Good Decisions: A Monthly Webinar for Enterprise AI Governance Insights

The Next Wave:

SLMs, Agentic AI, and the Future of Model Governance

Jim Olsen, CTO at ModelOp

March 26, 2025

ModelOp

SLMs, Agentic AI, and the Future of Model Governance

- LLMs for Agentic AI
  - Reinforcement Learning
  - Advantages and Disadvantages of LLMs for Agentic AI
- SLMs and Agentic AI
  - Advantages and Disadvantages of SLMs for Agentic AI
  - Model Distillation and SLMs
- An Agentic AI Architecture
  - A Basic LangGraph Example
  - The Challenges of Tracking Agentic AI Solutions
- Representing an Agentic AI Architecture in a Model Inventory
  - Agentic AI Use Case
  - Model Ensembles for Representing Agentic AI Implementations
- Tracking Agentic AI Solutions
  - Track Agent and Solution Performance
  - Track Risks, Reviews, and Compliance

ModelOp

What is Agentic AI?

- Agentic AI is an artificial intelligence system that acts autonomously using agents to perform tasks, make decisions, and achieve specific goals.
  - Goal-oriented behavior with minimal human intervention
  - Ability to analyze, plan, and take actions based on situational and acquired data
  - Adaptability to dynamically changing conditions or inputs
- Agentic AI solutions are designed to simulate human-like reasoning and decision-making capabilities in solving complex problems or executing tasks.
- Agentic AI solutions will use expert models inside of agents to accomplish their tasks.

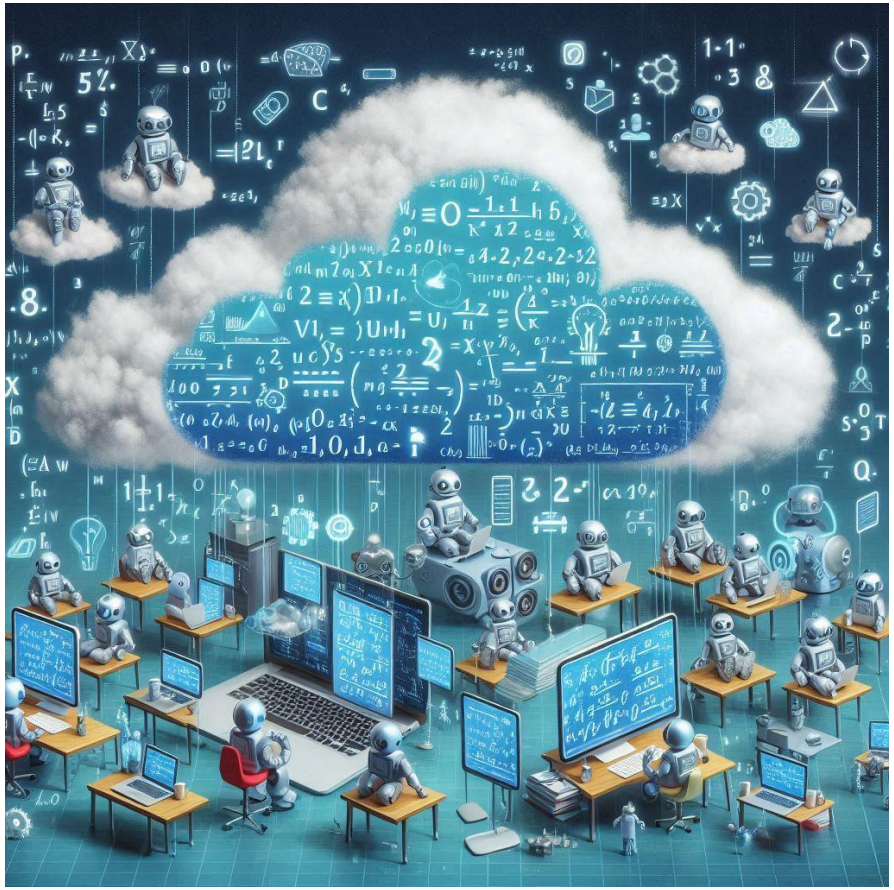

ModelOp

Using LLMs in Agentic AI

- Agentic AI requires expert models
  - Models must be experts in the domain they are dealing with
  - Agentic AI LLMs need to contain all of the knowledge for all tasks they must perform
  - LLMs are large in nature requiring a fair bit of resources
  - Only a few models would then be shared across agents
  - These models will be quite large, expensive to run, or can not be self hosted
  - Foundational models out of the box may or may not be up for the task

