

**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**

Deploying APEX Private Cloud with NVIDIA T4 GPUs can offer a lower three-year total cost of ownership (TCO) and faster data analysis compared to a set of comparable Amazon EC2 instances

Cities around the world are looking to data to solve serious transportation problems.<sup>1</sup> 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. To assess the performance of both solutions with two different workloads, we first deployed a database VM on-premises using a 1TB dataset on APEX Private Cloud. That database VM completed our online analytical processing (OLAP) workload in 10.5 hours. We then deployed a comparable Amazon EC2 m5.2xlarge instance and configured it similarly to the on-premises APEX Private Cloud database VM with the same 1TB database. The EC2 m5 instance needed 40 hours to complete the same OLAP workload. Being able to reduce data analysis time could help smart cities analyze data more quickly. That, in turn, could help put analysis in the hands of decision makers sooner to save time for commuters and tourists.

With the additional available hardware resources of the APEX Private Cloud solution, we also ran a machine learning (ML) workload of image classification tasks on a separate VM. We compared its performance to a similarly configured Amazon EC2 g4dn.xlarge instance that ran the same ML workload.

We also wanted to see if the APEX Private Cloud solution could handle the test workloads at scale. We set up twelve database VMs, each with a 1TB dataset, to maximize the available cluster CPU utilization. The APEX solution we tested had nine NVIDIA® T4 GPUs, and we used them to set up nine ML VMs capable of performing image classification tasks.

While running the database and ML workloads simultaneously and at scale, the twelve APEX Private Cloud database VMs completed the OLAP workloads in 17 hours. Scaling the image classification workload did not affect the VM query processing rate; the performance was on par with our single-VM ML tests. Predictable performance when scaling workloads could also help smart cities analyze data more efficiently and reduce the need to purchase additional compute resources.

In addition to the performance advantage, APEX Private Cloud could offer a better value than the EC2 solution. For our three-year total cost of ownership (TCO) scenario, we assumed the two solutions we tested had the same number of VMs (or instances for the Amazon solution) running 24/7. We estimate the APEX Private Cloud solution would cost \$311,863.86, which could save a municipality USD \$90,408.42 (or 22 percent) compared to the estimated \$402,272.28 in three-year costs for the EC2 solution (twelve EC2 m5.2xlarge instances and nine EC2 g4dn.xlarge instances in that scenario). Those savings could allow municipalities to continue investing in innovative projects or devote more funds to other departments.

## A brief overview of Dell APEX

Dell APEX delivers cloud resources as-a-service, wherever an organization needs them. We chose a co-location cloud option for this study, but organizations can choose to deploy APEX offerings in their data center, a co-location facility, or an edge location. IT staff configure and manage the lifecycle of their APEX solutions through the Dell Technologies APEX console, a unified web interface, which we used in this study.<sup>2</sup>

For more information about APEX, visit <https://www.delltechnologies.com/en-us/apex/index.htm>.

## APEX Private Cloud

APEX Private Cloud delivers virtualized compute and storage for on-premises and edge locations. Dell claims this service can “simplify your IT experience,” “accelerate your hybrid cloud adoption,” and “drive modernization” in the private cloud. According to the website, APEX Private Cloud works as a single platform for both traditional and containerized apps.<sup>3</sup>

## Complete data analysis in less time while running image classification workloads

The APEX Private Cloud solution we tested came equipped with three VMware® vSphere® nodes, nine NVIDIA T4 GPUs, 144 Intel Cascade Lake-based CPU cores, 2,304 GB of RAM, and 40 TB of usable vSAN storage. We configured one database VM with eight cores, 32 GB of memory, and 2.5 TB of storage to house a large OLAP database, database backup, logs, and operating system.

We also deployed a general purpose EC2 m5.2xlarge instance with eight cores, 32 GB of memory, and 2.5 TB of EBS gp2 storage. For more details on how we configured our test VMs and instances, see the [science behind this report](#).

