A Forrester Total Economic Impact™
Study Commissioned By NVIDIA
March 2018

The Total Economic Impact™ Of NVIDIA DGX-1

Artificial Intelligence And Deep Learning Cost Savings And Business Benefits Enabled By NVIDIA DGX-1

FORRESTER®

Table Of Contents

Project Director:
Sam Conway

ABOUT FORRESTER CONSULTING

Forrester Consulting provides independent and objective research-based consulting to help leaders succeed in their organizations. Ranging in scope from a short strategy session to custom projects, Forrester's Consulting services connect you directly with research analysts who apply expert insight to your specific business challenges. For more information, visit forrester.com/consulting.

FORRESTER®

Benefits And Costs

Faster product development:
$2,449,136

Improved deep learning and model training:
$435,797

Reduced implementation time:
2 months

Executive Summary

As unprecedented amounts of data are generated and collected every day, organizations seek to infuse the power of artificial intelligence (AI) and deep learning to more effectively derive insights and new knowledge from data to optimize operations, create new products and services, and automate decisions.

NVIDIA commissioned Forrester Consulting to conduct a Total Economic Impact (TEI)™ study and examine the potential return on investment (ROI) enterprises may realize by deploying DGX-1. The purpose of this study is to provide readers with a methodology to evaluate the potential financial impact of the DGX-1 on their organizations.

With its DGX product line, NVIDIA offers AI supercomputers that help its customers transform their business with AI-infused applications. The DGX-1 combines NVIDIA's latest GPU set with its high-performance computing software integrated with all of the leading deep learning frameworks, such as TensorFlow, Caffe, Caffe2, MXNet, Theano, Torch, Pytorch, Microsoft Cognitive Toolkit, and more. A single DGX-1 appliance can replace 400 CPU-based servers. To better understand the benefits, costs, and risks associated with this investment, Forrester interviewed three customers with over one year of experience using DGX-1.

Prior to using DGX-1, the customers used a mix of computer workstations that were built in-house. However, prior attempts at building and implementing do-it-yourself (DIY) platforms for AI have yielded only modest success. Customers found that building and maintaining DIY GPU and CPU AI platforms was complex because of the many unanticipated integration and performance issues, as well as higher operations and support costs with DIY platforms.

Key Findings

Quantified benefits. The following risk-adjusted quantified benefits are representative of those experienced by the companies interviewed for this study. The ROI figure included in this analysis is based on a composite organization of DGX-1 users interviewed for this study.

- **Accelerate model development, leading to an incremental revenue increase of $2.4 million.** With improved accuracy and training of models, organizations were able to develop, and put into service, new initiatives for winning business and attaining internal efficiencies sooner than they would with legacy DIY platforms.
- **Avoid legacy hardware costs of $1.1 million.** The DGX-1 provides a full stack of hardware and software, which can be quickly integrated into an organization's IT ecosystem. Interviewed organizations stated that they saved significant amounts of time and staff resources when they no longer had to build, test, and maintain DIY platforms.
- **Reduce deep learning model training time by three days.** Models were trained 72 hours faster, on average, using DGX-1 units than on previous hardware buildouts. This increased efficiency of salaried data scientists conducting model training resulted in a benefit of nearly $436,000 from faster model development.
- **Reduce implementation time.** DGX units were integrated and implemented two months faster than previous solutions. This efficiency saved organizations $102,000.

FORRESTER®

ROI
294%

Benefits PV
$4.05 million

NPV
$3.02 million

Payback
<6 months

Unquantified benefits. The interviewed organizations experienced the following benefits, which are not quantified for this study:

- **Ability to experiment with larger data sets.** Interviewed organizations stated that they were more likely to begin experiments with larger data sets rather than a subset of the data because of the increased model training performance.
- **Ability to tune models more completely.** Organizations were more comfortable expanding the range of hyperparameter tuning because of faster performance leading to faster experimentation and more accurate models.
- **Ability to attract and retain top talent.** Organizations using DGX-1 can attract top talent who want to work with state-of-the-art technology. The ease of deployment and reduced maintenance requirements also greatly improved morale, which allowed organizations to retain their top performers.

Costs. The interviewed organizations experienced the following risk-adjusted costs:

- **DGX hardware costs of $447,000.** While the composite organization needed the capability of three units, the average price for one DGX-1 unit was $149,000. Additionally, a support contract from NVIDIA would add $67,050 annually.
- **Support costs of $566,393.** This is the staffing cost of dedicated support for the deployment of all three DGX-1 units.
- **Implementation costs of $12,833.** These are the labor costs of implementing and fully integrating the DGX-1 deployment into the organization's ecosystem.

Forrester's interviews with four existing customers and subsequent financial analysis found that a composite company, based on the organizations interviewed for this study, should expect benefits of $4.95 million over three years versus costs of $1.14 million, adding up to a net present value (NPV) of $3.02 million and an ROI of 294%.

