The Last Engineer Who Knew

Network Firefighting Is No Longer a Strategy

A White Paper on Telecom's Knowledge Crisis

Shawn Ennis
Founder & CEO, Rapax
shawn@citus.cloud

rapax.app | sales@rapax.app

I. The SONET Story

I started my career building SONET rings. If you know what that means without Googling it, this paper is for you.

When I came up in this industry, fiber was an emerging technology. Running a new fiber route was an engineering event — it required planning, specialized skills, and people who understood the physics of light over glass. Today you can buy a fiber patch cable at Best Buy for twelve dollars. That's how much time has passed.

But here's the thing nobody talks about at industry conferences: the networks those technologies built didn't disappear when the next wave arrived. They're still running. Somewhere in your network right now, there's a shelf, a card, or a ring segment that went end-of-life before some of your engineers graduated high school. And somewhere in your organization — maybe — there's one person who still remembers how it works. How the protection switching actually behaves. Which tributaries are mapped where. Why that one node was configured differently from every other node in the ring and what happens if you normalize it.

This paper is about what happens when that person leaves.

It is about the convergence of three forces — retirements, downsizing, and expert flight — that are draining institutional knowledge from telecom operations faster than it can be replaced. It is about the compounding cost of that loss in a world of heterogeneous, multi-generational networks where expertise cannot be hired off the street or trained in a semester. And it is about why knowledge management, powered by a new generation of AI capabilities, represents the most practical and urgent investment a telecom operator can make today.

II. The Three-Vector Talent Crisis

The telecom workforce is contracting. That statement alone is not new — industry observers have been flagging demographic pressure for years. What has changed is the speed, the simultaneity, and the compounding nature of the talent drain. Three vectors are converging at once, and each one accelerates the damage of the other two.

Vector 1: The Retirement Wave

A large portion of the current telecommunications workforce is nearing retirement age, creating a knowledge and skill gap as seasoned professionals exit the industry.[1] This is not a projection — it is happening now. In major economies, more than 10,000 workers reach retirement age every single day. In telecom specifically, the engineers who designed, built, and maintained the networks that carry today's traffic are reaching the end of their careers. They carry decades of accumulated operational knowledge — knowledge that lives in their heads, not in documentation systems.

The scale of replacement demand is staggering. Industry groups have projected that over 200,000 fiber technicians alone will be required in the United States over the next five years.[2] And that projection addresses only one layer of the workforce. Network architects, operations center engineers, capacity planners, and infrastructure specialists face similarly acute shortages.

Vector 2: Downsizing and Restructuring

At the same time that retirements are draining the top of the experience curve, cost-reduction programs are cutting into the middle. Post-5G deployment cycles, many operators have entered restructuring phases, reducing headcount to manage margins.[3] The arithmetic of downsizing is particularly cruel for institutional knowledge: the most expensive employees — the ones with the deepest expertise and the longest tenure — are typically the first targets for reduction-in-force programs. Organizations shed their most knowledgeable people to save their most visible line item.

The remaining teams absorb the work, but they cannot absorb the expertise. A junior engineer can be assigned the tickets that a departing senior engineer handled. That junior engineer cannot be assigned the senior engineer's thirty years of context about why the network is configured the way it is.

Vector 3: Expert Flight

[1]Capcon Networks, Addressing the Broadband Workforce Shortage in the U.S., 2025; ClearlyIP Global Telecom Industry Jobs & Workforce Outlook 2025
[2]ClearlyIP, Analysis of 2025 Telecommunications Job Market Trends; Fiber Broadband Association workforce projections
[3]Broadstaff, Wireless Labor Shortage Is Over: Telecom Hiring Trends 2026

The skills most valuable in modern telecom operations — cloud architecture, software-defined networking, data analytics, automation, and AI — are the same skills that every other industry is willing to pay a premium for. Telecom operators are not just competing with each other for talent; they are competing with big tech, cloud providers, financial services, and a rapidly expanding AI sector.[4]

Nearly 70% of organizations reported difficulty filling full-time positions in 2025, reflecting persistent labor shortages and widening skills gaps.[5] In telecom, this is compounded by a perception problem. Talented engineers increasingly view telecom as a legacy industry with limited advancement opportunities. The result: the best mid-career engineers — the ones operators are counting on to replace the retirees — are leaving for industries that offer faster career trajectories, higher compensation, and more cutting-edge work.

