



# FV3 on AWS – Successes and Challenges

29 July 2020

Brian Etherton, PhD – Numerical Weather Prediction Scientist  
Stefan Cecelski, PhD, Christopher Cassidy, and Travis Hartman



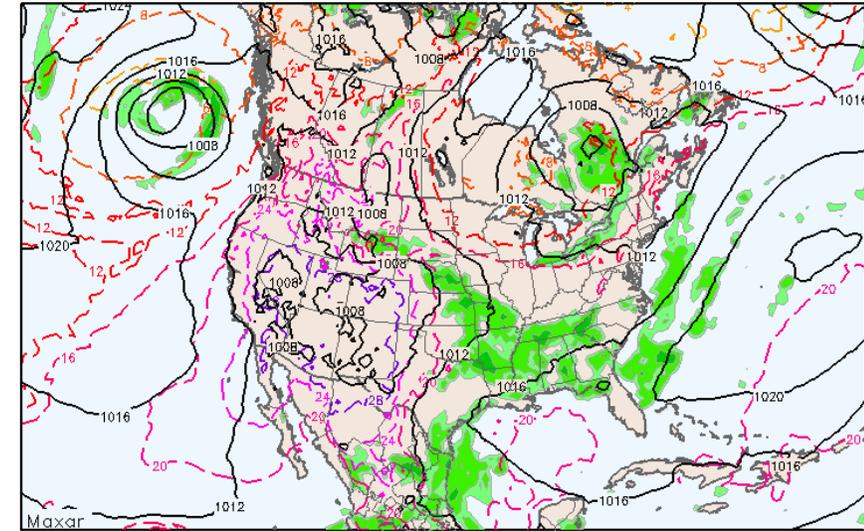
**MAXAR**



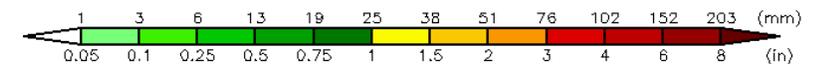
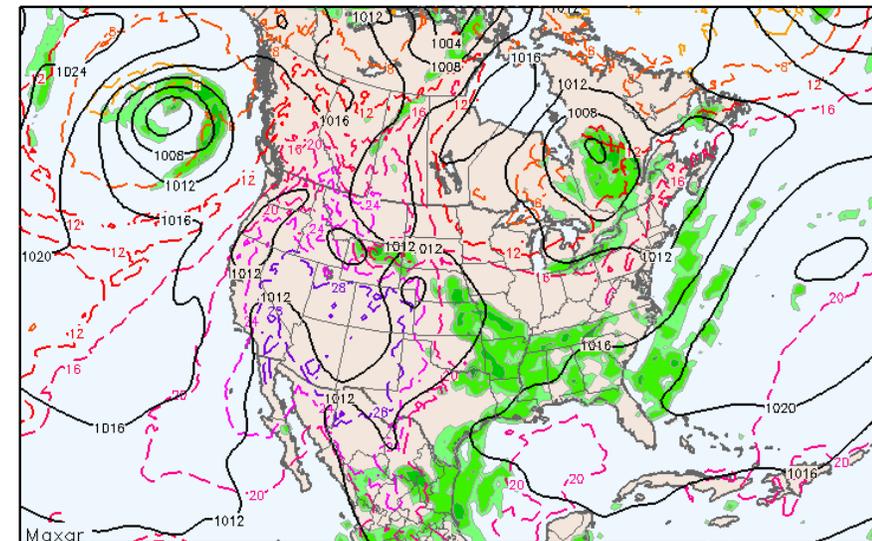
## Background

- For over a year Maxar has run, and continues to run, a numerical weather prediction (NWP) application on AWS cloud computing resources for a program called Cirrus.
- The success of the program relies on Maxar being able to run its NWP application faster than NOAA does on its supercomputer.
  - Allows Maxar to deliver the weather forecast generated by the NWP application to clients with greater lead times, guiding more informed and timely decisions.
- To achieve the required application performance, Maxar engaged with AWS to design and build a cloud computing environment as a cost-effective and scalable solution.
  - Cirrus leverages AWS EC2, ParallelCluster, EFA, FSx for Lustre and other AWS tools to build a cloud-based high-performance computing (HPC) cluster.

