FORRESTER®

The Total Economic Impact™ Of Red Hat Hybrid Cloud Platform For MLOps

Cost Savings And Business Benefits Of Red Hat OpenShift For Operationalizing AI/ML In Financial Services

MARCH 2022

Table Of Contents

Consulting Team: Sam Conway
Benjamin Corey

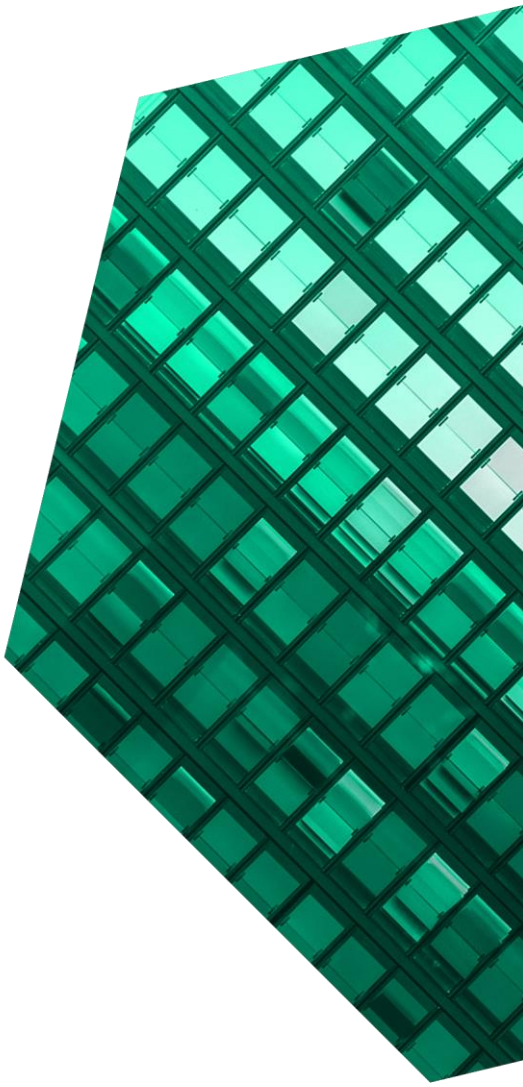

ABOUT FORRESTER CONSULTING

Forrester Consulting provides independent and objective research-based consulting to help leaders succeed in their organizations. For more information, visit forrester.com/consulting.

Executive Summary

Data scientists can create a new model in a matter of weeks. However, integrating the model into the business and recognizing value can be an arduous and often unaccomplished journey. MLOps helps teams provision infrastructure, stage models, manage dependencies, orchestrate model calling, and serve the model in a scalable fashion to accelerate the time-to-value of AI/ML applications for organizations.

In today's modern enterprise, artificial intelligence (AI) and machine learning (ML) have become critical business tools for industry leaders. However, organizations frequently experience challenges as AI/ML use is scaled from tactical or experimental use cases to enterprise-wide adoption. Machine learning operations (MLOps) are a necessity when operationalizing AI/ML models. MLOps comprises tools, technologies, and practices to enable organizations to deploy, monitor, and govern AI/ML models in production applications.[1]

Red Hat provides a secure, enterprise open hybrid cloud platform and tooling to enable self-service for data scientists and developers to integrate, streamline, automate, and simplify the MLOps process. With Red Hat OpenShift as the foundation (also available as a managed cloud service across major cloud providers), organizations can bring intelligent applications to market faster, better control their MLOps software stack, and simplify AI/ML lifecycle integration and deployment.[2] Organizations can also leverage Red Hat's extensive partner ecosystem and consulting services.[3]

Red Hat commissioned Forrester Consulting to conduct a Total Economic Impact™ (TEI) study and examine the potential return on investment (ROI) financial services enterprises may realize by deploying Red Hat as a foundation for accelerating MLOps.[4] The purpose of this study is to provide financial industry leaders with a framework to evaluate the potential financial impact of enabling MLOps on Red Hat for their organizations.

KEY STATISTICS

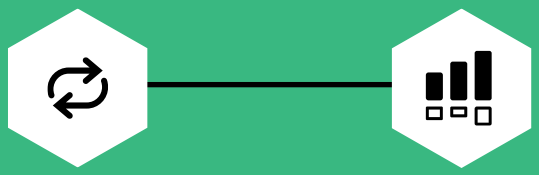

Return on investment (ROI)
210%

Payback period
13 months

To better understand the benefits, costs, and risks associated with this investment, Forrester interviewed three decision-makers with experience accelerating MLOps with Red Hat and speeding up delivery of AI-powered financial services applications. For the purposes of this study, Forrester aggregated the interviewees' experiences and combined the results into a single composite organization.

Prior to deploying Red Hat as part of their MLOps process, the interviewees noted how their organizations lacked a consolidated platform — using siloed and often redundant tools that did not scale to their enterprise needs. Prior solutions yielded limited success, leading to inefficiencies that held back their organizations from fully recognizing the benefits of AI/ML efforts.

After the investment in Red Hat, the interviewees deployed an enterprise-class cloud-native platform for MLOps, reducing manual operations efforts, accelerating the production of AI/ML models, and reducing support costs.

