

A Leader's Guide To Operational AI in Action




Table of Contents

01.	Cover Page	09.	AI-Assisted Humans
02.	Table of Contents	10.	Context Layer Case Study: Cummins
03.	What is Operational AI	11.	Agents
04.	The Building Blocks of Operational AI	12.	Conclusion
06.	A Maturity Framework for Building to Operational AI	13.	About ChapsVision
07.	Enterprise Search		
08.	Enterprise AI Search Case Study: Pfizer		

What is Operational AI?

There's a class of enterprise where the work you do carries real-world consequences: building an aircraft, protecting the financial well-being of millions, developing a molecule that saves lives. In these organizations, the most important decisions and work don't happen in email threads or slide decks. They operate on information inside specialized, complex, multimodal, often siloed systems of record: ERP, Product Lifecycle Management (PLM), Portfolio Management Systems (PMS), Laboratory Information Management Systems (LIMS), regulatory repositories.

The first wave of enterprise AI — chatbots, copilots, productivity assistants — was built for the office layer, not the operational one. For many teams, it delivered genuine value. But for mission-critical work, it never reached the decisions that matter most.

Operational AI is built for this type of mission-critical work. It connects to the systems where that work actually runs, retrieves complex, domain-specific knowledge those systems hold, and produces answers that can be trusted, defended, and acted on.

Trusted Operational AI is:

- **Grounded:** Every answer built from your own knowledge, traceable to the source.
- **Governed:** Every action, human or agent, permissioned and provable.
- **Strategic Autonomy:** Independence from a single LLM provider, independence from a locked-in hyperscaler, control over your data security, shielded from political decisions. You need control over your AI choices.

No matter where you are on your journey, these are the 3 things that need to be clear on Day 1.

When an answer that's almost right is still wrong, it can cost millions, endanger lives, put businesses and reputations at risk; in these cases productivity AI isn't enough. That's what Operational AI is built for.

The Building Blocks of Operational AI

Operational AI is a layered architecture. Four foundational components work together to make AI trustworthy in mission-critical environments.



01. Connection & Data Preparation


Before AI can reason about your business, it needs to reach your data. Enterprise knowledge and data can live across dozens of fragmented systems, in formats no one standardized: legacy CAD files, structured databases, decades of multimodal records. Connection means indexing and unifying those sources securely. Data Preparation means cleaning, transforming, and enriching that raw material, so AI outputs are accurate and defensible.




02. AI Context Layer

The Context Layer makes AI answers accurate, traceable, and trustworthy in the environments where being wrong has real consequences. It's the architecture that sits between your data and your AI, connecting sources, structuring knowledge, and ensuring every response is grounded in what your organization knows, not what a model guesses.

Retrieval-Augmented Generation (RAG) is how AI draws answers from your enterprise knowledge. By pairing Advanced RAG, combining multiple retrieval methods (vector, keyword, structured, multimodal), with ontologies that model how your business actually thinks — your parts, your programs, your terminology — simple document retrieval becomes genuine operational intelligence. Every answer is grounded, cited, and traceable.



The Building Blocks of Operational AI



03. Agentic Layer for Humans and Agents

This is where AI moves from answering to acting. An agentic operational surface lets human workers and AI agents collaborate on the same workflows, with agent orchestration managing multi-step tasks, and an agentic gateway connecting third-party AI tools. The agentic layer is where the work itself happens.



04. Governance Plane

As agents take on real work across real systems, visibility and control become non-negotiable. The governance plane spans the entire platform, enforcing data permissions for every human and agent, managing policies on what AI may and may not do, providing real-time observability across the agent fleet, and maintaining full audit trails so every action is traceable and explainable. This is how trust is engineered into an AI system.

A Maturity Framework for Building to Operational AI

Most organizations don't arrive at Operational AI overnight. In fact, a recent study by ChapsVision found only 10% of enterprises have deployed agentic AI in production. Agentic AI is a journey, and understanding where you are helps you invest in the right next step.

Stage 1: Enterprise Search Foundation


This is where most successful AI transformations begin. Before AI can reason about your business, you need to be able to find knowledge across the systems where it lives. Enterprise Search unifies fragmented sources into a single, queryable knowledge base, reducing time spent searching, improving answer quality, and establishing the connective tissue that all future AI depends on. It's immediately valuable and foundational to everything that follows.

