

Migration and performance of CMAQ & WRF-CMAQ in the public cloud with commercial images

Arturo Fernandez

odyhpc

20th Annual CMAS Conference

November 1-5 2021 (virtual)

Outline

- Public cloud environments
- Running HPC apps (CMAQ) in the public cloud
- Integrated product for CMAQ
- Performance evaluation
- Cost analysis
- Final remarks



Public Cloud Environments

What is the public cloud? → Rented IT (a la carte) from Cloud Service Providers (CSPs)

Pros

- Flexibility
- No upfront cost
- Unlimited storage space and other resources
- Availability of new hardware every so often
- What about performance?

Cons:

- New skills (very basic to highly specialized)
- Everything is billed
- Lower performance for HPC workloads?



Running HPC apps (CMAQ) in the public cloud

How do you run HPC apps in the public cloud?

- ❖ Choose the right IaaS
- ❖ Compile and build the app

What is the right IaaS?

There is no universal answer as each case is unique. In general, hungrier computational resource apps require more resources → Major CSPs

How do you compile and build the app?

- Start from scratch → Challenge: The hardware and software must work together
- Use integrated product



Integrated product for CMAQ

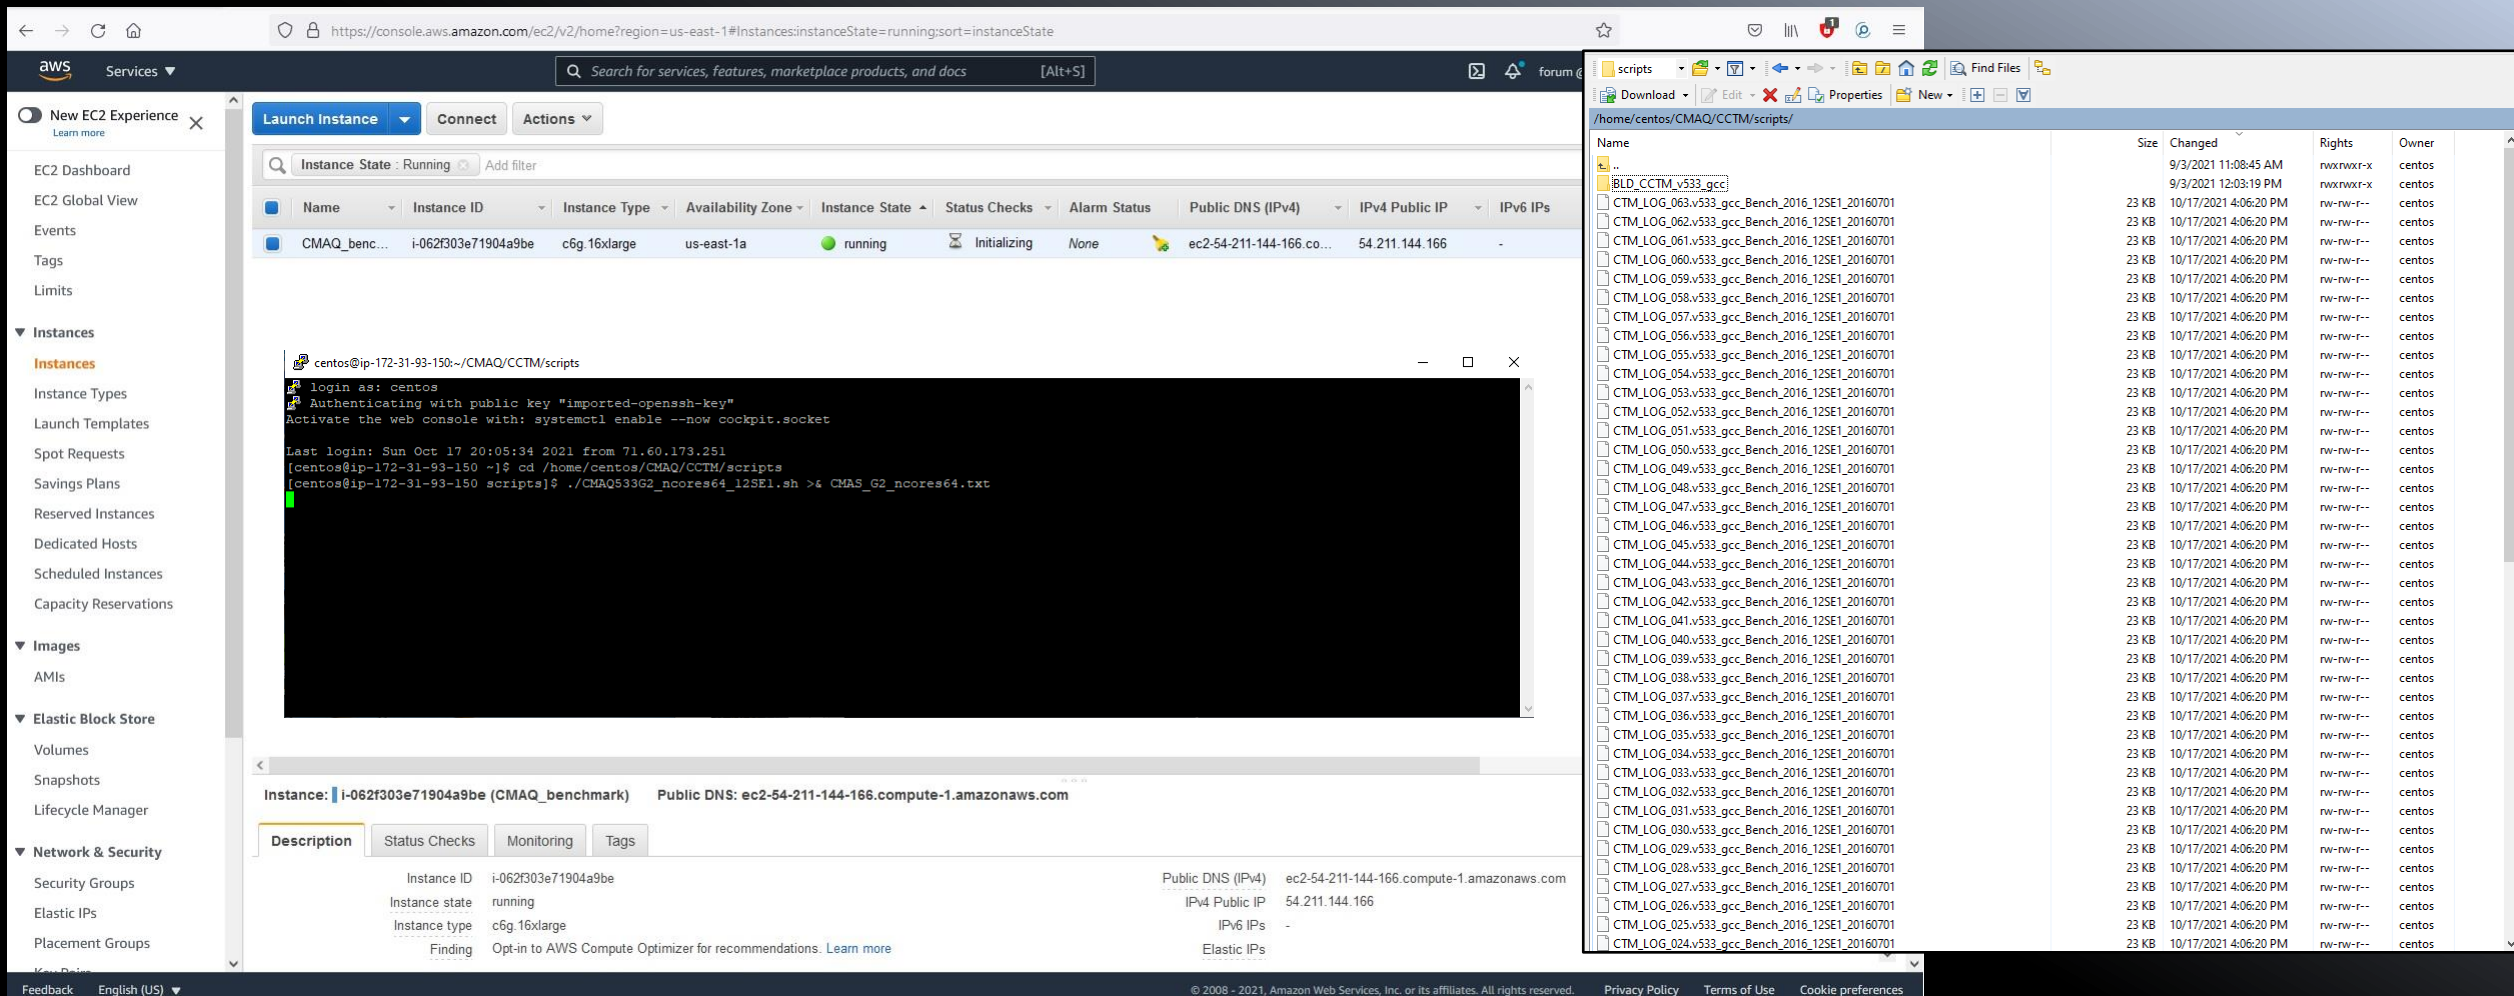
Integrated product: CMAQ & WRF-CMAQ

- It has precompiled executables optimized for public cloud infrastructure or Infrastructure as a Service (IaaS)
- Currently available from the AWS Marketplace (potential future availability from the Azure Marketplace)
- Dual functionality for single instances and clusters
- It also has postprocessing apps, which can even be used with instances with GPU capabilities



Integrated product for CMAQ

Launch and connect to an instance running CMAQ



The screenshot displays the AWS Management Console interface. The top navigation bar shows the AWS logo and a search bar. The left sidebar contains navigation links for EC2 Dashboard, EC2 Global View, Events, Tags, Limits, Instances, Images, Elastic Block Store, and Network & Security. The main content area shows the 'Launch Instance' button and a table of instances. The instance 'CMAQ_benc...' is highlighted, showing its details: Instance ID i-062f303e71904a9be, Instance Type c6g.16xlarge, Availability Zone us-east-1a, Instance State running, Status Checks Initializing, Alarm Status None, Public DNS (IPv4) ec2-54-211-144-166.compute-1.amazonaws.com, IPv4 Public IP 54.211.144.166, and IPv6 IPs -. Below the table, the instance details are shown, including the Description, Status Checks, Monitoring, and Tags. A terminal window is open on the instance, showing the execution of 'systemctl enable --now cockpit.socket' and 'CMAQ_G2_ncores64.txt'. A file explorer window is also open, showing a directory structure with many log files.

