Introduction

Like most industries, the travel industry is at the cusp of fundamental transformation, thanks to AI. Many, if not most, travel organizations are still navigating how to apply AI in ways that are impactful, scalable, and secure.

At Evoort Solutions, we have taken a focused, domain-first approach. **Compass AI** is a purpose-built, scalable AI platform for modern travel, integrating various AI applications using MCP (Model Context Protocol) and A2A (Agent-to-Agent) protocols. Whether the goal is automating customer interactions, assisting agents, surfacing insights, or refining decision-making, Compass AI delivers AI where it matters most.

This whitepaper introduces our vision, walks through the architectural foundations of Compass AI, and explores how it is already solving complex travel challenges — starting with our successful development of an AI-powered customer service solution for travel businesses.
We are not theorizing what AI *could* do in travel. **We are building it — one solution at a time.**

Problem Statement

Travel businesses operate in complex, fragmented environments with siloed data, legacy systems, and piecemeal automation. Traditional AI solutions, including NLP-based AIs, often fail to address these ground realities, lacking the contextual depth and operational fit required across diverse travel workflows. Meanwhile, rising demands for speed, personalization, and efficiency strain internal teams.

At Evoort, we believe the answer is not piecemeal fixes — complex travel businesses need a unified, domain-aware AI platform. **Compass AI** is our purpose-built solution designed to host interconnected AI applications that evolve use case by use case. It integrates seamlessly with enterprise systems, respects policy constraints, and enables intelligent interactions between internal tools and third-party agents via Model Context Protocol (MCP) and Agent-to-Agent (A2A) protocols.

Co-created with travel businesses, Compass AI delivers scalable automation across the travel value chain — enhancing both customer experience and operational performance.

Compass AI - An Overview

Compass AI is Evoort Solutions modular AI platform designed for the travel industry. Compass is a flexible, domain-native intelligence platform that enables travel enterprises to apply AI where it matters — across customer service, operations, content, and strategy.

At its core, Compass AI combines state-of-the-art large language models (LLMs) with Retrieval-Augmented Generation (RAG) and a robust, locally hosted verification pipeline — delivering context-aware, accurate, and secure outputs within the organization's infrastructure for maximum data control, accessibility, and integration fidelity.

Compass AI is structured as a modular platform, with each use case developed independently and integrated gradually based on business priorities.

Each component of Compass AI is built to solve a specific, high-value challenge — from handling complex customer queries to assisting support agents, triaging incoming tickets, summarizing internal knowledge, or generating campaign content. Travel enterprises can adopt what they need now and expand incrementally over time, without disrupting core systems or overhauling workflows.

Compass AI adapts to business needs, scales with business growth, and brings true human-like AI capabilities at lightning speed across the travel journey.

At Evoort, we have already built the first — and a key — AI automation for travel businesses: **customer service automation.**
A particular use case developed for airlines is already live and delivering measurable results. The solution is highly

adaptable for other sectors such as hotels, online travel agencies, bed banks, room aggregators, travel management companies, cruises, and destination activity providers.

Technical Approach and Intelligence Stack

Compass AI is engineered to bring advanced generative AI capabilities into the structured, high-stakes environment of the travel industry. Its technical architecture prioritizes modularity, data grounding, policy alignment, and security-first deployment — ensuring that every AI output is useful, traceable, and production-grade.

Here is a deeper look into the components that make Compass AI function intelligently and reliably across varied enterprise use cases:

LLM-Centric Architecture with Domain Adaptation

At the core of Compass AI are foundation-level large language models (LLMs), selected based on use case requirements — whether open-weight (e.g., Mistral, LLaMA) or API-based (e.g., GPT, Claude).
These models are enhanced through:

- Prompt engineering and dynamic context injection, shaped by Compass AI's Model Context Protocol (MCP) — a standardized framework for contextualizing prompts with business logic, user history, and policy layers.
- Instruction tuning for travel-specific intent categories and user behaviors.
- Lightweight finetuning (LoRA/PEFT) on anonymized internal query data (where permitted) to improve precision for domain-heavy tasks like airline bookings, loyalty FAQs, or policy clarifications.

Retrieval-Augmented Generation (RAG) with Secure Context Injection

To ensure that LLM outputs are factual, current, and non-hallucinatory, Compass AI uses a flexible RAG pipeline composed of:

- Indexed enterprise knowledge bases (e.g., FAQs, SOPs, regulatory policies) embedded using Sentence Transformers or OpenAI/LLM-compatible vector encoders.
- Real-time SQL or API connectors that securely fetch context-specific data (e.g., loyalty points, flight status, membership details).
- Query rewriting & vector reranking for better retrieval precision and disambiguation of user intent.