Improvement in data scientist efficiency:

20%

KEY FINDINGS

Quantified benefits. Risk-adjusted present value (PV) quantified benefits include:

- **Accelerated model deployment and increased revenue with new and improved products or services.** Deploying Red Hat to support MLOps activities enabled organizations to operationalize and streamline putting models into production applications faster. With a scalable system for model lifecycle governance, organizations improved the performance of activities such as loan underwriting and targeted marketing.
- **Improved data scientist efficiency, enabling time savings of up to 20%.** Deploying Red Hat provided data scientists with consistent tooling, practices, and a self-service environment. With the right tools in place, data scientists spent more time working and less time waiting on infrastructure provisioning.
- **Delivered software developer time savings of up to 60%.** With Red Hat OpenShift, organizations could create and deploy applications using continuous integration and delivery (CI/CD) tools. Development teams were able to complete and deploy models without the need for administrative teams to configure and provision infrastructure.

> **"[Our Red Hat platform] is faster, [is] more secure, decreases time-to-market, and is stable."**
>
> *AI architecture lead, financial services*

- **Provided infrastructure operations time savings of up to 60%.** Having self-service tools for data scientists reduced the burden on infrastructure teams to configure and deploy individual workstations. Architecture teams also spent less time on redundant work with a standardized platform and pipeline.
- **Enabled infrastructure savings of up to 30%.** Organizations deployed MLOps on the Red Hat technology using existing hardware. Organizations could run more applications concurrently and handle scheduling with Red Hat OpenShift. avoiding investment in additional infrastructure to support advanced workloads.

> **"The Red Hat team have always been a pleasure to work with. Our experience overall with Red Hat has been tremendously positive. They are a really good partner, regardless of your deployment size, and are very eager to work with you."**
>
> *Principal consultant of cloud engineering, professional services*

Unquantified benefits. Benefits that are not quantified for this study include:

- **Improved data scientist retention and acquisition.** Removing the need for teams to manage and configure hardware improved employee morale, while providing cutting-edge tools to data scientists increased the attractiveness of organizations to potential hires.
- **Democratization of data across the organization.** A dedicated MLOps platform enables a larger percent of the organization to work with data and make data-informed decisions.
- **Increased competitiveness.** Applying data to business problems improves the competitive position of organizations.

Costs. Risk-adjusted PV costs include:

- **Implementation and training.** Interviewees' organizations dedicated internal resources to the planning and deployment of Red Hat technology for their MLOps initiatives. Additionally, organizations engaged professional services to aid in planning and deployment.
- **Red Hat fees.** Organizations pay subscription fees for the use of the Red Hat OpenShift foundation technology.
- **Ongoing management.** Organizations dedicate internal resources to the ongoing management of their MLOps process and platform.

The decision-maker interviews and financial analysis found that a composite organization experiences benefits of $11.4 million over three years versus costs of $3.7 million, adding up to a net present value (NPV) of $7.7 million and an ROI of 210%.

ROI
210%

BENEFITS PV
$11.4 million

NPV
$7.7 million

PAYBACK
13 months

Benefits (Three-Year)

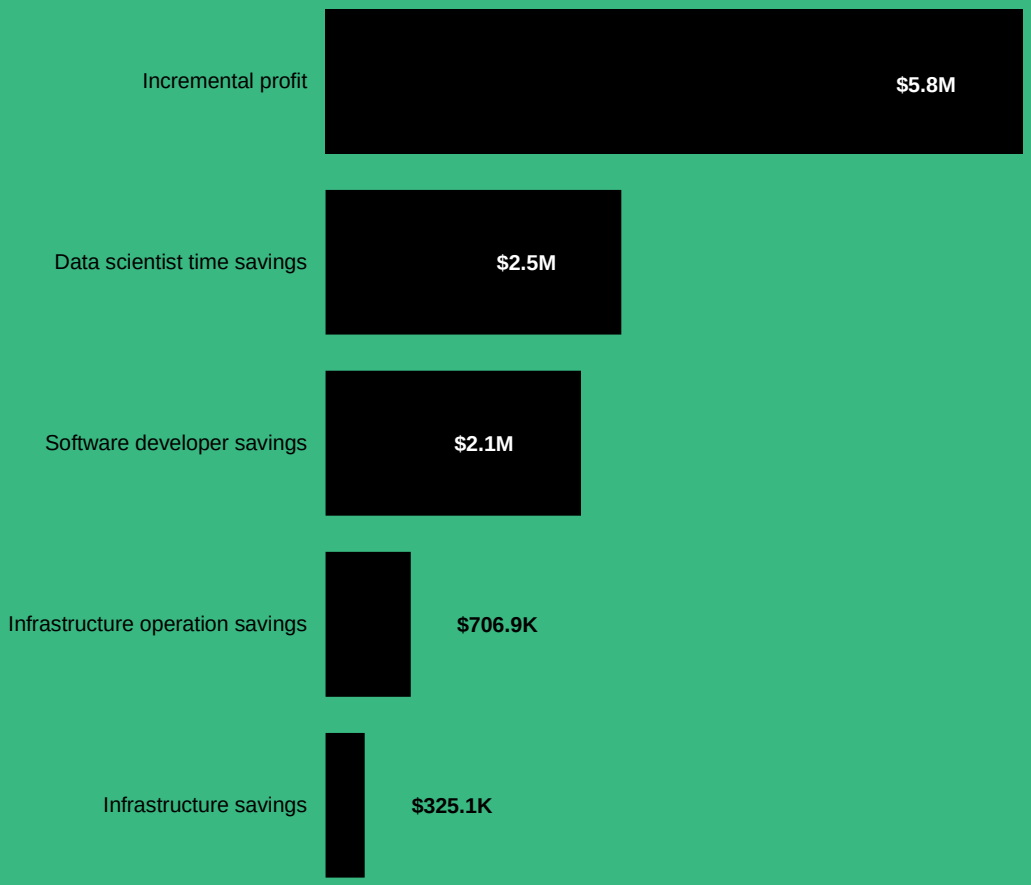

TEI FRAMEWORK AND METHODOLOGY

From the information provided in the interviews, Forrester constructed a Total Economic Impact™ framework for those organizations considering an investment in MLOps with Red Hat.

