



WHITE PAPER

AI Solution Development Costs:

A Comprehensive Guide to Cost Control for Traditional and Agentic AI

By Ramki Krishnamurthy, Data Analytics Lead, REI Systems



SEPTEMBER 2025

- 03** *Introduction*
The Challenge: Why Cost Management Matters
- 04** *Introducing FinOps*
The Lay of the Land
- 05** *FinOps and Data Management*
Core Factors Governing AI Solution Development Costs
- 07** *How Requirements and Design Choices Influence Each Cost Factor*
- 12** *Cost Ranges Based on Solution Type and Complexity*
- 14** *Tools for Rough Estimates*
- 15** *Conclusion*
References

THE
NEW
FIN
OPERATIONS
COST



Introduction

In an era where federal agencies face growing demands to deliver outcomes faster, smarter, and with fewer resources, Artificial Intelligence (AI) has emerged as a mission-critical enabler—driving automation, accelerating decision-making, and unlocking new levels of operational efficiency. While much debate exists about when to use AI, its impact on the workforce, and its responsible uses, the cost of implementing these solutions may preempt all other adoption drivers.


With federal agencies accelerating their adoption of AI – over 1,700 AI use cases have been reported across 37 agencies in 2024 alone (more than double the previous year)¹ – understanding and controlling development costs has become critical. This white paper is an expanded version of the concepts discussed in [“Three ways agencies can prepare before AI costs skyrocket - Nextgov/FCW.”](#)

AI projects can incur runaway expenses if not meticulously planned, especially as agencies explore advanced “Agentic AI” (autonomous, decision-making AI systems) that introduces new cost drivers. **This guide outlines the core cost factors in AI solution development and provides strategies to optimize costs, with examples and case studies relevant to government contexts.** The goal is to help federal stakeholders implement AI cost effectively while maintaining the high standards of performance, security, and ethics expected in the public sector.

The Challenge: Why Cost Management Matters

The promise of AI is great, but so are its costs. Federal CIO reports show agencies are leveraging AI to improve efficiency and mission execution (e.g. anomaly detection, streamlined processes, better decision-making)¹ – from **helping patent examiners at USPTO to fraud detection in VA benefits**¹. However, these gains can be undermined by unbridled spending. AI costs are multifaceted – spanning data, computing (training and ongoing inference), talent, integration, and more – and **often extend far beyond initial hardware/software investments**². As Army CIO Leonel Garciga noted, a big lesson learned is that AI is “expensive stuff to do,” which is driving the service to “tighten the guardrails on use cases” and ensure the “juice is worth the squeeze.”

Government-wide guidance on AI acquisition has been issued with [OMB Memorandum M-25-22](#), and agencies must actively implement its directives and proactively ensure careful planning and oversight to prevent budget overruns while meeting mission objectives. If done correctly, the **long-term benefits of AI will far outweigh the costs**² – but achieving this requires strategic cost management from day one.



¹AI in Action: 5 Essential Findings from the 2024 Federal AI Use Case



Introducing FinOps: A Framework for Cloud Cost Accountability

To effectively manage the variable and often unpredictable cloud expenses associated with AI, it is essential to adopt a FinOps (Finance + DevOps) framework. FinOps is a cultural practice that brings financial accountability to the cloud spending model, enabling organizations to maximize business value by helping engineering, finance, and business teams collaborate on data-driven spending decisions. For AI, where computational needs can spike unexpectedly during training or scale rapidly for inference, FinOps provides the necessary structure to prevent runaway costs while fostering innovation.

The Lay of the Land: Data-Related Costs

Data infrastructure and management are primary drivers of AI implementation costs. These costs can quickly become unmanageable without thorough planning. Data requirements vary significantly based on the specific AI use case an agency is considering. Therefore, a complete and comprehensive assessment of data needs is essential. This assessment should encompass:



Data Volume: The total amount of data required for training and the data that will be generated and processed during ongoing operations.



Data Variety: The various types of data being analyzed, such as structured, unstructured, or semi-structured data.



Data Velocity: The speed at which new data is generated, collected, and needs to be processed.



Data Storage: The choice of storage options, considering performance, scalability, and security to find a balance between high- and low-cost solutions.

