Model Risk Management (External Version)

Modernizing Model Risk Management (MRM)
Coaching, Challenging, and Improving Modeling Teams and Processes
Managing Artificial Intelligence (AI) Risk

February 2021

Jacob Kosoff
Head of Model Risk Management

Table of contents

Modernizing Model Risk Management

Strategic Vision for MRM

Towards Increased Efficiency & Effectiveness

Transparent | Consistent | Risk-Based | Collaborative | Talent Incubator | Automated

Risk-Based, Value-Added, and Timely

Validation activities and model risk management improve the model ecosystem while protecting the Bank and our customers from unintended consequences from model risk. These activities are tailored to the risk of the model and are value-add and timely to encourage faster deployment of models.

Enabling Innovation, Machine Learning (ML), and Artificial Intelligence (AI)

Model risk management should support and, where practical, enable use of innovative modeling to support the strategy of innovating with digital and data. By becoming knowledgeable on new technologies, machine learning, and artificial intelligence, MRM provides valuable feedback and reassures stakeholders that our challenge is effective in managing new and more complex risks coming from more modern data science approaches.

Automated Application of Policies and Standards

Components of model validation requirements and expectations around performance testing, monitoring, and documentation are automated and made repeatable in code, allowing for faster validation and review. By applying policy in code, developers can meet a minimum policy requirement threshold by default. The additional benefit includes repeatability and consistency across the model universe.

Model Risk Management (MRM)

2021 Priorities for Leading Banks

1. **Quality and Continuous Improvement:** MRM provides actionable, timely, and value-added model validation activities. Through process improvement, automation, and training, MRM has improved quality and efficiency of validations.

2. **Timely and Complete:** On track to complete 95%+ scheduled validation activities by 12/31/2021.

3. **Risk Management and Relationships:** Exceeding stakeholder expectations around effective model risk management for capital planning, current expected credit losses (CECL), pandemic model impacts, data controls, overlays, artificial intelligence/machine learning, and strategic priorities dependent on models and data products.

4. **Build the Best Team:** Building out and growing MRM talent to plant roots at the Bank, beyond MRM.
    a) Continue to develop and promote MRM associates to roles within MRM and roles throughout the Bank.
    b) Maintaining less than 10% annual voluntary external turnover.

5. **Innovate:** Maintaining and continuously improving core competencies for machine learning, artificial intelligence, data, digital/streaming models, Python, cloud/distributed computing, open-source libraries, and modern model platforms with software lifecycle management best practices.

Head of MRM

Evolving Responsibilities

Topic / Role	2011	2021	Enhanced Responsibilities
Chair of MGC	NO	YES	The Head of MRM is now typically the chair of the highest-level Model Governance Committee.
Critical beyond CCAR	NO	YES	The use of modeling and quantitative analytics has grown tremendously beyond capital planning.
Significant Analytical Tools	NO	YES	Head of MRM often oversees the examinations of all significant analytical tools and end-user computing tools.
Leads AI Risk Management	NO	YES	Head of MRM usually leads AI Risk Management at financial institutions, including providing leadership on ethical AI and model fairness.
Modernizes Model Risk & provides Guidance on Emerging Analytical Risks	NO	YES	Head of MRM now oversees and establishes standards for validations of models using opensource libraries, models on the cloud delivered through APIs, machine learning models, and other newer modeling approaches.
Agile Model Risk	NO	YES	At most financial institutions, MRM teams now provide agile model risk management within development sprints. In those cases, the Head of MRM is accountable for ensuring that their teams validate models at an earlier stage of their development, with resulting changes and remediation activity becoming less costly and quicker to implement.
Develops Talent to Lead MRM	NO	YES	At many financial institutions, the Head of MRM is explicitly expected to grow internal talent to fill future MRM managerial openings.
Develops Talent to Lead throughout the Organization	NO	YES	As one of the senior data scientists at the firm, the Head of MRM is often expected to coach and mentor other leaders at the financial institution on quantitative talent development strategies. The Head of MRM is often expected to develop and prepare quantitative talent to fill critical people management roles throughout the firm.
Provides Enterprise-wide Data Science Training	NO	YES	At a few financial intuitions, the Head of MRM takes accountability for internal Analytics Institutes, Quant Summits, and quantitative talent development. The Head of MRM strategizes on curriculum needed to continue to develop and retain best-in-class quantitative associates and data scientists.
Strong Reputation with Local and National Regulatory Agencies	NO	YES	At most financial institutions, the Head of MRM is expected to build and continuously strengthen trusting relationships with national and local leaders at the FRB, OCC, CFPB, and FDIC.
Builds and Strengthens Relationships with Internal and Industry Peers	NO	YES	The Head of MRM is usually expected to develop trust with internal peers and industry peers to facilitate the exchanging of industry best practices. Must be recognized as forward-thinking, knowledgeable, and as a peer to be invited into certain dialogues.

