AI as a New Leadership Muscle

The Executive Takeaway: Architecture & The Vendor BS Detector

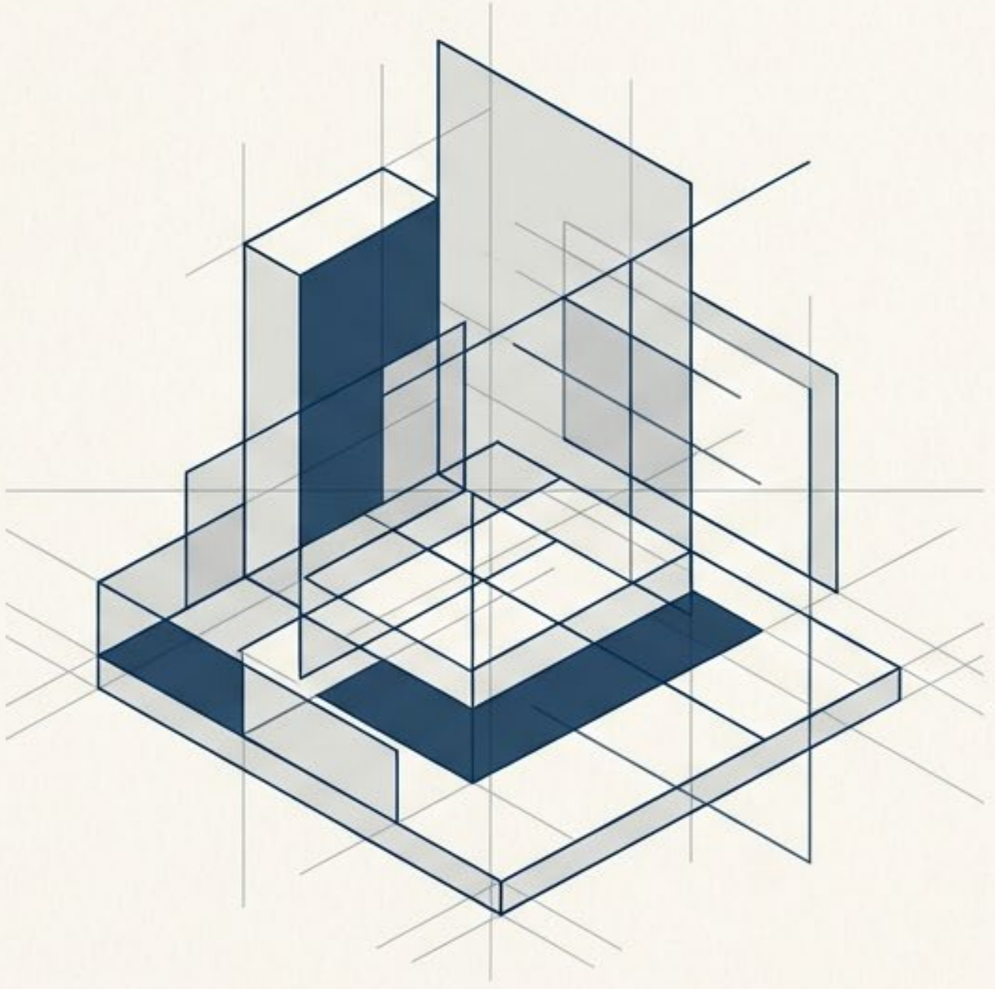

MASTERCLASS SYNTHESIS FOR THE PACIFIC COAST BANKING SCHOOL

The Anatomical Schematic

Layer 1: The Brain (Base Models)

The Commodity Engine

GPT-4, Claude. The electricity, not the appliance. They are "stochastic parrots" that know the internet but lack your bank's specific credit appetite. Relying solely on the Brain invites hallucinations.

Layer 2: The Context Wall (RAG)

The Library Card

Retrieval-Augmented Generation. The bridge to private data. Uses a Vector Database to look up your 2026 policy manual before the model speaks, anchoring output in reality.

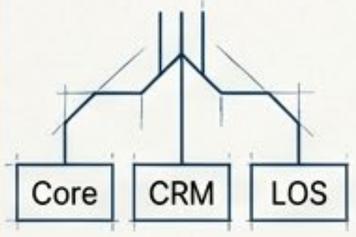

Layer 3: The Nerve System (MCP)

Hands and Feet

Model Context Protocol. The open standard enabling Agentic Workflows. The secure handshake that allows AI to securely do banking tasks instead of just talking about them.

Vibe Coding: Architect of Intent

The shift from writing syntax to describing intent. Empowers a Business Analyst to build functional internal tools in an afternoon, bypassing the 18-month IT roadmap.

Teaching AI: RAG vs. Fine-Tuning

✗ Fine-Tuning	✓ RAG
• Permanent brain surgery • Expensive & rigid • Gets stale instantly	• Dynamic lookup • Cites current reality • Required for banking

The Vandals Checklist: 3 Questions to Kill a Bad Pitch

How to cement your authority, evaluate architecture, and protect your bank.

	The 'Wonk' Question	The Strategic Translation	The Score
Data Residency	How are you handling Data Residency?	Is my bank's highly sensitive data leaving our secure tenant to train your proprietary model?	**RED FLAG:** Co-mingled training data or vague anonymization. **GREEN FLAG:** Zero-data retention guarantees strictly within our tenant.
Context Architecture	Are you fine-tuning, or what is your RAG stack?	How do you ensure the AI cites our current SOPs rather than hallucinating a 2023 internet vibe?	**RED FLAG:** "We built a custom proprietary brain." **GREEN FLAG:** "We use a RAG architecture with a Vector Database."
System Integration	Is this MCP-compliant?	Can this securely execute workflows in our Core, or are we just buying another isolated chatbot silo?	**RED FLAG:** Relies entirely on fragile, custom 1-to-1 API integrations. **GREEN FLAG:** Uses Model Context Protocol (MCP) as an open standard for system handshakes.