> ***"Nearly 70% of organizations reported difficulty filling full-time positions in 2025, reflecting persistent labor shortages and widening skills gaps that challenge competitiveness and productivity."***
>
> — SHRM, 2026 Top Five Workplace Issues

The net effect of these three vectors is not additive — it is multiplicative. Retirements remove the deepest knowledge. Downsizing removes the next tier of experienced replacements. Expert flight removes the ambitious mid-career engineers who might otherwise step up. What remains is a workforce that is simultaneously thinner, less experienced, and less equipped to manage the complexity of modern telecom networks.

[4] SDxCentral, Telecom Industry Must Address Its Talent Shortage, 2023

[5] SHRM, 2026 Top Five Workplace Issues Report

III. The Compounding Problem

If telecom networks were homogeneous — if every operator ran a single vendor, a single technology generation, and a single management platform — the talent crisis would be manageable. New engineers could be trained on one stack and deployed immediately. But that is not the world any operator lives in.

The HetNet Reality

Modern telecom operators run four to six generations of technology simultaneously. TDM, SONET, MPLS, GPON, XGS-PON, software-defined transport, 4G, 5G — each generation carries its own CLI syntax, alarm semantics, provisioning workflows, performance thresholds, and failure modes. Each vendor within each generation adds another layer of variation: a Calix AXOS OLT behaves differently from a Nokia ISAM. A Cisco IOS-XR router is a fundamentally different operational experience from a Cisco IOS-XE device, despite sharing a brand name.

The expertise required to operate these networks effectively spans tens to hundreds of vendor, device type, and technology generation combinations. No single engineer masters all of them. The engineers who come closest — the ones who have worked across multiple generations and can contextualize a new alarm based on experience with an older platform — are precisely the ones who are retiring, being downsized, or leaving.

The Firefighting Trap

When experts leave, remaining teams get pulled into reactive troubleshooting on the devices and technologies they understand least. Industry research shows that fault analysis and response accounts for 36% of error-induced outage time, with configuration changes contributing another 27%.[6] Both of these activities are deeply knowledge-dependent: the speed and accuracy of fault analysis depends on the engineer's familiarity with the specific device, its expected behavior, and its historical quirks.

This reactive posture creates a compounding trap. Engineers who spend their shifts firefighting legacy issues cannot execute on modernization projects — fiber overbuild, network consolidation, cloud migration, SDN deployment. The more time spent keeping legacy infrastructure alive, the less time available to build the networks of the future. And as modernization stalls, the legacy infrastructure ages further, requiring even more firefighting. The cycle accelerates.

The Human Error Escalation

The Uptime Institute, drawing on 25 years of longitudinal data, estimates that human error — directly or indirectly — contributes to between 66% and 80% of all downtime incidents.[7] But the critical detail is what kind of human error. Of organizations that experienced a major outage

[6] Network World, Human Error in Network Operations and How to Deal With It, 2025
[7] Uptime Institute, Annual Outage Analysis, 25-year longitudinal estimate

caused by human error, 85% traced the root cause to staff failing to follow procedures or to the procedures themselves being flawed, outdated, or nonexistent.[8]

That number is not stable. It is accelerating. The proportion of human-error outages caused by failure to follow procedures jumped ten percentage points year-over-year, rising from 48% in 2024 to 58% in 2025.[9] Eighty percent of operators surveyed said their most recent serious outage could have been prevented with better management, processes, or configuration.[10]

This is not a "people are careless" problem. This is a "the procedures don't exist, are outdated, or are unfindable" problem — which is to say, it is a knowledge management problem. And it is getting worse precisely because the people who carried those procedures in their heads are the ones leaving.

