



# AGENT INTELLIGENCE

How Vivun Models Expert Knowledge  
for **Scalable, Proactive** Work

**vivun.**

# Contents

Executive Summary	03
Introduction: The Promise of Agent Intelligence	04
Agent Intelligence	05
Modeling Expertise	06
From Model to Graph: Encoding Expertise as Agent DNA	07
Transplanting the Brain: Embedding Knowledge in the Agent	08
Memories: The Missing Link in AI Agent Design	09
Output That Matters: Introducing Accelerators	10
Interfaces and Modalities: Meeting Users Where They Work	11
The Mouth, Not the Brain: The Role of LLMs in Agent Intelligence	12
Use Case Deep Dives: Where Agent Intelligence Delivers	13
The Future of Sales is an Intelligent, Expert AI Agent	14

# Executive Summary

**Enterprise productivity is undergoing a significant paradigm shift. Traditional AI applications can offer assistance through summarization and search, but fall short when it comes to delivering real work.**

Vivun's Agent Intelligence introduces a new standard—where intelligent agents not only understand complex domains but take action on your behalf, proactively generating the work products that drive revenue.

At the core of our approach is a differentiated architecture that understands and separates expertise from language models. By starting with capturing how experts think and solve problems, Vivun builds structured knowledge graphs that serve as the “brain” of the AI agent it creates.

As these knowledge graphs become enriched with contextual data from CRM systems, conversations, email, Slack communications, and more. These persistent memories enable and empower AI Agents to understand your deals, your customers and your products with nuance and context.

**The result is an AI Sales Agent that faithfully executes on behalf of the expert it was modeled after, and does the work of the best Sales Engineers on the planet.**

This work includes, but is not limited to, the creation of stakeholder maps, handoff documents, solution documents, competitive analyses, and product feedback summaries— all of which are delivered to teams proactively, without prompting. This means you or your sales team need not be prompt engineers to get the most of our agent.

The work products generated are proactively created and customized to the seller and buyer needs, driving consistency, velocity, and scale across the sales process.

This white paper explores the core components of Agent Intelligence— **modeling, memory, modality, and work product generation**—and outlines how Vivun's agent gives more power to sales teams.



# Introduction: The Promise of Agent Intelligence

According to the 2024 State of Sales Report from Salesforce, sales teams utilizing AI are 1.3 times more likely to experience revenue growth compared to those not using AI, which highlights the technology's potential to enhance sales outcomes.

Many vendors will be announcing AI agents, but few will live up to the name.

As sales teams race to adopt AI, most solutions have failed to meet expectations. While many can generate content, surface summaries, or answer basic questions, they fall short when it comes to executing important work on behalf of the sales rep. This leaves reps to manually fill in the gaps, and overextend themselves, while they miss opportunities for scalable, repeatable execution.

Gartner defines **AI Agents** as

**“autonomous or semi-autonomous software entities that use AI techniques to perceive, make decisions, take actions and achieve goals in their digital or physical environments.”**

While we agree with this definition, but are choosing to extend it to set a new standard.

# Agent Intelligence

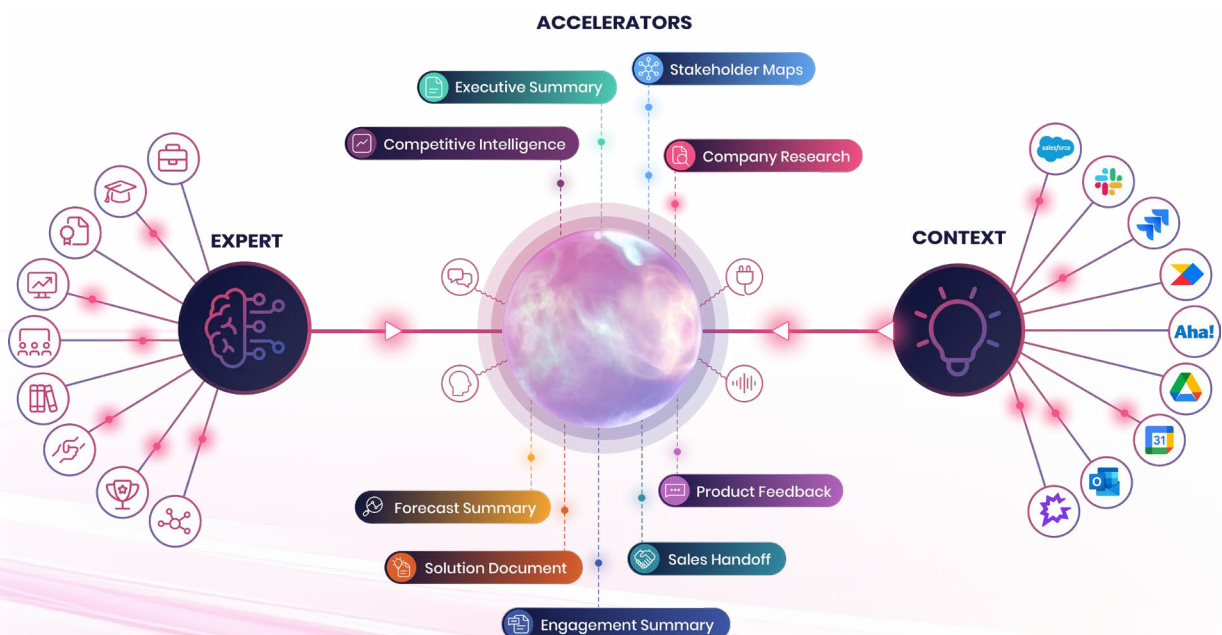
**Vivun builds AI Agents that not only assist—but does so proactively.**

