

# Generating fuel for growth

Migrating legacy data warehouse  
to cloud-native data warehouses



## Beyond COVID-19: Renewed focus on velocity to insights

The pandemic accelerated the push to digital, including virtual care, remote monitoring devices and wearables. As a result, the volume, velocity and variety of health care data have exploded. So, all stakeholders had to reassess their data strategies and data technology infrastructure. They needed to accelerate speed to insights and capture additional business value in time.

Payers and providers have a renewed appreciation for the importance of high-velocity analytics and speed to insights. These help to identify new care modalities and protocols to improve outcomes while keeping the total cost of care of populations in check. COVID-19 shrank the insights window from months to weeks to days.

The rolling 7-day averages and daily trends have become mainstream for population health management and clinical care coordination to guide care decision-making. Near real-time analytics reduce friction between payers and providers while improving risk, quality and cost of care.

## What is holding health care back?

Health care organizations want to unleash the power of insights at scale. They're looking to execute targeted and timely interventions for care coordination of at-risk populations. And they need to use payment models based on risk and quality while supporting interoperability.

These initiatives call for a new foundation for a data and analytics management platform that:

- Is flexible to support different data formats and structures
- Can scale on demand for compute-intensive, time-sensitive analytic workloads
- Is secure and cost-effective

Many health care organizations are still invested in cost-prohibitive, closed-architecture, legacy data and analytics management infrastructure. Examples include Teradata, Netezza and Oracle. This infrastructure only scales vertically. That means scaling of both storage and compute at the same time. This vertical scaling comes at a steep price and require significant lead time to execute. And even with most scaled solutions, the scalability is limited to the maximum capacity licensed and available in the data center. There are no options to address on-demand, unexpected business spikes.

As a result, many health care organizations are not fully capitalizing on the power of insights that could be possible on a more flexible and cost-friendly, cloud-native data and analytics management platform.

## A new insights modernization strategy

There's good news for health care organizations with investments in legacy data and analytics management platforms. Optum has demonstrated a reduction of up to 75% operating cost is attainable when migrating an existing, legacy data warehouse (DW) such as Teradata to a cloud-based DW platform. We chose one of the newer cloud-based DW platforms for the proof of value (PoV) case study after evaluating available platform options. We'll refer to it as "the Cloud DW" in the rest of this white paper.

Depending on the existing hardware and software footprint, we believe that migrating legacy, closed-architecture data warehouse and analytics to a cloud-native platform can yield annual operating cost savings in the range of 30%–75% for comparable workloads.

In fact, the 3-year ROI from migrating to the Cloud DW from platforms such as Teradata, Netezza and the like can pay for the migration of existing data infrastructure. It can also generate surplus savings to fund modernization, transform operations and fuel other growth initiatives.

## Teradata vs. Cloud DW: Proof of value case study

For a large payer organization, Teradata is one of the many platforms to support informatics, analytics and insights for specific business lines. Over the years, several business segment-specific data marts have been ported onto Teradata, making Teradata almost a mission-critical warehouse. For these lines of business, the resulting infrastructural costs for operating a Teradata footprint for managing data and delivering insights have grown year over year to more than \$10 million per annum.



**Up to 75%  
reduction in total  
infrastructure cost**

ROI from migrating to the Cloud DW from legacy platforms can generate surplus savings to fund incremental modernization, transformation or growth initiatives.

Amid this perpetual digital data explosion, business expectations have grown for on-demand analytics at speed, scale and budget. So Optum decided to execute a comparative performance and cost assessment to confirm Teradata upgrade versus migrating to the Cloud DW and using it as the default data management and insights platform.



## Challenge

**Proof of value assessment:** Benchmark and prove that using the Cloud DW to enable the existing Teradata use cases would result in predictable performance and be more cost-effective. Utilization patterns would be consistent across Teradata and the Cloud DW environments.

Targeted objectives for PoV included:

- 1 End user queries representative of frequent workloads on Teradata
- 2 Business Objects reports performance
- 3 Data extracts performance
- 4 Scalability
- 5 Total infrastructure cost

## Proof of value summary

### Teradata upgrade considerations

Upgrading Teradata would have required additional investment due to:

- Scale-up cost of upgrading the platform to a bigger and better Teradata (also known as floor sweep cost)
- Backup and recovery (BAR) cost is an add-on cost over base "scale-up" price
- High availability/disaster recovery (HA/DR) cost, which requires a second "paid" Teradata instance
- Over-sizing Teradata to support high end of workloads
  - The fixed compute and storage architecture of Teradata requires vertical scaling to accommodate worst-case operations scenarios. **Average utilization was only 30%** of a fully scaled environment.
- Specialized administration resources (i.e., overhead) to tune environment for workloads to perform at optimal levels

