










ModelOps

ModelOps Essentials

Best Practices for Success with Enterprise AI

ModelOps Essentials | Table of contents

- 1  Introduction
- 2  Lay the Foundation for Enterprise AI
- 3  Solve the “People” Challenge for Enterprise AI
- 4  Establish ModelOps Processes to Operationalize AI
- 5  Critical Requirements for a ModelOps Solution
- 6  Follow an Actionable Blueprint for ModelOps
- 7  Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI



1

Introduction

► **Introduction**

Lay the Foundation for Enterprise AI

Solve the "People" Challenge for Enterprise AI

Establish ModelOps Processes to Operationalize AI

Critical Requirements for a ModelOps Solution

Follow an Actionable Blueprint for ModelOps

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI



Introduction

A large majority of C-suite executives now see AI as essential to their organizations' survival, but significant obstacles are impeding their ability to scale AI. One obstacle is the lack of a structured, systematic way to deploy and manage the hundreds or thousands of models that an enterprise may rely on for the predictions and recommendations that guide decision-making and influence action.

ModelOps is the key strategic capability that enables enterprises to overcome this obstacle. This guide describes key points to consider when implementing ModelOps.

Who should read ModelOps Essentials?

- CIOs and AI Architects, who have the overall responsibility of model life cycle (MLC) in production and want to manage Shadow AI and eliminate Model Debt
- Executives who see the compelling need for AI and want their organization to implement it more broadly, and rapidly
- Business unit managers who want to start or accelerate their unit's or function's use of AI and/or gain greater benefit from models already in use
- Operations directors, ITOps, Data Ops, DevOps, and ModelOps teams looking for a more effective way to deploy and manage models
- Data scientists who want to devote more time to developing and improving models and less time to managing them
- Compliance teams who need to control and have full visibility on the risks introduced by AI initiatives

This guide contains practical insights designed to help organizations scale AI enterprise wide more rapidly and more confidently.



ModelOps: A Definition

ModelOps is the key strategic capability to **Deploy, Monitor** and **Govern** model life cycle in production, across the entire enterprise AI.

It goes beyond technology, as it is the organizing principle that also encompasses people, and processes involved in operationalizing AI models.

How each Professional will benefit from ModelOps

Business leaders rank enterprise AI as a high priority, but scaling remains challenging. One obstacle is structural. Enterprise AI touches many parts of the business, and it requires a unified strategy that gives each area of the business the flexibility to use the AI tools that best suit their needs. In most enterprises, the way teams are organized can create or reinforce conflicts among and between lines of business, data science, IT, DevOps, DataOps and risk/compliance teams. ModelOp minimizes departmental conflicts and operational frictions while ensuring that the outputs of all implemented AI tools flow into the business quickly, reliably, and with full compliance and accountability.



Data Scientists



ITOps



DevOps



DataOps



CIO



LOB Executive



LOB Analyst



Compliance



AI Architect



ModelOps Engineer

2

Lay the Foundations of Enterprise AI

Introduction

► **Lay the Foundation for Enterprise AI**

Solve the “People” Challenge for Enterprise AI

Establish ModelOps Processes to Operationalize AI

Critical Requirements for a ModelOps Solution

Follow an Actionable Blueprint for ModelOps

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI



Best Practices for Success with Enterprise AI

For more than a decade, business and technology leaders, analysts, consultants, and the media have watched the surge in artificial intelligence (AI) use. Most enterprises, however, have not yet seen the promised benefits from AI. Few have created a platform that supports the deployment of AI models for a variety of use cases in different functions, departments, and regions, and embeds model outcomes in products, processes, or decision-making. Even fewer have established effective governance for the models they have put in operation.

A 2019 [Accenture study](#) of 1,500 C-level executives found that although 84% believe they must employ AI to achieve their growth objectives, 76% said they struggle when trying to scale AI across the enterprise.

Common challenges highlighted in the research:

Leadership commitment to keep AI transformation a top priority throughout the entire journey

Gaps in culture and skills that create organizational barriers to enterprise AI transformation

Technology complexities across data science, data, and computational platforms that make it incredibly difficult for even the most advanced software engineering teams to build new AI-enabled applications

The first step for enterprises: Recognize Models as first-class enterprise assets

Most large organizations already make use of an array of machine learning (ML) models, as well as more traditional regression and rules-based models, all of which contribute to AI efforts.

The enterprises gaining the most from their AI investments elevate models to first class assets, investing the time and money to ensure that along with having a powerful model factory to create models, they also add the ModelOps capability, to efficiently drive the deployment, monitoring, and governance of their models and ensure they maximize return on their investments in models – no different than any physical or intellectual asset of the enterprise.

The People, Processes, and Technology sections of this guide include best practices for:

- Adding or adapting roles and aligning teams around ModelOps
- Designing processes to operationalize AI
- Defining critical requirements of a ModelOps technology solution

3

Solve the “People” Challenge for Enterprise AI

Introduction

Lay the Foundation for Enterprise AI

► **Solve the “People” Challenge for Enterprise AI**

Establish ModelOps Processes to Operationalize AI

Critical Requirements for a ModelOps Solution

Follow an Actionable Blueprint for ModelOps

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI



Solve the “People” Challenge for Enterprise AI

Business leaders rank enterprise AI as a high priority, but scaling results more often than not eludes them. One obstacle is organizational. Enterprise AI touches many parts of the business, and requires a unified strategy that gives each area of the business the flexibility to use the AI tools that best suit their needs while ensuring that the outputs of those tools flow into the business quickly, reliably, and with full compliance and accountability. But in most enterprises, the way teams are organized can create or reinforce conflicts between lines of business, data science, IT, DevOps, DataOps and risk/compliance teams. The addition of two key roles creates an organizational structure that preserves the independence of existing functional groups and paves the way for successful scaling of AI initiatives.

