.[4] [8] [9] By contrast, the paper's claims about media richness and forced interaction remain more interpretive and less strongly validated by the structural model. That is appropriate, because the Monte Carlo experiments were designed chiefly around industrial architecture behavior rather than around interpersonal media choice.

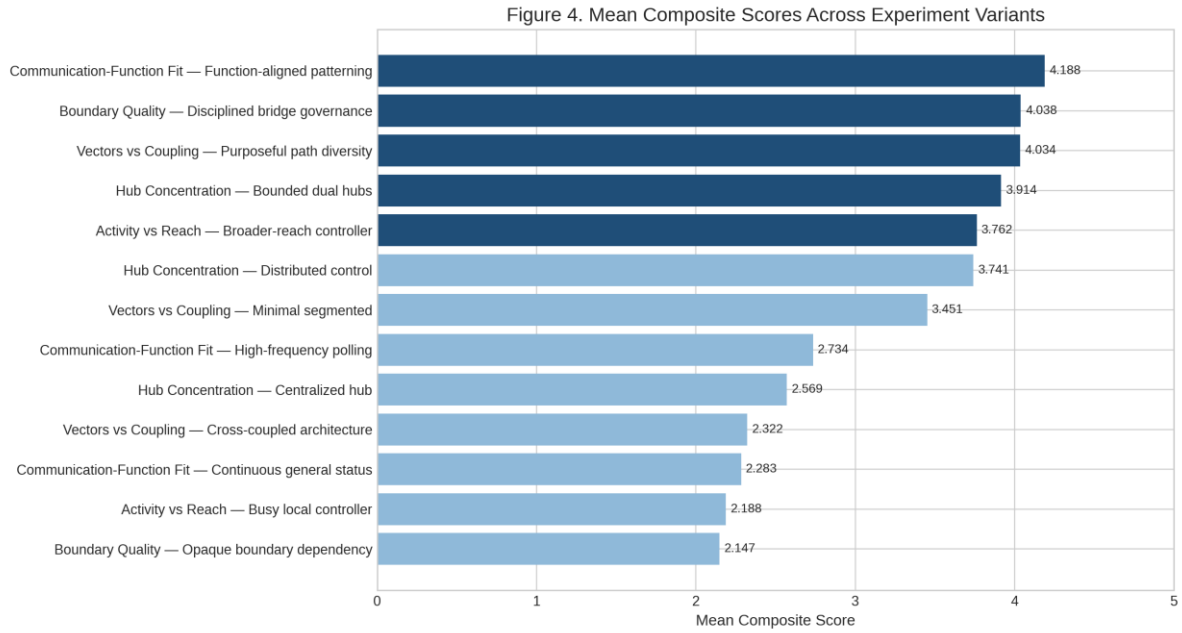


Figure 4 provides an aggregate rank ordering of all architecture variants, making the broader design logic of the framework visually explicit.

12. Implications for EtherNet/IP and Industrial Architecture

The practical payoff of the framework lies in how it can improve engineering judgment without displacing protocol rigor. EtherNet/IP architectures are already evaluated for conformance, interoperability, timing behavior, and segmentation. The present framework adds another layer of questions: Where does coordination actually concentrate? Which bridges are strategically critical? Which traffic patterns are functionally aligned, and which are inherited artifacts? Which boundaries are doing useful work, and which are merely concealing fragility?

Implication area	Engineering implication
Architecture review	Evaluate not only topology, but also where communication and coordination concentrate
Segmentation strategy	Judge boundaries by bridge quality, observability, and recovery behavior, not only by isolation logic
Gateway and controller placement	Treat concentration points as managed design choices requiring capacity and graceful degradation
Interoperability analysis	Study cross-domain communication as dependency structure, not only as a standards interface
Resilience engineering	Look for chokepoints, failure concentration, and hidden path dependency
Standards thinking	Extend protocol discussion with language that captures interaction behavior and coordination structure

For principal engineers, the value of this research is not that it alters the EtherNet/IP protocol stack. Its value is that it improves the interpretive layer used to reason about architectures built with it. The framework gives a disciplined vocabulary for identifying why a network that appears correct on paper may still be brittle in practice, difficult to troubleshoot, or overly dependent on a small number of nodes and paths. It also provides a better way to discuss architecture tradeoffs without collapsing them into simple pro-centralization or anti-centralization positions.