054 Hr EGFS Op SLP/850 mb Temp/6 Hr Pcp – Wed 29 Jul 18Z



054 Hr GFS Op SLP/850 mb Temp/6 Hr Pcp – Wed 29 Jul 18Z

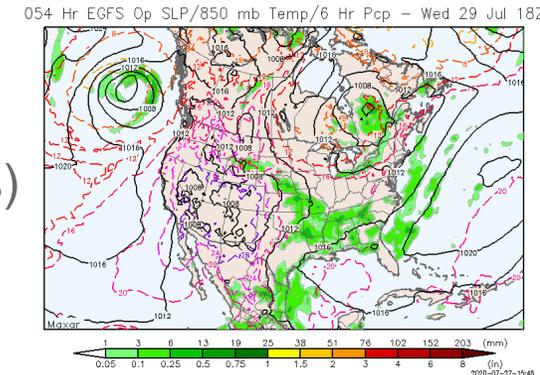


2020-07-27-15:42



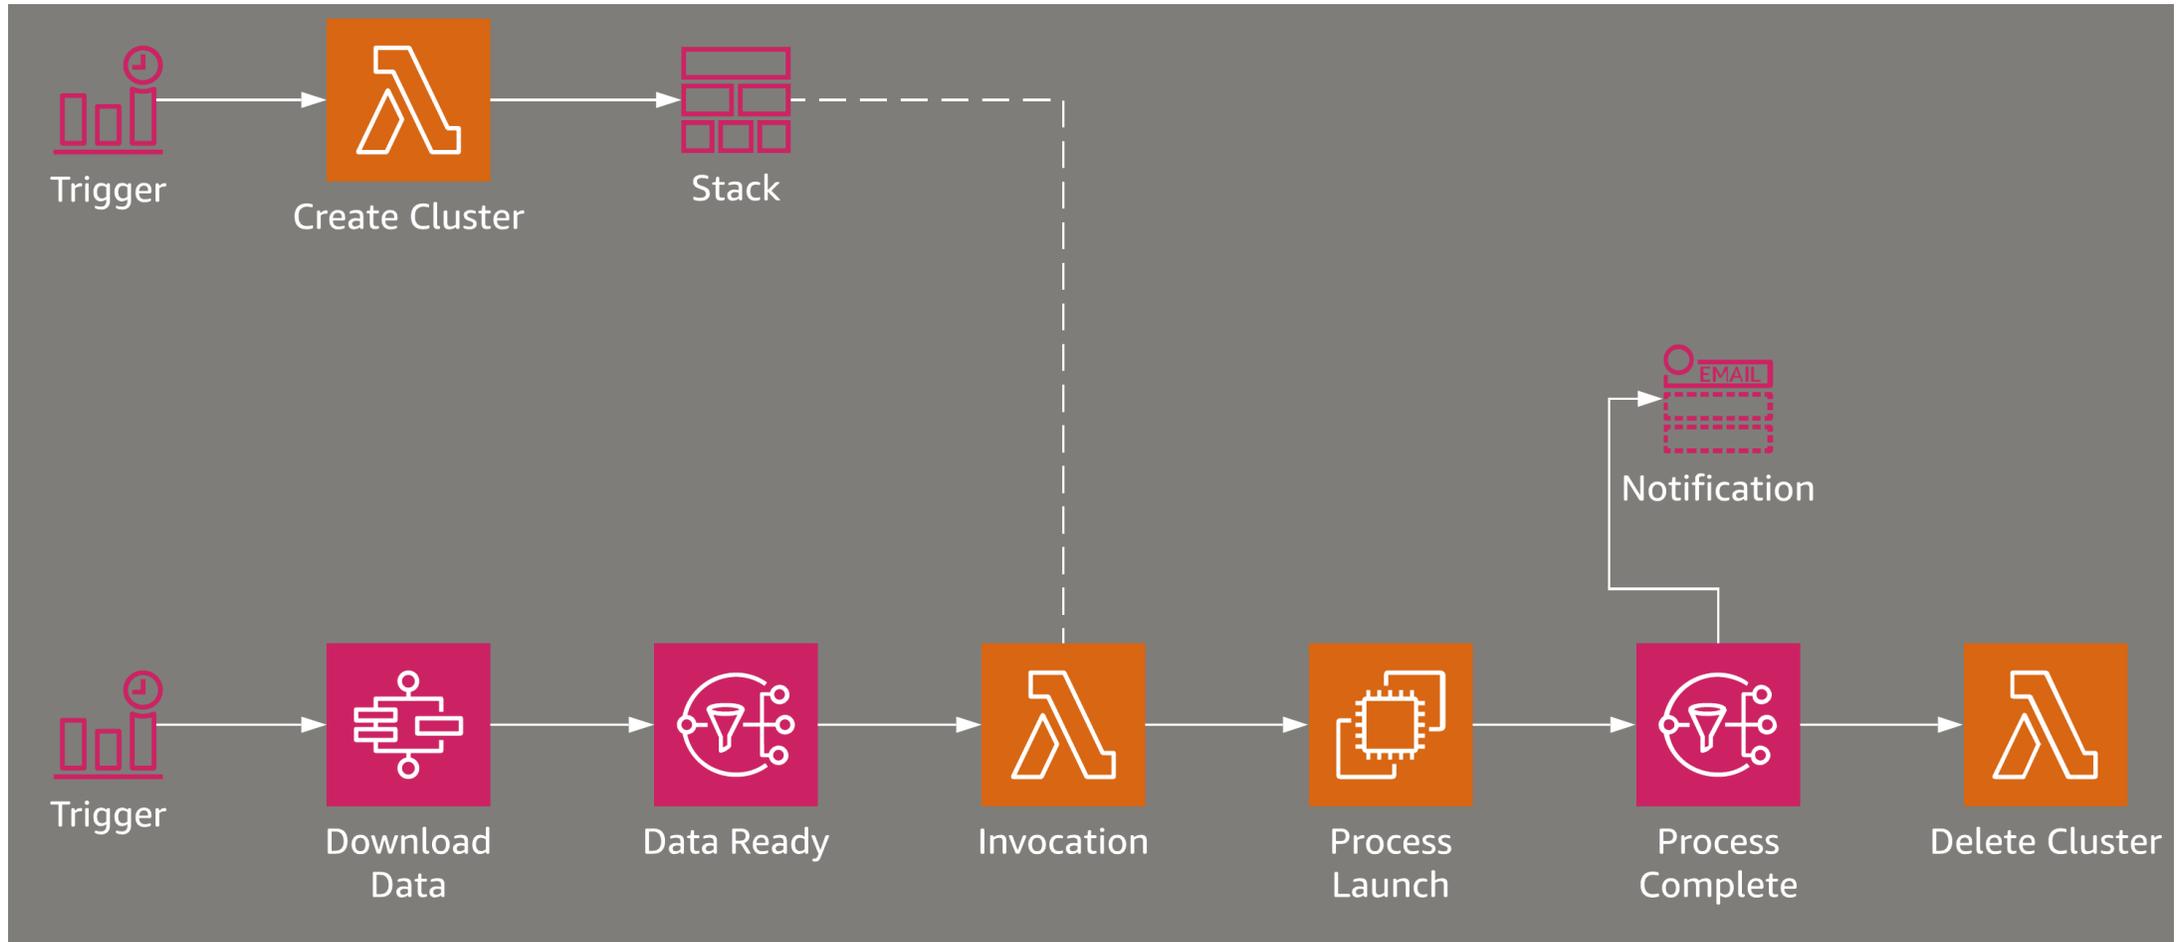
## AWS FV3GFS Workflow

- Initial effort was to refactor NCEP/VLAB code/workflow to AWS environment
  - Recompile code, tend to any issues that compilers found (Derived types)
  - Rework the NCEP workflow to be more hardware agnostic (Directories/paths)
  - Discern model layout best suited for AWS hardware (36 cores per node)
- The entire workflow is automated from cluster spin-up to cluster spin-down
  - HPC cluster resources (EC2 instances, EBS volumes, FSx Lustre) are re-allocated each day
- Execution workflow is automated using a series of Step and Lambda functions with SNS notifications
  - CloudFormation stacks build the AWS resources for the workflow – can be deployed in any AWS region with minimal modification
  - Infrastructure updates are automatically deployed to specified regions using CodePipeline
- Workflow designed and built with AWS Professional Services