Financial Summary

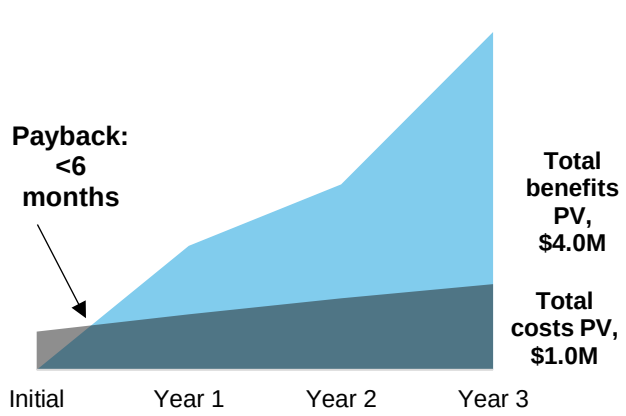

Benefits (Three-Year)

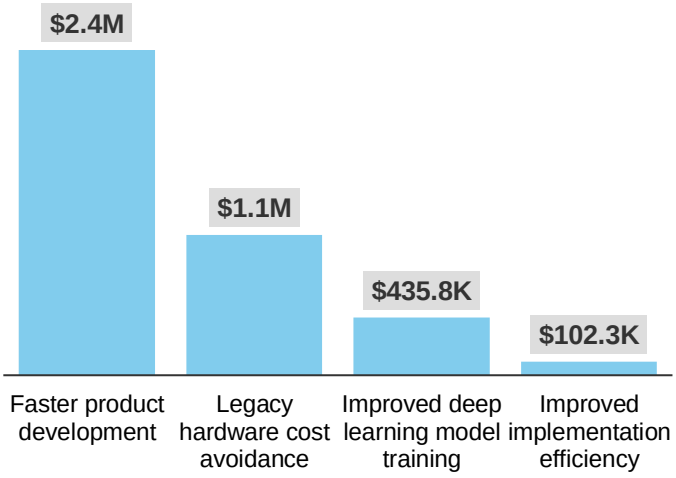

FORRESTER®

TEI Framework And Methodology

From the information provided in the interviews, Forrester has constructed a Total Economic Impact™ (TEI) framework for those organizations considering implementing NVIDIA DGX-1.

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 NVIDIA DGX-1 can have on an organization:

The TEI methodology helps companies demonstrate, justify, and realize the tangible value of IT initiatives to both senior management and other key business stakeholders.

DUE DILIGENCE
Interviewed NVIDIA stakeholders and Forrester analysts to gather data relative to DGX-1.

CUSTOMER INTERVIEWS
Interviewed four organizations using DGX-1 to obtain data with respect to costs, benefits, and risks.

COMPOSITE ORGANIZATION
Designed a composite organization based on characteristics of the interviewed 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 interviewed organizations.

CASE STUDY
Employed four fundamental elements of TEI in modeling NVIDIA DGX-1's impact: benefits, costs, flexibility, and risks. Given the increasing sophistication that enterprises have regarding ROI analyses related to IT investments, Forrester's TEI methodology serves to provide a complete picture of the total economic impact of purchase decisions. Please see Appendix A for additional information on the TEI methodology.

DISCLOSURES

Readers should be aware of the following:

This study is commissioned by NVIDIA and delivered by Forrester Consulting. It is not meant to be used as a competitive analysis.

Forrester makes no assumptions as to the potential ROI that other organizations will receive. Forrester strongly advises that readers use their own estimates within the framework provided in the report to determine the appropriateness of an investment in NVIDIA DGX-1.

NVIDIA reviewed and provided feedback to Forrester, but Forrester maintains editorial control over the study and its findings and does not accept changes to the study that contradict Forrester's findings or obscure the meaning of the study.

NVIDIA provided the customer names for the interviews but did not participate in the interviews.

The DGX-1 Customer Journey

BEFORE AND AFTER THE DGX-1 INVESTMENT

Interviewed Organizations

For this study, Forrester conducted four interviews with NVIDIA DGX-1 customers. Interviewed customers include the following:

INDUSTRY	REGION	INTERVIEWEE	KEY METRICS
Insurance	USA	Head of data science	• $20B+ revenue • 28K employees • 15 data scientists
Oil and gas	Global	Business intelligence manager	• $8B revenue • 5K employees • 12 data scientists
Inspection services	USA	Director of advanced analytics	• <$1B revenue • 50 employees • 8 data scientists
High performance computing infrastructure provider	Virtual, worldwide	CEO	• Private company • NVIDIA partner • Pay-per-use service

Key Challenges

Prior to implementing the DGX-1 data center, interviewees shared a number of challenges and pain points with their AI hardware platform. These issues slowed their ability to develop new deep learning models for business innovation.

- **Lack of support for implementation or maintenance.** Interviewees stated that when using DIY platforms, they had to provide their own support and maintenance. This led to surging expenses when trying to increase scale or add new workloads.
- **Searching for faster hardware.** Interviewees were driven by a constant search to improve processing power and performance. Rapid development in model training required consistent hardware upgrades to handle the increased workloads.
- **Using limited resources to innovate key business processes and products.** According to the experts interviewed for this study, organizations need to dedicate their data scientists to the development and training of new models, not the upkeep of hardware.

