// 2026 WHITE PAPER

The *Intelligent* MSP: Unlocking AI-Driven Growth Through Operational Excellence in Autotask

AI will reward the MSPs that got their data right first. How to build the operational intelligence layer that turns Autotask into an AI-ready, margin-driving foundation, before the competition does.

$401B
Global managed services market in 2025, heading to $847B by 2033

80%+
of AI projects fail, most due to poor data quality (RAND)

95%
of IT leaders say integration hampers AI implementation (MuleSoft)

● Trusted by 150+ MSPs Worldwide

Published by **Proxuma.io** · In partnership with **Dxfferent**
Edition **2026**

CONTENTS

00 Executive Summary

01 The MSP Landscape in 2026

02 Where the PSA Ends and Intelligence Begins

03 The Talent & Margin Crisis

04 The AI Adoption Paradox

05 The KPI Benchmarking Framework

06 The Proxuma Solution

07 ROI Business Case

08 The Future of MSP Operations

09 Conclusion

10 Get Started with Proxuma

· Sources & References

EXECUTIVE SUMMARY

00

The Case for Operational Intelligence

The MSP industry is at an inflection point. Growth is abundant. The operational tools to capture it are not.

$401B	$37B	95%	80%+
Global managed services market in 2025, heading to $847B by 2033	Enterprise AI investment in 2025, tripled in one year (Menlo Ventures)	of IT leaders say integration hampers AI implementation (MuleSoft)	of AI projects fail, most due to poor data quality (RAND)

The global managed services market reached **$401.15 billion in 2025 and is on a 9.9% annual trajectory toward $847.41 billion by 2033** (Grand View Research), and enterprise AI investment tripled in a single year to $37 billion (Menlo Ventures). Agentic AI, automation, and autonomous service delivery now dominate every keynote and vendor roadmap in the channel, and the race is on to buy it.

Here is the uncomfortable part. Most of that AI will not pay off, and the reason has nothing to do with the AI. Independent research finds more than **80% of AI projects fail** (RAND), and the cause is rarely the model. It is the data underneath it. An agent, an assistant, or a copilot can only act on the records it is given, and in most MSPs those records are scattered across Autotask, the RMM, billing, and a dozen other systems, entered inconsistently and never reconciled into one picture. Point modern AI at that and it returns the oldest result in computing: garbage in, garbage out. AI thinks and automation acts, but data decides.

Autotask itself is one of the most widely deployed PSA platforms in the channel for good reason: it captures the work an MSP does, ticket by ticket. What it does not do is turn that captured work into a picture you can act on. The data an MSP needs to run a better business already exists inside its tools, but it lives in spreadsheets alongside the PSA, in tribal knowledge for shift coverage, and in board reports assembled by hand hours before the client call. Cooper Copilot, Autotask's native AI assistant, can only be as good as that data, which is exactly the layer most MSPs have never had the time to keep clean and connected.

This is not a technology problem. It is an architecture problem. And it has a solution.

Proxuma is an operational intelligence platform built by Dxfferent, a team of MSP and Autotask specialists who advise and support MSPs every day, strategically and technically. It was not designed in a vacuum. Every capability traces back to a real question from a real MSP Dxfferent works with, which is the whole difference between generic SaaS and software built for MSPs, with MSPs, by people who do the work themselves. It is a platform, not a single-PSA tool: Autotask is its first and deepest integration, chosen because it is the most widely deployed PSA in the channel, and the architecture is designed to extend. It delivers three interlocking capabilities that transform how MSPs plan work, manage people, and understand their business:

> *"The data you need already exists. It is just scattered across a dozen systems, in a form you cannot act on. Operational intelligence is what connects it into one picture you can run the business on."*
>
> Proxuma.io

- **Proxuma Planning:** a constraint-optimization engine that builds the day's schedule across the entire queue, weighing skills, availability, SLA deadlines, capacity, and travel, then hands the dispatcher a drag-and-drop plan and real capacity visibility weeks and months ahead. One system of record for who does what and when, instead of spreadsheets and tribal knowledge.
- **Proxuma Shifts:** coverage management for NOC, helpdesk, and on-call that scores the real workload behind each shift, so load is balanced by data rather than by who is easiest to ask. It catches conflicts and rest-period gaps before the schedule goes live, and keeps engineers, on-call staff, and external contacts on one calendar through bi-directional Outlook sync.
- **Proxuma Intelligence:** it normalizes Autotask together with 12+ of the systems MSPs actually run day to day, the RMM, Microsoft 365, security, billing, and documentation, into one model, then ships 500+ pre-built measures and dashboards on top, through Power BI or hosted on your own stack. It is the AI-ready foundation for *chat with your data*: ask any operational question in plain language and get a grounded answer across every connected system, available today via an open MCP interface and deepening as Microsoft's Power BI Copilot and Proxuma's own AI layer mature.

With Proxuma Planning, scheduling a ticket by hand that used to take a dispatcher around five minutes now takes roughly thirty seconds: the system proposes an optimized plan that the dispatcher reviews and adjusts by drag-and-drop, rather than building the schedule from scratch. Across a full queue, that turns two to three hours of daily scheduling into well under an hour. MSPs using Proxuma Intelligence answer questions about revenue by contract type, technician utilization, SLA performance, and client profitability in seconds, instead of the days it takes to assemble a board report by hand.

The argument of this paper is simple. In the AI era, the MSPs that win will not be the ones that bought the most AI. AI and automation are leverage; the core is still the people, a scarce, expensive team whose time is the most valuable asset the business has. The winners will be the ones that build the data layer to deploy that team intelligently, the right person on the right work, honest utilization, capacity seen weeks ahead, with smart automation carrying the load underneath so human hours go where they matter. The pages that follow map where the PSA workflow ends and the intelligence layer begins, quantify what the gap costs, and show how Proxuma closes it. Operational intelligence does not replace the human judgment at the center of a strong MSP. It is what lets that judgment scale.

SECTION 01

The MSP Landscape in 2026: Abundance, Complexity, and the Operational Gap

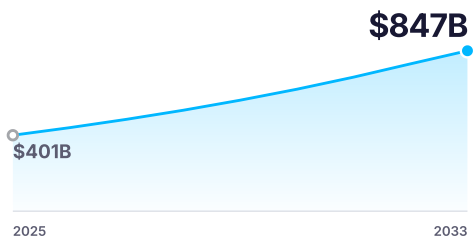
An industry at a generational inflection point, where growth is real but the ability to capture it is unevenly distributed.

A Market at Scale

The global managed services market reached **$401.15 billion in 2025**, according to Grand View Research, on a 9.9% compound annual growth rate (2026–2033) that carries it to **$847.41 billion by 2033**. In the United States, the largest single market, it stood at $88.13 billion in 2024 and is projected to reach $183.47 billion by 2030 at a 13.3% CAGR. These are structural shifts, not speculative projections, driven by enterprises offloading IT complexity, the mandatory evolution of cybersecurity services, the proliferation of cloud infrastructure, and the acceleration of digital transformation across every industry vertical.

Global managed services market

The market more than doubles, to $847 billion by 2033.



Source: Grand View Research.

For MSPs, this represents a generational opportunity. The addressable market is growing faster than the industry's capacity to serve it. Organizations that would once have maintained in-house IT departments are outsourcing at an accelerating rate. The open question is whether a given MSP has the operational infrastructure to capture it.

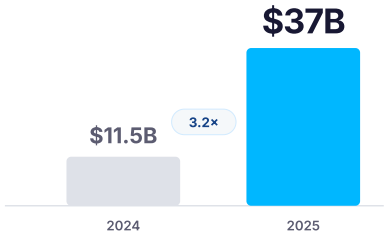
$401.15B Global MSP market, 2025	**$847.41B** Projected global market by 2033
9.9% CAGR, 2026–2033	**48%** of MSPs rank AI and automation as their #1 client need for 2026 (Kaseya)

The AI Wave Meets the Data Wall

Enterprise AI investment reached **$37 billion in 2025, roughly triple the $11.5 billion the year before** (Menlo Ventures), and it is flowing straight through the MSP channel. Organizations are turning to managed service providers to deploy, integrate, secure, and operate AI. But it cuts both ways: MSPs that cannot demonstrate AI fluency will be shut out of the highest-value contracts.

Enterprise AI investment

Spending roughly tripled in a single year, to $37 billion.



Source: Menlo Ventures, 2025 State of Generative AI in the Enterprise.

But the same wave is exposing a hard truth inside MSPs' own operations: adoption is racing ahead of results. More than 80% of AI projects fail (RAND), most of them not on the model but on the data, and Gartner expects organizations to abandon 60% of AI projects that lack AI-ready data through 2026. The platform vendors themselves now concede the point, describing clean, connected data as the prerequisite for any AI to work at all.

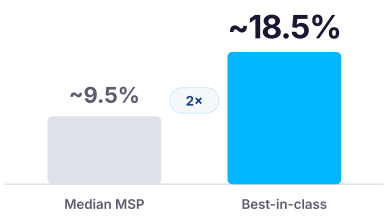
The irony is sharp, and it is splitting the field. The AI opportunity MSPs are asked to deliver for clients demands exactly the data foundation most have never built for themselves. The MSPs that have it compound their advantage, better insight, tighter operations, higher margin, more room to apply AI, while the majority stuck on fragmented data fall further behind, precisely as private-equity-backed consolidation accelerates across the sector. In the AI era, an MSP's data infrastructure is becoming its competitive destiny.

