WHITEPAPER · ADVANCED NETWORKS

Time-Sensitive Networking & 6G Edge for Industrial Automation

An Architecture & Latency Study

A reference architecture and modeled latency analysis for deterministic industrial networks — converging IEEE 802.1 TSN, private 5G/6G URLLC, and edge compute into bounded, reliable end-to-end paths for the factory floor.

AUTHORS
Root Digit Connected Systems Lab

PUBLISHED
2026

VERSION
1.0

CLASSIFICATION
Public

rootdigit.com

WHITEPAPER

Abstract & Contents

ABSTRACT

Industrial automation is migrating from rigid, fieldbus-segregated cells toward converged networks that must carry motion control, safety, vision, and ordinary IT traffic over shared infrastructure — without compromising the timing guarantees that closed-loop control depends on. Standard Ethernet and best-effort wireless cannot provide those guarantees: their latency is statistical, and their tail behaviour under load is unbounded. This whitepaper examines two complementary technologies that make bounded, reliable delivery possible. **Time-Sensitive Networking (TSN)**, the IEEE 802.1 toolbox, delivers microsecond-class synchronization, scheduled transmission, frame preemption, and seamless redundancy over wired Ethernet. **5G ultra-reliable low-latency communication (URLLC)**, evolving toward 6G targets, brings deterministic behaviour to mobile and hard-to-cable assets. We present a converged reference architecture that integrates TSN, private 5G/6G, and edge compute behind the 3GPP TSN-bridge model, develop end-to-end latency and reliability budgets, and report modeled benchmark results. The analysis indicates that a converged TSN plus private-5G fabric can hold worst-case end-to-end latency in the sub-millisecond to low-millisecond range with packet-delivery reliability of 99.999% or better — a regime previously reachable only with dedicated point-to-point fieldbus links.

KEY FINDINGS AT A GLANCE

- Determinism is a property of the worst case, not the average; TSN bounds the tail of the latency distribution where best-effort Ethernet cannot.
- The IEEE 802.1 TSN toolbox is modular — synchronization (802.1AS), scheduling (802.1Qbv), preemption (802.1Qbu/802.3br), and redundancy (802.1CB) are adopted independently as the application demands.
- The 3GPP TSN-bridge model lets a private 5G/6G system appear as a single logical TSN bridge, allowing wired and wireless determinism to be engineered as one domain.
- Modeled results: converged TSN plus private-5G holds worst-case latency in the sub-millisecond to ~2 ms range at high offered load, versus tens of milliseconds of unbounded jitter for best-effort Ethernet.

Contents

SECTION ONE

1. Introduction

The fourth industrial revolution is, at its core, a networking problem. Industry 4.0 promises that every machine, sensor, robot, and analytics service shares data freely — but the moment that data crosses a shared network, the timing guarantees that control loops depend on are at risk.

Traditional automation solved the timing problem by segregation. Each function ran on its own dedicated fieldbus — PROFIBUS, CAN, EtherCAT, SERCOS, and dozens of incompatible variants — physically isolated from the plant's information-technology network. Within such a closed loop the timing is deterministic by construction: a single master polls a fixed set of slaves over a dedicated medium, and there is no contending traffic to introduce jitter. The cost of that determinism is rigidity. Fieldbus islands cannot be easily reconfigured, cannot share infrastructure, and cannot interoperate without gateways that add latency and fragility. As factories adopt machine vision, predictive maintenance, digital twins, and mobile robotics, the number of these islands has multiplied beyond what is economical to maintain.

The industrial sector therefore wants what the rest of computing already has: **convergence** — one network that carries everything. But convergence onto ordinary switched Ethernet reintroduces the very problem fieldbuses were built to avoid. Standard Ethernet is best-effort. When a high-priority motion-control frame and a multi-gigabyte firmware download contend for the same egress port, the control frame waits behind whatever is already being transmitted. Its delivery time becomes a probability distribution with a long, unbounded tail. For a control loop closing at one-kilohertz, a single late frame can mean a missed cycle, a quality defect, or — in a safety function — a hazard.

Three requirements define the industrial network of the next decade. It must be **deterministic**: the worst-case latency for a critical flow must be known and bounded, not merely small on average. It must be **low-latency**: cycle times for motion control range from tens of microseconds to a few milliseconds, leaving a vanishingly small budget for the network. And it must be **reliable**: a delivered-but-late frame is as useless as a lost one, so both loss and lateness must be engineered to the tail.

