simactive
HIGH-END PHOTOGRAMMETRY SOFTWARE

PHOTOGRAMMETRY ON THE CLOUD WITH CORRELATOR3D™

GET IN TOUCH
www.simactive.com
contact@simactive.com

1 Introduction

Although cloud computing has been growing in importance in most fields, it only recently started being looked at as a viable option for photogrammetry. SimActive has designed Correlator3D to be as easy to use on the cloud as it is on desktops. This guide will look at the advantages as well as some of the challenges of processing imagery with virtual machines (VMs) on the cloud.

2 What is cloud processing and how does Correlator3D work on the cloud?

Cloud computing is the on-demand availability of IT resources over the internet. Here instead of buying hardware, like computers, and storage servers, these resources are accessed, on demand from data centers via the internet. Resources are paid for only during their time of use. In the case of Correlator3D, it must be installed on a compatible GPU enabled Windows virtual machine (VM) and is accessed via a remote desktop environment. It behaves just like when a user connects to an office computer from home, except here the computer is sitting in a data center hosted by a cloud provider like Amazon's AWS or Microsoft's Azure. Once set up, the user uploads images and processes them with Correlator3D just like on a physical workstation.

3 Supported virtual machine types

Since Correlator3D requires a computer with a GPU card running Microsoft Windows 10 or above, the VMs selected must also be similarly GPU enabled. Examples of VM types based on commonly used cloud providers are given in Table 1. As Correlator3D uses only 1 GPU, there is usually no need to choose more expensive VMs with multiple GPUs.

CLOUD PROVIDER	SERIES	VIRTUAL MACHINE TYPE	vCPUs	PROCESSOR	MEMORY (GiB)	GPU
AWS	G4 Series	g4ad.xlarge	4	AMD EPYC 7R32	16	AMD Radeon Pro V520
		g4dn.4xlarge	16	Intel Xeon Cascade Lake	64	NVIDIA T4
	G5 series	g5.xlarge	4	AMD EPYC 7R32	16	NVIDIA A10G
		g5.8xlarge	32	AMD EPYC 7R32	128	NVIDIA A10G
Azure	NV series	NV4as_v4	4	AMD EPYC 7V12	14	1/8th AMD Radeon Instinct MI25
		NV6ads_A10_v5	6	AMD EPYC 74F3V	55	1/6th NVIDIA A10G
		NV36ads_A10_v5	36	AMD EPYC 74F3V	440	NVIDIA 10G
	NC series	NC4as_T4_v3	4	AMD EPYC 7V12	28	NVIDIA T4

Table 1: *Virtual machines commonly used for photogrammetry processing with Correlator3D in AWS and Azure*

4 Benefits and challenges of photogrammetry on the cloud

BENEFITS	CHALLENGES
Secure: Industry standard security at the level of the cloud provider. At the user level, security governed by firewall rules set up by the user to allow user-defined network activity.	**Learning curve:** Setting it up the first time is challenging. Understanding what VM to use, where to locate it, how much storage required are a few things to consider.
Instant availability: VMs up and running in a few minutes. VMs can be created, projects processed and the VM shutdown / deleted quickly using scripts.	**Complex costing:** Many variables involved, each with an associated cost that can eat into the budget.
Reliable: Strong service level agreements from providers promising uptimes of 99.5%.	**GPU installation:** Depending on the provider and instance chosen, installing the correct GPU driver can be an involved process.
Scalable: Quickly increase VM processing power / storage space when needed. Or increase the number of VMs to enable distributed processing. *	**Access:** First-time users may not have access to the desired VM required for their projects. A request for a quote increase solves this in usually 48 hours.
No capital expenses: Pay only for what is used and for the duration resource is used / requisitioned.	
Easily quick data delivery: High bandwidth speeds at cloud provider level allow for quick delivery of results to clients.	
Archival: Raw data and results archival is quick to setup, economical, secure and reliable.	

**Each VM involved in distributed processing must possess a Correlator3D license.*

5 Will cloud processing result in time gains?

If the user does not have access locally to a powerful workstation, then time gains may be obtained. However, it does depend on project size and internet bandwidth. The time it takes to upload a couple of thousand images to a cloud VM and to download results is insignificant.

On the other hand, a project with 10,000 images @42 MP weighs upwards of a 150 GB. If processed without any compression, the resulting orthos would exceed 2 TB. Uploading and downloading these on a slower internet connection will decrease any time gained using high-end instances on the cloud for photogrammetry processing.