Margin Compression and the Best-in-Class Gap

The gap between an average MSP and a best-in-class one is wide, and it is measurable. Per the Service Leadership Index, the long-running benchmark of MSP financial performance, the best-run operators clear roughly two to three times the EBITDA margin of the median MSP, with the top tier sitting near 18 to 19% and the middle of the market around half that. The difference is not mainly pricing, sales volume, or client mix. The leaders do not charge more; they know, in real time, which engagements are profitable, which technicians are underutilized, which contracts are bleeding margin, and which agreements need repricing.

Service EBITDA margin

Best-in-class MSPs run about double the margin of the median.



Source: Service Leadership Index, the long-running MSP financial-performance benchmark.

An MSP without the measurement infrastructure to answer those questions runs on intuition, and on a thin-margin business intuition is expensive. A single contract drifting two points below target, unseen for a year, can quietly erase the profit of three healthy ones, and the operator never knows which one to fix.

Consolidation and the Data Readiness Premium

The MSP industry is in the middle of a significant M&A consolidation wave. Private equity-backed platforms, national consolidators, and strategic acquirers are aggregating MSP businesses at scale: Canalys and Omdia tracked 169 publicly-announced MSP deals in 2025, with private equity involved in roughly 70% of them. In that environment, clean, auditable, well-structured operational data has become a direct valuation driver, and the evidence is no longer anecdotal. EY's 2025 study of private-equity exit readiness found that **72% of firms name weak data and KPI reporting as their single biggest finance issue at exit**, 65% cannot reflect their value-creation story accurately in reported EBITDA, and 41% lack the data granularity to substantiate their equity story. EY's conclusion is blunt: assets that are not ready for the intensity of buyer due diligence see valuations come under pressure and timelines extend.

The numbers put a price on it. Private MSPs change hands at a median of roughly **8.9x EV/EBITDA** (Aventis Advisors, across 120 disclosed deals), with platform-scale operators reaching the low teens and sub-scale tuck-ins closer to 4 to 7x. What moves an MSP up that range is not size alone; it is the durability and provability of its margin, contract by contract. An MSP that can show real-time visibility into utilization, contract profitability, SLA performance, and technician productivity negotiates from a different position than one that needs a 90-day data extraction before a diligence team can even begin, because in a quality-of-earnings review, every point of EBITDA a buyer cannot verify is a point they negotiate off the price. The operational maturity Proxuma enables is, in that sense, a strategic financial asset, not just an operational convenience, and EY's advice is to build that readiness 12 to 24 months before a sale, long before most owners think they need it.

> **KEY INSIGHT**
>
> Growth is abundant; the ability to capture it is not evenly shared, and the advantage compounds. An MSP with clean, connected data prices sharper, runs leaner, and frees the capacity to grow, which produces more data and widens the lead. An MSP managing on fragmented data falls further behind each quarter. That gap does not close on its own; it widens, quietly, until a buyer, a competitor, or a churning client makes it impossible to ignore.

SECTION 02

Where the PSA Ends and Intelligence Begins

Autotask runs the operation. The layer above it, where capacity, margin, and AI-readiness are won, is a different job entirely.

Autotask as Infrastructure

Autotask PSA, built by Datto and now part of Kaseya, is among the most widely deployed Professional Services Automation platforms in the channel, and for good reason. It runs the daily workflow an MSP depends on: ticketing, time entry, contracts, billing, projects, and client communication. For most MSPs it is the operational core, the system the whole business is organized around.

That centrality is exactly why what sits above it matters. A PSA is built to run the workflow, not to model it, measure it, or see across it. Four jobs fall to the layer above, and they are the four that decide whether an MSP is merely busy or genuinely profitable.

One Connected View

An MSP's most demanding client and its most profitable client are rarely the same one. But tickets, time, and billing sit in separate places in the PSA, so telling them apart is a manual project, not a glance. The layer above joins them into one profitability picture, by client, by contract, by engineer, that you can actually read.

Analyst-Grade Reporting

The PSA's report builder was made for administrators, not board meetings. The moment a question needs a trend, a forecast, or margin cut by service line, it becomes an Excel export and a late night. Dimensional modeling and calculated measures, the analytics layer's job, turn that same question into a click.

Revenue Captured, Not Leaked

Complex contracts leak: an hour logged late, work done out of scope, a renewal never repriced. Each leak hides in the PSA because catching it means comparing delivered work to billed work, and those live in different systems. The layer above surfaces every unbilled hour as a line item to recover, not a vague shortfall discovered at quarter-end.

Forward Capacity

The PSA can tell you what the team did. It cannot tell you whether they can take the next project without breaking the SLAs on the last one. Capacity weeks ahead, by engineer and by project, is the difference between selling with confidence and over-committing blind, and it lives only in the layer above.

AI Is Only as Good as the Data Underneath

Autotask's native AI assistant, Cooper Copilot, is a genuinely capable tool, and Kaseya is investing heavily in agentic AI across its ecosystem. Within that ecosystem it does real work: smart ticket summaries, a writing assistant, resolution summaries, and triage that categorizes and associates incoming tickets. The same is true of every platform-native AI in the channel. Each is strong at reasoning over the data inside its own world.

But an MSP's world is bigger than any single vendor's. Beyond the Autotask and broader Kaseya ecosystem sits the rest of the stack the business actually runs on: a non-Kaseya RMM, Microsoft 365 and Google tenants, security platforms, documentation and billing tools. A platform-native assistant reasons over its own ecosystem's data, and the systems outside it never come through. There is no open, normalized data layer, a data lake the MSP owns and can query across every system at once, that brings the full picture together. The same holds for every PSA: the unified data each platform offers is unified only within its own walls.

This is the layer Proxuma adds, and it is complementary rather than competitive. By normalizing Autotask together with 12+ connected systems into one model the MSP owns, hosts, and can extend, Proxuma builds the cross-vendor picture no single platform's AI is positioned to see. Where the platform copilots make each system smarter, Proxuma makes the whole operation legible.

> *"Every platform's AI is sharp inside its own four walls. An MSP runs across a dozen walls at once, and the cross-system picture is the one no single vendor can give you."*

Running an MSP, and Running It on Data

There are two ways to run an MSP, and most reach the first without ever reaching the second. The first is running the work: tickets get closed, shifts get covered, invoices go out. A PSA runs this well, and the leading platforms run it very well. The second is running the business on data: knowing which clients actually make money, where capacity will break two weeks from now, which contract has quietly slipped below margin, and acting while it still matters. That second mode is a different discipline, and it does not live inside the PSA. It lives in a layer above it, and that is true of every PSA in the channel, not any one in particular.

The gap between the two is not closed by working harder in the tools you already run. Two MSPs can carry the same tickets, the same contracts, and the same stack, and still finish the year far apart, because one runs on memory and reaction while the other runs on data and anticipation. The edge is not effort. It is whether the operational data the business already generates, scattered across the PSA and a dozen other systems, becomes something the business can actually steer on. That is the layer Proxuma builds: for scheduling, for shift coverage, and for cross-system intelligence, sitting alongside the PSA and turning data the MSP already has into decisions it can make in time.

WHERE THE MARGIN HIDES

Two MSPs serve the same demanding client. The first sees a healthy monthly invoice and assumes the account is fine. The second sees that this client is being delivered at roughly $60 an engineer-hour while the rest of the portfolio earns $145, a contract quietly bleeding margin, and reprices it before the renewal. The data was sitting in both businesses. Only one of them could read it. That is where the margin hides: not in the work you are not doing, but in the data you already generate and cannot yet see.

SECTION 03

The Talent & Margin Crisis: Why Operational Chaos Is Not a People Problem

The forces compressing MSP margins are structural. So is the solution.

The Talent Shortage Is Real and Worsening

Talent is the binding constraint. In JumpCloud's SME IT Trends Report, **49% of IT teams say they lack the resources and staffing to secure their organization** against cyber threats, and security ranks as their single biggest operational challenge. The global cybersecurity workforce gap stands at roughly 4.8 million professionals (ISC2 2024 Cybersecurity Workforce Study). Its 2025 study found the constraint shifting from headcount to skills: the harder problem is no longer only hiring more people, it is getting more out of the skilled engineers a team already has. MSPs are not competing only with each other for engineering talent, they are competing with hyperscalers, enterprise IT organizations, and a venture-backed software industry that can outbid managed service providers on compensation at every seniority level.

The implications are direct: the engineers MSPs do have are too valuable to waste on scheduling logistics, manual timesheet reconstruction, or searching for the right ticket to assign. Every hour an engineer spends navigating operational chaos is an hour not spent resolving client issues, building skills, or contributing to billable work.

Margin Compression Is Structural

Margin compression in managed services is not a market mood. It is built into the model. Most MSP revenue is priced as a flat fee, per seat or per device, while the cost of delivering it is variable: a client who logs twice the tickets this quarter costs twice as much to serve, at the same price. That gap between fixed-price contracts and consumption-based delivery cost is where margin leaks quietly, and it widens as a relationship ages and scope creeps in unbilled. Cost only pushes the same way. The tooling stack keeps growing, and every annual price increase from a PSA, RMM, security, or documentation vendor lands straight on cost of delivery, while per-seat pricing, the most benchmarkable number a buyer has, sits under constant downward pressure.

The instinctive response is volume: more clients, more agreements, more tickets per engineer. It is also the trap. Utilization can be pushed for a while, but past roughly 75 to 80% service quality slips, SLA breaches climb, and the churn and credits on existing accounts quietly eat the revenue the new ones brought in. The deeper problem is the one almost no MSP can see: scaling on volume without knowing margin contract by contract scales the unprofitable contracts at exactly the same rate as the profitable ones. Revenue grows, margin shrinks, and the healthy blended number on the dashboard hides which accounts are doing the damage. That is not fixed by working harder. It is fixed by being able to see, per client and per contract, where the margin actually is, and is not.