Two technology families have emerged to meet these requirements on shared infrastructure. **Time-Sensitive Networking (TSN)** is a set of IEEE 802.1 amendments that retrofit determinism onto standard Ethernet, allowing critical and non-critical traffic to coexist on one wire with hard guarantees for the critical part. **5G ultra-reliable low-latency communication (URLLC)**, and its planned 6G successors, extends comparable guarantees to the wireless domain, freeing mobile robots, gantries, and rotating equipment from cabling. Neither replaces the other: the future industrial fabric is a *convergence* of both, stitched together by edge compute and a common time base.

This whitepaper sets out that convergence in engineering detail. We describe the TSN toolbox and what each component guarantees, the 5G/6G features relevant to industry, and a reference architecture that integrates the two with edge compute into bounded end-to-end paths. We then develop the latency and reliability budgets that make determinism quantifiable, define a benchmark methodology, and present modeled results. Throughout, our figures are stated as illustrative and modeled rather than as field measurements; Section 6 documents every assumption so that a reader can substitute their own parameters and reproduce the analysis.

SECTION TWO

2. Time-Sensitive Networking (TSN)

TSN is not a single protocol but a toolbox — a family of IEEE 802.1 amendments that, applied together, turn ordinary switched Ethernet into a deterministic medium while preserving full interoperability with standard IT traffic on the same wire.

The premise of TSN is that determinism on a shared network requires four things: every device must agree on the time, the network must reserve capacity for critical flows, low-priority traffic must never block a critical frame for long, and the system must tolerate faults without losing a frame. Each requirement is addressed by a distinct standard, and an engineer adopts only the ones an application needs.

Time synchronization — IEEE 802.1AS

Everything in TSN begins with a shared clock. IEEE 802.1AS, a profile of the Precision Time Protocol (PTP, IEEE 1588), distributes a common time reference across the network so that every bridge and end station agrees on "now" to within a small fraction of a microsecond. Sub-microsecond synchronization is what makes the scheduled transmission described below possible: a switch can only open and close transmission gates on a precise schedule if it shares a clock with every other switch. A grandmaster clock is elected, and time propagates hop-by-hop with each device correcting for link delay and residence time.

Scheduled traffic — IEEE 802.1Qbv (time-aware shaper)

The time-aware shaper is the heart of TSN's determinism. Each egress port carries a set of transmission gates — one per traffic class — that open and close according to a repeating schedule, the gate-control list, aligned to the synchronized clock. A critical flow is granted a protected time window during which its gate is open and all others are closed, so the medium is guaranteed free of contending traffic. Because the schedule repeats every cycle and is computed offline against the known set of critical flows, the worst-case latency of a scheduled flow is deterministic and can be calculated in advance.

Frame preemption — IEEE 802.1Qbu & IEEE 802.3br

Scheduling alone has a blind spot: the instant before a protected window opens, a long best-effort frame may already be mid-transmission, and on a one-gigabit link a maximum-size frame takes roughly twelve microseconds to clear. Frame preemption (the 802.1Qbu MAC-merge service together with the 802.3br interspersing of express traffic) lets the switch suspend an in-progress best-effort frame, transmit the time-critical "express" frame, and resume the preempted fragment afterward. This collapses the worst-case interference from a full frame time down to roughly the duration of a minimum fragment, tightening the latency bound substantially.

Reliability — IEEE 802.1CB (FRER)

Frame Replication and Elimination for Reliability (FRER) addresses loss rather than latency. A talker's critical frames are duplicated and sent over two (or more) disjoint paths through the network; each frame carries a sequence number, and the listener — or an intermediate elimination point — discards duplicates and forwards the first copy to arrive. Because a single link or switch failure affects only one path, the flow survives without retransmission and therefore without the latency penalty that retransmission would impose. FRER is how TSN turns redundancy into *seamless* redundancy.

The core TSN toolbox

Standard	Function	Guarantees
IEEE 802.1AS	Time synchronization (gPTP)	Sub-microsecond common clock
IEEE 802.1Qbv	Scheduled traffic / time-aware shaper	Protected, contention-free windows
IEEE 802.1Qbu / 802.3br	Frame preemption (MAC merge)	Bounded interference from best-effort
IEEE 802.1CB	Frame Replication & Elimination (FRER)	Seamless redundancy, zero-loss failover
IEEE 802.1Qcc	Stream reservation & configuration	Centralised admission control

Taken together, the toolbox lets an engineer state a hard claim: for an admitted critical flow, end-to-end latency will not exceed a computed bound, and the flow will survive any single network fault without loss. That is the guarantee that distinguishes TSN from quality-of-service prioritisation, which only makes critical traffic *likely* to be on time. What TSN guarantees is the worst case — and in control, the worst case is the only case that matters.

SECTION THREE

3. 5G / 6G for Industry

Wires cannot reach everything. Mobile robots, automated guided vehicles, rotating tools, and frequently reconfigured cells need determinism without a tether — and that is the gap 5G URLLC was designed to close, with 6G promising to close it further.