ModelOp

Reinforcement Learning in Large Language Models

- Expert models need to receive specialized training, reinforcement learning in LLMs provide this.
  - Augments the knowledge of a foundational model with specialized knowledge
  - Uses a reward model for a specific knowledge domain
  - Loss model quantifies the accuracy of the output and helps prevent divergence
  - Large language model will retain all of the foundational knowledge
- Disadvantages of LLMs for Agentic AI
  - May require sending all your data out to a vendor model
  - Self hosting and training are very GPU and memory intensive bringing high cost
  - Time to add knowledge and cost to do so are prohibitive to additional refinements

ModelOp

SLMs and Agentic AI

- Small Language Models (SLMs)
  - SLMs are designed with fewer parameters, resulting in lighter computational requirements and faster processing
  - They are typically tailored for specific tasks or domains and are less generalized compared to LLMs
  - Since they are smaller, they require less memory, storage, and processing power
  - Are limited in their capacity to learn thus are not a good general solution for different problems
- SLMs are a good fit for agentic AI
  - Can create experts for a specific task or domain
  - Can run locally on off the shelf hardware

ModelOp

Model Distillation

- Utilize a foundational LLM model to ‘teach’ a smaller model specialized knowledge
  - An automated form of supervised learning
  - Teacher sets soft targets for each input example
  - Student is trained to predict these values leveraging a loss model to predict deviation
  - Knowledge of specific areas is thus condensed into the smaller model
- Limitations
  - The larger model must already contain the knowledge of the target space
  - The smaller model may exhibit oversimplification of tasks due to the compression of knowledge.

ModelOp

Agentic AI Architecture

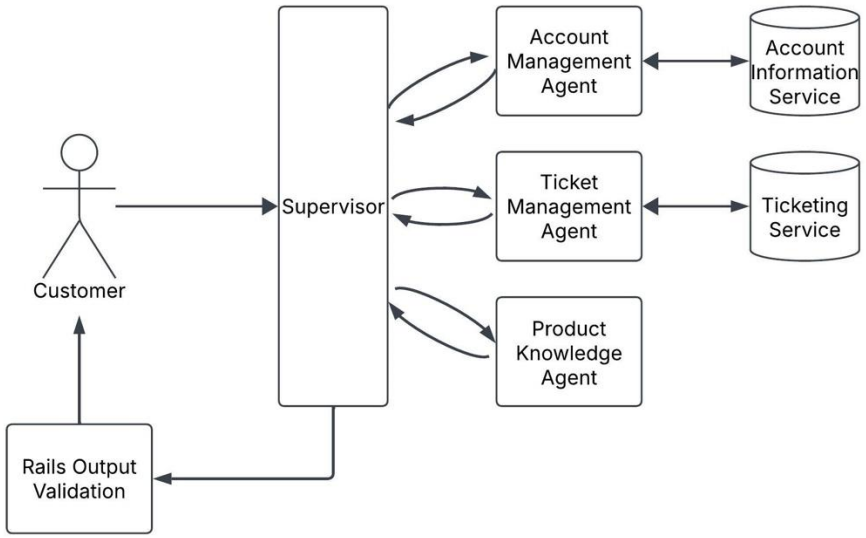

- Automated Customer Support
  - Multiple agents work together to gather information about the customer and its status
  - A supervisor implemented utilizing LangGraph manages information and requests to the agents
  - The agents are specialized in the knowledge and/or tool interactions
  - Supervisor utilizes agents freely to create response to customer
  - Final output is validated using Guardrails.ai to ensure response is not offensive

ModelOp

Tracking Agentic AI

- Many Expert Models
  - Expert SLMs provide value to the business
  - Must understand where expert models are being used
  - Track performance of the expert generically
  - Track performance of the expert as part of the solution
- Governance of your agents
  - Approvals of models for varying use cases
  - History of distillation techniques and any utilized datasets
  - Adherence to any state or regional regulations specific to your use case

