



**MATRIX
INTELLIGENCE**

AI-Powered Solutions, Accelerated Success

IBM Watson for Oncology Failure: SERVE Framework Analysis

TREATMENT PLAN:
HIGH-DOSE RADIATION
(FATAL)
CHEMO PROTOCOL:
USE TOXIC COCKTAIL
NO SURGERY

About This Analysis

This analysis applies the SERVE Framework to examine why IBM's initial Watson for Oncology AI implementation 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 IBM Watson for Oncology 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 IBM Watson for Oncology.



The SERVE Framework™

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



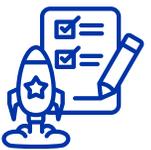
S. Spot the Struggle

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.

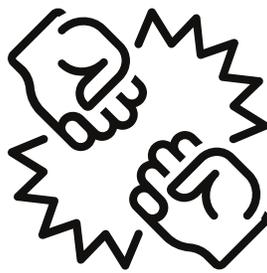
The SERVE Framework is more than a checklist. SERVE is a mindset. By starting with human struggles, enhancing strengths, and evolving through real-world feedback, organizations can ensure their AI solutions genuinely serve the people they are built for.

Case Overview

Between 2012 and 2018, IBM invested billions in Watson for Oncology, an AI system marketed with Memorial Sloan Kettering as a clinical-decision tool to assist doctors in creating cancer treatment plans. The product was positioned as a breakthrough in applying AI to medicine, promising to bridge the gap between research and clinical practice.

Instead, internal presentations later revealed that Watson for Oncology frequently generated “unsafe and incorrect” treatment recommendations, some of which contradicted national cancer guidelines and exposed physicians to malpractice risk. These issues triggered customer dissatisfaction, contract cancellations, and mounting skepticism among clinicians.

The collapse of Watson for Oncology carried significant consequences for IBM. Despite aggressive marketing, adoption stalled at only a few dozen hospitals worldwide, far below expectations. High-profile clients, including MD Anderson Cancer Center, exited partnerships over underperformance and cost concerns. Investigations by *STAT* and *The Wall Street Journal*, combined with critical physician feedback, eroded IBM’s credibility in healthcare AI. Ultimately, the effort produced poor returns on a multi-billion-dollar investment, contributing to restructuring and the eventual sale of key Watson Health assets.



Spot the Struggle

IBM prioritized a technology showcase over addressing oncologists' real clinical needs.

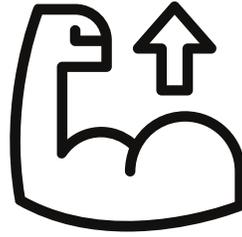
Watson for Oncology was marketed to “bridge the gap between research and clinical practice,” but its early design focused on showcasing AI capabilities rather than addressing oncology workflow needs such as guideline alignment, trial access, or complex patient data (Dolfing, 2024).

This gap between promise and functionality eroded clinician trust and limited adoption.

A human-centered approach would have started with oncology discovery—ethnographic research, shadowing tumor boards, and workflow mapping—to uncover real struggles. AI use cases should then have been green-lit only when tied to measurable burdens like faster guideline adherence or improved trial matching.



Require every AI initiative demonstrates clear evidence of stakeholder discovery before approving funding. Prioritize use cases where success can be measured in real reductions to user burdens, not in abstract AI capability.



Enhance Human Strengths

IBM positioned Watson to replace physician judgment instead of supporting it.

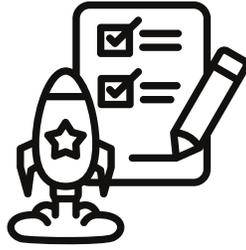
Watson produced ranked treatment lists “much like Amazon star ratings,” with some recommendations contradicting national cancer guidelines and exposing patients to unsafe advice (STAT News, 2018).

This framing treated AI as a substitute for medical expertise rather than a partner in clinical decision-making, eroding physician confidence and adoption.

A human-centered approach would have designed Watson as a co-pilot or junior employee focused on surfacing relevant literature, highlighting contraindications, and preserving physician authority in final treatment choices.



Require all AI systems be framed as augmentation, not replacement.
Success should be measured by how effectively AI enhances expert decision-making while keeping final responsibility with humans.