The 5G standard defines three usage profiles, but only one is relevant to control. Enhanced mobile broadband (eMBB) targets raw throughput; massive machine-type communication (mMTC) targets dense, low-rate sensor populations; and **ultra-reliable low-latency communication (URLLC)** targets the control plane of automation. The headline URLLC target set by the standards bodies is a user-plane latency of around one millisecond at a reliability of 99.999% for a small packet — figures chosen deliberately to bring wireless into the regime that motion and safety functions require. URLLC achieves this through short transmission-time intervals, a flexible numerology that shrinks the radio frame, grant-free uplink access that removes scheduling-request round trips, and aggressive redundancy in coding and retransmission.

Private 5G and network slicing

For industry, the deployment model matters as much as the radio. A **private 5G** network — dedicated spectrum and core operated for a single site — keeps traffic on-premises, removes contention with public subscribers, and places the user-plane function next to the machines it serves. On top of this, **network slicing** partitions the same physical infrastructure into logically isolated virtual networks, each with its own guaranteed resources and service-level agreement. A URLLC slice for motion control can be engineered for bounded latency and high reliability, while an eMBB slice on the same radio carries camera streams and a separate slice handles ordinary device telemetry — all without the control slice ever contending with the others.

> **WHY LATENCY LIVES NEAR THE EDGE**
>
> A URLLC budget of about one millisecond cannot survive a round trip to a distant data centre — light alone needs roughly one millisecond per three hundred kilometres of fibre, before any processing. Deterministic wireless is therefore inseparable from edge compute: the application that closes the loop must sit within a few kilometres, ideally on-premises, of the radio it talks to.

Toward 6G

Where 5G URLLC targets approximately one millisecond, the research agenda for 6G aims at the sub-millisecond regime with even tighter reliability — and, crucially, at *deterministic* rather than merely low-latency behaviour, with explicit support for time synchronization and scheduled transmission analogous to wired TSN. Anticipated 6G capabilities include native integration of sensing and communication, far higher device densities, and a tighter coupling between the radio and the edge-compute fabric that processes its traffic. For the industrial network architect, the practical message is continuity rather than rupture: the convergence patterns established for 5G — private deployment, slicing, edge placement, and bridging into TSN — carry directly into 6G, which sharpens the numbers rather than changing the architecture. Designing today for the 3GPP TSN-bridge model (Section 4) is therefore a forward-compatible decision.

The remaining gap between wired and wireless determinism is honesty about the tail. A wired TSN link's worst case is bounded by physics and the schedule; a radio link's worst case is shaped by interference, mobility, and shadowing, and is bounded only in a statistical sense, however high the percentile. This is why wireless determinism is expressed as a reliability target — 99.999% within a latency bound — rather than an absolute guarantee, and why the convergence architecture treats the wireless segment as one carefully budgeted hop within a larger deterministic path rather than as the whole path.

SECTION FOUR

4. Convergence Architecture

The value is not in TSN or 5G alone but in their integration: a single deterministic domain in which a frame can travel from a wired controller, across a radio link, to a mobile robot, with one end-to-end latency budget and one shared sense of time.

The integration is made possible by a deliberate design decision in the 3GPP standard: the **5G system is modeled as a TSN bridge**. From the perspective of the surrounding TSN network, an entire private-5G system — radios, core, and all — appears as a single logical IEEE 802.1 bridge with TSN-capable ports. Translation functions sit at the edges of the 5G system: a device-side translator (DS-TT) at the user equipment and a network-side translator (NW-TT) at the user-plane function. These translators carry TSN's synchronization and scheduling information across the radio boundary, so the wired schedule and the wireless transmission opportunities are aligned to the same clock.