The database VM on the APEX Private Cloud solution took 10.5 hours to analyze around 1 TB of data. That was 73.5 percent less time than the Intel Cascade Lake-based EC2 m5.2xlarge instance, which needed 40 hours to analyze the same 1 TB of data, due to EBS quotas and limitations. The sooner your OLAP applications can complete their workloads, the sooner decision makers can make evidence-based decisions to improve essential areas of the city, such as traffic flow and transit access.

Figure 1 shows the total time to complete the OLAP workload for the database VM of the APEX solution and the EC2 m5 instance.

### Time to complete data analysis workload



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.

## Run complex image classification tasks

On the APEX solution, we deployed one ML VM with one NVIDIA T4 GPU, four CPU cores, 16 GB of memory, and 1 TB of storage to store the image dataset. The inference ML workload processed images of different vehicles commonly found in cities (buses, cars, bikes, and others).

We also deployed an Intel Cascade Lake-based EC2 g4dn.xlarge instance with one NVIDIA T4 GPU, four CPU cores, 16 GB of memory, and 1 TB of storage. We ran the same image classification workload with the same dataset on the EC2 g4 instance that we ran on the APEX Private Cloud. As Figure 2 shows, the EC2 g4 instance handled slightly more queries per second (3 percent more).

### Queries per second



Figure 2: The queries per second that the machine learning VM of the APEX Private Cloud solution and EC2 g4dn.xlarge instance handled during the image classification workload we used in testing. Higher is better. Source: Principled Technologies.


# Scale VMs to support multiple data analysis and ML workloads with APEX Private Cloud

To assess OLAP performance at scale, we increased the number of database VMs on the APEX Private Cloud cluster to twelve to consume 96 CPU cluster cores. As Figure 3 shows, APEX Private Cloud with twelve VMs running concurrently took 17 hours to analyze 12 TB of data.

Time to complete data analysis workload at scale on APEX cluster



Figure 3: The total time, in hours and minutes, that the twelve database VMs needed to complete the OLAP workload we used in testing on the APEX solution. Source: Principled Technologies.



To test performance at scale of a mixed-workload environment, we also increased the number of ML VMs on the APEX cluster to nine to consume a total of 36 CPU cluster cores and all nine NVIDIA T4 GPUs. This time, we ran the hour-long image classification workload three times while the database data analysis workload ran in the background. As Figure 4 shows, all nine ML VMs on APEX Private Cloud supported 42,309 queries per second, or 4,701 queries per second per VM. Putting the cluster under heavy resource utilization while concurrently supporting twelve database VMs and nine ML VMs increased the database analysis time by less than 62 percent (10.5 hours vs. 17), still outperforming the single EC2 m5.2xlarge instance processing a single 1TB dataset.

These tests show that compared to our single-VM tests, the APEX solution showed no signs of ML performance degradation after we scaled up the number of ML VMs it supported and ran the data analysis and ML workloads together.