13. Specification Expansions for EtherNet/IP

The conceptual framework and the first set of EtherNet/IP-oriented experiments establish that interaction discipline, path governance, concentration visibility, and boundary quality materially shape industrial-network behavior. A natural next question is whether those validated claims can be translated into specific expansions of the EtherNet/IP adaptation of CIP. The answer proposed here is yes, provided the expansions are framed not as corrections to core transport semantics, but as extensions that make the specification more explicit about traffic-function discipline, multicast governance, boundary observability, concentration health, and segmented deployment behavior.

This section therefore presents five **specification expansions** grounded in the strongest claims of the paper. The purpose is not to claim measured field performance, but to provide a repeatable and inspectable analytical bridge between the paper's conceptual findings and plausible future specification work.

13.1 Expansion 1: Traffic Function Profile Expansion

Rationale

EtherNet/IP already provides defined traffic treatment categories through the QoS Object, including DSCP values for Class 0/1 priority categories, CIP explicit messaging, and IEEE 1588 traffic.2 The current categories describe traffic transport and priority, but do not provide a standardized way for an application profile to state the intended operational function of a traffic pattern—for example, local cyclic control, supervisory monitoring, commissioning, diagnostics, or time synchronization.

The proposed profile therefore does not attempt to inspect CIP application semantics or assign a new DSCP value to every use case. That would be ambiguous: a CIP explicit message can be used for diagnostics, commissioning, configuration, or other functions. Instead, the profile provides a declarative, profile-level way to document which existing EtherNet/IP traffic categories are intended for defined application functions. This supports architecture review and commissioning repeatability without changing packet formats or mandatory QoS behaviors.1 2

Recommended Location

Add an informative policy subsection after the Chapter 3 QoS requirements and a normative optional-profile subsection after Volume 2, Chapter 5, Section 5-7: QoS Object.

Proposed location	Proposed title	Function
3-7.x	Traffic Function Policy Guidance	Defines the non-normative relationship between operational function and existing traffic categories.
5-7.x	Traffic Function Profile Declaration	Defines optional read-only declaration and conformance behavior.

Proposed Specification Text

3-7.x Traffic Function Policy Guidance

EtherNet/IP traffic treatment is determined by the transport and priority mechanisms defined in this specification. An application profile may additionally define a traffic function for configuration, documentation, and architecture-review purposes.

A traffic function shall not alter CIP packet formats, connection semantics, encapsulation semantics, DSCP values, 802.1Q priority values, or QoS Object behavior unless a separate requirement in this specification explicitly requires that alteration.

A traffic function declaration shall not be interpreted as an assertion that the device can infer the operational purpose of a received CIP message.

The following traffic-function identifiers may be used by an EtherNet/IP application profile:

Identifier	Traffic function	Permitted EtherNet/IP traffic categories
0	Unspecified	Any category supported by the device

1	Local cyclic control	CIP Transport Class 0 or Class 1 connected data
2	Supervisory monitoring or coordination	CIP explicit messaging, or Class 0/1 connected data where defined by the application profile
3	Diagnostic or maintenance access	CIP explicit messaging
4	Commissioning or configuration access	CIP explicit messaging
5	Time synchronization	IEEE 1588 traffic as supported by the device

An application profile that uses a traffic-function identifier shall state the permitted EtherNet/IP traffic categories for that identifier. The profile shall not require a device to identify, inspect, prioritize, suppress, or reject traffic solely from a traffic-function identifier unless that behavior is separately specified.

5-7.x Traffic Function Profile Declaration

A device may implement the optional Traffic Function Profile Declaration.

A device that implements the Traffic Function Profile Declaration shall make available, through a read-only profile-defined attribute or application-profile declaration, the following information:

1. The Traffic Function Profile revision supported by the device.
2. The traffic-function identifiers supported by the device.
3. For each supported traffic-function identifier, the EtherNet/IP traffic categories that the device is capable of using for that function.
4. Whether the declaration is fixed by the device implementation or configurable through vendor-defined offline configuration.

A device shall not change the DSCP Value Attributes or 802.1Q Tag Enable attribute of the QoS Object solely because it implements this profile.

A device that does not implement this profile shall retain the behavior specified elsewhere in this specification.