KEY TAKEAWAYS

- ModelOps ownership has to be in the remit of the CIO to proactively manage Shadow AI and eliminate Model Debt
- Managing model life cycle is not the same as managing traditional IT technology platform, applications and software
- Effective ModelOps requires 2 new roles: AI Architect and ModelOps Engineer.

Benefits of Dedicated ModelOps Leadership

By centralizing responsibility for model deployment, monitoring, and governance, the enterprise can expect to benefit from:

- Consistent, repeatable methods for deploying models into business applications
- Greater speed in model deployment
- More cost-effective model management
- Risk control and processes to ensure ethical fairness
- Clearer visibility into model status, usage, and history, facilitating auditing and reporting

Even in the early stages of the enterprise AI journey, organizations with just a handful of models are finding it difficult to move models into production. This situation is going to become exponentially more challenging. Demands for using AI to gain advantage in the business are growing. Data science teams are expanding. Tools are becoming more powerful. And as Augmented ML capabilities are added to conventional applications (business intelligence, HR, CRM, etc.) a wave of “citizen data scientists” will take advantage of these capabilities to develop models. Even if models are used entirely within the business – e.g. a model used by HR to recommend salary raises and promotions – they will still be subject to model life cycle requirements to ensure efficacy, timely refresh, and compliance with

corporate and regulatory standards. ModelOps as an enterprise capability is essential to ensure that the business can leverage the transformational value of AI while remaining protected against the exposure from unmanaged models and Shadow AI.

Where should this new centralized ModelOps capability sit organizationally? Most organizations early in their transformation logically place it with a corporate data science team, often under the Chief Analytics Officer. That may have made sense when the focus was on figuring out if data science could drive value for the enterprise. After all, we are talking about capability around data science models. Now that most enterprises have decided that data science models not



ModelOps Essentials | Solve the “People” Challenge for Enterprise AI

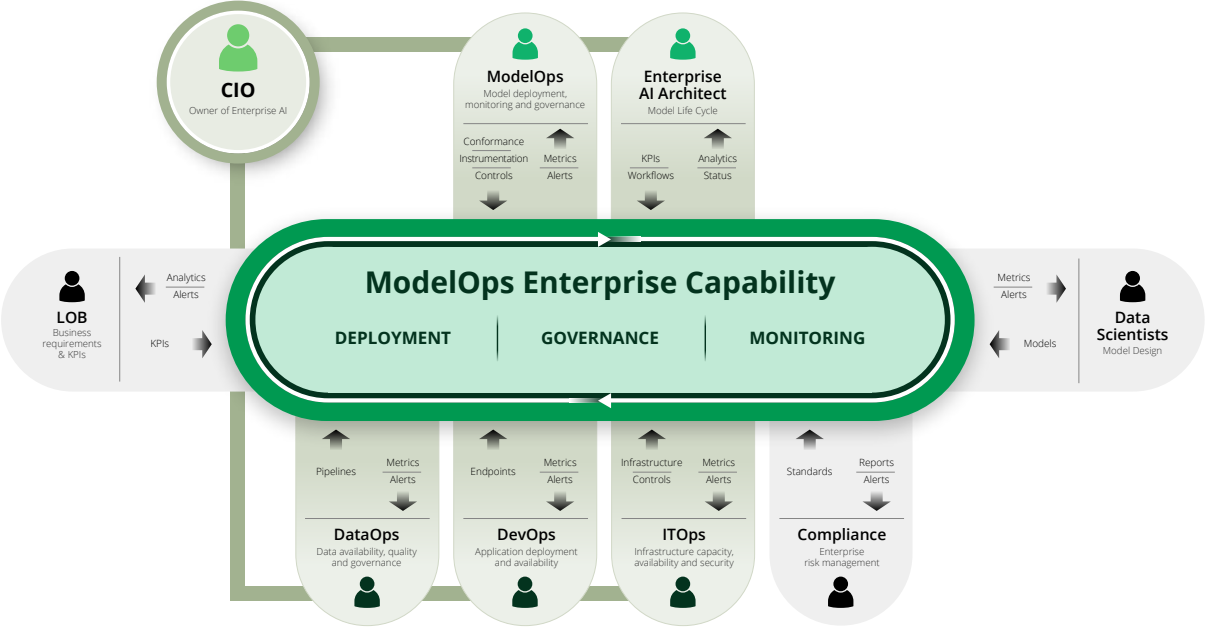
only add value but are critical to their future competitive standing, or even existence, it is time to revisit organizational accountability for ModelOps.

Models are technical artifacts, deployed into end-user applications, with complex operational requirements that become mission critical as more and more models are used to drive automated decisioning within the software that operates the business. For example running the manufacturing plants, routing supplies and finished product, ferreting out fraudulent transactions, or speeding-up back office processing. **Managing, monitoring and governing technology that does not include models is the remit of the CIO.** As more and more of the enterprise’s technology includes models, they are the only rationale choice to own ModelOps. And that frees the data science organization to do what they do best. Discover new models and improve existing ones to drive more and more value to the enterprise and its customers over time.



There are, however, important differences between managing and operating legacy technologies and model-driven ones. Managing models through their life cycle is not the same as managing traditional IT technology platform, applications and software- as explained in the next [Chapter](#)

The chart below illustrates how ModelOps could fit within an existing CIO organization, and introduces two new roles to effectively deliver this new capability – AI Architect and ModelOps Engineer – that we will discuss in detail shortly.



With this new organization in place under the CIO, with a process owner in place for the full model lifecycle, the enterprise is now in position to review and adapt its ModelOps processes a critical next step covered in next [Chapter](#).

Adapt or Add Roles and Align Teams

Recent studies by [McKinsey](#) and others show that enterprises that succeed at scaling AI show greater reliance on cross-functional teams. But who should lead such teams? Gartner proposes a new role: the enterprise AI Architect.