4.5% Measured annual revenue leakage, professional services; higher on complex contract books (SPI Research)	**66–69%** Average technician utilization, against a 70–75% best-in-class target (SPI Research)
2–3 hrs Daily dispatcher time on manual scheduling, before Proxuma	**49%** of IT teams lack the staffing to secure their organization (JumpCloud)
4.8M Global cybersecurity workforce gap (ISC2, 2024)	**8.9x** Median MSP valuation, EV/EBITDA, the prize for getting this right (Aventis Advisors)

Revenue Leakage: The Silent Margin Killer

Revenue leakage is the gap between the work an MSP performs and the work it actually bills, and it is the rare profitability problem that never announces itself. There is no write-off, no lost deal, no angry client. There is only a margin number that lands a few points lower than it should, month after month, and gets explained away as a tough quarter. Because almost no one puts performed-work and billed-work side by side, the leak is invisible by construction.

The measured benchmark, from SPI Research, puts leakage near 4.5% of revenue for well-run professional-services firms, with a target below 5%. Vendor estimates for MSPs with complex contract books run higher, commonly cited at **5 to 15%**, though those are estimates rather than measured benchmarks. Either way the money is real: on a $3 million MSP, even the disciplined 4.5% is roughly $135,000 a year of work performed and never recovered, and a leaky book loses several times that.

The sources are predictable and structural:

- **Orphaned time entries:** Hours logged against the wrong ticket or client, or never logged at all. Nothing flags an hour that was never recorded, so it simply does not exist.
- **Scope ambiguity:** Work performed outside contract scope but never escalated for billing, because the engineer solving the problem is not the person who reads the contract.
- **Contract complexity:** Agreements modified, stacked, and grown over years until the billing engine can no longer calculate what is actually owed.
- **Billing drift across systems:** The devices in your RMM, the licenses in Microsoft 365, and the seats on the signed contract quietly diverge. You manage and secure 215 endpoints while the contract bills for 180, or provision licenses that never reach an invoice. The work is being done; the units simply never reconcile, because each number lives in a different system.
- **Late time entry:** Hours logged days or weeks later, after the consumption window has closed and the contract math is already wrong.

The most common leaks are not exotic. They are reconciliation gaps hiding in plain sight: the endpoints in Datto RMM that outnumber the devices on the contract, the Microsoft 365 licenses provisioned but never billed, the client whose tickets per seat run at double the volume the contract was priced for, the SLA tier being delivered a notch above what was sold. Every one of these is invisible to any single system, because catching it means comparing two: what is actually being delivered, in the RMM, the tenant, or the ticket queue, against what is being billed, in the contract. That comparison is precisely what a normalized cross-system model makes routine, and precisely what a PSA looking only at its own data was never built to do.

None of these is a discipline failure you can fix by telling technicians to be more careful, and it is not one you can dashboard your way out of either. The evidence is clear on this. In SPI Research's benchmark of professional-services firms, the high performers leak materially less than the rest, and the difference is not a single tool. It is maturity: better estimating, tighter change control, and, tellingly, how often leaders put the numbers in front of their teams. High performers review operational metrics with their staff roughly twice as often as everyone else, and that accountability cadence, not the existence of a report, is what tracks with lower leakage. SPI states it plainly: visibility alone is insufficient.

The reason is simple once you see it. Most leakage is human and structural at the same time: an hour not logged, a scope change not raised, a contract term not enforced. A system can surface every one

of those, but only a person can act on them. This is where Proxuma earns its place. By pulling Autotask and 12+ connected systems into one model, it makes the leak visible at the root, which contract, which engineer, which pattern, in real time rather than a quarter late. That is what lets a service manager coach the technician, fix the upstream dispatch process, and build a team that understands its own profitability. The integration finds the money. The people, armed with it, recover it.

A small, real example shows how literal that is. An MSP we worked with at Dxfferent billed travel time at half rate, on the reasonable-sounding assumption that round trips between client sites were eating into the day. It felt fair. It was quietly expensive. Reading the actual time entries told a different story: almost every travel entry came in under thirty minutes, because technicians moved client to client rather than back and forth from the office, and true round trips were rare. So instead of discounting all travel by 50%, we capped billable travel at thirty minutes. That single change, read straight out of the data, added about $7,500 a month and moved the business from 5% to 6.2% EBITDA. On the surface that is barely a point. In reality it is close to a quarter of the company's entire bottom line, recovered from an afternoon spent reading time entries that had been sitting there the whole time.

No new tooling found that money. The data already existed. What was missing was someone looking at it the right way, which is exactly the capability Proxuma is built to give every MSP by default.

Cybersecurity and the NIS2 Reckoning

Cybersecurity is both the largest growth driver and one of the largest cost centers in modern MSP operations. It is among the most frequently bought add-on services, and it also demands the highest-skilled engineers, the most tooling, and the most rigorous documentation and reporting. That tension alone would make security a discipline an MSP has to run on data rather than goodwill. In the EU, regulation has now made it unavoidable.

The NIS2 Directive (Directive (EU) 2022/2555) raised the cybersecurity floor across the Union. Its transposition deadline was October 2024, and through 2025 and 2026 it has been entering force across member states, with Germany's national law live since December 2025. Three things make it matter to every MSP. First, MSPs and managed security service providers are now explicitly in scope, named in the EU's own implementing regulation, with mid-size providers treated as important entities and larger ones as essential. Second, accountability sits with the management body: directors must approve and oversee the security measures and can be held liable, as transposed into national law. Third, the penalties are real, up to €10 million or 2% of global annual turnover for essential entities, and €7 million or 1.4% for important ones.

For an MSP the deeper shift is not the rules themselves. It is what they demand: proof. NIS2 requires in-scope organizations to manage the security of their supply chain, which means an MSP's regulated clients are now obliged to vet their provider and write security evidence into the contract. At the same time the MSP, in scope itself, must be able to show a supervisor its risk-management measures, its incident records, and its reporting on demand, against an incident cascade that runs from a 24-hour early warning to a 72-hour notification. None of that can be assembled by hand the week an auditor or a client calls. It has to be standing operational data: continuous, structured, and ready to produce. Compliance, in other words, has become an operational-intelligence problem, the same one this paper has described throughout, now with a regulator and a contract attached. Concretely, a normalized model can produce the audit trail on demand, which controls are in place, how each incident was handled and within what window, drawn from Autotask and the connected RMM, security, and documentation systems, instead of a manual scramble the week the regulator or the client calls.

> **AUDIT-READINESS IS DATA-READINESS**
>
> The MSP that can produce its security posture, incident timelines, and SLA evidence on demand turns NIS2 from a threat into a moat: it keeps its own regulated clients and wins theirs. The MSP still assembling that evidence by hand is one audit, or one client questionnaire, away from losing a contract. The difference is not how secure you are. It is whether you can prove it, on the day you are asked.

The Scalability Ceiling

Growth does not break an MSP gradually. It breaks at thresholds, and the same ones recur across the industry, because the practices that work at one size stop working at the next.

Up to roughly 8 to 10 engineers, a lead can dispatch from memory. They know who is good at what, who is free, and what is on fire, and reporting is a spreadsheet the founder updates. It works because one person can hold the whole operation in their head. Past 15 engineers that stops being true. Dispatching by gut starts dropping tickets, because the number of ways to assign work grows with engineers times tickets times skills times SLA deadlines, not in a straight line. The half-day billing reconciliation that was fine across 30 clients becomes a week-long job across 100. The quarterly report the founder produced personally turns into a role someone has to be hired into.

Each threshold is the point where a manual practice has to become a system, or growth stalls. Around 15 to 25 engineers it is scheduling and capacity planning. Past 30 to 50 clients it is billing reconciliation and contract consumption. Above that it is real-time reporting, because no analyst can refresh by hand fast enough to keep up. MSPs that hit these walls usually respond by adding overhead, coordinators, dispatchers, an operations manager, just to keep the existing machine running, which is why margin so often compresses precisely when revenue grows. The alternative is to put the infrastructure in before the wall, so the same team absorbs more work instead of more management.

The Real Lever Is Leverage, Not Headcount

Put these forces together and the conclusion is not that MSPs need more people. It is that the scarcest, most expensive asset they have, skilled engineering time, is being spent on work no engineer should have to do, while the data that would let them fix it sits unused. Decades of management research point the same way. From the World Management Survey across more than 20,000 firms to US studies of data-driven firms, performance improves when data is paired with the skills and the management practice to act on it, not when a dashboard is bought and left to run. The dashboard is not the lever. The capability to act on it is.

This is what operational intelligence is for. The engineer stays the one who does the work and makes the call. The layer underneath carries the operational load, watches the whole operation at once, and surfaces the one thing that actually needs a person, while handing the manager a real-time picture to coach and decide on. It does not replace the expert. It makes a scarce, expensive team measurably more effective.

> **THE OPERATIONAL LEVERAGE POINT**
>
> Four problems, one root: operational data that exists but never becomes action. The answer is not more headcount. It is freeing your best people from the work that wastes them, and putting the numbers in their hands while they can still act on them.

SECTION 04

The AI Adoption Paradox: MSPs Are Adopting AI Faster Than They Can Prove It Works

Widespread AI adoption. Elusive AI ROI. The explanation is not in the models, it is in the data that feeds them.

The Adoption Illusion

Kaseya's 2026 State of the MSP Report shows AI moving into day-to-day operations: **53% of MSPs now use AI to automate ticketing, patching, and monitoring**, 30% use it to eliminate repetitive tasks, and 48% rank AI and automation as the number one client need for 2026. Yet only 13% report generating meaningful revenue from AI. Adoption and measurable ROI are very different propositions.