MRM Expertise

- At many financial institutions, MRM teams are responsible for being experts across a variety of topics, leading internal Analytics Institutes, Quant Summits, and quantitative talent development. The MRM team is crucial in Bank-wide development of best-in-class quantitative associates and data scientists.
- The MRM team leads the way in establishing standards for validations of models using opensource libraries, models on the cloud delivered through APIs, machine learning models, and other newer modeling approaches, thus, expertise is required:

Topic	Details
Core Foundations	Python, R, Git, CDSW, Open-Source
AI and Machine Learning	NLP, Random Forest, Deep Learning, Gradient Boosting, SVM, General Unsupervised, TensorFlow
Data	Unifi, Hadoop Ecosystem, Data Governance, RCIF, Pipelines, Controls
Architecture & Engineering	APIs / Openshift / MuleSoft, Containers, Bamboo, Cloud, HBase, Graph, Streaming
DevOps & Automation	Bamboo, CI/CD, Unit Testing, Integration Testing, Airflow, Sphinx, Linting, APIs
Special Topics: Streaming	Apache Kafka & Flink, PySpark, PyFlink
Special Topics: Graph	GraphX, Apache Giraph
Special Topics: Advanced Programming	Functional Programming, SOLID, Scala, Clean Code and Refactoring

MRM Automation

Process & Key Outcomes

Process: 2018-2020

- Products, Outcomes, Resources, and Talent needs defined for automation
- Proof of Concepts (POC) completed:
  - Ongoing Monitoring POC
  - Test Library POC
  - NLP Search POC
  - Automated publication of documentation POC
- Data Lake Mart
- API Links

<u>Testing Library</u>

- Functional, domain-specific validation testing library for integration testing

<u>Annual Reviews & Significant Changes</u>

- Initial annual reviews automated on new model development
- Significant change reviews automated for new or migrated models

<u>Ongoing Monitoring Automation</u>

- Data architecture & MRM standards
- Code template for model developers
- Active dashboard

Key Outcomes: 2021

- Complete Annual Review inventory migration
- Completeness of testing libraries and other tools
- Enhancements to existing tools
- Adoption of new domains or methodologies into testing libraries

Activity or Product	Obj 1 Automated Testing	Obj 2 Policy as Code	Obj 3 Automated Documents	Obj 4 Repeatable Monitoring	Net Benefit w/ Automation
A) Reusable testing code	✓	✓		✓	12%
B) Model submission tool	✓			✓	4%
C) Enterprise monitoring dashboard		✓		✓	6%
D) Automated Documentation		✓	✓		15%
E) Annual Review Automation	✓	✓	✓	✓	20%
F) Significant Change Review	✓	✓	✓	✓	10%

Coaching Modeling Teams

MRM coaches, challenges, and improves modeling teams

Modeling teams before MRM coaching

Monitoring programs are manual, custom, and cannot easily be aggregated for reporting or analysis.

Model development and coding is not standardized: Does not follow software development best practices, including modularization, source version control (git), unit testing, integration testing, a production pipeline with technology controls, team or enterprise coding standards.

Documentation is manual and is not usually version controlled.

Significant Changes are manual and require manual steps for build and testing by developers. Results in manual documentation updates and notifications.

Validation and Annual Reviews are manual and completed after development is completed or after manual submission of materials by model developers.

Model Submissions are manual and only experienced developers fully understand all requirements.

Modeling teams after MRM coaching

Monitoring programs are templated, are automated, and can be aggregated for enterprise reporting.

Model development and coding is standardized: By following best practices of software development, code updates are easier to make, easier to track, less prone to bugs (unit testing), can apply integration and validation tests in the build process, and applies standard code requirements to increase reusability an understanding across the Bank.

Data capture is standardized and centralized with additional support by feature stores to allow appropriate use of Big Data.

Documentation is partially automated with version control.

Significant Changes can be automated into the dev/test/prod build and promotion steps of model development.

Some **Validation Tests** are standardized, automated and completed as part of any new model development, model update, or annual review.

