Performance Analysis of a Numerical Weather Prediction Application in Microsoft Azure

Emmanuell D Carreño, Eduardo Roloff, Jimmy V. Sanchez, and Philippe O. A. Navaux

WSPPD 2015 - XIII Workshop de Processamento Paralelo e Distribuído

UFRGS
UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

GPPD
Grupo de Processamento Paralelo e Distribuído

CNPq
Conselho Nacional de Desenvolvimento Científico e Tecnológico

.inf
INSTITUTO DE INFORMÁTICA
UFRGS

Introduction

- Numerical Weather Prediction (NWP) is one area of science that have evolved continuously with the help of computing.
- Every year datasets grows larger and computing demand reaches the limit of deployed architecture.
- High Performance Computing (HPC) allowed performing simulations on less time, but also has increased upfront costs.

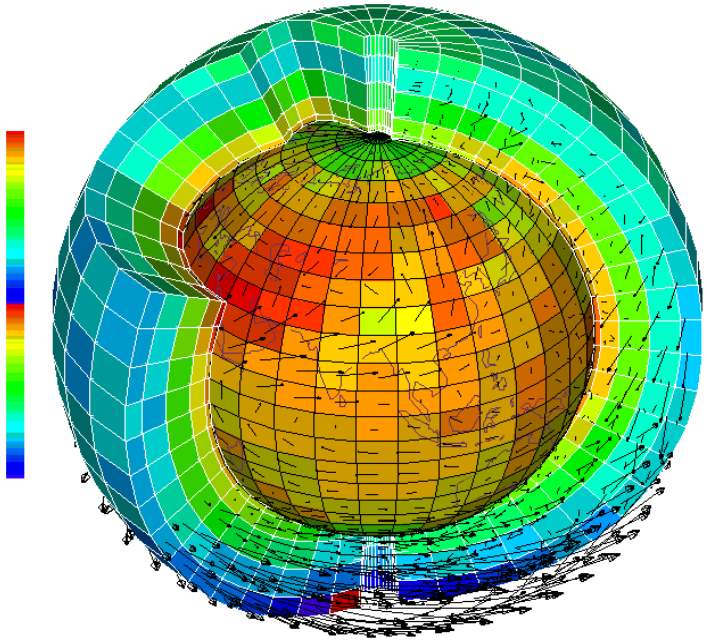

Motivation

Analyze the scaling performance of an HPC (NWP) application in Microsoft Azure Cloud Infrastructure.

Why?

HPC

- Low priority jobs to the cloud
- Hardware investment budget
- Need for energy efficiency

NWP

- eScience platforms
- On-demand severe weather alert systems
- Sharing results

Cloud

- Previous work: small HPC benchmarks
- Optimize costs
- Reduce complexity

Outline

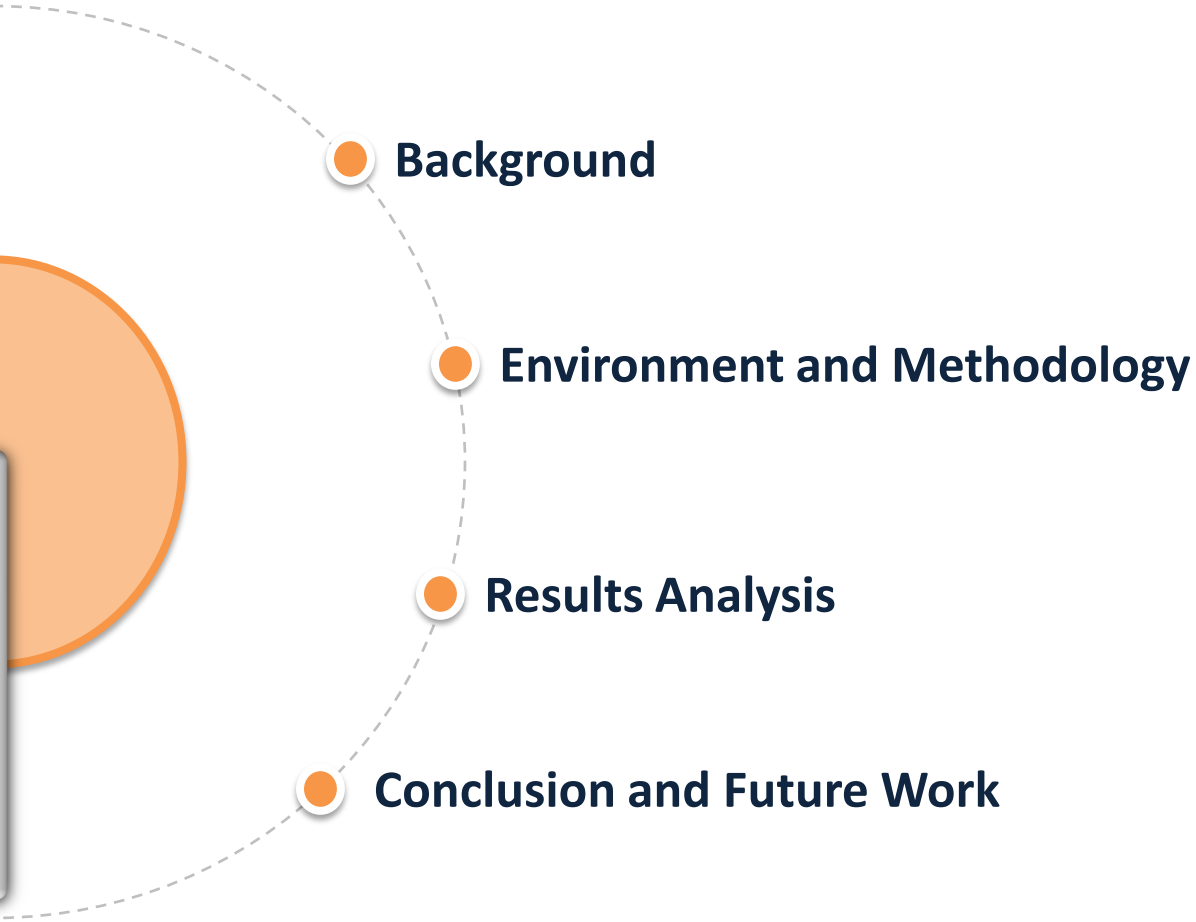
Background
Environment and Methodology
Results Analysis
Conclusion and Future Work

