Total Cost of Usage: The Key to Understanding the True Costs of a Cloud Data Warehouse

WHITE PAPER

early adopter research

ACTIAN™

CONTENTS

early adopter research

EXECUTIVE SUMMARY

Total Cost of Ownership (TCO) is a common method for evaluating data warehouse costs. When evaluating cloud data warehouses, however, this approach does not make sense. This paper introduces the concept of Total Cost of Usage (TCU) and presents a proforma comparison that illustrates the value of cloud data warehouses, highlighting the features and functions of Actian Avalanche Cloud Data Warehouse.

In our view, the entire idea of Total Cost of Ownership, which is a common way to understand costs, just doesn't apply to the cloud.

TOTAL COST OF USAGE - THE KEY TO UNDERSTANDING THE TRUE COSTS OF A CLOUD DATA WAREHOUSE

Understanding the new opportunities that the cloud presents is far more difficult than most people realize. When a powerful service like a data warehouse is made available in the cloud, it is not just a technology change, but a profound change in the division of labor and the economics driving costs. It is easy to miss subtleties that can have a huge dollar impact.

In our view, the entire idea of Total Cost of Ownership, which is a common way to understand costs, just doesn't apply to the cloud. For most end-users of cloud data warehouses, the cost of ownership drops to zero. The ability to rapidly use the cloud service and see results is more immediate. Ownership is a burden that is no longer required. Similarly, maintenance of hardware, software stack, and all the skills and infrastructure to operate it are not needed as well. Much of the TCO thinking about accepting, managing, owning, and amortizing those responsibilities and risks can be discarded. Instead, both the business and IT focus on the growth required, who in the organization requires access, and providing a valuable service, tied to achieving goals they are accountable for. The cloud therefore removes barriers to growth. For those who best understand the new cloud economic and technology opportunity, the cloud can better provide a service to the organization matched to the current planning and execution cycle, leading to a higher chance of competing and succeeding.

In this paper, we work through the details of the economics of a cloud data warehouse using a concept we think is more appropriate, Total Cost of Usage. While the goal of TCO and TCU analysis are the same, to understand the fully loaded cost of a technology, the emphasis is different. Discarding the idea of ownership helps better explain how

In this new world, the way that the data warehouse works, the way it charges for expanding usage, and the increased business value that can be obtained as part of an implementation all must be taken into account.

the money flows and how value is created. TCU is about what life is like when you have de-risked and de-owned your relationship with a data warehouse. In this new world, the way that the data warehouse works, the way it charges for expanding usage, and the increased business value that can be obtained as part of an implementation all must be taken into account.

This analysis also examines the differentiating advantages of the Actian Avalanche Cloud Data Warehouse, which is uniquely positioned to reduce total customer costs while improving performance, thanks to the Avalanche engine that delivers significantly higher performance from lower cost infrastructure. It is designed with a hybrid architecture that enables customers to run both on-cloud and on-premise to meet modern balanced requirements across multiple clouds and premises.

Our analysis will proceed in the following steps:

- First, we will examine an example of the differences between the TCO of creating a 100TB capacity using an on-premise data warehouse and the TCU of creating the same capacity using the Actian Avalanche Cloud Data Warehouse.
- Second, we will examine the properties of the Actian Avalanche Cloud Data Warehouse that have an impact on TCU.
- Finally we will look at how the hybrid use case relates to our thinking about TCU.

TCO VS TCU: A PROFORMA COMPARISON OF ON-PREMISE AND CLOUD DATA WAREHOUSES

There are a variety of motivations to use a cloud data warehouse rather than an on-premise installation, but cost tends to get the most attention. This is for good reason because the out-of-pocket costs to meet the same set of requirements using a cloud data warehouse can be a fraction of the on-premise costs.

Working through an example is the easiest way to understand why paying for usage and avoiding ownership drives costs down so dramatically.

This proforma comparison uses the following scenario:

- What is the three-year cost for supporting a 100TB data warehouse, assuming the on-premise software licenses have been paid for as a sunk cost?

Both the increased business value and the direct costs are analyzed.

