

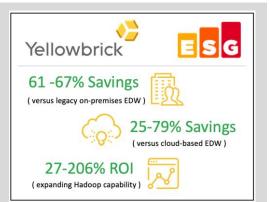
**ESG Economic Validation** 

# The Economic Benefits of Yellowbrick Data Warehouse for Distributed Cloud Environments

By Aviv Kaufmann and Kerry Dolan, Senior Validation Analysts; and Eve Falk, Validation Analyst October 2021

# **Executive Summary**

Recent events have brought about an unexpected fundamental shift in the way business functions. Organizations have unexpectedly had to rethink not only the way their employees work, but how the entire organization functions from procurement to delivery. Never before has timely insight been more critical to making quality informed decisions. Expanding this ability to different parts of the business to quickly query and extract actionable insights from real-time data feeds can help reduce risk and play an important role in overall company performance.



ESG research shows that almost half of organizations already considered their data as part of their core products and services, and an amazing 93% of organizations felt that their company was successful or very successful in utilizing data to create/discover significant, new revenue opportunities or to increase value in the future.<sup>1</sup>

ESG validated that, when compared with other on-premises and cloud-only solutions, Yellowbrick Data Warehouse for hybrid, multi-cloud, and distributed cloud environments **delivers improved and predictable performance and cost with improved operational simplicity and deployment**. ESG's modeled scenarios, based on validation with real-world Yellowbrick customers, predict significant savings for organizations looking to deploy a modern and agile analytics environment that delivers certain timely insights to answer today's uncertain questions. One customer summed it up very well: *"With Yellowbrick, we can now add more retail data, get faster insight, and do not have to add more people or hardware."* 

<sup>&</sup>lt;sup>1</sup> Source: ESG Research Report, *Data Storage Trends in an Increasingly Hybrid Cloud World*, March 2020.



#### Introduction

This ESG Economic Validation focused on the quantitative and qualitative benefits organizations can expect by deploying Yellowbrick's modern data warehouse for hybrid and multi-cloud environments instead of legacy enterprise data warehouse (EDW) or cloud-only enterprise data warehouse solutions and by augmenting and extending the value of existing Hadoop deployments.

#### Challenges

It is well known that we are generating and gathering more data than ever before. Just about every aspect of the business and the surrounding ecosystems provides bits and pieces of micro and macro intelligence that, when combined, make us more knowledgeable about the past, present, and future direction of operations. But actionable insight can only be derived through efficient delivery, comprehensive collection, and effective correlation between the data points that we collect. The EDW has provided organizations with an effective means to generate insight through data analytics for decades, but many of the top EDWs were designed around a rigid on-premises pipeline process derived long before the hybrid cloud technologies of today existed. Release after release has resulted in improved capabilities, but at the expense of trading off functionality for added complexity, and with less-than-ideal levels of performance, cost, and agility.

In an ESG research study, performance was overwhelmingly identified as the most common challenge experienced by organizations today with their enterprise data warehouse solutions,<sup>2</sup> which comes as no surprise, since faster performance allows for quicker and more inclusive real-time insight. A separate ESG research study from 2019 identified cost as the most common factor that was having an adverse impact on organizations' data analytics strategies and investments over the previous 24 months.<sup>3</sup>

What challenges have you experienced with your enterprise data warehouse solution? (Percent of respondents, N=310, multiple responses accepted)

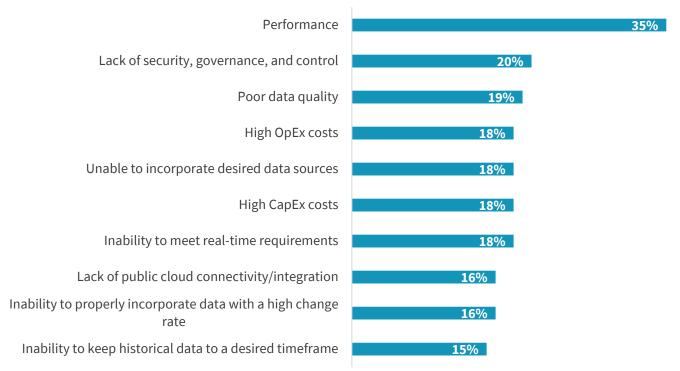


Figure 1. Top 10 Challenges Experienced with Enterprise Data Warehouse Solutions

Source: Enterprise Strategy Group

<sup>&</sup>lt;sup>2</sup> Source: ESG Brief, *Enterprise Data Warehouse Trends*, March 2020.

