Companion Animal Veterinary Software Part IV:

PIMS in the Age of AI: Weather the Storm or Wither?

Navigating Practice Challenges with Support of Technology and AI

February 10, 2026

Jon Ayers, Jeff Dixon, Adam Little, Adam Wysocki
With Special Guest Robert Sanchez

TL;DR:

- **A PIMS is two things: databases and an application layer.** Not all PIMS databases are equal. Some are systems of record (appointments, billing, client records) where a single authoritative source is essential; others are contributory databases (medical observations, radiology, communications) where distributed data across multiple sources is natural and functional. AI will replace or evolve the application layer, but a system-of-record function persists.
- **The PIMS's enduring role may be governance, not workflows.** As AI-native applications handle the workflows practitioners actually use, the PIMS pivots to managing permissions, audit trails, compliance, and conflict resolution. The moat is no longer the UI; it is trust. Winners will be radically open to integration and offer agent-safe execution boundaries.
- **Incumbents face pressure from every direction.** Horizontal AI platforms may understand practitioner workflows better than the PIMS itself; enterprise groups will build unified platforms independent of any single PIMS; and the legacy data moat is weaker than assumed. Whether incumbents can pivot from platform to infrastructure remains an open question.

What becomes of the Veterinary PIMS in the world of AI?

Bottom line: its role in containing and managing systems of record may not go away, and it may pivot to serve as a governance layer for those systems.

But the application layer workflows will either be replaced by AI-native applications or evolve with AI as part of the PIMS offering. Traditional workflow screens stop being "the product" and become, at best, a thin fallback layer behind agent-driven workflows.

The moat is no longer the UI; it is trust.

To understand why, we must consider the roles of a PIMS, pre-AI, in serving a practice.

There are, of course, multiple possible futures for the PIMS.

- Existing systems of record may become irrelevant as customers create and curate their own;
- System-of-record companies may successfully build orchestration and action layers on top of their data assets;
- Personalized AI assistants and horizontal solutions may encroach on enterprise workflows from the consumer side; or
- Entirely new companies may develop more robust alternatives.

These scenarios are not mutually exclusive. However, legal and regulatory constraints often reinforce the role of the system of record even as the application layer evolves, which is why the governance function deserves particular attention.

I. PIMS, Databases (Systems of Record vs. Contributory Databases) and Application Workflows

What is a PIMS once dissected into its constituent parts? At its most fundamental level, a PIMS can be characterized in two components:

First, it is a set of databases. Some of these databases are *systems of record*, which are authoritative, singular sources of truth, while others are *contributory databases* that contribute to a collective description of the items they contain, alongside data that may reside in other systems outside the PIMS.

Second, it is an application layer that sits on top of those databases. This application layer consists of a set of workflows that do one of two things: they either present data out of the databases for decision-making (such as a report, an analysis, or a patient's medical history assembled as part of a SOAP process) or they update one or more of the databases within the PIMS, such as booking an appointment, recording an exam finding, or processing a payment.

This two-layer framework (databases below, application workflows above) also provides the foundation for understanding the integration challenge. Third-party applications are, in essence, additional application layers that need to read from and (in many cases) write to the same databases that the PIMS's own workflows use. The nature of that access depends on what kind of database they are touching, and that is where the distinction between systems of record and contributory databases becomes critical.

