28 Build and Manage an AI Operating Model

Generative AI is on a roll, showing potential applications and benefits for nearly every industry. But before the novelty wears off, several long-term strategies must be addressed.

Sustainable success with AI requires a strategic approach that considers both business, technical and human aspects. The AI Operating Model is the reference architecture for the many moving parts involved.

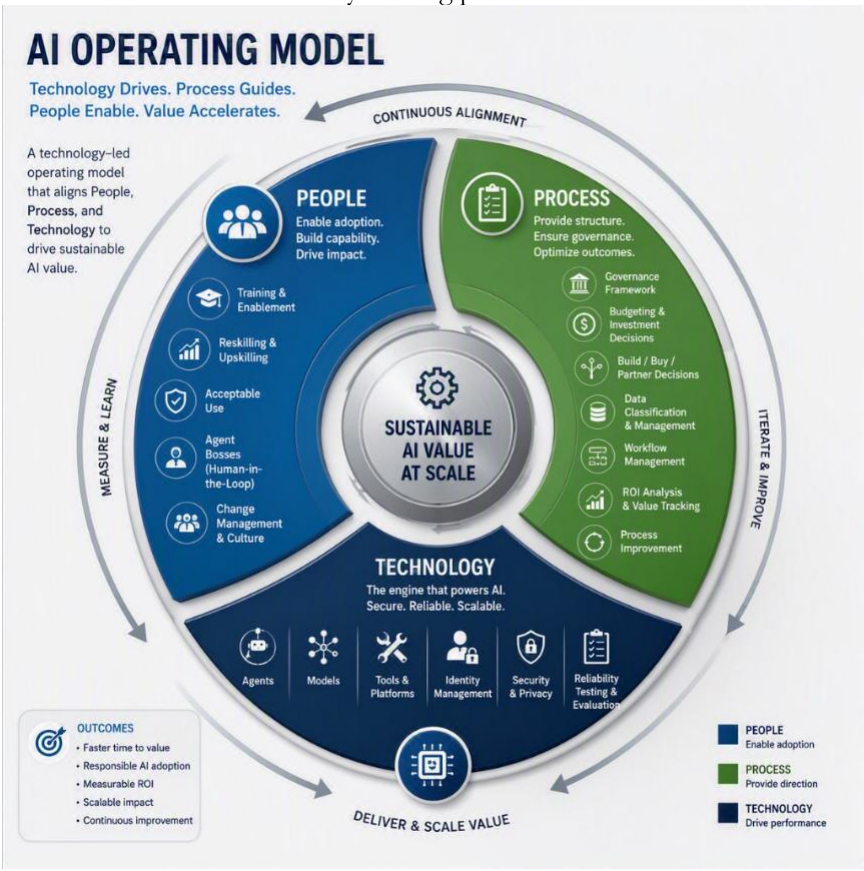

Who should lead that charge? Often the de facto AI leader becomes the CIO, but their technical vision isn't enough.

The Right Team for Sustainable Success

The first step of building a CoE for generative AI is to assemble the team.

The name of the team isn't important. In my work with hundreds of early adopters of Microsoft 365 Copilot, many have created AI "focus groups" or "launch teams." Some use tried and true "Champion Networks" or "Steering Committees," and one has even called it the "Digital Mindset Acceleration Team." Whatever the title, successful organizations are developing **a more sustainable team**.

For the rest of this chapter, I'll settle on the term "Center of Excellence" (or CoE). The CoE is a dedicated team that is built to last in the **long-term**. It differs from a launch team, which typically take narrower, short-term viewpoints.

CoEs should consist of a diverse mix of people who have relevant skills, expertise, and perspectives. The team should include the following roles:

- A leader or director, coordinating the CoE team and its activities.
- Champions, advocating for the CoE and GenAI. Often, these are power users who benefit from the solutions and data.
- AI experts with technical skills and the domain knowledge to design, develop, deploy, and maintain generative AI solutions and data.
- Enablers that provide support services. Think of IT, legal, compliance, cybersecurity, training, and corporate communications.

Make Your (AI) Statement

Several startup activities of the CoE will take time. But the organization's appetite to use AI won't wait. So quickly **pursue two quick wins:** an AI Statement and an AI Policy.

The **AI Position Statement** articulates the organization's philosophy and position on AI. At a high level, it acknowledges the responsibility of the organization and the user community to respect fairness, safety, privacy, and ethics. It also notes the limitations and uncertainties of AI and the need for human oversight and intervention. **It aligns AI efforts with the organization's mission and values.**

An **AI Appropriate Use Policy** details the guidelines for the use and development of AI in the organization. It documents the decisions to prevent or mitigate the risks and challenges associated with AI, such as bias, discrimination, misuse, or abuse. Pertinent parts of the policy should be added to the Employee Appropriate Use document so that staff agrees about their role and risks.

Executive Consideration: Some organizations' policies prohibit the use of public AI tools for business purposes. Others reduce risk by making usage of AI tools is optional, stating "By agreeing, you acknowledge that AI can be provide inadvertently incorrect or biased information and that the organization is not responsible for its results."

The statement should be only a page or so in length, while the policy may be several pages. HR, senior management, and legal should approve both.

Chart the Course

Larger organizations may task the CoE to create a charter. The charter should:

- Clarify the purpose, scope, and goals of generative AI in the organization.
- Align AI vision with the organization's values.
- Address human concerns, such as upskilling, inclusivity, and impact to jobs.
- Establish standards and expectations for generative AI development, deployment, and usage.
- Address quality, security, privacy, and fairness.

Importantly, it should be succinct enough to read, but thorough enough to build the trust and accountability of stakeholders and users.

