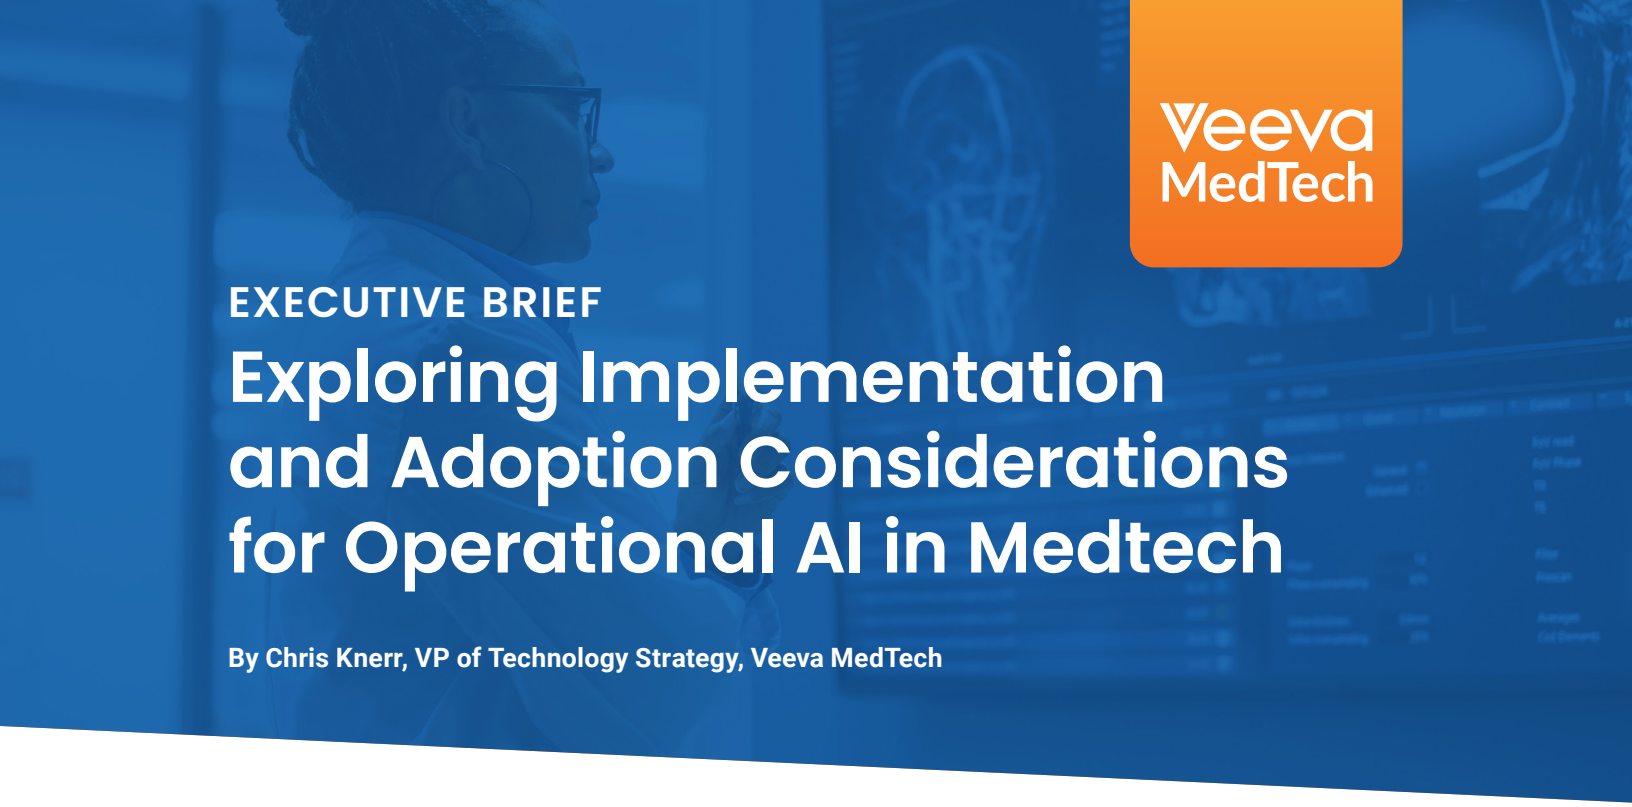

EXECUTIVE BRIEF

Exploring Implementation and Adoption Considerations for Operational AI in Medtech