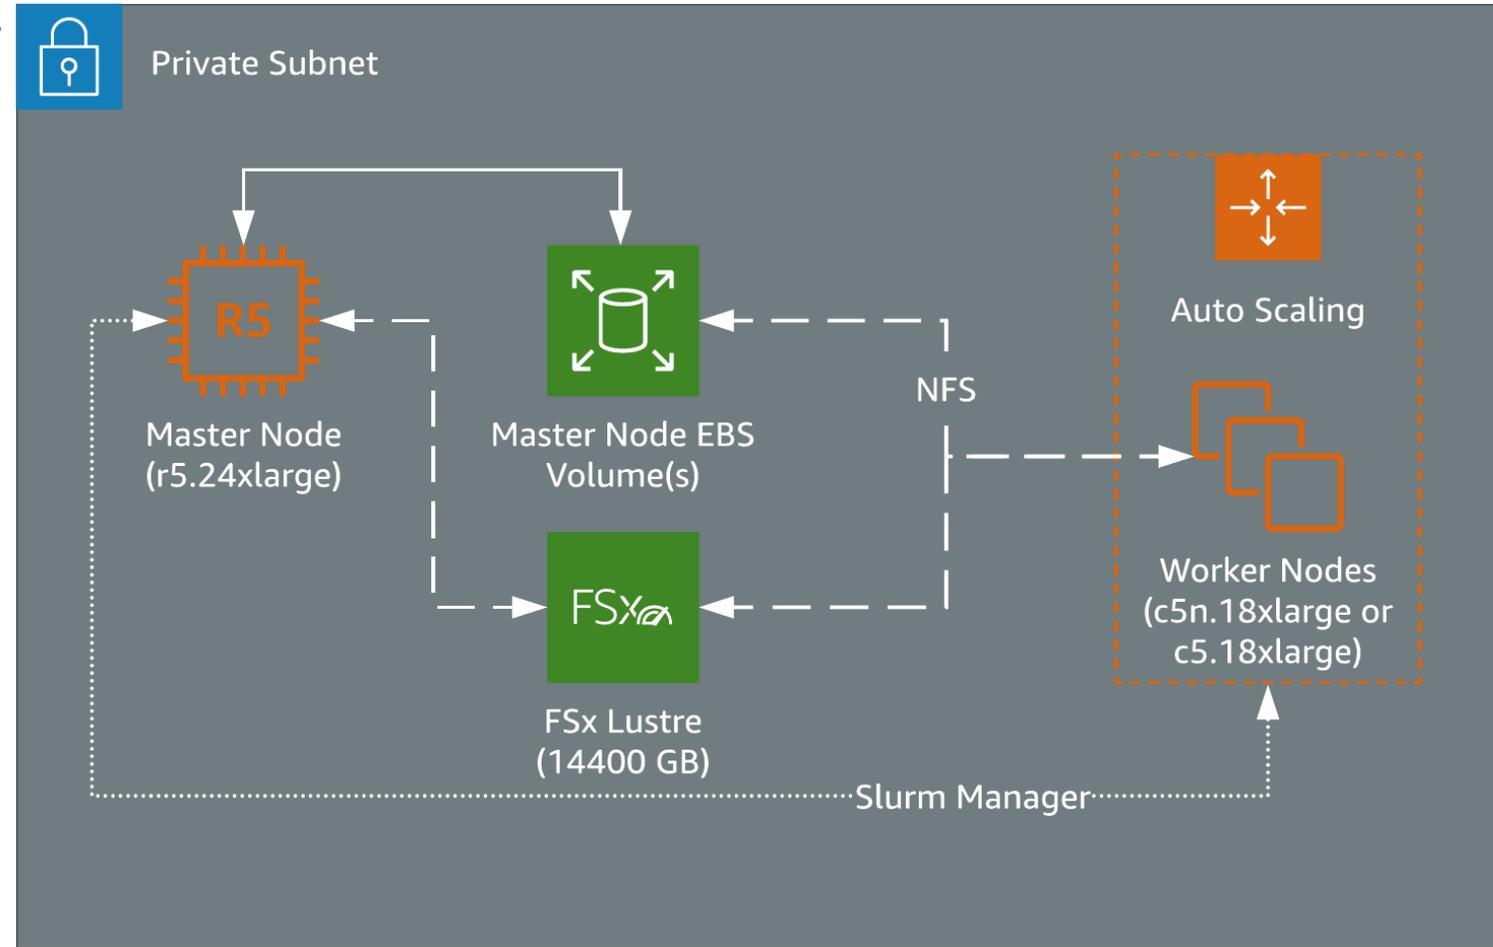
## Automated FV3GFS Workflow within AWS





## Maxar's AWS Cloud HPC Clusters

- Leveraging AWS ParallelCluster for cluster management
  - 2 clusters: production in us-east-1, backup in us-west-2
- Production cluster:
  - Master Node: r5.24xlarge
  - Worker Nodes: 156x c5n.18xlarge within an ASG
  - File System: 14400 GB AWS FSx Lustre (12 object storage targets [OSTs])
  - Network: AWS Elastic Fabric Adapter (EFA) with max bandwidth of 100 Gbps