MRM Recommended Modeling Practices

MRM coaches model development teams to use modern practices

Data Pipelines
Model Pipeline
Software Development Pipeline

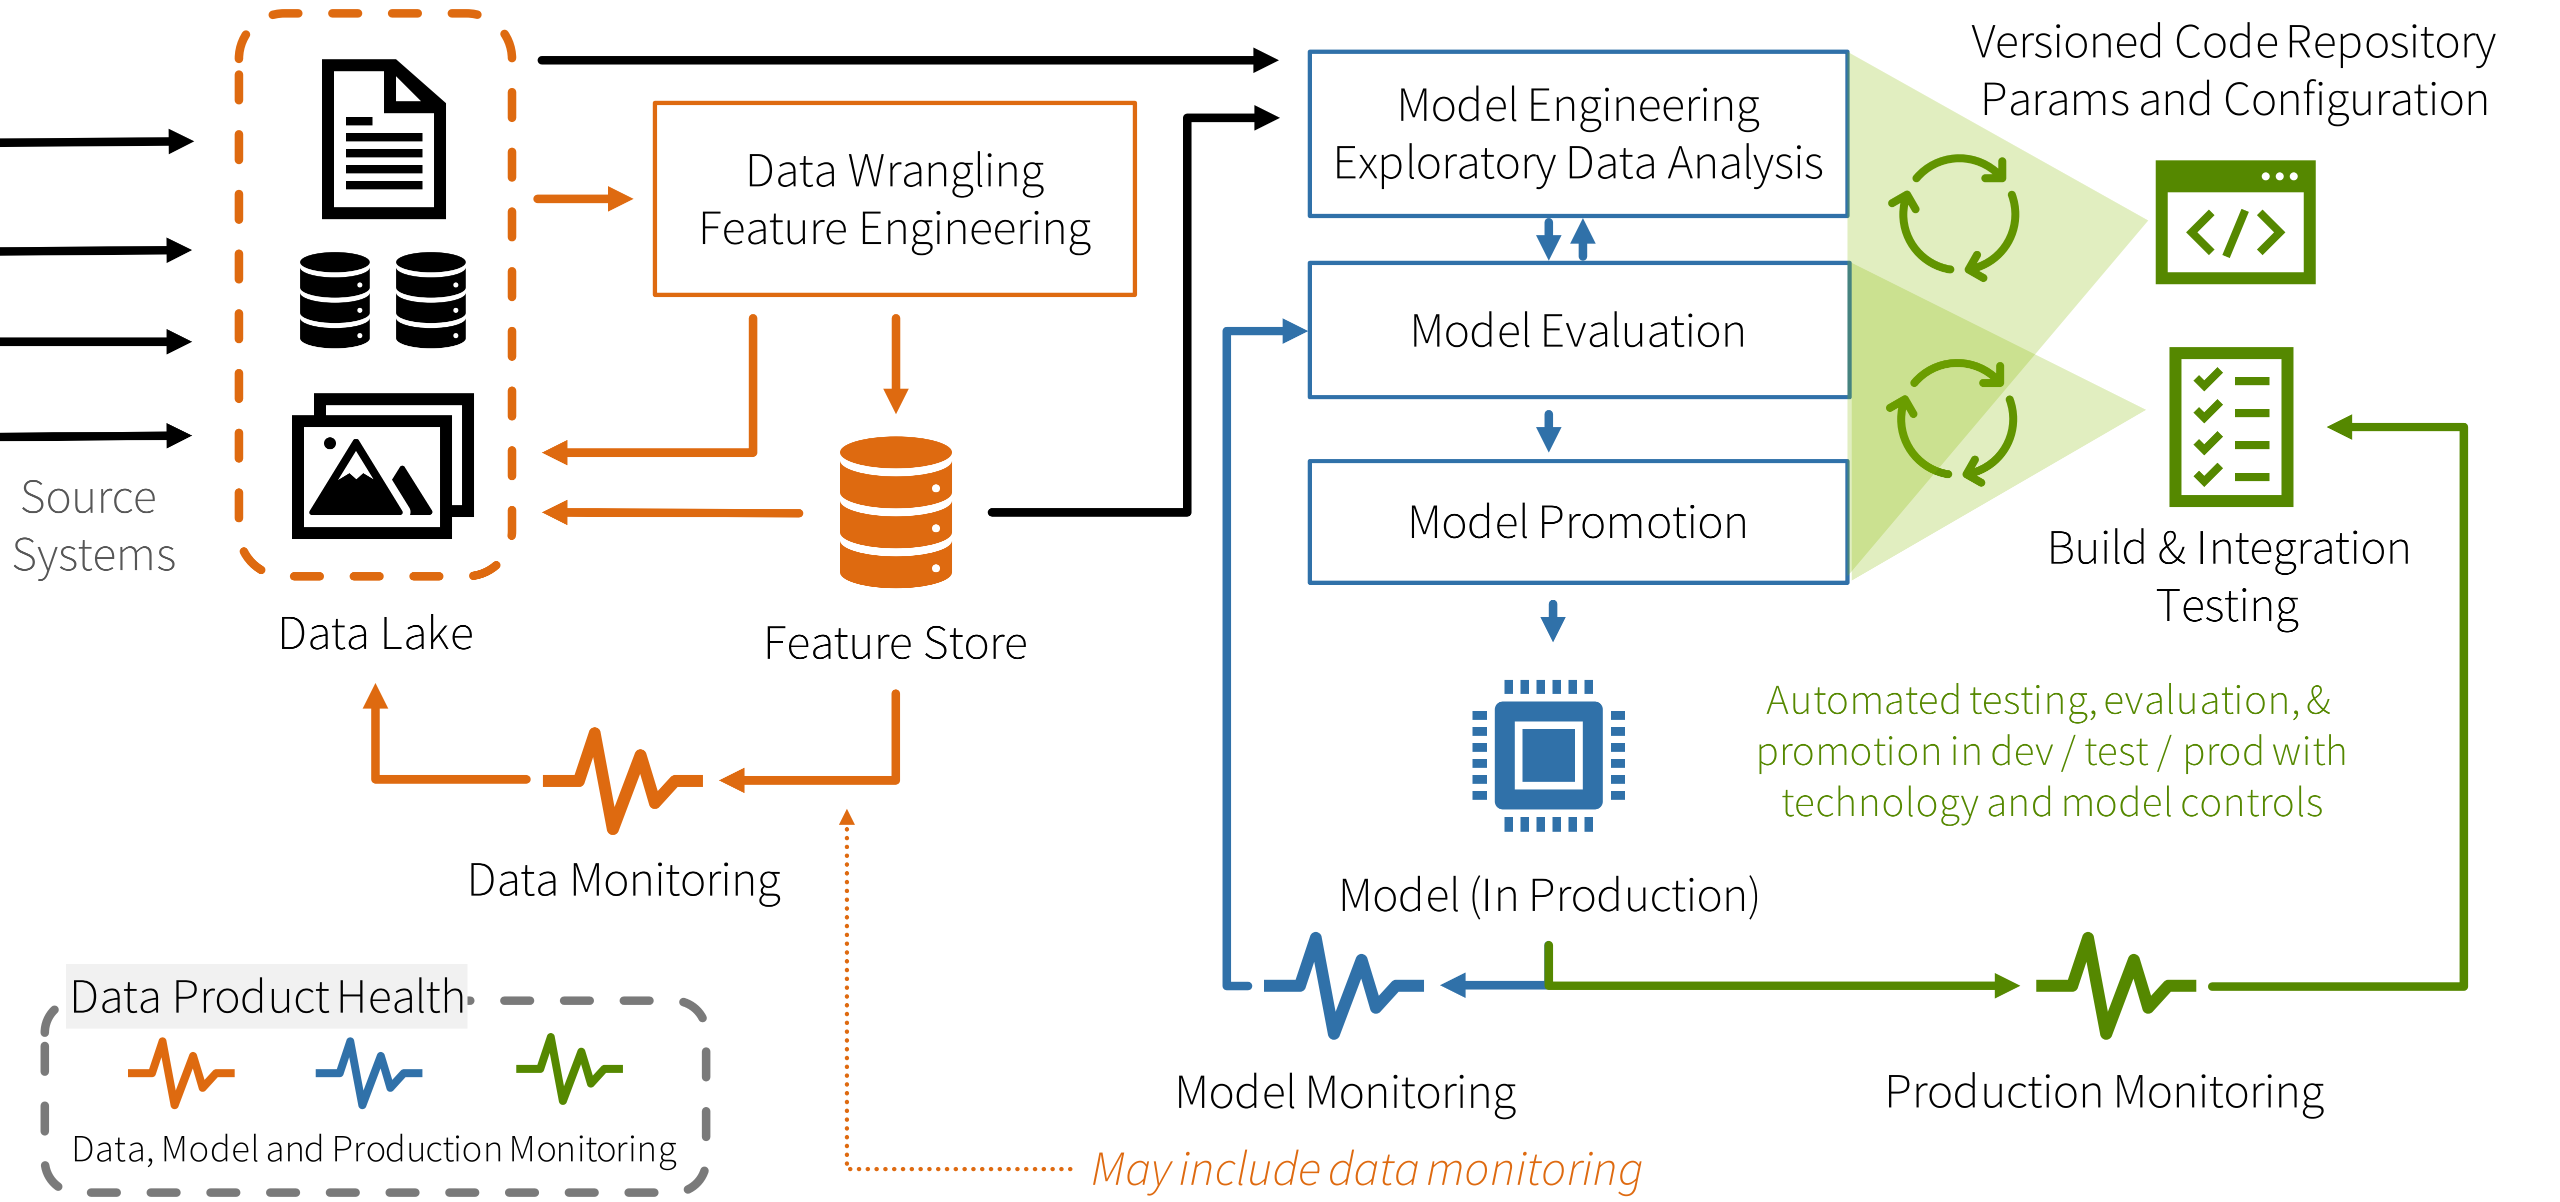

2021 Changes for Model Owners

Best practices have become expectations.

1. **Write findings to the new/updated model owner procedures.**
   MRM has updated the model procedures for model owners and developers and has trained all stakeholders on these procedures. These procedures reflect more modern and agile development practices. MRM writes findings when practices do not align with updated procedures.

2. **New models must be developed on the centralized data science platform.**
   Existing models must migrate to the data science platform and all new models must be built on the data science platform. Models developed outside the data science platform results in a finding.

3. **MRM is developing a centralized model monitoring standard.**
   In coordination with data risk in IT risk management (ITRM) and data governance and model platforms in the data analytics office (DAO), new monitoring requirements will be required for new development later in 2021.

4. **Automation activities will no longer consider compatibility with SAS models.**
   Automation will focus on models deployed to the model platform ecosystem with a focus on ML Ops.

5. **Production model processes are handling data from approved sources.**
   Local data marts, intermediate data marts, and models where owners/developers can overwrite model data is a finding.