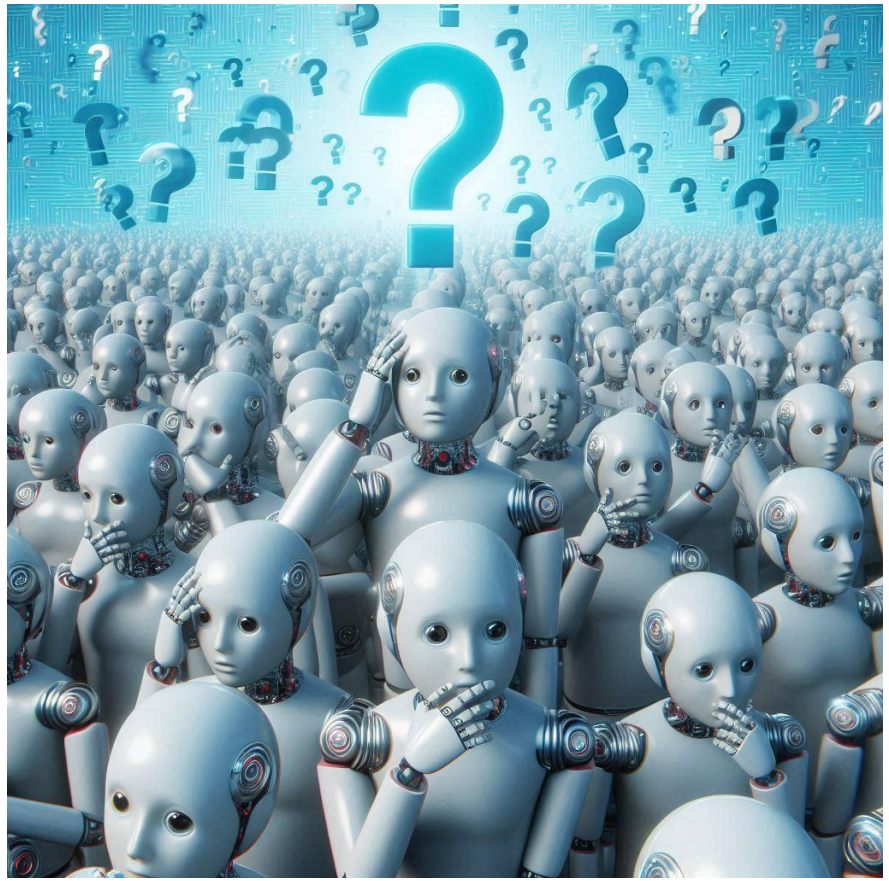

ModelOp

The Agentic AI Use Case

- Describes the solution
  - Includes risks, forms, documentation
  - Is implementation independent
- References the implementation
  - Ties to one or more implementations
  - Implementations contains the code and agents