> "I don't have to have our IT team go and develop an environment and add all of these different things that have different dependencies. With the DGX-1, all of those dependencies are already ironed out for us."
>
> *Business intelligence manager, energy*

Solution Requirements

The interviewed organizations searched for a solution that could:

- **Provide a fast, scalable out-of-the-box solution.** Organizations were looking for a solution which could be implemented immediately, without additional software or hardware purchases and complex IT configurations.
- **Accommodate popular open-source deep learning frameworks.** Data science teams wanted a platform that could support multiple open-source deep learning frameworks that take advantage of GPUs.
- **Rapidly increase model training and AI-innovation endeavors**. Interviewees wanted to train models faster and get them into production quicker.
- **Dependability**. Interviewees stressed that they wanted their data scientists to focus on experimentation and innovation — not worrying about hardware maintenance.

> "When we deployed the NVIDIA solution, they came in and did most of the ramp up with our engineers. On the day they left, our engineers felt comfortable with them leaving. It was awesome."
>
> *Head of data science, insurance*

Key Results

The interviews revealed that key results from the DGX-1 investment include:

- **Faster model training and production.** With the DGX-1, organizations were able to train models faster, thereby reducing the time data scientists needed to dedicate to each experiment (model version). This also allowed organizations to implement new internal initiatives or offer new products and services sooner.
- **Improved reliability with lower operating expense.** Organizations found the DGX-1 to be more reliable than prior hardware solutions. The ease of integration and reduced downtime eliminated costs associated with ramp, tuning, and maintenance.
- **Improved accuracy.** The DGX-1 was an integral tool in the creation of machine learning processes. With the DGX-1, organizations were able to rapidly improve the accuracy of models they developed.
- **Support.** Organizations reported that NVIDIA provided exceptional implementation and ongoing support for the DGX-1 and the software stack.

Composite Organization

Based on the interviews, Forrester constructed a TEI framework, a composite company, and an associated ROI analysis that illustrates the areas financially affected. The composite organization is representative of the four companies that Forrester interviewed and is used to present the aggregate financial analysis in the next section.

The composite organization that Forrester synthesized from the customer interviews has the following characteristics:

- **Description of composite.** The composite organization is a Global 2000 multinational company. The company provides testing, sampling, and logistical planning services for various large vendors in a wide range of industries. The company employs a team of data scientists who are tasked with the creation and training of algorithms to solve complex business objectives.

  Forrester constructed this composite with characteristics that can apply to analogous use cases in the agriculture and food, automotive, aviation, chemical, construction, consumer goods and retail, energy, finance, industrial manufacturing, life sciences, logistics, mining, oil and gas, and public sectors. Forrester and NVIDIA encourage users of this study to use the TEI methodology and calculations with their own variables and assumptions to create a custom ROI model.
- **Deployment characteristics.** The organization employed a team of 10 data scientists, working on three DGX-1 appliances. The organization uses a leading framework, such as Tensorflow, MXNet, Caffe2, Pytorch, among others, to train models on large datasets so that they can develop algorithms capable of delivering insights to their sampling, testing, and logistical planning services. Furthermore, the organization uses the DGX-1 to conduct various R&D experiments with the goal of improving business processes. Prior to using the DGX-1, the composite organization was using a mix of DIY appliances, with each being supplied the computing power of roughly 133 CPU-based servers.

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.

Key assumptions
10 data scientists

3 DGX-1 units

15 versions of model sent to production before final is ready

"Now we have the technology that can turn around and evaluate a deal in a matter of hours, or a couple of days, versus taking us months to evaluate and lose the deal ultimately when somebody beats us to the punch. That's how you get to be the first instead of last. DGX-1 technology can help that."

Head of business intelligence, oil and gas

Analysis Of Benefits

QUANTIFIED BENEFIT AND COST DATA AS APPLIED TO THE COMPOSITE

Total Benefits

REF.	BENEFIT	YEAR 1	YEAR 2	YEAR 3	TOTAL	PRESENT VALUE
Atr	Faster product development	$180,000	$720,000	$2,250,000	$3,150,000	$2,449,136
Btr	Legacy hardware cost avoidance	$1,164,375	$0	$0	$1,164,375	$1,1058,523
Ctr	Improved deep learning model training	$175,240	$175,240	$175,240	$525,721	$435,797
Dtr	Improved implementation efficiency	$112,500	$0	$0	$112,500	$102,273
	Total benefits (risk-adjusted)	$1,632,115	$895,240	$2,425,240	$4,952,596	$4,045,729

Faster Product Development

In pursuit of constantly improving their process efficiencies and competitive advantage, the composite organization staffs a data science team responsible for researching and developing the use of AI technology, specifically deep learning. The team's goal is to use large volumes of data gathered through business operations and external sources and use this to train models.

