master.of.code GLOBAL

The Agentic Shift:

How Enterprises Are Building the Next Wave of Conversational AI with vCons, MCP, and A2A

Ted Franz
VP Sales & Partnerships

Ted Franz

VP Sales & Partnerships

With over 20 years of experience in software and telecommunications, Ted Franz is a seasoned leader in Conversational AI and IT services. He has built a career on driving revenue growth, launching new markets, and forming high-impact partnerships that fuel digital transformation.

At Master of Code Global, Ted helps enterprise clients design and scale AI-driven solutions that elevate customer experience and deliver measurable business outcomes. Backed by dual MBAs and a global network, he brings together deep industry expertise and practical insight to transform emerging technologies into real-world value.

Everywhere I go, people are asking about AI. Executives want to know how to use it to scale customer support. Marketers want personalization without manual effort. Product teams want to plug in an LLM and call it done.

But here's what most aren't ready for: once you move past the first chatbot or prototype, everything starts to fall apart. You miss important info jumping between tools. Data gets siloed. Handoffs between bots and humans break down. And good luck attempting to measure performance across it all.

That's the gap I've been seeing — and it's the reason I put this guide together.

At Master of Code Global, I work with teams across industries who are serious about AI. They're not just experimenting with automation — they're trying to build long-term, scalable systems. And lately, **three techs have stood out to me as real enablers of that next leap:**

1. **vCons** – a standardized, portable way to capture and store AI-human conversations.
2. **MCP** – a framework for coordinating tools, APIs, and decision logic across agents.
3. **A2A** – Google's new open standard for agent collaboration across vendors and domains.

Individually, each solves a specific pain. But together, they lay the groundwork for what I see as **the future of CX: composable AI** — intelligent, orchestrated experiences that persist across sessions, tools, and platforms.

In this eBook, I'll walk you through **what's broken in today's AI landscape, what these technologies actually do, and how your team can start implementing them now.**

Let's get into it.

TABLE OF CONTENT

PART 1

vCons, MCP, and Google's A2A

The Building Blocks of the Next AI Wave

The landscape of Conversational and Agentic AI has been evolving rapidly, but for many organizations, progress has been constrained by foundational architectural problems. Despite advances in LLMs, NLU engines, and chatbot builders, the underlying systems that power AI-driven interactions have remained fragmented, opaque, and difficult to scale.

Over the past few years, **three persistent pain points** have surfaced for teams building real-world conversational AI systems:

1. **Loss of context** across tools, sessions, and handoffs — especially between AI agents and humans.
2. **Poor interoperability** between conversational platforms, vendors, and orchestration systems.
3. **Limited reusability and traceability** of interaction data for analytics, auditing, or retraining.

To address these, new technologies have emerged — not as new features, but as new infrastructure:

1. **vCons (Virtual Conversation Objects)** were introduced to standardize and persist conversations across systems, creating portable, analyzable records of AI-human interactions that include transcripts, metadata, and decision traces.
2. **MCP (Model Context Protocol)** was designed to solve the challenge of tool and data orchestration in agentic systems — giving agents a structured way to interact with external resources and manage long-running context.
3. And most recently, **Google introduced the A2A (Agent-to-Agent) protocol,** an open standard that enables cross-vendor agent collaboration — allowing AI agents from different domains or platforms to work together in real time, with shared state and secure communication.

These technologies weren't created in isolation — they were driven by the growing need for scalable, trustworthy, and composable AI architectures. As organizations move from isolated chatbots to distributed, task-based AI agents, there's an increasing demand for a common language, a persistent memory layer, and reliable protocols for collaboration.

In this section, I'll explore how vCons, MCP, and A2A work together to transform how we build AI solutions — from customer support to marketing journeys — and why they matter right now for enterprises looking to differentiate, scale, and future-proof their AI capabilities.

Today's Reality: Single-Agent Systems, Siloed Context, Limited Reuse

Let's consider a common use case: a customer asks for a refund through a chatbot.

How It Works Today:

- A chatbot tries to answer or handle the refund with backend APIs.
- If it fails, the case is escalated to a human agent, often without full context.
- Data is logged in proprietary formats, siloed across platforms.
- AI decisions are hard to audit, retrain, or replay.

Limitations:

- Context is lost between systems.
- AI and human agents work in isolation.
- No reusable, standardized conversation record.
- Low visibility into AI reasoning and performance.

Current AI Refund Use Case

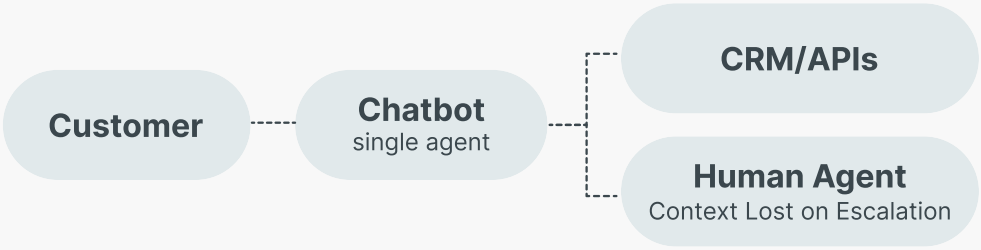

The Shift: Building with vCons, MCP, & Google's A2A

Now imagine the same use case, but **built with these foundational technologies:**

vCons (Virtual Conversation Objects)

- ✓ **Capture full interaction history:** transcript, metadata, AI decisions, handoffs.
- ✓ **Standardized and portable** — like the ".PDF of a conversation."