For **Agentic AI solutions**, the data considerations are significantly amplified. These solutions often require vast amounts of diverse data for training, including observational data from interactions, textual data for reasoning, and, potentially, sensory data from environments. Agentic systems typically operate in dynamic environments, which drives the need for high-velocity data infrastructure to support **real-time data ingestion and processing**. Furthermore, because Agentic AI learns from its actions and outcomes, it necessitates robust data collection mechanisms for **continuous feedback loops**, which adds to both the volume and complexity of the data that must be managed.

FinOps and Data Management

FinOps supports data governance at this level by aligning financial accountability with data lifecycle management, ensuring that both short-term operational needs and long-term archival strategies are cost-effective and compliant. In almost real-time, FinOps practices enable visibility into the costs associated with data ingestion, transformation, and storage across environments (e.g., data lakes, cloud object storage, analytics platforms). It encourages proactive rightsizing, tiering, and deletion policies to prevent unnecessary data sprawl and reduce spend on high-performance storage.

From a long-term perspective, FinOps facilitates cost-aware archival strategies by helping teams evaluate and transition cold or infrequently accessed data to lower-cost storage tiers (e.g., Azure Archive Storage, AWS Glacier). It ensures that governance policies—such as retention, encryption, and access control—are implemented within financial constraints and tracked for compliance. By integrating with tagging policies and usage attribution, FinOps enables traceability and accountability for data storage decisions, ensuring that governance is not only a technical requirement but also a financially optimized practice.

Agencies should develop a budget that accounts for data collection and storage, how to govern the data, the ongoing management of that data, and the capabilities and cost of cloud-based offerings in the market (for example, Amazon Web Services vs. Microsoft Azure vs. Google Cloud). This budgeting should have a development component, an operational component and a retention component.

For example, at the Federal Emergency Management Agency (FEMA), as part of the development of the ACT-IAC Igniting Innovation award-winning **FEMA Data Exchange (FEMADex)**, REI conducted a data and data maturity assessment prior to developing a full data strategy designed to help modernize how FEMA shares and uses data.

Crucially, agencies must prioritize data quality, security, and privacy by establishing a proper data governance framework with clear procedures and rules for who in the organization can access the data or have the responsibility to manage it. Additionally, data management tools are necessary to handle the day-to-day processing tasks. [Read our white paper on [Data Strategy and Governance](#).]

Core Factors Governing AI Solution Development Costs

While exact figures vary widely based on the specific AI solution, complexity, and industry, we can outline a general framework for how costs are typically allocated across components of AI solution development. These allocations are not static and will shift based on the design choices, scale, and operating model of the solution.

FinOps models can enable cost allocation to specific teams, models, or projects, and provide data-driven decision making around resource efficiency, right-sizing, and commitment purchasing (e.g., Reserved Instances or Savings Plans). This collaborative approach ensures that AI investments align with business goals while preventing cost overruns and enabling continuous optimization as models scale and evolve.

Note that Agentic AI solutions will generally skew the below proportions towards the higher end for data and computational aspects, and introduce a new "Safety, Governance, and Responsible AI" category to consider.



Data Infrastructure and Management: 30-50% (*potentially higher for Agentic AI*).

This encompasses the costs associated with data collection, storage, governance, quality assurance, and ongoing management tools. As highlighted above, this is often the most significant cost driver due to the sheer volume, velocity, and variety of data AI solutions consume.



Computational Infrastructure: 20-40%. The processing power needed for AI models, especially for training and inference, can be substantial. This includes the cost of servers, GPUs, specialized AI accelerators, and networking equipment or semi-structured data.



AI Model Development and Customization: 15-30%. This involves the costs of developing, training, testing, and fine-tuning AI models. This can include salaries for data scientists, machine learning engineers, and domain experts, as well as the cost of specialized software and platforms.



Integration and Deployment: 5-15%. Costs related to integrating the AI solution into existing systems, deploying it into production environments, and ensuring its compatibility and performance.




Monitoring, Maintenance, and Updates: 15-30%. Ongoing costs for monitoring AI model performance, maintaining the infrastructure, troubleshooting issues, and updating models to ensure continued accuracy and relevance.



Safety, Governance, and Responsible AI: 5-15%. This introduces significant costs related to ensuring safe operation, establishing robust governance frameworks for autonomous behavior, and addressing responsible AI use.

It's important to note that these are rough estimates, and factors like the availability of pre-existing data infrastructure or the use of highly specialized, custom models can significantly alter these proportions. If agencies can account for and model their costs in these categories, they will have much better and robust control over AI costs.



How Requirements and Design Choices Influence Each Cost Factor

This section explores how specific requirements and design choices influence cost factors, with a focus on Agentic AI and the application of FinOps principles.

