

White Paper The Agentic Revolution in Enterprise Software

March / 2026

Introduction

Cerebri AI Inc. is a B2B SaaS company that uses artificial intelligence to provide enterprise travel programs with state-of-the-art T&E analytics driven by curated customer data processed on our Cerebri AIQ data engineering platform. We also perform audit services for travel programs, and with the purchase of Interplx in Q4/2025, we now process expense reports and employee reimbursements.

In 2026, Cerebri AI will roll out agentic AI T&E agents, which we believe will revolutionize how enterprise travel managers work and how they manage their complex travel programs.

Cerebri AI T&E agents will expand our product offerings to include budget tracking, auto-authorization of T&E expenses, RFP management, automated audit services, and contract compliance tracking.

Our best-in-class T&E data engineering ensures this transition is seamless. In effect, only an AI-first architecture can power agents that decide, act, observe, and reflect in real time, delivering outcomes legacy systems were simply never built to achieve.

Agents Are Replacing Dashboards

For over two decades, dashboarding systems such as PowerBI and Tableau have been the "windows" of choice into a company's data. However, we are now seeing a shift from **visual monitoring** of dashboards to **agentic action**. Dashboards, while visually appealing, have become a bottleneck. The transition to agents is a fundamental change in how we work, where speed and "doing" take precedence.

From "Looking at What Happened" to "What Do I Do?"

The biggest weakness of traditional dashboards is that they require us to look at a chart or other visual treatments of our data, identify an event or trend, and then decide on a response:

- **The Dashboard Way:** Data shows that a department's off-program T&E spending (leakage) is higher than that of other departments.
- **The Agentic Way:** An off-program booking is identified in our TMC data feed, our agent sends a Teams message: *"You should review the need to book your hotel with Expedia, we have a discounted nightly rate including free parking at the hotel selected."*

Features	Traditional Dashboards	AI Agents
Effort →	Pull: You must log in, filter, and find insights	Push: Agents monitor data, notify you what matters
Insight →	Descriptive: Shows <i>what</i> happened in the past	Prescriptive: Explains <i>why</i> and suggests <i>what to do</i>
Speed →	Latent: Often relies on nightly batch updates	Real-Time: Acts on streaming data as events occur
Action →	Passive: You must leave the tool to execute a task	Active: Can trigger workflows (e.g., budget approval)
User Skill →	Specialized: Requires "software literacy" to interpret	Universal: Accessible via natural language (chat/voice)

Typical dashboard users must wade through multiple screens before getting the answers they need. AI agents act as an intelligent assistant, and by accessing large language models, they understand the context of the task at hand and only surface essential data, eliminating "hunting" for information across dozens of tabs.

Dashboards are **observers**; agents are **participants** that continuously monitor data, detect patterns, make decisions, and provide proactive insights.

In summary, dashboards show **what happened**, while dynamic agents explain **why it happened** and suggest **what we should do about it**. Agents can communicate with other agents, presenting a near limitless number of agentic solutions that can be quickly implemented as needs are identified.

Well-designed agentic systems anticipate agents "talking" to each other and use **orchestration agents** to enable greater automation and more complex problem solving. Orchestration of multiple agents and other tools enables smaller workflows to be created.

For instance, as outlined above, our travel program hotel agent could send a traveler the list of preferred properties and contracted rates for a location being visited when flagging an off-program booking detected by another agent.

Agents are built to handle **ad hoc discovery** by understanding the context of our questions or tasks. Developing dashboards requires our designers to be subject-matter experts to pre-stage or anticipate the answers we need.

Agents eliminate this need for pre-staging as they understand the context of the input directly from us and our enterprise environment.

In summary, dashboards will continue to be useful in doing deep dives into enterprise data repositories and for audit purposes, while agents will define our user experience with enterprise systems:

- **Dashboards:** quarterly reviews, strategic planning, checking performance of agentic solutions.
- **Agents:** day-to-day execution, anomaly detection, rapid answers.

Data Quality

Why are we inundated with so much "imperfect" data? The answer has traditionally been that bad data is expensive to fix, and in the end, humans are there to interpret it.

In **agentic AI**, data quality isn't a "nice to have," it is the **operating safety-net** of the system.

Automating decisions with agents yields more consistent results, provided the data inputs are of high quality. With dashboards, the rule was **"garbage in, garbage out."**

With agents, it's **"garbage in, we're stuck."** Low-quality data cannot be used with confidence in automating decision-making, especially at real-time speeds.

Dashboards vs. Agents: The Data Quality Bar

Requirements	Traditional Dashboard	Agentic AI
Tolerance for Error →	Higher: Humans can filter out errors, noise	Lower: Errors trigger inaccurate predictions, actions
Data Latency →	Higher: 24-hour data drops are very common	Lower: Ideally lag times, real-time ideal
Structure →	Flexible: Analysts often "clean" data in Excel	Stricter: Predictable & consistency data quality
Governance →	Variable: Difficult to rely on ad hoc data changes, reporting	Stricter: Automated decisions require consistency

Agentic AI System Design

In designing and bringing to market agentic AI solutions, industry consensus has shifted toward a **specialized, modular approach** as the foundation for robust vertical applications.

While a **"super-agent"** (often called a **meta-agent** or **supervisor**) may be feasible, building a single monolithic entity is now considered a significant technical risk.

Building enterprise-wide consensus that agentic AI solutions can drive bottom-line results requires consistently "profitable" results. Specialized agents drive higher accuracy, are easier for developers to debug, and usually offer lower latency. In most circumstances, specialized agents are far more likely to avoid "hallucinations."

AIQ Data

Cerebri AI agentic AI T&E solutions start with state-of-the-art **data engineering** and **AIQ Data**. T&E data is notoriously inconsistent and unreliable.

AIQ Data takes in T&E data in any format, from daily or weekly file transfers to real-time data feeds via high-volume APIs. Key datasets include card, expense, TMC, and HR hierarchy data.

AIQ Data feeds Cerebri AI agents and third-party AI booking agents. Cerebri AI does not book travel but expects that new agentic AI booking tools, such as BizTrip.ai, will make serious inroads in the traditional OBT marketplace for the same reasons agents will largely replace the use of dashboards.

Agentic booking tools need curated travel data for the same reasons travel managers do to manage other aspects of their travel programs.

Traveler preferences do not have to be "learned" by new agentic booking tools; they are already encapsulated in our AIQ Data repository.

In summary, AIQ Data will be accessible using our new AI T&E agents and third-party agentic booking tools.