Our agents are powered by domain-specific ontologies: structured representations of expert knowledge modeled after how the best Sales Engineers, product strategists, and technical sellers approach complex problems. This enables our agents to operate with deep domain expertise, making them capable of executing sophisticated tasks in a way that aligns with the real-world needs of enterprise sales teams.

While large language models (LLMs) have transformed how we interact with information, they are not sufficient on their own. LLMs are conversational engines—they can express ideas, draft content, and retrieve facts—but they lack the structured reasoning, memory, and domain fidelity needed to drive business-critical outcomes.

Vivun addresses this limitation by separating the LLM from the agent's core intelligence. In our architecture, the LLM is the interface—the mouth—not the brain. The brain is a knowledge graph derived from expert behavior, enriched with real-time context and reinforced with memory. This approach enables our agents to deliver output that is not just fluent, but fundamentally **correct, relevant, and actionable**.

We will explore how Agent Intelligence is modeled, how it's embedded within Vivun's architecture, and how it transforms the way work is done—ushering in a new era of AI that doesn't just support teams, but proactively performs expert work alongside them.



# Modeling Expertise

**At the foundation of Agent Intelligence lies a simple but powerful idea: to create AI agents capable of doing expert-level work, you must first understand how experts think.**

Vivun's R&D and leadership team brings deep experience in modeling expert thinking across diverse domains.

**We start by studying the behaviors, mental models, and workflows of top performers:**

- Sales Engineers who drive complex deals
- Producers who craft platinum records
- Finance leaders who manage billion-dollar portfolios.

From Bridgewater to Universal Music to the world's most sophisticated pre-sales organizations, we've worked alongside professionals who consistently outperform. What unites them is not just what they do, but how they approach their work: identifying key relationships, prioritizing inputs, navigating ambiguity, and executing with precision.

We start by breaking down these expert behaviors into their component parts. We've spent years capturing the nuances of how top Sales Engineers qualify deals, run discovery, align stakeholders, and transition information across teams. This process is neither guesswork nor abstraction—it's a rigorous, research-driven methodology rooted in behavioral modeling and domain-specific analysis.

Once this tacit expertise is mapped, we translate it into structured ontologies—organized representations of the relationships, concepts, and workflows that define expert performance in technical sales. These ontologies serve as the scaffolding for the agent's "brain," forming the foundation for all downstream intelligence. They define not only what the agent knows, but how it reasons, prioritizes, and acts.

This step is critical. Without a deep and accurate representation of expert knowledge, AI agents risk becoming generic assistants, capable of producing grammatically correct responses that lack substance or insight. Users perceive these tools as novelties or gimmicks, and don't leverage them to perform deeper work within a complex domain.

Our goal is different: to build agents that reflect the sharp judgment and high standards of the best human performers—and can scale that performance across an entire organization.

Modeling expertise is not a one-time event. It is a living discipline, continuously informed by customer feedback, evolving market conditions, and ongoing collaboration with thought leaders across sales, product, and engineering. It is the first—and perhaps most important—step in building agents that are not just intelligent, but trusted.



# From Model to Graph: Encoding Expertise as Agent DNA

Once expert behavior has been modeled, we translate that knowledge into a format that machines can reason over: the Knowledge Graph. This is where theoretical modeling becomes practical implementation.

A knowledge graph is a structured network of concepts and relationships that mirrors how an expert understands their domain. Unlike static documentation or even traditional databases, knowledge graphs are dynamic, contextual, and deeply interconnected. They reflect not only what matters, but how it matters—capturing the nuanced relationships between people, processes, products, and outcomes.