Users should also consider the distributed processing capabilities of Correlator3D. These can play a crucial role in significantly reducing the processing times of large projects by distributing portions of the project to multiple VMs, and then assembling the results automatically.

Furthermore, with cloud delivery, the inherent high bandwidth available with cloud instances can be advantageous. Users can send deliverables directly to a client without having to download the data locally first.

Certain Correlator3D modules like ortho generation are more CPU intensive while others, like the DSM generation module, are more GPU intensive. As shown in Figure 1, VMs with higher vCPUs will generate orthos faster than a VM with fewer vCPUs. However, speed gains will start to drop beyond 20 vCPUs.

With certain cloud providers, it is also important to ensure storage drives have high IOPS which are essential for writing orthos and DSM tiles quickly to the disks. Lower IOPs will result in longer processing times for these modules.

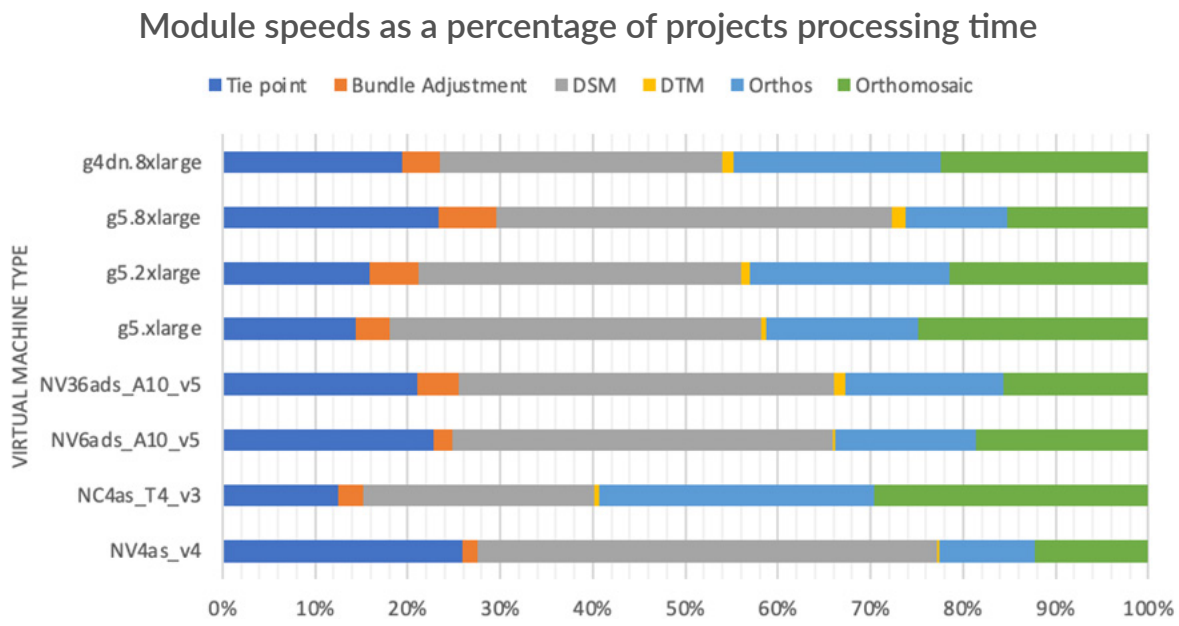

Figure 1: *Processing time taken per module by various virtual machines expressed as a percentage of total project processing time (g4 and g5 series from AWS, NV and NC series from Azure)*

6 Costs involved

Virtual machine (VM) prices depend on factors like number of vCPUs, amount of RAM, attached storage, GPU and expected network performance. However, the least expensive VM is not the best option since storage costs are charged at a fixed rate for the duration of the requisition period. Slower virtual machines equate to storage volumes being kept for longer periods.

In Figure 2, the horizontal axis shows commonly used virtual machines from Azure and AWS and compares them to two local machines, a basic gaming laptop and an i9 13th Gen workstation with 64 GB RAM and an NVIDIA RTX 4000 series GPU. The left vertical axis shows the time taken in hours while the right vertical axis shows the price of the VM.

