Agentic Sales & Marketing Systems: A Multi-Agent AI Framework

Author

Name: Lawrence Duru
Email: lawrence@duru.co.uk

Abstract

We present a modular, cloud-native multi-agent system for orchestrating autonomous AI agents in enterprise sales and marketing workflows. Our framework comprises a set of role-specialised agents: Lead Generation, CRM Integration, Pipeline Prioritisation, Offer Matching, and Cross/Up-Sell Recommendation, each capable of operating independently and collaboratively within real-world business environments.

Each agent interfaces with external systems (such as CRMs, campaign tools, and enrichment APIs) through secure connectors, and reasons through its own goal context using a hybrid of symbolic planning and generative language models. Agents are coordinated via a shared orchestration layer that integrates Semantic Kernel for agent-level memory and planning, and the Multi-Agent Collaborative Planning (MCP) Toolbox to support multi-step collaborative task decomposition. Planning and execution are augmented with Azure OpenAI for language understanding and content generation.

The system is deployable on a per-client basis and is designed for integration with heterogeneous enterprise systems. This framework provides a real-world testbed for multi-agent reasoning, partial autonomy, distributed planning, and cross-agent coordination in commercial contexts. We aim to share our architecture, early implementation insights, and MAS design challenges with the community, and to contribute a practice-led MAS application model for business-critical domains.

---

1. Introduction

Modern sales and marketing teams operate in highly fragmented digital ecosystems with multiple CRMs, campaign tools, customer data platforms, and manual workflows combining to produce inefficiencies and lost revenue. While AI has been applied to these environments, most tools remain narrowly focused, deterministic, and siloed.

We propose an agentic AI framework that applies the principles of multi-agent systems, including specialised roles, autonomous decision-making, shared planning and coordinated execution, to sales and marketing operations. Our goal is to support intelligent, partially autonomous systems that collaborate across tools and data sources to improve decision quality, efficiency, and responsiveness in revenue-generating workflows.

---

2. System Architecture Overview

The framework is designed as a modular, Azure-native system, deployed on a per-client basis to ensure data isolation and operational flexibility. The architecture consists of:

- **Agent Execution Layer** (Azure Container Apps / Functions): Hosts the agents and their orchestration logic.
- **Orchestration Layer**: Utilizes Semantic Kernel for memory, planning, and LLM integration, with optional support for Google's MCP Toolbox to manage inter-agent coordination.
- **Reasoning Layer**: Azure OpenAI GPT-4 for generative reasoning and dynamic content creation.
- **Integration Layer**: Secure connectors to CRM systems (Salesforce, Dynamics), campaign tools (Braze, Adobe Campaign), and data enrichment APIs (e.g., LinkedIn, Clearbit).
- **Context + Memory**: Cosmos DB, Redis, and Cognitive Search provide long- and short-term context.

Each agent performs a distinct role and communicates via shared planning goals or system events. A representative architecture diagram is shown in Figure 1.

While the framework is still under development, the architecture has been designed with deployment and reproducibility in mind, using open components such as Semantic Kernel, Azure OpenAI, and container-based orchestration. We aim to release implementation scaffolding and example workflows in future phases.

---

3. Agent Roles and Capabilities

Lead Generation Agent	Scans external data sources, scores and enriches leads, pushes to CRM
CRM Integration Agent	Synchronises contact, opportunity, and activity data
Pipeline Prioritisation Agent	Scores open deals, flags risks, and suggests follow-ups
Offer Prioritisation Agent	Matches customers to products/offers and personalises messaging
Cross/Up-Sell Agent	Identifies expansion opportunities from usage, behaviour, and history

Each agent can act independently or as part of a larger goal plan.

---

4. Agent Planning and Coordination

Agents coordinate using a hybrid approach. The MCP Toolbox supports plan decomposition and inter-agent goal sharing, while the Semantic Kernel handles function orchestration, memory access and tool invocation. LLMs assist with plan adaptation, content generation and explanation.

This model allows agents to act tactically while aligning with shared objectives such as lead conversion, account expansion or deal closure.

---

5. Real-World Deployment and Use Cases

This system is designed for enterprise deployments, supporting:

- Integration with heterogeneous marketing stacks
- Per-client data isolation and scaling
- Flexibility in toolchain (Salesforce, Braze, Adobe, Dynamics, etc.)

Use cases:

- Real-time lead triage
- Adaptive offer generation
- Intelligent pipeline health monitoring

For example, the Lead Generation Agent identifies and enriches a promising contact. The Pipeline Agent scores it high-risk due to limited engagement. The Planner diverts effort to the Offer Agent, which generates a personalised reactivation campaign.

---

6. MAS Contributions and Future Work

This framework provides:

- A modular, real-world MAS implementation using current LLM and planning infrastructure
- A testbed for goal decomposition, collaboration and adaptive learning under enterprise conditions.
- Opportunities to explore:
    - RL for offer agent adaptation
    - Emergent agent communication and messaging
    - Meta-agent oversight models
    - Adaptive agent incentives in business contexts

We seek feedback from the MAS community on coordination strategies, long-term memory, and learning mechanisms under real-world constraints. This approach bridges research and deployment, offering practical insights into MAS behaviour in enterprise environments.

Figure 1: System Architecture