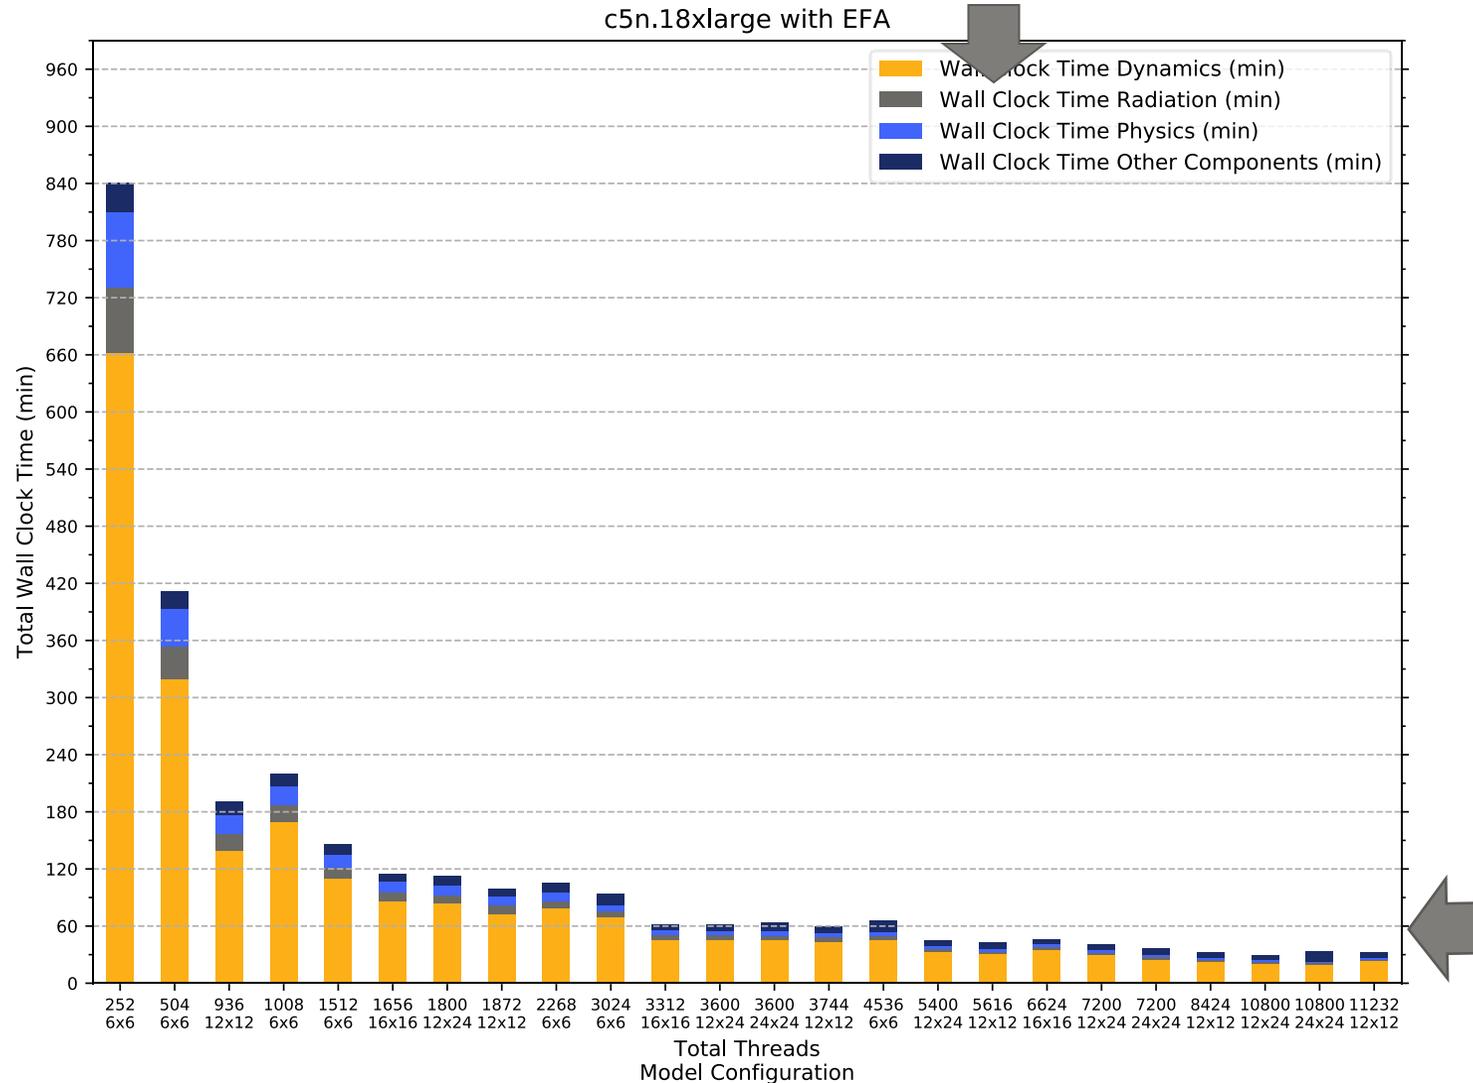
## Case study on performance

- **37** different HPC cluster configurations were built to test the FV3GFS application on as little as **252** cores to over **11,000** cores.
  - Each of these clusters was built using AWS ParallelCluster to ensure consistent configuration and deployment.
- Maxar currently runs the application on both c5.18xlarge and c5n.18xlarge AWS EC2 instances.
  - **24** of the case study cluster configurations utilize c5n.18xlarge with the EFA networking adapter (max 100Gbps).
  - The remaining **13** configurations use c5.18xlarge instances with TCP networking (max 25Gbps).
- All configurations leverage a **14TB** FSx Lustre file system with a progressive file layout (PFL) across the 12 object storage targets (OSTs).