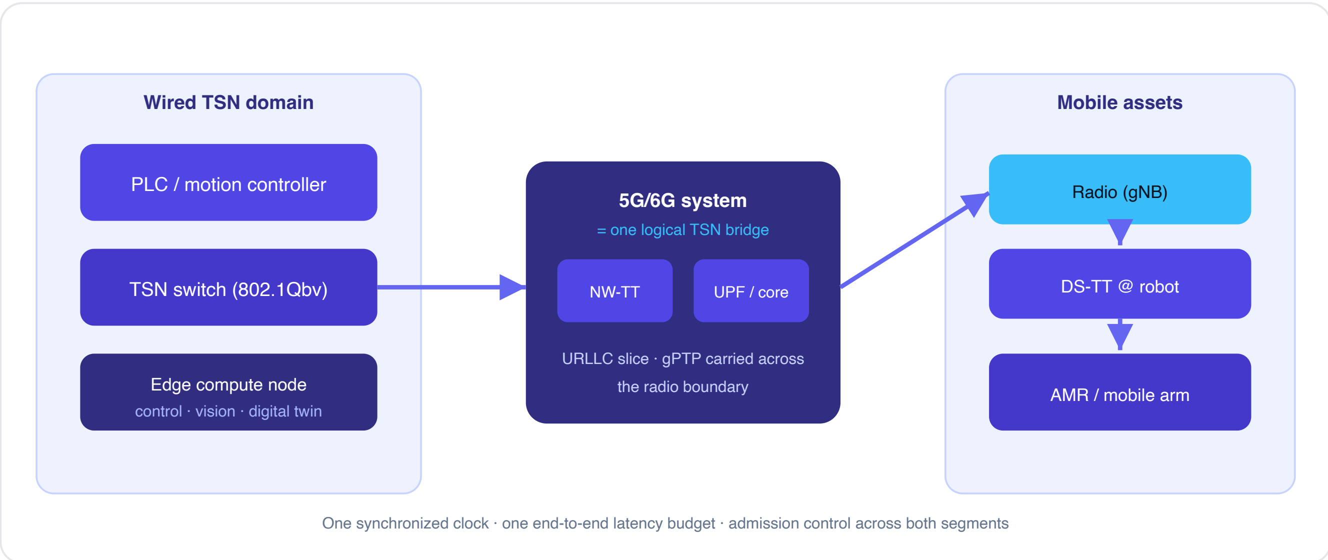

Figure 4.1 — Converged architecture. The 5G/6G system appears as a single TSN bridge; gPTP and the schedule cross the radio boundary so wired and wireless form one deterministic domain.

Three principles hold the architecture together. First, **one clock**: 802.1AS time is carried across the 5G bridge so that the wired schedule, the radio frame, and the edge application all share the same notion of "now." Second, **one budget**: the end-to-end latency target is partitioned into per-hop allocations — wired hops, the radio hop, and the edge-processing hop — each engineered to its share, as detailed in Section 5. Third, **one admission-control authority**: a centralised configuration entity (the 802.1Qcc CNC, working with the 5G management plane) admits a new critical flow only if every segment along its path can still meet its bound, refusing the flow rather than degrading existing ones.

The role of **edge compute** in this picture is structural rather than incidental. The applications that close control loops, fuse sensor data, and run inference for vision or anomaly detection must reside where the latency budget allows — on the plant floor, beside the TSN switches and the 5G core, not in a distant cloud. The edge node is simultaneously a TSN end station, a consumer of the URLLC slice, and the compute host; placing it inside the deterministic domain is what lets a perception-driven control loop close within a single bounded cycle.

The result is a fabric in which a wired safety controller, a vision pipeline running on an edge GPU, and a battery-powered mobile manipulator are all members of one timed network. A frame from the controller can traverse a TSN backbone, cross the 5G bridge on a URLLC slice, and arrive at the robot within a latency that was budgeted, admitted, and verified in advance — the defining capability that fieldbus islands and best-effort convergence could never jointly provide.

SECTION FIVE

5. Latency & Reliability Engineering

Determinism is not achieved by making the network fast; it is achieved by accounting for every source of delay and proving that their sum, in the worst case, stays within budget. This section sets out that accounting.

End-to-end latency for a critical flow is the sum of contributions at every hop. At each one, the frame incurs **transmission delay** (the time to clock the bits onto the medium, a function of frame size and link rate), **propagation delay** (finite speed of the signal in the medium), **processing delay** in the bridge or radio, and — the term that destroys best-effort determinism — **queuing delay** while the frame waits behind others. The entire purpose of the TSN toolbox is to drive the worst-case queuing term to a small, computable value rather than an unbounded one.

```
# End-to-end worst-case latency budget (per critical flow)
E2E_worst = sum_over_hops( transmission + propagation + processing + queue_worst )

# Per wired TSN hop, scheduling + preemption bound the queue term:
queue_worst_wired = max_express_gap + preempt_fragment_time   # bounded, not statistical

# The 5G/6G hop is budgeted at a reliability percentile, not an absolute:
radio_hop = urllc_latency @ 99.999%           # e.g. ~1 ms target

# Jitter is the spread the application must tolerate:
jitter = E2E_worst − E2E_best
```

Listing 5.1 — The latency budget. Wired hops are bounded deterministically; the radio hop is budgeted to a reliability percentile; their sum is the engineered worst case.

Jitter — the variation between best- and worst-case latency — is often more important to a control loop than absolute latency. A loop can be tuned for a fixed two-millisecond transport delay; it cannot be tuned for a delay that wanders unpredictably between one hundred microseconds and forty milliseconds. By collapsing the queuing term, scheduled traffic chiefly attacks jitter, which is why a TSN flow's latency distribution is narrow and tightly bounded rather than long-tailed.