<sup>&</sup>lt;sup>3</sup> Source: ESG Master Survey Results, <u>*The State of Data Analytics*</u>, August 2019.

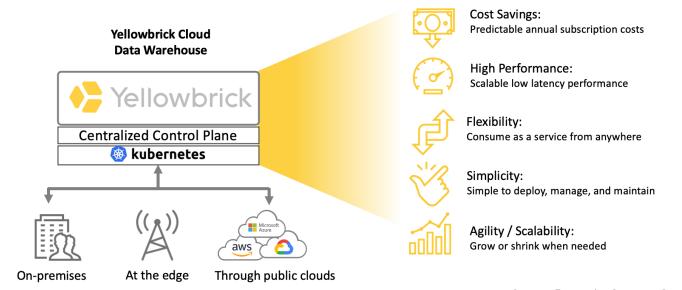
While organizations certainly understand the value of the insight that they are unlocking from their data, it is critical that they invest in an EDW that provides high levels of performance while minimizing both capital and operational costs (aka, price/performance). Legacy EDWs provide high levels of performance but can be extremely rigid and costly to purchase and operate. Open source Hadoop-based solutions can be more cost-effective and agile but can be complex to build and manage and offer limited performance and scalability for real-time analytics. Cloud-only EDWs offer excellent scalability and agility with reduced complexity but performance of a solution. A solution that can offer all three across private clouds, public clouds, and edge locations could prove a great platform for organizations.

## The Solution: Yellowbrick Data Warehouse

Yellowbrick is a modern data warehouse natively designed for hybrid cloud that offers the exceptional performance and high levels of concurrency of specialized on-premises systems and also the agility, scalability, and simplicity of cloud-only solutions at a low and predictable cost through a fixed-cost annual subscription. Yellowbrick can be consumed as a service anywhere with the same data and performance everywhere—either on-premises/in a private cloud, through any public cloud, at the edge/distributed environments, or a combination of these environments. Yellowbrick Data Warehouse offers:

- Industry-leading price/performance for mixed workloads on real-time and at-rest data.
- Flexibility to deploy in a data center/private cloud, public cloud(s), and at the edge/distributed environments, with the ability to manage all environments from a centralized control plane.
- Compatibility with common BI, data science, and data motion tools.
- High availability for tier-1 applications.
- ANSI SQL compliance/PostgreSQL front-end for streamlined migration and management.
- Enterprise-class security and compliance (encryption everywhere, HIPAA, SOC 2, PCI-DSS, and FedRAMP).
- The ability to take full advantage of Kubernetes. Yellowbrick Data Warehouse can run anywhere and on anything.

#### Figure 2. Yellowbrick Cloud Data Warehouse



# **ESG Economic Validation**

ESG completed a quantitative economic analysis on Yellowbrick Data Warehouse as compared with alternative solutions such as specialized on-premises EDWs, cloud-only EDWs, and Apache Impala implementations.

ESG's Economic Validation process is a proven method for understanding, validating, quantifying, and modeling the economic value propositions of a product or solution. The process leverages ESG's core competencies in market and industry analysis, forward-looking research, and technical/economic validation. ESG conducted in-depth interviews with end-users to better understand and quantify how Yellowbrick has impacted their organizations, particularly in comparison with previously deployed and/or experienced EDW solutions. In addition to having experience with on-premises Hadoop and specialized EDW solutions, some of the customers interviewed had experience with, or had run proofs of concept on, cloud-only EDW solutions. The qualitative and quantitative findings were used as the basis for a simple economic model comparing the expected costs of on-premises, cloud-only, and SQL-on-Hadoop solutions.