The objective of the framework is to identify the cost, benefit, flexibility, and risk factors that affect the investment decision. Forrester took a multistep approach to evaluate the impact that the MLOps solutions can have on an organization.

DUE DILIGENCE

Interviewed Red Hat stakeholders and Forrester analysts to gather data relative to MLOps capabilities from Red Hat.

DECISION-MAKER INTERVIEWS

Interviewed three decision-makers at organizations using Red Hat technologies as part of their MLOps process to obtain data with respect to costs, benefits, and risks.

COMPOSITE ORGANIZATION

Designed a composite organization based on characteristics of the interviewees' organizations.

FINANCIAL MODEL FRAMEWORK

Constructed a financial model representative of the interviews using the TEI methodology and risk-adjusted the financial model based on issues and concerns of the decision-makers.

CASE STUDY

Employed four fundamental elements of TEI in modeling the investment impact: benefits, costs, flexibility, and risks. Given the increasing sophistication of ROI analyses related to IT investments, Forrester's TEI methodology provides a complete picture of the total economic impact of purchase decisions. Please see Appendix A for additional information on the TEI methodology.

The Red Hat OpenShift For MLOps Customer Journey

Drivers leading to the Red Hat investment

Interviewed Decision-Makers			
Interviewee	**Industry**	**Region**	**Revenue**
AI architecture lead	Financial services	Headquartered in Asia	$3 billion
Principal consultant of cloud engineering	Professional services	Headquartered in Europe	$29 billion
Chief digital transformation officer	Financial services	Headquartered in the United States	$3 billion

KEY CHALLENGES

Forrester interviewed decision-makers at three organizations using the Red Hat OpenShift platform for MLOps. Prior to OpenShift, the organizations all lacked consolidated platforms and relied on disparate and siloed tools. The organizations struggled with common challenges, including:

- **Inability to fully utilize organizational data.** Interviewees acknowledged that their firms possessed a wealth of data, but without the right tools in place, the organizations could not operationalize it. The chief digital transformation officer for a financial services firm explained: "We're a bit behind digitally, but we are extremely data-driven as an organization. We have a ton of models, and there has been an initiative to enhance those models, making them more knowledgeable, more sophisticated in decision-making by cycling through our datasets from the past 25 to 30 years. We've always had the data — that has not been the issue. Historically, we've lacked in terms of having the tools both in hardware and software."
- **Need for a platform to support scalable and repeatable processes.** Prior to investing in their Red Hat solution, the interviewed organizations relied on disparate tools and ad hoc solutions. Not having a consistent platform meant that organizations lacked visibility across their enterprise, with many data science teams doing redundant work, making duplicate purchases, and causing headaches for infrastructure professionals. Moreover, infrastructure teams could not keep up with provisioning demands, slowing the work of data science teams.

The AI architecture lead for a financial services firm explained: "Before we invested in OpenShift and had a larger AI initiative, teams were using standalone images and running them independently without any Kubernetes or orchestration."

The principal consultant of cloud engineering stated: "One of the challenges we faced was disconnected processes. There was not a consistent way to have a streamlined, repeatable way to build, train, deploy, and consume models."

The chief digital transformation officer detailed: "We have multiple businesses that were acquired at different points, and we have been in the process of consolidating that. We have redundancy across the entities, and the idea as we embark on this initiative is to determine what is best-in-breed and retain that champion."

“When we selected Red Hat for the organization’s continuous platform, one of the main reasons was security. OpenShift has security features that are more advanced than in other Kubernetes platforms.”

AI architecture lead, financial services

- **Painstaking compliance processes and stringent security criteria.** Interviewees’ organizations were either in the financial services industry or serving clients within it. As such, their firms were bound by tight industry and government regulations, which made it difficult to adopt new technology, and compliance reviews drastically slowed development. The AI architecture lead explained: “Security is a major concern for us in this regulatory environment. Not all Kubernetes platforms are able to meet this criterion.”

“In this hybrid multicloud world, having the ability to have consistencies around how you are deploying on-prem into your cloud providers is critical, and it’s hard to achieve it without tools like OpenShift.”

Principal consultant of cloud engineering, professional services

COMPOSITE ORGANIZATION

Based on the interviews, Forrester constructed a TEI framework, a composite company, and an ROI analysis that illustrates the areas financially affected. The composite organization is representative of the three decision-makers that Forrester interviewed and is used to present the aggregate financial analysis in the next section. The composite organization has the following characteristics:

Description of composite. The composite organization is a data-driven financial services firm, with $5 billion in annual revenue from global operations. The organization develops ML models to develop and improve products and services. Over the three-year model, the organization expands usage of the platform as more product and service lines incorporate models into their work.

