

2026 Field Service Benchmark Report



Table of Contents

Executive Summary	03
Dataset Overview and Definitions	04
Key Findings	05
First-Time-Fix	06
Customer Experience Curve	07
First-Time-Fix By Industry	08
Mean Time to Resolution	09
Mean Time Between Visits	10
Cost of Failed Visits	11
Remote Resolution Potential	12
Preventative Maintenance vs Reactive Work	13
Preventative Maintenance Effectiveness	14
Skills Gap	15
Workforce Retention	16
Cost Per Success	17
Multi-Visit Cases	18
Cost Savings	19
Setting the Stage for 2026	20



Executive Summary

INTRODUCTION

Field service leaders are no longer winning by adding more tools to their stack, they are winning by unifying the tools and data they already have into a single intelligence layer.

The 2026 Field Service Benchmark shows a widening gap between organizations that have consolidated critical service information and those still operating in silos. Leading organizations resolve issues faster, in fewer visits, and at significantly lower cost. Others remain stuck in cycles of repeat work, longer downtime, and escalating spend.

Across hundreds of organizations and millions of service events, one pattern is clear: when service data, knowledge, and systems converge into one source of truth, AI can finally act on the full context of every customer and asset, helping teams scale top-performing outcomes beyond a few individuals.

This year's report is not just a snapshot of performance; it is a blueprint for how consolidating your organization's service data and systems into a single agentic platform—rather than adding another AI point solution—unlocks double-digit efficiency gains and a step-change in customer experience.



Dataset Overview



Aquant gathered and analyzed real, anonymized data from:

Nearly **30 million** service events spanning **7 million** assets

Nearly **\$8.3 billion** in service costs

Over **600k** technicians

161 service organizations, including service divisions within OEMs and third-party service organizations across manufacturing, medical devices, commercial printing, industrial machinery, food service equipment, and more

Nearly **3 years** of data

Definitions

Case	A customer-reported issue or service request tracked from open to close.
Service event	A logged service interaction tied to a case (e.g., remote troubleshooting session, onsite visit, follow-up action).
Onsite visit (truck roll)	A dispatched field visit to a customer site.
Successful resolution	A case closed as resolved.
First-time-fix (FTF)	A case successfully resolved on the first onsite visit.
Time to resolution (MTTR)	Elapsed time from case open/creation to case close/resolved, measured in days.
Failed visit	An onsite visit that does not result in resolution and requires follow-up work.
Remote resolution potential	Estimated share of onsite visits that could have been resolved remotely.
Mean time between visits (MTBV)	Average days between interactions for the same asset.
Mean time between field visits (MTBFV)	Average days between onsite visits for the same asset.
Cost per success (CPS)	Total service cost divided by number of successful resolutions.
Performance tiers	Organizations grouped into Top 20% of companies / Industry benchmark / Bottom 20% of companies.

Key Findings



The top tier is pulling away.

Top performer organizations run service on unified context - through an agentic platform, so results are consistent and repeatable. Bottom performers are still fighting siloed information, so problems drag on and repeat.



Failed visits are the biggest fixable cost problem.

Failed visits account for an industry benchmark of **25% of total service cost**—44% for bottom performers vs. 14% for top performers. One miss often triggers more trips and longer downtime.



Companies could reduce up to 26% in service costs.

If more teams could operate like top performers more often, the model shows **double-digit cost reductions—up to 26% in some segments**. The path to scaling that performance is consolidating data into a shared intelligence layer so AI can drive end-to-end execution.



Many field trips didn't need to happen.


Remote resolution potential is **2.2% of events for top performers** vs. **19.5% for bottom performers** (with cost impact of 1.4% vs. 18.1%). The difference is smarter troubleshooting and better access to asset and case history before dispatch.

The Foundation of Great Service

First-time-fix (FTF) is the clearest line between smooth customer experiences and frustrating repeat visits. The industry benchmark sits at 77%, but top performers reach 88%, while the bottom group lags at 60%—a nearly 30-point gap in “fixed right the first time.”

This gap translates directly into customer effort: bottom performing organizations require far more follow-up visits and longer timelines to reach the same outcome, turning what should be a single, seamless experience into multiple disruptions.