For Vivun, this means encoding key elements of technical selling: what makes a stakeholder influential, what signals risk in a deal, how product gaps are tracked and escalated, and what best practices look like for handoffs between Sales and Customer Success. Every concept is grounded in lived expertise and made navigable through the graph structure. This allows our agents to reason, reference, and respond with the kind of contextual awareness that mirrors human judgment.

It's important to distinguish this approach from more common techniques like vector embeddings or fancy UI wrappers around an LLM. While embeddings are useful for

identifying similarity or relevance within unstructured text, they lack the semantic depth required for structured decision-making. And while large language models excel at generating fluent text, they do not “know” anything — they rely on statistical correlations rather than grounded, persistent understanding.

Our agents are not built to guess. They are built to **know**. By externalizing knowledge into a graph—separate from the language model—we ensure that the Agent’s reasoning engine is transparent, auditable, and precise. This separation of concerns allows the LLM to function as the interface (how the agent interacts with the user) while the knowledge graph serves as the brain (how the agent thinks).

This architectural choice is what enables Vivun’s AI Sales Agent to move beyond hallucinated answers and generate work products grounded in real, verifiable knowledge.

It’s not just better output—it’s output you can trust.

# Transplanting the Brain: Embedding Knowledge in the Agent

**With the expert knowledge modeled and encoded into a knowledge graph, the next step is instantiating the agent—giving it the ability to function autonomously within an enterprise environment. This is where the knowledge graph is transplanted into an AI agent capable of perceiving, reasoning, and acting.**

This transplantation gives the agent its brain—a structured way to think based on how the best experts operate. But expertise alone isn't enough. For the agent to function in the real world, it must also have contextual awareness: of customers, conversations, and competitors in an evolving business context.

**Vivun's AI Sales Agent connects seamlessly with the core platforms where business happens:**

- Your CRM System
- Your Call Recording System
- Your Email and Collaboration Tools
- Your Files and Documents
- Web-Based Resources

These integrations allow the agent to ingest data continuously, updating its internal understanding of opportunities, personas, timelines, risks, and product gaps. Every data stream enriches the agent's view, providing not only static facts but dynamic signals that inform action.

To make this contextual intelligence durable, Vivun introduces the concept of memories. Our agents retain information over time—anchoring facts, decisions,

and interactions. It remembers who key stakeholders are in a deal, what concerns were raised in discovery, what feedback was provided to the product management team, and what assets have already been delivered to the buyer.

Memories are more than a convenience—they are the foundation for the proactive delivery of work products. They are what allows the agent to act without being prompted. Rather than wait to be told that a handoff is occurring between sales and post-sales teams, the document is proactively created from the beginning of the opportunity, and it iterates over the life-cycle of the opportunity. When it detects a complex stakeholder environment, it generates a stakeholder map. When it identifies recurring product concerns, it updates a feedback summary.

In this way, memory is the connective tissue between knowledge and action. It ensures that the agent is not just intelligent in theory, but capable of navigating real-world complexity with continuity and insight. Just as a great Sales Engineer brings not only skills but context and experience to every interaction, so does Vivun's AI Sales Agent – for every rep, for every opportunity.



# Memories: The Missing Link in AI Agent Design

**For most AI Agents, memories are either nonexistent or ephemeral. A user poses a question, the model generates a response, and the context is reset when they attempt to start the next agent task.**

But in the world of sales —where relationships evolve over time, deals progress through complex lifecycles, and historical context is essential—this stateless behavior falls short.

**Vivun changes that by making memory a first-class citizen.**

In our framework, memories are persistent, structured representations of what the agent has learned. They are not just cached interactions or temporary threads—they are stored knowledge elements tied directly to the agent's graph-based understanding of the world. Each memory is anchored in domain-specific concepts, such as stakeholder relationships, customer objections, feature requests, deal risks, and prior engagements.

These memories are continuously updated as new information flows into the system from connected sources like CRM entries, Slack threads, emails, calendar invites, and customer call transcripts. The agent doesn't just remember facts—it remembers their relevance, their relationships, and their implications.

This memory capability is what makes Vivun's AI Sales Agent **contextually intelligent**, not just generative.

It's the difference between an agent that answers questions and one that knows how to help—proactively, consistently, and with precision.

For example, if an account executive has an ongoing deal with multiple decision-makers, the agent will recall prior conversations, capture the key technical requirements, and automatically generate a stakeholder map. If a customer has expressed interest in a beta feature, the agent remembers that and includes it in future handoff documentation or product feedback summaries. If a competitor was mentioned in early discovery, the agent flags the risk in future internal deal reviews—without being asked.