> ***"80% of operators believe their most recent serious outage could have been prevented with better management, processes, or configuration."***
>
> — Uptime Institute, 2025

The Financial Spiral

Industry research consistently places the cost of IT downtime at $5,600 per minute or more for enterprise-scale operations — exceeding $300,000 per hour for many organizations.[11] More than half of all major outages cost over $100,000, and nearly one in five exceed $1 million.[12]

For telecom operators — where margins are already under inflationary pressure and subscriber growth is difficult to come by — every preventable outage directly erodes the capital available for network investment and competitive differentiation. The financial logic is circular and destructive: more knowledge loss leads to more outages, more outages lead to more firefighting, more firefighting consumes the budget and engineering bandwidth that should be going to growth, and less growth makes it harder to attract and retain the remaining talent. The spiral feeds itself.

[8] Uptime Institute, Annual Outage Analysis Report, 2025

[9] Uptime Institute, Annual Outage Analysis Report, 2025. Year-over-year increase from 48% (2024) to 58% (2025)

[11] Gartner estimate: average IT downtime costs $5,600/minute. ITIC 2024 survey: 90% of midsize/large companies report 1-hour downtime exceeding $300,000

[12] Uptime Institute, Annual Outage Analysis 2024: 54% of major outages cost more than $100,000

IV. Why Now: The Knowledge Management Moment

The idea that telecom operators need better knowledge management is not new. What is new is that the technology to deliver it has fundamentally changed.

The KM Graveyard

Operators have tried knowledge management before. SharePoint sites. Confluence instances. Custom wikis. Vendor-specific documentation portals. Binder-based standard operating procedures. Every one of these tools failed for the same structural reason: they required people under operational pressure to search correctly, maintain content proactively, and organize information consistently. Under the stress of a 3AM outage, none of those behaviors happen.

Research shows that employees spend up to 30% of their workday searching for information, and nearly half of those searches fail to find what they need.[13] In a NOC context, that statistic is more than an efficiency metric — it is an operational risk factor. An engineer spending 30% of a shift hunting for a runbook is an engineer not resolving an outage. A search that fails at 3AM doesn't just waste time; it forces the engineer to improvise, escalate, or guess. That is how the 85% procedure-failure statistic gets generated.

Why LLMs Change the Equation

Large language models, combined with retrieval-augmented generation (RAG), represent a structural break from every previous generation of knowledge management technology. The distinction is not incremental — it is architectural:

Previous KM tools required users to construct precise keyword searches against organized document repositories. If the engineer didn't know the right terminology, the right file name, or the right folder structure, the search returned nothing useful.

RAG-powered knowledge systems understand intent, not just keywords. An engineer who types "the ONT on Elm Street keeps dropping" gets the correct troubleshooting procedure even if they don't know the exact alarm name, device model number, or which document contains the answer. The system retrieves relevant content from across the entire documentation corpus and synthesizes a contextually appropriate response.

This is not a theoretical capability. It is the same underlying technology that powers the large language models millions of people use daily — applied to an operator's own proprietary documentation and operational procedures. The documentation doesn't need to be perfectly organized, consistently formatted, or centrally maintained before the system can extract value from it. It can work with what exists today, and improve as the knowledge base grows.

Three Capabilities That Matter

[13] Enterprise knowledge management research via Stratechi, 2025

For telecom operators evaluating AI-powered knowledge management, three capabilities distinguish solutions that deliver operational value from solutions that create another unused tool:

1. Capture what exists. The system must ingest unstructured documentation — vendor manuals, runbooks, SOPs, MOPs, escalation procedures, historical ticket data, tribal knowledge documents — and make it queryable immediately, without requiring upfront curation or reorganization. If the system demands a six-month content migration project before it delivers value, it will fail for the same reasons its predecessors failed.

2. Identify what's missing. The system should analyze the relationship between the operator's device inventory and its documentation coverage. Which device types have zero associated troubleshooting procedures? Which alarm categories have no triage runbooks? Which procedures haven't been updated in two or more years? Which queries from engineers consistently return no results? These gap analyses transform knowledge management from a passive repository into an active investment planning tool.