Composite Userbase

User type	Year 1	Year 2	Year 3
Data scientists	50	75	100
Developers	17	25	33
Infrastructure ops	8	13	17

Deployment characteristics. The organization has an MLOps platform built on Red Hat OpenShift, with 250 dual-core subscriptions.

Key assumptions

- **$5 billion annual revenue**
- **3:1 data scientist to developer ratio**
- **6:1 data scientist to admin ratio**

Analysis Of Benefits

Quantified benefit data as applied to the composite

Total Benefits						
Ref.	**Benefit**	**Year 1**	**Year 2**	**Year 3**	**Total**	**Present Value**
Atr	Incremental profit	$600,000	$1,800,000	$5,000,000	$7,400,000	$5,789,632
Btr	Data scientist time savings	$425,000	$956,250	$1,700,000	$3,081,250	$2,453,888
Ctr	Software developer savings	$459,000	$843,750	$1,336,500	$2,639,250	$2,118,719
Dtr	Infrastructure operations savings	$150,000	$281,250	$450,000	$881,250	$706,893
Etr	Infrastructure savings	$67,500	$135,000	$202,500	$405,000	$325,075
	Total benefits (risk-adjusted)	$1,701,500	$4,016,250	$8,689,000	$14,406,750	$11,394,207

INCREMENTAL PROFIT

Evidence and data. Interviewees reported that having consistent tooling improved the speed and quality of data scientists' work. By putting more models into production, organizations were capable of creating and or improving products and services. Accelerating the time-to-value for models helped organizations recognize incremental revenue. The interviewees' organizations used the Red Hat OpenShift software infrastructure to produce and refine models for a plethora of use cases.

- The chief digital transformation officer explained: "Whether you look at our credit models, pricing models, decision engines, [or] collection models — these all today incorporate machine learning of some sort, a variety of algorithms. We also use machine learning for segmentation and pricing optimization. The frequency of [our team] touching and executing changes to models that are passing through model risk management has certainly increased."
- The AI architecture lead stated: "We have models for marketing, churn models, credit card churn models, and ATM cash optimization models. We have reduced the time-to-market for use cases by one to two months."
- The principal consultant of cloud engineering detailed: "We brought model production time down by a couple of weeks, to probably a week or less. A lot of it is because there is now more consistency and tooling, and MLOps is now more fully integrated."

Modeling and assumptions. In modeling incremental profit, Forrester assumes:

- The composite organization has annual revenue of $5 billion. Of this revenue, 10% is impacted by AI/ML models in Year 1. This impact grows to 25% by Year 3 as more product teams incorporate AI/ML models.
- The revenue increase attributed to Red Hat is 1.5% to 5% over three years, primarily driven by an accelerated time-to-market.

- An average operating margin of 10% is applied, based on NYU Stern School of Business benchmarking for financial services organizations.[5]

Risks. Forrester recognizes that incremental profit may vary from organization to organization. Specific risk considerations include:

- Geography, vertical, and competitive landscape.
- The number of direct revenue-impacting models developed annually.
- Availability of data to train models.
- Internal talent.

Results. To account for these risks, Forrester adjusted this benefit downward by 20%, yielding a three-year, risk-adjusted total PV (discounted at 10%) of $5.8 million.

Incremental Profit

Ref.	Metric	Source	Year 1	Year 2	Year 3
A1	Annual revenue	Composite	$5,000,000,000	$5,000,000,000	$5,000,000,000
A2	Percentage of products/services impacted by Red Hat	Composite	10%	15%	25%
A3	Revenue increase attributed to Red Hat	Assumption	1.5%	3.0%	5.0%
A4	Increased revenue due to Red Hat MLOps	A1*A2*A3	$7,500,000	$22,500,000	$62,500,000
A5	Operating margin	NYU Stern School of Business	10%	10%	10%
At	Incremental profit	A4*A5	$750,000	$2,250,000	$6,250,000
	Risk adjustment	↓20%			
Atr	Incremental profit (risk-adjusted)		$600,000	$1,800,000	$5,000,000
	Three-year total: $7,400,000			**Three-year present value: $5,789,632**	

DATA SCIENTIST TIME SAVINGS

Evidence and data. Interviewees highlighted that with Red Hat OpenShift as the foundation for MLOps, their organizations could operationalize and streamline the work of data scientists. Organizations created self-service tools so data scientists could work on models instead of waiting for infrastructure provisioning. The Red Hat infrastructure also provided a catalyst for streamlining resource management, facilitating visibility and collaboration across teams and eliminating time wasted on redundant tasks.

- The principal consultant of cloud engineering stated: “[Red Hat] brings more consistencies into your processes, and those processes are defined more holistically across the organization on what good looks like, what it does not look like, how we should be building these models, how we should be deploying them, what types of reviews do we want to look like, how do we get all those nuts and bolts around biases, [and] explainability instead of having that very kind of decentralized and inconsistent process for what people are doing. We are able to standardize on that and are

able to get more governance and rigor around that process."

- The AI architecture lead detailed: "Before, when new data scientists were added to our team, they had to wait for notebooks to be deployed on bare metal servers for them. Now, with Kubernetes, we can dynamically add more notebooks. The data science team is no longer dependent on any other team to request new infrastructure."
- The chief digital transformation officer explained: "It's a little easier to view and maneuver within the tool. Traditionally, if the model risk management team took two months to approve and validate a model or change a model, in the new way of doing business, you can potentially chop that in half."