Context from RAG is injected into the LLM prompt at runtime via MCP, constrained by token budget and priority scoring logic (e.g., policy over FAQ or live system data over static records). The result is model output that remains grounded in operational truth while adapting fluidly to user context.

5-Layer Verification and Routing Pipeline

Compass AI applies a structured intelligence pipeline to gate and verify every interaction before generating a final response. This pipeline includes:

- **Policy Compliance Classifier**
   Determines whether the user query or response needs to be suppressed, redirected, or adjusted based on policy (e.g., visa assistance limitations, escalation rules).

- **Ambiguity Detection Module**
   Uses a shallow intent confidence threshold + semantic distance analysis to identify vague or partial queries, triggering clarification prompts when needed.

- **Intent-to-System Mapper**
  Maps structured user intent (booking, loyalty, policy, account, general) to corresponding backend systems or RAG indexes for retrieval.

- **Secure Context Builder**
  Assembles a prompt with retrieved documents, metadata, and tokens like PNR, customer email, or system flags in a pre-validated, obfuscated format, governed by MCP.

- **Human Handoff Layer**
  Flags queries that exceed model certainty thresholds, policy constraints, or fail verification – routing them to human agents with full session context.

Additionally, the recent launch of Agent-to-Agent (A2A) protocol has made it easier to build interacting multi-agent reasoning (e.g., combining loyalty eligibility with ticketing policies). Compass AI is now embedding a lightweight **Agent-to-Agent (A2A) coordination layer**, enabling specialized agents to exchange context and resolve complex workflows collaboratively.

Internal Deployment & API-First Integration

Compass AI is deployable in fully private or hybrid environments:

- Self-hosted on cloud VPCs (AWS/GCP/Azure) using containerized LLM runtimes (e.g., Ollama, vLLM, LM Studio).
- No customer data ever leaves your environment – API integrations are local, and logs are fully auditable.
- Frontend-agnostic design – can be embedded in voice IVRs, web chat, agent tools, or mobile apps.

All modules expose RESTful APIs or Lang Chain-style tool wrappers, enabling orchestration within broader automation pipelines. Agent interoperability is further supported through the A2A protocol, allowing individual modules or agents to operate independently while collaborating on multi-step tasks.

Extensibility for Emerging Use Cases

Compass AI is intentionally modular and pluggable, making it easy to:

- Add new RAG sources (e.g., knowledge base updates, partner APIs).
- Introduce new intents or workflow branches.
- Swap LLM providers based on pricing, latency, or model performance.
- Deploy new agents (e.g., contract summarizer, trip planner, review synthesizer) without architectural changes.

Each new module shares a common interface layer, allowing consistent context sharing through MCP and inter-agent collaboration via A2A – making Compass AI a unified platform rather than a patchwork of disconnected tools.

Security and Compliance

Security and data governance are central considerations when applying AI within the complex operational landscape of the travel industry. With systems that span bookings, loyalty programs, payments, and sensitive customer data, any AI capability introduced into this environment must be engineered with careful attention to control, transparency, and deployment integrity. Compass AI is structured with these priorities in mind. Its architecture is built to support **secure, locally hosted deployment models**, giving organizations full visibility into how data is accessed, processed, and used across AI modules. Whether integrated with booking systems, internal

APIs, or document repositories, Compass AI is designed to operate inside the enterprise perimeter – reducing exposure risk and enabling alignment with internal policies and infrastructure constraints.

Architectural Safeguards and Control [UY5] Mechanisms

- **Private Deployment Support**
  Compass AI modules can be hosted within your cloud environment or on-premises infrastructure, allowing inference, retrieval, and orchestration to run entirely within your control. External API calls or third-party LLM endpoints are never required unless explicitly configured.

- **Authentication and Access Integration**
  Each module is compatible with enterprise authentication standards, such as SSO, token-based access, and identity management systems. Role-based permissioning can be layered onto workflows to ensure controlled exposure across internal teams.

- **Event Logging and Monitoring Hooks**
  The system supports integration with internal logging and observability frameworks. All interactions – including query handling, data retrieval, and model decisions – can be surfaced for auditing, traceability, and iterative tuning.