By using the DGX-1, the composite organization should be able to accelerate the training and production of their proprietary machine learning models. The DGX-1 not only provides GPU power that is superior to previous hardware solutions, but it also delivers a fully integrated software stack. With this in hand, the composite organization can increase workforce efficiency and win new business with more accurate predictive models.

In modeling the impact of DGX-1 in accelerating the composite organization's product development, Forrester made the following assumptions:

- The percentage of annual revenue related to new business and operational efficiencies will increase over time because of accurate and continual model training.
- Similar to annual revenue, the acceleration factor will improve over time as new models are created, and existing models become more accurate.
- Gross margins for the organization are 40%.

Forrester realizes that product development acceleration and the associated profit impact will vary significantly across companies, countries, and industry verticals. Specific risk considerations include:

- Internal talent and resources dedicated to machine learning R&D.
- Volume and availability of rich data for model training.
- Possession of business processes or product/service lines which can benefit from automation.

The table above shows the total of all benefits across the areas listed below, as well as present values (PVs) discounted at 10%. Over three years, the composite organization expects risk-adjusted total benefits to be a PV of more than $4.05 million.

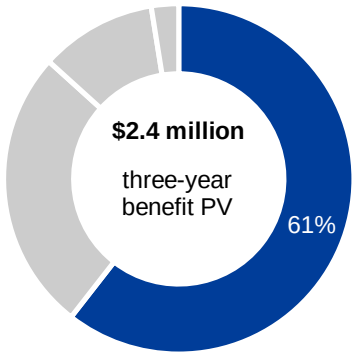

Faster product development: **61%** of total benefits.

FORRESTER®

To account for these risks, Forrester adjusted this benefit downward by 40%, yielding a three-year risk-adjusted total PV of $2,449,136.

Faster Product Development: Calculation Table

REF.	METRIC	CALC.	YEAR 1	YEAR 2	YEAR 3
A1	Annual increase in revenue		$100,000,000	$150,000,000	$300,000,000
A2	Annual revenue portion effected		15%	20%	25%
A3	Acceleration factor		5.0%	10.0%	12.5%
A4	Gross margin		40%	40%	40%
At	Faster product development	(A1*A2)*A3*A4	$300,000	$1,200,000	$3,750,000
	Risk adjustment	↓40%			
Atr	Faster product development (risk-adjusted)		$180,000	$720,000	$2,250,000

Legacy Hardware Cost Avoidance

The composite organization is able to replace legacy units, developed in-house. Prior to using the DGX-1, the composite organization primarily built their machine learning appliances in-house using components purchased from a variety of vendors. This approach was not cost effective, requiring the organization to deal with numerous vendors, as well as dedicate high-salary data scientists to building the appliances.

Based on the experience of the organizations interviewed for this study, Forrester assumes that for the composite organization:

- A single DGX-1 unit replaces three DIY appliances. The three legacy appliances had a combined compute power of 400 CPU-based servers.
- The time required to build and test a DIY appliance is 520 hours.
- The average fully burdened salary of a senior data scientist is $375,000.

Parts related to building a DIY appliance total roughly $50,000 in value. The reduction in legacy hardware cost avoidance will vary with:

- Internal skills and salaries.
- Relationships with hardware vendors and related pricing.
- Compute power of existing legacy appliances.

To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year risk-adjusted total PV of $1,058,523.

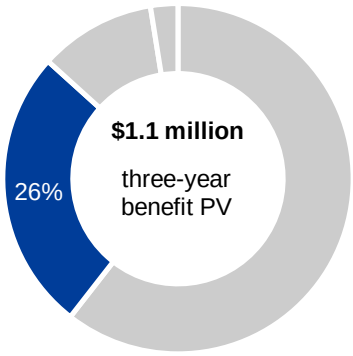

Legacy hardware cost avoidance:
26% of total benefits.

FORRESTER®

Legacy Hardware Cost Avoidance: Calculation Table					
REF.	**METRIC**	**CALC.**	**YEAR 1**	**YEAR 2**	**YEAR 3**
B1	Number of appliances replaced per DGX-1		3		
B2	DGX-1 units		3		
B3	Total appliances avoided building	B1*B2	9		
B4	Hours per appliance built		520		
B5	Average fully burdened salary		$375,000		
B6	Effective hourly rate	B5/2080	$180.29		
B7	Internal cost appliance built	B4*B6	$93,750		
B8	Parts cost per appliance		$50,000		
Bt	Legacy hardware cost avoidance	B3*(B7+B8)	$1,293,750	$0	$0
	Risk adjustment	↓10%			
Btr	Legacy hardware cost avoidance (risk-adjusted)		$1,164,375	$0	$0

Improved Deep Learning Model Training

The composite organization should find the DGX-1 offered superior computing power and reliability than previous DIY platforms. The fully integrated deep learning frameworks, high performance computing software and tools of the DGX-1 improve reliability and reduce downtime, meaning models can be trained without interruption. The composite organization was able to train models 72 hours faster and redeploy FTE assets who would otherwise be monitoring model training to other areas of the organization.