1. Data Infrastructure and Management

Cost Driver	Tradeoffs	Recommended Action
Data Volume & Velocity	Larger datasets and higher data ingestion rates directly translate to increased storage and processing costs. For Agentic AI, this is intensified by the need for real-time data streams and continuous feedback loops.	Choose cost-effective storage solutions (e.g., cold storage for archival data vs. high-performance storage for frequently accessed data) and scalable data pipelines to handle high-velocity data streams.
Data Variety & Complexity	Handling diverse data types (structured, unstructured, semi-structured) and complex data formats requires more sophisticated data processing tools and potentially more expensive data engineering expertise. Agentic AI systems often require vast, diverse datasets, including observational, textual, and sensory data.	Adopt a multi-modal database and a flexible data processing framework (e.g., Apache Spark) to handle diverse data in a unified and cost-effective manner.
Data Quality & Cleansing	Poor data quality necessitates extensive data cleaning and pre-processing, which adds significant labor and computational costs.	Implement automated data validation pipelines and establish clear data ownership roles early to reduce manual cleaning efforts, improve model performance, and contain long-term operational costs.
Data Governance & Security	Implementing robust data governance frameworks, access controls, and security measures is crucial but adds to operational overhead and software costs. Compliance with regulations (e.g., GDPR, HIPAA) can further drive up these costs.	Implement a risk-based security framework and leverage automated data governance tools to streamline tasks, ensure regulatory compliance, and mitigate data breaches.
Cloud vs. On-Premise Data Storage	Cloud-based storage offers scalability and pay-as-you-go models, but costs can escalate with large data volumes and frequent access.	Employ a hybrid or multi-cloud strategy to balance cost, performance, and control. Use cloud storage for scalability and on-premise solutions for predictable, high-volume workloads or sensitive data.

2. Computational Infrastructure

Cost Driver	Tradeoffs	Recommended Action
AI Model Complexity	More complex models (e.g., large language models, deep neural networks) require significantly more computational power and memory, often demanding expensive GPUs or specialized AI accelerators.	Match model complexity to the mission need. Avoid overengineering by selecting the smallest model that meets performance requirements. Consider fine-tuning smaller pre-trained models instead of building large models from scratch.
Training Data Size	Larger datasets extend training duration and compute needs, even for relatively simple models, impacting cost and delivery timelines.	Balance data volume with quality and marginal value. Use sampling, data deduplication, and feature selection to optimize dataset size while maintaining model accuracy.
Inference Frequency & Latency Reqs	Real-time or high-frequency inference demands always-on, high-performance infrastructure, driving up costs. This demand is often sustained and higher for Agentic AI systems.	Classify workloads by latency tolerance. Use edge deployment or asynchronous processing where possible, reserving real-time compute for genuinely time-sensitive tasks.
Centralized vs. Distributed Models	Centralized models may reduce infrastructure duplication, while distributed setups require shared platforms to manage costs. Geographic dispersion adds complexity.	Align infrastructure strategy with organizational structure. Evaluate the cost implications of centralized vs. federated deployment models and optimize for shared infrastructure in multi-office or hybrid environments.
Cloud-Based Computing & FinOps	Leveraging on-demand processing power from cloud providers allows flexible scaling, but sustained high usage can lead to substantial recurring costs.	Implement AI-aware FinOps practices. Tag compute resources by project, automate instance right-sizing, and leverage spot/preemptible instances where tolerance allows. Plan for sustained workloads using reserved capacity or savings plans.
Simulation Environments (Agentic AI)	Sophisticated simulation environments required for training and evaluating autonomous agents are resource-intensive to build and operate.	Scope simulation fidelity to mission need. Use tiered testing environments—lightweight for iteration, high-fidelity for final evaluation. Leverage shared simulation platforms to amortize setup costs.