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP	IPv6 IPs
CMAQ_benc...	i-062f303e71904a9be	c6g.16xlarge	us-east-1a	running	Initializing	None	ec2-54-211-144-166.compute-1.amazonaws.com	54.211.144.166	-

```
centos@ip-172-31-93-150:~/CMAQ/CCTM/scripts
login as: centos
Authenticating with public key "imported-openssh-key"
Activate the web console with: systemctl enable --now cockpit.socket

Last login: Sun Oct 17 20:05:34 2021 from 71.60.173.251
[centos@ip-172-31-93-150 ~]$ cd /home/centos/CMAQ/CCTM/scripts
[centos@ip-172-31-93-150 scripts]$ ./CMAQ533G2_ncores64_12SE1.sh > CMAQ_G2_ncores64.txt
```

Name	Size	Changed	Rights	Owner
CTM_LOG_063.v533_gcc_Bench_2016_12SE1_20160701	23 KB	9/3/2021 11:08:45 AM	rw-rw-r--	centos
CTM_LOG_062.v533_gcc_Bench_2016_12SE1_20160701	23 KB	9/3/2021 12:03:19 PM	rw-rw-r--	centos
CTM_LOG_061.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_060.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_059.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_058.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_057.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_056.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_055.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_054.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_053.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_052.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_051.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_050.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_049.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_048.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_047.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_046.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_045.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_044.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_043.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_042.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_041.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_040.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_039.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_038.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_037.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_036.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_035.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_034.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_033.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_032.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_031.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_030.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_029.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_028.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_027.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_026.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_025.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos
CTM_LOG_024.v533_gcc_Bench_2016_12SE1_20160701	23 KB	10/17/2021 4:06:20 PM	rw-rw-r--	centos

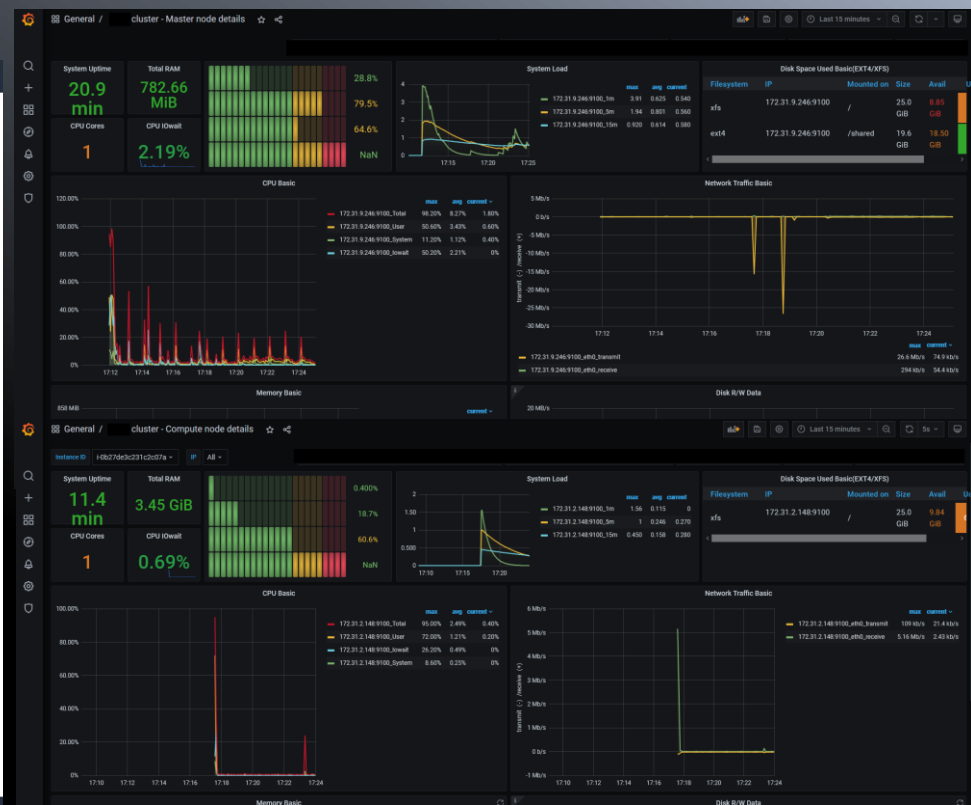
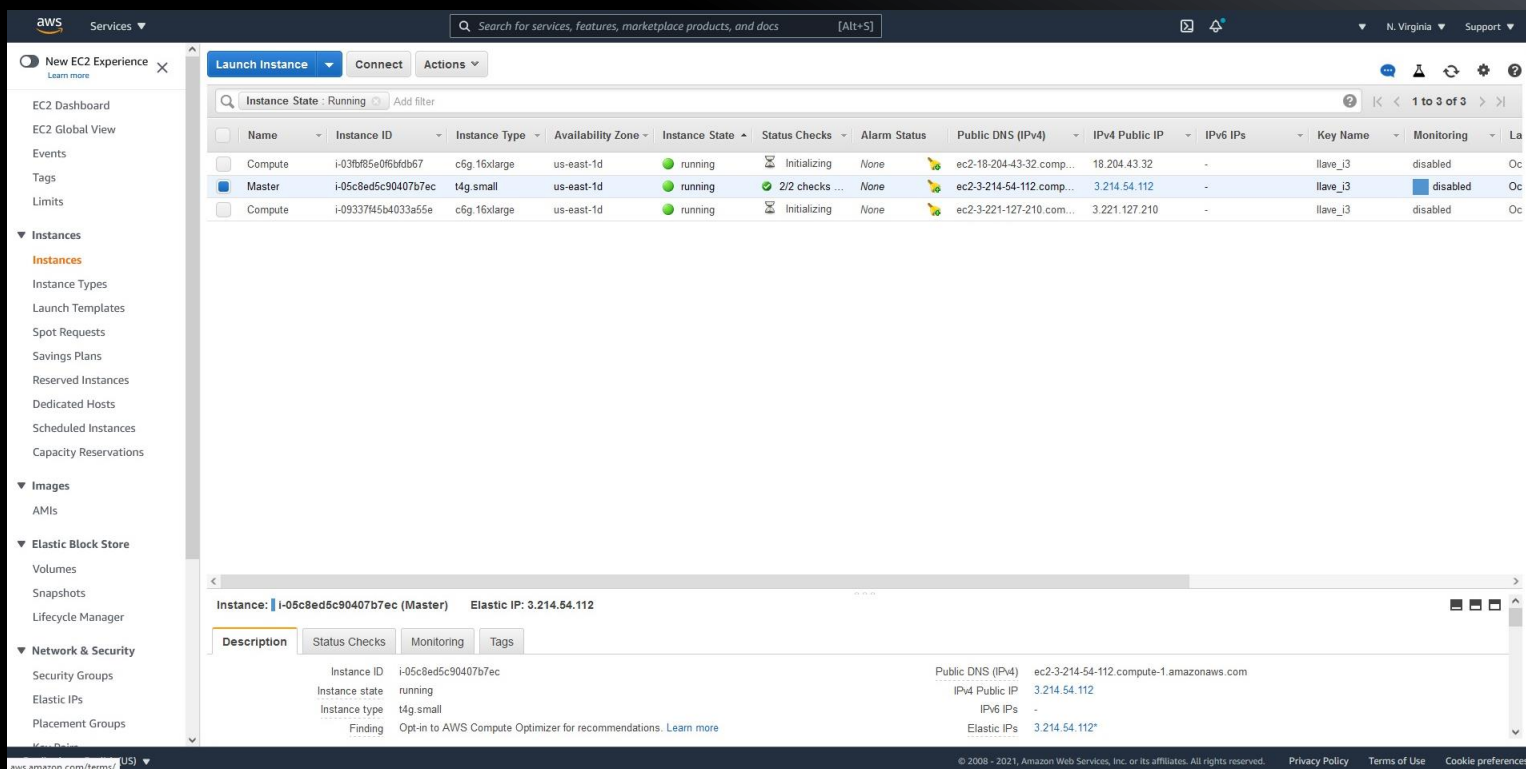
Migration and performance of CMAQ & WRF-CMAQ in the public cloud with commercial images