Expectations for Model Developers

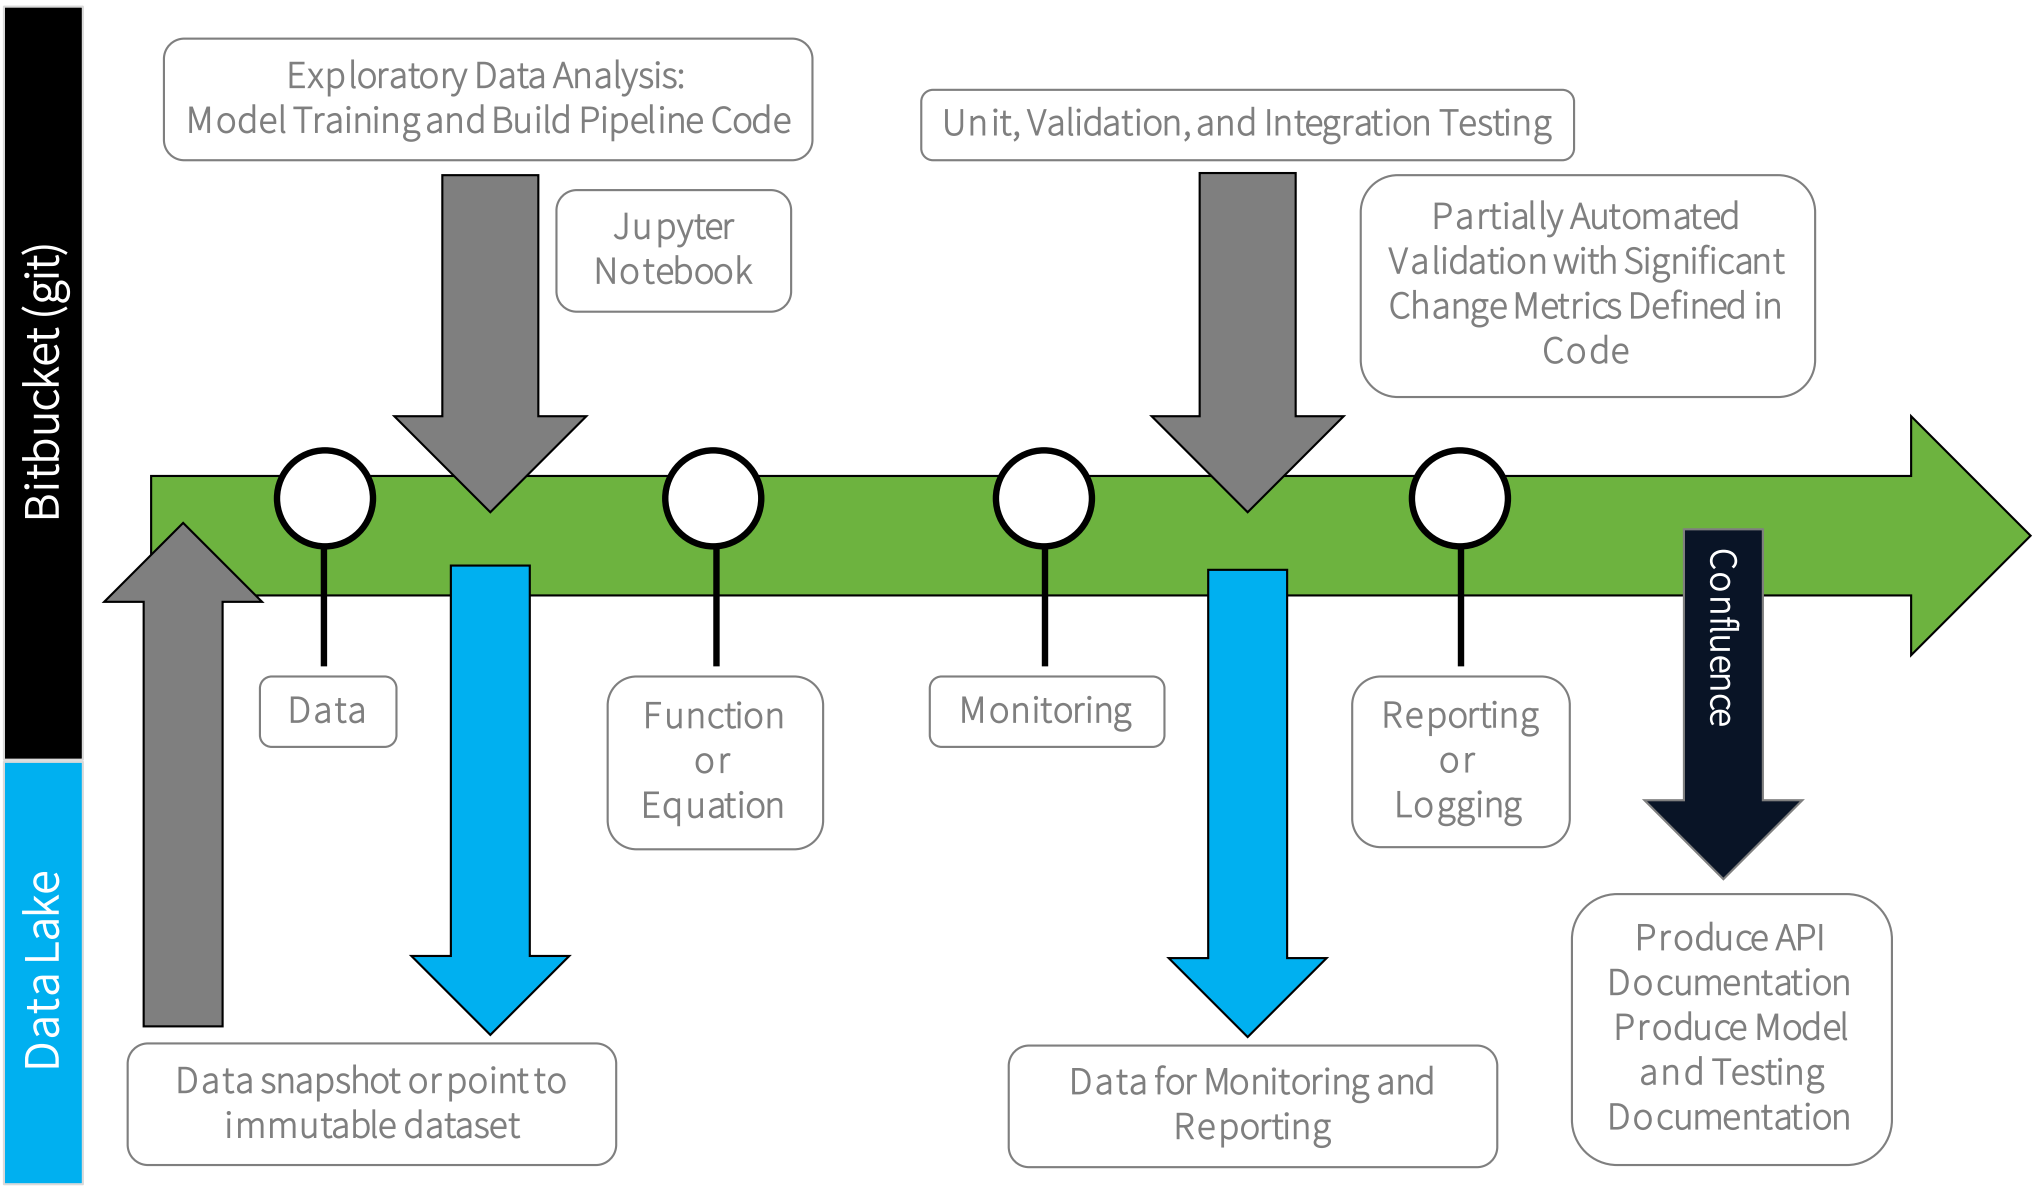

Effective coaching to modeling teams leads to faster model validations

Automated Testing: Validation requirements, expectations, and model testing is automatically and thoughtfully built into the model process.

Policy as Code: Templates reduce reliance on reading hundreds of pages of policies, procedures, and PowerPoint decks.

Automated Documentation in Pipeline: Documentation is embedded with model design, is version controlled with code, and produces repeatable language.

Modularization and Cataloguing: Reusability & Extensibility:

- PySpark functions can be reused by other teams and extended to other solutions (Rapid Innovation).
- Standardize function design/documentation, create discoverable functions, and catalogue.
- Create digital style guides and standards for developers to facilitate adoption of best practices: Stylesheets can be updated in real time.
- Create and validate once, then reuse (efficiencies for developers and model validation).
- Migrate to PySpark, which is portable across teams, takes advantage of distributed computing in the data lake, has numerous libraries, and lends itself to automation.
- Modularization allows for division of tasks, facilitates re-use of code, improves change management, allows for automated testing, and reduces MRM burden.

Feature Store Approach: Feature stores improve collaboration, enable rapid development, simplify automation, capture data lineage, speed up the transition to production, and encourage re-usability.

Models as APIs: Banks use output across teams to create automated execution and reporting of CECL, CCAR, and other aggregate model processes.

Agile Approach: Agile teams are typically cross functional and consist of members from different teams, encouraging a collaborative approach. Agile lends itself to projects with a sense of urgency, significant complexity, and those which require an iterative process.

Managing Artificial Intelligence (AI) Risk

Regulatory Environment for AI

Regulatory scrutiny of financial institutions' use of artificial intelligence (AI) and machine learning (ML) continues to increase with a focus on bias, explainability, and fairness.