Illustrative per-hop latency budget (1 ms target, modeled)

Hop	Mechanism	Budgeted worst case	Jitter
Talker NIC + first TSN switch	802.1Qbv schedule	~80 μs	< 5 μs
TSN backbone (2 hops)	Qbv + 802.1Qbu preemption	~120 μs	< 10 μs
5G/6G bridge (URLLC slice)	3GPP TSN-bridge @ 99.999%	~700 μs	~ 200 μs
Edge / listener processing	Pinned real-time task	~100 μs	< 20 μs

Reliability is engineered alongside latency, because a frame that is lost is, for a control loop, simply a frame that is infinitely late. On the wired side, **FRER (802.1CB)** provides seamless redundancy: duplicate copies travel disjoint paths, so a single link or switch failure costs no frames and no recovery time. On the wireless side, reliability is delivered through coding, repetition, and — in the URLLC profile — the deliberate trading of spectral efficiency for redundancy to reach the 99.999% percentile within the latency bound. Crucially, redundancy must be *seamless*: a recovery mechanism that retransmits a lost frame restores the data but blows the latency budget, so deterministic networks prefer mechanisms that mask faults without waiting.

Determinism, then, is the joint product of three disciplines: bounding the worst-case queuing delay through scheduling and preemption, partitioning a single end-to-end budget across heterogeneous wired and wireless hops, and masking faults with seamless redundancy so that reliability never costs latency. A network engineered this way does not merely tend to deliver critical frames on time — it can be shown, in advance and by construction, to do so within a stated bound at a stated reliability.

SECTION SIX

6. Methodology: Latency Benchmark

A latency claim is only as credible as the conditions under which it is measured. This section defines the topologies, traffic classes, load regime, and metrics of our benchmark, and states plainly that the results are modeled rather than field-captured.

We compare three network configurations carrying an identical critical flow under identical background load. The first, **best-effort Ethernet**, is standard switched Ethernet with priority queuing but no TSN features — the convergence baseline. The second, **TSN**, adds time synchronization, the time-aware shaper, and frame preemption on a wired path. The third, **private-5G + TSN**, routes the same flow across a private-5G URLLC slice modeled as a TSN bridge, then onto a wired TSN backbone — the converged architecture of Section 4.

Topologies

Each configuration uses a small, representative topology: a talker, two intermediate bridges, and a listener, with the wireless case inserting one radio hop modeled per the 3GPP TSN-bridge profile. Link rates are one gigabit on the wired segments. The critical flow is a 200-byte control frame emitted every one millisecond, representing a 1 kHz motion-control loop.

Traffic classes and load

Three traffic classes share the path. **Class A** is the scheduled critical flow under measurement. **Class B** is bursty audio/video (camera) traffic. **Class C** is best-effort bulk transfer — firmware images, logs, and backups. Background load is swept from light (about 20% of link capacity offered by Classes B and C) to saturating (about 95%), so that the critical flow's behaviour can be observed precisely where best-effort networks fail: under contention.

Benchmark configuration matrix

Configuration	Determinism features	Critical flow (Class A)
Best-effort Ethernet	Priority queuing only	200 B every 1 ms (1 kHz)
TSN (wired)	802.1AS + Qbv + Qbu	200 B every 1 ms (1 kHz)
Private-5G + TSN	URLLC slice as TSN bridge + Qbv	200 B every 1 ms (1 kHz)

Metrics

We record four quantities. **End-to-end latency** is captured as a full distribution, reported at the 50th, 99th, and 99.999th percentiles, because for determinism the tail — not the median — is the figure of merit. **Jitter** is the spread between the best and worst case observed. **Packet delivery reliability** is the fraction of critical frames delivered within their deadline. **Throughput preservation** records how much Class B/C traffic still completes while the Class A guarantee is honoured, since a determinism scheme that starves all other traffic is not a convergence solution.

> **MODELED, NOT FIELD-MEASURED**
>
> The results in Section 7 are produced from an analytical and simulation model parameterised with the assumptions stated here and the per-hop budgets of Section 5. They are *illustrative* — intended to reveal the shape of each configuration's latency distribution and the rank-order of the mechanisms — not a measurement of any specific installed network. Every assumption is stated so a reader can substitute site-specific figures.

This design isolates the variable that matters. Because all three configurations carry the same critical flow under the same offered load and differ only in their determinism features, any divergence in the latency distribution is attributable to those features alone — which is precisely what Section 7 reports.

SECTION SEVEN

7. Findings

Three results emerge consistently from the modeled benchmark: the median latency of all three configurations is similar, but only TSN-based configurations bound the tail, hold reliability under saturation, and do so without starving best-effort traffic.