Traffic Function Profile Conformance Statement

A device claiming conformance to the Traffic Function Profile Declaration shall identify the supported profile revision and each traffic-function identifier that it implements.

Scope Boundary

This expansion declares intended use of existing traffic categories. It does not define a universal rule that diagnostics are low priority, that control is always urgent, or that all explicit messages are alike. It leaves DSCP and 802.1Q treatment under the existing QoS mechanisms and the application profile's defined requirements.2

13.2 Expansion 2: Multicast Governance Expansion

Rationale

Volume 2 already specifies multicast connection handling, recommends unique multicast addresses per active connection, provides multicast-address allocation mechanisms, and

defines a multicast TTL configuration whose default value is one. It also cautions users about increasing multicast TTL beyond one.2

The gap is not the absence of multicast mechanics. The gap is the absence of a concise deployment-level declaration that tells an engineer whether an implementation can expose and govern the relevant multicast controls in a routed or segmented deployment. The proposed expansion does not change multicast allocation or make unique address allocation mandatory; it makes the existing configuration state and its deployment implications more explicit.1 2

Recommended Location

Add a new optional conformance profile after Chapter 3, Section 3-3.9: multicast connection rules, with supporting read-only disclosure associated with the TCP/IP Interface Object, Class Code 0xF5.

Proposed location	Proposed title	Function
3-3.9.x	Multicast Governance Profile	Defines the deployment and configuration disclosure requirements.
5-4.x	Multicast Governance Capability Declaration	Defines optional TCP/IP Interface Object disclosure attributes.

Proposed Specification Text

3-3.9.x Multicast Governance Profile

This section defines an optional Multicast Governance Profile for an EtherNet/IP device that supports IP multicast for CIP Transport Class 0 or Class 1 connections.

A device that implements the Multicast Governance Profile shall disclose, for each applicable TCP/IP Interface Object instance, the multicast configuration controls supported by that interface.

The profile shall not change the multicast address-selection rules, multicast connection semantics, or consumer packet-screening requirements defined elsewhere in this specification.

A device that implements the profile shall support read access to the TTL Value and Mcast Config attributes for each applicable TCP/IP Interface Object instance.

If the device supports setting either the TTL Value or Mcast Config attribute, it shall document whether an existing multicast connection is affected immediately, affected after reconnection, or requires device restart or other configuration application behavior.

A device shall not automatically increase the multicast TTL Value in response to establishing a connection, receiving a message, detecting a router, or detecting a multicast listener.

When the configured multicast TTL Value is greater than one, the device shall make a read-only indication available through the profile-defined disclosure mechanism. The indication shall not imply that cross-subnet multicast is appropriate, authorized, or supported by the surrounding infrastructure.

5-4.x Multicast Governance Capability Declaration

A TCP/IP Interface Object instance that implements the Multicast Governance Profile shall make the following read-only information available:

Attribute	Required behavior
Multicast Governance Profile Revision	Identifies the profile revision supported by the interface.
TTL Configuration Support	Indicates whether the TTL Value attribute is supported and whether it is settable.
Multicast Allocation Support	Indicates whether the Mcast Config attribute is supported and whether the implementation supports default allocation, configured allocation, or both.
Configured Multicast Scope Indication	Indicates whether the configured TTL Value is one, greater than one, or not supported. This indication shall report configuration state only.
Configuration Application Behavior	Indicates whether a change to multicast configuration applies immediately, applies to subsequently established connections, requires reconnection, or requires restart.

The profile shall not require discovery of multicast routers, IGMP queriers, multicast listeners, remote subnets, switch configuration, or end-to-end multicast delivery state.

Multicast Governance Profile Conformance Statement

A device claiming conformance to the Multicast Governance Profile shall identify the supported profile revision and the TCP/IP Interface Object instances to which the profile applies.