By Chris Knerr, VP of Technology Strategy, Veeva MedTech

Key Takeaways

- **Operational AI is now a key C-level strategic roadmap initiative for medtech organizations—in addition to clinical AI**
- **Maturity varies by company, but CIOs are broadly experimenting with and/are actively operationalizing a range of advanced technologies and use cases**
- **Implementation considerations—data strategy, security, and compliance—are paramount for scaling and driving business value.**

At the recent Veeva MedTech Summit in Austin, TX, I had the honor of leading an executive roundtable with several CIOs and senior IT leaders on the impact of AI in medtech operations. The group spanned mid- to large-cap segments with wide-ranging device and diagnostics portfolios, and our conversation centered around five key considerations from Veeva's recent whitepaper "AI is a Starting Point, Not a Magic Wand":

1. *Are we clear on the fit of the technology to potential use cases?*
2. *Do we have the right data to train AI models?*
3. *How can AI eliminate or reduce manual work?*
4. *What risks does AI introduce?*
5. *How can we leverage AI and ML technology to spur further innovation?*

The dialogue focused on pragmatic implementation and operationalization including how changes in the market—post-pandemic acceleration of digitization and the pace of technology and engineering change, in particular—are driving C-level appetite for broader and faster adoption of AI in medtech.

Here is a summary of the key findings.

Build the foundation, then pace and scale

CIOs are rolling out AI for operations in a phased manner. A typical pattern is to deploy in-house versions of current technologies, e.g., large language models (LLMs) available from OpenAI, Meta etc., and provision them behind enterprise firewalls.

In-house LLMs and scaled machine learning (ML) models enable experimentation with both business users, for enhanced enterprise search, particularly in large bodies of documents and enterprise unstructured data stores, as well as technically oriented assisted coding for IT teams. Paced rollouts help manage the pace of change thoughtfully, allowing business and technical users to adapt to the new capabilities without overload.

With the foundation built, IT leaders are enabling their organizations to start envisioning the efficiencies and growth opportunities that operational AI can provide—an internal crowdsourcing of use cases that address specific opportunities and pain points, as opposed to generic AI "hype cases" which may or may not prove valuable.

SPECIFIC USE CASES

- **Augmented Coding / Development**
- **Open-ended requests: Q&A with specific ` source attribution to trusted enterprise or third-party data, such as:**
  - **Who are our global product vendors for specialty plastics? By region?**
  - **What are our operating margins by product family?**
- **Automation of routine, but sophisticated tasks, such as creating draft clinical study reports based on summaries of large data sets**
- **Helpdesk ticket automation to reduce effort and cycle time**
- **Simplification and best practice enforcement of business workflows**

In tandem, IT organizations are building out teams of data scientists and data engineers to develop, deploy, support, and industrialize AI applications. To support this, governance and executive sponsorship across the business and IT are critical.

While CIOs have some breathing room around value realization, we are at a stage where pressure is increasing to identify and quantify real business value delivered to the CEO and the Board by monetizing data.

2 Data rules

The struggle to collect, curate and validate data remains acute. Data interoperability among systems of record, legacy systems, and third-party data sources presents a continued challenge, as does the speed of data refresh updates.

Security is a growing concern and raised a number of specific challenges within our senior IT leader group. In particular, because data extracts into a data lake denormalize the data, it's often difficult to preserve underlying security permissions from source to target. Thus, least access permissions, commonly defined at the role-object level, are not inherited in the data lake target without additional complex design and implementation work. By default, any user with access can see all the data, which is a genuine problem for data loss prevention programs and data privacy compliance. Region- and country-specific data protection rules further compound the issue.

To address this challenge, there are multiple foundational security approaches. These should exist in parallel, constituting a bimodal, or multi-modal, AI/data lake security model.

- In the platform out pattern, allow broad data access with security only at the source extract level to data scientists and engineers, with stringent controls on access and on exports. (This is along the lines of a client-server "Sys Admin" model.)
- In the use case in pattern, construct tight security at the user level, composed of objects and roles in the access and presentation layers. This pattern significantly narrows down what users can do and see in order to drive data security controls.
- Organizationally, the combination of platform out and use case in works well to enable exploration of novel use cases and open-ended interrogation of data using unsupervised learning for a subset of AI explorers and designers, while avoiding excess access for ordinary business users.
- Protected health information (PHI) is typically deprecated or stripped out of enterprise data lakes to minimize PHI data protection risk.
- PHI and other ultra-sensitive data can be "tokenized," which entails replacing the atomic records with a cryptographic token, but at an added financial cost and additional data architectural complexity. (This is a common commercial technology in multi-party credit card and financial transaction processing, but not yet widely adopted within the enterprise otherwise due to cost and complexity.)
- Master data in systems of record (product master, HR data, marketing authorizations, etc.) and interoperability is a key enabler of controls in the data lake and AI models.
- Avoid public models, which have limited or no IP controls and data privacy protections. Because this is a new area, CIOs are unclear about data breach remedies, which merits caution.