And who should own the responsibility of operationalizing all models across the organization? Consensus is growing that the CIO appears to be the best suited to successfully operationalize Model Life Cycle (MLC) in production.

However, while it may seem desirable to vest this responsibility in an existing functional group, doing so can have a number of potential drawbacks:

Many teams have a part to play in ensuring success with enterprise AI, as described [here](#). But a best practice for an enterprise that wants to accelerate its ability to scale AI is to assign responsibility specifically for designing and overseeing MLC management functions. While it may seem desirable to vest this responsibility in an existing functional group, doing so can have a number of potential drawbacks.

IT teams

Traditionally, IT roles take the lead in many AI efforts, studies show, but giving existing IT teams leadership over AI is not ideal, for the following reasons:

1

The KPIs for ITOps typically center on application performance, availability and security, as well as infrastructure cost. IT teams don't typically have the skillset to address things like model effectiveness or compliance.

2

ITOps typically do not drive cross-functional processes similar to MLC management.

3

While ITOps is often accountable to lines of business for application availability and cost, they are rarely accountable for business value of an application, which is a major component of accountability for models.

4

The demands placed on IT to deliver nonstop availability of new and existing applications at increasing scale while holding the line on costs taxes even the most well-resourced IT team and leaves little time for developing and managing a comprehensive ModelOps capability.

Data science teams

Data science teams often are expected to take the lead, particularly when an enterprise has established a center of excellence for analytics. However, data scientists are not ideal leaders for ModelOps because:

- Data scientists are as their title implies – scientists. They are trained to apply complex mathematics to extract value from data. Operationalizing the models they create is simply not what they do. Using data scientists to do operational tasks squanders their unique talents.
- Data scientists like experimenting. They seek new and novel outcomes. In contrast, operationalizing models at scale requires establishing and adhering to rigid practices that deliver predictable performance - without surprises.
- To do their jobs well, data scientists need maximum flexibility to experiment with different tools and techniques. If they're forced to spend time on operationalizing their models, what they implement will likely be tailored to work with the particular type of model produced by their tool of choice and won't be suitable for any other type of model or tool. This leads to fragmented and ultimately an unwieldy MLC management situation that can't scale.

DevOps teams

DevOps teams are often called upon to bridge the gap and write the scripting that moves models from a data science tool or platform into business applications. There are drawbacks to relying on DevOps, however:

- DevOps specialists are accustomed to working with the type of coding, tools, and processes used in software development, which are quite different from those used in AI model creation.
- DevOps teams know how to track versions and upgrades for software applications. But models are actually much more varied and unstable than software. As mentioned earlier, unlike conventional software, AI models decay over time and have to be refreshed at cadences ranging from quarterly to weekly. Every model has different requirements depending on the application in which it's deployed.
- Each application may use dozens or hundreds of models, meaning the scale of MLC management can be orders of magnitude larger than DevOps.
- Like ITOps and other teams, DevOps teams have many competing priorities fulfilling their role delivering and operating reliable, functional applications to the business.

Business teams

Business teams have the most to gain from scaling AI model deployment. Yet these teams are not well equipped to lead ModelOps either, because:

- Business teams identify the type of business application and value needed but they usually rely on others to procure or develop it.
- Business teams may control some of the data that analysts require, but they do not oversee the enterprise infrastructure and must rely on IT to do so.
- Business teams are not experts in model development or MLC management. They may not even know such management is required.

Why ModelOps requires new roles

Considering the above, the question is: **who should lead such teams?** Gartner proposes a new role: the Enterprise AI Architect, and best practices from the field suggest that the ModelOps Engineer is an additional required new role.



Enterprise AI Architect (EAIA)

The Enterprise AI architect oversees MLC management from model inception through refreshes to eventual replacement. By designing and implementing processes that clearly define how models move through the MLC - including the data generated, collected and reported and the responsibilities of each participant at each step - the EAIA **establishes the technical and organizational scaffolding that unites data scientists, data engineers, developers, IT operations, risk managers and business unit leaders around concrete, clearly defined processes.**

The EAIA also **defines and implements the systems that automate the MLC**, enabling maximum velocity and efficiency with visibility, accountability and control at scale.



An EAIA **must understand how the technical and business requirements of each model drive its life cycle.**

This includes knowledge of the trends in model creation tools and the needs of different data science teams working on different projects. They also need to understand how consuming applications drive different MLC requirements. For example, a model deployed in an application subject to governmental or industry regulations will have different requirements for compliance reporting than a similar (or even the same) model deployed in an unregulated application. Additionally, different models will require different refresh rates depending on how rapidly the “ground truth” of the incoming data departs from training data and the impact of drift on the efficacy of the model.




Through the EAIA's ability to enable systematic MLC management, the enterprise gains greater efficiency, gets new models into production faster and reduces the risk of incorrect, outdated or non-compliant models being put into or remaining in production. Responsibility for communicating to senior management might fall within the scope of the Enterprise AI Architect.

ModelOps Engineer

Another new role, ModelOps Engineer, likely reporting to the Enterprise AI Architect, **acts as the technical liaison between business, data science, ITops, DataOps DevSecOps and compliance teams.**

The ModelOps Engineer works with the extended cross-functional teams to translate the blueprints and templates provided by the EAIA into a specific MLC for every model under development. **This ensures that all models meet relevant standards for efficacy, data usage, interpretability, auditability, fairness and of course, ROI – and does so in a consistent way.**



The ModelOps Engineer may work initially in a single business unit, but over time they can become part of the

core team that provides ModelOps services across the enterprise. If implemented properly, the centralization of ModelOps as a discipline provides maximum flexibility for business units and data science teams to satisfy their unique business requirements without being locked into a single, restrictive approach, while maintaining visibility and accountability across the enterprise.



Together, the Enterprise AI Architect and the ModelOps Engineer form the core of a function that ensures that models are operating and conforming to business, technical and compliance KPIs 24x7.