Modeling and assumptions. In modeling data scientist time savings, Forrester assumes:

- The organization's initial deployment is used by 50 data scientists, growing to 100 by Year 3 as more products, services, and functions incorporate AI/ML models.
- Data scientists have an average fully burdened annual salary of $200,000.
- Of time saved, 50% is reallocated to value-generating activities.

> **"The time we save helps us spend more time following the latest trends in AI technologies."**
>
> *AI architecture lead, financial services*

Risks. Forrester recognizes that data scientist time savings may vary from organization to organization. Specific risk considerations include:

- Rate of MLOps adoption.
- Geography, vertical, and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 15%, yielding a three-year, risk-adjusted total PV of $2.5 million.

Data Scientist Time Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
B1	Number of data scientists	Interviews	50	75	100
B2	Time savings with Red Hat platform	Assumption	10%	15%	20%
B3	Data scientist fully burdened annual salary	TEI standard	$200,000	$200,000	$200,000
B4	Productivity recapture	Assumption	50%	50%	50%
Bt	Data scientist time savings	B1*B2*B3	$500,000	$1,125,000	$2,000,000
	Risk adjustment	↓15%			
Btr	Data scientist time savings (risk-adjusted)		$425,000	$956,250	$1,700,000
	Three-year total: $3,081,250			**Three-year present value: $2,453,888**	

SOFTWARE DEVELOPER SAVINGS

Evidence and data. Interviewees found that creating a streamlined pipeline for integrating models into applications greatly improved the efficiency of their developer teams. Prior to investment in Red Hat OpenShift as the foundation for MLOps, interviewees struggled with inefficiencies when integrating models into applications. Development teams often waited for multiple infrastructure teams to provision and refine the required resources.

- The AI architecture lead explained: “We deploy our models as APIs. We can scale horizontally or vertically in OpenShift to handle the increasing load. However, in the legacy platform, we would have to add new servers or memory as we needed it. In OpenShift, it is just a configuration. We’ve seen a big speedup with this, and our developers don’t have to waste their time waiting.”

  The AI architecture lead added: “There is much shorter time spent in provisioning [and] prepping deployments and configurations. Before, you would have to wait for a new starter setup provisioning for our application and then add service to our CI/CD pipeline, prepare deployment scripts — this took time and added one to two months in deployment.”
- The chief digital transformation officer stated: “If you look at the benefit for us, whether it’s time-to-get-started or time-to-market end-to-end, you’ll just see a tremendous lift there in terms of how much we can push through in a condensed time frame.”
- The principal consultant of cloud engineering detailed, “The way that the MLOps process is working and how we are able to go into those different stages, I’m saving a few months on a productionalized deployment.”

Modeling and assumptions. In modeling developer time savings, Forrester assumes:

- A 3:1 support ratio of data scientists to developers in AI/ML efforts.
- Developers have an average fully burdened annual salary of $150,000.
- Of time saved, 50% is reallocated to value-generating activities.

Risks. Forrester recognizes that developer time savings may vary from organization to organization. Specific risk considerations include.

- Prominence of AI/ML models in business applications.
- Rate of MLOps practice adoptions.
- Geography, vertical, and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of $2.1 million.

Year 3 developer time savings:

60%

Software Developer Savings					
Ref.	**Metric**	**Source**	**Year 1**	**Year 2**	**Year 3**
C1	Number of developers supporting model integration	B1/3	17	25	33
C2	Time savings with Red Hat platform	Interviews	40%	50%	60%
C3	Developer fully burdened annual salary	TEI standard	$150,000	$150,000	$150,000
C4	Productivity recapture	Assumption	50%	50%	50%
Ct	Software developer savings	C1*C2*C3*C4	$510,000	$937,500	$1,485,000
	Risk adjustment	↓10%			
Ctr	Software developer savings (risk-adjusted)		$459,000	$843,750	$1,336,500
	Three-year total: $2,639,250			**Three-year present value: $2,118,719**	

INFRASTRUCTURE OPERATIONS SAVINGS

Evidence and data. Interviewees highlighted that deploying Red Hat OpenShift and leveraging its capabilities to operationalize MLOps environment greatly reduced the burden on infrastructure operations teams serving data scientists and developers working in AI/ML practices. Self-service tools for data scientists and automations for deployment and scaling reduced operational overhead in activities that previously required effort from multiple admin teams.

Year 3 infrastructure operations savings:

60%

- The AI architecture lead stated: “Before, we had to deal with more than five different teams — including storage admins and database admins — and many other roles that had to be included in provisioning. Currently one of them is in the process because we can do everything with configurations on OpenShift. Someone can just put up a YAML file, and we state how many CPUs, [state] how much memory, and define how much storage that we need to allow. We don’t need to work with any admins.”
- The chief digital transformation officer detailed, “The self-service removes a lot of the frustrations related to making requests and having the IT teams create and provision users and identify the appropriate access and all that.”

Modeling and assumptions. In modeling infrastructure operations savings, Forrester assumes:

- A 6:1 support ratio of data scientists to admins in AI/ML efforts.
- Admins have an average fully burdened annual salary of $100,000.

- Of time saved, 50% is reallocated to value-generating activities.

Risks. Forrester recognizes that infrastructure operations savings may vary from organization to organization. Specific risk considerations include.