The majority of MSP AI usage is concentrated in shallow applications: AI-assisted email drafting, chatbot-based first-line client interaction, and vendor-provided AI features bundled into existing tooling (including Cooper Copilot in Autotask). These applications have value, but they do not represent the operational intelligence transformation that AI can deliver, and that competitive MSPs increasingly require.

The Data Quality Root Cause

Integration is where AI adoption stalls. In MuleSoft's 2025 Connectivity Benchmark Report (1,050 IT leaders), **95% say integration challenges hinder their ability to implement AI**, and only 29% of enterprise applications are integrated, despite organizations running an average of 897 applications. In the MSP context, the underlying data problem takes specific, documented forms:

Inconsistent Ticket Classification

The same issue type (a password reset, a VPN failure, a server alert) is categorized differently by different technicians, or by the same technician in different contexts. Machine learning systems trained on this data learn nothing useful. Pattern recognition on inconsistent inputs produces inconsistent outputs.

Orphaned Time Entries

Time logged against the wrong ticket, the wrong project, or not logged at all. This data corruption propagates through every downstream calculation: utilization rates are wrong, contract profitability calculations are wrong, and billing is incomplete. No AI system can compensate for systematic input failure.

Inconsistent Contract Structures

Contracts that have evolved organically, with addenda, modifications, bundled services, and informal scope expansions accumulated over the client relationship, create a billing environment that defies systematic analysis. The contract structure recorded in the PSA frequently does not reflect the agreement that actually governs the relationship.

Missing Integrations

Work performed through RMM platforms, documentation systems, project management tools, and security platforms that is never reflected in the PSA. AI systems that can only see PSA data are making decisions with a fraction of the operational picture. The result is recommendations that are systematically incomplete.

Why Most AI Projects Fail

Independent research is consistent on the scale of AI project failure: RAND reports that, by some estimates, **more than 80% of AI projects fail**, roughly twice the rate of non-AI IT projects, and Gartner expects organizations to abandon 60% of AI projects that lack AI-ready data through 2026. The failure mode is almost never the AI model itself. Modern large language models are capable enough. The failure is almost always in the data: too little volume, too little quality, or too little structure for the model to learn anything useful.

For MSPs, the implication is stark: the $37 billion AI investment wave is predominantly flowing toward organizations that have spent years building clean, structured, integrated data infrastructure. MSPs that have not made this investment are not AI-ready, regardless of what AI tools they have licensed.

AI as a Multiplier of Discipline, Not a Fix for Chaos

AI does not fix a messy operation. It multiplies whatever discipline is already there. Point it at clean, consistent data and it earns its place: it catches the client whose ticket volume quietly doubles two months before they churn, or the contract that slipped below margin three weeks ago. Point it at the half-entered, scattered data most MSPs actually run on, and it produces confident nonsense.

The failure mode that should worry an MSP is not AI that breaks loudly. It is AI that is confidently wrong. A hallucinated answer does not arrive with a warning; it arrives as a clean number on a clean dashboard. You do not get an error, you get false certainty, and you price a contract or staff a project on data that was never real. That is worse than having no AI at all, because it replaces an honest "we don't know" with a convincing wrong answer.

> *"Buying AI is the easy part. The MSPs that win do the hard part first: they clean their data, so they can trust what the AI tells them."*

Picture it concretely. Across a dozen clients, the same misconfigured update is quietly generating tickets, a handful here, a few more there, logged under different categories by different technicians and never obviously connected. Handled the usual way, that is fifty separate tickets and fifty separate closes. Read as one dataset, with consistent data underneath, the pattern is unmistakable: one root cause, one fix, pushed once through the RMM, and the whole class of tickets stops coming back. That is what AI and automation are genuinely good at, finding the signal in volume, but only when the data is clean and connected enough for the pattern to be real. On fragmented, inconsistently tagged data, the same tools see fifty unrelated incidents and help with none.

Proxuma fixes the data first. Every integration it builds, across Autotask and all 12+ connected systems, feeds one normalized data model. Data from across the MSP stack, the PSA, monitoring and backup, documentation, CRM, security, HR, networking, and the Microsoft environment, maps into a single consistent schema, so the same client, contract, or engineer means the same thing in every measure. On top of that grounded model sit the anti-hallucination safeguards that make AI usable in a business you have to answer for: answers are tied back to the underlying records and flagged when the data does not support them, MSP operators from Dxfferent stay in the loop to check what the system surfaces, and the data and models run on infrastructure you control, including the option of Proxuma's own dedicated hardware. The 500+ pre-built measures sit on the same clean model, so every key MSP metric is calculated the same way every time. The result is AI you can steer a business by, rather than AI bolted on for a demo.

It is worth being precise about how that works, because anti-hallucination is an easy phrase and a hard thing. A general AI chatbot generates its answer word by word, and when it does not actually know, it will still produce a confident, plausible number, because producing something is what it does. Proxuma works the other way around. The AI does not invent the figure. It translates the question into a query against the governed data model, and the model returns the number from the real records. The language model handles the language; the data handles the truth. Every answer traces back to the rows it came from, and when the data cannot support a question, the honest result is "we don't have that," not a confident guess. That is the difference between asking a chatbot, which answers, and asking your data, which goes and looks. It is also why the foundation matters more than the model on top: any copilot, Microsoft's or Proxuma's, is only as trustworthy as the semantic layer beneath it, and that layer is the part most MSPs have never built.

The Integration Imperative

For an MSP, the integration gap is not abstract. The PSA, the RMM, the documentation system, the security stack, billing, and accounting each store data in their own schema, with their own taxonomy and no native way to analyze across them. A question as basic as which clients are unprofitable once security tooling and after-hours work are counted touches five systems that do not talk to each other. The broader IT world has measured the same wall: in MuleSoft's enterprise research, 95% of IT leaders say integration challenges hold back their AI. For MSPs the gap is sharper, because the operational picture is split across more disconnected tools than almost any other kind of business.

Proxuma's Data Connectors and Unified Data Model layers close that gap by pulling every system into one coherent analytical layer. The payoff is organization-wide data readiness: the MSP can put Microsoft Copilot, vendor-embedded AI, or any future intelligence layer on data that is already clean, connected, and structured, instead of discovering on day one that the foundation underneath is not ready.

SECTION 05

The KPI Benchmarking Framework: What Gets Measured Gets Managed

The gap between the KPIs MSPs should track and the KPIs they actually track in real time is where margin goes to die.

Managing by Anecdote

Ask the average MSP owner how their most profitable client is doing this month. They will tell you a story. They will describe a recent ticket escalation, a positive client conversation, or a renewal discussion that went well. What they will not tell you, because they cannot tell you, is the exact gross margin on that account for the current period, the engineer hour rate being realized against the contract value, the utilization rate of the resources deployed on that account, or the trend line on SLA performance over the last six months.

This is managing by anecdote. And it is the default operating mode for the majority of MSPs, not because their leaders lack sophistication, but because their operational infrastructure does not make real-time KPI visibility possible.

The Metrics That Decide Profitability, and What It Takes to See Them

Each metric below drives a real profit decision, and most MSPs cannot see them in real time. The inputs are scattered across systems, they need business-hours and time-intelligence logic, and they have to be normalized before they mean anything. Producing them live is the work Proxuma Intelligence's 500+ pre-built measures and modeled relationships do. Here is what each metric measures, what good looks like, and why it decides whether an MSP makes money. The targets are directional, what a healthy operation aims for, not a published industry benchmark, because for several of these no consistent benchmark exists.

METRIC	WHAT IT MEASURES	WHAT GOOD LOOKS LIKE	WHY IT MATTERS
Engineer Hour Rate (EHR)	Revenue generated per engineer hour logged	Tracked per contract, comfortably above loaded cost	Bill rate is what you quote; EHR is what you keep after non-billable hours, and the gap is where margin hides. With talent scarce and expensive, it is the line between a healthy desk and one that quietly loses money.
Technician Utilization	Each engineer's true planned load, combining shifts, regular time, Outlook items, tickets, and tasks, not just logged billable hours	A full but realistic load, with no one silently booked past 100%	Billable-hours utilization misses the real risk: the engineer whose agenda runs well past 100% once shifts, regular time, Outlook, tickets, and tasks are added together. It is usually your best people, pulled into everything, and it stays invisible until every commitment sits in one view.
Mean Time to Resolution (MTTR)	Average time from ticket open to resolution	< 4 hours (P1)	Resolution speed is what clients actually feel, and in the US it is increasingly written into SLAs with financial penalties. Slow priority resolution is the leading early signal of the churn that follows.
Contract Profitability	Gross margin by contract/client after direct costs	≥ 50% (service GM)	Most MSPs know their overall margin but not which contracts earn it and which destroy it. In a US market being rolled up by private equity, per-contract margin is exactly what an acquirer scrutinizes, and what sets your exit multiple.
Revenue Leakage Rate	% of performed work not captured in billing	< 2%	Work delivered but never billed is margin you already paid to produce. At US contract sizes, even a few percent is six figures a year walking out the door unnoticed.
Recurring Revenue & Retention	Recurring revenue per month, and whether the base expands faster than it churns (net revenue retention)	Net revenue retention above 100%	Recurring revenue makes an MSP predictable; net revenue retention asks the sharper question, whether the base grows through upsell faster than it churns. Above 100% means the book expands before you win a single new client, exactly what buyers pay premium multiples for.
Account Health & Churn Risk	A composite early-warning signal per client, combining rising tickets and escalations, falling satisfaction, SLA slippage, late or unpaid invoices, and shrinking license usage	At-risk accounts flagged months early, while there is still time to save them	By the time churn shows in the numbers, the client has usually already decided. The early signals, rising tickets, falling satisfaction, slipping SLAs, late invoices, unused seats, each sit in a different system and only become a reliable warning when one model reads them together. A quarter's early warning is worth far more than the churn rate measured after the client is gone.
SLA Compliance Rate	% of tickets resolved within contracted SLA windows	≥ 95%	SLA attainment is the contractual promise an MSP is actually judged on. In the US enterprise and security space, missed SLAs trigger credits, penalties, and lost renewals on the highest-value accounts.