Scope Boundary

This expansion exposes configuration capability and state, not end-to-end multicast correctness. It does not mandate IGMP snooping, multicast routing, unique multicast addresses, or a specific switch configuration. Those are deployment and infrastructure decisions outside a single EtherNet/IP endpoint's complete visibility.2

13.3 Expansion 3: Bridge and Boundary Diagnostics Expansion

Rationale

A segmented architecture may remain protocol-conformant while a gateway, router, or forwarding device becomes an unobserved dependency. The paper calls this concern boundary quality: cross-boundary communication should be transparent enough to review, diagnose, and govern.1

Chapter 10, Bridging & Routing, is the correct primary location because Volume 2 explicitly reserves it for EtherNet/IP-specific additions to CIP bridging and routing and currently contains no such additions.2

Recommended Location

Add Sections 10-2 through 10-2.7 after the Chapter 10 introduction. Add an optional informative SNMP mapping in Appendix G for devices that also implement an SNMP Agent.

Proposed Specification Text

10-2 Bridge and Boundary Diagnostics Profile

10-2.1 Scope

This section defines an optional EtherNet/IP Bridge and Boundary Diagnostics Profile for a device that implements EtherNet/IP-specific CIP bridging or routing between two or more local network interfaces.

For the purposes of this section, a bridge instance is a local forwarding function within the device that forwards CIP communication between a defined set of local network interfaces. A bridge instance is not required to represent a security zone, a firewall, a complete route, or a complete network path.

A device that does not implement EtherNet/IP-specific CIP bridging or routing shall not be required to implement this profile.

10-2.2 Profile Disclosure

A device that implements this profile shall make the following information available through an explicit-message diagnostic service or a profile-defined read-only object:

1. The number of bridge instances implemented by the device.
2. For each bridge instance, a locally unique bridge-instance identifier.
3. The local network-interface identifiers associated with that bridge instance.
4. The supported encapsulation transport mode or modes used by the bridge instance, where applicable.
5. The diagnostic attributes defined in Section 10-2.3.

A device shall indicate an unsupported attribute using the standard CIP service or attribute-access behavior applicable to the object or service through which the profile is implemented.

10-2.3 Diagnostic Attributes

For each bridge instance, the device shall provide the following read-only diagnostic attributes:

Attribute	Required behavior
Forwarded CIP communication count	A monotonically increasing count of CIP communications forwarded by the bridge instance since the last counter discontinuity. The implementation shall state whether the unit is a message, packet, or connection event.
Forwarding discard count	A monotonically increasing count of locally detected forwarding discards attributable to the bridge instance since the last counter discontinuity.
Forwarding error count	A monotonically increasing count of locally detected forwarding errors attributable to the bridge instance since the last counter discontinuity. A remote-path or external-switch condition shall not be reported as a local forwarding error unless the device can determine that condition locally.
Counter discontinuity indicator	An indicator or timestamp that changes when one or more profile counters are reset, restarted, or otherwise lose continuity.
Administrative state	A read-only state indicating whether the local forwarding function is administratively enabled, administratively disabled, or not determinable by the device.

The profile shall not require a device to report end-to-end latency, upstream switch utilization, remote topology, remote peer identity, security-zone identity, route completeness, or an inferred health score.

10-2.4 Access and Performance Constraints

Profile attributes shall be read-only.

The profile shall not require periodic implicit messaging, cyclic production, unsolicited notification, or modification of existing EtherNet/IP real-time transport behavior.

A device may limit the rate at which diagnostic attributes are updated, provided that the update behavior is documented and the device does not represent the attributes as instantaneous measurements when they are sampled or aggregated.

10-2.5 Management Mapping

A device that implements both this profile and an SNMP Agent may map profile attributes to an SNMP MIB in accordance with Appendix G. Such mapping is optional and shall not alter the semantics of the CIP-visible profile attributes.

10-2.6 Non-Security Statement

Implementation of this profile shall not be interpreted as evidence that a device provides firewall functionality, security-policy enforcement, security-zone separation, cybersecurity conformance, or end-to-end path validation.

10-2.7 Conformance Claim

A device claiming conformance to the EtherNet/IP Bridge and Boundary Diagnostics Profile shall identify the supported profile revision, the number of bridge instances implemented, and each profile-defined diagnostic attribute supported by each bridge instance.

13.4 Expansion 4: Concentration and Chokepoint Health Expansion

Rationale

A segmented architecture may remain protocol-conformant while a gateway, router, or forwarding device becomes an unobserved dependency. The paper calls this concern boundary quality: cross-boundary communication should be transparent enough to review, diagnose, and govern.[1]

Chapter 10, Bridging & Routing, is the correct primary location because Volume 2 explicitly reserves it for EtherNet/IP-specific additions to CIP bridging and routing and currently contains no such additions.[2]

Recommended Location

Add Sections 10-2 through 10-2.7 after the Chapter 10 introduction. Add an optional informative SNMP mapping in Appendix G for devices that also implement an SNMP Agent.