For the composite organization, Forrester assumes:

- 15 versions of a model are sent into calibration before it is fully trained and tested.
- Average fully burdened senior data scientist salary of $375,000.

Improved deep learning model training results will vary based on:

- Complexity of proprietary models.
- Volume and quality of available training data.

To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year risk-adjusted total PV of $435,797.

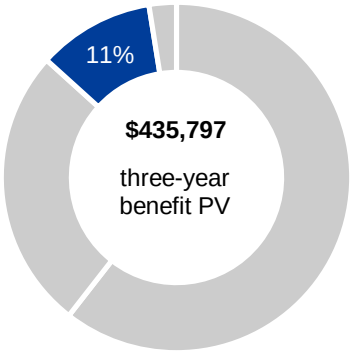

Improved model training: **11%** of total benefits.

FORRESTER®

Improved Deep Learning Model Training: Calculation Table					
REF.	METRIC	CALC.	YEAR 1	YEAR 2	YEAR 3
C1	Versions of model sent to production		15	15	15
C2	Hours saved training model per version		72	72	72
C3	Average fully burdened salary		$375,000	$375,000	$375,000
C4	Effective hourly rate	C3/2080	$180.29	$180.29	$180.29
Ct	Improved deep learning model training	C1*C2*C4	$194,712	$194,712	$194,712
	Risk adjustment	↓10%			
Ctr	Improved deep learning model training (risk-adjusted)		$175,240	$175,240	$175,240

Improved Implementation Efficiency

The composite organization should find that the DGX-1 was easy to implement and integrate with their current IT ecosystem. Furthermore, NVIDIA offered training and implementation support, ensuring that the organization was able to ramp up use of the DGX-1 two months faster than prior solutions. Due to DGX-1's ease of implementation, the composite organization was also able to have their units up and running with two fewer FTEs.

For the composite organization, Forrester assumes:

- Average fully burdened salary for senior data scientist of $375,000.
- Improved implementation efficiency results will vary based on:
- Complexity of existing IT ecosystem
- Ability and experience of FTE resources

To account for these risks, Forrester adjusted this benefit downward by 10%, yielding a three-year risk-adjusted total PV of $102,273.

Impact risk is the risk that the business or technology needs of the organization may not be met by the investment, resulting in lower overall total benefits. The greater the uncertainty, the wider the potential range of outcomes for benefit estimates.

Improved Implementation Efficiency: Calculation Table					
REF.	METRIC	CALC.	YEAR 1	YEAR 2	YEAR 3
D1	Reduced FTE required to implement		2		
D2	Reduction in months to implement		2		
D3	Avg. fully burdened salary		$375,000		
Dt	Improved implementation efficiency	D1*((D3/12)*D2)	$125,000	$0	$0
	Risk adjustment	↓10%			
Dtr	Improved implementation efficiency (risk-adjusted)		$112,500	$0	$0

Flexibility

The value of flexibility is clearly unique to each customer, and the measure of its value varies from organization to organization. There are multiple scenarios in which a customer might choose to implement DGX-1 and later realize additional uses and business opportunities, including:

- **Move less complex workloads to the cloud.** Organizations can use the DGX to perform high-intensity training. Once this is performed, the organization can move the simplified model into the cloud for production.
- **Redirect senior data scientist resources to R&D projects.** The reliability of the DGX-1 eliminates the need to have salaried data scientists dedicated to hardware maintenance and support. These FTEs can be redirected towards efforts to create and innovate products and services.
- **Continuous innovation and more model development.** Organizations can tune existing models faster and pursue additional models quicker compared to the DIY platforms.
- **Improve morale and attract top talent.** Implementing the DGX-1 allows organizations to let their data scientists focus on what they do best – cutting edge AI R&D. Organizations with the DGX-1 are better prepared to attract top talent who wish to work with leading technology and improved resilience and reduced downtime help them can help them retain existing talent.

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

"One of the things that I always see is you'll initially leverage technology to solve one problem, but then you see it down the road solving things you never even thought of."

Head of business intelligence, oil and gas industry

Flexibility, as defined by TEI, represents an investment in additional capacity or capability that could be turned into business benefit for a future additional investment. This provides an organization with the "right" or the ability to engage in future initiatives but not the obligation to so.

Analysis Of Costs

QUANTIFIED DATA AS APPLIED TO THE COMPOSITE

Total Costs

REF.	COST	INITIAL	YEAR 1	YEAR 2	YEAR 3	TOTAL	PRESENT VALUE
Etr	Hardware cost	$447,000	$0	$0	$0	$447,000	$447,000
Ftr	Support cost	$0	$227,755	$227,755	$227,755	$683,265	$566,393
Gtr	Implementation cost	$12,833	$0	$0	$0	$12,833	$12,833
	Total costs (risk-adjusted)	$459,833	$227,755	$227,755	$227,755	$1,143,098	$1,026,226

Hardware Cost

This category represents the cost associated with purchasing the DGX-1 hardware.