3. AI Model Development and Customization

Cost Driver	Tradeoffs	Recommended Action
Talent Acquisition & Retention	The scarcity of skilled data scientists, machine learning engineers, and AI governance experts drives up salaries and recruitment costs. For Agentic AI, more specialized skills are required, leading to higher costs due to talent scarcity.	Prioritize strategic workforce planning. Focus on hiring for core capabilities while leveraging academic partnerships or outsourcing specialized tasks. Upskill internal teams through structured training programs.
Open Source vs. Proprietary Tools	Using open-source AI frameworks reduces software licensing costs but may require more internal expertise for implementation and support. Proprietary tools often have higher fees but offer more comprehensive features and vendor support.	Evaluate tooling strategy based on internal capability and lifecycle needs. Use open-source for innovation and experimentation and proprietary tools for enterprise-scale reliability. Hybrid models are often the most cost-effective.
Model Training & Iteration	The iterative nature of model development, with multiple training runs and hyperparameter tuning, consumes significant resources. For Agentic AI, iterative design and extensive testing in simulated and real-world environments increase development timelines and costs.	Optimize training workflows and automate experimentation. Use tools like AutoML and MLOps platforms to reduce manual effort. Prioritize early-stage feasibility validation before large-scale model investment.
Custom Model Development vs. Off-the-Shelf Solutions	Developing custom models is significantly more expensive and time consuming than adapting or fine tuning existing pre-trained models or utilizing off-the-shelf AI services.	Choose the lowest-effort solution that meets performance needs. Start with pre-trained models or third-party APIs when feasible. Consider custom development only when mission requirements cannot be met otherwise.



4. Integration and Deployment

Cost Driver	Tradeoffs	Recommended Action
Complexity of Existing Systems	Integrating AI solutions with legacy systems can be challenging and costly, requiring custom APIs, data transformations, and extensive testing. For Agentic AI, a special approach is needed to enable agents to coordinate with complex legacy systems.	Conduct a thorough systems readiness assessment. Adopt Model Context Protocols to allow agents to dynamically interpret and adapt to existing components through standardized, context-aware interactions.
Deployment Environment	The choice of deployment environment (e.g., on-premise, cloud, edge devices) impacts infrastructure costs, network latency, and security. Agentic AI deployment into real-world or dynamic virtual environments presents challenges related to connectivity, security, and continuous operation, requiring robust infrastructure.	Select the deployment environment based on operational and compliance requirements. Use hybrid or multi-cloud strategies where flexibility is needed. Leverage containerization and infrastructure-as-code to ensure portability and consistency.
Testing & Validation	Thorough testing is crucial for successful deployment but adds to project timelines and costs. For Agentic AI, ensuring safe, reliable, and robust operation is paramount, involving rigorous and expensive testing, verification, and validation.	Embed testing throughout the development lifecycle. Adopt CI/CD with integrated testing frameworks, simulation-based validation, and structured user feedback loops. Allocate resources for both pre-deployment and post-deployment testing.

5. Monitoring, Maintenance, and Updates

Cost Driver	Tradeoffs	Recommended Action
Ongoing Model Performance Monitoring	Continuous monitoring is critical for detecting performance degradation, model drift, or anomalies. This requires investment in observability infrastructure and skilled personnel. For Agentic AI, this is particularly crucial as it requires constant monitoring to prevent unintended actions.	Establish real-time monitoring linked to operational and financial KPIs. Track accuracy, cost-per-inference, latency, and error rates. For agentic systems, deploy multi-layered anomaly detection to identify and intervene when agents deviate from expected norms.

5. Monitoring, Maintenance, and Updates

Cost Driver	Tradeoffs	Recommended Action
Data Refresh & Re-training	Models degrade over time if not re-trained with updated data. Regular refresh cycles incur costs related to data processing, labeling, and compute. Agentic AI learns from its actions, necessitating robust data collection for feedback loops, which adds to the volume and complexity of data to be managed.	Schedule re-training based on model performance triggers, not arbitrary timelines. Use active learning, semi-automated labeling, and MLOps pipelines to reduce friction. For agentic systems, design for auditability and explainability from the start by capturing decisions and inputs in structured logs.
Software Licenses & Support	Licenses, vendor support plans, and patching requirements contribute to recurring costs. Long-term commitments can reduce flexibility and lock in costs.	Conduct regular license usage audits and support contract reviews. Ensure license tiers and support levels match the current maturity and criticality of the AI system.
Troubleshooting & Bug Fixes	Resolving issues, debugging models, and addressing performance bottlenecks contribute to maintenance costs. For Agentic AI, troubleshooting may involve intricate behavioral analysis to understand why an agent made a specific decision.	Implement structured logging and observability for debugging. Adopt tiered support and post-mortem analysis practices to identify root causes and reduce recurrence.