Cloud Computing

- Developed with the combination and evolution of distributed computing and virtualization.
- With contributions from grid and parallel computing.

Essential Characteristics:

- On demand self-service
- Broad network access
- Resource Pooling
- Rapid elasticity
- Measured Service

Infrastructure as a Service (IaaS):

- Control over operating system
- User managing the VM
- Storage
- Limited control networking
- Deployed applications

Azure

Cloud computing infrastructure from Microsoft

Second cloud provider by market share

19 Datacenters and Diversity of Instance Sizes

Why Azure?

Better performance than other IaaS Providers (Roloff et al, 2012)

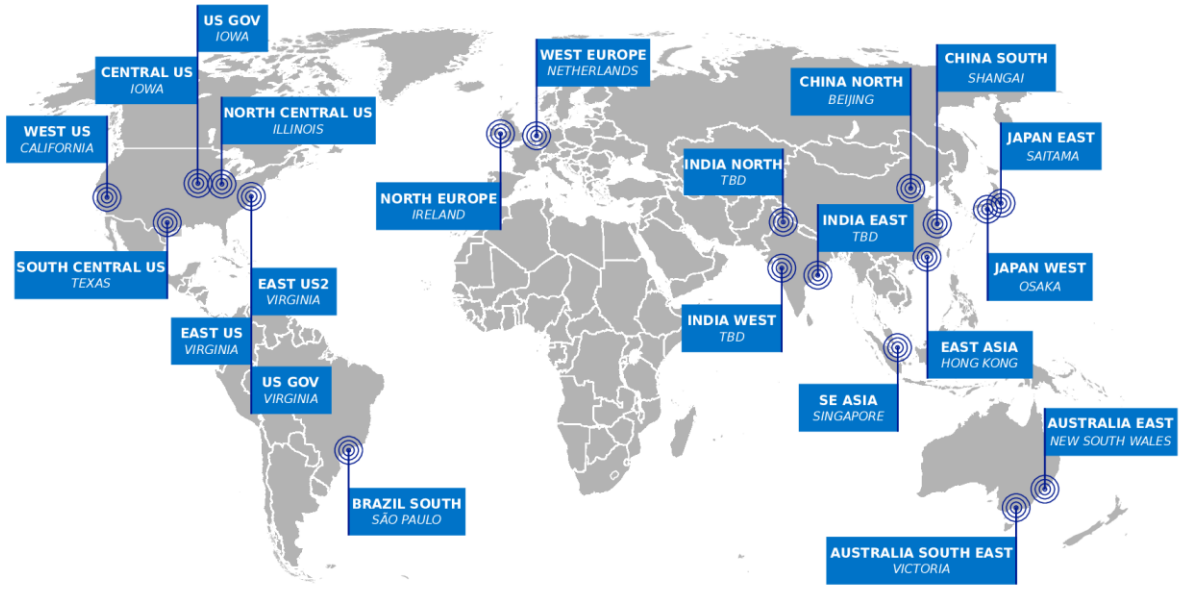

Adapted from Microsoft Ignite 2015

BRAMS

- Numerical weather prediction model used to simulate atmospheric conditions.
- Modified to simulate in a better way the tropical atmosphere, specially Brazil climate characteristics.
- Developed by the CPTEC-INPE.
- Used by Brazilian and South American regional weather centers .