- The composite organization would purchase three DGX-1 units for an average price of $149,000.
- The DGX-1 unit comes with a fully integrated software stack containing GPU optimized frameworks, tools, and libraries.

Some organizations may face varying costs based on the number of units purchased.

Forrester did not risk adjust fees for hardware, as these were fixed quotes from NVIDIA. With a 0% risk adjustment, this yields a three-year risk-adjusted total PV of $447,000.

The table above shows the total of all costs across the areas listed below, as well as present values (PVs) discounted at 10%. Over three years, the composite organization expects risk-adjusted total costs to be a PV of nearly $1,026,226.

Hardware Cost: Calculation Table

REF.	METRIC	CALC.	INITIAL	YEAR 1	YEAR 2	YEAR 3
E1	DGX-1 units		3			
E2	Price per unit		$149,000			
Et	Hardware	E1*E2	$447,000	$0	$0	$0
	Risk adjustment	0%				
Etr	Hardware (risk-adjusted)		$447,000	$0	$0	$0

Support Cost

This category accounts for the resource cost related to continued upkeep and support of the DGX-1 units.

- The composite organization would require one FTE resource to support and maintain their DGX-1 appliances.
- The company would purchase a support contract from NVIDIA for $67,050 per year commencing with the purchase of the DGX-1 units.

Forrester assumes that the average fully burdened salary for a FTE supporting the DGX-1 is $140,000.

Some organizations may face more difficult challenges when supporting the DGX-1, due to a deficit of internal support skills, or complex environment.

To account for these risks, Forrester adjusted this cost upward by 10%, yielding a three-year risk-adjusted total PV of $566,393.

Support Cost: Calculation Table

REF.	METRIC	CALC.	INITIAL	YEAR 1	YEAR 2	YEAR 3
F1	FTE required		0	1	1	1
F2	Avg. fully burdened salary		$0	$140,000	$140,000	$140,000
F3	Support contract		$0	$67,050	$67,050	$67,050
Ft	Support cost	(F1*F2) +F3	$0	$207,050	$207,050	$207,050
	Risk adjustment	↑10%				
Ftr	Support cost (risk-adjusted)		$0	$227,755	$227,755	$227,755

Implementation Cost

This category accounts for the resource cost related to implementing the DGX-1.

- The composite organization required one FTE for one month to implement the DGX-1 into their environment.
- NVIDIA provided on-site support during implementation, drastically improving ramp-up speed and reducing the need for FTEs.

Forrester assumes that the average fully burdened salary of the junior data scientist or senior IT operations staff required for implementation is $140,000.

Forrester acknowledges that there may be some risks involved in implementation which interviewed organizations did not experience. Factors that may impact implementation are:

- Organizational preparedness and agility.
- Complexity of current IT ecosystem.

To account for these risks, Forrester adjusted this cost upward by 10%, yielding a three-year risk-adjusted total PV of $12,833.

One month
of total implementation and deployment time.

Implementation risk is the risk that a proposed investment may deviate from the original or expected requirements, resulting in higher costs than anticipated. The greater the uncertainty, the wider the potential range of outcomes for cost estimates.

Implementation Cost: Calculation Table						
REF.	**METRIC**	**CALC.**	**INITIAL**	**YEAR 1**	**YEAR 2**	**YEAR 3**
G1	FTE required		1			
G2	Months to implement		1			
G3	Avg. fully burdened salary		$140,000			
Gt	Implementation cost	G1*((G3/12)*G2)	$11,667	$0	$0	$0
	Risk adjustment	↑10%				
Gtr	Implementation cost (risk-adjusted)		$12,833	$0	$0	$0

Financial Summary

CONSOLIDATED THREE-YEAR RISK-ADJUSTED METRICS

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.

Cash Flow Chart (Risk-Adjusted)