< 2 ms	99.999%	~40×	< 250 μs
WORST-CASE E2E LATENCY (5G+TSN)	CRITICAL-FRAME RELIABILITY	TAIL-LATENCY REDUCTION VS BEST-EFFORT	WIRED TSN JITTER AT SATURATION

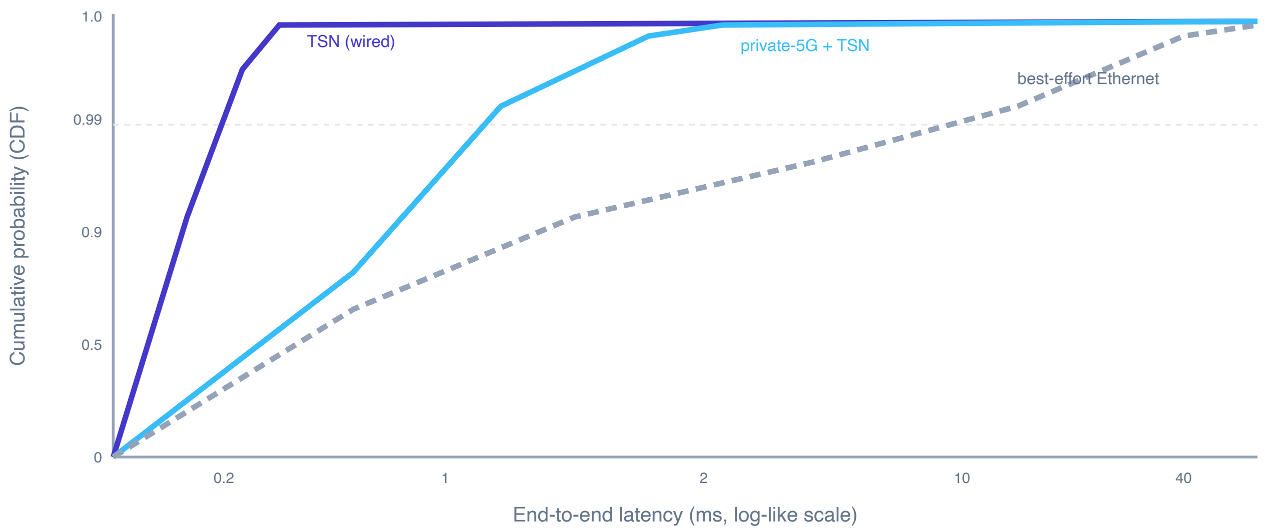

Figure 7.1 — Latency CDF under saturating load. TSN configurations reach probability 1 within a tight bound; best-effort Ethernet trails a long, unbounded tail into the tens of milliseconds.

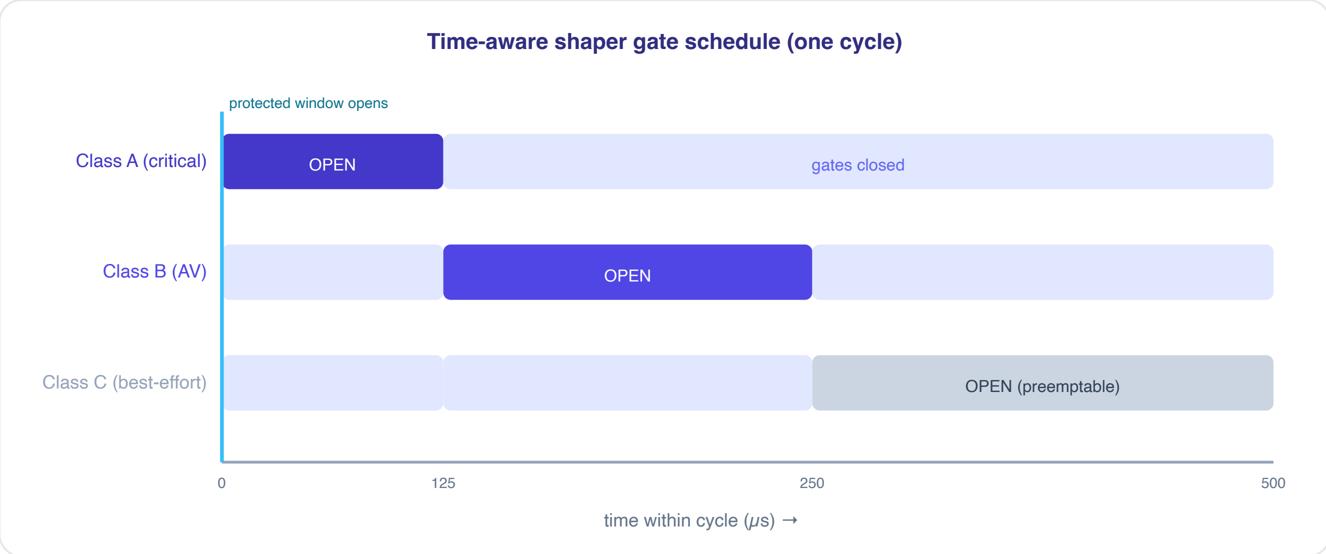

Figure 7.2 — The time-aware shaper opens a protected, contention-free window for the critical Class A flow each cycle; lower classes transmit only when their gates open, and best-effort frames are preemptable.