Stage 2: AI-Assisted Humans

The next step is applying the knowledge retrieved through Enterprise Search and RAG architectures to Generative AI. With enterprise context, an AI assistant or AI subject matter expert can now answer complex, multi-part questions with cited, traceable answers. They can even execute prompted workflows to automate tasks. Humans remain at the helm: better informed, faster, and more confident in the decisions they make.

Stage 3: Agents and the Future

With the context layer proven in stages 1 and 2, autonomous agents become viable. Agents don't just answer questions; they reason and take action across your systems. Multi-agent workflows can have far-reaching impact. This is where agentic governance becomes critical, to ensure agents can be trusted with data security, model use, costs, and more. The key insight: the quality of your context layer determines the ceiling of your agents' value. Foundations first, autonomy second.



Enterprise Search: The First Stage of AI Transformation

Enterprise AI Search is the ability to ask a question and receive a precise, cited answer by connecting to all your organization's systems of record: PLM, ERP, SharePoint, quality databases, and more. Instead of matching keywords and returning a list of links, it understands your question, retrieves the most relevant knowledge, and synthesizes a grounded, cited response.

Three Example Use Cases

Certified Part Reuse

Match new part requirements against PLM, CAD, and qualification records to reuse an approved part instead of re-validating. Instead of navigating five systems and sending emails to three colleagues, find an answer in minutes, with citations to all documents and requirements.

Regulatory & Compliance Lookups

A quality professional needs to verify which certifications apply to an updated component. AI Search surfaces the applicable standards, prior audit findings, and any open compliance gaps, across regulatory repositories and internal quality records, in a single, traceable response.

Onboarding and Institutional Knowledge Transfer

A new engineer joins a team with decades of accumulated domain expertise spread across legacy systems. Enterprise AI Search gives them immediate, accurate access to that institutional knowledge, compressing months of tribal knowledge transfer into days.

Enterprise AI Search Case Study: Pfizer



Profile

Global leader in pharmaceuticals and biotechnology.



Challenge

Inability to reach critical knowledge and data fragmented across dozens of research, regulatory, medical affairs, legal, and security systems. Teams spent days on tasks that should take hours.



Solution

Sinequa by Argonos was selected to power a unified, permission-aware enterprise search platform. Sinequa was deployed across more than 12 production use cases.

Use Cases

- Medical Affairs: Surfacing scientific insights across fragmented sources, enabling faster access to medical knowledge and better decision support.
- Clinical/R&D: Searching and analyzing clinical data, accelerating access to trial information, and reducing manual effort.
- Regulatory/GxP: Accessing compliance and regulatory content, reducing compliance risk, and expediting audits.
- Cross-Enterprise Impact: Providing a unified knowledge layer and reusable AI foundation for consistent user experience.

Impact

- Reduced meta-analysis work from two days to an hour
- Saved approximately 9,000 biostatistician hours annually
- Improved decision quality
- Reduced compliance risk
- Elimination of inconsistencies that had created regulatory exposure

AI-Assisted Humans

Once your enterprise knowledge is connected and searchable, the next step is making AI a genuine thought partner for the people doing your most complex work. Powered by the AI Context Layer, AI Assistants and AI Subject Matter Experts move beyond finding information and begin helping employees reason through it, answering multi-part questions with cited evidence, surfacing relevant context across systems, and enabling faster, more confident decisions. Humans stay in control; AI makes them dramatically more effective.

Three Example Use Cases

Change Impact Analysis

Trace a change across EBOM, MBOM, test procedures and certification documents to catch downstream impacts before costly late rework is needed. With full source traceability, in a single query.

Regulatory Dossier Intelligence

Life sciences teams preparing submissions need to rapidly surface relevant clinical data, prior safety findings, and applicable regulatory guidance. The context layer connects internal trial data, regulatory repositories, and external literature into a unified, queryable knowledge base that supports fast, defensible decisions and GxP compliance.

Material Shortage Impact

Connect a delayed shipment to affected work orders and customer deliveries to see revenue-at-risk before builds slip.