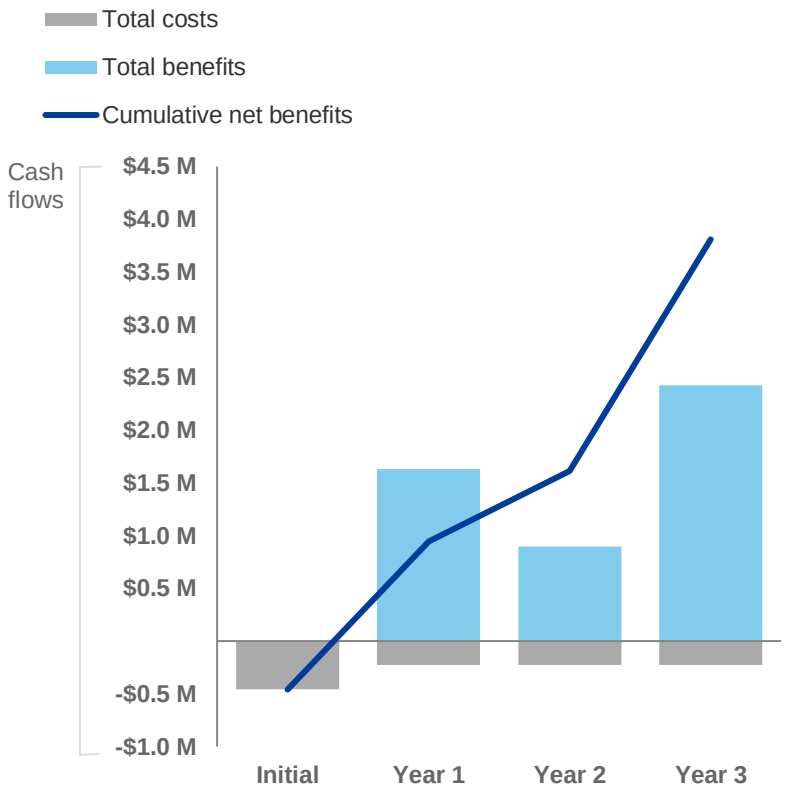

These risk-adjusted ROI, NPV, and payback period values are determined by applying risk-adjustment factors to the unadjusted results in each Benefit and Cost section.

Cash Flow Table (Risk-Adjusted)

	INITIAL	YEAR 1	YEAR 2	YEAR 3	TOTAL	PRESENT VALUE
Total costs	($459,833)	($227,755)	($227,755)	($227,755)	($1,143,098)	($1,026,226)
Total benefits	$0	$1,632,115	$895,240	$2,425,240	$4,952,596	$4,045,729
Net benefits	($459,833)	$1,404,360	$667,485	$2,197,485	$3,809,498	$3,019,503
ROI						294%
Payback period (months)						<6

FORRESTER®

NVIDIA DGX-1: Overview

The following information is provided by NVIDIA. Forrester has not validated any claims and does not endorse NVIDIA or its offerings.

The NVIDIA DGX-1 with Tesla® V100 is an integrated system for deep learning. DGX-1 features eight NVIDIA Tesla V100 GPU accelerators that connect through NVIDIA NVLink™, the NVIDIA high-performance GPU interconnect, in a hybrid cube-mesh network. Together with dual-socket Intel Xeon CPUs and four 100 GB InfiniBand network interface cards, DGX-1 provides unprecedented performance for deep-learning training. Moreover, the DGX-1 system software, powerful libraries, and NVLink network are tuned for scaling up deep learning across all eight Tesla V100 GPUs to provide a flexible, maximum performance platform for the development and deployment of deep learning applications in both production and research settings. DGX-1 with V100 GPUs achieve dramatically higher throughput than DGX-1 with previous-generation NVIDIA Tesla P100 GPUs, achieving up to 3.1x faster deep-learning training for convolutional neural networks. High-performance NVLink GPU interconnect improves scalability of deep-learning training, improving recurrent neural network training performance by up to 1.5x compared to slower PCIe interconnect. More productivity and performance benefits come from the fact that DGX-1 is an integrated system, with a complete optimized software platform aimed at deep learning that ensures DGX-1 outperforms similar off-the-shelf systems.

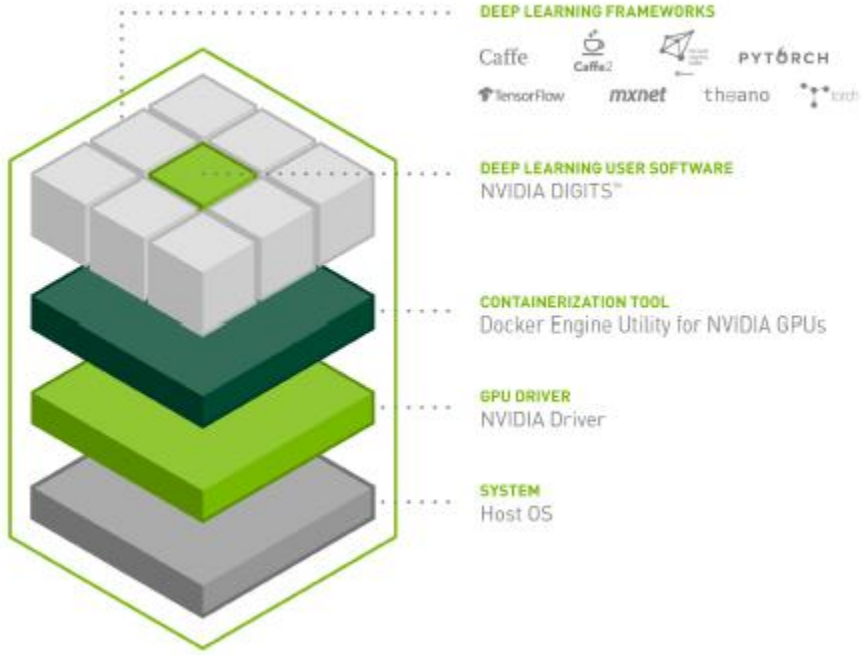

DGX-1 Software

The DGX-1 software has been built to run deep learning at scale. A key goal is to enable practitioners to deploy deep learning frameworks and applications on DGX-1 with minimal setup effort. The design of the platform software is centered around a minimal OS and driver install on the server, and provisioning of all application and SDK software in Docker containers through the DGX Container Registry5, maintained by NVIDIA. Containers available for DGX-1 include multiple optimized deep learning frameworks, the NVIDIA DIGITS deep-learning training application, third-party accelerated solutions, and the NVIDIA CUDA Toolkit. The figure below shows the DGX-1 deep learning software stack.