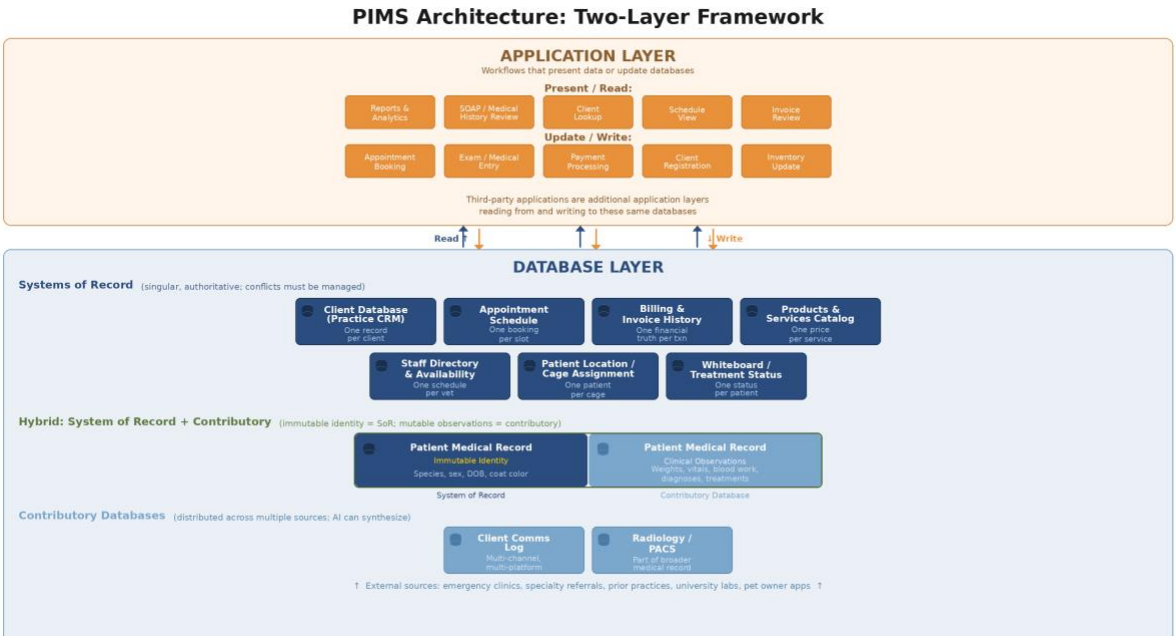

Figure 1: PIMS Two-Layer Architecture. Application workflows above, databases below

A PIMS holds multiple databases, but not all of them serve the same structural role. The critical distinction is whether a given database must function as a ***system of record***: a single, authoritative source where only one entry can exist for a given entity, and where duplicate or conflicting entries across multiple databases would be dysfunctional.

Consider the appointment calendar. Each time slot can hold exactly one appointment. When a slot is booked, it is consumed. If two systems, say an online booking application and a receptionist on the phone, attempt to book the same slot simultaneously, a conflict must be detected and resolved.

The system of record is the single arbiter of truth: it determines who got the slot and rejects the duplicate. The same logic applies to the client database (one record per client), billing (one financial truth per transaction), the products and pricing catalog (authoritative pricing rules and service definitions), staff availability (one schedule per veterinarian), cage assignments (authoritative current location assignment, with history), and the treatment whiteboard (authoritative current care-plan and task state, with audit trail and timestamps). In each case, duplication or conflict across databases would create operational dysfunction: double bookings, billing errors, misallocated patients, or dangerous clinical mistakes.

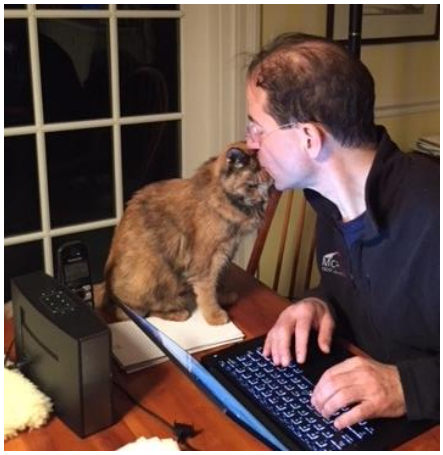

The patient medical record is not purely one or the other; it is a hybrid. Consider a cat named Reese, owned by Mr. Ayers. Certain attributes of Reese are

"immutable"[1] she is a feline, female, with a specific date of birth (from which her age is derived), and she has a particular coat/color (tortoiseshell/"tortie"). These identity attributes cannot change, and there can be only one authoritative entry. If two databases disagreed about whether Reese is a cat or a dog, that would be dysfunctional. For these stable identity attributes (which should change only via correction, with audit trail), the EHR functions as a system of record.

However, the vast majority of the EHR consists of mutable clinical observations that accumulate over time: Prince's weight at each visit, his blood pressure readings, blood work analyte values, exam findings, diagnoses, treatment notes, and prescribed therapeutics. These observations are not exclusive to one database.

Prince's medical history may sapn the practice's own EHR, a separate PACS for radiology, emergency clinic records from an after-hours visit, specialty referral notes, university lab results, or records from a prior practice or a low-cost spay/neuter clinic visited before adoption.

