

Travtus+

From Data Pulls to Data Operability

Why interoperability determines
what your stack can actually achieve



The housing industry has spent the last five years solving the reporting problem. Data has been extracted, consolidated and made visible across portfolios. That work was necessary. It is not sufficient.

The gap between intent and infrastructure is not a product problem, it is a structural one. Systems are no longer expected to simply surface insight, but to generate decisions and act across cross-functional workflows in real time. This places a fundamentally different demand on infrastructure.

None of these workflows exist within a single system. They depend on data moving continuously across many. When that movement breaks, the operation breaks with it.

Interoperability is not an enhancement. It is the condition that determines whether this model can function at all. Without it, AI remains a reporting layer with better language. With a focus on interoperability, it becomes operational.

No major property management
platform scores above

6.5 out of 10

for openess.

— Thesis Driven 2026



01.

The Problem We Thought We Solved

For much of the past decade, the dominant conversation in housing technology centred on a single frustration: operators could not get their own data out. Property management systems were built in an era where owning the workflow often meant controlling the data. APIs were limited, exports were manual, and moving data between systems required either custom development or significant operational effort.

That frustration was legitimate, and it prompted real change. Pressure on legacy property management systems, cloud-native CRMs, and a generation of PropTech vendors comfortable with API-first architecture meant that, by the early 2020s, most sophisticated operators could at least extract their data, aggregate it into a warehouse, and build cross-functional dashboards that previously would have required months of consultancy work. Data freedom, in the narrow sense of access, began to look achievable.

The industry, however, treated this progress as a solution when it was only a precondition. The challenge now is what happens after.

\$12.9M

average annual cost of poor data quality per organisation.

— Gartner

For multifamily operators, the downstream cost in AI underperformance, reconciliation effort, and delayed decisions is likely higher.

Can systems act on data in real time? Can insights trigger workflows across platforms without human intervention? Can the stack respond to events such as a lease expiry, a maintenance escalation or a pricing signal as a coordinated system rather than a collection of tools?



I need to architect my data differently for AI to really leverage it. It's beyond the individual data suppliers. It's what happens after I have the data.

**JAMES HAMRICK, SVP INFORMATION TECHNOLOGY,
GABLES RESIDENTIAL**

This reflects a broader shift. The bottleneck has moved from data extraction to data operationalisation, the layer that sits between raw data and action.

AI-native operations represent a different model entirely. Not AI that surfaces a dashboard or flags an anomaly for a human to act on, but AI that works everyday without manual handoffs at each step. That model does not work on a stack built for reporting. It requires systems that can share context, pass instructions and respond to events in real time across the entire operation. The infrastructure question is not a technology consideration. It is the precondition for whether this model is achievable at all.

There is no agreement yet on how to close that gap. The advantage will go to those who recognize the problem and solve it, rather than layering piecemeal AI applications.

02.

Integration vs. Interoperability: Why the Distinction Matters

These two words are used interchangeably throughout the industry. When there is a step change, the vocabulary needs to change with it. They describe fundamentally different capabilities, and that confusion continues to drive misplaced investment.

	Integration (SAAS Ecosystem)	Interoperability (AI Native Ecosystem)
What it does	Moves data between two specific systems	Enables behaviour across an entire stack
Direction	Typically one-way or scheduled	Bidirectional, real-time
Scale	10 systems = Up to 45 connections	One orchestration layer connects all, easy to scale
Error handling	Limited	Robust, with feedback loops
AI readiness	Cannot support agentic workflows	Designed for AI agent-based actions
Fragility	Breaks when commercial relationships shift	Resilient to vendor changes
Flexibility	Changes require rework across integrations	Systems can evolve without re-architecting the stack
When it matters	Works for reporting and isolated workflows	Required for real-time operations, cross-system execution and AI

Consider a single operational task. Qualifying a prospect, resolving a maintenance issue, or managing a renewal risk requires coordination across multiple systems. Availability from the PMS, pricing from revenue management, resident history from the CRM, workflow execution in maintenance systems, and logging outcomes. And all this needs to inform the investment decisions for the future. A chain of integrations cannot reliably support this. An interoperable stack can.

This is where vertical AI platforms solve the problem. The Model Context Protocol (MCP), introduced in 2024 and rapidly adopted across the AI ecosystem, is designed to enable systems to share context and execute actions across environments. It reflects a broader shift towards architectures where systems are expected to operate together, not just exchange data. But in order to adopt these standards, technology must be built for it. Data extraction and reporting can not engage in this use case.

For housing operators, the implication is not theoretical. Many PropTech solutions were built for a world where integrations were sufficient. In an environment where systems are expected to act across workflows, those assumptions begin to break down.

This is why many solutions fail to scale. They solve for sharing, but not for coordination. As a result, there is a growing shift towards more general-purpose enterprise infrastructure, but what is missing is a clear platform layer that can orchestrate workflows across systems.

The question is no longer what a system does in isolation, but how it behaves as part of a stack that needs to operate as a whole.

Operational complexity is increasing. Connectivity determines what can scale.



03.