System Hardware Architecture

This software architecture has many advantages: Since each deep learning framework is in a separate container, each framework can use different versions of libraries like libc, cuDNN, and others, and not interfere with each other. As deep learning frameworks are improved for performance or bug fixes, new versions of the containers are made available in the DGX Container Registry. The system is easy to maintain, and the OS image stays clean, since applications are not installed directly on the OS. Security updates, driver updates, and OS

patches can be delivered seamlessly. The deep learning frameworks and the CUDA Toolkit include libraries that have been custom-tuned to provide high multi-GPU performance on DGX-1.

System Hardware Architecture

The NVIDIA® DGX-1 is architected for high throughput and high interconnect bandwidth to maximize neural network training performance. The core of the system is a complex of eight Tesla V100 GPUs connected in the hybrid cube-mesh NVLink network topology. In addition to the eight GPUs, DGX-1 includes two CPUs for boot, storage management, and deep learning framework coordination. DGX-1 is built into a three-rack-unit (3U) enclosure that provides power, cooling, network, multisystem interconnect, and SSD file system cache, balanced to optimize throughput and deep-learning training time.

The figure below shows the hardware configuration on the DGX-1:

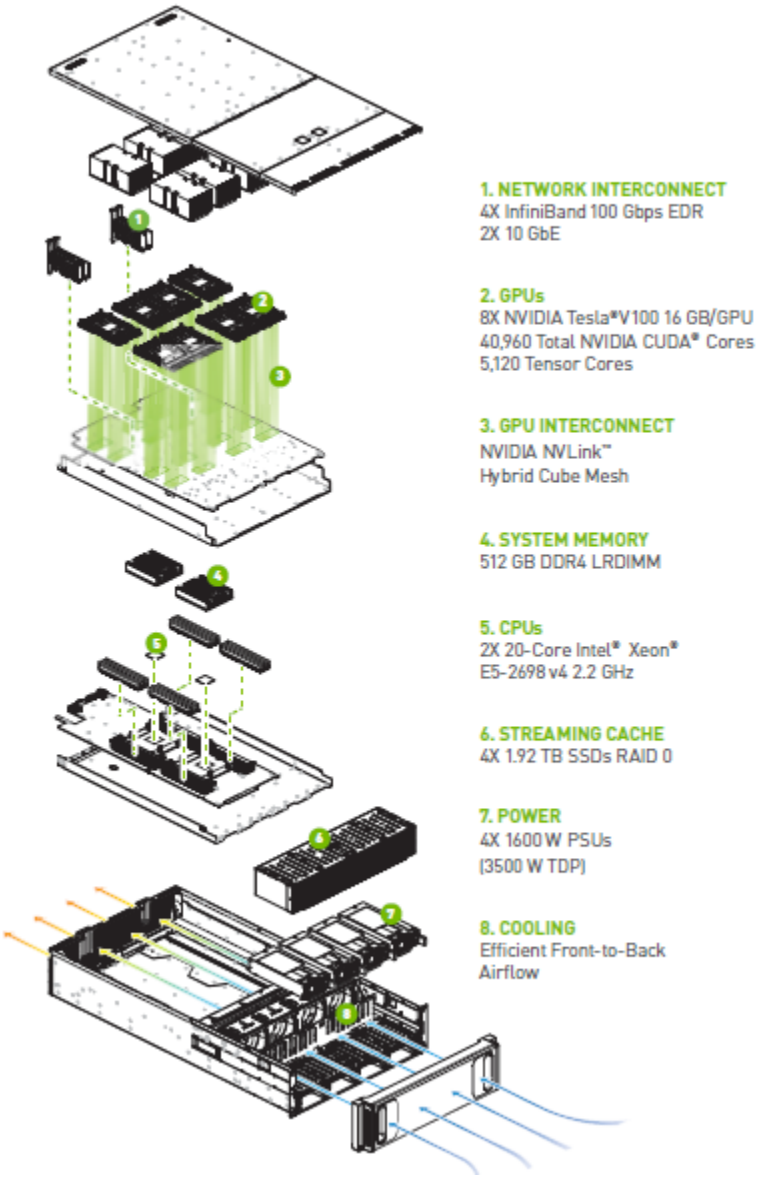

For more information on NVIDIA DGX-1, please visit www.nvidia.com/dgx-1.

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.

FORRESTER®