In an interesting innovation, some enterprises are setting up simulation companies that are air-gapped and share no data in common with the parent. The simulation company is a test bed for AI business case development and wireframe testing.

3 Standardizing the technology stack

Implementing LLMs and scale ML models only behind the enterprise firewall is critical to the data strategy outlined above.

A typical best practice technology stack includes:

- Standardized systems of record: ERP, CRM, clinical/regulatory
- Cloud first or cloud strongly preferred
- Cloud infrastructure on one of the major cloud providers, hosting data lakes
- Contained LLMs and ML models
- Standardized consumption/presentation layers
- Standardized middleware/tech integration tool
- Pervasive security components/services

Of note, hosted LLM models typically have limited ability to further train the models on new data. Because the training stage is intensely compute/storage intensive, major training is not feasible outside of core large scale technology companies, which limits the extensibility of the models to protected enterprise data and use cases. In effect, the trained model is only as good as the non-proprietary data on which it was trained. Retrieval Augmented Generation (RAG) models are of some help here, but do not have the full capabilities of LLMs.

As opensource models available to deploy behind enterprise firewalls proliferate, this situation will likely improve and bears monitoring, as model efficiency is a keen focus of the broader tech sector and is likely to improve.

4 Establish a clear controls and risk management framework

Veeva's earlier whitepaper explored the concern that clear regulatory guidance on AI for operations is lacking. There is broad agreement that this is a risk and that developing, approving, and implementing a framework is paramount.

Since the FDA has provided strong guidance on Software as a Medical Device (SaMD) and ML change control,[1] adopting these guidelines for operational AI is an excellent starting point, especially absent other guidance.

THE FDA'S GOOD MACHINE LEARNING PRACTICE LAYS OUT THE FOLLOWING GUIDING PRINCIPLES:

1. **Multi-disciplinary expertise is leveraged throughout the total product life cycle**
2. **Good software engineering and security practices are implemented**
3. **Clinical study participants and data sets are representative of the intended patient population**
4. **Training data sets are independent of test sets**
5. **Selected reference datasets are based upon best available methods**
6. **Model design is tailored to the available data and reflects the intended use of the device**
7. **Focus is placed on the performance of the human-ai team**
8. **Testing demonstrates device performance during clinically relevant conditions**
9. **Users are provided clear, essential information**
10. **Deployed models are monitored for performance and re-training risks are managed**

Once adjusted for operational needs, these are a valid starting point for an operational AI controls framework. For example, in the sixth guiding principle, replacing "device" with "model" or "use case." Use cases and models can also be risk-stratified according to the enterprises existing CSV or CSA framework.

Similarly, the guiding principles on predetermined change control plans (PCCPs) can be modified for operations.

1. Focused and bounded
2. Risk-based
3. Evidence-based
4. Transparent
5. Total product lifecycle (TPLC) perspective

[1] FDA, Good Machine Learning Practice for Medical Device Development: Guiding Principles, 2021and Predetermined Change Control Plans for Machine Learning-Enabled Medical Devices: Guiding Principles, 2023

More important than the exact framework is the management and compliance control process to create a framework. While some of the SaMD guiding principles may not be needed for operational AI, taking them as a starting point and following a structured process to arrive at an approved controls framework will better manage and mitigate compliance and audit risk for operational AI.

5 Looking forward

While operational AI isn't yet delivering massive value for medtech organizations, industry CIOs are confident that solid use cases will eventually emerge with thoughtful oversight and paced investment. The situation is somewhat reminiscent of the internet in the late 1990s, when it was clear the internet was "something" but not "what."

Despite a broad consensus that there's not a "do nothing" strategy—and increasing pressure to monetize investments—significant runway remains to industrialize compliant AI capabilities at scale. While momentum and interest in AI for medtech operations continues, the complex work to establish appropriate risk, control, security, and data strategy capabilities still needs to mature in parallel. As this occurs, the scope of deployment and business benefits will emerge and come into focus.