- The FRB hosted an "Ask the Regulators" on the use of AI/ML on December 16th, 2020.
- On Jan 12, Federal Reserve Governor Brainard's speech focused on AI and Fairness.
- On Jan 12-13, the Federal Reserve Board of Governors hosted a two day "AI Symposium" addressing topics such as AI interpretability, explainability, bias and equitable outcomes.
- The Annual Model Risk Forum, hosted by the Federal Reserve on Jan. 26-27, provided a supervisory view update on the use of artificial intelligence (AI) and machine learning (ML).
  Key takeaways include:
  - Important for Board and Senior Management to think carefully about the risks from AI/ML and how to manage them
  - Key challenges of AI and ML include data quality, accuracy, protection, bias, explainability, fairness, unlawful discrimination, consumer compliance, scarcity of talent or expertise, adoption based on hype rather than informed decision-making.

Top Risk Areas

AI and Fairness

Key areas for Fairness, Explainability, Consumer Protection, and Bias

- Consumer Credit – Mortgage
- Consumer Credit and Small Business (Non-Mortgage)
- Marketing for Credit Products
- Marketing and Fraud (Consumer/Retail)
- AI and ML – Credit origination and portfolio monitoring
- AI and ML – Very Low/Negligible Risk models that may need to be upgraded
- AI and ML and uses of Alternative Data
- Tools for explainability

AI Risk & Lessons Learned

- AI may not improve the accuracy and fairness of decisions on consumer credit.
- There are concerns are around predictive power, stability, overfitting, performance through the cycle.
- AI expands model stakeholders as AI can be used for automation and other uses not associated with traditional bank models.
- AI and ML are types of models that fall under SR 11-7 and are governed by institutional model risk policies/procedures.
- Banks must be committed to ethical AI.
- AI solutions may be part of automated systems that require innovative approaches to monitoring and validation testing.
- AI and ML requires upskilling, new processes, new documentation standards, and other adjustments that address the unique characteristics of AI.
- Banks must establish best practices for open-sourced programming.
- While AI & ML can offer some improved model results, a bank should consider whether it is appropriate to accept this additional complexity and testing and monitoring involved:

RMA Journal "Understanding and Validating the Uses of ML Models in Banking"

https://www.rmahq.org/understanding-and-validating-the-uses-of-ml-models-in-banking/

Managing Artificial Intelligence Risk

AI Strategy

- Formalized AI Strategy: MRM directs AI Risk Management as a subset of MRM under the Model Governance Committee and is partnering with key stakeholders, including Compliance and IT Risk Management.

Risk Appetite Statement (RAS), KRIs

- Risk appetite statement includes data, models, data products, machine learning, fairness, or artificial intelligence.

AI Working Group and Risk Partnerships

- Formalizing AI working group and partnerships with Compliance and Regulatory Risk Management .

Standardized Testing Workpapers

- AI and Machine Learning standardized testing workpapers to supplement existing workpapers and form a basis for consistently evaluating AI.

Talent Training for Artificial Intelligence (AI) and Machine Learning (ML)

- Standardized training incorporated into Degreed pathways and onboarding programs to form a basis of understanding for emerging regulatory and administrative risks tied to AI and ML.

MRM's AI/ML 2nd Line Framework

MRM and Model Governance Committee oversee 2nd Line Framework for AI/ML Risk, collaborating with key partners across the bank.

Enterprise Risk

	Framework
Process Governance	• Integrate view of risks and set effective processes for coordination and accountability across multiple stakeholders. • Ensure AI enablement using a consistent and risk-based approach, while ensuring business continuity. • Brings visibility and awareness of AI Risk to the Bank.
Conceptual Design	• Consider criteria for enhanced model risk tiering and assessment considering intrinsic risk of AI/ML methods. • AI/ML-specific considerations such as: Use-case specific applicability; data quality, stability, validity and bias, interpretability of black-box algorithms, unintended bias, assessment for unintended outcomes, benchmarking, augmented KPIs and change management.
Technology & Data	• Assess risk associated with integration of vendor solutions, technology architecture, and data aggregation from third parties. • Incorporate data engineering and preparation to consistently address usefulness, quality, format, and value of data • Incorporate AI data risk into scope of existing risk data controls, aggregation, and reporting. • Strengthen technology capabilities to address aspects such as explainable AI and data drift. • Reconcile internal and external touchpoints such as cyber security, vendor solutions, data pipeline, and data sourcing.
Regulatory Compliance & Responsible AI	• Ensure transparency of AI systems in view of ethics and fairness. • Manage consumer perception and reputational risk effectively. • Align risk assessment with compliance and regulatory frameworks such as Fair Lending rules and Data Privacy (e.g. - FCRA, FERPA, GLBA, HIPAA, and GDPR). • Establish accountability for AI/ML ethical risk monitoring and risk appetite.
People & Culture	• Establish culture that ensures holistic risk management. • Build competency and skills to understand, assess, and manage AI-related risk.