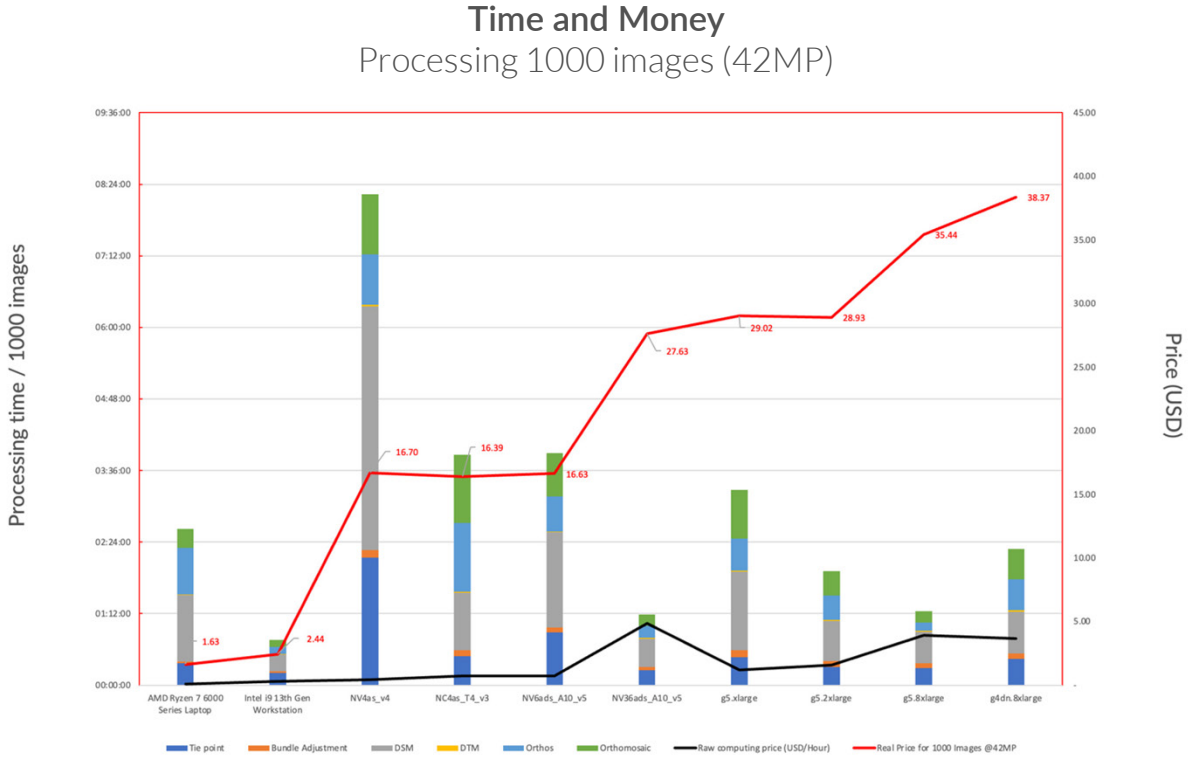

Figure 2: *Time taken to process 1,000 images (@42 MP/image) by common VMs and the associated costs involved*

While similar virtual machines between cloud providers are closely matched in price / hour, the real price can vary significantly depending on the costs of transferring results out of the VM. The black line is the price per hour of each VM. However, it really does not give as complete a picture as the red line which is the real price that takes into consideration the price of the VM for the duration of processing a 1,000 image project (@42 MP / image), the time it takes to clean the project, the storage price for the duration of the project and the price to transfer the result out of the VM. Using this 1,000-image example, using a VM like the NV36ads_A10_v5, the project could be completed in just over an hour at a total cost of under 30 USD.

Figure 2 also shows that, often smaller VMs of newer series outperform an older, larger series and are cheaper as well. For example, the older but larger g4dn.8xlarge with 32 vCPUs (Intel Xeon) running at 2.5GHz with 128 GB RAM with 16GB video memory on an NVIDIA T4 tensor core GPU takes almost twice the time to process 1,000 images as the newer, smaller and lower priced g5.2xlarge which has 8 vCPUs (AMD EPYC 7R32) running at 2.8 GHz with only 32 GB RAM and 24 GB video memory on an NVIDIA A10g GPU card.

7 Conclusion

If the modest learning curve can be surmounted and internet connectivity does not pose a challenge, processing imagery on the cloud presents a powerful, secure and economically viable option. With variety of options, there is a virtual machine to suit most project budgets and timeframes. As Correlator3D is equally intuitive to use on the cloud as it is on desktops, users who prefer not to invest in expensive hardware will find using Correlator3D on the cloud an attractive option.