20th Annual CMAS Conference – November 1-5 2021 (virtual)

Integrated product for CMAQ

Launch and connect to a cluster using AWS-Parallelcluster

- Uses slurm as scheduler
- Cluster monitoring tools are also available



Migration and performance of CMAQ & WRF-CMAQ in the public cloud with commercial images
20th Annual CMAS Conference – November 1-5 2021 (virtual)

Performance evaluation

Software

Benchmark: 2016 Southeast U.S.

Domain size: 100 x 80 x 35

Species tracked: 218

Input files > 6 GB

Not very (RAM) memory hungry: Was able to run on instances with 8GB

Hardware

Intel Xeon E5-2697 v4
32 cores (dual socket)–
14 nm (2016)



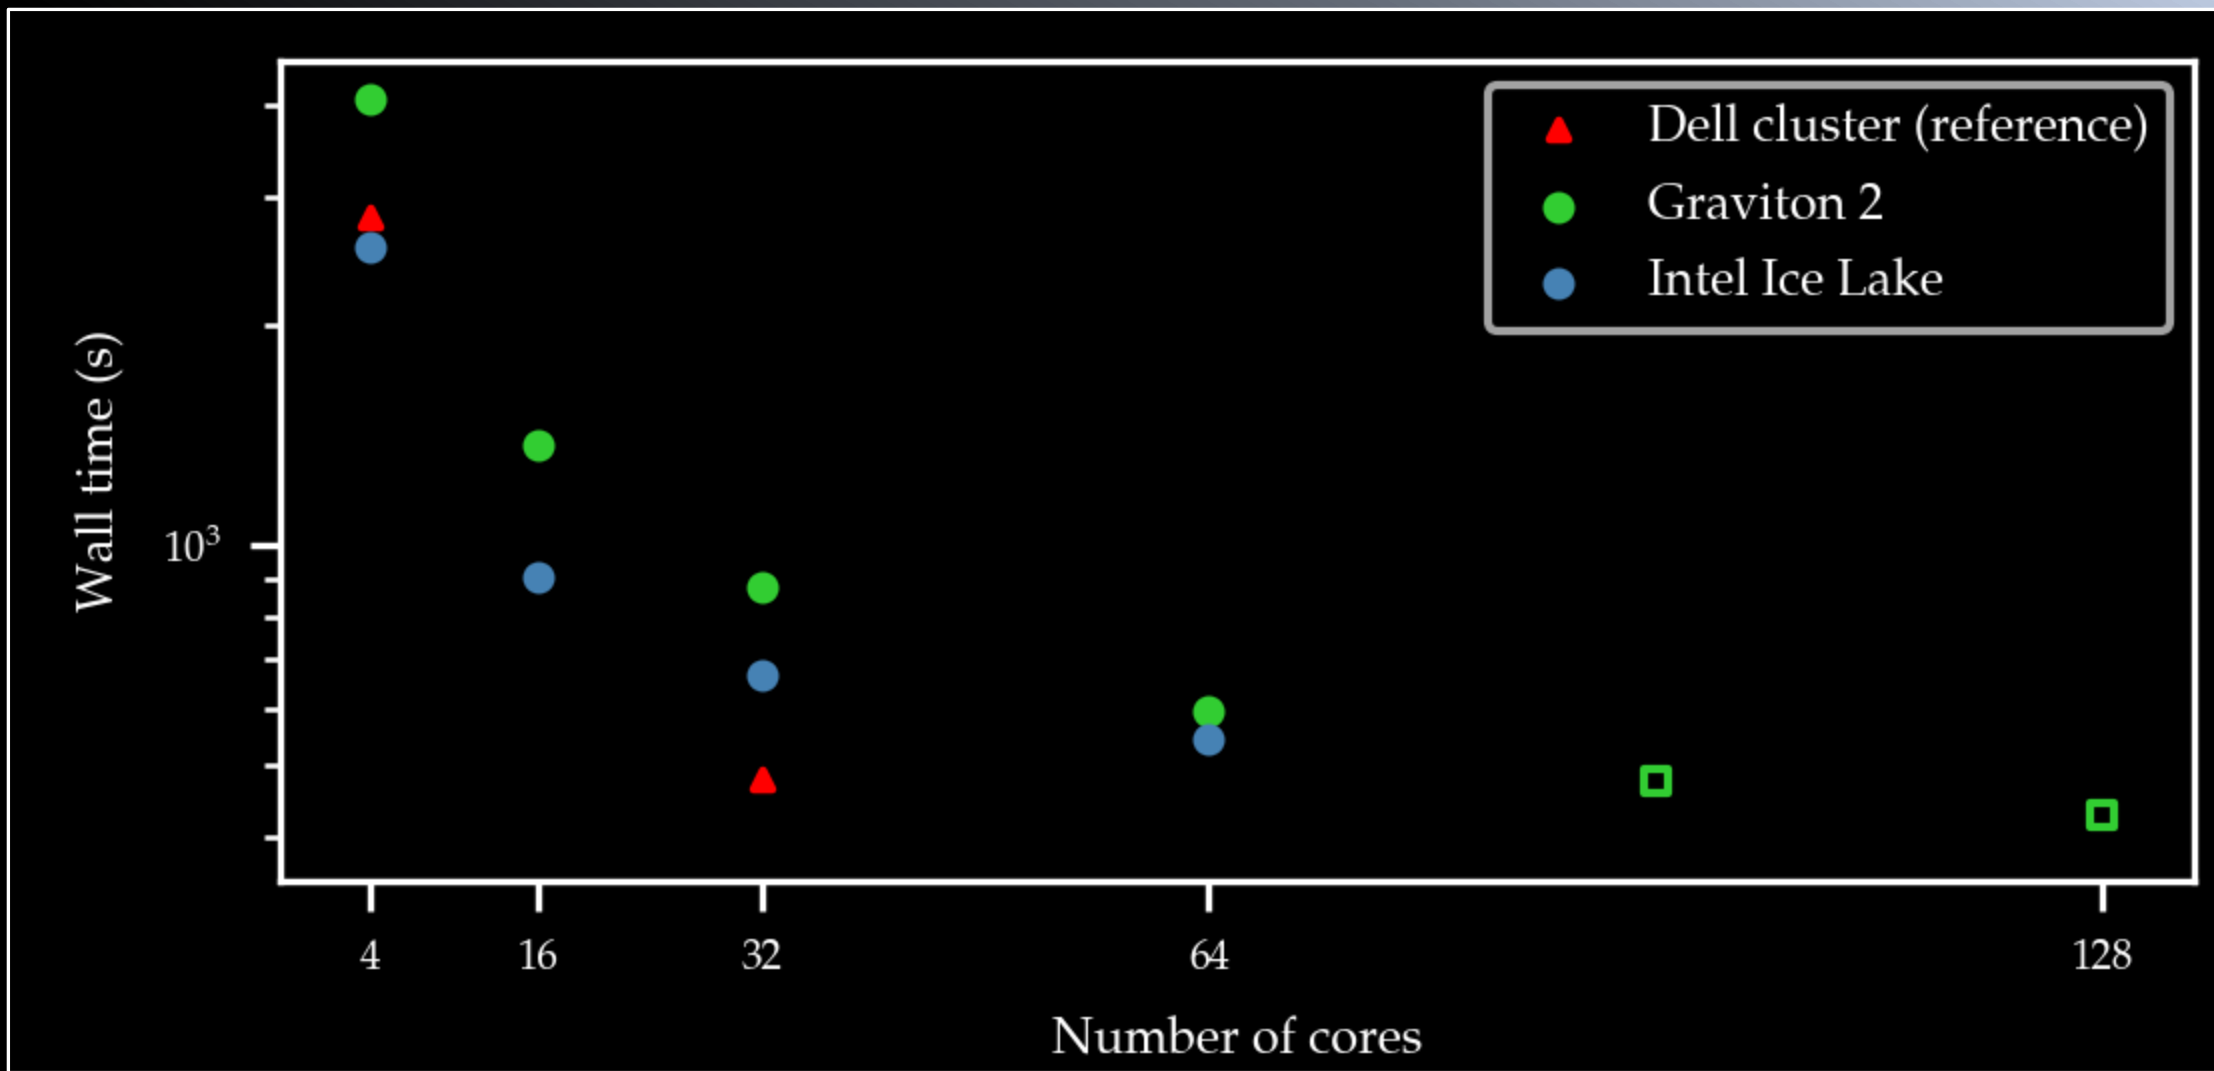
AWS Graviton2
64 cores (single socket)
– 7 nm (2019)