MCP (Model Context Protocol)

- ✓ **Manages** how AI agents access tools and context.
- ✓ **Enables** agents to retrieve relevant data and APIs securely.

Google's A2A (Agent-to-Agent Protocol)

- ✓ **Allows** different AI agents from different vendors to collaborate in real time.
- ✓ **Enterprise-grade** authentication, authorization, and agent identity.
- ✓ **Enables** cross-agent collaboration in multi-vendor environments.

How the Refund Scenario Works Now:

The Orchestrator Agent (via MCP) delegates to:

- ✓ A Refund Agent
- ✓ A Policy Compliance Agent
- ✓ A Sentiment Agent

If needed, **a 3rd-party agent** via A2A handles **specialized tasks.**

The entire session is captured in a vCon for auditing, training, or compliance.

A human agent joins with full context intact.

Future AI with vCons, MCP, and A2A Refund Use Case

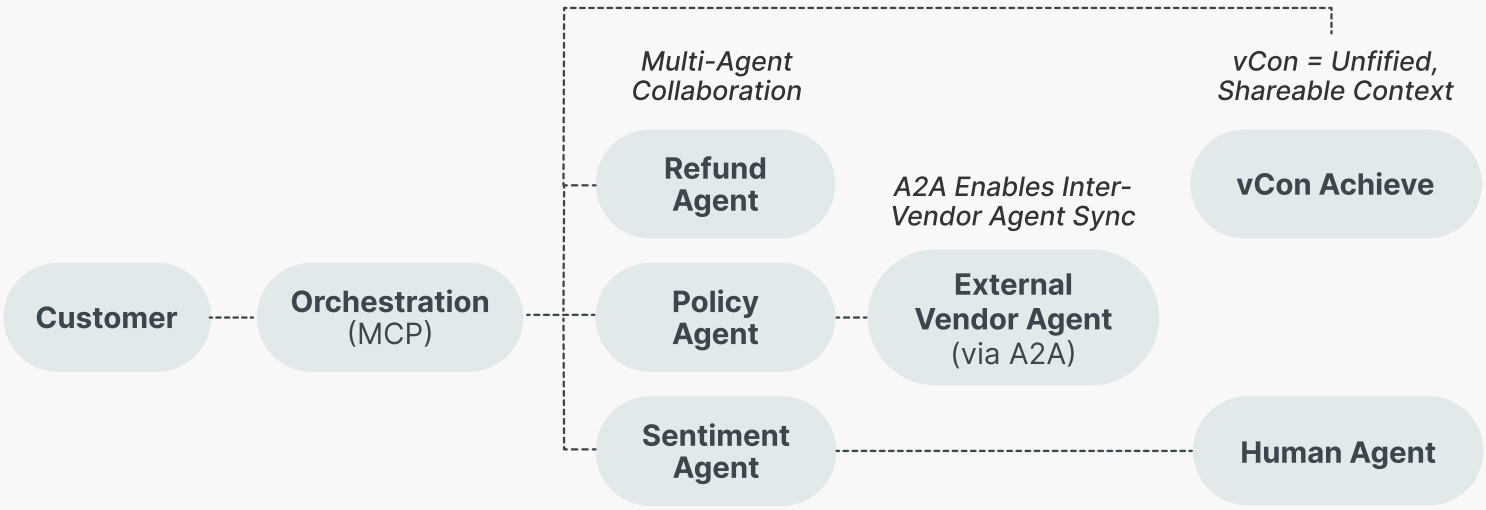

Benefit Area	Impact with vCons, MCP & A2A
ROI	Faster development, better reuse of agents, smarter automation
Scaling	Modular architecture, reusable components, orchestration-ready
Trust & Compliance	Transparent AI reasoning, traceable conversations, vCons = audit trail
Vendor Flexibility	A2A = Best-of-breed agent ecosystem without lock-in
Faster Time to Market	No need to build all agents in-house; plug-and-play AI partners
Analytics	Rich metadata, cross-platform insights, ideal for model retraining

Why It Matters Now

Google's release of A2A into the open standard ecosystem confirms that agentic interoperability is no longer a "nice-to-have" — it's the future. When combined with vCons and MCP, it unlocks a modular, intelligent, and highly traceable ecosystem that today's enterprises desperately need.

Enterprises building with this stack will not only scale AI more effectively — they'll gain measurable advantages in compliance, trust, and agility.