First Time Fix Rate by Performance Tier



Top-performing organizations resolve issues on the first visit far more often, with first-time-fix rates nearly 30 percentage points higher than bottom performers, turning repeat work into rare exceptions rather than the norm.




When Cases Stay Open, FTF Falls

FTF doesn't just vary by performance tier, it degrades over time, and the drop-off is steepest for bottom performers. On Day 1, even lagging organizations are relatively close to the benchmark. But as cases remain open longer, first-time-fix declines across the board and the gap between leaders and laggards widens dramatically.

This curve is the customer experience story underneath the operational KPIs: the longer a case stays open, the more likely it is to require repeat effort (more visits, more scheduling friction, more downtime, and more customer frustration). By Day 30, the bottom tier's first-time-fix has collapsed to nearly half of cases—while top performers remain comparatively resilient.

Customer Experience Curve — First Time Fix Rate vs Days Open



First-time-fix declines the longer work stays open—but top performers hold steady far longer than bottom performers. The gap between top and bottom widens from 5 points on Day 1 (99% vs 94%) to 33 points by Day 30 (87% vs 54%), showing how delays compound into repeat work and a degraded customer experience.




FTF Varies By Industry

Not all industries operate under the same conditions. Asset complexity, parts ecosystems, compliance requirements, and field constraints all influence how achievable a high FTF rate is. Looking at FTF by industry provides a useful lens for goal-setting, helping teams benchmark performance against peers facing similar operating realities.

The benchmark shows meaningful variation across verticals. Some segments sustain first-time-fix above 80% at Day 30, while others sit closer to 70%, reinforcing that “good” performance should be calibrated by industry context, not just internal ambition.

First-Time Fix Rate by Industry



FTF varies by industry, reflecting differences in complexity and service constraints. Some segments sustain 80%+, while others cluster closer to 70%.






How Fast Is “Fast Enough”?

Customers feel time more than any internal KPI. The industry benchmark for MTTR is 4.6 days, but top performers close cases in about 2.5 days, while bottom performers stretch to just over 10 days—turning what could be a quick fix into a multi-week ordeal.

This time penalty compounds with each missed first-time-fix: every failed visit adds roughly 13.7 days to resolution, making repeat work not just an operational burden but a major driver of customer dissatisfaction.

Mean Time to Resolution (MTTR) by Performance Tier (Days)



Leading organizations close cases in a fraction of the time, resolving issues in just a few days while bottom performers stretch into double-digit days—turning quick fixes into drawn-out customer experiences.

How Often Are You Returning?

Leading companies don’t just fix problems faster—they reduce repeat visits by resolving issues more completely the first time.

The industry benchmark for MTBV is 83 days, but top performing companies push that to 133 days, and bottom performers sit at

just 48 days. The same pattern holds for onsite visits: top performers average 136 days between field visits, while bottom performers are back in just 51 days, reinforcing a cycle of repeat truck rolls that strain capacity and frustrate customers.

Mean Time Between Visits (Days)



Leading companies average 133 days between service visits and 136 days between field visits because they resolve issues more completely—while bottom performers are back in just 48 days (51 days field-only), reinforcing a costly cycle of repeat truck rolls.






The Cost of Failed Field Visits

Failed field visits are one of the biggest—and most fixable—sources of waste in service operations. They represent about 25% of total service cost on median, but only 14% for top performers and a staggering 44% for bottom performers.


Each failed visit typically triggers about 1.5 additional visits and adds nearly two extra weeks to the resolution timeline, compounding labor, scheduling, and customer-impact costs with every miss.

Service Budget Lost (%) to Failed Field Visits




Failed field visits quietly drain service budgets—about 25% of total service cost across the industry and nearly half for bottom performers—making repeat truck rolls one of service’s biggest, most fixable cost leaks.

The Remote Resolution Opportunity



Remote resolution opportunity (% of visits)



Potential service cost savings from remote resolution

Remote resolution potential measures the share of service events where a field visit could likely have been avoided and the issue resolved remotely.

The benchmark shows a sharp separation between leaders and laggards. Top performers have remote resolution potential of just 2.2% of events, suggesting that most field visits are truly necessary. At the median, 10.0% of events have avoidance potential. Bottom performers climb to 19.5%—nearly 1 in 5 field visits representing a missed opportunity for remote resolution.