### Cloud DW solution

- Cloud DW chosen based on feature-based evaluation of available options
- 10 terabytes (TB) of sample data from enterprise data warehouse (EDW) moved from on-premises to the Cloud DW
- Minimal modification to DDL to create same data structures in the Cloud DW
- Workloads executed from on-premises laptops connecting to the Cloud DW performance results

### 1. End user queries

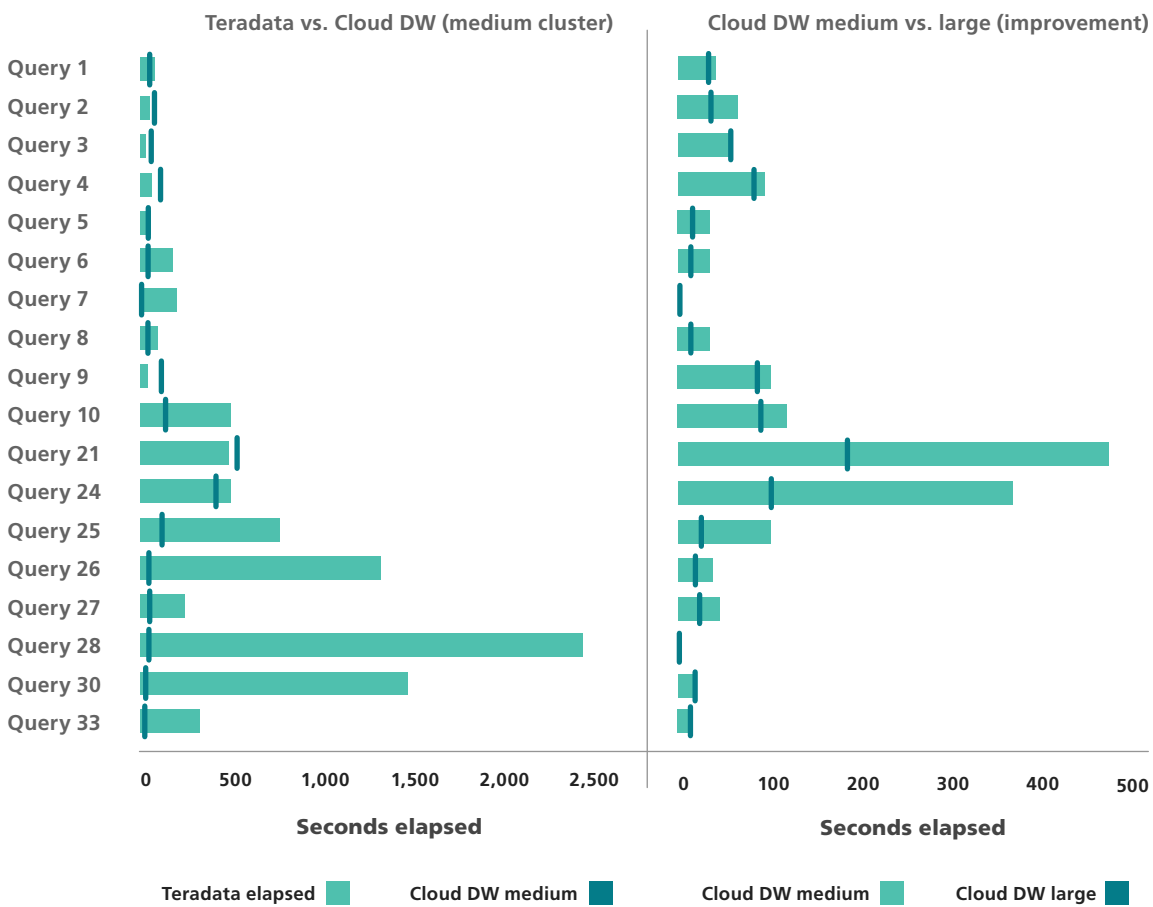
A subset of end user queries often submitted on Teradata were run on both Cloud DW and Teradata to measure and compare results.

- **No changes** to SQL syntax needed to run existing queries on Cloud DW
- Queries ran **as well or better** on the Cloud DW using a medium compute cluster
  - Use of large compute cluster improved response even further



Queries ran **as well or better** on the cloud DW using a medium computer cluster.