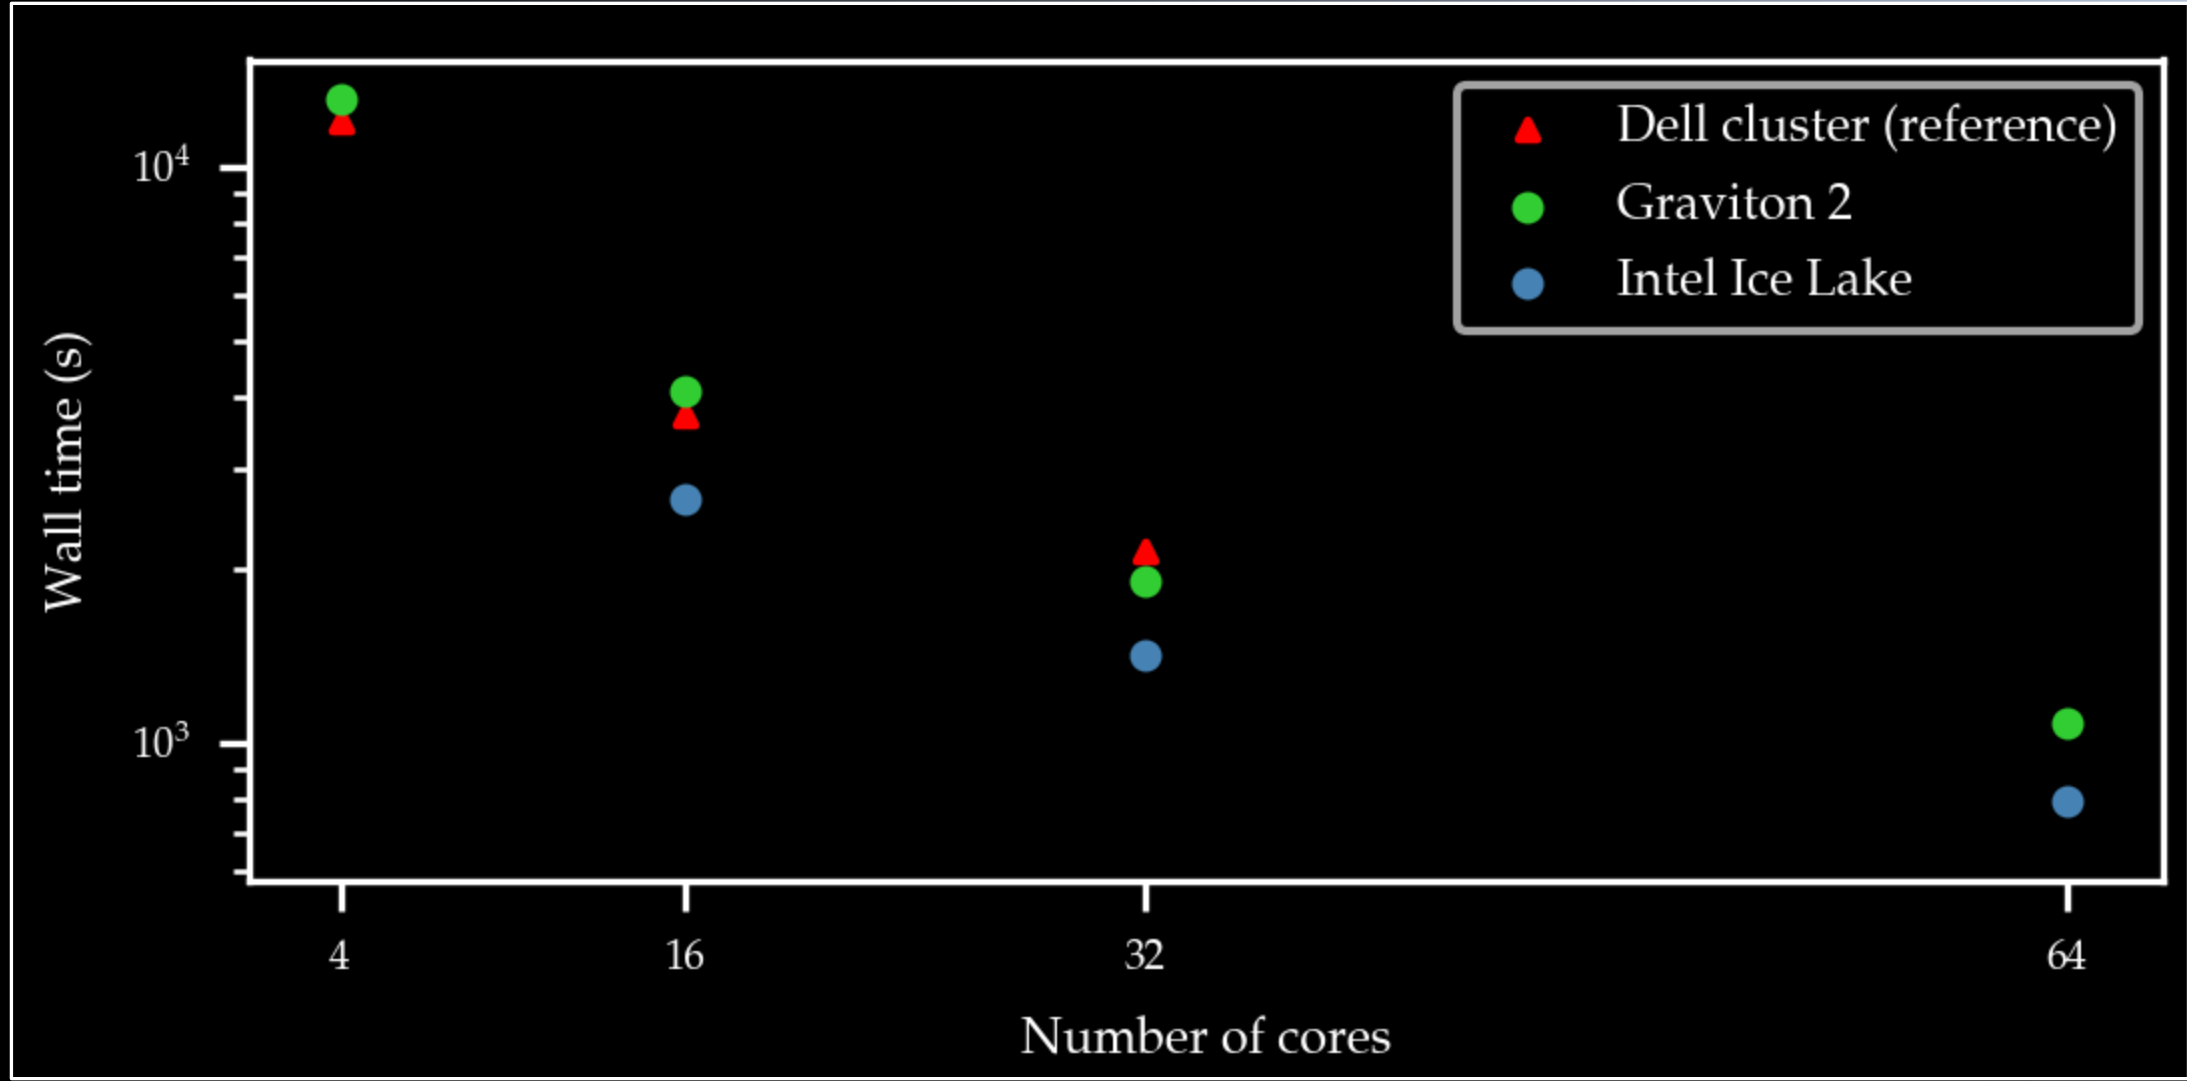
Intel Xeon Platinum
8375C -64 cores (dual
socket)– 10 nm (2021)