METRIC	WHAT IT MEASURES	WHAT GOOD LOOKS LIKE	WHY IT MATTERS
Engineer Efficiency vs Revenue	In an MRR model priced per seat or asset, the billable hours an engineer logs against the recurring revenue of the clients they serve, expressed as an effective rate	More client revenue supported per engineer-hour, and that rate rising over time	In a per-seat MRR model, raw revenue-per-head is meaningless. The sharper read is the recurring revenue an engineer's clients pay divided by the hours logged against them, an effective rate per engineer. It is hard to build, joining contract revenue, client assignment, and time across systems, but it is the truest test of whether your scarcest resource serves the clients that pay for it.
Repeat-Issue Leverage	Tickets closed per root-cause fix: how many repeat tickets trace back to a single underlying problem	A high and rising ratio, one problem fix retiring dozens of tickets	The biggest leverage is rarely closing tickets faster. It is noticing that two hundred of them share one root cause and killing it with a single script, instead of handling each by hand, a pattern that stays invisible until tickets across every client are read as one dataset.

Best-in-class figures are the targets Proxuma holds from MSP operations. Engineer Hour Rate (EHR) is Proxuma's term for revenue realized per engineer hour, what the industry also calls the client effective rate. Each calculation above is produced by Proxuma Intelligence's pre-built measure library, running across the normalized, connected data model instead of a single system's reports.

Three Numbers That Decide Your Margin

Contract Profitability: the number that survives diligence

Ask an MSP owner for their gross margin and most can give you a company-wide figure. Ask which of their forty contracts are actually profitable and the answer turns into a guess. That gap is the whole problem, because margin is not made at the company level. It is made or lost one contract at a time. A flat-fee agreement signed three years ago, never repriced while the client doubled their endpoints and bolted on a security stack, can quietly run at a loss while the dashboard still shows a healthy blended number. The profitable contracts subsidize the sick ones, and nobody can see which is which.

Computing real per-contract margin is harder than it looks, which is why so few do it. You take the revenue on each agreement and subtract the true cost of delivering it: engineer time at loaded cost rather than bill rate, the product and license COGS that sit in a different system, and the tooling allocated to that client. Those inputs live in the PSA, the billing platform, and procurement, and they only mean anything once they are joined to the same contract and the same period. Proxuma performs that join continuously, so margin by contract is a live number rather than a quarterly archaeology project.

In the US market this has stopped being a back-office nicety. The sector is consolidating fast, and private-equity-backed buyers run diligence on exactly this: not the headline margin, but how durable that margin is contract by contract. An MSP that can show clean per-contract profitability negotiates from strength and earns a better multiple. One that cannot is asking the buyer to take its word for it, and buyers discount what they cannot verify.

Client Churn: the number acquirers price you on

Most MSPs track churn after the fact. A client gives notice, the recurring revenue drops, someone updates a spreadsheet. By then the decision is already made. The expensive part of churn is not the lost monthly fee, it is everything that went into winning that client in the first place: the sales cost, the onboarding, the months before the account turned profitable, all written off at once. Keeping a marginal client is far cheaper than replacing a lost one, yet most MSPs cannot see which accounts are drifting until the notice lands.

Reading churn properly means tracking each contract's lifecycle, separating true loss from downgrades and from the upsell that masks it, and connecting it to the early-warning signals that precede it: a client whose ticket volume is climbing, whose satisfaction is slipping, whose SLA performance has quietly degraded. Those signals sit in different systems and are rarely tied back to the revenue at risk. In the US, where MSPs are valued and acquired on net revenue retention, churn is not only an operational metric, it is a valuation input. A book that retains 95% of its revenue is worth materially more than one that retains 85%, and a buyer will find that number whether the seller tracks it or not.

The economics are not subtle, and the recent numbers are blunt about it. Growth from the base you already have is far cheaper than new logos: in recurring-revenue businesses, winning a dollar of brand-new customer revenue costs roughly four times what it costs to expand an existing account (KeyBanc Capital Markets, 2024). And the firms whose base expands faster than it churns grow dramatically faster for it: across private software companies, the highest-retention cohort grows at nearly double the median rate (SaaS Capital, 2025). For an MSP the lesson is exact: the cheapest growth you will ever book is the client you already have, kept and quietly expanded. The leak is the most expensive line on the books, and the one fixed earliest by seeing it coming.

Technician Utilization: the lever hiding in plain sight

Utilization is the share of an engineer's paid, available time that turns into productive, billable work. It sounds simple and is almost always measured wrong, because "available" has to net out PTO, holidays, training, and internal work, and "billable" has to be classified the same way across managed-services and project work. Get the definition loose and the number is noise. Get it right and it becomes one of the strongest levers an MSP has.

The reason is arithmetic. In the US, where senior engineering talent is among the most expensive and hardest to keep anywhere, a fully loaded engineer is a large fixed cost whether they are productive or idle. Moving utilization a few points costs nothing extra in salary; it simply converts cost you are already carrying into delivered value. For most MSPs that is a faster and cheaper route to margin than chasing new logos, which arrive with acquisition cost and onboarding drag. The catch is that you cannot manage utilization you cannot see in real time, by individual and by team, and that live view is exactly what a PSA does not provide on its own.

Utilization also has a ceiling that has nothing to do with arithmetic and everything to do with people. Push a team sustainably past the healthy band and you do not just risk service quality, you burn out the engineers you most need to keep. That matters more here than in almost any other business: professional-services firms ran staff attrition around 12% in 2024 (SPI Research), every departure walks out with client knowledge, and with a global cybersecurity workforce gap near 4.8 million (ISC2) the replacement is slow and expensive. So the very metric that drives margin, pushed blindly, destroys the asset it depends on. This is where data-driven operational intelligence earns its name. Proxuma watches load and utilization per engineer in real time and flags the individual drifting toward overload before they burn out, not after they resign. The goal is not to run people hotter. It is to hold utilization in the productive band, fairly distributed, so the team stays profitable and intact at once.

The Real Gap Is Latency, Not Knowledge

Here is the uncomfortable part: every metric in this section is trackable, and almost no MSP tracks them in time for it to matter. The data exists, scattered across the PSA, the RMM, and billing, but turning it into a live number takes manual extraction and assembly no operations team can sustain weekly, so it gets done quarterly, if at all. The cost of that delay is not abstract. MIT Sloan Management Review puts the price of bad and stale data at **15 to 25% of revenue** for the typical company. And the decisions made on it are only as good as the data underneath: McKinsey finds just one organization in five believes it makes good decisions, and that managers waste well over half the time they spend deciding, much of it because the information arrives late or unclear. A metric you see once a quarter is not a management tool, it is a post-mortem. By the time the report shows a contract that slipped below margin, you have already delivered three more months of it at a loss.

That is the real gap between average and best-in-class, and the evidence says it is speed, not knowledge or talent. The leaders run on the same tickets, the same contracts, and the same tools as everyone else. What they do differently is look more often and act sooner. In SPI Research's professional-services benchmark, the top performers review operational metrics with their teams nearly twice as often as the rest, and score materially higher on running the business by data. McKinsey's research points the same way: organizations that decide both quickly and well are roughly twice as likely to report superior financial returns. The average MSP runs a quarter behind its own numbers; the best-in-class runs a week behind, close enough to catch the margin slipping, the engineer tipping into overload, the client starting to drift, and the hours going unbilled while each is still small and cheap to fix. The advantage was never knowing more. It is knowing in time, and that comes down to how fast raw data becomes a number you can trust. A spreadsheet refreshed once a quarter cannot do it. A live model can.

Engineer Hour Rate as the Profitability Compass

Of all the KPIs in the framework, the effective hourly rate, what Proxuma calls Engineer Hour Rate (EHR), is the most useful and the least often tracked in real time. It captures the efficiency of the whole service delivery operation in a single number: how much revenue is generated for every hour an engineer invests in client work.

Consider the illustration: an MSP billing a client at an effective rate of $62.50 per engineer hour, because the contract is underpriced, the scope has expanded without adjustment, or time is being miscategorized, while the same engineer generates $145 per hour on a different account. The business case for repricing, renegotiating, or transitioning the underperforming contract is clear. But it is only visible if EHR is tracked by contract and by client.

Proxuma Intelligence's pre-built measure library includes EHR calculation at the contract, client, engineer, and portfolio level, updated in real time as time entries flow through Autotask. This is not a report that requires a financial analyst to produce. It is a live operational instrument that any service manager can access.

> **REPORTING VS OPERATIONAL INTELLIGENCE**
>
> This is the line between reporting and operational intelligence. A report tells you what happened last quarter. Operational intelligence tells you what is happening now, in time to change it: which contract to reprice before the renewal, which engineer to move before they burn out, which client to call before they churn. The number itself is not the point. Acting on it while it still matters is, and that is only possible when the measure is live, not assembled by hand after the quarter has already closed.

SECTION 06

06

The Proxuma Solution: Three Modules, One Operating System

Built by MSP specialists, starting with Autotask. Live on the PSA you already run within 24 hours, not months.