6. Safety, Governance, and Responsible AI

Cost Driver	Tradeoffs	Recommended Action
Development of Safety Protocols & Ethical Guardrails	Designing and implementing robust safety mechanisms and ethical guidelines directly into agentic systems is complex and adds significant development and testing costs.	Build safety and ethics into the system architecture from the outset. Use alignment techniques, maintain testbeds for edge-case evaluation, and engage ethicists during design.
Legal & Compliance Costs	Navigating the evolving legal and regulatory landscape for autonomous systems, including liability concerns, can incur substantial legal and compliance expenses.	Engage legal, policy, and compliance experts early. Map regulatory requirements to system capabilities. Document decision logic, auditability measures, and fallback mechanisms.

6. Safety, Governance, and Responsible AI

Cost Driver	Tradeoffs	Recommended Action
Human-Agent Collaboration Interfaces	Designing intuitive interfaces for humans to monitor, interact with, and potentially override agentic systems is critical.	Prioritize a human-in-the-loop design. Invest in clear status dashboards, override capabilities, and transparent feedback loops. Test interfaces with real users in operational settings.
Specialized Hardware for Embodied Agents	If the Agentic AI is embodied (e.g., robotics), there are significant hardware costs for the physical robots, sensors, actuators, and their maintenance.	Assess embodiment needs early in the design phase. Where feasible, simulate extensively before hardware deployment. Partner with hardware providers that offer modular, serviceable systems.
Public Trust & Acceptance Initiatives	Assess embodiment needs early in the design phase. Where feasible, simulate extensively before hardware deployment. Partner with hardware providers that offer modular, serviceable systems.	Proactively build public engagement plans. Use pilots, open reporting, and community involvement to demystify capabilities and highlight safeguards.



Cost Ranges Based on Solution Type and Complexity

To effectively plan and budget for artificial intelligence initiatives, a clear understanding of the potential costs is essential. The following table breaks down the estimated cost ranges for **AI solution development** based on two key factors: **solution type** and **complexity**. It covers a spectrum of modern AI, including **Traditional AI**, **Generative AI (GenAI)**, and **Agentic AI**. For each category, the table provides descriptions for low, medium, and high complexity projects to help frame the scope of work and associated financial investment. These ballpark figures are designed to guide organizations in making strategic decisions and allocating resources appropriately for their AI goals.

Solution Type & Complexity	Estimated Annual Cost Range (USD)	Description
Traditional AI - Low Complexity	\$10,000 - \$150,000	Simple, rules-based systems, basic ML models (e.g., regression, classification on small datasets). The cost reflects a small team and minimal compute/data infrastructure needs.
Traditional AI - Medium Complexity	\$150,000 - \$500,000	More sophisticated ML models (e.g., support vector machines, random forests), moderate data volumes, some data integration, and some cloud computing resources.
Traditional AI - High Complexity	\$500,000 - \$2M	Deep learning models, large and complex datasets, real-time processing, extensive data pre-processing, custom model development. These costs are driven by a larger team, significant GPU usage, and extensive data preparation.
Generative AI (GenAI) - Low Complexity	\$50,000 - \$250,000	Fine-tuning pre-trained GenAI models using existing APIs and smaller, proprietary datasets. The cost is primarily for talent and API usage.
Generative AI (GenAI) - Medium Complexity	\$250,000 - \$1M+	Training a custom GenAI model or heavily modifying a pre-trained one. The high-end of this range is driven by substantial compute requirements and a large, specialized team.
Generative AI (GenAI) - High Complexity	\$1M - \$200M+	Developing a novel GenAI architecture or training a foundational model from scratch. The cost is dominated by massive-scale compute (e.g., thousands of GPUs for months) and top-tier research talent.
Agentic AI - Low Complexity	\$20,000 - \$100,000	Simple autonomous agents for constrained environments with clear rules (e.g., basic task automation within a defined digital system). Limited human oversight. Costs are low due to limited scope and predefined rules.
Agentic AI - Medium Complexity	\$100,000 - \$500,000+	Autonomous agents in semi-structured environments requiring some degree of planning, learning, and adaptation. May involve human-in-the-loop oversight. The cost increases with the need for more complex reasoning and integration with multiple systems.
Agentic AI - High Complexity	\$500,000 - \$1.5M+	Highly autonomous agents in complex, dynamic, and potentially unpredictable real-world environments. The cost is high due to the need for advanced reasoning, continuous learning, robust safety mechanisms, and extensive testing.