#### Yellowbrick Economic Overview

ESG's economic analysis revealed that Yellowbrick provided its customers with significant savings and benefits in the following categories:

- **Reduced price/performance of EDW** Yellowbrick provided customers with low and predictable annual costs and greatly improved performance when compared with previously used solutions, resulting in significant savings from a price/performance perspective and providing expanded analytics capabilities to the organization.
- Architectural flexibility and simplicity Yellowbrick is a natively hybrid cloud EDW designed to eliminate constraints, provide flexibility of choice, and reduce complexity for organizations looking to modernize their analytics environment across the distributed cloud while removing many of the limitations and headaches they may have experienced with previous solutions.
- **Operational savings** Unlike many systems, Yellowbrick provides increased performance and flexibility of choice without adding complexity around planning, deploying, managing, and maintaining the solution. Daily management is simplified through a single interface and automation or elimination of many traditional management tasks.



ESG found that customers reported significant savings to achieve the same levels of performance and/or reported significant performance increases over their previous deployments that led to the ability to perform jobs and queries faster. Customers reported savings in the following areas:

Improved performance – ESG previously validated the results of performance testing that demonstrates a significant advantage over cloud-only EDW (up to 5x faster response time), legacy HDD-based on-premises EDW (up to 100x faster), and Impala (2.64x faster). ESG also validated the ability of the latest Yellowbrick second-generation platform to better meet the needs of high-

performance use cases with up to a 3x performance improvement versus the first-generation platform. While some flash-based specialized on-premises EDWs may also be able to provide high levels of performance, this comes at an extremely high cost. The most intriguing aspect of the Yellowbrick performance advantage is that the

"We got roughly a 3x-5x performance improvement versus the on-premises solution that we replaced."

5

solution is often cheaper or on par with other solutions, resulting in a huge price/performance advantage for Yellowbrick (as high as 25x).

 Predictable performance – Yellowbrick customers reported that not only did the Yellowbrick Cloud Data Warehouse provide improved performance regardless of the underlying architecture it was running on, but the loading and query time of Yellowbrick remained consistently low, where other cloud-only and on-premises solutions often did not provide the same levels of consistency—especially at higher levels of concurrency.

"We saw cube build times reduced from 24 hours down to 8 hours compared to our cloud-only EDW."

• **Reduced and predictable cost** – Customers reported substantial savings since migrating their analytics workloads to Yellowbrick. Whether deployed on-premises or in the cloud, all Yellowbrick deployments are priced competitively through an all-inclusive and predictable annual subscription or can be consumed flexibly with on-demand pricing. Unlike on-premises infrastructure, the hardware does not have to be overprovisioned or purchased with a large upfront capital expense and does not require additional licensing, support, and

"[Cloud DW] costs were going up and up – we were jumping through hoops like turning the service on and off at certain hours throughout the day to minimize costs. We tried to make a schedule, but users would break the schedule; it did save hundreds of dollars a day, but it became quite a headache for the operations team. Also, scaling was not quick with [Cloud DW] at that time. Resizing a cluster took time and downtime." maintenance payments. In addition, when consumed as a service in the cloud, Yellowbrick provides significant savings and more cost predictability when compared with most cloud-based EDW pricing models that may have additional cost for data egress, bandwidth tiers, or streaming services and limitations on number of users or costs to increase parallelism.

• Elimination of constraints – Yellowbrick enabled customers to achieve much more than they had been able to with their previous legacy on-premises or cloud-only solutions, based on the fact that many of the performance,

concurrency, and/or technical constraints were no longer issues. This enabled organizations to run a higher volume of batch queries and include a more holistic data set in their analysis. Although the systems that they were comparing to may have been equivalent from a storable capacity perspective, the cost to make these systems and technologies capable of processing the same amount of addressable data with the equivalent level of in-memory performance provided by Yellowbrick would have been extremely prohibitive where and if possible. *"For filtering and returning a few or thousands of rows out of millions or billions, that's where Yellowbrick shined."* 

• Faster time to insight – The price/performance offered by Yellowbrick enabled organizations to run more queries and generate reports much faster. It also allowed them to ask new questions and include more comprehensive data sets in their analysis than they were able to with previous solutions in the same window of time. *"We saw lots of feedback around response times primarily, saying 'I'm running this in Yellowbrick now, and it runs 10x faster than it ran in big data natively, so I can do many more queries during a normal workday than I could before."* 

# Architectural Flexibility and Simplicity

As organizations look to modernize their data centers, processes, and applications, it is important that they make the right choices to ensure success going forward. As they assess and make decisions between private, public, or hybrid cloud solutions, they must make sure that continuity is provided for the operations that the business has relied on for decades while also improving and modernizing analytics capabilities to seamlessly integrate with Kubernetes and modernized

services across distributed cloud environments. The simplicity and architectural flexibility offered by Yellowbrick can make this transition quick and simple, while reducing some of the risk involved.