4

Establish ModelOps processes to Operationalize AI

Introduction

Lay the Foundation for Enterprise AI

Solve the "People" Challenge for Enterprise AI

▶ **Establish ModelOps Processes to Operationalize AI**

Critical Requirements for a ModelOps Solution

Follow an Actionable Blueprint for ModelOps

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI



Design Processes and Automations to Operationalize AI

The processes needed to deploy, monitor, and govern models might seem straightforward to those unfamiliar with how models are used in enterprises today. At FANG and digital natives companies with limited business complexities, models guide hundreds of decisions and automate both human and machine-driven action. Within a single business unit, dozens of models may be operating simultaneously. "Model factories" use an array of data science tools and techniques to generate and update all these models. For most enterprises, however, this scale of AI operations is not easy and adopting AI technologies and hiring top AI talents is not enough, and they risk falling further behind innovators who leverage their ability to scale the use of models in business to drive further competitive advantage.

So, how do enterprises that are not operating at this scale close the gap, and ensure that such valuable intellectual property is leveraged properly, by ensuring that models are put into the business in the most efficient manner and then maintained in a way that meets standards set by the company, its stakeholders, and its regulators?

ModelOps is the key.


Although many models contain scripting, models are not software. This means they cannot be managed following typical DevOps practices and processes. Each model potentially has a unique life cycle process. More details at the end of this chapter.



KEY TAKEAWAYS

- MLC Operationalization should not be left to chance or to manual effort due to its Scope, Complexity, and high Business Risk: it has to be process-driven
- Every single model has a different life cycle process
- These processes must be automated for the enterprise to drive AI at scale

ModelOps Processes and Automations



Deploy, Monitor, and Govern



1. Deploy models consistently and efficiently to any consuming application. The first step is to register an abstraction of all models in a central “catalog.” All the elements that compose the model—such as source code, tests, input and output schemas, training data, metadata, as well as outputs of training—should be included, along with all the elements required to execute it, including libraries. The next step is to apply low-code techniques to deploy the model into the desired business application. By providing visibility into all models, the catalog serves as a springboard for more efficient model deployment according to each model’s unique deployment path. The catalog should also generate and persist a robust lineage for the model, showing how it was put into production, where it was used in the business, and any later versions.



2. Monitor models continuously. Unlike software, models decay over time. Model performance should be tracked in three dimensions: statistically, technically, and from a business perspective:

- Statistically, is the use of input data and the output inferences performing as designed, or is there evidence of data “drift”? And how do the model’s inferences compare with one or more “better” models?
- Technically, is the model delivering inferences with the originally specified load and lag time and burden on operational systems?
- Finally, is the model continuing to provide useful business insights, or is the model at risk of exhibiting bias or violating business rules? And are new models being put into production quickly and efficiently?

If any of these metrics fall outside pre-set parameters, the best practice is to automate the updating of the model—including any necessary approvals—so that an optimized version can be quickly returned to production.



3. Govern models like any corporate asset. Models are a form of intellectual capital. Like other forms of capital, they should be inventoried and assessed using tools and techniques that make auditing and reporting as efficient as possible. With a robust lineage maintained for each model within a central catalog, management can see which model or models were used in specific business applications at any decision point, explain what inferences were delivered and how, as well as the impact those inferences had on the business. With automatic updates to model metadata, management can reproduce the model at any point in its life cycle. This is particularly important in industries or functional areas that require clear model interpretability and rigorous adherence to regulations.

Fortunately, establishing ModelOps best practices can be done incrementally, although there are [pitfalls](#) to avoid, as spelled out in the [Appendix](#) to this Guide. An organization might want to start with business units that use models the most, for example, or those experiencing the strongest threat from competitors. Model usage in other business units can continue uninterrupted and then be woven into the ModelOps process in a phased, orderly fashion.

DevOps and ModelOps

Many enterprises have tried to simply use or extend DevOps tooling for their AI/ML workflows but are challenged as DevOps tools was designed to handle technical implementations, whereas with models:

- 1. Many of the decisions related to model operationalization are not about technical matters, but business and ethical considerations: how will the model be used, is this proper use, is there partiality for protected entities, etc.
- 2. Tight relationship between data and the model itself that needs to be continuously monitored once in production.
- 3. Given the strategic nature of models and especially the implications with consumer decisioning, governance of these models is imperative--and required in regulated industries. DevOps tooling was aimed to provide automation and efficiency in technical implementations, not to be an enterprise's model governance repository.

Traditional Software	AI/ML Models
Deterministic	Probabilistic
Code is separable from data	Encoded with data
Does not degrade over time	Drift over time, and need refresh

5

Critical Requirements for a ModelOps Solution

Introduction

Lay the Foundation for Enterprise AI

Solve the "People" Challenge for Enterprise AI

Establish ModelOps Processes to Operationalize AI

► **Critical Requirements for a ModelOps Solution**

Follow an Actionable Blueprint for ModelOps

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI

ModelOp

Critical Requirements for a ModelOps Solution

The technology environment for enterprise AI is both complex and rapidly evolving. There are massive investments being made across the enterprise AI stack by both new entrants and legacy technology providers. Forward-thinking leaders should choose a ModelOps Solution that allows them to abstract that complexity and continue to take advantage of technological innovations in the market for both model factory tools and runtime environments.

At the same time, regulations on enterprises use of AI will become more stringent and a ModelOps solution must offer strong governance capabilities for enterprises to meet these compliance requirements.

As a result, when implementing ModelOps, leaders should consider the following technological requirements, for model deployment, execution and monitoring, continuous improvement, and governance.

