



Executive summary



Faster data analysis

Analyze data 3.8x as fast as a comparable Amazon EC2 m5.2xlarge instance



Better return on investment

22% lower three-year TCO for 24/7 usage vs. similarly configured Amazon EC2 m5.2xlarge and g4dn.xlarge instances



Support a variety of workloads

APEX Private Cloud simultaneously supported both database and GPU ML workloads

Deliver better return on investment and faster data analysis while performing image classification tasks with Dell Technologies APEX Private Cloud compared to a set of comparable Amazon EC2 instances

Cities around the world are looking to data to solve serious transportation problems.¹ Data from public transportation, traffic signals, and surveillance cameras can improve traffic, for example. Furthermore, having robust and reliable hardware resources that can process and store all this information can enable officials to plan resources based on the behavior of commuters and tourists, help people to better navigate and engage a city's important landmarks, help manage crowds during large events, and more.

We tested two cloud solutions—one private and one public—to see how they might handle workloads representative of “smart city” transportation projects. Here's what we found:*

- We calculate that the APEX Private Cloud solution we tested would cost 22 percent less over three years compared to a similarly configured Amazon EC2 solution. These savings could allow municipalities to invest in innovative projects or devote more funds to the departments that need them.
- An APEX Private Cloud VM required just 10.5 hours to complete an online analytical processing (OLAP) workload. A comparable Amazon EC2 m5.2xlarge instance needed 40 hours for the same workload. Reducing OLAP data analysis time could help cities make key decisions sooner.
- On a separate VM, the APEX Private Cloud solution had a slightly better query processing rate for a machine learning (ML) workload compared to an Amazon EC2 g4dn.xlarge instance.
- We scaled up the APEX Private Cloud solution to have twelve VMs doing OLAP work and nine VMs doing ML simultaneously. It completed all twelve OLAP workloads within 17 hours, and the ML processing rate for each VM was on par with the single-VM tests. Predictable performance at scale can help cities analyze data more efficiently.

*For full specs and detailed cost calculations, see the [science behind the report](#).

Faster data analysis

To test the data analytics capabilities of each environment, we used an APEX Private Cloud database VM with eight cores, 32 GB of memory, and 2.5 TB of storage to house a large OLAP database. We configured a general purpose Amazon EC2 m5.2xlarge instance with comparable specifications.*

While the APEX Private Cloud VM took 10.5 hours to analyze a 1TB dataset, the Amazon EC2 instance required 40 hours—a full workweek—due to EBS quotas and limitations.

Time to complete data analysis workload Lower is better


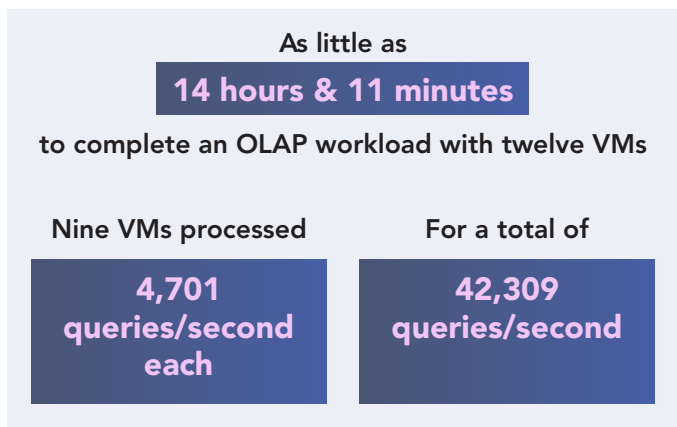


Figure 1: The total time that a single database VM of the APEX Private Cloud solution and the EC2 m5.2xlarge instance needed to complete the data analysis workload we used in testing. Lower is better. Source: Principled Technologies.

Scale VMs to support multiple complex workloads

To test each solution's ability to scale, we increased the number of database VMs/instances to twelve and ran a machine learning (ML) workload at the same time. When we increased the number of VMs to twelve and added a concurrent machine learning workload, the APEX solution showed no signs of ML performance degradation. It completed the OLAP workload in 17 hours, and performed an average of 4,701 queries per second on all VMs for the ML workload.



*For full specs on the solutions we tested and a full breakdown of our calculations for the TCO, see the [science behind the report](#).

Spend less over three years

We calculated a three-year total cost of ownership (TCO) for the APEX and Amazon solutions we tested. We estimate that the three-node, GPU-enabled APEX solution would cost \$311,863.86 over three years—22 percent less than the \$402,272.28 we calculated for the Amazon solution.*

Estimated TCO in USD over three years Lower is better

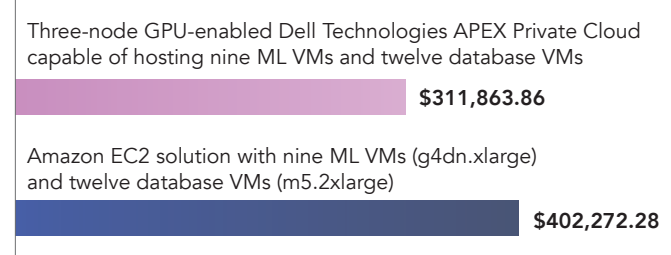


Figure 2: The TCO, in US dollars, for the two solutions we tested over three years. Lower is better. Source: Principled Technologies.

Conclusion

A fast cloud solution can help officials make timely decisions and, in the end, help residents and tourists navigate a city. In our testing, a database VM on APEX Private Cloud completed a data analysis workload in 10.5 hours, nearly a quarter of the time that an Amazon EC2 m5.2xlarge instance needed to run the same data analysis workload on the same dataset. APEX Private Cloud can also offer a better investment. We found that when maximizing resource availability, the APEX solution could deliver a 22 percent lower three-year TCO compared to a similarly configured Amazon EC2 solution of twelve general purpose m5.2xlarge instances for OLAP workloads and nine g4dn.xlarge instances for ML workloads, all at 24/7 resource utilization per month.

1 Otonomo, "8 Smart Cities Lead the Way in Advanced Intelligent Transportation Systems," accessed September 22, 2021, <https://otonomo.io/blog/smart-cities-intelligent-transportation-systems/>.

Read the report at <https://facts.pt/IOpaHrs>



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the report.