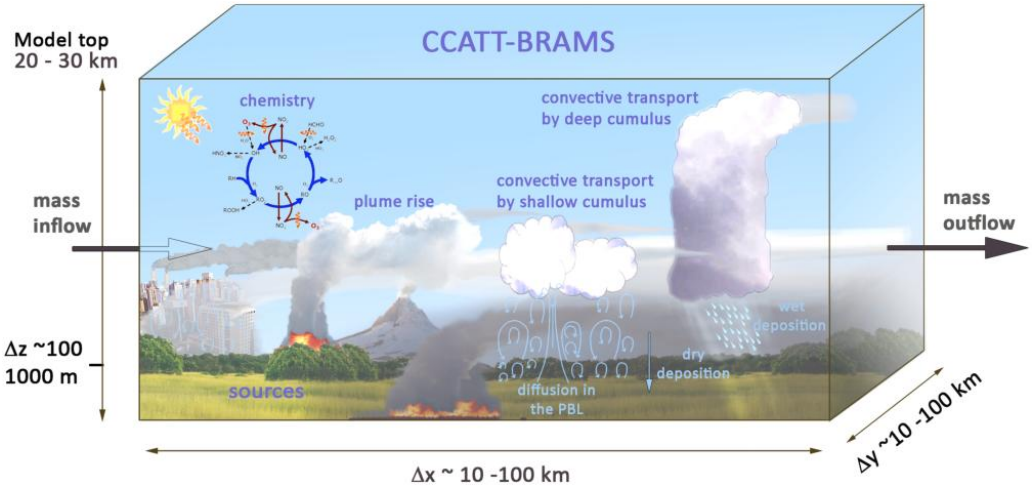

Source: Longo et al. 2013

BRAMS Workflow

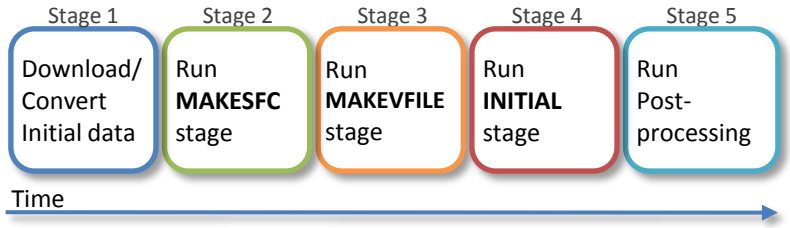

Content	Size [GB]
CLIMATOLOGY	0.52
DP	86.00
EMISSION-DATA	3.60
FIRES-DATA/DSA	0.05
FIRES-DATA/MODIS	0.18
SOIL-MOISTURE	23.00
SURFACE-DATA	8.80
Total	122.15

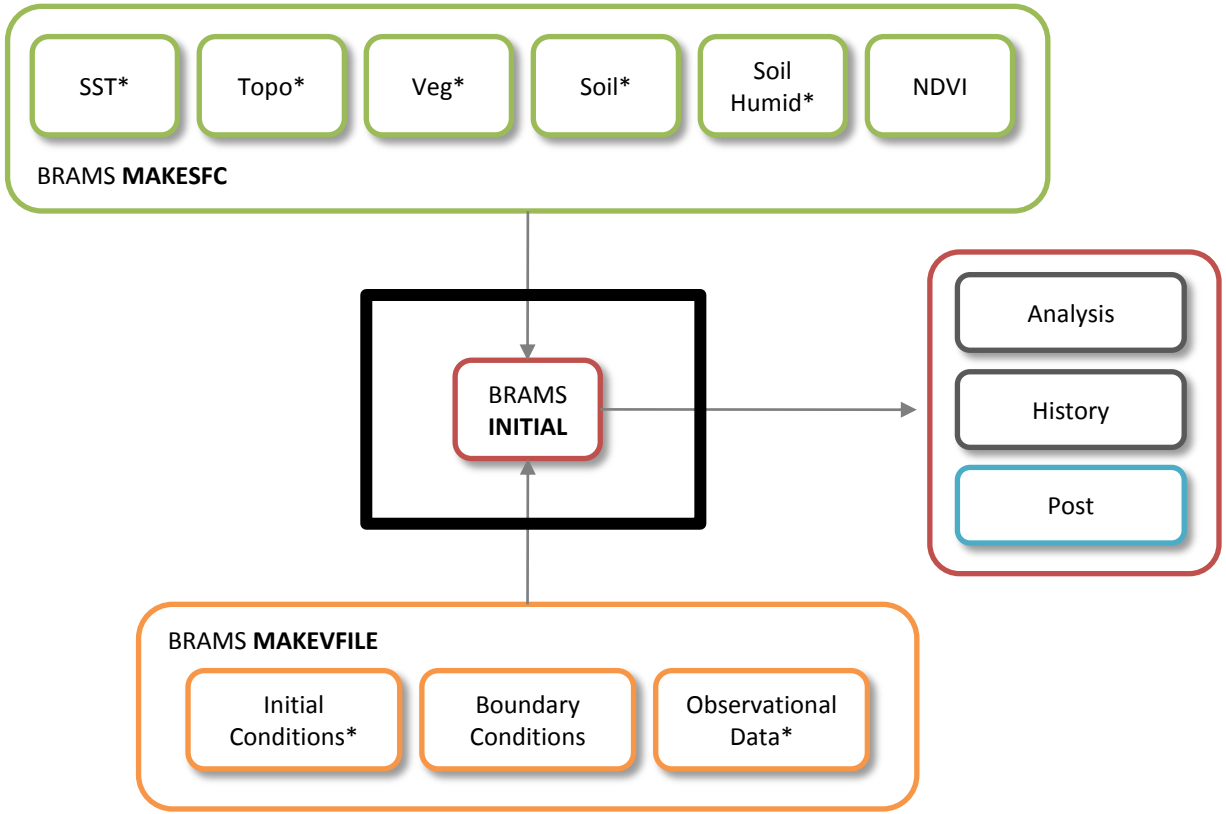

Output data	Size [MB]
MAKESFC & MAKEVFILE	12
Analysis	1100
History	147
Post-process	8
Total	1267