Proposed Specification Text

10-2 Bridge and Boundary Diagnostics Profile

10-2.1 Scope

This section defines an optional EtherNet/IP Bridge and Boundary Diagnostics Profile for a device that implements EtherNet/IP-specific CIP bridging or routing between two or more local network interfaces.

For the purposes of this section, a bridge instance is a local forwarding function within the device that forwards CIP communication between a defined set of local network interfaces. A bridge instance is not required to represent a security zone, a firewall, a complete route, or a complete network path.

A device that does not implement EtherNet/IP-specific CIP bridging or routing shall not be required to implement this profile.

10-2.2 Profile Disclosure

A device that implements this profile shall make the following information available through an explicit-message diagnostic service or a profile-defined read-only object:

1. The number of bridge instances implemented by the device.
2. For each bridge instance, a locally unique bridge-instance identifier.
3. The local network-interface identifiers associated with that bridge instance.
4. The supported encapsulation transport mode or modes used by the bridge instance, where applicable.
5. The diagnostic attributes defined in Section 10-2.3.

A device shall indicate an unsupported attribute using the standard CIP service or attribute-access behavior applicable to the object or service through which the profile is implemented.

10-2.3 Diagnostic Attributes

For each bridge instance, the device shall provide the following read-only diagnostic attributes:

Attribute	Required behavior
Forwarded CIP communication count	A monotonically increasing count of CIP communications forwarded by the bridge instance since the last counter discontinuity. The implementation shall state whether the unit is a message, packet, or connection event.
Forwarding discard count	A monotonically increasing count of locally detected forwarding discards attributable to the bridge instance since the last counter discontinuity.
Forwarding error count	A monotonically increasing count of locally detected forwarding errors attributable to the bridge instance since the last counter discontinuity. A remote-path or external-switch condition shall not be reported as a local forwarding error unless the device can determine that condition locally.
Counter discontinuity indicator	An indicator or timestamp that changes when one or more profile counters are reset, restarted, or otherwise lose continuity.
Administrative state	A read-only state indicating whether the local forwarding function is administratively enabled, administratively disabled, or not determinable by the device.

The profile shall not require a device to report end-to-end latency, upstream switch utilization, remote topology, remote peer identity, security-zone identity, route completeness, or an inferred health score.

10-2.4 Access and Performance Constraints

Profile attributes shall be read-only.

The profile shall not require periodic implicit messaging, cyclic production, unsolicited notification, or modification of existing EtherNet/IP real-time transport behavior.

A device may limit the rate at which diagnostic attributes are updated, provided that the update behavior is documented and the device does not represent the attributes as instantaneous measurements when they are sampled or aggregated.

10-2.5 Management Mapping

A device that implements both this profile and an SNMP Agent may map profile attributes to an SNMP MIB in accordance with Appendix G. Such mapping is optional and shall not alter the semantics of the CIP-visible profile attributes.

10-2.6 Non-Security Statement

Implementation of this profile shall not be interpreted as evidence that a device provides firewall functionality, security-policy enforcement, security-zone separation, cybersecurity conformance, or end-to-end path validation.

10-2.7 Conformance Claim

A device claiming conformance to the EtherNet/IP Bridge and Boundary Diagnostics Profile shall identify the supported profile revision, the number of bridge instances implemented, and each profile-defined diagnostic attribute supported by each bridge instance.

13.5 Expansion 5: Segmented Deployment Profile Expansion

Rationale

A segmented architecture may remain protocol-conformant while a gateway, router, or forwarding device becomes an unobserved dependency. The paper calls this concern boundary quality: cross-boundary communication should be transparent enough to review, diagnose, and govern.1

Chapter 10, Bridging & Routing, is the correct primary location because Volume 2 explicitly reserves it for EtherNet/IP-specific additions to CIP bridging and routing and currently contains no such additions.2

Recommended Location

Add Sections 10-2 through 10-2.7 after the Chapter 10 introduction. Add an optional informative SNMP mapping in Appendix G for devices that also implement an SNMP Agent.