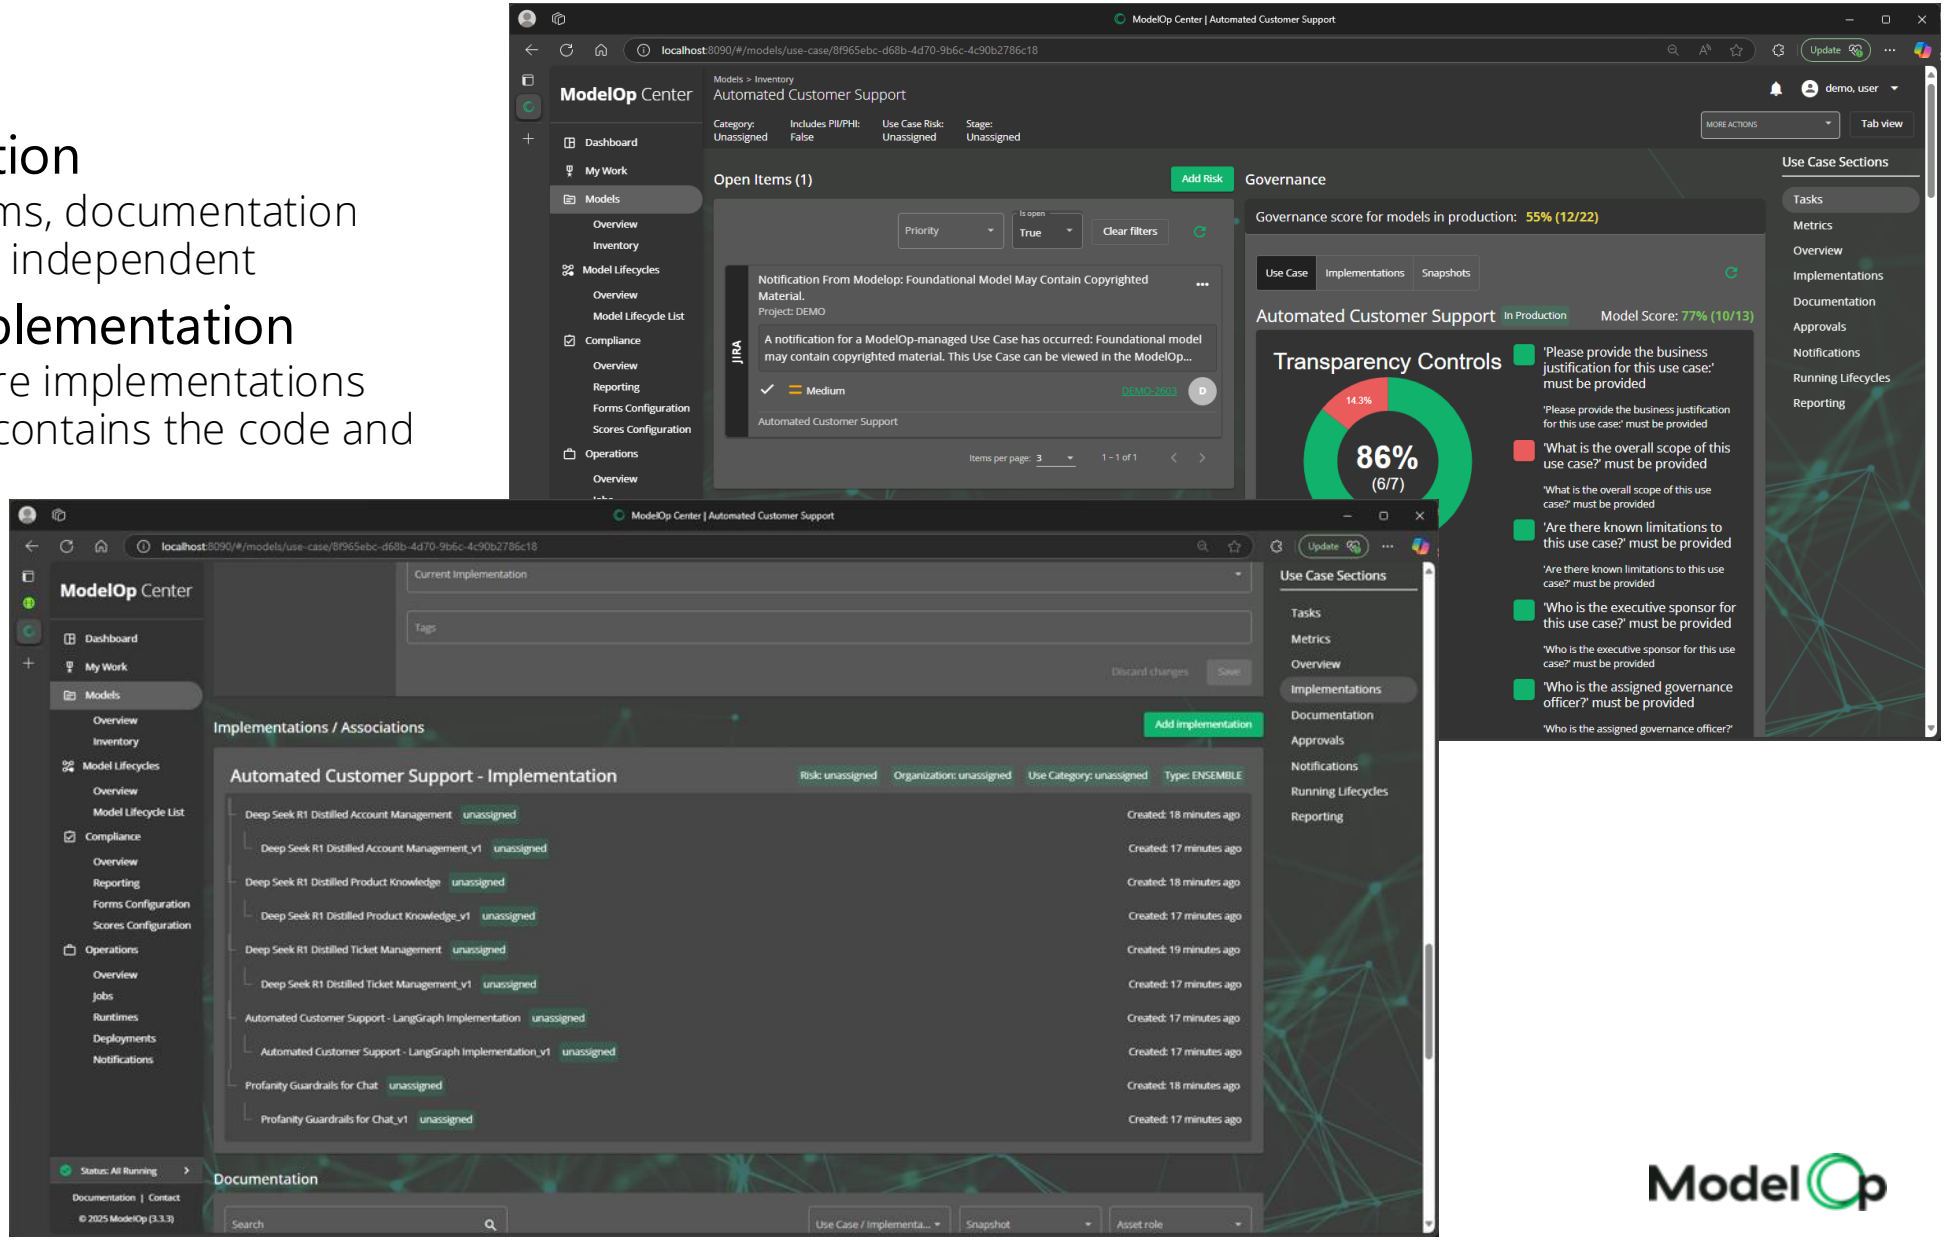

Model Ensembles

- Model ensembles represent the implementation in the inventory
  - Pull together all parts of the implementation
  - Ensembles can share referenced models
- An Agentic AI Solution is made up of many parts
  - A Solution should pull together multiple implementation models to represent the overall solution
  - Agents are one part of the solution
  - Rails files
  - LangGraph Python code and resource