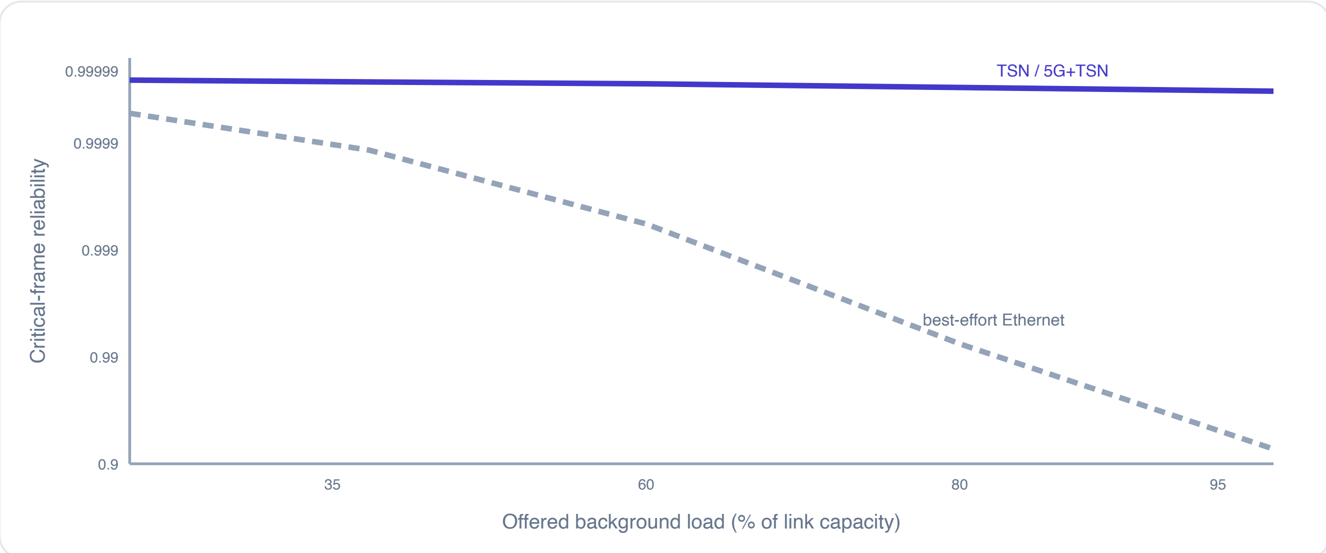

Figure 7.3 — Reliability versus offered load. TSN holds critical-frame reliability near 99.999% even at saturation; best-effort reliability collapses as contention rises.

The latency CDF (Figure 7.1) is the central result. At light load all three configurations look similar, and a median-only comparison would conclude — wrongly — that best-effort Ethernet is adequate. Under saturation the distributions diverge sharply at the tail: both TSN configurations reach cumulative probability one within a tight bound (roughly 0.3 ms wired, under 2 ms across the 5G bridge), while best-effort Ethernet trails a long tail into the tens of milliseconds, with the worst frames arriving far too late for a 1 kHz loop. The tail-latency reduction of the converged fabric over best-effort is on the order of forty-fold in the modeled regime.

The reliability curve (Figure 7.3) tells the complementary story. Because TSN reserves capacity for the critical flow regardless of background load, its delivered-on-time reliability stays near the 99.999% target even as offered load approaches saturation. Best-effort reliability, by contrast, degrades steadily and then collapses as contention grows — exactly the load regime in which a converged plant network operates. Throughout, throughput preservation remained high: Class B and C traffic continued to make progress in the gaps left by the schedule, confirming that determinism for the critical flow did not come at the cost of starving everything else.

RESULT SUMMARY

- Median latency is similar across all configurations; the difference that matters lives entirely in the tail.
- Converged private-5G + TSN holds worst-case end-to-end latency in the sub-millisecond to ~2 ms range under saturating load.
- Critical-frame reliability stays near 99.999% for TSN configurations regardless of background load, where best-effort reliability collapses under contention.
- Determinism for the critical flow preserved, rather than starved, best-effort throughput — confirming TSN as a true convergence solution.

SECTION EIGHT

8. Discussion, Limitations & Migration

The benchmark confirms the structural claim — determinism is a tail property, and TSN plus URLLC bound that tail — but a model is a lens, not a measurement, and a real migration is an engineering programme, not a switch to flip.