KEY TAKEAWAYS

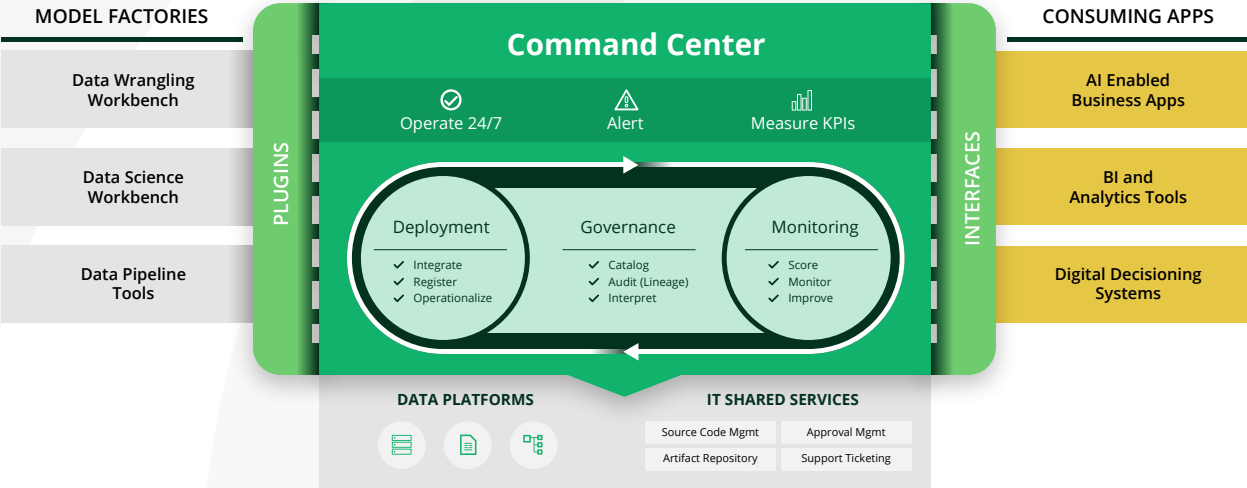
The three Critical Requirements for a ModelOps Tech Solution, relative to models in production, are:

- Ability to Deploy, Monitor and Govern any model across all enterprise AI, regardless of their development tool or runtime environment
- Ability to abstract the complexity of the enterprise AI stack, driving agility and scale in the enterprise's operationalization of models in business
- Ability to automate the model life cycle in the enterprise with repeatability, resilience and scale

Deployment	<ul style="list-style-type: none"> • Support a variety of data science tools and workbenches • Integrate with standard IT platforms and systems • Interconnect with a range of consuming applications • Include a central, searchable catalog for models and related metadata • Automate model approval process, including the application of business rules • Automate any security scans required by enterprise policy infrastructure
Execution and Monitoring	<ul style="list-style-type: none"> • Support a variety of model frameworks and languages • Support any scoring mode (batch, request/response, streaming) • Execute models built in a variety of different environments, inside and outside of the enterprise • Manage MLC KPIs to identify and avoid performance issues • Monitor data quality and data "drift" based on pre-set thresholds • Evaluate statistical performance of models • Monitor each model's impact on infrastructure • Measure model latency to meet business SLAs
Continuous Improvement	<ul style="list-style-type: none"> • Trigger model refresh based on monitoring alerts • Support for multiple modes of retraining (incremental, batch) • Support A/B testing of model performance and champion/challenger model comparison
Governance	<ul style="list-style-type: none"> • Enable the collection and management of all model training output artifacts along with the model metadata • Provide robust management of the frameworks, libraries, and other dependencies that are required to execute the model • Support model version management • Support data pipeline management • Collect and persist each step in the model's life cycle for full reproducibility and auditability • Include tools and mechanisms to test for and eliminate model bias

ModelOps Essentials | Critical Requirements for a ModelOps Solution

The below image shows how a ModelOps Solution abstracts the complexity of the AI technology stack and how it allows freedom to integrate with existing environments that will change over time.



Gain Benefits at Every Stage

With a rich set of ModelOps features, enterprises benefit at every stage of MLC management.

Deploy

During development and deployment:

- Data scientists can be innovative in developing models in response to business needs.
- DevOps teams/software engineers can be less involved in packaging models.
- IT does not need to create a unique environment for each model yet retains control of data pipeline configuration and infrastructure optimization.
- Model review, testing, and approvals are automated, with workflows visible to all participants.
- Business unit managers see models deployed faster.

Monitor

Execution is efficient and consistent and is followed by **continuous monitoring**:

- Model accuracy, performance, data quality, and the demands placed on enterprise infrastructure are all regularly assessed so that modifications can be made promptly.
- Continuous model improvement is supported through retraining and redeployment.

Govern

With robust **governance**, the enterprise can be confident that correct versions of models are deployed but that earlier versions are reproducible, if necessary, for audit or compliance purposes.



Implementing ModelOps makes it possible to manage the entire MLC, from inception through any retraining to retirement.

6

Follow an Actionable Blueprint for ModelOps

Introduction

Lay the Foundation for Enterprise AI

Solve the "People" Challenge for Enterprise AI

Establish ModelOps Processes to Operationalize AI

Critical Requirements for a ModelOps Solution

► **Follow an Actionable Blueprint for ModelOps**

Glossary

ModelOps Essentials

Best Practices for Success with Enterprise AI

ModelOp

Follow an Actionable Blueprint for ModelOps

As mentioned previously, ModelOps can be implemented incrementally, in a phased, orderly fashion. This means an organization's journey toward enterprise AI does not need to be delayed; in fact, it should not be.

- 1** The first step is a **ModelOps Capability Maturity Assessment**. At this phase, the business unit, department, or functional area's model management efforts are compared to industry benchmarks. The assessment then:
 - Identifies key organizational gaps—from technical, process, and governance perspectives
 - Prioritizes recommendations based on critical business initiatives
 - Establishes key metrics and processes for the program
 - Evaluates the ROI on a ModelOps implementation
- 2** In the second phase, organizational roles and responsibilities related to ModelOps are defined and mapped. The organization then receives a comprehensive roadmap of actions to be taken during the next 12-18 months to reach the agreed-upon target state.
- 3** During the third phase, a full communications plan is drafted to establish common expectations across the organization and socialize the value, requirements, and expectations for the ModelOps capability. After this phase, the organization may begin one or more pilot projects to apply ModelOps capabilities to its operations, very likely in a carefully controlled environment.