### End user query results (seconds)



#### Average query time

Teradata: **432.55 seconds**

Cloud DW medium: **91.11 seconds**

#### Average improvement moving from

**medium to large configuration in Cloud DW**

Cloud DW medium: **91.11 seconds**

Cloud DW large: **45.81 seconds**

## 2. Business Object reports

Same reports run multiple times

- **No change** to reports needed to run against the Cloud DW
- **Consistent response time** for the Cloud DW
  - Response time varied on Teradata due to queuing

### Business objective performance test comparison

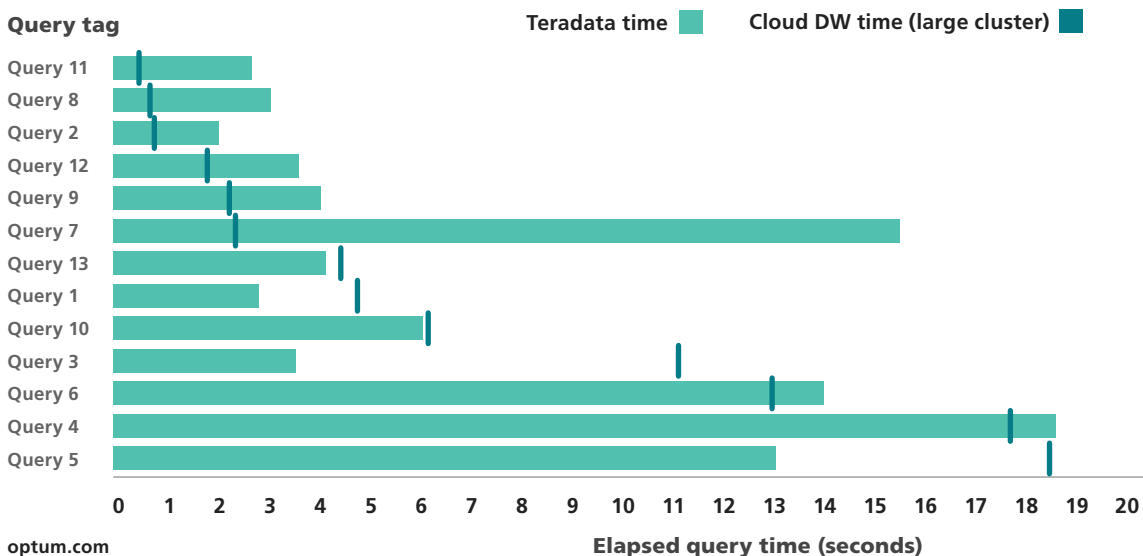


## 3. Data extraction

The existing Teradata extraction processes involve the extraction of 2TB of data from 100+ base tables. The same queries were executed non-optimized on the Cloud DW and the extract performance compared well relative to Teradata.

- **No changes** to SQL syntax needed
- **Extractions ran faster** or as good as Teradata

### EPR query times



#### 4. Scalability

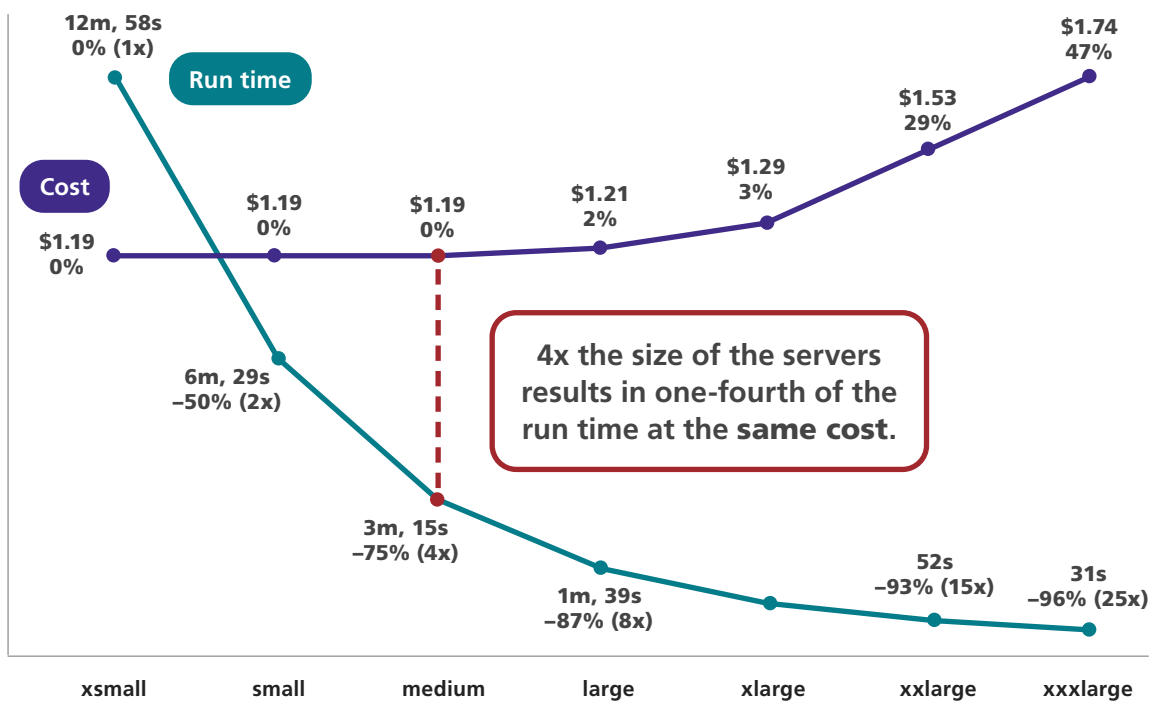
Since Teradata expansion would require vertical scaling, a vertical scalability test was done for the Cloud DW by loading 1 billion records for varying compute capacity.