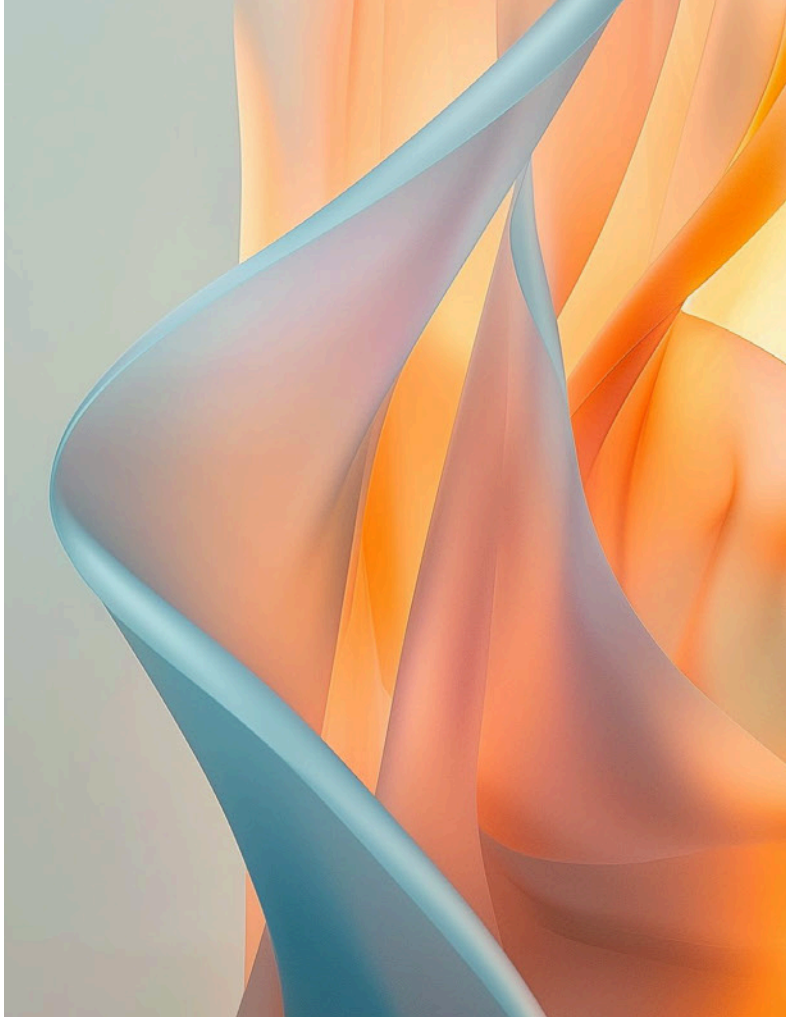

Conclusion: This Is the New AI Infrastructure Stack

1. **vCons**
   Durable, structured conversation memory

2. **MCP**
   Context-aware tooling and orchestration

3. **A2A**
   Cross-platform agent collaboration at scale

Together, these tools represent the new AI infrastructure layer. They solve the messiness of multi-agent systems, handoff confusion, and vendor fragmentation.

In the next 12 months, expect more platforms and services to adopt these standards, and early adopters to reap the rewards.

To help you and your readers delve deeper into these technologies, here are some authoritative resources:

vCons (Virtual Conversation Objects)

- ✓ **IETF Working Group:** The Internet Engineering Task Force (IETF) has established the vCon Working Group to develop a standardized format for representing digital conversations across various platforms. GitHub+3conserver.io+3IETF Datatracker+3
- ✓ **vCon Container Specification:** This draft outlines the JSON-based standard container for conversational data, detailing its structure and components. IETF Datatracker

MCP (Model Context Protocol)

- ✓ **Official Specification:** Anthropic has introduced MCP as an open standard to enable seamless integration between AI applications and external data sources. The comprehensive specification provides in-depth details on its architecture and implementation. Model Context Protocol+1The Verge+1
- ✓ **Introduction by Anthropic:** For an overview and the rationale behind MCP, Anthropic's official announcement offers valuable insights. The Verge

Google's A2A (Agent-to-Agent Protocol)

- ✓ **Google's Announcement:** Google has unveiled the A2A protocol to facilitate seamless communication and collaboration between AI agents across different platforms. Their developer blog provides an in-depth introduction. Home- Google Developers Blog
- ✓ **Media Coverage:** For a concise summary and analysis of A2A's implications, this article from The Verge offers a comprehensive overview. The Verge

These resources should provide a solid foundation for understanding the development, purpose, and applications of vCons, MCP, and A2A.

What are some practical steps an organization can take to implement these technologies now?

Great question — here are **practical, near-term steps** an organization can take to begin **implementing vCons, MCP, and A2A today,** even if you're just starting to evolve from traditional chatbot or conversational AI architectures:

1 Start Capturing Conversations as vCons (or Close Equivalents)

Practical Actions:

- **Audit your current chat/voice interaction logs** — see what metadata you're already capturing (timestamp, intent, response time, agent actions).
- **Begin structuring logs using vCon-like formats:** JSON objects that include speaker turns, sentiment, NLP outputs, escalation status, etc.
- **Use open-source tools or middleware** (e.g., adapters or logging layers) to generate vCon-compatible exports.

Outcome:

- Create a **single source of truth** for every customer conversation.
- Lay the foundation for **training, auditing, and cross-platform reuse.**

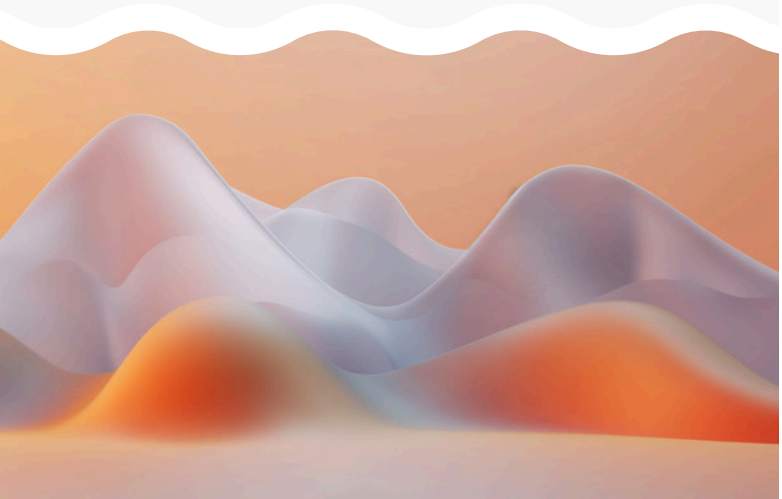

2 Design Your Agent Workflows Using MCP-Like Principles

Practical Actions:

- **Map current tasks** (e.g., order lookup, refund validation, loyalty status check) to tool or data access patterns.
- Design your bot/agent flows to **explicitly separate reasoning** (LLM/NLU) from action (tool use).
- If you're using tools like OpenAI Function Calling, Azure AI Studio, LangChain, etc., **use structured tool registries** to simulate MCP.