More details about the requirements for a robust ModelOps feature set are provided [here](#).

ModelOps Maturity Model

		Leading
Strategy & Enablement	Ad hoc ModelOps strategy not defined or documented <ul style="list-style-type: none"> • Containerization & micro services strategy in infancy • No ModelOps strategy, org, nor processes 	Maturing ModelOps strategy defined; Isolated buy-in <ul style="list-style-type: none"> • ModelOps defined as a capability in overall Analytics strategy • Initial micro services and containerization strategy defined and being piloted • Just-in-time ModelOps org/roles/RACI being implemented as an afterthought, causing issues and delays • ModelOps support and training team not in place
	Disparate & ad hoc approaches <ul style="list-style-type: none"> • Disparate approaches to developing models & data pipelines • Manual recoding or rework when migrating to different environments, including 3rd party environments • Single-modal scoring (e.g. Batch) • Limited or no model asset management • Manual testing, validation, approvals, and deployment of models 	Initial Enterprise standards; limited adoption <ul style="list-style-type: none"> • Attempts to unify by enforcing standard Workbenches/Platforms • Model execution inconsistent across environments • Some multi-modal scoring examples (Batch, REST) • Pockets of code and asset management, but no unification • Promotion processes and checklist defined • Ad hoc scripts created to semi-automate testing & deployment • Models integrated with CI/CD tooling • Initial containerization/micro services implemented
	Lack of effective, consistent monitoring <ul style="list-style-type: none"> • Narrow focus for model "monitoring" • Inability to measure model performance against target KPIs • No ability to track the end-to-end model deployment process 	Limited visibility through ad-hoc approaches <ul style="list-style-type: none"> • Lacking traceability to Business KPI's • Lacking holistic view: Data, Model, Infra, MLC, Business • Models monitored using ad hoc scripting • Ad hoc back testing, champion/challenger • Limited A/B/C testing, typically through ad hoc scripting • Limited integration for re-training
	Limited ModelOps governance policies <ul style="list-style-type: none"> • No ModelOps SLA's • Isolated, ad-hoc model development & deployment processes • Immature infosec and change management standards for Modeling • No model traceability, container, micro services, or model exchange policies • Limited model interpretability for new models 	Initial ModelOps governance policies <ul style="list-style-type: none"> • ModelOps SLA's are being piloted • Model traceability semi-automatically captured • Model usage/access policies are documented, but not implemented • Model environment multi-tenancy policies are designed, but not implemented • Initial container governance policies drafted • Limited model interpretability for new models
		ModelOps adopted across the org & fully supported <ul style="list-style-type: none"> • ModelOps strategy defined and being executed • Containerization strategy defined & being executed • Micro services strategy defined & being executed • ModelOps Org structure in place • ModelOps RACI in place • ModelOps support team and structure in place • ModelOps training & enablement program in place • ModelOps metrics & ROI tracking in place
Deployment		Model deployment fully optimized/automated <ul style="list-style-type: none"> • Standardized model abstraction across DS frameworks • Consistency of execution across environments • Consistency with vendor environments • Multi-modal scoring (streaming, batch, & on-demand) • Robust model asset management • Automated testing, approvals, & deployment of models • Configurable MLC reports Enterprise-wide • Full CI/CD integration • Mature container & micro services architecture
Monitoring		Comprehensive Enterprise-wide visibility <ul style="list-style-type: none"> • Traceability to business KPI's with feedback loop • Holistic monitoring: Data, Model, Infra, MLC, Business • Programmatic visibility into ModelOps KPI's & SLA's • Configurable MLC reports Enterprise-wide • Automated thresholds for alerting • Robust champion/challenger and A/B testing • Configurable rules for automated re-training
Governance		Comprehensive ModelOps governance <ul style="list-style-type: none"> • Robust Model inventory & auditability • Full model traceability capture • Implemented model usage & access policies • Implemented model multi-tenancy policies • Implemented container governance & security processes • Optimized/automated model validation process • Implemented bias and explainability processes • Implemented partner/supplier model exchange policies

Conclusion

A 2019 PwC global, cross-industry study estimates that AI could contribute up to \$15.7 trillion to the global economy by 2030. AI use cases vary from sector to sector, but in a 2019 study of the top benefits of AI, analyst firm Gartner identified 100 AI use cases in 40 industries.

It should be clear by now that AI is not the wave of the future: it's the wave of the present.

Organizations that succeed in implementing AI out-perform their peers. Yet numerous studies confirm that scaling AI efforts beyond pilot projects or siloed efforts has proven challenging.

Model-driven enterprises require a model-centric architecture to leverage and future-proof their existing investments in data, data science, and AI. ModelOps is the key to implementing this essential architecture.

Contact us to make Enterprise AI real with ModelOps

Learn more about our company and what we offer



© Copyright 2020 ModelOp, Inc.



About ModelOp

ModelOp enables large enterprises to address the critical scale and governance challenges necessary to fully unlock the transformational value of enterprise AI and Machine Learning investments. The ModelOp Center platform is the essential business accountable software solution that automates the complete life cycle for models, regardless of where they are created or deployed. Fortune1000 companies in financial services, manufacturing, healthcare and other industries rely on ModelOp to put their models into business.