- Prior-state solution and infrastructure requirements.
- Geography, vertical, and prevailing labor rates.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of nearly $707,000.

Infrastructure Operations Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
D1	Number of admins supporting data scientists	B1/6	8	13	17
D2	Time savings with Red Hat platform	Interviews	40%	50%	60%
D3	Admin fully burdened annual salary (blended)	TEI standard	$100,000	$100,000	$100,000
D4	Productivity recapture	Assumption	50%	50%	50%
Dt	Infrastructure operations savings	D1*D2*D3*D4	$166,667	$312,500	$500,000
	Risk adjustment	↓10%			
Dtr	Infrastructure operations savings (risk-adjusted)		$150,000	$281,250	$450,000
	Three-year total: $881,250			**Three-year present value: $706,893**	

INFRASTRUCTURE SAVINGS

Evidence and data. Interviewees reported that they were able to deploy Red Hat OpenShift on existing hardware, with no additional investment. By bringing data science and support teams into a single MLOps practice, organizations were able to eliminate spend on some legacy solutions — some teams were paying for the same solution but unaware of that when working in silos. Additionally, organizations improved resource utilization and could run more applications concurrently using Red Hat OpenShift for scheduling.

- The AI architecture lead explained: “We can run more applications concurrently compared to before. OpenShift handles scheduling of applications more precisely, so it can run more applications in parallel. We need less infrastructure to run the same workload.”
- The principal consultant of cloud engineering stated: “There were different processes that people were following. You have some folks doing things on the cloud, some on-prem — different ways that people were achieving the same objectives. Consolidating that to using [Red Hat OpenShift] helps us realize cost savings [in licenses, hardware, and cloud computing costs].”

Modeling and assumptions. Based on customer interviews, Forrester assumes that estimated spend with legacy solutions to run current workloads is valued at 10% to 30% over current Red Hat spend.

Risks. Forrester recognizes that infrastructure savings may vary from organization to organization. Specific risk considerations include:

- Legacy solutions.
- Size of deployment and infrastructure demands.

Results. To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year, risk-adjusted total PV of more than $325,000.

Infrastructure Savings

Ref.	Metric	Source	Year 1	Year 2	Year 3
E1	Current spend	Red Hat	$750,000	$750,000	$750,000
E2	Estimated additional infrastructure needed to run same number of workloads	Interviews	10%	20%	30%
Et	Infrastructure savings	E1*D2	$75,000	$150,000	$225,000
	Risk adjustment	↓10%			
Etr	Infrastructure savings (risk-adjusted)		$67,500	$135,000	$202,500
	Three-year total: $405,000			**Three-year present value: $325,075**	

UNQUANTIFIED BENEFITS

Additional benefits that customers experienced but were not able to quantify include:

- **Improved data scientist retention and acquisition.** Data scientists and engineers want to spend more time analyzing data, creating models, and solving business problems. Removing the need for teams to manage and configure hardware improved employee morale, while providing cutting-edge tools to data scientists increased the attractiveness of organizations to potential hires.

  The principal consultant of cloud engineering stated: “[Red Hat] has helped in job satisfaction. In our organization, people can be forced to wear multiple hats. Whether it’s a scientist becoming an engineer, or an engineer becoming a scientist — if you do not have a good way to do what they like and they are spending a lot of time on the nuts and bolts trying to enable something, you are going to have dissatisfaction. It’s harder to hire.”

  The AI architecture lead explained: “The [Red Hat OpenShift] platform has increased [employee] happiness. They were spending too much time with provisioning before. New employees are happy with it since they can do many things from a self-service platform. It helps us find new people to hire.”

- **Increased competitiveness.** The ability to utilize data and apply it to business problems improved the competitive position of organizations. AI/ML models could be applied to business processes and product or service ideation to better compete in challenging markets.

  The principal consultant of cloud engineering explained, “As we are able to generate insights quicker, we can be more competitive in the offerings we provide our clients.”

FLEXIBILITY

The value of flexibility is unique to each customer. There are multiple scenarios in which a customer might implement MLOps and later realize additional uses and business opportunities, including:

- **Expanding AI/ML model development to new teams and processes.** Interviewees noted that they had plans to expand model development to new teams and lines of business. For example, one financial services firm that had previously focused model usage on actions like loan origination planned to develop models to improve its customer support center.
- **Providing portability to allow teams to realize value in other solutions.** Interviewees noted that Red Hat's open source nature gave them the flexibility to move work between different tools and deployments. The chief digital transformation officer detailed: "The value proposition we find in Red Hat is [that] you could develop to recreate models. Once we have perfected a model, we can essentially lift and shift it to another geography. We're not restricted to working within one platform. Teams can experiment within that sandbox, and if they still want [to], they can also configure it to deploy straight into production within the Red Hat ecosystem. Or you could segregate it, say in a container, if you'd like to go that route, or you can deploy it on a private cloud or public cloud."

Flexibility would also be quantified when evaluated as part of a specific project (described in more detail in Appendix A).