Tools to explore:

- LangGraph for agent flow design
- Anthropic's MCP Spec for inspiration
- Custom middleware for request/response tracking

Outcome:

- Build AI agents that are **modular, maintainable, and testable.**
- Pave the way for **context-aware decision-making.**

3 Pilot Multi-Agent Interoperability (Early A2A Concepts)

Practical Actions:

- **Identify adjacent agents or services you want to coordinate:** e.g., refund approval bot ↔ policy engine ↔ fraud check service.
- **Start with loosely coupled service-to-service APIs** using shared schemas and webhook-style communication.
- If building with LLMs, use frameworks that support **agent collaboration or delegation** (like CrewAI, AutoGen, or OpenAI Assistants).

Outcome:

- Begin developing **multi-agent playbooks** for tasks like refund handling, lead qualification, or onboarding.
- Prepare your stack for **future A2A adoption** once open standard protocols mature and gain tool support.

4 Run a Low-Risk Use Case (e.g., Refund Request, FAQ Triage)

Choose a self-contained, high-volume use case where you can test:

- **Logging** conversations as vCons
- **Delegating** tasks via tool/function calls
- **Collaborating** across multiple micro-agents (internal or external)

Example: "Request a refund → check eligibility → suggest store credit → escalate if needed → log everything in vCon format"

5 Set Up Analytics + Governance Early

Practical Actions:

- **Build** a vCon analytics layer (or even just dashboards from enriched logs).
- **Capture:**
  - Which agents were involved
  - Sentiment/CSAT drift
  - Human fallback rates
- **Define** data retention, auditing, and observability practices.

6 Find the Right Partner (Optional but Helpful)

If your internal team doesn't have deep expertise in:

- **Agent orchestration**
- **LLM tuning**
- **Data interoperability**
- **vCon architecture**

...then bring in a trusted services partner to help stand up a working reference architecture, prototype or roadmap to a proof of value.

PART 2

The Conversation Design Revolution —

How Marketers and CX Teams Win with vCons, MCP, A2A, and Rich Messaging Channels

As the AI stack evolves with technologies like **vCons, MCP, and Google's A2A protocol,** the role of marketers and conversation designers is being redefined — not replaced, but elevated. These aren't just back-end plumbing or abstract protocols. They're the foundation for building **richer, more context-aware, omnichannel experiences** that can finally live up to the promises of personalization, automation, and CX at scale.

For marketers, this means **moving beyond disconnected campaigns and channel silos** to orchestrated journeys that adapt to context, behavior, and emotion — across SMS, RCS, Apple Messages, web chat, and more. For conversation designers, it means crafting interactions that are **modular, data-driven, and agent-aware,** where conversation flows are no longer locked into rigid scripts but are **dynamic, tool-integrated, and reusable.**

This shift introduces **immense creative potential,** but also a new layer of complexity: mapping intent to orchestrators, managing multi-agent choreography, structuring reusable vCon outputs, and aligning conversation design to enterprise data strategies. While tools are catching up quickly, many teams are finding that a new kind of architecture-savvy collaboration is needed to make it all come together, not just to build experiences, but to maintain and scale them effectively.

The upside? Those who lean into this shift early will be able to create **truly differentiated experiences** that are not only high-performing but also easier to measure, optimize, and scale.

The New Foundation for Rich, Personalized Conversations

Before diving into use cases, let's revisit some key concepts and what these technologies bring to the marketing/CX stack:

Tech	What it Enables
vCons	Standardized, portable conversation objects with metadata, transcripts, outcomes — like the "CRM of conversations."
MCP	Contextual decision routing — AI agents can access business rules, data, and tools to respond meaningfully.
A2A	Inter-agent collaboration across vendors and domains — enabling modular, best-of-breed experiences.
RCS / Apple Messages for Business	Rich UI elements like carousels, maps, secure payments, calendar, photos, etc. — all inside a conversation.

Let's revisit our refund use case from part 1 of our discussion.

Use Case 1: Customer Refund in a Rich Channel (RCS or Apple Messages)

Without vCons, MCP, A2A:

- A scripted chatbot replies with plain text via web chat or SMS.
- Escalations require re-authentication and repeated explanation.
- Agents have to guess customer intent or manually re-check transaction data.
- Limited opportunity for upselling, CX recovery, or proactive outreach.

With vCons, MCP, A2A + RCS/ Apple Messages:

- The user taps “Request Refund” in an RCS button UI.
- **A Refund Agent** (MCP) is triggered, which:
  - Validates the order,
  - Checks refund eligibility,
  - Detects sentiment issues via **a Sentiment Agent.**
- **A Policy Agent** collaborates via A2A for legal/regional rules.
- The refund is approved, and a **rich carousel** of store credit offers is shown within the conversation.
- All context is logged in a **vCon** for compliance, retraining, and campaign attribution.