3. Document as you go. Knowledge capture must be integrated into daily operational workflows. When an engineer resolves an issue that required tribal knowledge, the system should capture the resolution as a draft procedure that can be reviewed and approved by a shift lead. Every resolved incident becomes a potential knowledge asset. This is how the knowledge base grows organically with operational tempo rather than depending on periodic documentation sprints that never get prioritized.

The Analyst Perspective

Major analyst firms are converging on the same conclusion. Gartner predicts that AI agents will evolve to discover tacit knowledge, and that interactions with them will become the process itself — with hidden knowledge utilized by these agents leading to new value assets.[14] Gartner further projects that by 2029, 70% of enterprises will deploy agentic AI as part of IT infrastructure operations, up from less than 5% in 2025.[15] Forrester's 2026 predictions for technology infrastructure and operations identify autonomous AI preventing outages as a critical trend.[16]

The timing argument is straightforward. The senior engineers who carry your organization's institutional knowledge are still here today. Their knowledge can be captured, structured, and operationalized. In three to five years, many of them will not be. The cost of capturing knowledge now is a fraction of the cost of reconstructing it later — if reconstruction is even possible.

[14] Gartner, Top Predictions for IT Organizations and Users in 2026 and Beyond, October 2025
[15] Gartner, Predicts 2026: AI Agents Will Transform IT Infrastructure and Operations, December 2025
[16] Forrester, Predictions 2026: Technology Infrastructure and Operations, November 2025

> ***“By 2029, 70% of enterprises will deploy agentic AI as part of IT infrastructure operations, up from less than 5% in 2025.”***
>
> — Gartner, Predicts 2026: AI Agents Will Transform IT Infrastructure and Operations

V. The Implementation Roadmap

Value Day One. Compounding Value Over Time.

Knowledge management initiatives fail when they are framed as multi-year transformation programs that require extensive upfront investment before delivering results. The roadmap below is designed around a different principle: start extracting value from existing documentation immediately, and let the system's value compound organically as the knowledge base grows and integrates with operational workflows. The timelines below are baseline estimates; operators with dedicated project resources can compress each phase significantly.

Phase 1: Capture What You Have (Weeks 1–2)

Import existing documentation into the AI knowledge engine — vendor manuals, runbooks, SOPs, MOPs, escalation procedures, tribal knowledge documents, and ticket history exports. No upfront curation or reorganization is required; the system indexes what exists and makes it queryable via natural language immediately. Engineers can begin asking questions in plain language instead of searching file shares, SharePoint sites, or binder-based SOPs.

Success metric: Time-to-answer on the twenty most common troubleshooting queries, measured against the baseline of existing search tools.

Phase 2: Identify What's Missing (Weeks 2–4)

The system performs coverage analysis across the operator's device inventory against its documentation inventory. Which device types have zero associated procedures? Which alarm categories have no triage runbooks? Cross-reference with alert frequency data to prioritize gaps by operational impact, not alphabetical order. The output is a prioritized "knowledge debt" report that gives leadership visibility into where documentation investment will deliver the highest operational return.

Success metric: Number of critical coverage gaps identified and ranked by operational frequency.

Phase 3: Document as You Go (Weeks 4–8)

Integrate knowledge capture into shift workflows. When an engineer resolves an issue that required tribal knowledge — knowledge that existed only in someone's head — the system captures the resolution as a draft procedure. Shift leads review and approve. Every resolved incident becomes a potential knowledge asset. This is where the organic value curve begins: the system gets measurably smarter with every shift, and the rate of "no-result" queries begins to decline.

Success metric: New procedures created per week; knowledge base growth rate; declining no-result query rate.

Phase 4: Consolidate and Decommission (Weeks 6–10)

As the AI knowledge base matures, identify which legacy documentation systems can be retired. Consolidate to a single source of truth. Use staleness detection to flag procedures that reference decommissioned equipment, deprecated firmware versions, or retired vendor platforms. Reduce documentation tool sprawl and the maintenance burden associated with multiple parallel systems.