- **Compute cost** is somewhat flat as the VW size is increased 4x (from XSmall to Medium)
- 75% reduction in overall run time for the same workload for flat compute cost promises **increased analytical scalability at the same cost**



**75% reduction** in overall run time for the same workload

#### Scalability



#### 5. Total infrastructure cost

Comparing the 90-day workload utilization on Teradata versus the Cloud DW, the benchmark study identified an estimated reduction in total cost of infrastructure in the range of 70%–85%. Key factors influencing the lower Cloud DW infrastructure cost are attributable to:

- **Cloud-native architecture**
- **3x–8x data compression** (PoV demonstrated 5x storage reduction due to Cloud DW's smart compression)
- **No additional infrastructure** or cost for backup and recovery.
  - Cloud DW data automatically backed up and accessible for 90 days
- **High availability/disaster recovery (HA/DR)** — only pay for storage cost (low due to high data compression) and egress until consumption compute is needed
  - Data can be automatically replicated to different cloud platform regions or to different cloud platforms entirely



**3x–8x data compression**

**No additional infrastructure or cost**

### PoV conclusion

1. The Cloud DW offers a compelling **high-performance and scalable platform** with a very attractive runtime infrastructure cost.
2. A 3-year ROI view offers substantial reduction in annual infrastructure cost. It makes the legacy migration cost-neutral with a possibility for surplus savings generated to invest back in other priorities and growth initiatives.
3. A fully SQL-compatible Cloud DW architecture offers **backward compatibility**. It also **lowers the migration risk** somewhat for legacy data warehouse migrations should the desire be to simply “lift and shift” with minimal, must-have modifications made to data architecture, ETL and reporting.

## Conclusion

Our PoV demonstrates that health plans (and other health care entities) vested in legacy data technology infrastructure have a viable, cost-effective path to migrate to a cloud-native DW platform. This platform provides on-demand scalability and a data architecture that reduces the data maintenance activities and related cost by design.

Building on our Teradata to Cloud DW PoV pilot, it's evident that a 3-year business case (typical across health care) could:

- Pay for migrating legacy data warehouse to cloud-native DW platform
- Generate additional savings to invest in other growth, margin and experience initiatives

A “simple backward compatible migration” to the Cloud DW may be a preferred first step for some health plans. But Cloud DW offers many rich architecture options. These include support for semi-structured data (such as JSON, AVRO and Parquet) and access to complementary data management tools (for ETL, Catalogs, Governance, etc.).

Optum recommends evaluating the enterprise data strategy to broadly align with in-flight and planned business transformation initiatives. That way, you'll get the most value from your Cloud DW and analytics investment in a planned manner.

If planned and executed right, migration to the Cloud DW could yield the targeted business case ROI and facilitate the organization's journey to become a data-driven organization. Like any other on-demand, subscription-based cloud infrastructure, migration to the Cloud DW requires organizations that proactively and preemptively strategize their change management and platform usage governance framework.

Optum recommends starting with the assessment of type of usage patterns in the existing enthronelement, workloads to be moved and imminent business needs requiring insights at scale. Additionally, we recommend planning early for a governance infrastructure with a focus on measuring, monitoring and utilization management to get the most value from Cloud DW investments in a cost-effective manner.

## About Optum Advisory Services

Optum Advisory Services can help your organization accelerate and optimize your Cloud DW adoption journey. We have the people, processes and technology to support your legacy data warehouse modernization journey. This includes strategic assessment, planning, migration and optimization of Cloud DW utilization after launch. We have the maturity of experience to meet health care clients where they are in the legacy data modernization journey. We can partner with you to help you get the most value from data modernization investments.

## Meet our experts



### John Shin

Principle Architect/Director  
Optum Advisory Services  
Email: john.shin@optum.com



### Dinesh Malhotra

Vice President  
Optum Advisory Services  
Email: dinesh.malhotra@optum.com

---

Learn how Optum® Advisory Services can help you determine and reach your organization's goals.



Learn more

### Contact us:

empower@optum.com

1-800-765-6807

optum.com/advisoryservices

---



11000 Optum Circle, Eden Prairie, MN 55344

Optum is a registered trademark of Optum, Inc. in the U.S. and other jurisdictions. All other brand or product names are the property of their respective owners. Because we are continuously improving our products and services, Optum reserves the right to change specifications without prior notice. Optum is an equal opportunity employer.

© 2022 Optum, Inc. All rights reserved. WF6327062 03/22