Architecture Philosophy

Operational intelligence is not a dashboard you bolt on. It is the layer that turns the work a PSA records into decisions you can act on, and in time into actions you can safely automate. That is what Proxuma is, and it is why it sits above Autotask rather than replacing it. Every module connects directly to the Autotask environment and reads, normalizes, and surfaces the data that already exists, in the form needed to act on it, with no rip-and-replace, no migration, and no months-long implementation. Two engines do the heavy lifting underneath: a constraint-optimization solver that computes schedules algorithmically, which most MSP tools never attempt, and a normalized, AI-ready data model that can be queried in plain language. The same principle runs through both. The machine carries the operational load and surfaces what matters; the decision stays with the person.

The three modules, Planning, Shifts, and Intelligence, each address a distinct operational gap, and their combined effect is the transformation of an Autotask-based MSP into a data-intelligent, AI-ready operation. Most MSPs begin with Planning and Shifts to get service delivery running cleanly, then add Intelligence to turn that clean operational data into financial and AI-ready insight.

01 Proxuma Planning

Algorithmic scheduling, project planning, and resource capacity, from today's queue to years ahead

A dispatcher building the daily schedule by hand typically spends **2 to 3 hours** processing the queue, assigning tickets, checking engineer availability, and re-juggling the plan every time something changes. Proxuma Planning collapses that work with a constraint-optimization engine, the kind of true algorithmic scheduling most MSP tools do not attempt, that weighs skills, availability, SLA deadlines, capacity, and travel across the entire queue to compute an optimized schedule in seconds. The dispatcher reviews the result and adjusts it by drag-and-drop, staying in control of the call. In Dxfferent's deployments, work that took roughly five minutes per ticket by hand now takes about thirty seconds, which turns hours of daily scheduling into well under an hour. And Planning is not only the daily queue: the same environment runs multi-project Gantt planning and multi-year resource capacity, so the data that schedules today also answers whether the team can take on next quarter's project without breaking the work already committed.

Live Drag-and-Drop Scheduling

Real-time ticket assignment with instant visibility into engineer availability. Changes in the queue propagate immediately: no refresh cycles, no manual queue reconciliation.

Project Gantt & Multi-Project View

Open any project in a full Gantt, every phase, task, and service call on one timeline, or see every project at once, filtered by team, project lead, or client, and reschedule by dragging or stretching a bar. The multi-project, multi-year project view Autotask does not provide natively.

Multi-Year Resource Capacity

Capacity weeks, months, and years ahead, per resource. Colored bars show planned hours, spare capacity, and overbooked time at a glance, calculated from real workload across both Autotask service calls and Outlook, with PTO, holidays, and internal meetings included, not just who happens to have tickets assigned.

Plan Bucket

Staged ticket management for pre-planning and queue organization. Work items can be organized into planning buckets before formal assignment, enabling more deliberate, strategic dispatch rather than reactive queue management.

External Invites

Scheduling coordination for partner resources, subcontractors, and vendor specialists, extended into the same Proxuma Planning environment without requiring separate workflow management.

OPERATIONAL IMPACT

MSPs using Proxuma Planning report dispatcher time on scheduling tasks falling from 2 to 3 hours to roughly 30 minutes per day, recovering 1.5 to 2 hours of senior operational capacity daily that can be redirected to client escalations, team management, or strategic planning.

02 Proxuma Shifts

Intelligent coverage management for NOC, helpdesk, and on-call operations

Coverage is where three problems collide: keeping the service desk staffed, keeping SLAs intact, and not burning out the people who carry the load. Most MSPs manage it by hand, on a whiteboard or a shared calendar, which is exactly why it goes wrong. The senior engineer quietly ends up on every hard rotation, a gap opens at the worst possible time, and no one sees it until a client does. Proxuma Shifts treats coverage as an operational-intelligence problem instead. It scores the real workload behind each shift from actual ticket volume and complexity, not just hours on a calendar, catches conflicts and coverage gaps before they reach the live schedule, and keeps the whole rotation visible and fairly balanced across the team, so load is distributed by data rather than by who is easiest to ask.

Intensity-Based Balancing
Not every shift is equal: nights and weekends weigh heavier than daytime cover. Proxuma scores that intensity and spreads the hard shifts fairly, and you can set how much capacity a shift consumes, say 80%, so an engineer on coverage still has room for planned work. Load is distributed by data, not by who is easiest to ask, which is what keeps a team from burning out.

Automatic Rotation
Set rotation rules once and shifts cycle on their own, fairly and predictably, instead of a manager rebuilding the rota by hand every week.

Conflict & Coverage Detection
Overlapping assignments, insufficient rest periods, and skill-coverage gaps are caught in planning, before they reach the live schedule. A problem caught in planning costs nothing; one that reaches service delivery costs clients.

Autotask + Outlook Aware
Proxuma plans against both Autotask activities and Outlook, with PTO, holidays, and meetings included, so a shift never lands on top of work or time off that already exists. Assignments flow back to engineer calendars automatically, and everyone sees who is on duty, on standby, or off in one view.

03 Proxuma Intelligence

The MSP intelligence layer: 500+ pre-built measures, run it on Power BI or your own stack, AI-ready for chat-with-your-data

Proxuma Intelligence is the layer that turns a clean operation into a self-aware one, and it is where Proxuma's longest-lasting advantage sits. It takes the operational data already flowing through Autotask and 12+ connected systems and shapes it into one connected, analytical, AI-ready model: 500+ pre-built measures on a four-layer architecture you can run through Power BI or host on your own stack, so no third party ever stores your operational data. It is built so the answer to almost any operational question is already modeled and waiting.

LAYER 1 — Data Connectors
Direct, maintained integrations with Autotask and 12+ connected systems across the MSP stack: the PSA, monitoring and backup, documentation, CRM, security, HR, networking, and the Microsoft environment. Data flows automatically from source systems into the analytical layer: no manual exports, no scheduled batch jobs, no data engineering required from the MSP.

LAYER 2 — Unified Data Model
A normalized, dimensional data model that structures raw operational data into coherent analytical relationships. Tickets, time entries, contracts, engineers, clients, and SLA data are connected through hundreds of modeled relationships in a consistent schema that enables cross-entity analysis: the foundation of genuine operational intelligence.

LAYER 3 — 500+ Pre-Built Measures & Dashboards
More than 500 pre-calculated measures and ready-to-use dashboards covering the KPIs that decide an MSP's numbers: team and individual utilization, contract and project profitability, Engineer Hour Rate, SLA compliance, MRR trends, pipeline, and quarterly reviews, plus dozens of financial report packs for client profitability, margin, and cost of service. Work that would take an analyst weeks to build, live out of the box.

LAYER 4 — Chat With Your Data AVAILABLE VIA MCP
Ask any operational question in plain language, across every connected system, and get a grounded answer: "Which contracts are under our margin target this quarter?" "Who is overloaded next week?" "Where did we leak hours last month?" Because every answer comes from the one normalized model, it is anti-hallucination by design, reading from a single source of truth rather than guessing across inconsistent copies. This is not only a roadmap promise. Through a Model Context Protocol (MCP) interface, an MSP can already connect the AI of its choice to the Proxuma model and query it directly, on its own terms and within its own compliance requirements, instead of being locked to a single vendor's assistant. The clean, fully modeled semantic layer is the hard part, and it is already built; the turnkey conversational experience matures alongside Microsoft's Power BI Copilot and Proxuma's own AI engine.

One Capability, Concretely: Revenue by Contract Type

To compute it, the model joins time entries, contract terms, and delivery cost across systems, then splits revenue by contract structure, flat-fee managed, time-and-materials, block hours, project, and recurring, to show which structures earn margin and which

quietly consume it. Without that joined model it is a multi-week data exercise that happens, at best, once a year. Here it updates continuously. The payoff is concrete: an MSP walks into a renewal already knowing which contract types lose money and by how much, and reprices from evidence instead of instinct.

Why Proxuma, and Not the Alternatives

Operational intelligence is not a feature you can buy in pieces, and that is where the common alternatives fall down. Each delivers a fragment, a model, a report, a calendar, a new platform, and leaves the connective layer that makes data act as one still missing.

🧱

Build it yourself in Power BI

A capable BI consultant can model Autotask data, given months of work and a standing maintenance burden. The model needs rework whenever the source schema changes, and it carries none of the MSP-specific logic. Proxuma ships 500+ maintained, MSP-specific measures from day one.

📋

PSA-native reporting

Autotask reports and Cooper Copilot work inside the PSA, on PSA data, in the form the PSA stores it. They are not designed to connect scheduling, capacity, EHR, and SLA into one real-time, cross-system picture. That cross-entity layer is what Proxuma adds.

🔁

Rip-and-replace platforms

Enterprise field-service or ITSM suites can deliver intelligence, after a multi-month migration, a new system of record, and a price tag to match. Proxuma sits on the Autotask you already run, live within 24 hours, no migration.

🎯

Point-solution AI schedulers

Standalone AI scheduling tools optimize the calendar, but on their own they leave you without capacity visibility, EHR, contract profitability, or a connected data layer. Proxuma treats scheduling as one part of one operating system, tuned per MSP.

The pattern is the same in every case. Each alternative solves one fragment and leaves the cross-system layer, the part that turns scattered data into one operating picture, unbuilt. That layer is the whole point of operational intelligence, and it is what Proxuma is, on the Autotask you already run.

A Tuesday, Before and After

Before. The dispatcher rebuilds the day's schedule by hand off a spreadsheet. A senior engineer quietly takes the third hard ticket in a row. The service manager learns a flagship client slipped below margin when the quarterly report lands, six weeks too late. Forty endpoints in the RMM never make it onto an invoice. None of it is visible until something breaks.

After. The schedule is proposed in seconds and adjusted by drag-and-drop. Load is balanced by data, so no one silently burns out. The margin slip surfaces the week it happens, in time to reprice at renewal. The billing gap shows up as a line item to recover. Same team, same tools, same clients. The only difference is that the operation can finally see itself, in time to act.

