SoftServe

SoftServe Insights Platform

Unstructured content goes in.
Trusted, structured intelligence comes out.

FOCUS INDUSTRIES

- MANUFACTURING
- LIFE SCIENCES
- HEALTHCARE
- FINANCIAL SERVICES
- ENERGY
- CROSS-INDUSTRY

Organizations handling large volumes of documents know the problem already: critical knowledge is scattered across systems, manual review is the bottleneck, and standard AI tools generate answers no one trusts enough to act on. SoftServe Insights Platform is built to close that gap.

WHERE DOCUMENT INTELLIGENCE BREAKS DOWN

Most enterprises are not short of data. They are short of a way to surface what the data says accurately, traceably, and without assigning a team of specialists to read everything first.

Fragmented content, slow decisions
Critical knowledge sits across SharePoint, Confluence, ERP systems, and local drives. Teams spend hours searching and validating before they can act on anything.

Manual review at scale fails
Contracts, compliance records, vendor submissions, SOPs, and audit files require expert review that does not scale. The bottleneck is structural, and it compounds with every new document volume increase.

AI answers that no one fully trusts
Generic AI tools provide answers without grounding. No citations. No traceability. Results that are hard to audit cannot be used in regulated or high-stakes workflows.

AI stays stuck in pilots
Poor integration with existing security and compliance models keeps AI deployments isolated. Without governance, access control, and auditability, AI cannot move from proof of concept to production.

WHAT SOFTSERVE INSIGHTS PLATFORM DOES

SoftServe Insights Platform is an enterprise document intelligence platform that turns unstructured content into clear, structured, and evidence-backed insights. It delivers explainable results, integrates with existing business tools, and is built from the ground up around the governance requirements of regulated and operationally complex organizations.

It's not a chatbot. It's not a one-off implementation. It's a modular, production-grade platform designed to run across departments, use cases, and documents without being rebuilt for each one.

BUSINESS IMPACT

5x
faster time to value vs. bespoke development

5x+
lower CAPEX compared to custom build approaches

90 GB
of assets processed and indexed in approximately 25 minutes

✳ Performance figures are based on SoftServe pilot projects and benchmark data. Individual results depend on document volume, complexity, and deployment configuration.

HOW IT WORKS

Powered by NVIDIA's full technology stack and brought to life by SoftServe, this is your engine for enterprise transformation. We combine advanced AI models, accelerated computing, and high-fidelity simulation through the NVIDIA Omniverse platform to process enterprise signals at lightning speed. This allows you to model workflows, test scenarios, and predict outcomes accurately. It moves your AI initiatives out of isolated pilot phases and into powerful, repeatable production.

CORE CONTENT INTELLIGENCE

SIP processes documents, extracts structured data, compares against standards, and summarizes at scale. Every output is grounded, citations are built in, confidence is visible, and every answer traces back to the source text. The system works across text, tables, and images simultaneously, which are in a document.

ENTERPRISE READINESS

Custom PII filtering, role-based access control, and multi-tenant architecture are standard, not optional. The platform deploys on-premises, in the cloud, or in hybrid configurations. IaC setup, a custom evaluation framework, and API-first design mean integration follows existing security models rather than requiring organizations to adapt to the platform.