In practical terms, memory powers the proactive generation of work products—documents and assets that previously took hours of manual effort, but now arrive automatically, grounded in up-to-date knowledge.

Without memory, AI output is generic—resembling a search result or a best guess. With memory, Vivun's agents deliver **personalized, situation-aware responses** that reflect the same depth and continuity you'd expect from a seasoned Sales Engineer.

# Output That Matters: Introducing Accelerators

**The true measure of an AI agent isn't how well it converses—it's what it produces. Vivun's AI Sales Agent doesn't just produce static content, but rather high quality work products, that we call Accelerators.**

Accelerators are AI-generated, contextually-aware, and proactively created assets that sales reps use to accelerate decision making and shorten sales cycles.

Unlike passive summaries or static reports, Accelerators are dynamic, actionable, and deeply tied to the deal context the agent continuously learns and remembers.

## **Examples of Accelerators generated by Vivun's AI Sales Agent include:**

- Solution Documentation which outlines what products to sell and the value case behind it
- Stakeholder Maps that identify and organize the key personas involved in a deal
- Sales Handoff Documents that ensure seamless knowledge transfer between sales and post-sales teams
- Product Feedback Summaries that aggregate real buyer insights and escalate them to Product
- Company Research Briefs that prepare sellers with key insights before a discovery or qualification call

These outputs are not prompted ad hoc—**they are generated proactively** as your AI Sales Agent monitors interactions across Slack, CRM, email, call recordings, and more, it detects signals that trigger the creation of specific work products. As the Solution becomes clear through discovery, a Solution Document is created. If a new stakeholder is identified, the stakeholder map is refreshed. If a feature concerns surfaces repeatedly, product feedback is logged and organized for review.

## **The value of these outputs is immediate and measurable:**

- As much as 40 hours of time saved per opportunity
- Higher win rates from a standardized, repeatable sales process
- Shorter cycle times as all sales teams can operate independently

With Work Products, Vivun transforms AI from a passive assistant into an active contributor.

# Interfaces and Modalities: Meeting Users Where They Work

**An intelligent agent is only as valuable as its accessibility. That's why Vivun's approach to Agent Intelligence includes a deliberate focus on interfaces and modalities the channels through which your users engage with their AI Sales Agent.**

Our AI Sales Agent is designed to meet sales teams where they already work. Whether it's through chat interfaces, embedded experiences in Slack, or the dedicated AI Sales Agent UI, users can engage with the agent across their existing workflows. Need to surface a stakeholder map mid-conversation? Your AI Sales Agent can drop it directly into a thread. Preparing for a call? Pull up company research in the workspace. No context-switching required.

**While many vendors rush to implement flashy features, Vivun takes a measured, quality-first approach.**

R&D efforts are actively exploring voice interfaces and avatar-based interactions\*\* that simulate live meetings or video-based assistance.

But these innovations will only be released when the underlying technology can match the quality of the expert-modeled knowledge that defines our AI Sales Agent's intelligence. Voice and avatars are not gimmicks—they are modalities of trust, and they must reflect the same rigor and reliability as every other part of the Agent.

By aligning modality with context and capability, Vivun ensures that your AI Sales Agent is always accessible, always useful, and always grounded in expertise—no matter how or where users choose to engage.



# The Mouth, Not the Brain: The Role of LLMs in Agent Intelligence

**Large Language Models (LLMs) are one of the most transformative technologies in modern AI—but without structure, direction, or memory, their outputs often fall short in enterprise environments.**

They're capable of generating fluent text, but they don't understand the nuances of your business, your product, your customers, or your goals.

**At Vivun, we take a fundamentally different approach: LLMs are the “mouth” of the agent, not the brain.**

This distinction is critical. The brain of our agent is the expert-modeled knowledge graph: a structured, contextual representation of how top technical sellers think and work. It's where reasoning happens, where priorities are set, and where complex relationships are understood and remembered. This brain gives the agent its identity, its standards, and its ability to take meaningful action.

The LLM, by contrast, is the interface layer — a powerful tool for retrieving, communicating, and interacting with the knowledge encoded in the graph. It enables natural language exchanges, allows users to ask follow-up questions, and makes the agent feel human. But without a structured brain behind it, the LLM alone is prone to hallucinations, inconsistencies, and generic advice.