Establish Some Goals

While the goals of each business will differ, some of the (iterative) activities of an AI CoE are listed below.

1. Align AI vision and goals with business strategy.
2. Establish an implementation and adoption roadmap.
3. Identify and evaluate the feasibility of use cases and opportunities.
4. Approve standards and decide whether to build/buy and go/no-go about new solutions.
5. Establish and monitor business metrics and outcomes.[1]
6. Manage aspects such as quality, security, ethics, fairness, and inclusivity.
7. Facilitate knowledge sharing and AI culture amongst the broader community.

Establishing a RACI (Responsible/Accountable/Consulted/Informed) will allow expeditious but balanced decision making. For instance, HR should be accountable for facilitating knowledge transfer, while IT handles standards and project feasibility.

While AI can help create a competitive advantage, the CoE must have staying power to foster sustainable innovation and improvement.

Set Milestones and Metrics

A critical task of the CoE is to develop a roadmap that outlines the organization's AI objectives, milestones, and metrics, and how to pursue them. A strategic roadmap should:

- Identify and prioritize the most impactful and feasible generative AI opportunities.
- Outline the LOB leaders who will allocate resources and assign responsibilities.
- Monitor project teams' AI activities, tasks, and deliverables.
- Evaluate outcomes and benefits.

The plan should be developed and overseen by the CoE, but execution relies on IT, stakeholders, and the extended user community.

Identify and Manage Risks

One of the **ongoing functions** of the COE is to deal with risk, and an early assessment should be the CoE's next step.

Generative AI has unique risks and challenges to assess and mitigate. Early adopters are showing a wide range of responses to such concerns. Some are racing towards the "rewards" of AI, seeking an early mover advantage; others are first tackling the "risks." Either way, the CoE should decide the organization's philosophy about where on the spectrum the organization will fall.

Specific steps to recognize and plan for risk include:

- Conducting a risk assessment and building a mitigation plan. Frameworks like NIST's AI Risk Management Framework aren't specifically built for enterprises but are good to scan to understand the scope of impact AI will have.[2]
- Implementing and enforcing the principles, guidelines as defined in the AI charter.
- Collecting and analyzing feedback and incidents from stakeholders, users, and community, using channels and mechanisms such as Teams channels, surveys, and small group feedback sessions.
- Communicating the risks and challenges to executive and HR management, and extolling the importance of quality, security, privacy, ethics, and safety across your organization.
- Creating an Acceptable Use and Expectations document for users to sign. Be clear with stakeholders: since GenAI systems aren't always correct and can sometimes be offensive.[3] IT help desks can only offer best effort support, with few SLAs.

The long-term risk of AI needs to be shared by the legal, IT, and senior leadership representatives in a CoE, not solely by the organization's technology team.

Help the CIO Prepare for Tough Decisions

If resources were unlimited, everyone who asked would receive Copilot. Realistically, the CoE (and the CIO) should agree on how to say "No" (or "Not yet") to licensing requests.

Few people get miffed, for instance, if they are assigned an F3 (Frontline) license instead of an E3. But then again, they don't read headlines about E3 creating a massive workforce shift! But they do see press that infers that AI may threaten job security.[4]

Savvy **CoEs shouldn't underestimate the potential flak from unlicensed people.** Imagine if an already skeptical associate scans a headline that says, "AI Won't Replace Humans — But Humans with AI Will Replace Humans Without AI." They then attend a meeting where someone is using Copilot to take meeting minutes. No wonder they'd feel at risk of being left behind!

CIOs will be in the hot seat if they're the only ones to say "No." Here are suggestions to stay cool.

- Use the CoE and its cross-functional team members. If the business unit doesn't have immediate use cases, their leadership should acknowledge that and support saying "No." If they do have a use case, their leadership, the CIO, and CFO should slot them into the roadmap.
- Decide as a committee (using pilot groups' data and budget constraints) which departments, roles, or business processes have enough return to **justify the investment**.
- Do **budget modeling**. Consider a model to share the cost with heavy-usage departments. See Chapter 26 on funding.
- Consider some **discretionary exceptions** (i.e. a two-level up approval exception process (boss + boss's boss)). But the exclusionary problem is worsened when some people in a department get Copilot, while others wait.
- Decide **who pays for tokens**. Recent trends have led AI providers to move to consumption based (rather than flat billing models), throwing ROI evaluations for a spin. Does the LOB or the IT department pay? The CoE should help decide.

Aim at the Target

AI Operating Models are in their infancy. At its best, a CoE can help create a competitive advantage and a positive impact from generative AI. At a minimum, it reduces the risk of reactionary headaches along the way.

However, building an operating model for generative AI is not a one-time or a one-size-fits-all process. It requires a tailored and iterative approach.

We are in the early days of defining good practices. AI CoEs are providing leadership, governance, support, and training for AI initiatives. They are also beginning to realize they should reduce liabilities by understanding, addressing, and compensating for risk.

So once the first committees, policy makers, and pilot groups finish their initial sprints, it's time to create a CoE to work toward an Operating Model for AI. **It's one way to achieve sustainable AI success in your organization.**

Coming up next, learn about some of the more complex and long-term challenges that the CoE should tackle.

References and Resources

[1] https://sloanreview.mit.edu/article/what-leaders-should-know-about-measuring-ai-project-value/
[2] https://www.nist.gov/itl/ai-risk-management-framework
[3] https://arstechnica.com/information-technology/2024/02/googles-hidden-ai-diversity-prompts-lead-to-outcry-over-historically-inaccurate-images/
[4] https://hbr.org/2023/08/ai-wont-replace-humans-but-humans-with-ai-will-replace-humans-without-ai