```
[aws]
aws_region_name = us-east-1
[global]
cluster_template = AZ-a
update_check = true
sanity_check = false
[cluster AZ-a]
base_os = alinux
master_instance_type = r5.24xlarge
compute_instance_type = c5.18xlarge
placement_group = DYNAMIC
vpc_settings = AZ-a
master_root_volume_size = 3000
compute_root_volume_size = 100
initial_queue_size = 100
maintain_initial_size = true
max_queue_size = 100
fsx_settings = fs
ec2_iam_role = EC2-Cluster-Generatio
extra_json = { "cluster" : { "cfn_sc
scheduler = slurm
[fsx fs]
shared_dir = /UFS
storage_capacity = 14400
imported_file_chunk_size = 1024
```



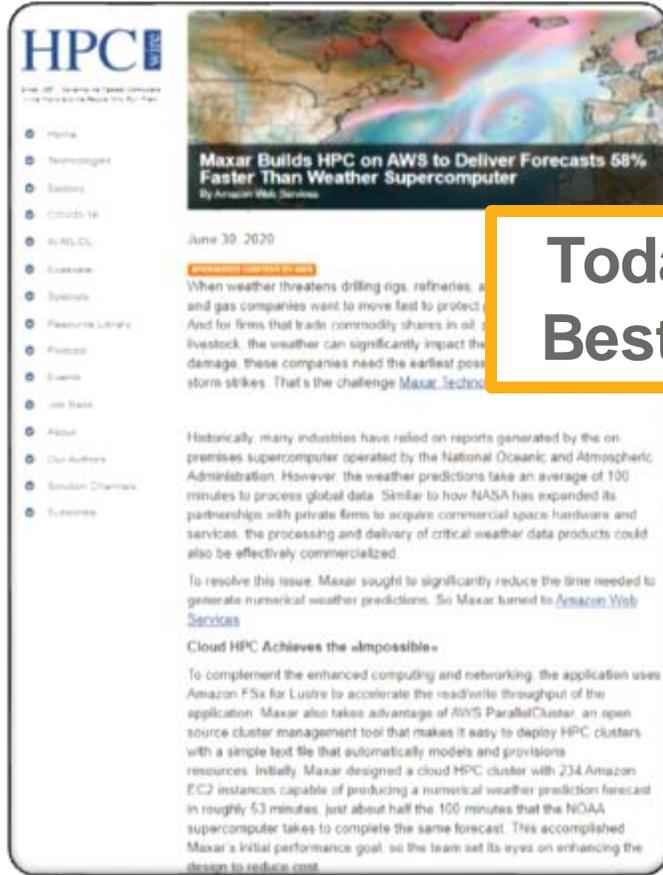
## Results – Wall clock times by application component