SECTION 07

The ROI Business Case: Putting Numbers to the Opportunity

An illustrative MSP analysis: $3M revenue, 15 engineers. A simplified estimate of where the value sits.

Model Parameters

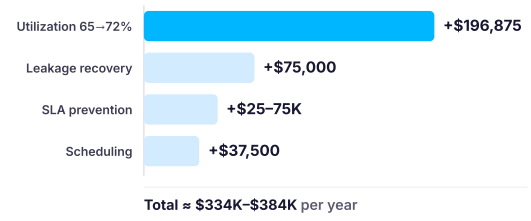
This is a deliberately simple model of a representative mid-market MSP: $3 million in revenue, 15 engineers, metrics at the industry average. It is conservative by construction. It counts only directly attributable gains, the drivers overlap rather than stack (recovering leaked revenue and lifting utilization are themselves what move gross margin, so margin is not summed again as a separate line), and it ignores the strategic upside entirely, competitive differentiation, M&A readiness, lower management overhead. Every input is grounded in a benchmark cited elsewhere in this paper or in Dxfferent's own deployment data.

It is an estimate, not a promise, and it will vary with each MSP's size and contract mix. And it still lands between $334,000 and $384,000 a year on a $3 million business, more than a tenth of revenue, recovered from data the MSP already owns.

VALUE DRIVER	CURRENT STATE	PROXUMA-ENABLED STATE	ANNUAL VALUE
Revenue Leakage Recovery ~5% leakage on $3M (SPI Research measures ~4.5%; vendor estimates run 5–15%), about half recovered through tighter time entry and contract visibility	~$150K leaking annually	~50% recovered	+$75,000
Utilization Improvement 65%→72% utilization across 15 engineers at $125/hr, ~1,500 billable-available hrs/yr each (best-in-class technician utilization sits around 70–75%); assumes recovered capacity is sold against existing or near-term demand	65% utilization (~14,625 hrs billed)	72% utilization (~16,200 hrs billed)	+$196,875
Scheduling Efficiency Dispatcher savings: 1.5 hrs/day × 250 days × $100/hr coordinator rate (Dxfferent deployment data)	2 hrs/day on manual dispatch	~30 min/day with Proxuma Planning	+$37,500
SLA Breach Prevention Estimated value of preventing avoidable, scheduling-related SLA breaches (penalties, credits, churn). Most breaches have non-scheduling causes; this counts only the schedulable subset.	Several avoidable scheduling-related breaches/yr	Conflict-aware scheduling prevents most of them	+$25,000–$75,000
Total Illustrative Annual Value Drivers partly overlap and are not additive guarantees. An estimate, not a promise.			**$334,000–$384,000**

Where the annual value comes from

About $334K–$384K a year on a $3M MSP, most of it from utilization.



Total approximately $334,000–$384,000 per year. Drivers partly overlap and are not additive guarantees. An estimate, not a promise. Source: model above (Dxfferent deployment data and cited benchmarks).

The Engineer Hour Rate Lens

The headline margin hides the detail that matters most: which clients actually make money. Engineer Hour Rate (EHR) shows it, the revenue a client generates for every engineer-hour it consumes. The trap is that two clients can pay almost the same monthly fee and still be worlds apart in profitability, because one quietly eats far more engineering time than the other. Take three clients whose contracts look similar on paper:

- **Client A** pays $8,500 a month, but consumes 136 engineer-hours. That works out to just **$62.50 per hour**.
- **Client B** pays $9,000 a month, and consumes only 48 hours: **$187.50 per hour**.
- **Client C** pays $11,600 a month across 80 hours: **$145 per hour**.

On the invoice, all three look like healthy accounts, and Client A even looks fine because it pays nearly as much as Client B. Through the EHR lens it is the problem: it earns barely a third of what the others do per engineering hour, because it consumes roughly three times the support its fee assumes. That is the contract to reprice or rescope, and without EHR measured by contract no one would ever have flagged it, because the invoice looks perfectly normal.

Revenue per engineer-hour, by client

Three clients on a similar fee, three very different realities.



Illustrative, from the worked example above. All three pay a comparable monthly fee.

The Strategic Dimension the Table Does Not Price

There is value here the ROI model above leaves out. In a market consolidating fast, an MSP whose margin, utilization, and contract profitability are provable in live data is worth materially more than one that needs a 90-day data scramble before a buyer can even assess it, and it can fold its own acquisitions onto one model from day one rather than inheriting another black box. Section 8 quantifies that. The point here is simpler: the same infrastructure that runs the business better is what makes it worth more.

INVESTMENT PERSPECTIVE

The build-versus-buy math has quietly shifted. AI makes it tempting to think you can just build this yourself, and you can stand something up in a weekend. The first version was never the hard part. The hard part is making it production-grade and anti-hallucination, keeping it alive as every connected system changes its schema, and having the right people to do that instead of billing clients. So the real choice is not Proxuma versus doing nothing. It is Proxuma versus the options that quietly cost more: vibecoding a model you then maintain forever, hiring a data team to run it, or stitching together a drawer of point tools that never become one picture. Proxuma is the unified command center, at the price of a single operational tool.

SECTION 08

Where Operational Intelligence Takes You

From today's baseline to an investment-grade MSP, and what each step along the way is worth.

Where You Are, and Where the Steps Lead

These five levels are the data-and-operations slice of Dxfferent's MSP maturity model, a framework that scores an MSP across nine pillars, from this one through organisation, governance, portfolio, and go-to-market. We focus on data and operations here because it is the pillar AI depends on, and the one Proxuma was built for. Maturity tracks profitability closely: MSPs that run on data reach on the order of 2.6 times the profitability of those that do not (Service Leadership OML). The ladder is the honest version of where most MSPs sit today, and being low on it is not a failing: for most, the infrastructure to climb was simply never within easy reach. The encouraging part is that you do not climb it one rung at a time over years. The step that matters is a single, accessible jump, and it starts from wherever you stand right now.

LEVEL 1 — Ad Hoc

Characteristics: No reliable operational numbers. KPIs are looked up by hand when someone asks for them. Decisions run on instinct and the loudest voice in the room.

AI Readiness: None. Data is too sparse and inconsistent to support any meaningful analytics.

LEVEL 2 — Scattered Spreadsheets

Characteristics: A handful of KPIs tracked in separate spreadsheets, inconsistently maintained and rarely shared. Reporting is manual and occasional.

AI Readiness: Minimal. Numbers exist but are siloed across people and tools.

LEVEL 3 — Central Reporting

Characteristics: Core KPIs reported monthly through one central dashboard. The whole team finally sees the same numbers, but after the fact.

AI Readiness: Limited. Data is centralized, but still backward-looking and not yet connected across systems.

LEVEL 4 — Real-Time, Acted On PROXUMA PLANNING + INTELLIGENCE

Characteristics: Live dashboards whose KPIs actively drive day-to-day decisions, by client, engineer, contract, and service line. Scheduling and coverage run on capacity data, not instinct.

AI Readiness: High. A clean, connected, real-time model the AI can sit on top of.

LEVEL 5 — Predictive AI + COPILOT + CLEAN DATA

Characteristics: Predictive analytics on historical data steer proactively, before the problem is visible: churn risk, capacity shortfalls, and margin drift surfaced early, with natural-language queries via Copilot.

AI Readiness: Full. AI amplifies every operational decision, on a foundation it can trust.

The Top Rungs Are Closer Than They Look

Level 5 reads like the future, but its building blocks are here now, and they all rest on the Level 4 foundation: a clean, connected, real-time model. Microsoft's Copilot for Power BI is generally available today, so on a clean model an operations manager can already ask, in plain language, which engineers were under-utilized last quarter or which contracts sit below margin, and get an answer in seconds that used to take an analyst days. Predictive operations, acting on a churn signal or a margin slip before it lands, are not a product to buy either; they are pattern recognition on connected data. In both cases the prerequisite is identical, and it is the only hard part: a normalized model the AI can trust. The MSP that has built it switches these capabilities on the day they arrive, from Proxuma, from Microsoft, or from anywhere else. The one that has not watches them work for everyone but itself.

The Endgame: An Investment-Grade MSP

Most MSPs sit at Level 2 or 3 today: data in spreadsheets, or a monthly report that lands too late to act on. The jump to Level 4, where operations are seen and steered in real time, is the most accessible, highest-return step a typical Autotask MSP can take, and it is exactly the jump Proxuma is built to make, in months rather than the years of manual maturation it would otherwise take. It is not a tool you graduate into once you are already advanced; it is the lever that gets you there.

You do not have to get your house in order before you start, because putting it in order is exactly the job it does. That same clean foundation is what Level 5 is later built on. The operational gains, though, are only half the reason to take the step. The other half is what that foundation is worth.

The MSP market is consolidating fast, with 169 publicly-announced deals in 2025 and private equity involved in roughly 70% of them (Omdia), and it rewards scale steeply: where a sub-$5M MSP changes hands near 5x EBITDA, platform-scale operators reach the low teens (Aventis Advisors). Operational intelligence is the price of entry to the two ways an MSP wins that game. You can become the asset everyone wants to buy: the MSP whose margin is provable contract by contract, whose churn is low and visible, and whose data survives due diligence without a 90-day scramble, which is exactly what EY finds separates a premium exit from a discounted one. Or you can become the platform that does the buying, folding each acquired MSP onto one normalized model and seeing its true performance from day one post-close instead of inheriting another black box to untangle.