The EHR and PACS within the same practice already illustrate this reality: the medical record is spread across two databases by design. Having multiple databases contribute to Prince's clinical picture is not dysfunctional; it is the natural state of veterinary medicine. AI is well-suited to synthesize across these distributed sources to assemble the complete clinical picture needed for a case assessment. For these mutable observations, the EHR is a contributory database, not a system of record.

The same contributory logic applies to the client communications log, where messages legitimately originate from and are stored across multiple platforms, including the PIMS, a texting service, an email marketing tool, and a phone system. The same applies to radiology, where images may reside in the practice's PIMS, PACS, referral institutions, or cloud-based AI radiology platforms.

This distinction has significant implications for the integration debate. Applications that need to interact with a system of record (booking an appointment, updating a client record, processing a payment, or establishing a pet's immutable identity) require tightly managed, conflict-aware API access to the PIMS. Open APIs are the baseline; governance is the safety layer: scoped permissions, idempotent writes,[2] auditability, and practice-controlled consent, *and not* vendor-controlled lock-in. Applications that consume or contribute to clinical observations have a fundamentally different integration profile: they need to read from and write to the medical record, but the existence of parallel data sources is a feature, not a flaw.

II. Governance: What's Left?

[1] "immutable" is not precisely accurate. When we use this term here, we mean more precisely "stable identity attributes that should change only via correction with audit trail."

[2] "Idempotent" in computing refers to an operation or function that produces the same result regardless of whether it is applied once or multiple times. This means repeated executions do not change the state of the system beyond the initial application, ensuring safety, reliability, and consistency.

If AI-native applications increasingly replace the PIMS application layer, and if contributory databases like the medical record can legitimately be distributed across multiple sources, what enduring role does the PIMS play? The answer may lie in governance.

There is a version of this story where systems of record survive by pivoting. The PIMS becomes the "governance layer": the place where permissions are defined, audit trails are kept, and compliance is enforced. Multi-user synchronization is still hard. Permissioning is still hard. Knowing who changed what, and when, still matters. When a receptionist and an online booking app compete for the same appointment slot, something has to arbitrate. When a veterinarian updates a patient's record while a technician is entering lab results, something has to manage that concurrency. These are governance functions, and they are tightly bound to the systems of record identified above.

The winners in this model will be radically open to integration and offer a practice-configurable rules engine that other tools and agents can safely leverage. The real moat is not CRUD[3] screens, but trust: identity and permissions, consent, audit trails, data integrity, and an "agent-safe" execution boundary: least-privilege access, approvals, rollback, and visibility into what changed.

Ironically, the more intriguing questions around context orchestration and verification become easier to address with a strong system of record, precisely because these systems already capture essential changes, permissions, and logs. The governance layer is not merely a fallback; it is the natural foundation for the agentic workflows that are replacing the traditional UI.

But a set of databases with a governance layer is arguably a smaller business than being the platform, and more easily commoditized. The PIMS of today is architected to be the center of the universe: the hub through which all practice data flows, the application that veterinarians and staff interact with all day long. A database with governance is something fundamentally different, less visible to the user. It is a compliance sidecar, quietly ensuring data integrity and access control while AI-native applications handle the workflows that practitioners actually see and use. In this world, differentiation shifts away from the application layer entirely and toward the rails that must remain reliable: openness, uptime, recoverability, permissions, communications, and compliance.

This dynamic is compounded for enterprise veterinary groups operating with fragmented PIMS deployments. Rather than attempting to integrate multiple PIMS vendors or swap out systems across hundreds of locations, these organizations have a strong incentive to build a unified orchestration and data layer that sits above fragmented PIMS deployments, gradually centralizing selected system-of-record functions where feasible, while treating each site PIMS as a transactional contributor during a transition period. The PIMS database, in this scenario, is further demoted from platform to data

[3] CRUD: Create, Read, Update, Delete. These are the four fundamental operations that formed the backbone of how we interact with persistent data in applications such as a PIMS.

contributor. There is already a precedent for this in radiology, both veterinary and human, called a Vendor Neutral Archive (VNA).[4]

This pressure will intensify as AI coding capabilities move from standalone tools (Lovable, Bolt, Cursor, Copilot, Codex) into the operating system itself. When anyone can spin up agents and lightweight interfaces with simple voice commands (much as web browsers went from standalone software to built-in components of the operating system), workflows will increasingly migrate into native platform layers: email, chat, marketing, and office suites, with agents coordinating across them. A PIMS that cannot plug into those layers cleanly and securely becomes progressively less usable, regardless of how "complete" its internal UI appears. Consumer behavior is a powerful catalyst for forced innovation; once capabilities become built into everyday personal and professional workflows, the systems that cannot participate get replaced fast. Netscape and Adobe Flash are instructive precedents.