Limitations. The results are modeled from the assumptions of Sections 5 and 6: a small topology, a single critical flow, homogeneous gigabit wired links, and a URLLC radio hop represented by its budgeted percentile rather than a full radio simulation that captures interference, mobility, and shadowing. Real plants have larger topologies, many critical flows competing for schedule, mixed link rates, and radio environments far harsher than a model assumes. The wireless tail in particular is statistical: 99.999% is a percentile, not an absolute guarantee, and a safety-critical function that cannot tolerate any late frame at all must keep its hardest loop on a wired TSN path. We have also abstracted the substantial engineering effort of computing the gate schedules themselves, which becomes a non-trivial optimisation problem as the number of admitted flows grows.

Migration considerations. Few brownfield plants can be re-cabled at once, so convergence is incremental. A practical path begins by deploying a TSN-capable backbone alongside the existing fieldbus islands, gatewaying legacy traffic onto it, then migrating cells to native TSN as equipment is refreshed. Time synchronization (802.1AS) is the first capability to establish because everything else depends on it. Wireless is introduced where cabling is most painful — mobile robots and rotating equipment — using a private-5G slice bridged into the established TSN domain. Throughout, centralised admission control (802.1Qcc) must be in place before the first guaranteed flow is admitted, because retrofitting admission control onto a network that already carries critical traffic is far harder than designing it in.

> IN PRACTICE
>
> The most common migration failure is treating TSN as a faster switch rather than a scheduled system. Determinism is configured, not purchased: it requires an accurate inventory of every critical flow, a computed gate schedule, and admission control that refuses flows it cannot guarantee. A TSN switch installed without that discipline behaves like a slightly better best-effort switch — and the tail returns.

Recommendations

1. **Establish a common time base first.** Deploy 802.1AS synchronization across the backbone before any scheduling; every other guarantee depends on it.
2. **Adopt the toolbox incrementally.** Add the time-aware shaper and preemption where cycle times demand it, and FRER only on flows that cannot tolerate loss — do not pay for guarantees an

application does not need.

3. **Design for the 3GPP TSN-bridge model** so that wired and wireless determinism form one domain, and so today's 5G deployment is forward-compatible with 6G.
4. **Keep the hardest safety loops wired.** Reserve the statistical wireless tail for flows that tolerate a percentile guarantee; give absolute-deadline loops a bounded wired TSN path.
5. **Put admission control in place before the first guaranteed flow,** and inventory every critical flow; an unmanaged schedule degrades back to best effort.

Conclusion

The deterministic industrial network is no longer a research aspiration; it is an engineering discipline with standardised tools. TSN retrofits bounded, reliable delivery onto ordinary Ethernet, and private 5G — evolving toward deterministic 6G — extends those guarantees to mobile and uncabled assets, the two stitched into a single timed domain by the 3GPP TSN-bridge model and anchored by edge compute. Our modeled benchmark shows the consequence in the only place it matters: the tail of the latency distribution, where the converged fabric holds worst-case latency in the sub-millisecond to low-millisecond range at 99.999% reliability, while best-effort convergence fails under exactly the load that real plants impose. The technology is ready; the work that remains is design — inventorying flows, computing schedules, partitioning budgets, and enforcing admission control. Operators who treat determinism as something configured and proven, rather than purchased, will converge their networks without surrendering the timing that automation depends on.

References

1. IEEE 802.1AS — Timing and Synchronization for Time-Sensitive Applications (generalized Precision Time Protocol).
2. IEEE 802.1Qbv — Enhancements for Scheduled Traffic (time-aware shaper); IEEE 802.1Qcc — Stream Reservation Protocol and configuration enhancements.
3. IEEE 802.1Qbu & IEEE 802.3br — Frame Preemption and Interspersing Express Traffic (MAC merge).
4. IEEE 802.1CB — Frame Replication and Elimination for Reliability (FRER).
5. 3GPP TS 23.501 — System architecture for the 5G System; 5G as a logical TSN bridge (DS-TT / NW-TT) and URLLC service requirements.
6. Root Digit Connected Systems Lab — Converged TSN / private-5G latency benchmark methodology (2026).

Determinism,
end to end.

Root Digit designs, integrates, and operates deterministic industrial networks — TSN backbones, private 5G/6G URLLC slices, and the edge compute that closes control loops on time. We inventory your critical flows, compute the schedules, partition the latency budget, and prove the worst case. You keep the cycle.

Talk to our network engineers →

rootdigit.com

General: info@rootdigit.com · **Consultation:** rootdigit.com/contact/consultation
Dubai International Financial Centre (DIFC) · Wyoming, USA