Analysis Of Costs

Quantified cost data as applied to the composite

Total Costs							
Ref.	**Cost**	**Initial**	**Year 1**	**Year 2**	**Year 3**	**Total**	**Present Value**
Ftr	Implementation and training	$735,000	$0	$0	$0	$735,000	$735,000
Gtr	Red Hat subscriptions	$0	$787,500	$787,500	$787,500	$2,362,500	$1,958,396
Htr	Ongoing management	$0	$393,750	$393,750	$393,750	$1,181,250	$979,198
	Total costs (risk adjusted)	$735,000	$1,181,250	$1,181,250	$1,181,250	$4,278,750	$3,672,594

IMPLEMENTATION AND TRAINING

Evidence and data. Interviewees' organizations deploying Red Hat OpenShift for MLOps incurred internal labor costs for IT and development teams involved in planning usage, identifying dependencies, and allocating required infrastructure. Additionally, these firms engaged Red Hat consulting services to ensure a speedy and issue-free deployment.

Modeling and assumptions. In modeling implementation and training costs, Forrester assumes:

- The composite organization dedicates 20 internal resources to initial deployment. These resources include stakeholders from IT, data science, and applications teams. These resources spend 15% of their time during the initial deployment period on Red Hat OpenShift-related activities.
- The average fully burdened annual rate for internal resources involved in the initial deployment is $150,000.
- The organization engages Red Hat for professional services related to deployment and training.

Risks. Forrester recognizes that implementation and deployment costs may vary from organization to organization. Specific risk considerations include:

- Organizational size and internal bureaucracies.
- Internal skills.

Results. To account for these risks, Forrester adjusted this cost upward by 5%, yielding a three-year, risk-adjusted total PV (discounted at 10%) of $735,000.

Implementation And Training						
Ref.	**Metric**	**Source**	**Initial**	**Year 1**	**Year 2**	**Year 3**
F1	Internal resources dedicated to initial deployment	Composite	20			
F2	Percentage of time working on Red Hat integration and implementation	Interviews	15%			
F3	Average fully burdened annual rate for implementation FTE	TEI standard	$150,000			
F4	Professional services fees for implementation and training	Composite	$250,000			
Ft	Implementation and training	(F1*F2*F3)+F4	$700,000			
	Risk adjustment	↑5%				
Ftr	Implementation and training (risk-adjusted)		$735,000			
	Three-year total: $735,000			**Three-year present value: $735,000**		

RED HAT SUBSCRIPTIONS

Evidence and data. Organizations deploying MLOps on Red Hat OpenShift incur subscription costs based on the number of cores used. Red Hat OpenShift can be consumed on the cloud or deployed on-premises for more sensitive workloads.

Red Hat pricing varies based on a number of factors, and readers should connect with a Red Hat representative to get accurate pricing for specific solution needs.

Modeling and assumptions. In modeling Red Hat subscription costs, Forrester assumes that the composite organization has 250 dual-core subscriptions.

Risks. Forrester recognizes that subscription costs may vary from organization to organization. Specific risk considerations include the type and size of deployment.

Results. To account for these risks, Forrester adjusted this cost upward by 5%, yielding a three-year, risk-adjusted total PV of $2 million.

Red Hat Subscriptions						
Ref.	**Metric**	**Source**	**Initial**	**Year 1**	**Year 2**	**Year 3**
G1	Red Hat subscriptions	Composite		$750,000	$750,000	$750,000
Gt	Red Hat subscriptions	G1		$750,000	$750,000	$750,000
	Risk adjustment	↑5%				
Gtr	Red Hat subscriptions (risk-adjusted)		$0	$787,500	$787,500	$787,500
	Three-year total: $2,362,500			**Three-year present value: $1,958,396**		

ONGOING MANAGEMENT

Evidence and data. Interviewees' organizations dedicated limited internal resources to the management and administration of their MLOps platform.

Modeling and assumptions. In modeling ongoing management costs, Forrester assumes:

- The composite organization has five full-time resources involved in MLOps platform management. Resources dedicate 50% of their time to Red Hat OpenShift related activities.
- The average fully burdened annual salary for a member of the platform management team is $150,000.

Risks. Forrester recognizes that subscription costs may vary from organization to organization. Specific risk considerations include:

- Type and size of deployment.
- Skill set of internal resources.

Results. To account for these risks, Forrester adjusted this cost upward by 5%, yielding a three-year, risk-adjusted total PV of $1 million.

Ongoing Management

Ref.	Metric	Source	Initial	Year 1	Year 2	Year 3
H1	Internal resources dedicated to MLOps platform management	Composite		5	5	5
H2	Platform management FTE fully burdened annual salary	TEI standard		$150,000	$150,000	$150,000
H3	Percentage of time dedicated to Red Hat	Interviews		50%	50%	50%
Ht	Ongoing management	H1*H2*H3		$375,000	$375,000	$375,000
	Risk adjustment	↑5%				
Htr	Ongoing management (risk-adjusted)		$0	$393,750	$393,750	$393,750
	Three-year total: $1,181,250			**Three-year present value: $979,198**		

Financial Summary

CONSOLIDATED THREE-YEAR RISK-ADJUSTED METRICS

Cash Flow Chart (Risk-Adjusted)