Performance evaluation – CMAQ measurements



Performance evaluation – WRF-CMAQ measurements



Cost analysis

Performance evaluation is relatively straightforward but cost analysis is more complex. A full analysis requires a total cost of ownership (TCO), which is unique for each organization. As a minimum, it must include:

- Computational power account for most of the cost but not all of it
- Storage
- Outbound traffic (can also add up for HPC apps)
- Commercial AMI fees must also be considered

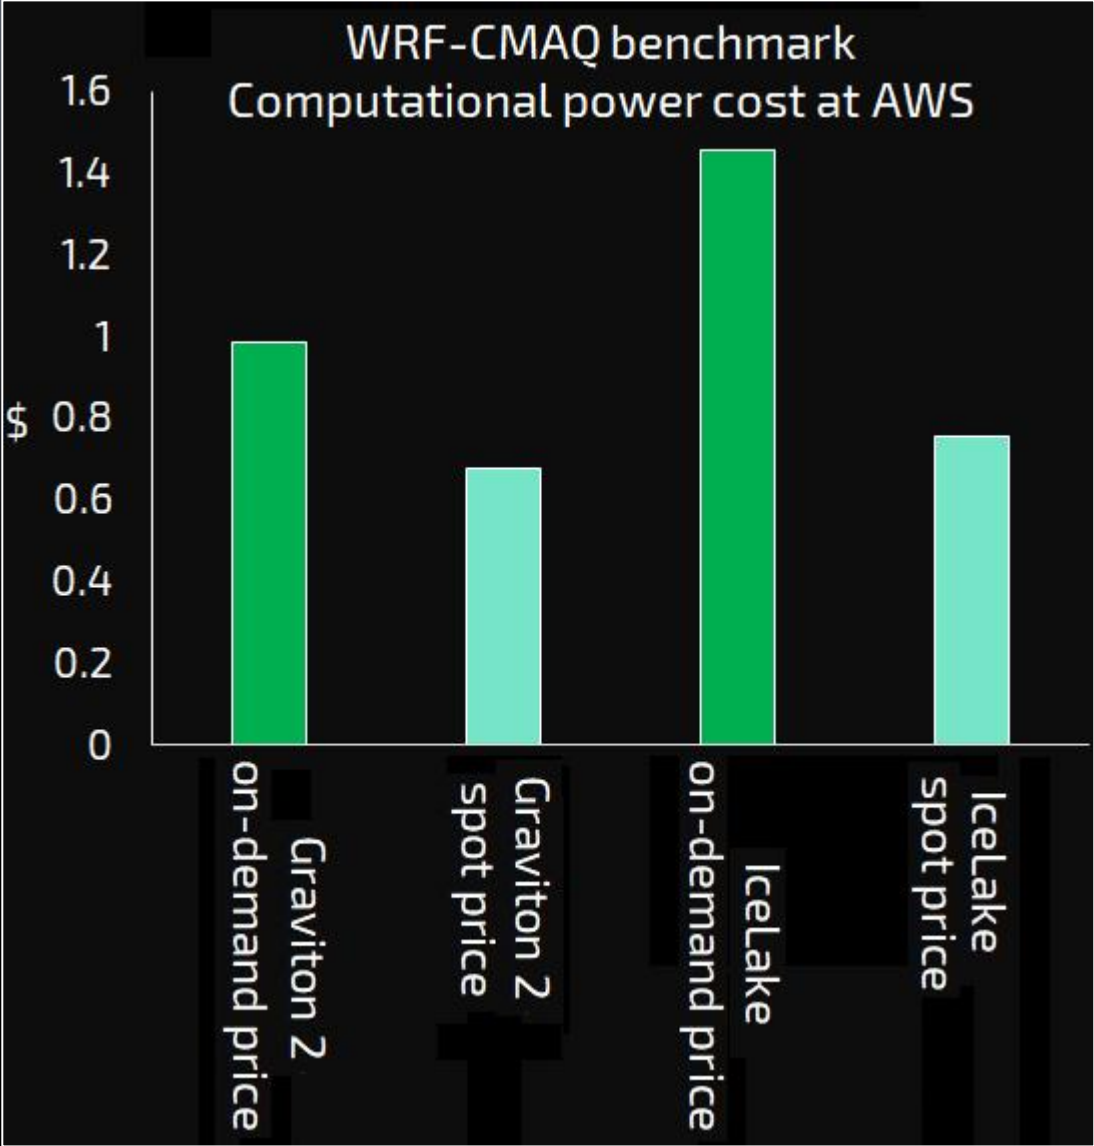
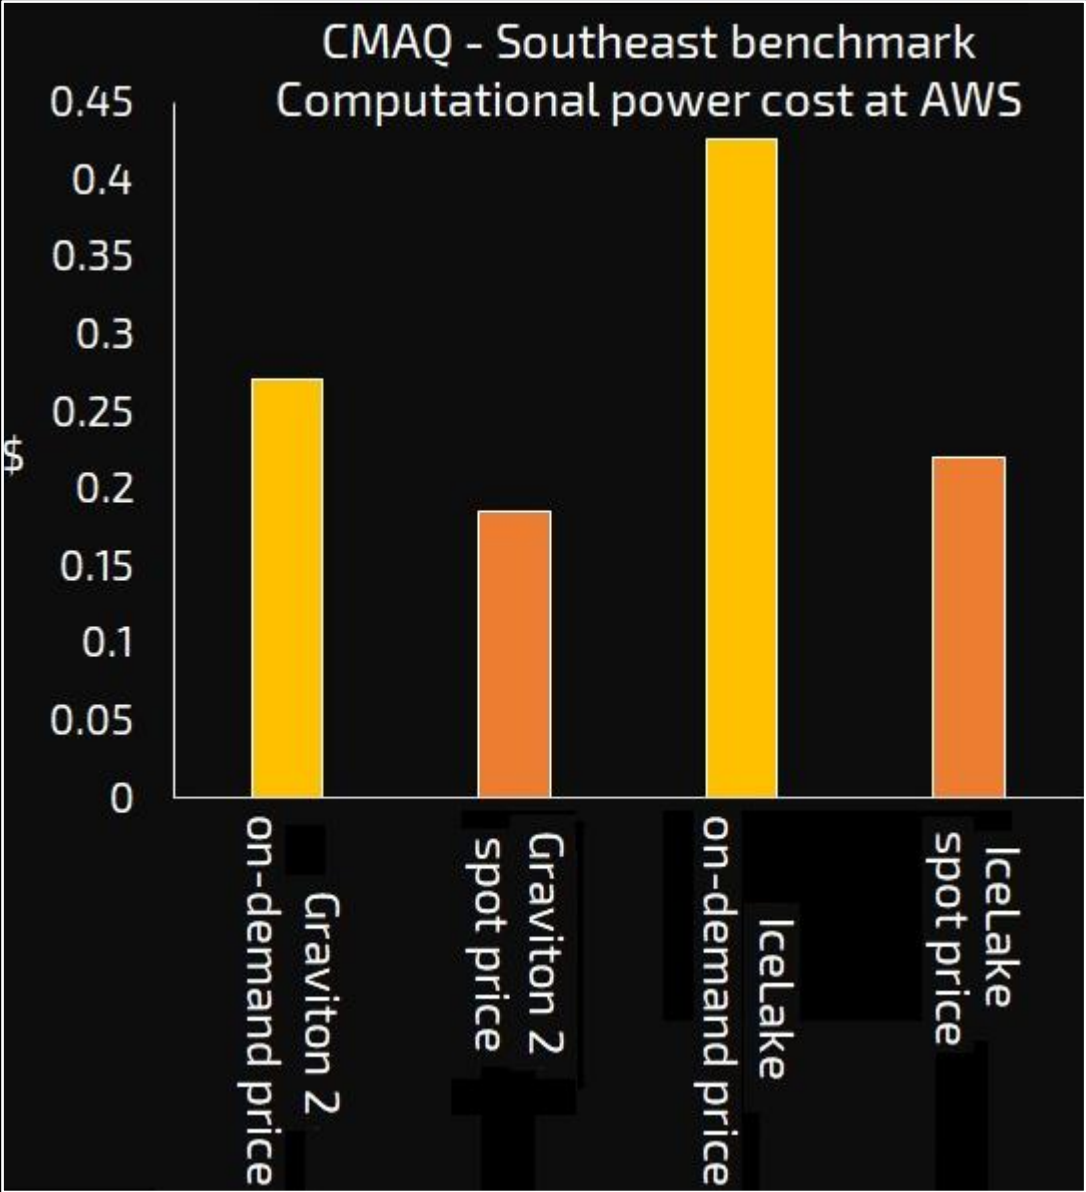
Many categories have several tiers and conditions

Preliminary cost analysis

- It focuses on computational power plus AMI fees
- It categorizes cost associated with computational power into 2:
 - On-demand prices (maximum cost)
 - Spot prices (maximum savings)



Cost analysis – cost estimates



Conclusions

The main conclusions from the study are the following:

- CMAQ & WRF- CMAQ are available as an integrated image from the AWS Marketplace
- Performance is good and they can be run with many cores using clusters of several instances
- Two choices: Intel IceLake vs. Graviton2 (performance vs. cost)
- Matching the CMAQ case with the right hardware is important to maximize performance and keep expenses low

Questions?

