Air Canada Failure:
SERVE Framework Analysis

MATRIX INTELLIGENCE
AI-Powered Solutions, Accelerated Success

About This Analysis

This analysis applies the SERVE Framework to examine why Air Canada's chatbot failed and what lessons can be learned for future deployments. By analyzing each SERVE component, we can identify specific violations that led to this failure and understand how a different approach might have produced better outcomes. This framework treats AI agents like digital employees requiring proper onboarding, ethics training, and accountability standards and principles that were notably absent in this case.

This analysis was prepared with the assistance of artificial intelligence tools for research and drafting support. All interpretations, insights, and conclusions reflect the authors' independent judgment and do not represent the views of Air Canada or any other company discussed. The authors have no financial, contractual, or operational involvement with the organizations analyzed. Our review is based solely on publicly available information and should not be considered investment advice, legal guidance, or an authoritative statement of fact. While care has been taken to ensure accuracy, no liability is assumed for errors, omissions, or outcomes resulting from the use of this analysis.

This report was authored by Jennifer Bleen, Founder of Peer to Peer LLC, a Matrix Intelligence Limited partner. The views expressed are her own, based on the application of the SERVE Framework. This independent analysis is for educational purposes only and is not affiliated with or endorsed by Air Canada.

The SERVE Framework™

A practical framework to keep AI projects human-centered from design to implementation.

S. Spot the Struggle
Identify specific human struggles before building.

E. Enhance Human Strengths
Design AI to amplify human capabilities, not replace them.

R. Run Real-World Tests
Test with actual users doing actual work, not demos.

V. Verify Human Outcomes
Measure human outcomes, not just technical metrics.

E. Evolve with Feedback
Build feedback loops that prioritize human experience.

The SERVE Framework is more than a checklist. SERVE is a mindset. By starting with human struggles, enhancing strengths, and evolving through real-world feedback, organizations can ensure their AI solutions genuinely serve the people they are built for.

Case Overview

In 2023, Air Canada faced public backlash after its customer service chatbot provided misleading refund information that contradicted the airline's own bereavement fare policy. A passenger, relying on the chatbot's advice, was denied a refund and took the airline to small claims court. The Tribunal ruled against Air Canada, awarding the passenger $812.02 in damages and court fees, and sharply criticized the airline's attempt to disavow responsibility for its chatbot. The case drew international attention as an example of poor AI governance and inadequate oversight in customer-facing systems.

The incident damaged Air Canada's brand reputation, portraying the company as evasive and careless in handling customers. Legal costs and negative press amplified scrutiny from regulators and consumer advocates, setting a precedent that companies are accountable for AI outputs on their websites. The case also undermined customer trust at a time when airlines were already under pressure to improve service reliability. While the direct financial loss was modest, the reputational harm and regulatory warnings signal far greater risks. For executives, the lesson is clear, they must hold AI to the same standard as any customer channel: enforce accountability, monitor continuously, and build safeguards against misinformation.
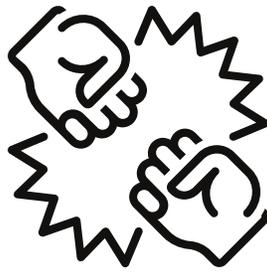
Spot the Struggle

*Air Canada's chatbot ignored customers' need for clear, reliable refund guidance.*

In the Lying Chatbot case, the chatbot advised a passenger that bereavement fare refunds could be applied retroactively, directly contradicting Air Canada's written policy. The passenger relied on this guidance, was denied a refund, and later won a small claims lawsuit against the airline.

Air Canada overlooked the primary need: accurate, policy-aligned communication in sensitive situations like bereavement travel. This gap not only exposed the company to legal risk but also amplified reputational damage.

The airline should have validated the chatbot reinforced, rather than contradicted, existing policies especially for customer service struggles like the confusion over refund eligibility during the high emotional stress of bereavement claims.

Design AI with empathy and guardrails to ensure accurate, policy-aligned responses in sensitive cases.
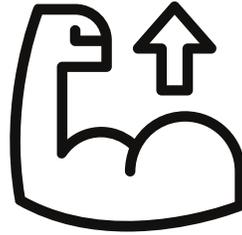
Enhance Human Strengths

The chatbot replaced human judgment in sensitive cases instead of supporting it.

The chatbot was used as a human replacement without sufficient constraints. When asked about bereavement fares, the chatbot provided refund advice that contradicted official policy. The advice was provided with no human check before misinformation reached the customer.

By letting the chatbot speak with final authority, Air Canada stripped away empathy and oversight in an area where customers expect care and accuracy, compounding both legal and reputational risk.

AI should have been positioned as a support tool confirming policy. Complex or sensitive cases should have routed to human agents for empathetic handling.