Rich Refund Use Case with MCP, vCons, A2A, and RCS/Apple

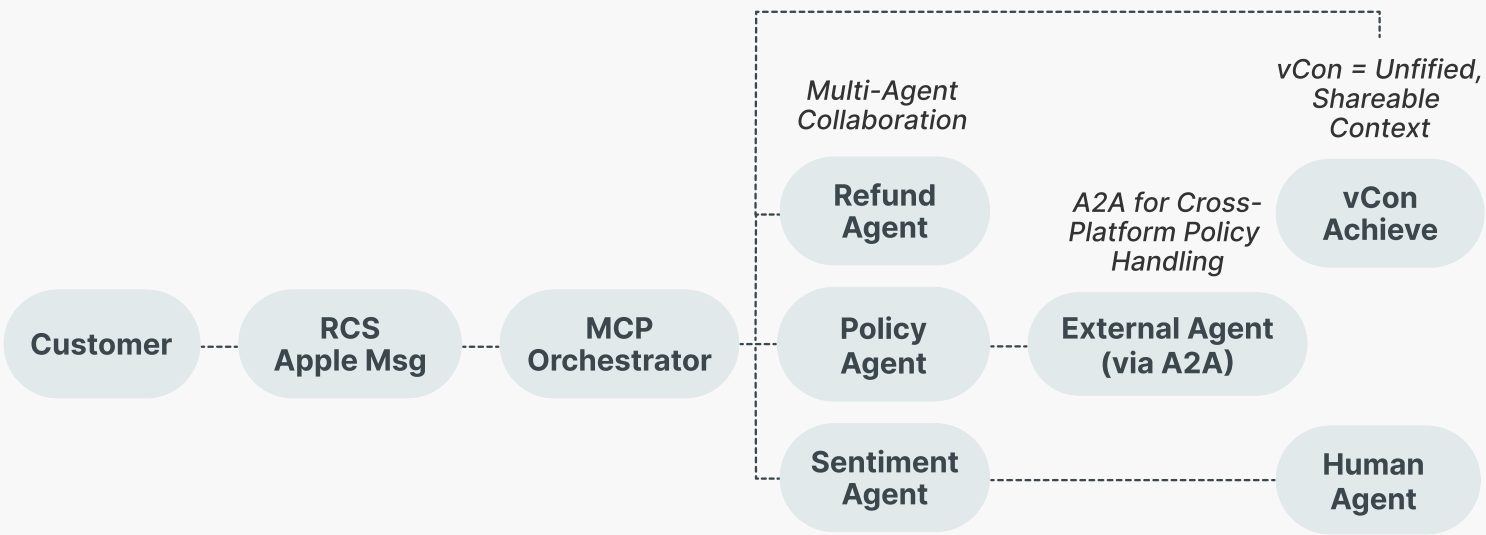

Impact for Designers & Marketers:

Drag-and-drop flows can be **channel-aware** without rewriting logic for each platform.

Responses are **data-driven and context-aware,** reducing scripting and guesswork.

Analytics from vCons = conversion tracking, sentiment, agent performance, CSAT — **all in one place.**

Now let's think of how marketers can use this for an outbound campaign use case. In this example, for retail and e-commerce.

Use Case 2: Personalized Re-Engagement for Retail/E-commerce

Imagine sending a reminder to a customer who abandoned a cart two days ago.

Today (Without MCP, vCons, A2A):

- The customer gets a generic SMS or email.
- If they reply, a chatbot offers limited support and can't access past behaviors or cart details.
- No rich product images or checkout links in the same thread.
- No easy way to update or reuse this flow across campaigns.

With the New Stack:

1. Customer receives an **RCS message:** "Still interested in your picks?"
2. **The MCP-based Marketing Orchestrator Agent:**
   - Pulls the user's cart and browsing data.
   - Invokes a **Promo Agent** to tailor incentives.
   - Checks for **stock availability or alternate SKUs via another agent.**
3. In-message carousel shows items + personalized offers.
4. Customer taps "Buy Now" → in-line checkout, Apple Pay, or Google Pay.
5. Conversation is wrapped into a **vCon,** tagged with:
   - Campaign ID
   - Delivery info
   - Product interest
   - Agent path
   - Offer accepted
6. A2A agents from logistics, loyalty, or product support can continue the conversation if needed.

Result: The conversation becomes the journey, not a broken series of disconnected tools.