7

Glossary

Introduction

Lay the Foundation for Enterprise AI

Solve the “People” Challenge for Enterprise AI

Establish ModelOps Processes to Operationalize AI

Critical Requirements for a ModelOps Solution

Follow an Actionable Blueprint for ModelOps

► **Glossary**

ModelOps Essentials

Best Practices for Success with Enterprise AI



ModelOps Essentials | Glossary

Term	Definition
Abstraction	The art of replacing specific details about a model with generic ones
Artificial intelligence (AI)	A computer engineering discipline using mathematical or logic-based techniques to uncover, capture, or code knowledge and sophisticated techniques to arrive at inferences or predictions to solve business problems.
Asset	Any individual component that is used and required during model deployment, such as model source code, schemas, dependencies, serialized objects, etc.
Data drift	The evolution of data over time, potentially introducing previously unseen variety and/or new categories of data.
Deployment (aka Productionization)	The process of making a model available for use by the business
Enterprise AI	Enterprise AI encompasses the end-to-end business processes by which organizations incorporate AI into 24x7 business functions that are accountable, manageable and governable at enterprise scale.
Governance	The management and mitigation of model risk to provide full transparency and auditability of all models across the enterprise
Inferences	Descriptions of the relationship between the independent variables and the outcomes in a dataset
Interpretability	The ability of a human to retrace how a model generates its inferences or predictions
Lineage	All human and system interactions (code changes, testing, promotions, approvals, etc) that have occurred throughout a a model's entire life cycle
Machine learning (ML)	A subset of AI that uses algorithms to parse data, capture knowledge, and develop predictions or determinations. ML models are first trained on datasets; then, once in production, use a closed-loop process to "learn" from experience and improve the accuracy of their predictions or determinations. Some ML models are both complex and opaque, making it difficult to explain how the models arrive at specific predictions or determinations.
Model	A set of code that represents functions, actions, and predictions important to the business
Model Debt	The implied cost of undeployed models and/or models deployed without proper monitoring and governance
Model decay	A change in model performance that makes it less accurate in its inferences or predictions
Model life cycle (MLC)	A model's journey from creation through testing, deployment, monitoring, iteration, and retirement
ModelOps	The key strategic capability for operationalizing enterprise AI. ModelOps encompasses the systems and processes that streamline the deployment, monitoring, governance, and continuous improvement of data science models, but its fundamental role is to improve business results.
Monitoring	The act of observing statistical, technical, and ethical aspects of a model's performance in operation
Predictions	Descriptions of the relationship between the independent variables and the outcomes in a dataset which are used to estimate outcomes for new data points
Schema	The definition of a model's expected data inputs or outputs expressed in a standard way.
Shadow AI	The implied cost and risk of deployment of AI initiatives and models in production with no accountability to IT or governance organizations. It is expected to be the biggest risk to effective and ethical decision
Training	Tuning model parameters to optimize performance on a particular dataset, with the typical output being a trained model artifact

Appendix

Pitfalls to Avoid when Implementing ModelOps

The journey toward implementing ModelOps capabilities can be interrupted or even derailed by the same types of pitfalls encountered in many other change efforts.

One that always makes the list: lack of executive buy-in. Fortunately, there seems to be a high level of executive awareness that success at AI may be critical to business survival. Documenting the importance of ModelOps to scaling AI—perhaps through a Maturity Assessment helps avoid this pitfall.

Unfortunately, other pitfalls exist, summarized below.

Pitfalls	Recommendation
Lack of an empowered and accountable model operations owner for the enterprise	The organization should choose an owner for model operations early and hold them accountable for key performance indicators (KPIs) such as mean time from “model ready” to “model in business,” model freshness and availability, continuing cost of model operations, the model governance process, and, ultimately, business metric improvement and overall ROI from models.
Postponing an enterprise model operations capability investment until after data strategy has been decided and implemented	Management should assume AI/ML processes will eventually become one of the most critical consumers and producers of data and immediately allow AI/ML investments to drive models into business while providing continuous feedback to the enterprise’s data/DataOps capability.
Assuming that an organization’s DevOps capabilities will suffice for ModelOps	Treat Data science and AI/ML models as first class and unique enterprise assets. Adopt a clear definition of such a model that will be relevant throughout its entire life cycle from ideation through useful operation to retirement.
Confusing “data-driven” business intelligence with “model-driven” AI/ML efforts	Recommendation: Clearly separate investments, operational cadence, expected results, accountability, and ownership of BI from data science and AI/ML. They are both valuable but fundamentally distinct.
Hoping that data science/ML workbenches will provide enterprise-ready model governance	The enterprise should not become tied to the model governance capabilities provided by any particular DS/ML/AI development tool but instead seek ModelOps solutions that are specifically designed to deliver all advertised capabilities for all models independent of their source.

For additional details, see this “ModelOps Pitfalls” white paper.
<https://www.modelop.com/resources/ai-transformation-with-modelops/pitfalls>

Data Scientists



Design, build, and test models that enable substantial business value for a particular use case

Design, build, and test experiments intended to improve models and to support models in business

Enterprise AI role:

- Deliver analytic models that create value when deployed in business applications.
- Develop tests for determining efficacy of deployed models and driving improvements.

ModelOps requirements:

- Freedom to use the most effective model development tool for each application and use case.
- Minimal overhead and constraints on model development imposed by model deployment considerations.
- Automated packaging and delivery of models from development workbench to production environment.
- Real-time visibility to performance of deployed models.

ITOps



Maintains the core infrastructure and associated services needed to support models running in business applications

Enterprise AI role:

- Provide highly available, secure infrastructure to operate performant models in enterprise applications at scale, 24x7.
- Deliver performance metrics for deployed models and applications.
- Plan and execute infrastructure evolution to support AI technologies

ModelOps requirements:

- Centralized catalog of all models, model runtime requirements, lineage and operational history.
- Standardized models that can be deployed, monitored and controlled in a consistent manner, in any infrastructure (on prem, cloud, hybrid) regardless of the tools used to create them.
- Automated alerts regarding model performance and behavior.
- Automated processes and approvals for deployment, testing, refresh, and monitoring.
- Real-time visibility to performance of deployed models.