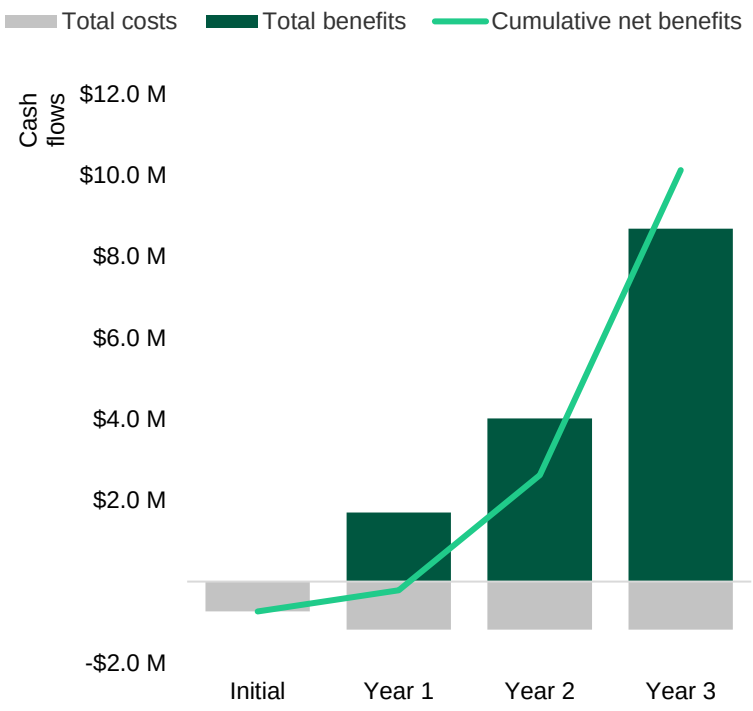

The financial results calculated in the Benefits and Costs sections can be used to determine the ROI, NPV, and payback period for the composite organization's investment. Forrester assumes a yearly discount rate of 10% for this analysis.

These risk-adjusted ROI, NPV, and payback period values are determined by applying risk-adjustment factors to the unadjusted results in each Benefit and Cost section.

Cash Flow Analysis (Risk-Adjusted Estimates)

	Initial	Year 1	Year 2	Year 3	Total	Present Value
Total costs	($735,000)	($1,181,250)	($1,181,250)	($1,181,250)	($4,278,750)	($3,672,594)
Total benefits	$0	$1,701,500	$4,016,250	$8,689,000	$14,406,750	$11,394,207
Net benefits	($735,000)	$520,250	$2,835,000	$7,507,750	$10,128,000	$7,721,613
ROI						210%
Payback (months)						13

Appendix A: Total Economic Impact

Total Economic Impact is a methodology developed by Forrester Research that enhances a company's technology decision-making processes and assists vendors in communicating the value proposition of their products and services to clients. The TEI methodology helps companies demonstrate, justify, and realize the tangible value of IT initiatives to both senior management and other key business stakeholders.

TOTAL ECONOMIC IMPACT APPROACH

Benefits represent the value delivered to the business by the product. The TEI methodology places equal weight on the measure of benefits and the measure of costs, allowing for a full examination of the effect of the technology on the entire organization.

Costs consider all expenses necessary to deliver the proposed value, or benefits, of the product. The cost category within TEI captures incremental costs over the existing environment for ongoing costs associated with the solution.

Flexibility represents the strategic value that can be obtained for some future additional investment building on top of the initial investment already made. Having the ability to capture that benefit has a PV that can be estimated.

Risks measure the uncertainty of benefit and cost estimates given: 1) the likelihood that estimates will meet original projections and 2) the likelihood that estimates will be tracked over time. TEI risk factors are based on "triangular distribution."

The initial investment column contains costs incurred at "time 0" or at the beginning of Year 1 that are not discounted. All other cash flows are discounted using the discount rate at the end of the year. PV calculations are calculated for each total cost and benefit estimate. NPV calculations in the summary tables are the sum of the initial investment and the discounted cash flows in each year. Sums and present value calculations of the Total Benefits, Total Costs, and Cash Flow tables may not exactly add up, as some rounding may occur.

PRESENT VALUE (PV)

The present or current value of (discounted) cost and benefit estimates given at an interest rate (the discount rate). The PV of costs and benefits feed into the total NPV of cash flows.

NET PRESENT VALUE (NPV)

The present or current value of (discounted) future net cash flows given an interest rate (the discount rate). A positive project NPV normally indicates that the investment should be made, unless other projects have higher NPVs.

RETURN ON INVESTMENT (ROI)

A project's expected return in percentage terms. ROI is calculated by dividing net benefits (benefits less costs) by costs.

DISCOUNT RATE

The interest rate used in cash flow analysis to take into account the time value of money. Organizations typically use discount rates between 8% and 16%.

PAYBACK PERIOD

The breakeven point for an investment. This is the point in time at which net benefits (benefits minus costs) equal initial investment or cost.

Appendix B: Endnotes

[1] Source: “Implement ModelOps To Operationalize AI,” Forrester Research, Inc., August 13, 2020.

[2] This can be augmented with Red Hat OpenShift Data Science (a managed cloud service providing a fully supported sandbox to rapidly develop, train, and test models in the public cloud before deploying in production), Red Hat Integration, and/or Red Hat Runtimes.

[3] For example, organizations can leverage an AI/ML architectural review or an AI/ML residency Red Hat Consulting Service.

[4] Total Economic Impact is a methodology developed by Forrester Research that enhances a company’s technology decision-making processes and assists vendors in communicating the value proposition of their products and services to clients. The TEI methodology helps companies demonstrate, justify, and realize the tangible value of IT initiatives to both senior management and other key business stakeholders.

[5] Source: “Margin/ROIC by Sector (US),” NYU Stern School of Business (https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/mgnroc.html).


FORRESTER®