The bottom line:

- The cloud data warehouse delivers a total of $5.6 million in net benefits over three years
- $1.419 million of this benefit comes from increased business performance
- $4.182 million comes from cost savings

This table shows a summary of the analysis:

The cloud data warehouse delivers a total of $5.6 million in net benefits over three years

				NPV rate	8%
$K	Year 1	Year 2	Year 3	Total	Present Value
Total Business Improvements:	**$450**	**$472**	**$496**	**$1,419**	**$1,216**
Business Performance	$450	$472	$496	$1,419	$1,216
Total Costs:	**($2,406)**	**($872)**	**($904)**	**($4,182)**	**($3,693)**
Traditional License & SW Upgrade Costs	($921)	($967)	($1,016)	($2,905)	($2,489)
Hardware Upgrade Costs	($1,566)	$0	$0	($1,566)	($1,450)
Cloud Data Warehouse Service	$333	$360	$389	$1,082	$926
IT Staffing Costs	($251)	($264)	($277)	($793)	($679)
Net Benefits:	**$2,856**	**$1,344**	**$1,400**	**$5,600**	**$4,908**
Cumulative benefit	$2,856	$4,200	$5,600	$5,600	
Breakeven Year:	year 1				

Negative figures (red) represent cost savings

On-premise data warehouses usually have worse compression of data than cloud data warehouses, so as data volumes grow, upgrades are needed more often.

Here are the details:

- The **business performance benefits** come from the ability to execute 10 more revenue generating projects per year due to reduced workloads on the data warehouse and faster handling of real-time data on a single platform. It is assumed that the $450,000 first year revenue due to faster time-to-production increases five percent per year.

Cost savings come from reductions in:

- License and upgrade fees related to the lifecycle expansion of on-premise data warehouses as well as the lack of compression that may lead to the need for more storage and capacity and increased license costs. These were $921,000 in the first year and increase 5 percent per year. The inevitable need for upgrades is driven by a variety of forces:
  - On-premise data warehouses are often based on the appliance model, which has a three- to five-year lifecycle.
  - On-premise data warehouses must be built to handle peak loads.
  - On-premise data warehouses usually have worse compression of data than cloud data warehouses, so as data volumes grow, upgrades are needed more often.
- Hardware upgrade costs are related to increasing the computing and storage capacity needed to handle the increased workloads. These amount to a capital expense of $1.566 million in the first year.
- Both for the appliance model and for commodity hardware, a refresh is required for both capacity- and lifecycle-driven upgrades.
- Cloud Data Warehouse Service costs are the fees that are paid for the use of cloud data warehouse. These are $333,000 in the first year and are predicted to increase 5 percent per year.
- IT Staffing Costs represent the savings on costs for skilled personnel for ETL, database management, and business analytics consulting and services who need to attend to the on-premise data warehouse but who are not needed for a cloud data warehouse. They are $251,000 in the first year and assumed to increase 5 percent per year.

As with any proforma analysis, the point is not that this specific scenario matches any particular company perfectly. Rather, the value lies in revealing the shape of the costs. The proforma analysis clearly shows that the increased efficiency, reduced complexity, and lower operational costs of the cloud-based solution can lead to massive savings.

Of course, the real world will never be as simple as a proforma. Most of the time, companies will find a need to have both parts of this proforma for many years to come. Hybrid models will seek to move workloads to their natural home and find ways to minimize the on-premise cost differential.

Our vision of TCU is to focus on what is present and what is missing from the on-premise vision of TCO.

TCU AS A METRIC FOR EVALUATING CDWS

Now we can use the findings from the proforma analysis as a framework for understanding what helps cloud data warehouses to drive value.

Our vision of TCU is to focus on what is present and what is missing from the on-premise vision of TCO.

The proforma analysis taught us that the value of cloud data warehouses comes from these characteristics:

- Storage Efficiency
  - Gaining an advantage in compressing data, which reduces costs for storage capacity.
- Compute Efficiency
  - Processing workloads faster with fewer computing resources.
- Compute Elasticity
  - Being able to expand or contract compute capacity as needed.
  - Being able to use as many compute engines as desired.
  - Being able to have elasticity in licensing as well as capacity.
- Operational Simplicity
  - Reducing complexity of configuration and admin leads to increase in workloads processed.

The idea of TCU is to focus on these dimensions, which highlight the changes from the on-premise era. The point of TCU is to understand your total costs, how those costs are shaped differently, and how workloads can be better served by elastic compute models.

- Operational Offloading
  - De-owning, that is, outsourcing the care of, the hardware, operating software, networking, and data center infrastructure.
  - De-owning the staff needed to manage the complexity of the hardware, operating software, networking, and data center.
  - De-owning part of the staff that managed the complexity of the data warehouse.
- Subscription Licensing
  - Replacing the large upfront costs of a perpetual software license with a subscription payment.

The idea of TCU is to focus on these dimensions, which highlight the changes from the on-premise era. The point of TCU is to understand your total costs, how those costs are shaped differently, and how workloads can be better served by elastic compute models. For example:

- Because of the elasticity of a cloud data warehouse, the system does not have to be designed for peak loads. It can grow when needed.
- Because compression reduces the storage footprint, the storage volumes grow much slower.
- Because much of the work of administration and operations is offloaded to the cloud vendor and the cloud data warehouse vendor, far fewer expert staff are required.
- Because workloads can be processed much faster, business benefits related to faster results and cost savings are achieved.