- **Purpose-built design** Yellowbrick was purpose-built from the ground up with a goal in mind of solving modern analytics problems, without the limitations often imposed when incrementally improving legacy hardware and software platforms that have existed for generations. The optimized software and specialized hardware were designed to support the latest technologies but also make optimal use of whatever organizations are running on to offer the ideal balance between functionality, price, and performance for today's distributed analytics workloads.
- **Distributed cloud flexibility** Customers liked the idea that Yellowbrick was designed for the distributed cloud, which allowed them to choose between running the data warehouse on-premises for control and compliance and/or consuming it as a service for operational flexibility. Many legacy on-premises EDWs were simply ported to run on virtual infrastructure in the cloud, and modern cloud EDW solutions do not offer an on-premises option. Yellowbrick was designed to run on Kubernetes and provides single tenant consistency between deployment

choices, as they both run on dedicated Yellowbrick hardware and can run on any cloud or multi-cloud environment with identical sets of data and the same high levels of performance. By disaggregating the storage from the compute, Yellowbrick also allows organizations to take advantage of any compute resources while keeping data where it is to avoid painful or costly migrations.

"We went from 20 full racks down to 3 Yellowbricks in less than a Rack – we are saving more than \$100K a year in power and cooling alone."

- Simplified transition The transition to Yellowbrick proved much simpler than expected for many of the customers we spoke with and was certainly less painful than other data warehouse transitions that they had made previously. The tools and services provided by Yellowbrick allowed a smooth transition from previous platforms, and most of the SQL compliant data did not have to be modified at all. Yellowbrick was able to ensure that their box met all of the security, compliance, availability, and ecosystem compatibility required at the customer site. *"It's a pretty smooth process—the machine comes in, pre-tested—you plug it in and load data, and the thing works. It was mostly lift and shift from an end-user standpoint, since they use PostgreSQL language. And the support is great, very responsive."*
- Seamless interoperability Yellowbrick was built with a Postgres front-end, is ANSI SQL compliant, and provides support for open APIs like Kafka and Spark. This means that no specialized skills are required for writing queries, and Yellowbrick integrates easily with all of the specialized BI, data science, and data motion tools that customers are used to using, extending previous investments in technologies and training. In addition, Yellowbrick can provide enhanced real-time analytics capabilities at a lower cost to augment any Hadoop deployment by replacing any SQL-on-Hadoop query engine like Hive and Impala with improved performance and less operational complexity.
- **Reduced EDW footprint** When using Yellowbrick in a private cloud, customers reported up to a 97% reduction in footprint when replacing racks and racks of their previous legacy EDW solutions with Yellowbrick. The ultra-dense all-flash Yellowbrick platform not only provides significant footprint savings over commodity hardware and legacy systems on a per-node basis but can provide improved performance with fewer nodes, fewer compute cores, and less memory, leading to significant reduction in power, cooling, and data center costs. This advantage is even greater today with the latest second-generation Yellowbrick platform, which can provide the same levels of performance with far fewer nodes and cores than the first generation or up to 3x greater performance in the same footprint.

# Operational Savings

Yellowbrick was designed with a goal of minimizing complexity; making deployment, management, and operation far simpler; and allowing more users to access more of the organization's data faster. This results in significant savings and benefits provided by faster time to value and improved operational efficiency.

• Faster time to value – End-users reported that they were able to deploy and POC the Yellowbrick solution in a fraction of the time and were able to completely transition to the Yellowbrick system in as little as a few months, providing significantly shorter time to value for analytics operations when compared with some of the legacy data warehouses they had encountered. Improved capabilities

with less time spent in transition directly translated into earlier enablement of improved business processes, operational savings, and direct increases in revenue.

"You can see the team at Yellowbrick thought about how to implement things in a way that was simple for the user."