- **Policy-Aware Processing Pipeline**
  A multi-layered verification system allows organizations to encode internal logic for query validation, escalation handling, and information access control before any response is generated. This logic is extensible and can be tuned to reflect internal governance standards or operating thresholds.

- **Ephemeral Data Handling by Default**
  Compass AI modules are configured for transient processing unless otherwise specified. Inputs are not stored, cached, or reused outside the immediate context of the task, aligning with common principles for controlled, short-lived data workflows.

These architectural principles reflect its current and future design direction: one that prioritizes integration without exposure, insight without overreach, and intelligence without compromising control.

This approach positions Compass AI to operate confidently within the security-conscious environments of modern travel organizations – while giving stakeholders the tools to shape how AI functions inside their systems.

Use Cases and Real-World Applications

Compass AI is built to deliver value incrementally – one use case at a time. Its modular design allows organizations to start with targeted capabilities and scale naturally into broader workflows. Whether enhancing customer experiences or driving back-office efficiency, Compass AI enables AI adoption that is measurable, context-aware, and built for long-term integration.

Below are practical applications already driving value across the travel industry:

Customer Service Automation

Compass AI powers a production-ready assistant that handles common traveler questions, such as:

- "What is my booking status?"
- "How many points do I have?"
- "Can I upgrade?"
- "What is the baggage policy for United?"

R Responses are grounded in real-time data using RAG and escalated to agents when needed, complete with session context.

Agent Assist Tools

These lightweight, API-driven modules plug into existing CRMs and support platforms to:

- Summarize chat histories for agent handover.
- Suggest context-aware responses.
- Condense lengthy ticket threads for faster resolution.

Intelligent Ticket Routing and Email Triage

Compass AI automates message classification and routing by extracting key metadata (intent, urgency, customer identity). It also flags edge cases for manual escalation – reducing triage time and error rates.

Knowledge Interpretation & Search Assistants

Users can query internal SOPs, policy documents, or FAQs conversationally. Compass AI retrieves and summarizes relevant passages using semantic search + RAG, enabling questions like:

- "What is our refund policy for weather delays?"
- "Summarize this 20-page vendor agreement."

Marketing Content Generation

Compass AI supports marketing teams by generating:

- SEO-optimized travel guides and blog posts
- A/B ad variants tailored to specific segments
- Campaign-ready social captions and review digests

It ensures the content is contextually relevant, tonally consistent, and ready to be reviewed – not rewritten.

Strategic Insights and Executive Briefs

Modules mine and condense unstructured data (feedback, market signals, competitive updates) into:

- Sentiment snapshots
- Weekly intel digests
- Internal leadership briefs

These outputs help decision-makers move faster with less noise.

Designed to Expand, Not Just Deploy

Each use case can be activated independently or integrated with others. There is no all-or-nothing commitment – organizations can start small and scale into a broader intelligence layer as needs evolve.

Compass AI is beyond just a toolkit. It is a platform – and it is equipped to prove its value across real operations.

Implementation and Integration

AI solutions often fail not because of technical limitations, but because they are hard to implement, inflexible to integrate, or misaligned with enterprise workflows. Compass AI was designed to avoid those pitfalls.

From day one, Compass AI has been built to **fit within existing infrastructure**, operate securely inside enterprise environments, and provide value in the lowest possible turnaround time. Whether you are starting with a single

module or planning a broader rollout, Compass AI offers a frictionless integration path tailored to your organization's readiness and tech stack.

Fast, Focused Deployments

Compass AI modules are designed to be set up and validated quickly.

Each deployment begins with:

- **Scope alignment** on the use case, data access, and operational environment
- **System integration mapping** with your APIs, CRMs, ticketing systems, or content stores
- **Prompt and model configuration** for tone, compliance handling, escalation paths, etc.
- **Onboarding, QA, and go-live** – with human-in-the-loop workflows where applicable

API-First, Platform-Agnostic Design

Compass AI is built to **plug into your ecosystem**, not replace it. Every module exposes REST APIs and supports event-based or callable workflows via:

- Webhooks
- JSON endpoints
- LangChain / agent wrappers

It can integrate with:

- CRMs and support platforms (e.g., Salesforce, Zendesk, Freshdesk)
- Booking and loyalty systems (via secure API connectors or SQL readers)
- Document repositories and knowledge bases (SharePoint, Notion, Google Drive, etc.)