Proposed Specification Text

10-2 Bridge and Boundary Diagnostics Profile

10-2.1 Scope

This section defines an optional EtherNet/IP Bridge and Boundary Diagnostics Profile for a device that implements EtherNet/IP-specific CIP bridging or routing between two or more local network interfaces.

For the purposes of this section, a bridge instance is a local forwarding function within the device that forwards CIP communication between a defined set of local network interfaces. A bridge instance is not required to represent a security zone, a firewall, a complete route, or a complete network path.

A device that does not implement EtherNet/IP-specific CIP bridging or routing shall not be required to implement this profile.

10-2.2 Profile Disclosure

A device that implements this profile shall make the following information available through an explicit-message diagnostic service or a profile-defined read-only object:

1. The number of bridge instances implemented by the device.
2. For each bridge instance, a locally unique bridge-instance identifier.
3. The local network-interface identifiers associated with that bridge instance.

4. The supported encapsulation transport mode or modes used by the bridge instance, where applicable.
5. The diagnostic attributes defined in Section 10-2.3.

A device shall indicate an unsupported attribute using the standard CIP service or attribute-access behavior applicable to the object or service through which the profile is implemented.

10-2.3 Diagnostic Attributes

For each bridge instance, the device shall provide the following read-only diagnostic attributes:

Attribute	Required behavior
Forwarded CIP communication count	A monotonically increasing count of CIP communications forwarded by the bridge instance since the last counter discontinuity. The implementation shall state whether the unit is a message, packet, or connection event.
Forwarding discard count	A monotonically increasing count of locally detected forwarding discards attributable to the bridge instance since the last counter discontinuity.
Forwarding error count	A monotonically increasing count of locally detected forwarding errors attributable to the bridge instance since the last counter discontinuity. A remote-path or external-switch condition shall not be reported as a local forwarding error unless the device can determine that condition locally.
Counter discontinuity indicator	An indicator or timestamp that changes when one or more profile counters are reset, restarted, or otherwise lose continuity.
Administrative state	A read-only state indicating whether the local forwarding function is administratively enabled, administratively disabled, or not determinable by the device.

The profile shall not require a device to report end-to-end latency, upstream switch utilization, remote topology, remote peer identity, security-zone identity, route completeness, or an inferred health score.

10-2.4 Access and Performance Constraints

Profile attributes shall be read-only.

The profile shall not require periodic implicit messaging, cyclic production, unsolicited notification, or modification of existing EtherNet/IP real-time transport behavior.

A device may limit the rate at which diagnostic attributes are updated, provided that the update behavior is documented and the device does not represent the attributes as instantaneous measurements when they are sampled or aggregated.

10-2.5 Management Mapping

A device that implements both this profile and an SNMP Agent may map profile attributes to an SNMP MIB in accordance with Appendix G. Such mapping is optional and shall not alter the semantics of the CIP-visible profile attributes.

10-2.6 Non-Security Statement

Implementation of this profile shall not be interpreted as evidence that a device provides firewall functionality, security-policy enforcement, security-zone separation, cybersecurity conformance, or end-to-end path validation.

10-2.7 Conformance Claim

A device claiming conformance to the EtherNet/IP Bridge and Boundary Diagnostics Profile shall identify the supported profile revision, the number of bridge instances implemented, and each profile-defined diagnostic attribute supported by each bridge instance.

13.6 Interpretation for Specification Revision

Across all five experiments, the recurring pattern is that the specification expansions create efficiency gains by making EtherNet/IP architectures more **intentional, visible, and governable**. The modeled benefit does not come from adding more traffic, more reach, or more concentration indiscriminately. It comes from adding discipline to how traffic roles are defined, how multicast is bounded, how bridges are observed, how chokepoints are diagnosed, and how segmented deployments are profiled.

This is important because it translates the paper's conceptual contribution into a standards-facing proposal set without forcing the stronger and much harder claim that the core transport semantics of EtherNet/IP are themselves defective. Instead, the results support a narrower and more defensible conclusion: **the specification can be expanded so that architectural behavior becomes easier to interpret, govern, and reproduce in segmented industrial systems**.

14. Limitations

This paper remains conceptual rather than empirical in the strict experimental sense. The reported simulations are **model-based comparative results**, not plant measurements or conformance-test outcomes. Their inputs were deliberately theory-informed priors chosen to represent plausible contrasts among architecture alternatives, and they should be interpreted accordingly.