 Elimination of tasks – Customers reported that Yellowbrick eliminated several of the time-consuming tasks of having to tune the hardware, database, and queries for optimal performance: vacuuming or grooming the data; or making cubes, index

performance; vacuuming or grooming the data; or making cubes, indexes, and partitions, all of which can consume significant amounts of time. *"We have probably 10 Yellowbrick machines for dev and test and production, and I'd say we use two people to manage that."* 

- Simplified administration and queries Administrators were able to take advantage of the simplicity and advanced functionality provided by Yellowbrick's user interface and automation capabilities to reduce total administration time by up to 33% versus previously deployed data warehouse solutions. This freed up administrators to work on net new functions that helped to improve the overall analytics capabilities of the organization. *"Yellowbrick has given us better workload management. We now can have analysts, developers, QA, and production running off the same machine without necessarily having to duplicate data across multiple machines."*
- Improved support experience In addition to reporting fewer issues overall, one often overlooked advantage that resonated with all of the customers who we spoke with was the level of care and support they received from Yellowbrick. Customers reported that, as a relatively new organization, Yellowbrick was far more responsive to their questions and needs, often implementing quick workaround or feature support solutions for organizations without added cost. Most of these organizations were used to dealing with large DW support organizations where they did not receive the same level of attention and responsiveness. This improved support can help save the organization valuable time and resources if and when issues arise.
- Faster data availability Yellowbrick provides up to 10TB/Hr bulk data loading and can stream inserts in parallel at a rate of millions of per second, with new data available for queries near instantly without the need for data grooming. This means that business insight was generated faster, enabling organizations to take action earlier with improved confidence and resulting in increased revenue, reduced risk, lower operating cost, and improved customer satisfaction. For one credit card company, this results in a reduction of risk: *"We can correlate more data to see where there are interactions, which can mean earlier detection of patterns, which helps limit losses."*

# **ESG Analysis**

ESG leveraged the information collected through vendor-provided material, public and industry knowledge of economics and technologies, third-party testing, and the results of customer interviews to create a series of three-year TCO/ROI models that compare the costs and benefits of Yellowbrick Data Warehouse with alternative on-premises and cloud-only data warehouse solutions as well as implementations where Yellowbrick Data Warehouse is used to augment existing Hadoop environments. ESG's interviews with customers who have recently made the transition, combined with experience and expertise in economic modeling and technical validation of analytics solutions, helped to form the basis for our modeled scenarios.

## Modeled Scenario #1: Yellowbrick versus Legacy On-premises Enterprise Data Warehouse Solutions

The first scenario compared the expected costs of providing data warehouse services for a large enterprise organization with roughly 280 TB of user data. ESG based the sizing of alternative solutions on a combination of end-user-reported deployments that had been improved and replaced with an on-premises Yellowbrick instance and modeled and published sizing and performance relationships against common competitors. ESG compared Yellowbrick against a legacy solution

based on HDD technology (legacy) and the latest offering from a leading EDW solution with flash storage technology (performance equivalent).

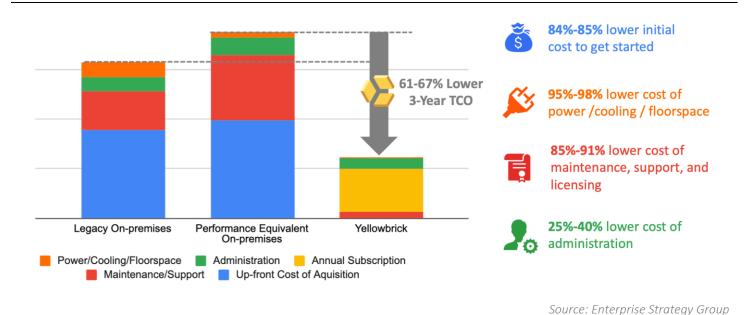
Both of the legacy solutions require a very large upfront investment in capital while the Yellowbrick solution, in addition to being more cost-effective, spreads the operational subscription costs out equally over the three-year period, allowing the organization to be up and running faster with as little as 15% of the upfront investment. The dense 10U Yellowbrick solution required only a portion of a single rack while the alternative solutions required two to eight full racks. ESG's assessment found that the Yellowbrick solution would save organizations up to 98% on power, cooling, and floorspace costs when compared with the larger and less efficient alternative solutions. Using conservative assumptions that were validated through customer interviews, ESG modeled the administrative benefits provided from simpler planning/purchasing, deployment and installation, hardware administration, and daily EDW administration costs, freeing up resources to work on finding ways to get more value from the analytics environment.