The PropTech Trap

The housing industry has spent a decade assembling best-of-breed technology stacks. There are now specialist tools covering virtually every workflow: leasing automation, revenue management, maintenance coordination, resident communications, market analytics and fraud detection, among many others. Some of these tools are genuinely excellent at their specific function. The problem is not the quality of the individual parts, but what happens when they are assembled together.

Point solutions are structurally optimised to own a workflow, not to share one. Their business models depend on becoming embedded in operational processes, and their data architectures reflect that priority. Data is retained as a competitive asset rather than shared as an organisational resource. Integration access, the APIs and connectors that allow other systems to communicate with them, is controlled, often monetised and subject to change when commercial relationships shift.

This dynamic is not only technical. It is structural to the PropTech market. Many vendors operate in a constrained market, backed by venture capital that prioritises growth and near-term returns. Products are optimised for rapid adoption and ownership of specific workflows, rather than long-term interoperability across a broader ecosystem.

This creates a misalignment. Operators need systems that can evolve together over time. Vendors are often incentivised to optimise for speed, differentiation and retention within a single product boundary.

This tension was manageable when integrations were primarily about data transfer. It becomes more problematic when they are also the mechanism through which workflows operate.

The commercial tension between Funnel Leasing and EliseAI earlier this year illustrates how quickly bilateral dependencies can become operational liabilities. Whatever the merits of either position, operators caught in the middle had no architectural fallback. Similar issues became public with the Entrata and Yardi lawsuit a decade ago. That is the structural risk. It is about what happens when any two vendors disagree and your workflows sit between them.



We maintain a deliberately flexible vendor strategy. Our agreements are limited to one-year terms, and we operate exclusively on providers' standard APIs with no customization.

We integrate their data as-is, ensuring we can pivot to an alternative partner quickly if performance or alignment falls short.

PAT KLEIN, SENIOR EXECUTIVE AND FORMER CTO SPEEDBAY, ESSEX, CAO QUIDELORTHO

The discipline Pat describes, API-only contracts, no customisation, short terms, minimal dependency, is the rational response to a market structure defined by vendor lock-in, where relationships cannot be assumed to remain stable. It is effective as a defensive strategy, but it does not solve the underlying problem. It manages dependency rather than removing it.

Avoiding lock-in depends on more than contractual discipline. It requires a procurement process that demands flexibility and interoperability from the outset. Without a platform layer that standardises how data moves and how

workflows are orchestrated across systems, operators remain dependent on the same fragile integration model, regardless of how contracts are structured.

50%

**of AI agents currently operate in isolated silos,
unable to share context with agents in adjacent systems.**

— Mulesoft Connectivity Benchmark, 2026

Even when everything appears to work, a deeper problem remains. A well-integrated best-of-breed stack still cannot provide the cross-functional intelligence that enterprise AI requires.

Each system sees only its own domain. Leasing understands prospect conversations. Maintenance tracks work orders. Revenue management models pricing. None of them independently understands what is happening across the full lifecycle of a resident.

The insight that matters lives between systems.

04.

The Second Trap: General Purpose AI Is Not Enough

The industry is not starting from zero. Across sectors, organisations are already adopting general purpose AI infrastructure. Foundation models, cloud platforms and enterprise AI tooling have become the default starting point.

In housing, this often shows up through tools like Claude, OpenAI or Microsoft's AI stack. These systems are powerful, flexible and improving rapidly. They solve for capability, but they do not solve for coordination. Many larger enterprises have been forced into building their own inhouse technology teams. Often initiated as a data aggregation team, the closed nature of the industry solutions have forced much heavier internal tech spend.

For every dollar spent per unit on a closed proptech solution, there is an additional two dollars spent on derisking from lock-in and lack of data. This cost will only increase when wanting to expand the role of the internal teams to in fact solve for interoperability. Every closed vendor in an AI native expectation costs the business exponentially more in internal tooling, talent and compute.

This is the second trap. The first was the proliferation of point solutions. The second is the assumption that general purpose AI can unify them. It cannot. What general purpose infrastructure provides is a foundational capability but it needs infrastructure to productionise it for enterprise level usage and governance.

What is missing is a layer designed for how housing actually operates. An AI infrastructure for the housing industry. This is not a new problem. Other industries have already solved it.

05.

What Other Industries Have Figured Out

The structural challenge facing housing operators isn't new. The same pattern of fragmented legacy systems, proliferating point solutions and AI requiring cross-system coordination has already played out in other industries. In each case, the transformative value came not from individual applications, but from an orchestration layer above the existing stack, the AI platform.



The Legal Sector: **Harvey**

The leading AI platform for the legal industry, Harvey, did not displace the software that law firms already used. It sat above it, routing tasks across document management systems, legal research databases, and productivity tools simultaneously. More than 100,000 lawyers across 1,300 organisations now use the platform.



Healthcare: **Abridge AI**

The sector's equivalent of the property management system, the Electronic Health Record, is deeply embedded, largely closed, and controlled by a small number of dominant vendors. Abridge AI did not attempt to replace these systems. It integrates into clinical workflows, converting patient-physician conversations into structured documentation within the EHR. It is now deployed across more than 200 health systems within HIPAA constraints.