A second limitation is translational. Social-network and organizational-communication concepts are being borrowed into industrial analysis. The paper argues that such translation is useful, but only if the industrial analogue is made explicit and the concepts are not applied uncritically. The strongest claims in this paper concern **bridges, concentration, modularity, and boundary effects**; the weakest concern universal statements about media richness and the equivalence of forced and congruent interaction.[6] [7] [9]

A third limitation concerns reproducibility across platforms. The original model was implemented in Excel using volatile random-number functions, while the reported results in this paper were reproduced in Python with a fixed seed to create a stable reference output. The reported tables therefore represent a reproducible rendering of the same model assumptions, not a claim that one specific Excel recalculation state is uniquely authoritative.[11] [12]

15. Conclusion

This paper has argued that industrial networks should be redefined not only as topologies of connected devices operating over compliant protocols, but also as **interaction systems** whose resilience, scalability, diagnosability, and coordination behavior depend on interaction density, path structure, concentration points, communication-function fit, and boundary quality. The literature supports the use of such heuristics, but only when they are treated as conditional descriptors rather than fixed laws.[1] [4] [6] [8]

The reported experiment method and results strengthen that conclusion. Across five structured Monte Carlo comparisons, the model consistently favored broader architectural reach over local busyness, purposeful path diversity over unmanaged cross-coupling, function-aligned communication over indiscriminate traffic growth, bounded dual hubs over uncontrolled centralization, and disciplined bridge governance over opaque boundary dependency.[11] Those findings do not prove protocol laws, but they do show that the framework can generate stable and engineering-relevant comparative judgments.

The strongest parts of the framework are therefore the claims about **vector structure, brokerage, concentration, modularity, and cascading dependence**, because these are supported by the broader literature, aligned with recognized network-science concepts, and reinforced by the reported comparative experiments.[5] [8] [9] The paper's added value for industrial protocol specialists is methodological: it offers a more precise language for architecture review, segmentation analysis, interoperability planning, and resilience reasoning in systems where protocol correctness alone does not fully explain behavior.

The broader conclusion is not that topology matters less than previously thought. It is that topology matters differently when interpreted through patterns of interaction. The industrial network is not only the topology; it is the topology plus the pattern of coordination and dependency that moves across it.

References

[1] UC Berkeley D-Lab: Concepts and Measurements in Social Network Analysis

[2] Wu, Waber, Aral, Brynjolfsson, Pentland. Mining Face-to-Face Interaction Networks Using Sociometric Badges

[3] Open University OpenLearn: The effects of distance on communication

[4] Hossain and Wu. Communications network centrality correlates to organisational coordination

[5] Granovetter. The Strength of Weak Ties

[6] Brashears and Quintane. The weakness of tie strength

[7] El-Shinnawy and Markus. The poverty of media richness theory

[8] Ding, Shi, and Xiao. Unveiling the Impact of Communication Network on Engineering Project Team Performance

[9] 10 Social Network Heuristics Transferable to IT, Logistics, and Industrial Networks

[10] Conway, Melvin E. How Do Committees Invent?

[11] The CIP Networks Library Volume 2: EtherNet/IP Adaptation of CIP, Edition 1.31

The ideas, opinions, and recommendations expressed herein are intended to describe concepts of the author for the possible use of ODVA technologies and do not reflect the ideas, opinions, or recommendations of ODVA per se. Because ODVA technologies may be applied in many diverse situations and in conjunction with products and systems from multiple vendors, readers and those responsible for specifying ODVA networks must determine for themselves the suitability of the ideas, opinions, and recommendations expressed herein for their intended use.
Copyright ©2026 ODVA, Inc. All rights reserved. For permission to reproduce excerpts of this material, with appropriate attribution to the author, please contact ODVA: TEL +1 734-975-8840; FAX +1 734-922-0027; EMAIL odva@odva.org; WEB www.odva.org.
CIP, Common Industrial Protocol, CIP Energy, CIP Motion, CIP Safety, CIP Sync, CIP Security, CompoNet, ControlNet, DeviceNet, and EtherNet/IP are trademarks of ODVA, Inc. All other trademarks are property of their respective owners.