Assistive AI Case Study: Cummins

Profile

Global powertrain and technology leader with more than 69,000 employees, including 15,000 engineers.

Challenge

Critical knowledge was distributed across multiple specialized systems — PLM tools, enterprise document repositories, warranty databases, and more. Engineers were spending over 10% of their time just searching for information. Generic AI tools like Copilot and ChatGPT lacked the deep enterprise context to be truly useful, and poor knowledge access meant limited AI to low-value use cases with no real path to agentic AI.

Solution

Sinequa by Argonos deployed as the foundation for Cummins' AI and digital knowledge initiatives. The platform connected Cummins' engineering systems, structured domain-specific knowledge using advanced RAG, and delivered context-aware, fully-cited answers across even the most technical product and design queries.

Impact

- Moved from POC to enterprise production deployment in ~12 weeks
- Multiple hours per engineer per week saved
- 5X improvement in engineering answer quality
- Highly attractive ROI delivered

Cummins now has the knowledge foundation needed to pursue truly agentic AI use cases.

Agentic AI: Capturing Outsized Value from Operational AI

Agentic AI is where the promise of AI becomes real for the enterprise. Instead of giving your teams better answers, it automates the complex, multi-step work. Compliance reviews that took days are done in minutes. Engineering change analyses that required five systems and three colleagues are now handled autonomously. The value isn't just speed; it's the ability to scale high-judgment work.

Why Context Is Key

Agents are only as good as what they know. An agent operating on incomplete, stale, or unstructured knowledge will produce unreliable results, and in regulated or mission-critical environments, that's not a rounding error. The context layer and great search isn't just a prerequisite to agentic AI; it's what determines how far your agents can be trusted to act.

Three Example Use Cases

Drawing Compliance Agent

Given a technical drawing, an agent identifies applicable standards, analyzes compliance across each one, flags non-conformances, and produces a complete compliance report with recommended resolutions — a process that previously took days, now done in minutes with full auditability.

Pharmacovigilance and Signal Detection

A governed, agentic surveillance fleet unifies case data, literature, and regulatory sources, surfacing signals automatically, and scaling safety team work efficiently with full auditability.

Test Plan Agent

Given a set of product requirements, an agent analyzes prior test efforts for similar situations, specifies the steps needed to meet each requirement, and flags unnecessary or redundant protocols, creating a thorough, efficient test plan grounded in the organization's actual testing history.

Conclusion

The organizations that will get the most from AI agents in the next five years are the ones investing in knowledge infrastructure right now. Not because agents aren't valuable — they are — but because the context layer determines the ceiling of what any agent can safely do.

The path to Operational AI is a journey, not a leap. Start with enterprise search: connect your systems, make institutional knowledge findable, and demonstrate immediate value. Extend to assistive AI: build on the context layer to ground AI responses and workflows in domain-specific knowledge so answers are accurate, cited, and defensible. Then expand to agents: give them the trusted knowledge foundation they need to act autonomously in the workflows where your organization is ready.

Across every step of this journey, trust is non-negotiable. AI must be grounded, governed, flexible, and maintain strategic autonomy. These requirements aren't constraints on AI value; they're the conditions that make value possible. An agent that can't be trusted can't be deployed. A knowledge foundation that can't be verified can't power decisions that matter.

The enterprises that will lead the future are not the ones chasing autonomy the fastest. They're the ones thoughtfully building the foundations for trusted, accurate, reliable agentic AI. That's Operational AI.

Start with search. Build that context layer. Scale to agents.

About ChapsVision

ChapsVision is a global leader in agentic AI and data solutions. Trusted by over **2000 customers** across **40 countries**, ChapsVision delivers industry-specific software solutions powered by agentic AI to meet the complex needs of data-intensive sectors. Its innovative and scalable Data & AI platform is powered by proven technology enablers that accelerate data acquisition, preparation, and analysis.

Driven by significant R&D investments in large-scale data and AI and supported by a targeted strategy of acquisitions and international expansion, ChapsVision has rapidly established itself as the trusted partner in agentic AI for business and government organizations around the world.

More information:

www.sinequa.com

www.chapsvision.com