Secure, Configurable Infrastructure Options

Compass AI supports multiple hosting models:

- **Cloud deployment within your VPC**
- **On-premise containers** using Docker/Kubernetes
- **Hybrid cloud inference** (e.g., self-hosted RAG + cloud-based LLM)

No data leaves your environment unless explicitly configured. All processing can happen behind your firewall, and identity integration (SSO, OAuth, IAM) ensures tight access controls.

Minimal Overhead, Maximum Control

Each module comes with:

- Constant deployment support
- Logging and observability hooks
- Versioning controls and rollback options
- Built-in fallbacks for edge case handling and human intervention

This allows your IT and operations teams to test, monitor, and scale Compass AI without introducing significant operational risk or complexity.

Compass AI is built to meet you where you are – technically and operationally. You do not need to overhaul your stack or commit to a long-term roadmap upfront. Start with the use case that matters most, integrate on your terms, and expand only when you are ready.

Future-Proofing and Scalability

AI in travel is not a one-time implementation – it is a journey. Compass AI was designed to evolve alongside that journey, offering organizations the ability to scale thoughtfully, adapt to new challenges, and continuously derive value as business needs and customer expectations change.

Where many AI tools become obsolete, inflexible, or unmaintainable, Compass AI is engineered for long-term relevance.

Designed for Long-Term Adaptability

Our solution is modular architecture ensures flexibility, allowing your organization to easily add new capabilities, expand into additional service areas, or accommodate evolving customer preferences without significant redevelopment or operational disruption. This modular design future-proofs your investment, ensuring sustained relevance and performance as market conditions shift.

Modular Growth, Not Lock-In

Every capability in Compass AI – from conversational assistants to agent tools, from content generators to strategic summarizers – is packaged as a discrete module. You can deploy one, several, or all of them, without being forced into a monolithic architecture.

This allows you to:

- Start with a narrow use case (e.g., knowledge search, booking assistant)
- Add new modules over time as confidence grows
- Extend across departments (e.g., operations → marketing → leadership tools)
- Make AI automations interact with each other and with existing systems to deliver seamless experience to business users and customers alike.

Model and Toolchain Flexibility

Compass AI is **LLM-agnostic and infrastructure-neutral**, meaning:

- You can switch between different open-source or commercial models (Mistral, GPT, Claude, etc.)
- You can host models internally, use API-based access, or run hybrid setups
- You are not locked into a single cloud provider, framework, or vendor

This flexibility ensures you can adapt your AI stack over time without rebuilding core systems.

Continuous Learning and Optimization

Compass AI's retrieval pipelines and prompt strategies are continuously tuned based on real-world usage:

- Query logs inform prompt optimization and UX tweaks
- Feedback loops enable rapid iteration on ambiguous or failed queries
- Modules can evolve through LoRA/PEFT finetuning or retraining when applicable

This ensures your AI capabilities improve the more you use them – without needing full reimplementation.

Built to Absorb New Use Cases

As your organization grows, so will the number of touchpoints where AI can add value. Compass AI is ready for:

- New domains: travel insurance, hospitality, destination services
- New channels: WhatsApp, mobile apps, airport kiosks, internal dashboards

Its pluggable architecture makes it possible to **extend intelligence into new frontiers without breaking what already works**.

Compass AI is not a fixed product with a shelf life – it is a dynamic platform meant to grow with your enterprise. Whether your goal is to automate more, serve faster, learn deeper, or innovate ahead of the curve, Compass AI is structured to scale intelligently and sustainably.

Conclusion

AI in travel is both possible and an imperative. But real impact comes not from one-off pilots or siloed tools. It comes from a holistic approach: modular, secure, and built to evolve.

Compass AI is that solution. Purpose-built for the travel industry, it brings together advanced language models, secure retrieval systems, and a flexible architecture that adapts to the way your organization works – not the other way around.

We have proven its value through successful deployments – starting with a conversational assistant that resolves real booking, loyalty, and account queries in production environments. And we are just getting started.

Whether you are exploring AI for the first time, looking to unify fragmented initiatives, or scaling what already works, Compass AI offers a clear, structured path forward.

Let us build the future of travel intelligence – together.

Take the next step:

- **Talk to Us** – Experience first-hand how our Compass AI can transform your business.
- **Talk to a Solutions Architect** – Speak directly with our engineering team and explore how we can solve your travel business specific problems.

Let us shape the future of travel support together.

CONTACT US

+ 1 909-494-6213 | 8530 Eagle Point Blvd, Suite 100, Lake Elmo, MN, 55042

Reach out to us here.