Defense and Intelligence: **Palantir**

Palantir's AIP platform did not replace existing enterprise systems. It sits above them, creating a semantic model that connects data, logic, and actions across the organisation. Its U.S. commercial revenue grew 71% year-over-year in Q1 2025.

The pattern is consistent across all three. The companies that captured disproportionate value did not compete with the existing software stack or even internal teams. They orchestrated it and supported industry ambitions.

06.

Building the Orchestration Layer

The need for an orchestration layer is no longer theoretical. The question is how it is implemented. The build versus buy conversation has always been on the table for large enterprises. The scaling cost by unit count for license fees on static software often built the case for an internal build. The case for ongoing maintenance and upgrades required dedicated teams.

However, with AI orchestration, the data is always moving and agents are constantly computing. The optimisation of the platform is essential for budgets to remain predictable. So, while in-house teams can replace single purpose use cases and applications with greater ease than before, there is a challenge of staffing and underwriting a full AI platform. This needs sophisticated understanding of model selections, self hosting and row level security which are technical concepts that are often underestimated in the industry.

In the absence of an industry platform, there is no choice but to take on this responsibility. However, if a platform exists for orchestration of common uses and intelligence, a partnership model can bring the benefit of “build and buy”. A platform foundation handles connectivity, orchestration and infrastructure, while internal teams build the business logic that differentiates the operation.

Adopting a platform provides speed and proven infrastructure, but requires careful evaluation to avoid replicating the same lock-in dynamics the industry is trying to move beyond.

The distinction is not simply build versus buy. It is where the complexity sits and who owns it over time.

07.

The Action Plan for C-Suite

Most operators are not starting from a blank slate. They have a PMS they cannot easily replace, a leasing stack that is three or four integrations deep, and AI pilots running in isolated workflows that have not connected to each other. The question is not what the ideal architecture looks like. It is how to move toward it from where you are.

01

Move procurement from departmental to enterprise

The fragmented stacks most operators are running today are not the result of bad technology decisions. They are the result of good departmental ones. Leasing bought for leasing. Maintenance bought for maintenance. Each tool solved the problem in front of it without accountability for what it could not connect to.

AI changes the accountability structure. When workflows are expected to operate across systems, a tool that cannot participate in that model is not a departmental problem. It is an enterprise one. Procurement decisions that were previously made at the functional level now need to carry an architectural sign-off. The question is not only whether a solution solves the immediate problem, but whether it can operate as part of a stack that needs to function as a whole.

Audit dependency before adding capability

Before the next technology purchase, map where your data actually lives and who controls access to it. Identify which vendor relationships are load-bearing, where a commercial shift or an API change would break an operational workflow. That audit will tell you more about your AI readiness than any RFP process.

Once that picture is clear, operators face a decision about how to build toward the orchestration layer. That decision should follow the audit, not precede it.



AI hasn't taken hold the way many expected. What the industry has built so far is a shelf of one-offs, valuable in isolation, but difficult to scale or expand.

ELIZABETH BRAMAN, CEO, REVOLUTION RE

Build the foundational layer on open infrastructure

The foundational layer of the stack should sit on platforms whose openness is structural rather than strategic. Microsoft Azure and Fabric, MS365, Twilio and their equivalents are the enterprise infrastructure that allow for interoperability. They have established ecosystems, documented APIs and no incentive to restrict how data is used. The data warehouse, communication infrastructure and identity layer belong here.

Industry-specific platforms then connect to this foundation as a purpose-built layer, extending what enterprise infrastructure cannot do on its own.

04

Choose an industry AI platform or build one

The question is no longer whether this layer will exist. It is who will build it, and how operators should evaluate it.

The most durable architecture is one where your own data infrastructure sits at the centre and every vendor connects to it as a spoke. This does not require replacing your PMS. It requires ensuring that data flows out of it on your terms, into infrastructure you own. This only works if the operator remains the hub.

A vendor that cannot deliver clean data and receive instructions as part of a wider system is a dependency with an exit cost, not a long-term partner.

The orchestration layer is not a tool you buy once. It is the AI infrastructure for the housing industry, compounding over time as workflows mature and operational logic accumulates. Operators face a clear choice: build this capability internally or work with a platform designed to provide it. What is not viable is continuing to add point solutions and assuming integration partnerships will hold.

The operators who move first on architecture will not just run more efficient operations. They will be the ones with the data infrastructure in place when AI capabilities make the gap between them and everyone else difficult to close. The defining criterion is interoperability: whether the platform can operate across systems, not within one.

The housing industry is at a similar point in its evolution to legal, healthcare and financial services when those sectors established the orchestration layers that now define them.

In each case, the operators who recognised the shift early and made decisions accordingly were the ones who captured the value. These industries chose to partner with an AI Platform partner but audit them for interoperability.

The same choice is now in front of housing.



Interested in AI platforms for the industry?

Visit www.travtus.com or email contact@travtus.com