DevOps



Develops business applications and ensures that they remain operational and available

Enterprise AI role:

- Create and implement agile processes and resources for incorporating models into business applications and deploying applications into the infrastructure.
- Report on business metrics for applications.

ModelOps benefits:

- Abstractions that enable models of different types and from different tools to be incorporated into applications and connected with any type of data pipeline (triggered, batch, real-time) with standardized monitoring and control functions.
- Consistent visibility for every model regardless of source.
- Automated model management and refresh without the need for ad-hoc, custom scripts.

DataOps



Develops and tests data pipelines for input to a model and receives the output of model inferences

Enterprise AI role:

- Provide relevant, curated data from multiple sources to all models.

ModelOps requirements:

- Complete, centralized model inventory with all relevant metadata regarding data sources accessed.
- Ability to quickly detect and correct changes in model behavior.
- Visibility and controls to ensure and verify that data protection guidelines for security and privacy are enforced.



CIO



Provides appropriate technologies to support the business strategy

Monitors, plans, and budgets for infrastructure capacity

Ensures the efficiency of teams responsible for data security, network utilization, application development

Enterprise AI role:

- Enterprise AI strategy, operational infrastructure, processes and oversight.

ModelOps requirements:

- Centralized, consistent tooling that allows automation of the processes that implement the MLC with the flexibility to accommodate each model and application use case.
- Real-time and trended visibility to the performance of all constituencies (Data Science, DevOps, DataOps, ITOps, Compliance) against their respective KPIs.
- Elimination of ad-hoc, fragmented approaches to ModelOps that compromise business value and expose the organization to risk.
- Visibility and controls to prevent the spread of “shadow AI” as AI capabilities become available to “citizen data scientists”.
- Data to drive realistic planning for infrastructure and staff needs and to protect existing infrastructure.

LOB Analyst



Defines the requirements for leveraging AI/ML models that can provide uplift for their business processes

Monitors KPIs and reports to LOB executive

Enterprise AI role:


- Work with the data science team to translate business requirements in model KPIs.
- Continuously monitor and report on model performance in terms of business KPIs.
- Ensure that variances from KPIs are discovered and resolved quickly.

ModelOps requirements:

- Clear articulation and automation of processes that implement the MLC.
- Real-time and trended visibility to model performance, especially at the business-KPI level.
- Ability to drive and track automated response processes when alerts show performance diverging from KPIs. Elimination of ad-hoc, fragmented approaches to ModelOps that compromise business value and expose the organization to risk.
- Visibility and controls to prevent the spread of “shadow AI” as AI capabilities become available to “citizen data scientists”.
- Data to drive realistic planning for infrastructure and staff needs and to protect existing infrastructure.



LOB Executive



Defines and executes strategies to leverage advanced analytics to increase the profitability of their business line

Enterprise AI role:

- Take maximum advantage of the business value available from use of AI technologies.
- Ensure that business KPIs drive the development, deployment and operation of models.
- Mitigate technical and organizational bottlenecks that limit the ability to derive value from AI investments.

ModelOps benefits:

- Conformance with corporate AI standards with minimum restrictions on data scientists and developers to create value.
- MLC automation to minimize use of staff time and delays and ensure compliance with internal and external approvals and regulations.
- Automated model updates (retraining) to limit decay and maintain maximum business value output.
- Continuous visibility to model performance in terms of business-level KPIs.
- Automated processes and approvals for deployment, testing, refresh, and monitoring.
- Real-time visibility to performance of deployed models.

Compliance Manager



Ensures that models, their data and consuming applications meet internal and regulatory standards for efficacy, privacy, fairness, and other parameters

Enterprise AI role:

- Minimize and mitigate risks and exposures created by the use of AI.
- Monitor and audit models and data for conformance with regulations and corporate standards.
- Communicate regulatory compliance to LOB, Data Science, DataOps and other teams to ensure models and data pipelines are developed to be conformant.

ModelOps requirements:

- Centralized model catalog with comprehensive metadata.
- Automated reporting on the complete history and lineage of any model to satisfy audit requirements.
- Alerts that trigger when models depart from compliance specifications or process approvals.



AI Architect



Designs the cross-functional processes that implement a responsive and effective MLC for all models, business units and functional organizations.

Architects the tooling used to automate the MLC and integrate it with the enterprise IT stack

Enterprise AI role:

- Design the technical standards, process templates and KPIs for the end-to-end MLC for all models used in all applications.
- Design and oversee implementation and operation of the enterprise ModelOps platform
- Report to all constituencies and executive management on the status and progress of the organization's enterprise AI journey.

ModelOps requirements:

- Tools that enable design and implementation of processes that define the MLC for each type of model in each application.
- Standardization of model deployment, monitoring and control capabilities.
- Freedom to empower Data Science teams to utilize the model creation tools best suited their unique requirements.
- Process automation that facilitates coordination of all constituencies with minimum friction.
- Ability to mitigate the impacts of Shadow AI as AI capabilities become embedded in common enterprise applications.
- Reporting that demonstrates the value of AI to the business

ModelOps Engineer



Serves as the primary, hands-on interface between all groups that implement and benefit from Enterprise AI

Enterprise AI role:

- Assists in the integration of ModelOps capability into enterprise IT systems and business applications
- Monitors model performance and ensures rapid response to any issues that arise

ModelOps requirements:

- Ability to reflect technical and business KPIs in an automated ModelOps framework
- Ability to enforce technical standards that conform models from any source to common standards for deployment, monitoring and governance, without imposing undue restrictions or requirements on Data Science, DataOps, DevOps, ITOps or compliance teams.
- Ability to implement processes that fit the unique requirements for consistent adherence to the SLA's of the business through mutually agreed-upon rules, triggers, and/or alerts
- Ability to quickly spot, resolve and report on issues with any part of the MLC process (technical, business, approval, reporting, etc.).