Outline

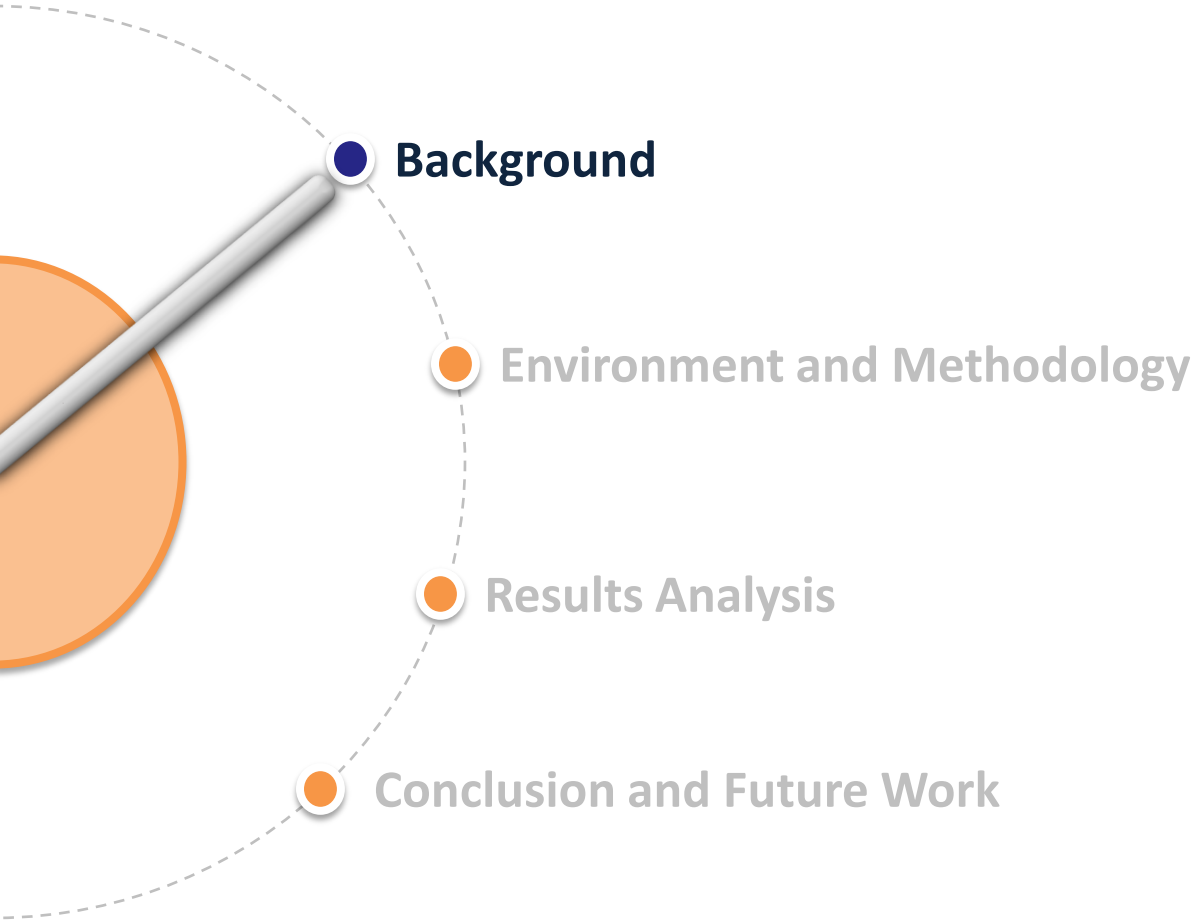
Background
Environment and Methodology
Results Analysis
Conclusion and Future Work

Controlling the cloud service deployment of BRAMS

- Initial cloud parameters
- *Instantiable* BRAMS VM
- Deploying VMs
- Workflow control
- Removing parts

- Reproducible experiments
- Reduce deployment time
- Reduce costs
- Automate service
- Usable in parallel applications