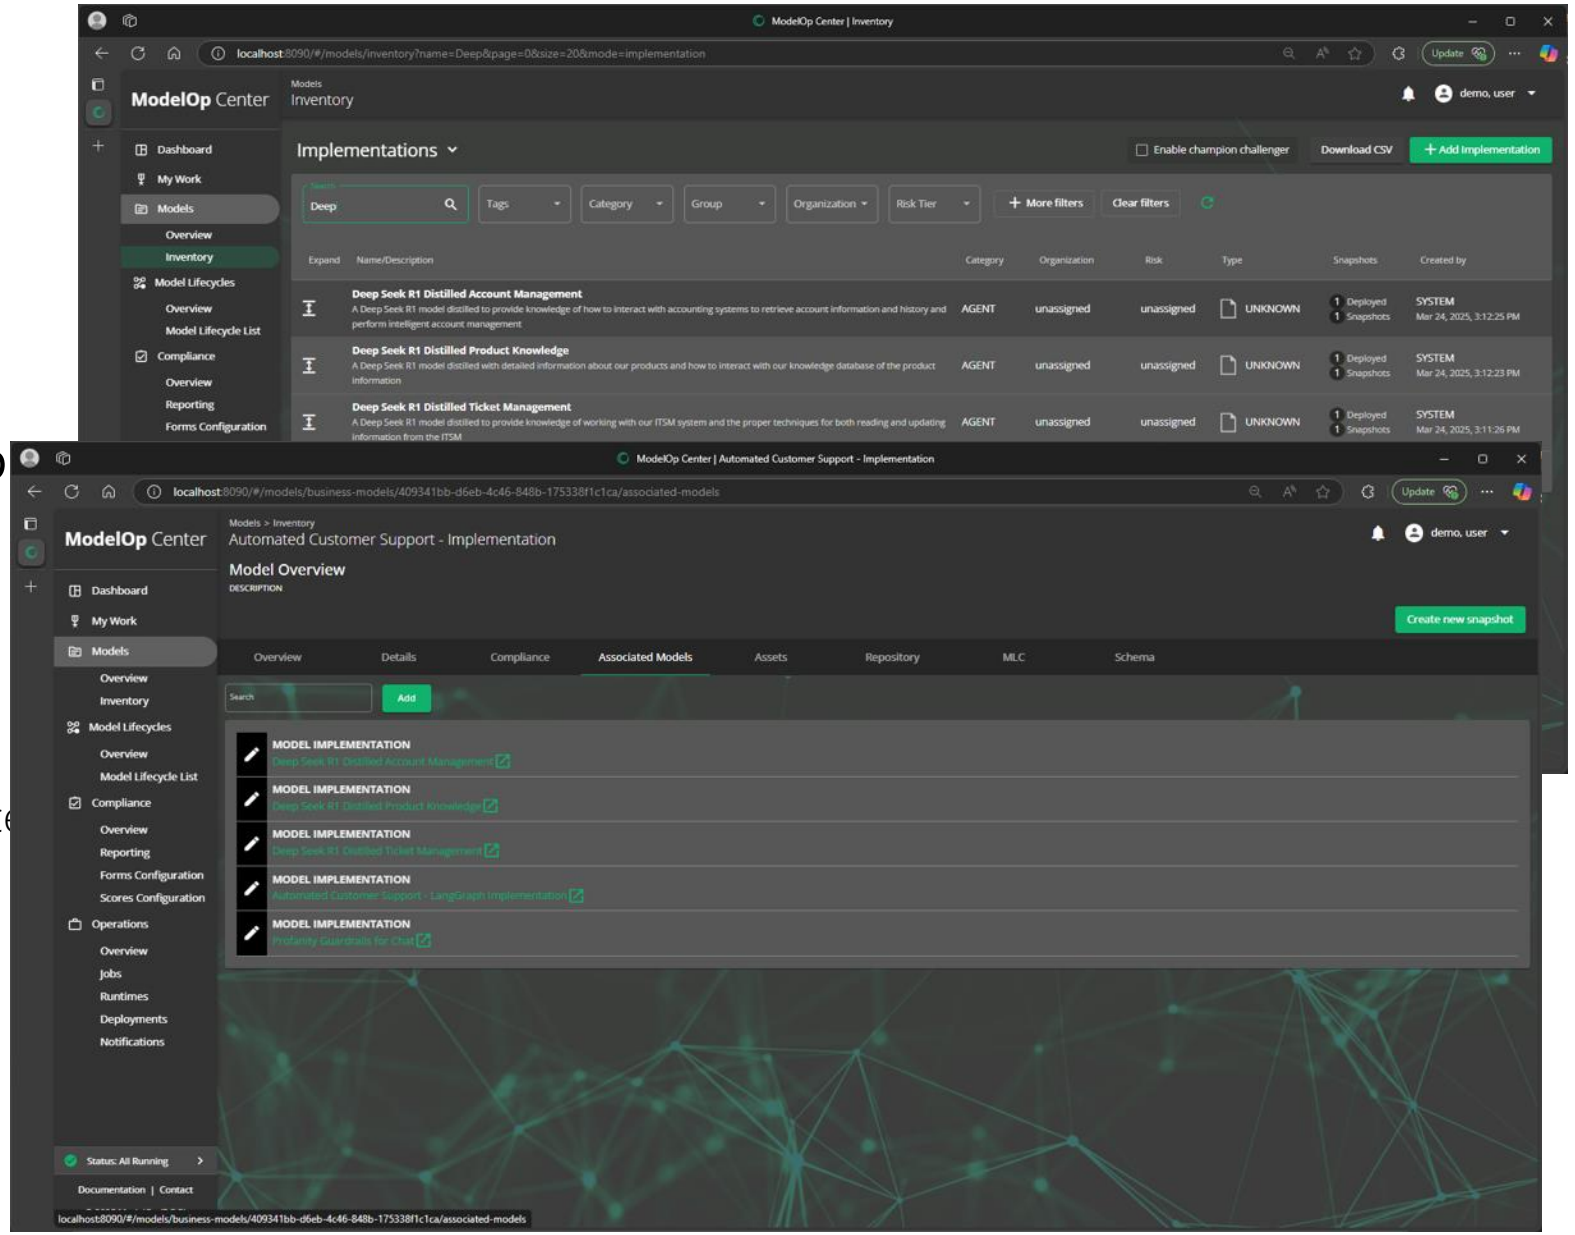

Track Agent and Solution Performance

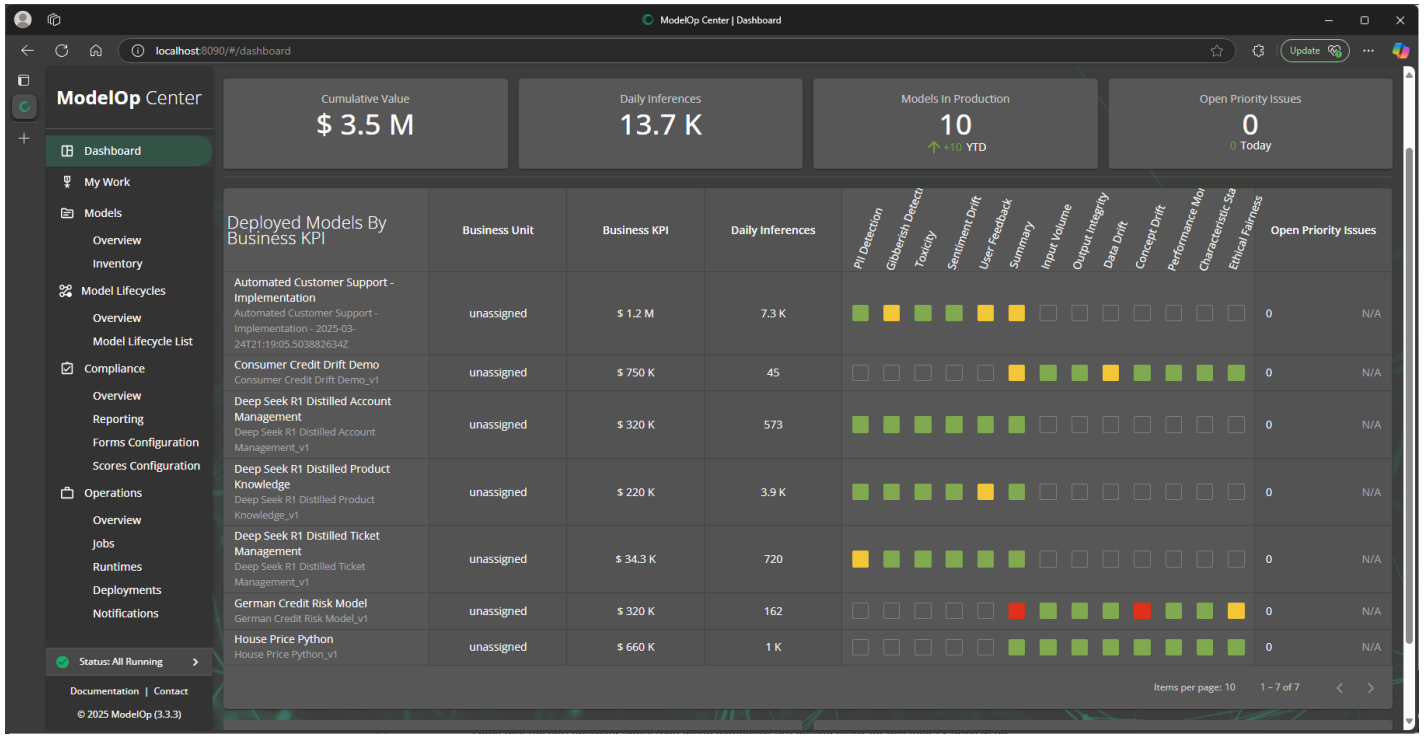

- A dashboard of how agents are performing
  - Shows individual performance of agents in performance of all of their tasks
  - Shows overall performance of agentic solutions
  - Identifies at risk agents, and ties them to their usage

Automate Risk Determination

- Scale risk determination using automated controls
  - Configurable methodology to determine potential risks to the company
  - Assign risks to use cases or models
  - Automatically resolve risks as conditions are met
  - Create a record of all steps taken to address risks
- Automation is key to scaling to agentic AI solutions!