**Number of queries per second at scale on APEX cluster**

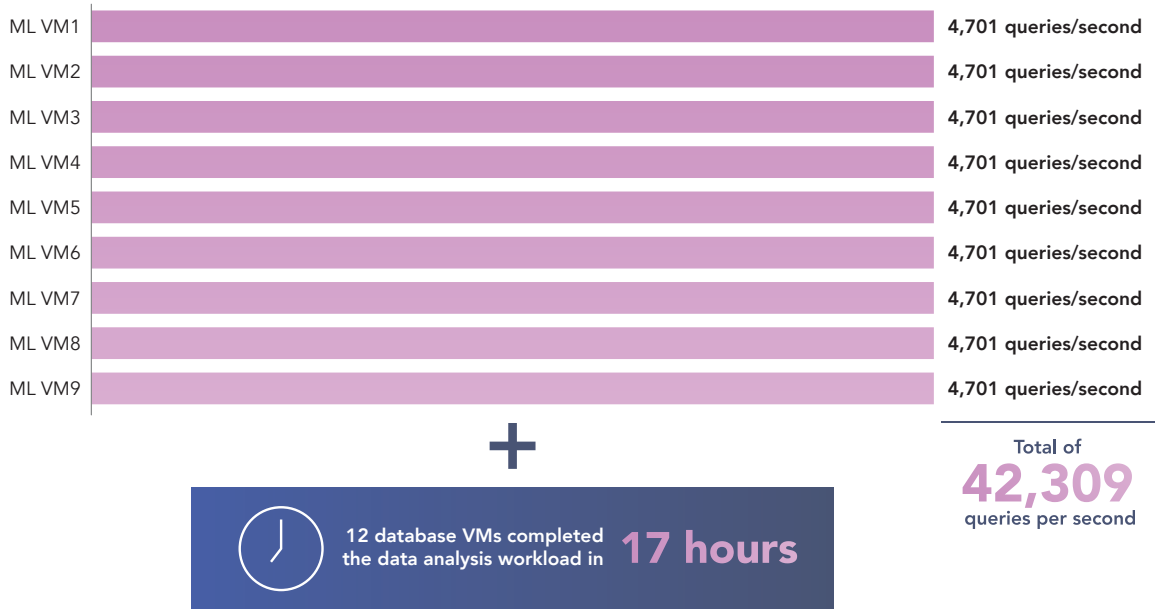


Figure 4: The queries per second that the nine machine learning VMs of the APEX Private Cloud solution handled during the image classification workload we used in testing. Source: Principled Technologies.



# Spend less over three years utilizing available resources with APEX Private Cloud

We assumed Amazon EC2 instances run independently (i.e., using unshared resources) and that those instances or VMs deliver predictable performance, regardless of instance scale. As such, we assumed that testing independent Amazon EC2 m5.2xlarge and independent EC2 g4dn.xlarge instances presented their performance in the best light. We then used those assumptions to estimate the performance for a scaled-up set of EC2 m5.2xlarge and g4dn.xlarge instances (twelve and nine instances, respectively), thus giving us a fair base performance that justified our cost comparison to the APEX Private Cloud solution, whose VMs ran on a single cluster with shared resources.

We calculated a three-year TCO analysis for the APEX Private Cloud solution we tested based on 24/7 resource utilization to account for systems that need to process and storage data around the clock. The APEX Private Cloud cluster we tested had 144 Intel Cascade Lake-based CPU cores, 2,304 GB of RAM, and 40 TB of usable vSAN storage, which allowed us to deploy twelve database VMs and nine ML VMs. To exactly match what could be deployed on the APEX Private Cloud solution at nearly 100 percent CPU, memory, and storage resource allocation, we also calculated a three-year TCO for an Amazon solution with twelve EC2 m5.2xlarge instances and nine EC2 g4dn.xlarge instances with 24/7 resource utilization per month. We found that APEX Private Cloud could offer a 22 percent lower TCO compared to the similarly configured Amazon EC2 solution of twelve general purpose m5.2xlarge instances for data analysis workloads and nine g4dn.xlarge instances for ML workloads. Note that we only tested performance on two individual EC2 instances, but we believe performance would increase linearly regardless of the number of instances under test. We used a three-year pricing model with no money up front for both solutions. We also did not include database software costs that would be the same for both solutions. For more pricing and configuration information for this three-year TCO scenario, see the [science behind this report](#).

Figure 5 shows our estimated costs for the two solutions for three years. Costs of your solution may vary depending on usage.

### Estimated TCO in USD over three years

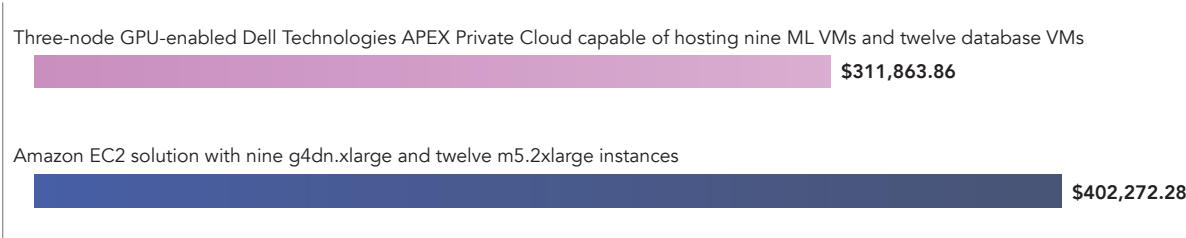


Figure 5: The TCO, in US dollars, for the two solutions we tested over three years. Lower is better. Source: Principled Technologies.