Figure 3 shows the results of ESG's three-year modeled TCO comparison for on-premises enterprise data warehouse solutions.

# Why This Matters

Over the years, organizations have come to rely heavily on the on-premises data warehouse as a core foundation of the business, but high-performing EDW systems are extremely expensive to purchase, operate, and maintain.

ESG found that the highperformance Yellowbrick Data Warehouse avoids a huge upfront capital expense with faster deployment as well as significant ongoing savings in administration, support and maintenance, and data center operating costs.



#### Figure 3. Expected Three-year Total Cost of Ownership for On-premises EDW

# Modeled Scenario #2: Yellowbrick versus Cloud-only Enterprise Data Warehouse Solutions

ESG next compared Yellowbrick's cloud subscription model to several alternative cloud-only enterprise data warehouse solutions. ESG based the sizing of alternative solutions on validation of performance testing results and known sizing guidelines relative to on-premises deployments. For the alternative solutions, ESG compared Yellowbrick against several

deployment sizing options and against several payment models to get an understanding of the range of options that a customer might choose between.

# "We saw savings of \$4K to \$32K per month for each cloud data warehouse cluster we replaced with Yellowbrick."

Just like the cloud-only data warehouses we compared it to, Yellowbrick's cloud deployment model offers substantial

operational savings as there is no large upfront investment required and there is no physical hardware to deploy, manage, and maintain. In addition, there are no data center-related costs for power, cooling, and floorspace. But unlike the cloud-only EDWs, customers can expect the same level of performance and concurrency as an on-premises deployment delivered from the single tenant Yellowbrick cloud deployment, and Yellowbrick can be accessed from any cloud.

The first cloud-based alternative (Cloud EDW 1) offered a large variety of infrastructure and payment options to accommodate different user preferences. By paying everything upfront, three years in advance, organizations can achieve significant savings. But paying this far in advance also comes at a cost in the form of an upfront investment that ties up funds that

# Why This Matters

Cloud-based EDW solutions provide organizations with improved business agility and significantly reduced operational costs versus on-premises deployments, but cost and performance can be unpredictable, and solutions often impose concurrency limitations.

Yellowbrick Data Warehouse, when deployed as a service, provides the exact same levels of performance and concurrency as when deployed on-premises with multi-cloud agility and simplicity for a fixed annual subscription rate with no surprises. could otherwise be put to use elsewhere. For this reason, ESG recommends building in a "cost of capital," which is essentially the cost of debt for an organization, until the benefits are realized (ESG used a conservative 8%). ESG also modeled annual agreement pricing as well as on-demand pricing.

The second cloud data warehouse (Cloud EDW 2) provides a simpler all-inclusive pricing model based on the number of hours a data warehouse was active, allowing users to save money during idle periods. Although the Yellowbrick and Cloud EDW 1 are available and priced for operations 24 hours per day, ESG assumed conservatively that Cloud EDW 2 would on average only be processing workloads 12 hours per day. Cloud EDW 2 is built on shared architecture and offered in a variety of preconfigured data warehouse sizes. ESG reviewed performance tests that indicated that it would require three to five separate data warehouses to provide similar levels of performance achieved by Yellowbrick. For this exercise, ESG priced a single data warehouse built on a cluster with 4 times as many nodes as the other cloud solutions (functional equivalent), as well as three to five separate data warehouses, each built on smaller clusters (performance equivalent).

ESG's cost analysis predicted that the Yellowbrick cloud solution would be 25% to 79% lower cost than the alternative solutions. While the majority of the cost for all solutions comes in the form of cloud-based subscription costs, additional incurred costs for data storage, network egress, data loading, and streaming inserts were factored in where required. While these costs are not large proportionately, they are costs that can be unpredictable and are not incurred by the Yellowbrick solution.

Figure 4 shows the results of ESG's three-year modeled TCO comparison for cloud-based enterprise data warehouse services.