It's important to note that the above figures are broad estimates and can be influenced by several factors:

- **Talent:** The salaries of specialized AI engineers and data scientists are a major cost driver, often accounting for **40-60%** of a project's budget.
- **Infrastructure:** The cost of cloud computing resources, especially high-end GPUs for training large models, can be **substantial**. For example, a single A100 GPU can cost over \$1 an hour to rent.
- **Data:** Acquiring, cleaning, and labeling large datasets can be a significant expense, ranging from a few thousand to over **\$100,000** depending on the domain.
- **Maintenance:** Ongoing costs for maintenance, monitoring, and model retraining can add an additional **15-30%** of the initial development cost annually.

The cost ranges provided are for the development and initial deployment of the AI solution. This typically covers the expenses incurred during the initial project lifecycle, which often spans the first year. It's important to note that these figures do not include ongoing operational costs, such as model maintenance, monitoring, and updates, which can add an additional 15-30% of the initial development cost annually.

Tools for Rough Estimates

To build rough estimates for AI solution costs, agencies can leverage a variety of tools, ranging from readily available cloud provider calculators to more specialized, internal frameworks.

Cloud Provider Pricing Calculators

These tools offer detailed cost breakdowns for services, allowing for preliminary budget planning based on infrastructure and service usage.

- **[Amazon Web Services \(AWS\) Pricing Calculator](#):** Offers detailed cost breakdowns for various AWS services, including compute (EC2), storage (S3), and machine learning services (SageMaker). This can help estimate infrastructure and some AI service costs.
- **[Microsoft Azure Pricing Calculator](#):** Similar to AWS, Azure's calculator allows users to estimate costs for Azure Virtual Machines, Storage, Azure Machine Learning, and other related services.
- **[Google Cloud Pricing Calculator](#):** Provides cost estimations for Google Cloud's compute, storage, and AI/ML services (e.g., Vertex AI).

Specialized Cost Estimation Approaches

For more customized or complex AI solutions, organizations can go beyond standard calculators to create more tailored estimates.

- **Internal Cost Estimation Models:** Agencies can develop custom spreadsheets or tools that factor in labor costs, software licenses, and hardware depreciation based on historical project data.
- **AI Expert-Provided Frameworks:** Engaging external specialists in digital transformation and AI deployment can provide access to established frameworks and tools for optimizing AI solutions and estimating costs.



Conclusion

AI offers federal agencies enormous potential to improve citizen services, operational efficiency, decision-making, and mission effectiveness. However, this potential comes with the responsibility of pursuing AI in a **cost-conscious**, **transparent**, and **accountable** manner. Successful AI adoption is not merely a technical endeavor, but fundamentally a planning and management challenge. In the scramble to take advantage of AI benefits, agencies must not lose sight of the need for adequate strategizing and planning to guard against the very real possibility of runaway costs.

The key to making AI projects compelling and palatable to oversight bodies lies in demonstrating that prudent planning and a mature financial operations framework like FinOps can control costs, ensuring benefits outweigh expenses. The Army's experience of carefully reviewing use cases to prevent excess costs and moving towards more cost-effective COTS platforms is a prime example of this strategic mindset. Federal technology leadership has also highlighted agencies' efforts to share code and leverage existing platforms to **avoid needless costs** and **accelerate AI maturity**. Concurrently, they emphasize risk management and safeguards to build public trust and prevent costly failures.

By embracing these practices, federal agencies can develop powerful, sustainable AI solutions, even advancing into cutting-edge Agentic AI for complex missions, without risking uncontrolled costs. Success depends on a clear understanding of **data needs**, **identification of key cost drivers**, **intentional design decisions**, and **the use of reliable estimation tools** to guide planning and execution. Engaging experienced digital transformation experts who have successfully delivered AI at scale can further accelerate impact while mitigating risk. REI Systems offers proven strategies and hands-on frameworks for FinOps and AI cost management tailored to the federal landscape. Finally, by rigorously documenting and sharing outcomes and cost frameworks, each agency contributes to a collective foundation of insight—advancing the responsible, cost-effective adoption of AI across government.



Ramakrishnan (Ramki) Krishnamurthy

Data Analytics Lead, REI Systems

Mr. Krishnamurthy is an accomplished technology leader with 25+ years of experience defining and executing transformative data strategies for government organizations, including Fannie Mae, HRSA, and FEMA.

Contact him at: rkrishnamurthy@reisystems.com

References

1. <https://www.cio.gov/ai-in-action/>
2. <https://blog.dol.gov/2024/12/20/boosting-efficiency-with-ai-real-world-success-stories>