Partners

Compliance Risk Management Committee

Credit Risk Management

Data & Analytics Office

Internal Audit

IT Risk Management

Model Risk Model Validation

New Initiative Risk Assessment Committee

Operational Risk Committee

Strategic Initiative Approval Committee

Technology Risk Oversight Committee

Third Party Risk Oversight Committee

Risks

<u>Compliance Risk</u>
- Regulatory: Consumer Protection
- Financial Crimes: BSA/AML
- Fair Lending Compliance
- UDAAP Compliance
- Privacy Compliance
- Financial Crimes Risk Management

<u>Non-Financial Risk</u>
- Execution, Delivery and Process: Monitoring & Reporting
- IT Risk Management
- Third Party Risk Management
- RCSA Program
- Scenario Analysis
- Operational Loss
- SOX

<u>Reputational Risk</u>
- Reputational Risk following Compliance, Non-Financial, and Legal Risk events or issues
- Media Relations
- Issues Management
- New Initiative Risk Assessment

<u>Legal Risk</u>
- Litigation, Enforcement Actions, and Dispute Resolution largely driven by Compliance and Operational Risk

Aligned Risk Quantitative Resources

Modernized approach to managing AI risk requires alignment of quantitative resources within a Bank's Risk function.

- Clearly documented and understood roles & responsibilities for quant resources within Risk Management.
- Prioritized development and use of all models in Risk Management, aligning to key projects/bank-wide initiatives while ensuring the most efficient model risk management and validation activities.
- Rationalizing models that are in place, need to be rebuilt, or are inefficient in their use of resources.
- Setting clear and consistent guidelines for the use of tools, data, and modeling methodologies in Risk Management.
- Increase model effectiveness, enhancing model monitoring approaches to better track effectiveness and efficiency.
- Improve quantitative career development and architecture, including consistent, structured, and exceptional onboarding and training experience.

Artificial Intelligence

Risks and Guardrails

AI Model Risks	Guardrails to Mitigate the Risk
Unnecessary complexity leading to overfitting of the model and/or lack of transparency.	Perform benchmarking using simpler or alternative solutions, like regression or decision trees. The goal is to identify the marginal predictive power of a more complex AI method over a more traditional transparent model and determine is the emerging risks outweigh the benefits.
Data Integrity - Inaccurate, inappropriate or missing training data, or data whose predictive potential changes or evolves over time.	Ensure that the training data is appropriate for the model's purpose and use (e.g., inclusive training data). Ensure that any data used to develop and tune the model is properly and consistently labelled. Perform ongoing monitoring for data.
Lack of model explainability and the most important variables influencing the output of the model.	Use statistical or graphical methods in parallel to AI solutions to isolate the most important variables and their connection to the output of the model.
Model/Data Bias - Bias/inaccuracies in the model due to biased training data.	Run statistical tests on the training data to ensure that the dataset is representative of the data it is intended to model and does not have inherent disparities. (e.g., Data used to train lending models should be inclusive of, and free from bias against, all protected classes: race, gender, age, etc.). Care should be exercised in the use of data that may be a proxy for or correlated with race or other prohibited bases. Data bias may not manifest until the output is created; thus, output is tested for potential statistical disparities. One can then use this information to adjust the training data in a way that may be more representative – or assess whether the training data is even appropriate for the model.
Lack of model transparency due to black box vendor restrictions.	Ensure sufficient performance statistics and analyses are available to conduct the ongoing performance monitoring necessary to evaluate overall performance and emerging behavioral changes.
Lack of model replicability and/or stability.	The stability of the variable effects can be evaluated based on hyper parameter sensitivity analysis and/or refitting the model on a different random data sub-sample. Feature and model drift can be assessed.
Lack of robustness of model fit.	Developers can test multiple ML methods, choices of assumptions, and different segmentation schemes using both expert judgment and data analysis.
Environmental changes affecting model performance.	Performance monitoring.

Cloud-Specific Model Risk

More vendor machine learning models are on the Cloud and are delivered as a “Models-as-a-Service” through APIs.

Model Risk teams must adapt processes around the three versions of Models as a Service:

1. Standard model served as a static API.
2. Dynamic and retraining model based on our and other banks’ data.
3. Customized product tailored for the Bank.

Validation Expectations for AI/ML Models

Validation Expectations

- Strong monitoring programs for all AI and ML models.
- Alternative design assessment for new methodologies applied on a team or at the Bank.
- Additional scrutiny for open-source libraries that are new to the Bank or not widely adopted.
- Evidence of subject-matter-expertise to apply AI and ML models and to limit key-person risk.
- Appropriate assessment of data for potential for bias and business justification for variable use (where appropriate).
- Hyperparameter sensitivity analysis and evidence of appropriateness of data for hyperparameters.
- ML models can overfit; purpose of regularization as part of the hyperparameter set.

Natural Language Processing (NLP) and Understanding (NLU) – Specific Expectations:

- Quantitative monitoring metrics.
- Manual review thresholds for critical processes.
- Stored data on speech-to-text and text understanding to independently monitor both engines (if applicable).
- Comparison to simpler techniques using Regular Expressions and SME rules (benchmarking, risk-based).
- Re-training from user feedback should be monitored and setup appropriately.
- NLP/NLU techniques (model and non-model) are evolving. New applications should consider relevance of new techniques to proposed methodologies based on risk of model.
- NLP/NLU techniques often involve various forms of neural networks: sufficient data and relevance of data must be established.