In thinking from the TCU perspective, the focus shifts from what you own to efficiency, how you are using the software, and how that usage is metered.

THE UNIQUE DESIGN AND VALUE OF THE ACTIAN AVALANCHE CLOUD DATA WAREHOUSE

From a TCU perspective, all cloud data warehouses are not the same, especially in the first four dimensions: storage efficiency, compute efficiency, compute elasticity, and operational simplicity.

There are wide variations in the performance and structure of cloud data warehouses in all of these dimensions. In this section, we will examine how the Actian Avalanche Cloud Data Warehouse measures up.

Avalanche's vectorized engine maximizes power of CPUs and allows workloads to be processed faster using fewer compute resources.

Storage Efficiency

- Avalanche has industry-leading compression that compresses data on average by a factor of 5 or 6, saving significant amounts of storage compared to traditional data warehouses.
- Avalanche can also federate queries across data in many repositories, allowing data to be stored in the optimal place for cost and performance.

Compute Efficiency

- Avalanche's vectorized engine maximizes power of CPUs and allows workloads to be processed faster using fewer compute resources.
- Avalanche is architected to support hundreds of concurrent interactive users.
- Avalanche does not require loading all data in memory.

Compute Elasticity

- Both compute and storage in Avalanche can be scaled up and down to meet the needs of workloads.
- Avalanche can also scale down to zero, allowing compute engines to be turned completely off.
- Avalanche separates storage and compute, allowing each to be sized according to need.
- Avalanche allows multiple instances of Avalanche or other compute engines to work on the same data.

Avalanche offers up to 9x cost savings over other cloud data warehouses thanks to faster compute engine performance. Faster compute means fewer compute engines are required and more workloads can be processed by a smaller set of compute engines.

Operational Simplicity

- Avalanche supports zero overhead updates.
  - Some data warehouses don't allow writes when supporting queries, forcing an update cycle in off hours.
  - Others allow writes but with slower performance or inconsistent results.
  - Avalanche allows updates to data while supporting high speed, consistent queries, without impacting performance. There is no need to design around the architecture.
- Avalanche comes with productized integration with cloud-native data sources.
  - Avalanche supports multiple popular BI and ETL tools.

Operational Offloading

- Avalanche runs on multiple cloud vendors managed and supported by Actian.
- Avalanche avoids the need for a large staff of experts to attend to the infrastructure and operations processes.

Subscription Licensing

- Avalanche is licensed on a subscription model with several tiers of capability to best fit different use cases.
- Avalanche provides complete cost predictability; there are no cost "surprises" or sticker shock when the cloud bill comes.
- Avalanche offers up to 9x cost savings over other cloud data warehouses thanks to faster compute engine performance. Faster compute means fewer compute engines are required and more workloads can be processed by a smaller set of compute engines.

LIVING IN A HYBRID MULTI-CLOUD WORLD

Actian Avalanche goes beyond current on-premise and cloud data warehouse vendors in that it offers the same software to run on-premise or in multiple clouds. This is a powerful offering when you consider that without it, an IT staff would have to learn to support multiple data warehouses in multiple hosting environments.

A January 2019 Actian survey of IT decision makers makes it clear that the trend toward hybrid is powerful:

- On-prem and cloud solutions reign. In total, 72% of IT decision makers reported that their data is stored on-prem and in one or more clouds. Migration to one platform will not be fast or easy or even optimal, meaning that we will be living in a hybrid world for a number of years.
- The desire for a hybrid environment is even higher. Some 87% agreed that they want a hybrid solution with both cloud and on-prem deployments, showing that the 72% currently using this approach will increase.
- Hybrid is preferred long-term. Some 70% said that on-prem and in single or multiple clouds is their ideal approach to computing and data analytics.
- When it comes to data storage today, almost half of IT decision makers (46%) say their data is on-premise and in multiple clouds while one-quarter (26%) say it's on-premise and in a single cloud. Almost one in four are not on-premise at all, either using a single cloud (5%) or multiple clouds (10%). More than one in 10 (13%) are on-premise only.

Actian Avalanche goes beyond current on-premise and cloud data warehouse vendors in that it offers the same software to run on-premise or in multiple clouds.

CONCLUSION

The challenge facing most companies is to:

- Make sense of their existing data warehouse landscape.
- Understand how to run it optimally with everything where it is now located.
- Then, based on a good understanding of the workloads, migrate workloads to where they will perform well for the lowest cost.

Actian Avalanche Cloud Data Warehouse offers a powerful tool in this process, one that allows workloads to be run where they need to be, not where they happen to be stuck.

This paper was written by Early Adopter Research and sponsored by Actian. Learn more about **Actian Avalanche** and **get started for free with a 30 day trial**.

Connect with us

early adopter research