- As core counts increase, a larger percentage of total wall clock time is spent in other parts of the application.
- Larger core counts require more communication, which partly explains more time spent in other components than the main three (dynamics, physics, and radiation).
- The scalability of the application is highly dependent on model configuration and the balance between MPI and openMP threading.



# Awareness of Maxar's Cloud Based HPC for NWP



HPC Wire Article

Today: MAXAR wins AWS award for Best HPC Solution for Public Sector



AWS Published Case Study

Maxar Blog on Cloud Based HPC for NWP





# Taking on the GEFS v12

## Toward a simpler, UFS-based Production Suite (notional)



NPS Modeling System	Current Version	Q1 FY 20	Q2 FY 20	Q3 FY 20	Q4 FY 20	Q1 FY 21	Q2 FY 21	Q3 FY 21 - Q2 FY 22 MORATORIUM	Q3 FY 22	Q4 FY 22	Q1 FY 23	Q2 FY 23	Q3 FY 23	Q4 FY 23	Q1 FY 24	Q2 FY 24	Q3 FY 24	Q4 FY 24	UFS Application
Global Weather & Global Analysis	GFS/ GDA Sv15						GFSv16												UFS Medium Range & Sub-Seasonal
Global Waves	GWMSv3																		
Global Weather Ensembles	GEFSv11																		
Global Wave Ensembles	GWESv3																		
Global Aerosols	NGA C v2																		UFS Marine & Cryosphere
Short-Range Regional Ensembles	SREFv7																		
Global Ocean & Sea-Ice	RTOFsv1.2																		UFS Seasonal
Global Ocean Analysis	GODASv2																		
Seasonal Climate	CDA S/ CF Sv2																	SF Sv1	
Regional Hurricane 1	HWRv12																		UFS Hurricane
Regional Hurricane 2	HMONv2																		
Regional High Resolution CAM 1	HIRes Window v7																		UFS Short-Range Regional HiRes CAM & Regional Air Quality
Regional High Resolution CAM 2	NAM_nests/ Fire Wxv4																		
Regional High Resolution CAM 3	RAPv4/ HRRRv3																		
Regional HiRes CAM Ensemble	HREFv2																		
Regional Mesoscale Weather	NAMv4																		UFS Air Quality & Dispersion
Regional Air Quality	CMAQv5																		
Regional Surface Weather Analysis	RTMA/ URM A v2.7																		UFS Air Quality & Dispersion
Atmospheric Transport & Dispersion	HySPLITv7																		
Coastal & Regional Waves	NWPSv1.2																		UFS Coastal
Great Lakes	GLWUv3.4																		
Regional Hydrology	NWMv2																		UFS Lakes
Space Weather 1	WAM/PEv1																		
Space Weather 2	ENLILv1																		UFS Hydrology
																			UFS Space Weather

Disclaimer: This is a plan; all decisions will be driven by evidence

- The GEFS ensemble is set for an upgrade in September 2020.
- Running the FV3/UFS 31 times at lower resolution is a new challenge for AWS resources.



## Concluding Thoughts

- For over a year, Maxar has been running FV3 in AWS.
- CI/CD achieved via CodePipeline, for relatively easy updates to the executable code.
- AWS resources allow for scalability to core counts greater than NCEP NCO without having to purchase the physical hardware of that magnitude.
- AWS resources and tools enable creativity in costing for testing/development. User has better control of matching budget to requirements, no huge investment upfront.
- We ran FV3 using numerous configurations, with core counts from 252 to over 10000
  - We found the configuration that worked best for us, but that is NOT the configuration that is necessarily best for everyone.
  - Members of the community can use the cloud to do FV3 simulations using resources that match their budgets and scopes.
  - Using the cloud, UFS users can get just the compute they need for their work.

**MAXAR**

[MAXAR.COM](https://maxar.com)