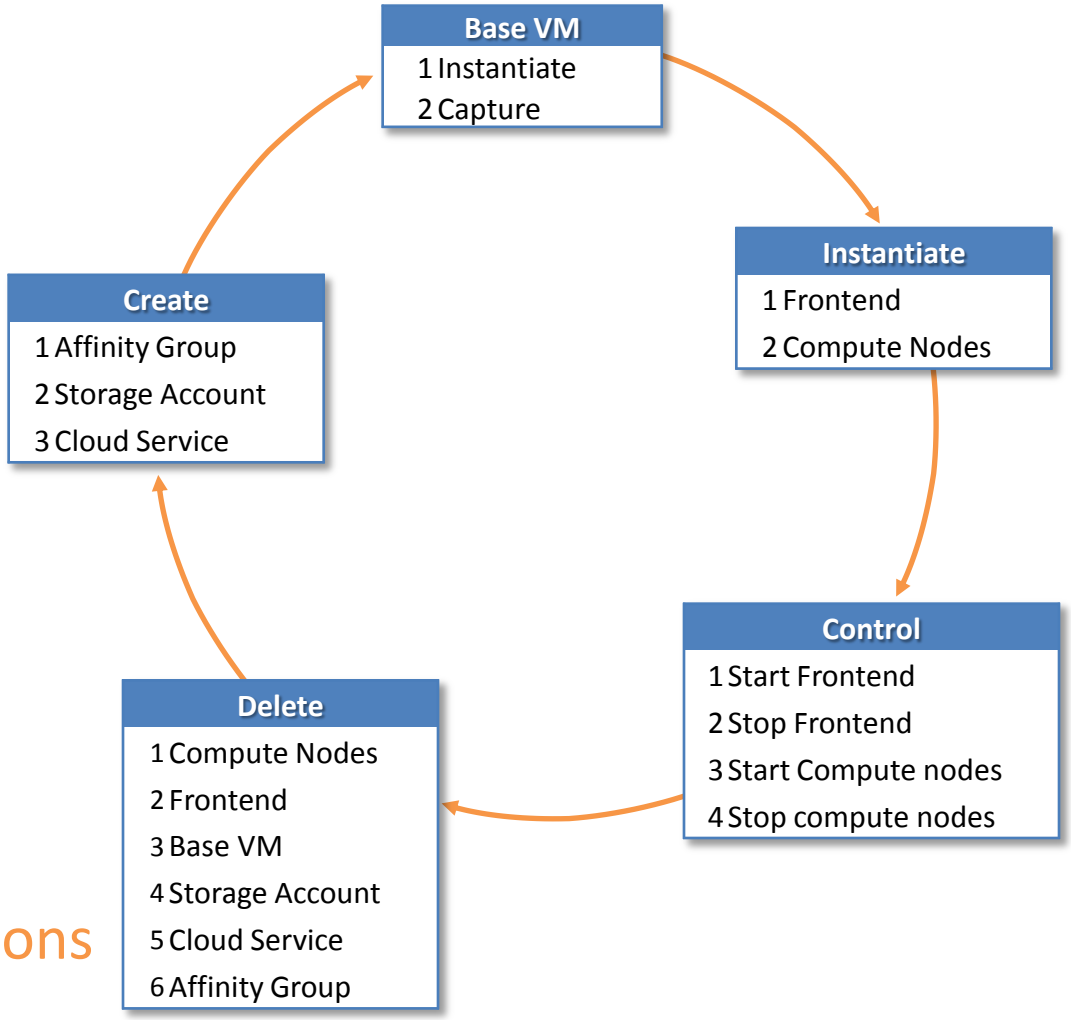

BRAMS Deployment Architecture

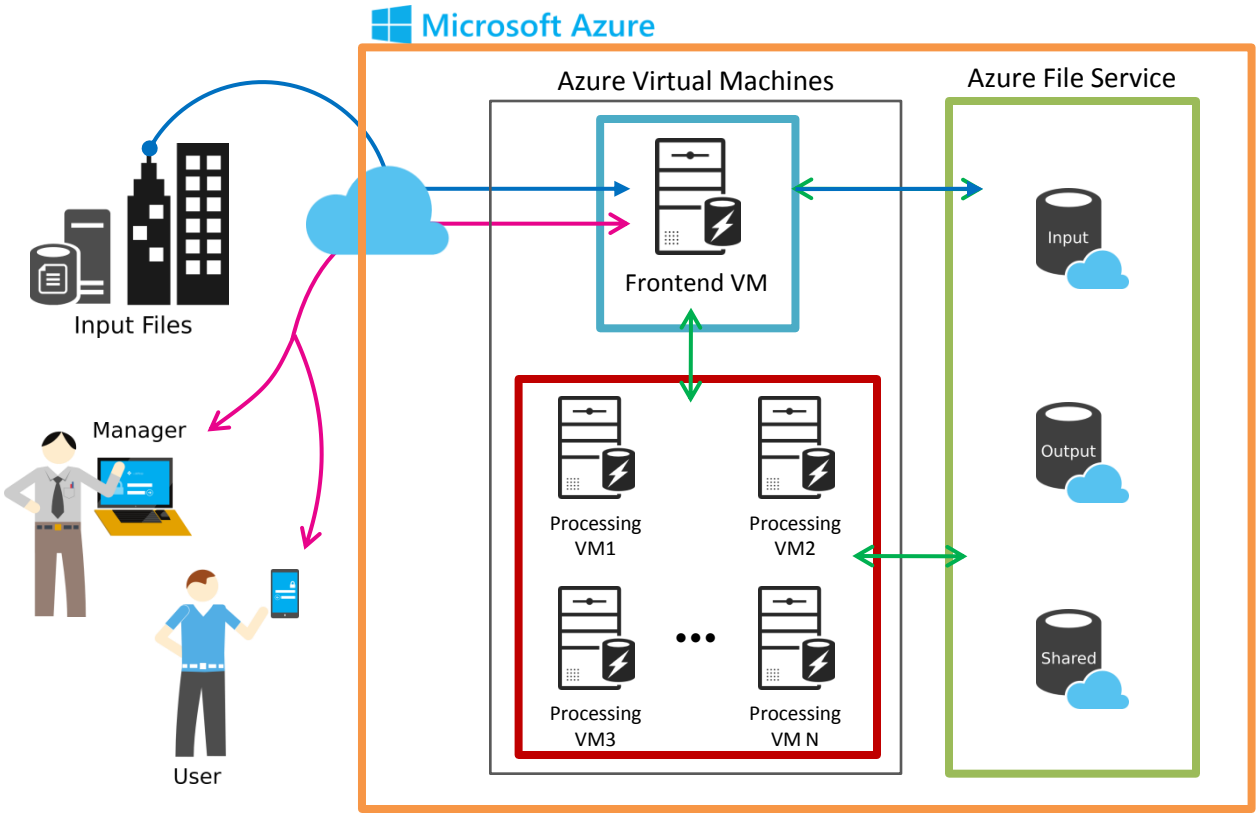

- **Frontend** gets input data
- **Storage** using Azure File Service
- **Processing** VMs created on-demand
- **Networking** in the same datacenter location

Environments

Characteristics	Cloud Instance Sizes			
	A4	A7	A8	A9