Personalized Retail Marketing with MCP, vCons, A2A, and RCS/Apple

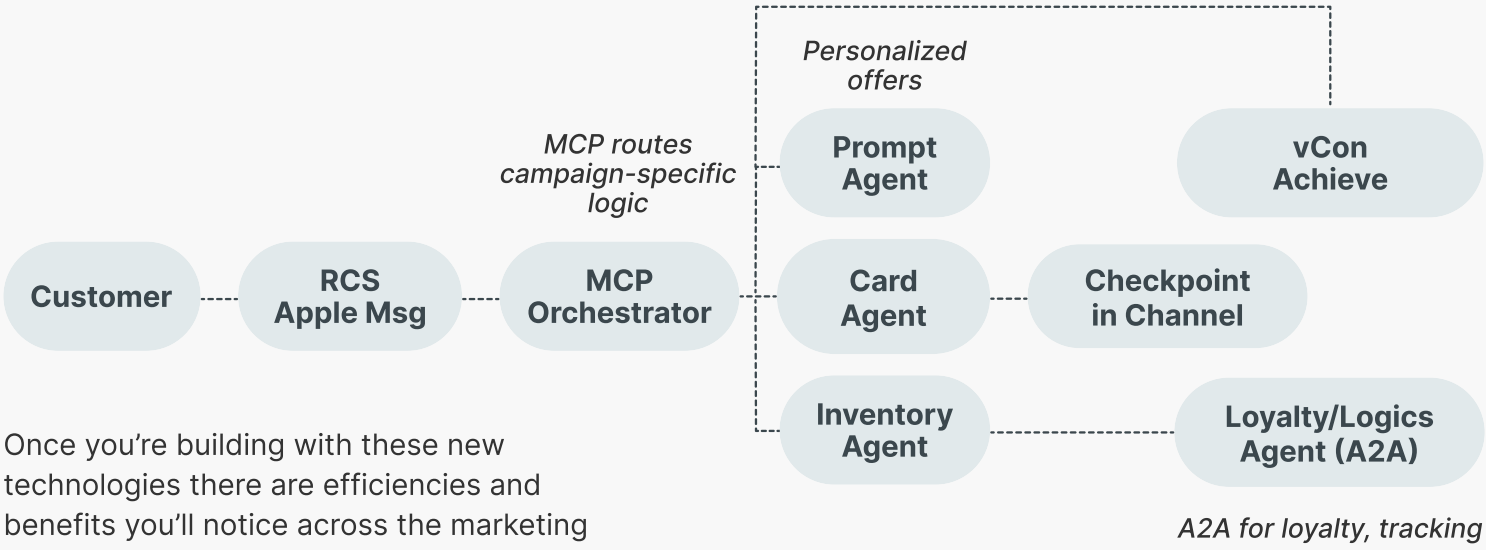

Once you're building with these new technologies there are efficiencies and benefits you'll notice across the marketing lifecycle which are important to note.

Creation, Optimization, and Maintenance: What Changes for the Better

Process	Without vCons / MCP / A2A	With the New Stack
Design	Custom for each channel, scripting heavy	Abstract flows + reusability across channels
Data Access	Manual integrations or API spaghetti	MCP-based contextual data routing
Channel Features	Underused due to flow rigidity	Fully leverages rich messaging like RCS & Apple Business Chat
Testing	Slow to iterate across platforms	Unified simulation + A/B flows via vCons
Analytics	Fragmented, hard to connect to outcomes	Every conversation becomes an analytics artifact
Handoff Logic	Manual, often breaks	Seamless, context-rich escalation via A2A
Maintenance	Duplicated logic per campaign/ platform	Centralized orchestration, plug-and-play agents

At Master of Code Global, we're very excited to see how these new technologies will be used and embedded into everyday tooling. Differentiation at scale is the name of the game and as we are all seeing, the impacts of AI continue to drive innovation.

Differentiated CX at Scale

Marketers want impact.
Designers want consistency.
AI wants context.
This new stack delivers all three.

vCons + MCP + A2A + Rich Channels give us:

- Conversational journeys that convert
- Multi-agent collaboration that scale
- Context that persists
- Rich UI with native platform feel
- Faster time to market for every campaign
- Unified analytics from ad click to checkout to refund

This Is Not Just AI Infrastructure — It's Brand Infrastructure

Every brand wants to personalize CX at scale. But without the right architecture, you're just duct-taping together bots, APIs, and UI kits.

This new model turns conversations into applications. It turns interactions into data. And it turns marketers and designers into orchestration engineers — without the code.

If you're designing, managing, or scaling AI-driven brand experiences, now is the time to rethink how you build — and what's finally possible.

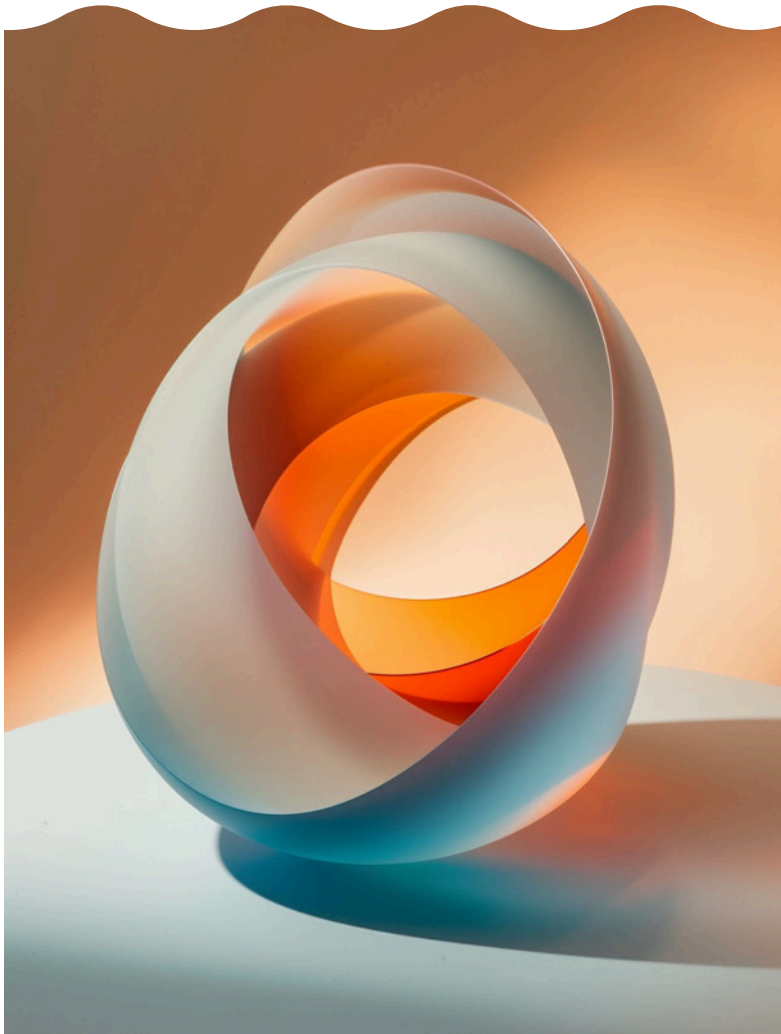

PART 3

Securing Conversational AI

How A2A, MCP, and vCons Must Evolve to Protect the Customer Experience

Conversational AI is moving from simple bots to Agentic ecosystems — autonomous agents collaborating to serve the customer better. New standards like MCP, A2A, and vCons are the backbone of this transformation.