PLATFORM CAPABILITIES

CAPABILITY	WHAT IT DOES
Deep Search Conversational AI	Natural-language queries across all ingested documents. Responses are grounded in source content with citation management built in. Not a generic chatbot, every answer is traceable to the document and passage it came from.
Multimodal Document Summarization	Structured, template-driven summaries generated from documents containing text, tables, and images. Output formats are configurable for reporting, briefing, and workflow automation use cases.
Multimodal Document Comparison	Rule-based and AI-driven comparison of documents against standards, templates, or prior versions. Surfaces gaps, deviations, and changes precisely — not just text differences.
Custom Structured Data Extraction	Reasoning-driven extraction of defined fields from documents of the same type (invoices, purchase orders, contracts, etc.) using a textual field description rather than hard-coded schema. Output is CSV-ready for downstream analytics.
Grounded Responses and Evidence-First Reasoning	Every answer includes the source passage, document, and page that produced it. Accuracy and observability are built into the response of "Deep Search Conversational AI" module architecture, not added as an afterthought.
Custom PII Filtering	Automated detection and handling of personally identifiable information across all ingested content. Configurable per use case, department, and regulatory requirement.
Access Control Management	Role-based access control with fine-grained permissions at the Knowledge base (documents) level. Multi-tenant architecture supports multiple teams or clients working from isolated environments within a single platform instance.
Integration Connectors	ERP, SharePoint, Confluence, and Jira integration supported. API-first design allows connection to additional enterprise systems without rebuilding the platform.
Domain-Specific LLM Integration	Model selection and configuration per use case and domain. LLM monitoring and observability via LangFuse. On-premises and private cloud LLM deployment supported.
Modular and Agentic Fit	Each capability is a standalone module. The platform supports agent integration and multi-agent architectures, allowing SIP to operate as a component within larger agentic workflows.

WHERE ORGANIZATIONS USE IT

Executive Briefing Generation - Operational reports, board packs, and multi-source document sets are summarized into structured executive briefings automatically. The analyst who used to spend two days consolidating sources now reviews a ready output.

Compliance Readiness Assessment - Audit records, policy documents, and regulatory submissions are processed against compliance frameworks. Teams get a structured readiness report rather than a manual gap analysis.

Purchase Order Intelligence and Validation - Invoices and purchase orders are processed for structured data extraction and validated against reference data. Exceptions surface automatically. Manual reconciliation stops being the process.

Contract Review Against Standard - Contracts are compared against master templates or regulatory standards. Deviations, missing clauses, and non-standard language are flagged. Legal and procurement teams stop reading every page looking for the gaps.

Vendor Due Diligence Review - Vendor documentation across multiple submissions is compared, scored, and summarized at scale. Risk analysts work from a structured output rather than a reading stack.

On-Site SOP Assistant - Technical personnel access equipment manuals, maintenance SOPs, and troubleshooting guides through conversational queries. In one production deployment, query response time dropped from minutes to under 30 seconds, with an estimated 10x cost reduction per query versus employee lookup time.

SIP VS. THE ALTERNATIVES

The choice between SIP, bespoke development, and low-code/no-code platforms comes down to three variables: how fast the team needs to see results, how much governance the use case requires, and how much ownership the organization wants over its own AI infrastructure.

SIP VS. BESPOKE DEVELOPMENT

Bespoke code-first approaches provide maximum architectural freedom. SIP delivers comparable customization flexibility with proven architecture, guardrails, and a deployment roadmap that starts at month one rather than month four. Governance, observability, and scalability are built in rather than developed from scratch.

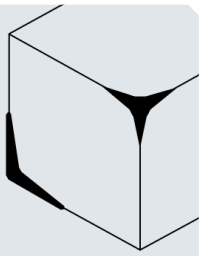

SIP VS. PRODUCTIVITY SOLUTIONS

Tools like Microsoft 365 Copilot are built for daily knowledge work. SIP is built for enterprise-level document operations: batch processing, repeatable workflows, auditability, grounded responses, and on-premises deployment. The use cases are different. So is the architecture.

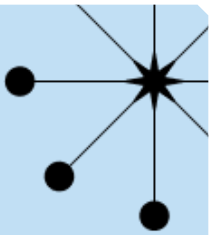

WHY SOFTSERVE

Production-grade from the start - SIP is built on reusable components from SoftServe's Insight Engine ecosystem - components that have already run in enterprise production environments, not components being stress-tested in a pilot for the first time.

Regulated domain depth - SoftServe has delivered large-scale AI systems in healthcare, life sciences, energy, manufacturing, and financial services. Governance, traceability, and security requirements in these sectors are not handled by adding a compliance layer. They are built into the architecture.

The platform grows with the use case - SIP is modular and API-first. New capabilities, connectors, and domain configurations extend the platform without rebuilding the core. Organizations that start with contract review can add compliance assessment, vendor due diligence, and SOP assistant without starting over.

End-to-end delivery -AI, ML, back-end engineering, and enterprise integration under one team. No handoffs between specialists who have never worked together on the same system.