Processor Model	Xeon E5-2660		Xeon E5-2670	
Processor Speed (GHz)	2,2		2,6	
Number of CPU Cores	8	8	8	16
Memory (GB)	14	56	56	112
Networking (Mbps)	800	2000	5000	10000

Forecast: 72h

Resolution: 50 km

Perimeter: 6000 km

Area: 1500x1500 km

Outline

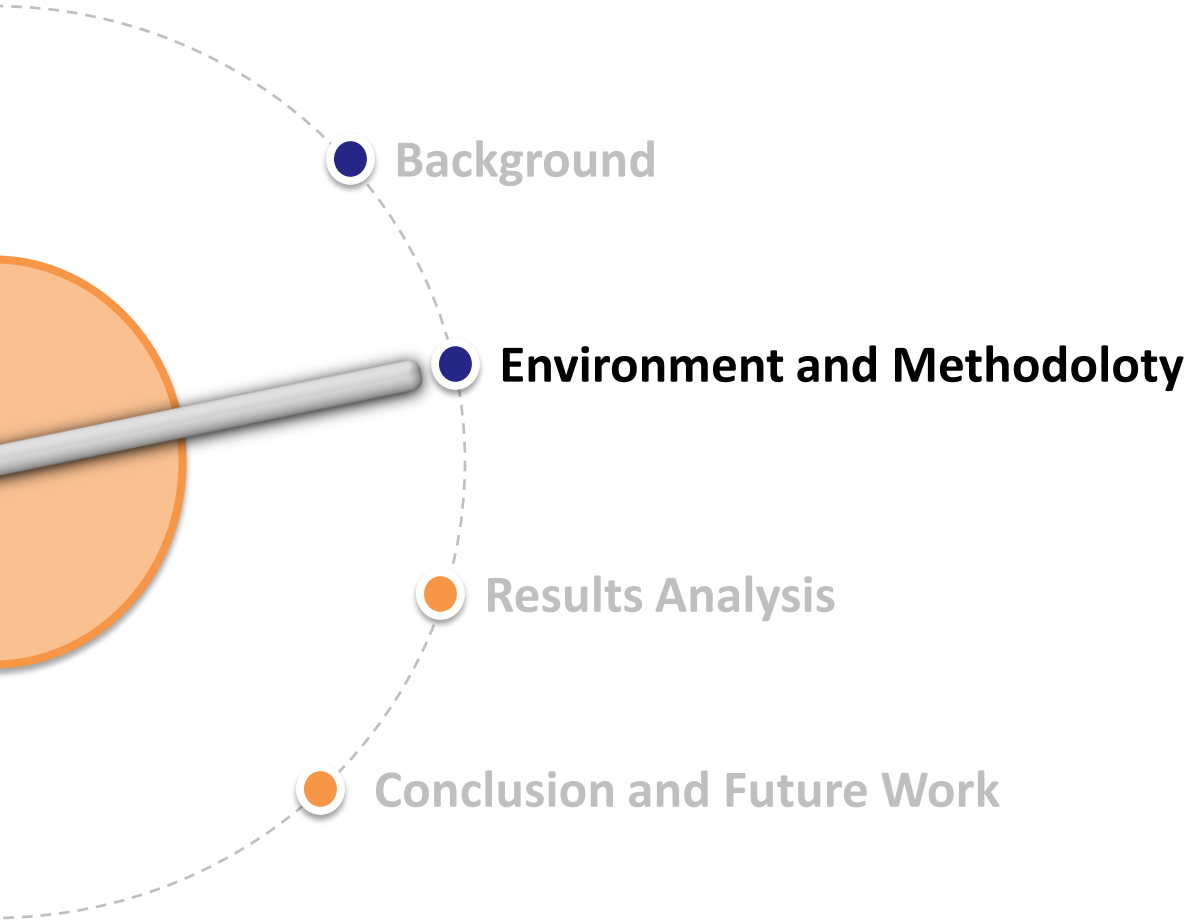
Background
Environment and Methodoloty
Results Analysis
Conclusion and Future Work