But each unlocks new risks: data tampering, prompt injection, unauthorized memory access, and broken trust chains. Security can no longer be an afterthought. It must be designed at the protocol level.

MCP — The Dangerous Double-Edged Sword

1 What MCP Enables:

- **Agents share** tools, context, and memory fluidly.
- **Rapid integration** of capabilities without custom API code.

2 Security Risks (summarizing the uploaded analysis):

- **Shared Memory Poisoning:** One compromised agent can corrupt global memory and mislead the entire system.
- **Tool Invocation Attacks:** Malicious or misrepresented tools can hijack workflows.
- **Version Drift:** Lack of strict versioning allows "silent failures" or exploits.
- **Confused Deputy Problems:** Agents unknowingly trust malicious intermediaries.

3 Solutions Required:

- **Scoped Context Access** (signed, auditable memory writes)
- **Tool Schema Sanitization**
- **Formal Versioning and Compatibility Enforcement**
- **Sandboxed Execution** of Invoked Tools

4 Takeaway:

- Without an Agent Security Framework layered on MCP, Conversational AI becomes an exploitable free-for-all.

A2A — Interoperability Needs Zero Trust

1 What A2A Enables:

- **Direct agent-to-agent communications,** even across vendors.
- **Modular, composable, dynamic workflows** between disparate agents.

2 Security Risks (based on Caleb Sima's insights):

- **Rogue Agent Interactions:** Without strict authentication, unauthorized agents could join conversations or workflows.
- **Credential Leakage:** Poor isolation could allow agents to leak API keys or private data during collaborations.
- **Trust Assumptions:** Assuming a partner agent is "safe" based on metadata rather than active proof.

3 Solutions Required:

- **Mutual mTLS Authentication** Between Agents
- **Policy Enforcement Points** (PEPs) on Every Action
- **Fine-Grained Access Control** by Agent Identity and Context
- **Agent Health and Behavior Monitoring** in Real-Time

4 Takeaway:

- In A2A, every agent must prove who they are — and what they're allowed to do — for every interaction.

vCons — Recording Conversations Without Compromising Privacy

1 What vCons Enable:

- **A portable, standardized "container"** for conversation records.
- **Supports audits, compliance** (GDPR, HIPAA), customer transparency.

2 Security and Compliance Risks (from DataTrails and IETF):

- **PII Leakage:** Transcripts can unintentionally store names, addresses, payment details.
- **Record Tampering:** Without tamperproofing, vCons could be altered to hide fraud or mistakes.
- **Unauthorized Access:** vCon files moving between systems without strict controls expose major breach risks.

3 Solutions Required:

- **End-to-End Encryption** of vCon Data
- **Consent Management:** Explicit logging of user permission to store or share conversations.
- **Digital Signatures:** Cryptographic integrity validation for every vCon file.
- **Selective Redaction:** Ability to remove sensitive fields pre-storage or pre-sharing.

4 Takeaway:

- vCons must be treated as sensitive artifacts — not casual conversation logs.

Toward a Unified Agent Security Framework

Principles Across MCP + A2A + vCons:

- **Zero-Trust by Default:** No implicit trust across agents or systems.
- **Scope and Minimize Access:** Always enforce least-privilege principles.
- **Immutable Audit Trails:** Everything must be logged immutably — for internal trust and external compliance.
- **Dynamic Policy Enforcement:** Fine-grained policies enforced per interaction, not per session.

Emerging Best Practice:

- **MCP = Secure Context & Tools**
- **A2A = Secure Collaboration**
- **vCons = Secure Records**

Only together, with security woven through each layer, can Conversational AI deliver the customer experience enterprises demand — without sacrificing trust.

Implementing Security in a Real-World Refund Use Case

To ground these concepts, let's walk through how security would be layered across an actual Conversational AI refund journey using MCP, A2A, and vCons — based on current best practices.

Step 1

Customer Initiates Refund via Conversational AI Agent

Security Actions:

- Customer **identity is authenticated** via OAuth 2.1 / OpenID Connect.
- **Conversation context sanitized** to remove unexpected instructions.
- **Session tokens scoped** to only refund workflows (least privilege).

Prevents unauthorized refund attempts or injection attacks at the start.

Step 2

Conversational Agent Prepares Refund Request (MCP Context)

Security Actions:

- **Context data** (order ID, reason) signed with cryptographic hash.
- **Scoped sharing:** Only essential fields passed forward.

Ensures that context traveling through MCP cannot be silently altered.

Step 3

Conversational Agent Calls Order Verification Agent (via A2A)

Security Actions:

- **Mutual TLS authentication** enforced between agents.
- **Signed, scoped A2A request** ("VerifyOrderEligibility" only).
- **Response includes provenance metadata.**

Enforces trust boundaries and prevents rogue eligibility approvals.