## Smart cities transportation use cases

For those unfamiliar with smart cities, the term describes urban centers that use “sensors and connected devices to collect and analyze data.”<sup>4</sup> Smart cities (or surrounding municipalities such as counties) can then use the data to improve city operations, manage resources, or make the daily routines of residents and tourists better.

Improving traffic congestion and mobility is just one example of how smart cities can both boost efficiency in city operations and improve residents’ daily lives. (Mobility includes “everything from traffic management to public transportation to infrastructure systems.”<sup>5</sup>)

As previously noted, cities around the world use data to solve traffic and mobility problems. One key example involving image recognition comes from New York City. The system there can “detect traffic jams, weather patterns, parking violations, and more, sending real-time alerts to city officials along the way.”<sup>6</sup>

Fueled by data, smart cities could work to solve transportation problems by building:

- **Adaptive traffic signals:** In a blog post from April 2021, NEC notes, “[t]raditional traffic signals use timers and inductive loops to determine when the light should change. While this is effective, it doesn’t allow cities to adjust the signal cycle when there are changes in traffic levels, leading to congestion.”<sup>7</sup> Adaptive traffic signals could allow municipalities to make those kinds of adjustments in near real-time.
- **Parking systems:** Cameras could scan and read license plates to obtain type, make, and model information of parked vehicles.<sup>8</sup> All that information could then help identify motorists and measure the time spent in a parking spot, for example. The technology could also make it possible for motorists to learn the location and number of available parking spaces near them.
- **Smart corridors:** These are portions of roads that use technologies to monitor traffic conditions, including vehicle collisions, weather events, and other obstacles.<sup>9</sup> Once they receive the data, motorists can then use the information to plan their commutes.





## Conclusion

Data and technology can offer many paths to solving commercial and localized transportation problems that stem from urban density. 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. We also increased the number of database VMs running the data analysis workloads to twelve and the number of ML VMs running the image classification workloads to nine to show how the APEX Private Cloud solution could handle a mixed-workload environment at scale.

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/>.
- 2 Dell Technologies, "Dell EMC VxRail Hyperconverged Infrastructure," accessed September 22, 2021, <https://www.delltechnologies.com/en-us/converged-infrastructure/vxrail/index.htm>.
- 3 Dell Technologies, "APEX Private and Hybrid Cloud," accessed February 25, 2022, <https://www.dell.com/en-us/dt/apex/cloud-services/private-hybrid-cloud.htm>.
- 4 Dataversity, "How Big Data Impacts Smart Cities," accessed September 22, 2021, <https://www.dataversity.net/how-big-data-impacts-smart-cities/#>.
- 5 Otonomo, "8 Smart Cities Lead the Way in Advanced Intelligent Transportation Systems," accessed September 22, 2021.
- 6 DZone, "Smart Cities and Image Recognition," accessed September 22, 2021, <https://dzone.com/articles/smart-cities-and-image-recognition>.
- 7 NEC NZ, "5 examples of smart city transportation solutions," accessed September 22, 2021, <https://www.nec.co.nz/market-leadership/publications-media/5-examples-of-smart-city-transportation-solutions/>
- 8 Ni, X., Huttunen, H., "Vehicle Attribute Recognition by Appearance: Computer Vision Methods for Vehicle Type, Make and Model Classification," accessed September 22, 2021, <https://link.springer.com/article/10.1007/s11265-020-01567-6>.
- 9 NEC NZ, "5 examples of smart city transportation solutions," accessed September 22, 2021.

Read the science behind this report at <http://facts.pt/vQfSwqS> ►



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 science behind this report.

This project was commissioned by Dell Technologies.