Performance - Scalability

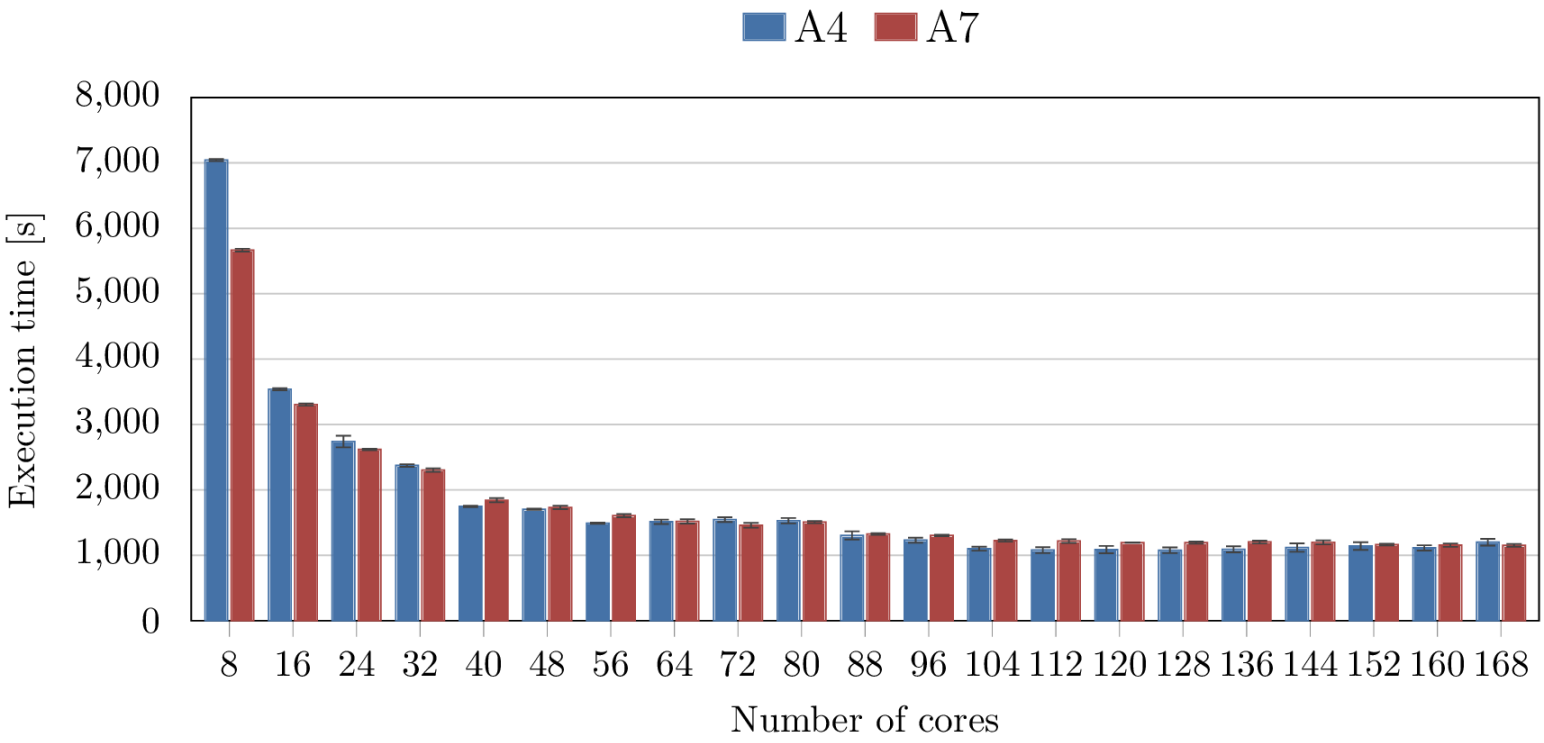

168 cores, 20KM

- Low performance variation.
- Not scaling well beyond 88 cores.
- After 56 cores time reduction is less than two percent
- Caused by Networking?