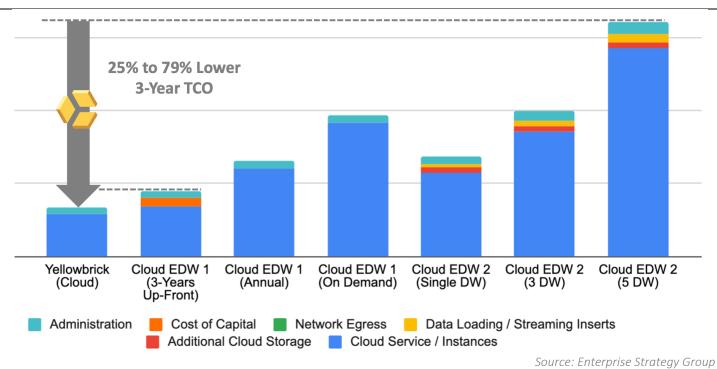


Figure 4. Expected Three-year Total Cost for Cloud-based EDW Services

## Modeled Scenario #3: Yellowbrick versus Apache Impala for Hadoop Deployments

Yellowbrick Data Warehouse can be used to augment existing investments that organizations have made by moving massive amounts of data to Hadoop solutions, which provide good batch analytics processing but struggle to provide adequate functionality and performance for ad hoc queries and real-time analytics. Customers can add SQL abstraction layers like Impala and Hive to achieve some functionality for these workloads, but these solutions can be complex to

operate and offer limited transactional performance that is restricted by the large block layout of the HDFS file system. By quickly migrating data from HDFS into the Yellowbrick Data Warehouse at extremely high speed and easily integrating feeds of new data, Yellowbrick can be used to provide real-time and ad hoc analytics capabilities to the organization with close to 3x the performance of Impala. One customer whom we interviewed stated that Impala was very memory-intensive and thus the price/performance was not viable, so they never made it to the POC stage. *"I think Impala would have been at least 2x more expensive. And there was always a question about whether it could really scale."* 

Based on a real-world use case, ESG modeled the expected cost and benefits for an organization to deploy a Yellowbrick on-premises Data Warehouse to augment an existing 25-node Hadoop cluster with a goal of providing improved analytics capabilities to the business instead of deploying Impala. In order to provide a reasonable performance solution with Impala, the organization estimated it would have to double the number of nodes in their cluster (to 50 nodes) to provide the added compute, memory, and disk resources to handle the added services and workloads.

By making an investment in Yellowbrick over three years, ESG's analysis predicted that the organization could avoid the need to grow its cluster and invest in expensive Impala licenses. ESG's analysis predicted these costs to be near equivalent, but the 5-node, 320-vCore Yellowbrick solution would offer far better performance, would be far less complex to manage and maintain (60% savings), and would cost far less to operate (97% savings) over three years compared to the 50-node Impala solution. ESG estimated the operational and administrative savings to total \$835K over three years, resulting in an expected three-year return on investment of 27%.

However, because the Yellowbrick solution would be expected to outperform the Impala solution by more than 3x, organizations can gain additional benefits by deploying Yellowbrick to provide new and/or improved analytics capabilities that would directly impact revenue. To model this, ESG assumed that the modeled organization currently operated at \$50M of annual revenue.

# Why This Matters

Hadoop solutions provide organizations with a cost-effective solution to consolidate data and perform batch processing to generate important insight, but expanding analytics capabilities through the addition of SQL abstraction layers like Hive or Impala can be costly, complex, and limited in performance.

ESG's cost-benefit analysis shows that adding a Yellowbrick Data Warehouse can provide significant operational savings and enhanced analytics performance and capabilities that positively impact the bottom line. Using conservative assumptions, ESG assumed that a typical organization can expect roughly a 5% annual revenue increase generated through processes directly benefitting from the use of analytics. ESG assumed that the addition of ad hoc and real-time analytics could account for about 30% of this expected revenue impact, and the faster time to insight of the Yellowbrick solution would allow for greater inclusion of data and provide an increase in accuracy of all generated insight, resulting in an additional improvement of about 20% across all query types. In total, this resulted in an expected annual revenue increase of roughly \$990K, a conservative total of only a 1.98% increase in expected revenue due to the increased capabilities provided by Yellowbrick. It should be noted that for many organizations, Yellowbrick may make a much larger impact to revenue, depending on a great number of variables. Factoring in this expected revenue impact over 3-years to the predicted savings resulted in an increased expected ROI of 206%.