The more provocative challenge may come not from veterinary-specific competitors but from horizontal AI platforms. What happens when a practitioner's browser-based AI assistant understands their PIMS usage patterns better than the PIMS vendor itself? In other words, when veterinary-specific tooling performs worse than general-purpose solutions like ChatGPT for the same clinical tasks?

As personal AI assistants grow more capable at understanding individual goals, roles, and motivations, the boundary between personal and professional technology stacks will blur. Practitioners will increasingly route around enterprise applications through unofficial "grey-use" of consumer AI tools, and companies will find it difficult to enforce adoption of their own systems in such cases. The question of whether companies like OpenAI or Anthropic will dominate both the personal and professional technology stack is not hypothetical; it is already unfolding.

---

[4] A *Vendor Neutral Archive* (VNA) is a key technology in healthcare that provides a centralized, standardized storage solution for medical images and associated data. VNAs are designed to be agnostic of the systems or vendors that generate the data, making them ideal for long-term storage and retrieval of medical images across various systems, including PACS systems. This helps ensure seamless data access and sharing across departments and even different healthcare organizations.
Differences Between VNAs and PACS: While Picture Archiving and Communication Systems (PACS Systems) are designed to store and manage medical images, VNAs offer a broader solution that focuses on standardizing and centralizing data storage.

This pattern is visible in how major AI platforms are already positioning.

OpenAI's Frontier offering,[5] for example, provides a clear visualization of the stack they are enabling: Frontier positions an agent execution and governance layer integrated with existing systems of record, adding shared context, identity/permissions, and auditable actions, while agents orchestrate workflows across those systems.

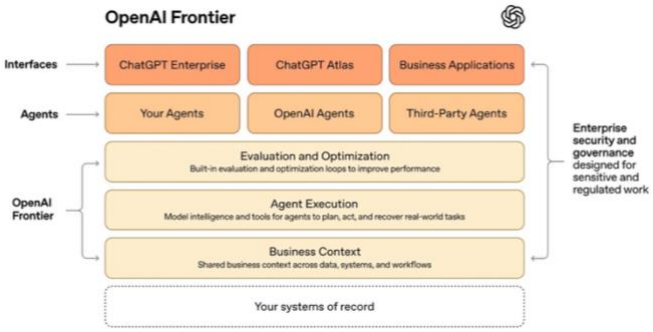

This is precisely the architecture this paper describes for the future of the PIMS, suggesting the trajectory is not speculative but already being built by the largest technology companies in the world.

It is not clear that the current PIMS incumbents are well-positioned for this pivot. Their entire architecture, their business model, and their relationship with practices all assume they are the center. Rebuilding as a governance layer would mean accepting a fundamentally different role: not the application veterinarians work in, but the infrastructure underneath the applications they choose to work in. It remains an open question as to whether incumbents can make that transition, or whether governance emerges from a new entrant or from the AI platforms themselves.

One further consideration weakens the incumbents' position: the data itself. Data quality and structure vary widely across practices, often shaped by clinic workflows and UI constraints, but the bigger failure is that practices can't easily access, move, and improve their own data. We must criticize the assumption that historical PIMS data represents an unassailable competitive moat for the PIMS vendor.[6] As AI-driven tools generate richer, more structured clinical data from diverse sources, and as distributed contributory databases grow in both volume and quality, the relative value of legacy data within the PIMS may diminish rather than appreciate.