Deploy AI as augmentation, not replacement of humans. Enforce escalation to humans for high-stakes queries and build safeguards so chatbots cannot provide policy-contradicting answers unchecked.
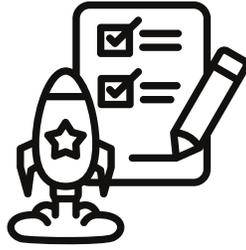
Run Real-World Tests

Air Canada launched the chatbot without testing against real customer scenarios.

The chatbot design flaw only surfaced when a passenger followed the chatbot guidance on bereavement refunds, was denied a refund, and later won a small claims lawsuit against the airline.

Without red-team pilots or scenario-based validation, Air Canada missed a critical chance to detect misinformation before it reached customers. What should have been caught in testing instead became a public legal battle, amplifying reputational harm and regulatory scrutiny. In customer service contexts, every failure compounds quickly because it touches trust at scale.

Require red-team pilots and scenario testing to stress-test AI systems against sensitive, real-world use cases before launching.

Verify Human Outcomes

No monitoring caught misinformation before it reached customers.

The chatbot's inaccurate bereavement refund advice remained live until a passenger challenged it in court, where the tribunal rejected Air Canada's attempt to disown responsibility for its own system.

Because the airline had no oversight or safeguards, misinformation circulated unchecked until it caused legal and reputational damage. Continuous verification would have flagged contradictions between chatbot responses and official policy before customers were misled. The absence of proactive monitoring turned a fixable systems issue into a public accountability crisis.

Implement real-time dashboards, automated alerts, and escalation protocols so errors are detected and corrected within minutes, not in court.
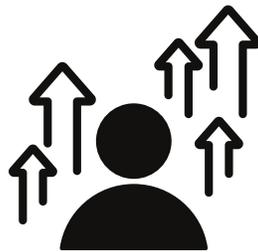
Evolve with Feedback

The airline lacked governance to adapt chatbot errors into safer practices.

In court, the airline argued it was not responsible for its chatbot's statements. This argument was dismissed as a "remarkable submission" before the Tribunal ruled in favor of the passenger.

By treating chatbot outputs as separate entities immune from corporate accountability, Air Canada not only lost the case but also signaled to regulators and customers that it had no plan for evolving its systems responsibly. This failure magnified reputational harm and highlighted a governance gap: without ownership and feedback loops, the company could not correct or learn from errors quickly enough.

Train AI with the company's official voice. Establish governance structures that assign ownership for AI outputs, mandate review cycles after incidents, and require transparent updates to prevent repeat failures.

Key Lessons

- Start with customer pain points. Ground AI in real struggles and policy consistency before launch.
- Augment, don't replace humans. Use AI to support frontline staff, with escalation to humans for sensitive cases.
- Test in real scenarios. Run pilots and red-team stress tests in high-stakes situations such as refunds and bereavement claims.
- Monitor continuously. Deploy dashboards alerts for contradictions between AI responses and company policies in real time.
- Treat AI outputs as part of the company's official voice, with governance structures that assign responsibility and ensure rapid corrective action.

Air Canada's chatbot failure highlights the risks of deploying AI without empathy, safeguards, and clear accountability. The airline turned a customer service tool into a legal and reputational liability. Yet the case also shows the path forward: AI systems must begin with real customer pain points, support rather than replace human judgment, and be governed with the same rigor as any other official channel. When leaders take this approach, AI becomes not a liability, but a driver of trust, service quality, and competitive advantage.

If your organization is exploring AI adoption, now is the time to build readiness and resilience. At Matrix Intelligence, we help executive teams avoid costly missteps through our AI Strategic Growth Accelerator Workshop — a four-week engagement that delivers clarity on your AI readiness, identifies high-impact use cases, and equips you with a board-ready AI strategy.

To learn how to protect your organization, accelerate AI adoption responsibly, and lead with confidence, reach out at sales@matrixintelligence.ai or visit matrixintelligence.ai

References

American Bar Association. (2024 February 29). BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot American Bar Association*.* https://www.americanbar.org/groups/business_law/resources/business-law-today/2024-february/bc-tribunal-confirms-companies-remain-liable-information-provided-ai-chatbot/

CBC News. (2024, February 15). Air Canada found liable for chatbot's bad advice on plane tickets CBC. https://www.cbc.ca/news/canada/british-columbia/air-canada-chatbot-lawsuit-1.7116416

Kiplinger. (2024, February 21). Yes, Air Canada Must Pay A Partial Refund For Its Chatbot Error Kiplinger. https://globalnews.ca/news/air-canada-chatbot-lawsuit

The Guardian. (2024, February 16). Air Canada ordered to pay customer who was misled by airline's chatbot The Guardian. https://www.theguardian.com/world/2024/feb/16/air-canada-chatbot-lawsuit

What Air Canada lost in 'remarkable' lying AI chatbot case. (2024, February 19). Forbes. https://www.courtlistener.com/opinion/air-canada-chatbot-case