Success metric: Number of legacy documentation systems retired; reduction in documentation tool licensing and maintenance costs.

Phase 5: Operational Integration (Weeks 8–12)

Connect the knowledge system to operational workflows. Knowledge should flow to the engineer in the context of their work, not require them to leave their operational tools to go search a separate system. When an alert fires, relevant troubleshooting procedures should surface alongside the alert data. When a device is provisioned, relevant configuration guides should be available at the point of action. This is the transition from “searchable documentation” to “AI-assisted operations.”

Success metric: Measurable MTTR reduction on knowledge-dependent incident categories.

Phase 6: Continuous Improvement Loop (Ongoing)

The system tracks which queries produce high-confidence answers versus low-confidence or no-result answers. That analytics feed drives the next round of documentation investment. Leadership gains visibility into knowledge health as a KPI alongside network availability and MTTR. Knowledge gaps stop being invisible and become measurable, prioritizable, and trackable.

Success metric: Query confidence scores trending upward; zero-result rate trending downward; knowledge health dashboard adopted as an operational KPI.

Addressing the Objections

“We’ve tried knowledge management before.” Previous tools required perfect search behavior and proactive content maintenance from people under operational pressure. RAG-powered systems understand intent, not keywords. The engineer describes the problem; the system surfaces the procedure. No correct search syntax required.

“Can we trust AI-generated answers for critical procedures?” Human-in-the-loop by design. The system provides procedures with source references. Engineers validate and execute. This is augmentation, not automation. The AI retrieves and presents; the engineer decides and acts.

“What about our proprietary data?” On-premises and private cloud deployment models ensure that proprietary procedures, network configurations, and operational data remain within the operator’s environment. No documentation leaves the operator’s control.

VI. Wade and the Rapax Agent Ecosystem

The framework described in this paper is not theoretical. It is the operating design of Rapax, an AI-native network service assurance platform built by a 25-year telecom operator for telecom operators.

At the center of Rapax's knowledge architecture is **Wade** — a RAG-enhanced AI agent purpose-built for network operations knowledge management. Wade ingests vendor manuals, runbooks, SOPs, MOPs, and operational procedures. It surfaces answers through natural language queries enriched with live network context. It performs the coverage analysis, staleness detection, and gap identification described in the roadmap above.

Wade does not operate in isolation. It is part of a multi-agent architecture where knowledge flows to the point of operational need:

Oscar provides operational status awareness — morning reports, alert forensics, device performance summaries — drawing on Wade's knowledge base for contextual enrichment.

Nora manages device and service provisioning, querying Wade for configuration guidance and best practices during lifecycle operations.

Grace manages alert notifications and subscriptions, enriching every notification with relevant troubleshooting context from Wade so that engineers receive actionable intelligence, not just alarm data.

Bruce builds and manages network integrations, leveraging Wade's knowledge base for integration documentation and troubleshooting guidance.

The architecture is designed around a simple principle: engineers should not have to go to the knowledge base. The knowledge base should come to them, in the context of their work, at the moment they need it. That is the difference between a documentation repository and an operational knowledge system.

Start the Conversation

The knowledge crisis in your network is not coming. It is here. The question is whether you are going to address it while the people who hold that knowledge are still around to capture it — or after they are gone.

If anything in this paper described your operation, let's talk. Fifteen minutes. No prep needed. No slide deck. Just a conversation between two people who have been in the trenches.

Book a call: cal.com/shawn-ennis

Email: sales@rapax.app

Learn more: rapax.app

About the Author

Shawn Ennis is the Founder and CEO of Rapax, an AI-native network service assurance platform built for the operators who keep networks running. He holds 12 patents in network management and service assurance, previously founded Assure1 (acquired by Oracle in 2021), and has spent 25 years in telecom operations — from building SONET rings to building the AI agents that will keep the next generation of engineers from having to learn what a SONET ring was.

He co-hosts the *Transformation Leaders Podcast* and writes about the intersection of network operations, institutional knowledge, and AI at rapax.app.