Appendix

[5] According to OpenAI, Frontier is a new platform that helps enterprises build, deploy, and manage AI agents that can do real work. Frontier gives agents the same skills people need to succeed at work: shared context, onboarding, hands-on learning with feedback, and clear permissions and boundaries. That's how teams move beyond isolated use cases to AI coworkers that work across the business. https://openai.com/index/introducing-openai-frontier/

[6] In business strategy parlance, a "competitive moat" is a business's durable, structural advantage that protects it from competitors, allowing it to sustain high profits and market share over the long term, much like a castle's moat protects it from invaders. These advantages can stem from factors like strong brands, patents, cost advantages, *high customer switching costs*, or network effects, making it difficult for rivals to compete effectively.

The table below classifies each of 10 PIMS databases accordingly.

Database	System of Record?	Comments
Patient Medical Records (EHR)	Hybrid	The EHR contains two fundamentally different types of data. **First**, there are immutable attributes of the pet (species, sex, date of birth, coat color) that are definitional and cannot change. These attributes require a system of record; there can be only one authoritative answer to “is Prince a cat?” or “what is Prince’s date of birth?” Once characterized, it can be duplicated to other databases. **Second**, there is the accumulation of clinical observations over time: weights, blood pressures, exam findings, blood work, diagnoses, treatment notes. These mutable observations are not exclusive to one database. A patient’s clinical history may span the practice’s own EHR, a separate PACS, emergency clinic records, specialty referrals, and prior practice records. Having these observations distributed across multiple sources is not dysfunctional; AI can synthesize across them to present a unified clinical picture. For these reasons, the EHR is a hybrid: *system of record* for the pet’s immutable identity, *contributory database* for the pet’s evolving medical history.
Client Database (Practice CRM)	Yes	Each client must have a single, authoritative record within the practice. Duplicate or conflicting client entries (different phone numbers, misattributed pets, inconsistent contact details) create operational dysfunction in communications, billing, and scheduling. The practice needs one definitive record per client.
Appointment Schedule / Calendar	Yes	Each appointment slot can only hold one booking. When that slot is filled, it cannot be booked a second time. The system of record must manage concurrent booking attempts and resolve conflicts; a receptionist on the phone and an online booking app cannot both book the same slot. This conflict-management requirement is the hallmark of a system of record.
Billing & Invoice History	Yes	Financial transactions must be authoritative and singular. Each invoice, payment, and estimate must exist in one definitive record. Duplicate or

		conflicting billing entries would create accounting errors, compliance issues, and client disputes. There can be only one financial truth for each transaction.
Products & Services Catalog (with Pricing)	Yes	The practice must have one authoritative list of what it offers and at what price. If multiple databases held conflicting prices or service definitions, the practice would generate inconsistent estimates, confuse clients, and create billing discrepancies. One catalog governs what is sold and charged.
Client Communications Log	Sometimes, depends on the practice	Communications with clients may legitimately originate from and be stored across multiple platforms: the PIMS, a separate texting service, an email marketing tool, a phone system. Having records across multiple databases is not dysfunctional; it reflects the reality of multi-channel communication. These records can be aggregated for a complete view without requiring a single authoritative source. However, some practices are sensitive to regulation/legal liability pressures, and often force communication log entries to be consolidated into a single system of record, with a lot of cutting and pasting
Veterinarian / Staff Directory & Availability	Yes	Staff availability must be managed from a single authoritative source to prevent scheduling conflicts. If multiple systems held conflicting availability data, the practice could double-book a veterinarian or schedule patients when a doctor is unavailable. One definitive record of who is working and when is essential.
Radiology / PACS	No	Like the EHR, radiology images and reports are part of the broader medical record and can exist across multiple systems. A practice may have its own PACS, receive images from referral institutions, or use cloud-based AI radiology platforms. Multiple sources of radiographic data are not dysfunctional. In fact, the EHR and PACS together illustrate how the medical record is already spread across two databases by design.
Patient Location / Cage Assignment	Yes	A hospitalized patient can only physically be in one cage at one time. Conflicting records about patient location would create dangerous confusion in a hospital setting; the wrong patient could receive the wrong treatment. This database must be singular and authoritative.

Electronic Whiteboard / Treatment Status	Yes	A patient's current treatment status and progress within the hospital must be maintained in one authoritative view. Conflicting status records could lead to missed treatments, duplicated procedures, or dangerous clinical errors. The care team must work from a single source of truth for in-hospital patient management.

Note: These 10 may not be a complete list. For example, boarding and grooming functionalities are not considered in this PIMS analysis, although they are frequently offered as additional functional components of a PIMS.