Step 4

Refund Triggered via Finance/ERP Agent (via A2A)

Security Actions:

- **Policy Enforcement Point** (PEP) checks identity, scope, and thresholds.
- **Transaction digitally signed** and recorded in immutable audit logs.

Eliminates overreach or unauthorized financial transactions.

Step 5

Customer Notified of Refund

Security Actions:

- Notification triggered **only after confirmed refund execution.**
- Minimum necessary info shared via **secure A2A event subscription.**

Prevents misinformation or premature customer updates.

Step 6

Conversation and Actions Logged in vCons

Security Actions:

- **vCon container created** for full transcript + metadata.
- **Encryption applied** (AES-256+ standards).
- **Selective redaction** of PII.
- **Tamper-evident signatures** on vCons.

Ensures compliance and protects customer data integrity.

Final Best Practice Checklist (June 2025 Standards)

Area	Best Practice
MCP Context Handling	Signed payloads, scoped access, provenance metadata
A2A Agent Communication	Mutual TLS, scoped permissions, signed requests, least privilege
vCons Storage	End-to-end encryption, selective redaction, tamper-evident signatures
Event Handling	Secure event subscriptions, PEP enforcement, audit trail creation
Authentication	Federated identity, session scoping, reauthentication for sensitive actions

Why This Matters

- **Prevents Refund Fraud:** Attackers cannot hijack refund workflows.
- **Enhances Trust:** Customers receive secure, verifiable service.
- **Regulatory Ready:** Logs and compliance artifacts are automatically protected.
- **Future-Proof Systems:** Supports evolving standards like A2A+ and future MCP versions.

Secure Refund Flow

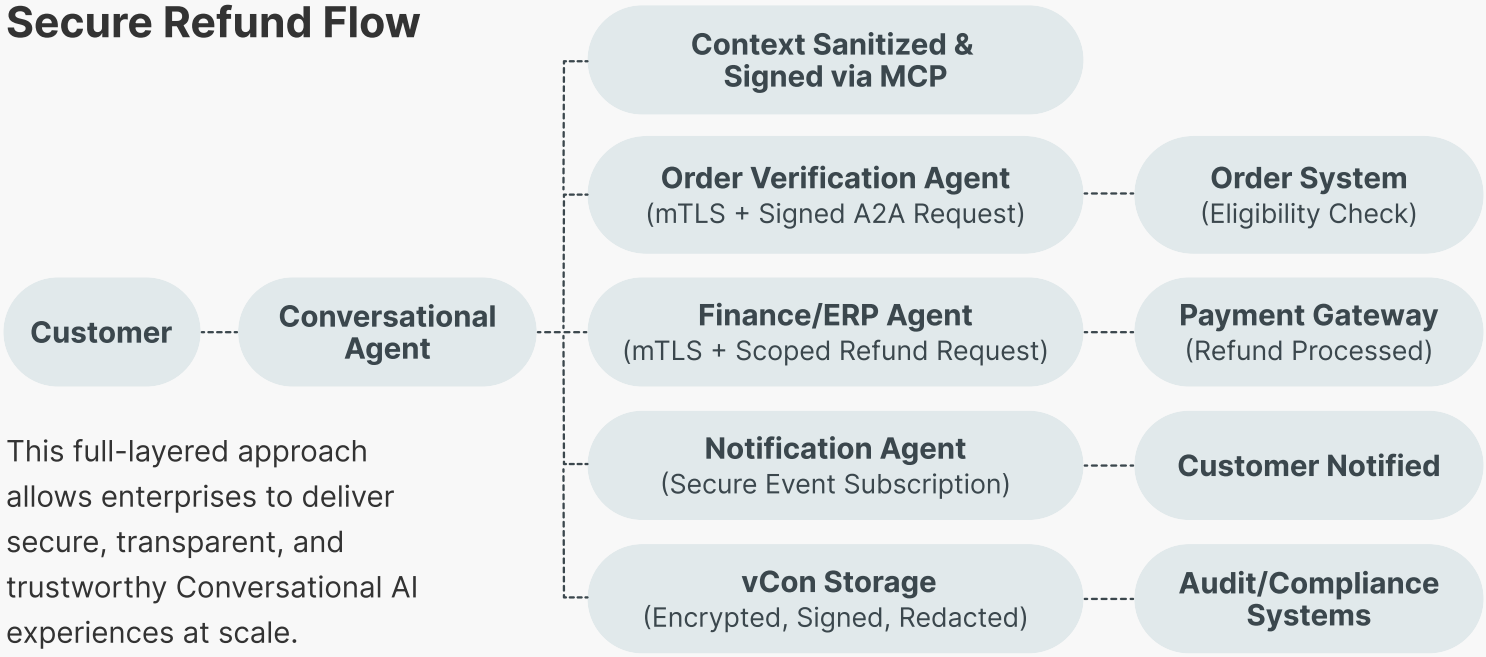

This full-layered approach allows enterprises to deliver secure, transparent, and trustworthy Conversational AI experiences at scale.

Conclusion

The future of Conversational AI is not just smarter — it's more complex and interconnected. MCP, A2A, and vCons are powerful enablers, but they will also be our greatest liabilities unless security is rethought at every interaction, memory, and recording. Secure systems will win trust. Unsecured systems will collapse under the weight of their vulnerabilities.

If you're designing agent ecosystems for Conversational AI, let's connect. DM me to discuss how to future-proof your customer experience against emerging threats.

master.of.code

GLOBAL