Run Real-World Tests

Watson was trained on synthetic scenarios instead of diverse real-world patient cases.

Internal documents revealed training relied on “synthetic” cases created by one or two physicians without statistical grounding, limiting Watson’s accuracy and clinical relevance (STAT News, 2018).

This shortcut undermined credibility because outputs failed to reflect the complexity and diversity of real cancer care across institutions and patient populations.

A human-centered approach would have piloted Watson with de-identified, real-world cases from multiple hospitals and diverse populations, paired with adversarial testing for edge cases.



Mandate real-world pilot testing with diverse, representative data before scaling any AI product. Training only on synthetic scenarios is insufficient for clinical-grade deployment.



Verify Human Outcomes

IBM lacked systematic monitoring to ensure recommendations were safe and aligned with medical standards.

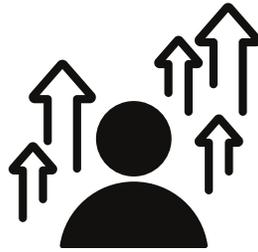
A physician at Jupiter Medical Center told IBM, “This product is a piece of s—. We bought it for marketing ... We can’t use it for most cases,” showing the system was marketed despite ongoing accuracy failures (Healthcare Dive, 2018).

Because verification was reactive rather than proactive, unsafe recommendations persisted until flagged by frustrated clinicians.

A human-centered approach would have implemented continuous monitoring against clinical guidelines, with immediate escalation for questionable outputs and real-time physician oversight.



Require continuous monitoring and escalation protocols for every AI deployment. Build dashboards and guardrails that surface unsafe outputs within minutes, not months.



Evolve with Feedback

IBM prioritized favorable optics over objective evaluation of Watson’s clinical value.

Internal presentations admitted studies were designed “to generate favorable findings,” prioritizing market perception instead of transparent validation of effectiveness (Advisory.com, 2018).

This choice eroded credibility in the medical community, as Watson appeared to evolve for marketing gains rather than patient safety and clinical trust.

A human-centered approach would have established medical ethics boards, preregistered evaluation protocols, and external advisory committees to guide responsible AI deployment.



Always require governance structures such as ethics boards, advisory panels, and transparent evaluation protocols before bringing AI products to market. Marketing outcomes must never outweigh patient safety.

Key Lessons

- Start with clinicians' jobs to be done. Fund discovery first; don't ship a demo as a product.
- Augment, don't automate judgment. Mandate human-in-the-loop and escalation for high-risk outputs.
- Test in the wild. Use multi-site, de-identified real-world data; include adversarial and regional reviews pre-launch.
- Monitor continuously. Build real-time guardrails and guideline-concordance checks with rapid human escalation.
- Govern like a medical device. Independent ethics review, preregistered evaluations, and publishable evidence before scale.

The failure of IBM's Watson for Oncology illustrates the risks of applying AI in high-stakes fields without grounding in human-centered safeguards. By emphasizing marketing and technical showcase over clinician needs, real-world testing, and governance, IBM eroded trust in healthcare AI. Yet the case also demonstrates what's possible when organizations start with clinician pain points, augment rather than replace expertise, and ground AI in ethics. Done right, AI can become a driver of trust, safety, and strategic 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

Advisory Board. (2018, July 27). IBM's Watson recommended 'unsafe and incorrect' cancer treatments, STAT report finds. [Advisory.com](https://www.advisory.com).

<https://www.advisory.com/daily-briefing/2018/07/27/ibm>

Dolfing, H. (2024, December 7). Case study 20: The \$4 billion AI failure of IBM Watson for Oncology. Henrico Dolfing.

<https://www.henricodolfing.com/2024/12/case-study-ibm-watson-for-oncology-failure.html>

Healthcare Dive. (2018, July 26). STAT: IBM's Watson gave 'unsafe and incorrect' cancer treatment advice. Healthcare Dive.

<https://www.healthcaredive.com/news/stat-ibms-watson-gave-unsafe-and-incorrect-cancer-treatment-advice/528666/>

STAT News. (2018, July 25). IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show. STAT.

<https://www.statnews.com/2018/07/25/ibm-watson-recommended-unsafe-incorrect-treatments/>