Either way, the MSP that runs on operational intelligence sets the terms in the consolidation rather than being priced by a buyer who can read its numbers better than it can. And the advantage compounds, because it cannot be bought back late. Every month of clean, connected data is a month of history a competitor who starts later will never recover. In a market being repriced in real time, starting now is itself the edge.

CONCLUSION

09

The Intelligent MSP: What Winning Looks Like

The managed services market will nearly double to $847 billion by 2033, but that growth will not lift every MSP equally. It is already concentrating toward the operators who run with precision at scale, and in a market consolidating this fast, the distance between them and everyone else widens every quarter. The MSPs that win the AI era will not be the ones that bought the most AI. They will be the ones whose data was ready for it.

The MSPs that lead this market in 2030 share a set of defining operational characteristics visible today. They know, precisely and in real time, which clients are profitable, which contracts need repricing, which engineers are overloaded, and which service lines are generating return on capacity. They can answer capacity questions not just for this week but for the next quarter. Their scheduling is conflict-free and data-driven, not reactive and manual. Their data is clean enough for AI to act on. Their reporting tells a coherent story that supports confident pricing, hiring, and investment decisions. On the maturity ladder, this is Level 4 and 5. Most of the market is still at Level 2 or 3, which is exactly why the window is open now.

These are not aspirational characteristics. They are the operational outcomes Proxuma is built to deliver, inside the Autotask environment that most competitive MSPs already run.

What It Takes, Concretely

Proxuma is the layer that makes this real on the Autotask an MSP already runs. Planning turns hours of daily dispatch into minutes and shows capacity a quarter ahead. Shifts makes NOC and helpdesk coverage reliable instead of improvised, balanced by real workload rather than by who is easiest to ask. Intelligence turns scattered systems into one model with 500+ measures and the clean, AI-ready foundation everything above depends on. It does all of this without replacing Autotask, without a multi-month implementation, and without the data-engineering team that building it yourself would demand. That is the real choice facing an MSP: not whether you could assemble something like this, but how long it would take, what it would cost, and who would keep it alive while you are trying to run the business.

Proxuma was built by Dxfferent, a team of Autotask and MSP specialists who advise and run managed-service operations every day. Choosing it means more than software. It means practitioners who have solved these problems before, who built the tool because they kept seeing the same ones, and who get you live in 24 hours instead of handing you a manual.

> **THE PROXUMA PROMISE**
>
> Your best engineers did not get into this work to reconstruct timesheets, hunt for the right ticket, or rebuild the same board report every quarter. The data to run a better MSP already exists inside your tools. What is missing is the layer that turns it into something a person can act on, and trust. Give your team that layer and the scarce, expensive talent you fought to hire spends its hours on clients instead of operational friction, while AI finally has a foundation solid enough to build on. The engineer stays the expert. Operational intelligence is what lets that expertise reach the whole business at once.

AI thinks. Automation acts. Data decides. The MSPs that win the next decade are the ones getting that data right now, while it still sets them apart.

START TODAY

Ready to Become an Intelligent MSP?

The AI era rewards the MSPs that got their data right first. Start with the operational layer that makes it possible: full platform access, live onboarding from the Dxfferent team, live within 24 hours.

Start 14-Day Free Trial | Schedule a Live Demo

- ✓ 14-day free trial, full access
- ✓ Full platform access: Planning, Shifts, and Intelligence
- ✓ Live onboarding from the Dxfferent team, no onboarding costs
- ✓ No lengthy implementation: live within 24 hours

Need deeper Autotask consulting support?

The Dxfferent team offers free intake calls for MSPs looking to optimize their Autotask environment, data architecture, and operational workflows, beyond the Proxuma platform. If your Autotask challenges run deeper than scheduling and reporting, speak directly with a Dxfferent specialist.

Book a Free Dxfferent Intake Call

Proxuma.io

Proxuma is an operational intelligence platform built by the Autotask and MSP specialists at Dxfferent, with Autotask as its first integration. Trusted by 150+ managed service providers worldwide, Proxuma delivers algorithmic scheduling, coverage automation, and Power BI–powered analytics through three purpose-built modules: Planning, Shifts, and Intelligence.

Built on a 500+ measure library and an AI-ready data architecture, Proxuma transforms Autotask from a workflow tool into a complete operational intelligence platform, enabling MSPs to schedule smarter, cover shifts without conflict, and see their business clearly in real time. Autotask is where Proxuma goes deepest today, chosen because it is the most widely deployed PSA in the channel and the heart of the Dxfferent team's expertise. It is also the first of a connected-system library that already spans 12+ platforms and keeps growing, on a normalized model built to extend well beyond a single PSA.

Planning | Shifts | Intelligence | AI-Ready | Autotask Native

Dxfferent

Dxfferent is the specialist MSP consulting and technology firm behind Proxuma. With deep expertise in Autotask PSA configuration, data architecture, and managed service operations, the Dxfferent team combines software with practitioner experience.

Dxfferent provides free intake consultation for MSPs seeking to optimize their Autotask environment, improve operational data quality, or assess their readiness for AI-augmented service delivery.

Learn more at dxfferent.com →

Sources & References

This white paper draws on industry benchmark reports and publicly available research from the following sources. Figures are drawn from the most recent available edition of each source as of this 2026 edition; first-party Proxuma and Dxfferent figures are noted as such.

01 Grand View Research
Managed Services Market Size, Share & Trends Analysis Report, 2025–2033. Global market: $401.15B (2025), projected $847.41B (2033) at a 9.9% CAGR (2026–2033).
grandviewresearch.com

02 Grand View Research
U.S. Managed Services Market Size, Industry Report. U.S. market: $88.13B (2024), projected $183.47B (2030) at a 13.3% CAGR (2025–2030); the U.S. is roughly 26% of the global market.
grandviewresearch.com

03 Menlo Ventures
2025: The State of Generative AI in the Enterprise. Enterprise AI investment reached $37B in 2025, roughly triple the $11.5B invested in 2024.
menlovc.com

04 MuleSoft (Salesforce)
2025 Connectivity Benchmark Report (survey of 1,050 IT leaders). 95% say integration challenges hinder their ability to implement AI; only 29% of enterprise applications are integrated, with an average of 897 apps per organization.
mulesoft.com

05 RAND Corporation & Gartner
RAND, Why AI Projects Fail and How They Can Succeed (2025): by some estimates, more than 80% of AI projects fail. Gartner: organizations will abandon 60% of AI projects that lack AI-ready data through 2026.
rand.org

06 Kaseya
2026 State of the MSP Report (1,000+ MSPs). 53% use AI to automate ticketing, patching, and monitoring; 13% generate meaningful AI revenue; 48% rank AI and automation as the #1 client need for 2026.
kaseya.com

07 ISC2
2024 Cybersecurity Workforce Study: global workforce gap of approximately 4.8 million professionals, the last year ISC2 published a headcount gap. The 2025 study shifted focus from headcount to skills.
isc2.org

08 JumpCloud
SME IT Trends Report (Q3 2024). 49% of IT teams say they lack the staffing to secure their organization; security ranks as the top IT challenge.
jumpcloud.com

09 SPI Research
Professional Services Maturity Benchmark (2025/2026). Measured revenue leakage of about 4.5% of revenue, against a target below 5%.
spiresearch.com

10 Canalys / Omdia
MSP M&A Tracker and MSP Market Update. 169 publicly-announced MSP deals in 2025, with private equity involved in roughly 70%.
omdia.tech.informa.com

11 EY
How data readiness can enhance private equity exit value. Clean, well-structured operational data as a determinant of M&A valuation and due-diligence speed.
ey.com

12 Datto / Kaseya
Cooper Copilot documentation for Autotask. Native AI assistant features: smart ticket summaries, writing assistant, resolution summaries, and ticket triage.
datto.com

13 Microsoft Learn
Power BI MCP server documentation and Copilot for Power BI overview (public preview / general availability). Natural-language query over Power BI semantic models.
learn.microsoft.com

14 Aventis Advisors
MSP Valuation Multiples (analysis of 120 disclosed transactions). Median private MSP deal at roughly 8.9x EV/EBITDA, with platform-scale assets reaching the low teens.
aventis-advisors.com

15 European Commission & EUR-Lex
NIS2 Directive (EU) 2022/2555. MSPs and MSSPs explicitly in scope; management-body liability; 24h/72h incident reporting; supply-chain security; fines up to €10M or 2% of global turnover (essential entities). Transposition from October 2024.
digital-strategy.ec.europa.eu

16 KeyBanc Capital Markets & Sapphire Ventures
Annual SaaS Survey. Winning a dollar of new-customer revenue costs roughly four times what it costs to expand an existing account; expansion payback is far shorter than new-logo payback.
sapphireventures.com

17 SaaS Capital
Private SaaS Growth Rate Benchmarks. Net revenue retention drives growth; the highest-retention cohort grows at nearly double the median rate, and lifting NRR a band adds several points of growth.
saas-capital.com

18 Proxuma.io & Dxfferent (first-party)
Proxuma product documentation: Planning, Shifts, and Intelligence capabilities; 500+ measure library; AI-ready data model. Dxfferent operational research: per-ticket scheduling reduced from about 5 minutes to about 30 seconds with Proxuma Planning.
proxuma.io

proxuma.

Ready to Become an Intelligent MSP?

The AI era rewards the MSPs that got their data right first. Start with the operational layer that makes it possible: full platform access, live onboarding from the Dxfferent team, live within 24 hours.

Start 14-Day Free Trial

Schedule a Live Demo

proxuma.io · dxfferent.com

Get the whitepaper online:

proxuma.io/intelligent-msp-unlocking-ai-driven-growth-through-operational-excellence

© 2026 Proxuma.io & Dxfferent. All rights reserved. · Published 2026 Edition · hello@proxuma.io