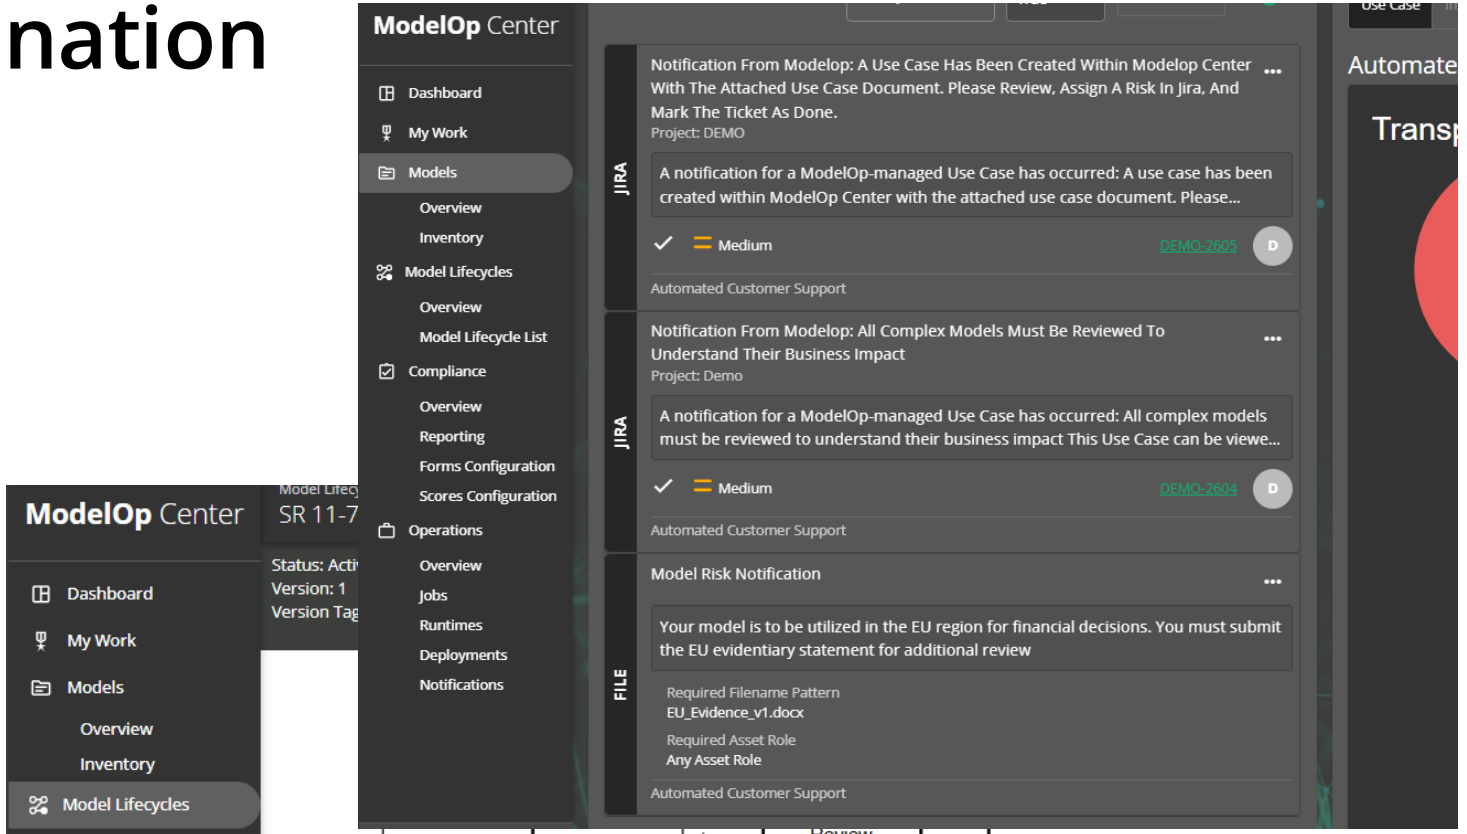

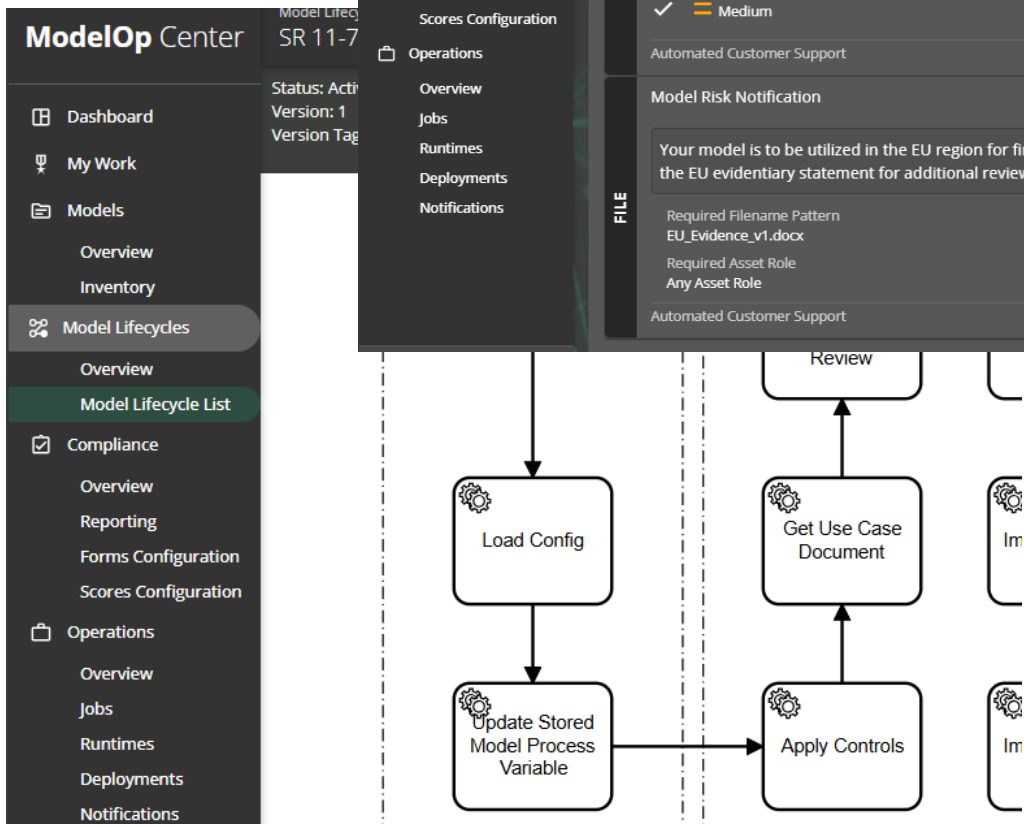

Are you ready for Agentic AI?

- Agentic AI will create an explosion of business-critical models in your enterprise.
  - Models can dynamically interact in unpredictable manners
  - You must know what solutions in your enterprise are using which agents
  - You must understand how each agent is behaving
  - You must understand how your overall solution is performing
  - You must also understand if proper risk management has been performed
- ModelOp Center delivers these capabilities and more
  - Designed from the ground up to be a pure model governance solution
  - Designed with Agentic AI and LLMs in mind
  - Handles all models from Excel spreadsheets to the latest reasoning models

ModelOp

THANK YOU

modelop.com | sales@modelop.com | linkedin.com/company/modelop

ModelOp is the leading AI Governance software for enterprises and helps safeguard all AI initiatives — including generative AI, Large Language Models (LLMs), in-house, third-party, and embedded systems — without stifling innovation. Through automation and integrations, ModelOp empowers enterprises to quickly address the critical governance and scale challenges necessary to protect and fully unlock the transformational value of enterprise AI — resulting in effective and responsible AI systems.