



WE

Enable Enterprise Efficiency

Agentic AI. NLX. Orchestration.



Faster time-to-value.
Go live in weeks, not quarters.



Natural-language (voice/chat)
experiences for all stakeholders.



No persistent data store
Transient caching for performance



Secure by default
Insurance grade compliance by design

Prewired to InsuranceNow, pre-built to scale.

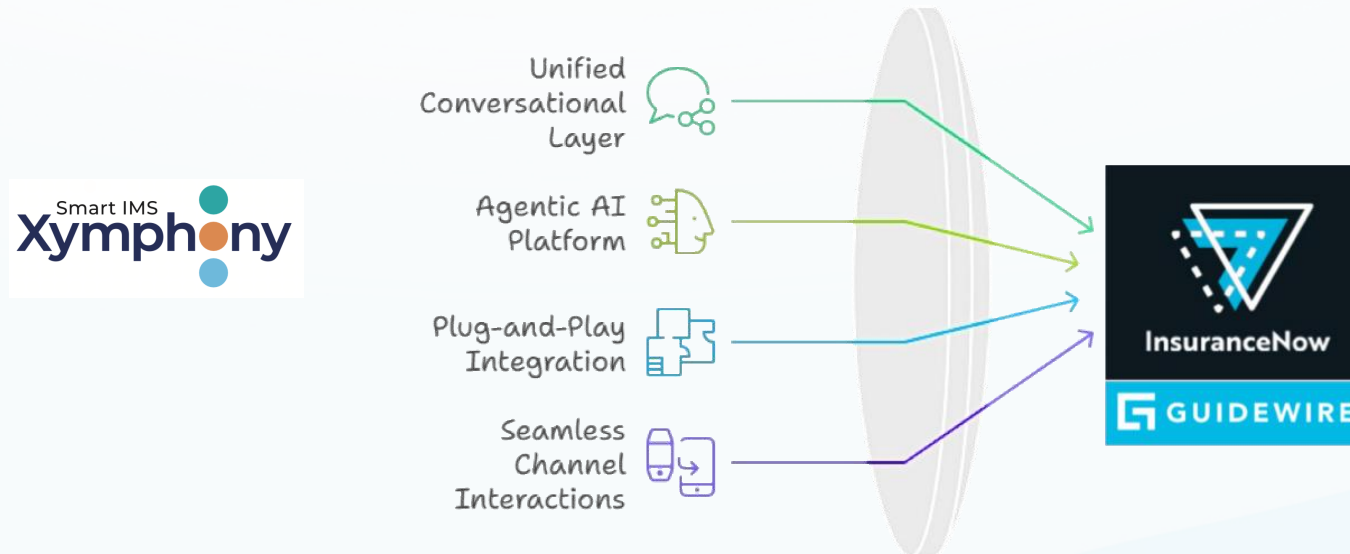
WHERE XYMPHONY FITS

Xymphony provides the intelligence + NLX layer

- No data duplication. Data stays in InsuranceNow.
- Pulls the right data when a user asks for it.
- Explains and highlights what needs attention.
- Respects user access rules and keeps an audit trail.
- Deep-links users back into InsuranceNow for action.

InsuranceNow stays the core platform

- System of record for policies, submissions, and claims.
- Runs the core processes and enforces rules.
- Remains the place where work gets done.



Xymphony complements InsuranceNow. It does not compete with core features.

WHAT XYMPHONY ENABLES

Natural Language Experience (NLX)


Ask in plain English (or speak it). Get answers you can act on.

- “Show me the policies issued yesterday.”
- “Which submissions were rejected in the last 7 days, and what were the top reasons?”
- “List the claims that are at risk of missing service targets this week.”
- “Why was claim 123 approved (or denied)? Summarize the key events and notes.”

Configurable assistants by role

Different jobs. Different questions. Same core data.

- Underwriting: pipeline overview, reject patterns, referrals, quick record summaries.
- Claims: claim status at a glance, reopen reasons, follow-ups, exception lists.
- Operations: backlogs, aging work, workload balance, items that need intervention.
- CIO / Executives: daily briefing, week-over-week changes, top exceptions, adoption signals.
- Start read-only. Expand capabilities as trust grows.



Interaction

Chat

Voice

API

AG-UI

A2A

Identity & Access Management,
Authorization Controls, RBAC

Agents



Observability

- Token Usage
- Latency
- Tracing
- Metrics
- Telemetry
- Audit Trails
- Reliability

Guardrails

- Hallucination
- Toxicity
- Jail breaking
- Code Exploits
- Factuality
- Formatting
- Etiquette
- PII / PHI

Models

- Public Models
- Hosted Models
- Finetuned

Evaluation

- Custom Datasets
- Knowledge Evals
- Prompt Playground
- Semantic Compare
- LLM as judge
- Relevancy Metrics
- Blue, Cosine, Rouge
- Sentiment Analysis
- Faithfulness

Knowledge

- Chunking
- Ingestion
- Retrieval
- Reranking
- Search
- Generation
- Agentic RAG
- Graph RAG
- Pipelines

AI Agents implemented on open protocols to ensure interoperability, scalability, and future-proof architecture



A2A

Agent-to-Agent Protocol

Enable seamless communication between autonomous agents with standardized message formats, discovery protocols, and handshake mechanisms.

- Inter-agent messaging
- Service discovery
- Protocol negotiation
- State synchronization



AG-UI

Agent-Generated UI

Allow agents to dynamically generate user interfaces based on context, user preferences, and task requirements for optimal interaction.

- Dynamic UI generation
- Context-aware layouts
- Adaptive components
- User preference learning



MCP

Model Context Protocol

Standardize how agents access and share context, memory, and knowledge across different AI models and systems.

- Context sharing
- Memory management
- Knowledge graphs
- Cross-model compatibility

**Built for real P&C operations. Open by design.
Ready when you are.**