**Fine-tuning models will not make the LLM have a brain.**

This is why traditional LLM applications often deliver average outputs. When asked to define what makes a good sales engineer, an LLM might give you a college-level answer—grammatically perfect, but strategically hollow. That's because it's regurgitating what's statistically common, not what's strategically correct. It has no model of excellence, no memory of your unique context, and no way to distinguish between what's true and what's just probable.

We overcome this by anchoring every interaction in expert-modeled knowledge. The LLM becomes a translator, not a thinker. It helps users access memories, explore concepts, and generate work products in a way that feels fluid and intuitive—but it never drives the logic on its own. This architecture is by design. It ensures that every output from Ava—the AI Sales Agent—is not only articulate, but accurate, contextual, and strategically aligned. By combining the fluency of LLMs with the precision of a graph-based brain, Vivun delivers an agent that doesn't just talk—it understands, remembers, and works.

Without a brain, an LLM is just noise. With an expert brain, it becomes a powerful voice for change.

# Use Case Deep Dives: Where Agent Intelligence Delivers

**Vivun's Agent Intelligence isn't theoretical—it's already transforming how sales teams operate. By proactively generating work products grounded in expert-modeled knowledge, your AI Sales Agent supports a range of high-impact use cases across the sales cycle. Below are four areas where her value is immediate, measurable, and game-changing.**

## Solutions

One of the most important outputs of a Sales Engineer is a clear, tailored recommendation on **what to sell—and why**. Without structure, solutioning becomes inconsistent: value drivers are misaligned, product fit is unclear, and business cases lose credibility. Your AI Sales Agent solves this by generating a dynamic Solution Document that outlines the recommended product set, mapped directly to customer pain points and desired outcomes. It includes business rationale, value justification, and technical fit—sourced from discovery notes, product data, and past deal intelligence.

## Stakeholder Maps

Identifying and managing deal stakeholders is a time-consuming, error-prone task—especially in complex buying groups with cross-functional decision-makers. Ava tackles this proactively. As conversations unfold across systems, your AI Sales Agent discovers stakeholders, understands their roles, and builds dynamic stakeholder maps that reflect the real structure of the deal. These maps are updated automatically as new personas are identified or org structures shift, enabling sellers to manage influence, track alignment, and plan multithreaded engagement strategies—without manual upkeep.

## Proactive Product Feedback Capture

In traditional sales motions, product feedback is either never captured, lost in a spreadsheet, or buried in Slack channels. It's anecdotal, inconsistent, and difficult to route back to product teams in a structured way. Your AI Sales Agent fixes that by detecting product feedback in real time, organizing it against known features or gaps, and creating **living** feedback summaries that are always up-to-date. Product managers get a clear, actionable view of what prospects and customers are asking for—backed by real deal context and citations. Sellers no longer have to remember to log it. It's just done.

## Sales Handoff

An often overlooked moment in the sales process is the handoff from pre-sales to post-sales teams. Without structure, vital context is lost: deal history, technical validation, stakeholder dynamics, and key objections all risk falling through the cracks.

These are just four of the 9 Accelerators – and are just the beginning. As Agent Intelligence evolves, so does its utility—enabling a world where work products are always delivered, always relevant, and always one step ahead.

# The Future of Sales is an Intelligent, Expert AI Agent

**The sales process is being redefined—not by incremental tooling, but by a fundamental shift in how work gets done. At the heart of this transformation is the rise of the AI Sales Agent —a new class of intelligent agent that doesn’t just support your team, but proactively contributes.**

Agent Intelligence represents the future of an AI Agent workforce: where expertise is captured, encoded, and scaled. In this new model, employees are empowered to move faster and further with less dependency. Functions are elevated, freed from repetitive work, and able to focus on strategic impact. Sales teams are transformed, operating with clarity, consistency, and unprecedented technical coverage.

The power behind this shift is a bold architectural choice—separating language from logic, interface from intelligence. By building agents grounded in domain-specific knowledge graphs, reinforced with memory, and activated across systems and channels, Vivun delivers an agent that faithfully executes the work of experts.

## **The path forward is clear:**

- Modular knowledge that evolves with your business.
- Agent-first design that meets users where they work
- Enterprise scale that delivers value across every role and every deal.

The AI Sales Agent isn’t just a product. It’s a partner. And with Vivun, the future of B2B sales is no longer a vision. **It’s here.**





## About Vivun

Vivun delivers the world's most advanced AI Sales Assistant—automating the critical work required to move complex deals forward. By generating sales-ready outputs that help every rep engage stakeholders, build value, and drive urgency, Vivun empowers sales teams to close more deals, faster.

To learn more visit [www.vivun.com](https://www.vivun.com).