This gap is even more pronounced in cost impact. Top performers show only 1.4% cost avoidance potential, while the median sits at 6.7%.


Bottom performers reach 18.1%, indicating that avoidable dispatch isn't just a volume problem—it's a major cost leak, often driven by repeatable failure modes such as incomplete triage, missing asset history, unclear symptoms, or lack of remote troubleshooting pathways.

The takeaway is straightforward: reducing avoidable dispatch isn't about pressuring teams to "go remote." It's about building the conditions that make remote resolution reliable—complete service context, strong diagnostic workflows, and the ability to guide agents and technicians with the right information at the right moment.


Remote resolution potential estimates the share of field visits that likely could have been resolved remotely. For top performers, that figure is just 2.2% of events (1.4% of cost)—suggesting most field visits are truly necessary. The benchmark is 10% (6.7% of cost), while bottom performers reach 19.5% of events (18.1% of cost).

A Snapshot of Maturity


Preventative Maintenance vs Reactive Work by Performance Tier



Top 20% of companies
● Reactive 45% ● Preventative 55%



Industry benchmark
● Reactive 84% ● Preventative 16%



Bottom 20% of companies
● Reactive 94% ● Preventative 6%

Top-performing organizations have nearly reached a 50/50 split between preventative maintenance and reactive work.

The bottom performers are still reacting to nearly every issue—missing the opportunity to prevent problems before they occur.

Despite growing awareness around the value of preventative maintenance, the vast majority of service organizations still operate in a reactive mode. At the benchmark, 84% of all visits are reactive, triggered only after something goes wrong. This means only 16% of visits are preventative—designed to prevent issues before they occur. But the contrast between top and bottom performers reveals the potential for transformation.

For top performers, only 55% of visits are reactive, meaning nearly half of their service activity is preventative by design. These organizations are not just fixing problems—they are preventing them, extending asset life and reducing emergency costs.


Bottom performers, on the other hand, remain locked in a break-fix cycle. A staggering 94% of their visits are reactive, highlighting a missed opportunity to reduce unplanned downtime, smooth capacity planning, and improve overall service efficiency.

When PM Works, Everything Gets Easier


Preventative maintenance (PM) is one of the clearest indicators of service maturity—because effective PM reduces surprise failures, smooths scheduling, and prevents “break-fix firefighting” from consuming capacity.

The benchmark shows a stark difference in PM effectiveness (days) across tiers. Top performers sustain substantially longer PM effectiveness windows, while bottom performers see far shorter windows—often signaling that PM programs are inconsistent, mis-targeted, or not translated into lasting reliability improvements.

PM Effectiveness (Days) by Performance Tier



Top performers achieve nearly 3x longer PM effectiveness windows than bottom performers (162 vs. 62 days), reflecting stronger preventative programs that reduce reactive work, improve uptime, and protect capacity, while bottom performers remain stuck in short-cycle recurrence.



The Skills Gap by First-Time Fix

First-time fix often hinges on a small group of “hero” technicians. When only a few can consistently resolve issues on the first visit, customer outcomes vary based on who gets dispatched.

This chart shows the FTF advantage of hero technicians vs. everyone else, measured in percentage points. At the benchmark, the gap is 6.5 pp (≈ 6–7 more first-time fixes per 100 jobs).

A smaller gap means more technicians can fix issues on the first visit. A larger gap signals expertise is concentrated in a few people—driving more repeat visits, longer resolution times, and higher cost when they’re not on the job.

Difference in First-Time Fix Rate Between Hero and Non-Hero Technicians




Leading companies keep the FTF gap to ~3 percentage points, reflecting broadly distributed expertise. Lower-performing companies show a 10-point gap, where critical knowledge remains concentrated among a few technicians.

Stability is a Service Advantage

Service performance doesn't scale when experience churns out of the organization. Workforce retention is a quiet but powerful driver of consistency, because repeatable outcomes depend on technicians who know the equipment, the customer environments, and the "tribal knowledge" that never makes it into a manual.

The benchmark shows a meaningful separation by performance tier: top performers retain their workforce at substantially higher rates, while bottom performers lose a much larger share of employees—creating constant re-training cycles, longer ramp times, and higher risk of failed visits.