Figure 5 shows the results of ESG's three-year modeled cost-benefit analysis for deploying the Yellowbrick Data Warehouse instead of Impala to augment an existing Hadoop-based deployment with improved analytics capabilities.

# Savings / Benefit Cost Expended Total Expected Savings and Benefits 27% to 206% Expected Return on Investment (ROI) Cost Of Adding Avoided Cost of Adding Additional HW and Impail a Licences

#### Figure 5. Three-year Cost-benefit Analysis to Expand Hadoop Analytics Capabilities with Yellowbrick

Source: Enterprise Strategy Group

While ESG's models are built in good faith upon conservative, credible, and validated assumptions, no single modeled scenario will ever represent every potential environment. ESG recommends that you perform your own analysis of available products and consult with your Yellowbrick representative to understand and discuss the differences between the solutions proven through your own proof-of-concept testing.

# **The Bigger Truth**

An effective analytics strategy has become a critical cog in the operational engine for modern data-driven organizations. The value of timely and comprehensive insight has been proven time and time again to provide positive and quantifiable improvements and advantages across nearly every part of the business from procurement and operations to revenue optimization and customer satisfaction. ESG research shows that through better understanding of their data, organizations have realized a wealth of benefits—with the top five being improved security and compliance, faster decision making, improved employee productivity, improved products and services, and improved ability to adapt to changing market conditions.<sup>4</sup>

ESG validated the overwhelmingly positive experiences that customers have had with the Yellowbrick Data Warehouse through a series of customer interviews. While the interviews covered many areas of focus, it was clear that customers agreed strongly about three things regarding Yellowbrick:

- First, Yellowbrick provided impressive performance increases (including lower response times, faster load times and scan rates, and improved network performance) and substantial operational savings compared to the on-premises and cloud technologies that they had replaced.
- Second, although any change is difficult, Yellowbrick was much quicker to deploy and far easier to learn to use and operate.
- Finally, Yellowbrick had provided them with far better customer service and was more responsive to their needs than the large vendors they were used to dealing with. As a result, they were able to successfully transition off of

<sup>&</sup>lt;sup>4</sup> Source: ESG Research Report, *Data Storage Trends in an Increasingly Hybrid Cloud World*, March 2020.

aging platforms and modernize and expand the analytics capabilities of the business quickly and without adding major risk or interruption to the organization.

The design of the Yellowbrick Data Warehouse for hybrid and distributed cloud environments provides organizations with improved and predictable performance, operational simplicity, and deployment flexibility with a cost-effective, fixed-cost, subscription-based licensing model. Yellowbrick was built from the ground up to work with Kubernetes and can be deployed on and make the most of any hardware or cloud infrastructure while realizing the true benefits of recent technologies available in the cloud like NVMe and FPGAs (most CDWs were not built for this). Through these advantages, ESG's modeled scenarios predicted that Yellowbrick can:

- Lower expected three-year total cost of ownership by up to 67% when compared with alternative on-premises enterprise data warehouse solutions.
- Save customers 25-79% compared to cloud-only data warehouse solutions.
- Provide a 27% to 206% return on investment by delivering additional analytics capabilities to Hadoop environments.

Today's distributed organizations are generating more data in more locations (private clouds, public clouds, and at the edge) than they have the time or money to analyze. The modern design, deployment flexibility, and economic advantage that Yellowbrick offers enables more people to ask more questions, ultimately impacting the bottom line of the business. One customer reported that the insight derived from Yellowbrick paid for itself in the first use case analyzing gambling data that directly led to an increase in profit. Others reported that Yellowbrick allowed insurance reports to be generated in only 45 seconds versus the 3.5 hours that it took previously. The performance improvement is impressive, but what is really impressive is that **this improvement ultimately led to a fundamental shift in operational ability** since insurance questions and approvals could now be answered in near real time during a call with a customer, rather than requiring multiple calls and follow-up discussions. As is the case with many rearchitected technologies, testimonials and modeled predictions may sound too good to be true. If that is the case, ESG suggests that you set up a POC to see for yourself and let Yellowbrick prove it.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.