Workforce Retention by Performance Tier



Leading companies retain far more of their workforce (87% vs. 66% for bottom performers), protecting institutional knowledge and enabling consistent service execution, while bottom performers face ongoing churn that fuels variability, repeat work, and higher cost.




Poor Performers Drive Up Costs Big Time

Hero technicians are defined by a blend of quality, speed, and customer outcomes, not just cost. They still achieve successful outcomes at a dramatically lower cost than their peers: hero technicians run at a 10% cost per success rate versus 59% for the bottom 20%, and well below the 26% industry benchmark.

This means a successful outcome from a hero technician costs almost six times less than one from a bottom-tier technician, proving that better service does not require higher spend. Investing in developing more heroes meaningfully lowers cost per success while also improving the service experience.

Cost per Success by Hero vs. Non-Hero




Hero technicians deliver successful outcomes at about one-sixth the cost of bottom performers—10% cost per success versus 59% for the bottom tier—highlighting how wide the cost gap becomes when service quality lags.

When Things Get Complicated

Not every issue can be resolved in a single visit—but how organizations handle multi-visit cases separates leaders from laggards. Top performers keep multi-visit cases leaner, with fewer events per case and far shorter time to resolution than bottom performers.

For multi-visit incidents, the metric ranges from around 12 days for top performers to 28.5 days for the bottom group, showing that complexity amplifies differences in process discipline, diagnostic quality, and technician enablement.

Multi-Visit Cases — Time to Resolve (Days)



For top performers, most issues are resolved with fewer touchpoints, while bottom performers require far more visits per case, turning what should be one clean resolution into a series of callbacks and revisits

When cases require multiple visits, leaders still keep control—resolving multi-visit issues in under 2 weeks, while bottom performers see those same cases stretch toward a month, increasing customer frustration and cost.






What If Everyone Performed Like Top Performers?

The most powerful insight from the 2026 benchmark data is the modeled upside. If every technician and team operated at top-performing levels, companies could unlock double-digit percentage reductions in service cost—typically in the low teens, with potential savings as high as 26% for some companies.

These savings are not theoretical; they reflect real-world performance already achieved by today’s best companies and technicians. The opportunity now is to scale knowledge and expertise across the workforce so every technician can perform like the top 20%.

Cost Savings (%) if Everyone Performed Like Top 20% of Technicians



If every technician operated at top-performing levels, companies could unlock double-digit percentage reductions in service costs—typically around the low teens—with bottom-performing companies seeing potential savings of up to 26%.

From Assistive AI to Operational AI

2026 is the year service leaders stop adding more tools and start unifying what they already have.

While the core challenges remain familiar, the path forward is no longer about expanding the tech stack—it’s about consolidating it, so performance becomes repeatable and scalable.

Asset-centric organizations are sitting on a wealth of fragmented data.


Work orders, asset histories, parts, knowledge articles, contracts, IoT signals, technician expertise, work summaries, tribal knowledge and customer feedback scattered across multiple systems. When that context remains siloed, point AI tools can only deliver narrow, assistive value inside a single workflow—like resolution in isolation—while missing the critical inputs (and downstream outputs) that determine whether the fix actually sticks.

The next step is an agentic approach.

Think: a village of specialized agents powered by one agentic platform and grounded in a shared intelligence layer. Instead of deploying a single-purpose “resolution tool,” service leaders need a platform where agents collaborate across the entire company—supporting troubleshooting and resolution, capturing and summarizing calls, learning from work performed in the field, updating knowledge based on outcomes, and feeding insights back into planning, parts, training, and customer experience.

Unify the intelligence layer. Then operationalize it with agents.

2026 won’t reward teams with the most AI tools. It will reward teams with the most connected context—so AI can drive consistent execution, not isolated assistance.






READY FOR ANYTHING

Aquant is the Agentic AI platform built for organizations that build, sell, and service complex equipment. Unlike generic AI tools, Aquant combines deep domain expertise with intelligent agents that understand the language of service. The platform captures both documented knowledge and institutional know-how, delivering expert-level guidance to every role—from technicians and contact center agents to leaders across the business.

aquant.ai ✨



"I need help with callibration"



What's the part number?





✨ Let's get started...

Revenue



80
60
40
20

Feb Apr Jun Aug Oct Dec