Performance - Scalability

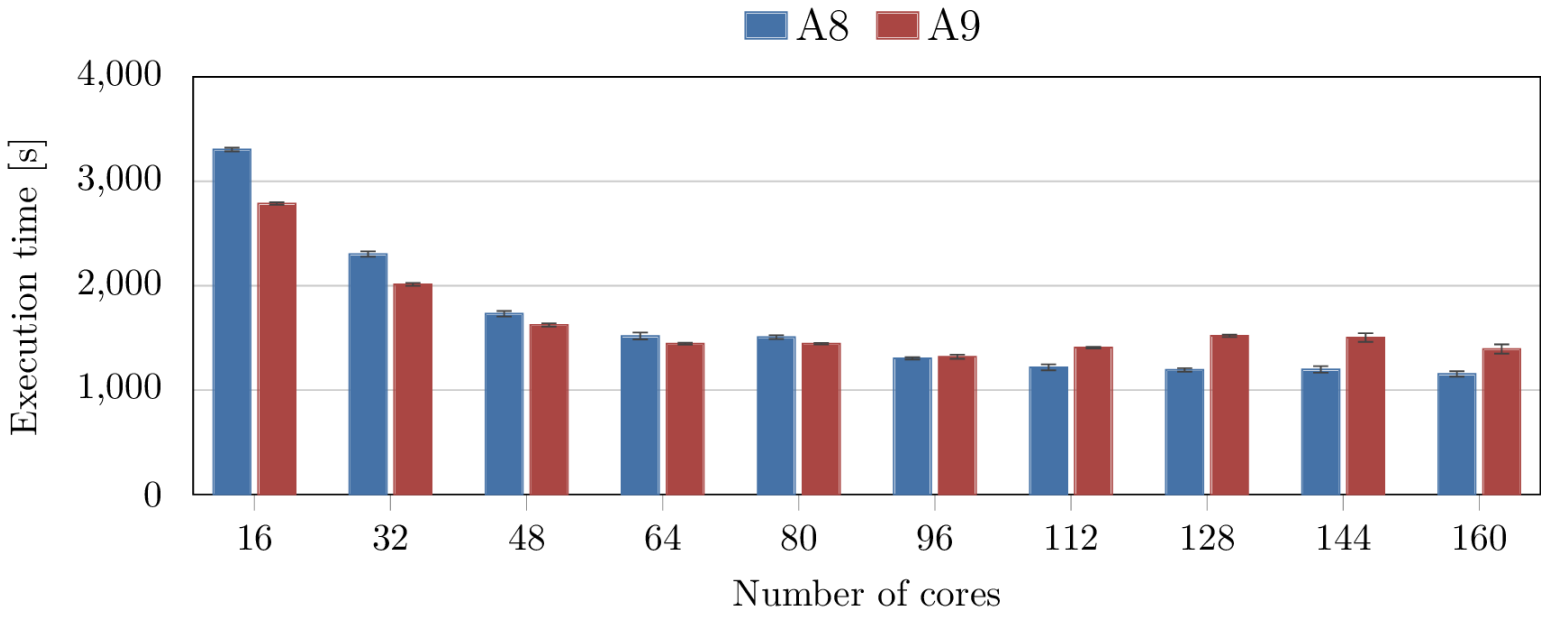

160 cores, 5KM

- Similar behavior
- Scaling after 80 cores minimal
- Application limitation, communication.

Post-processing

- Modifying the final part of the workflow.
- Using GrADS to generate contours or shaded plots.
- Using BING maps to obtain background
- Forecast: 72h
- Resolution: 50 km
- Area: 1500x1500 km

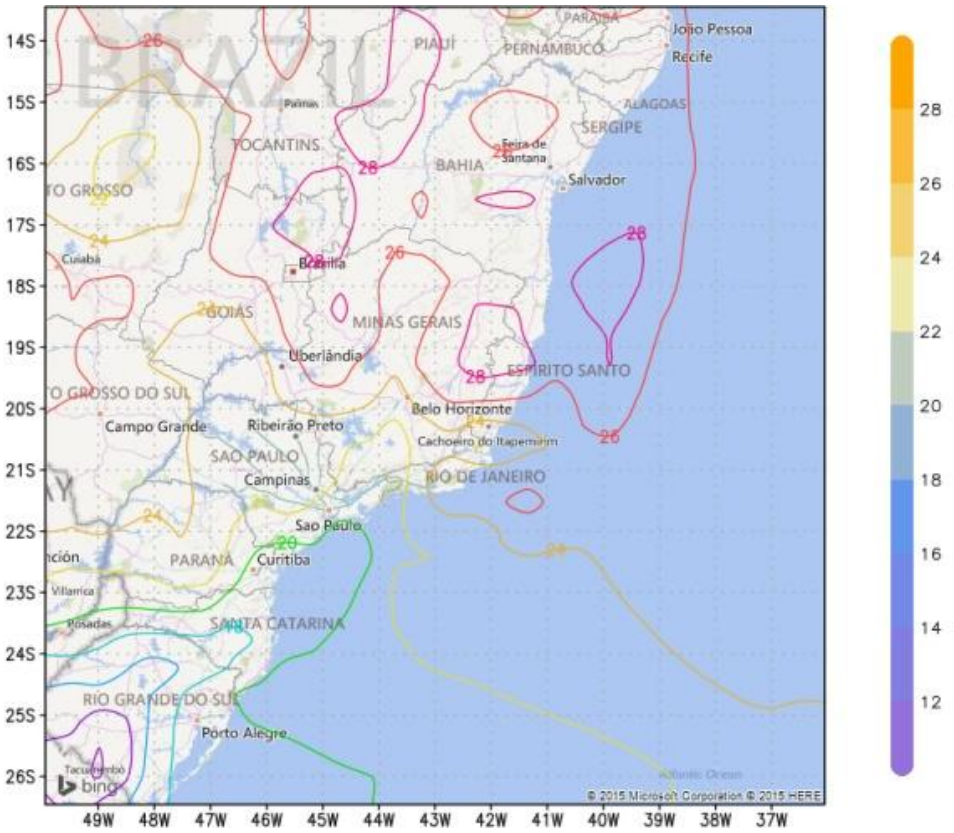

Outline

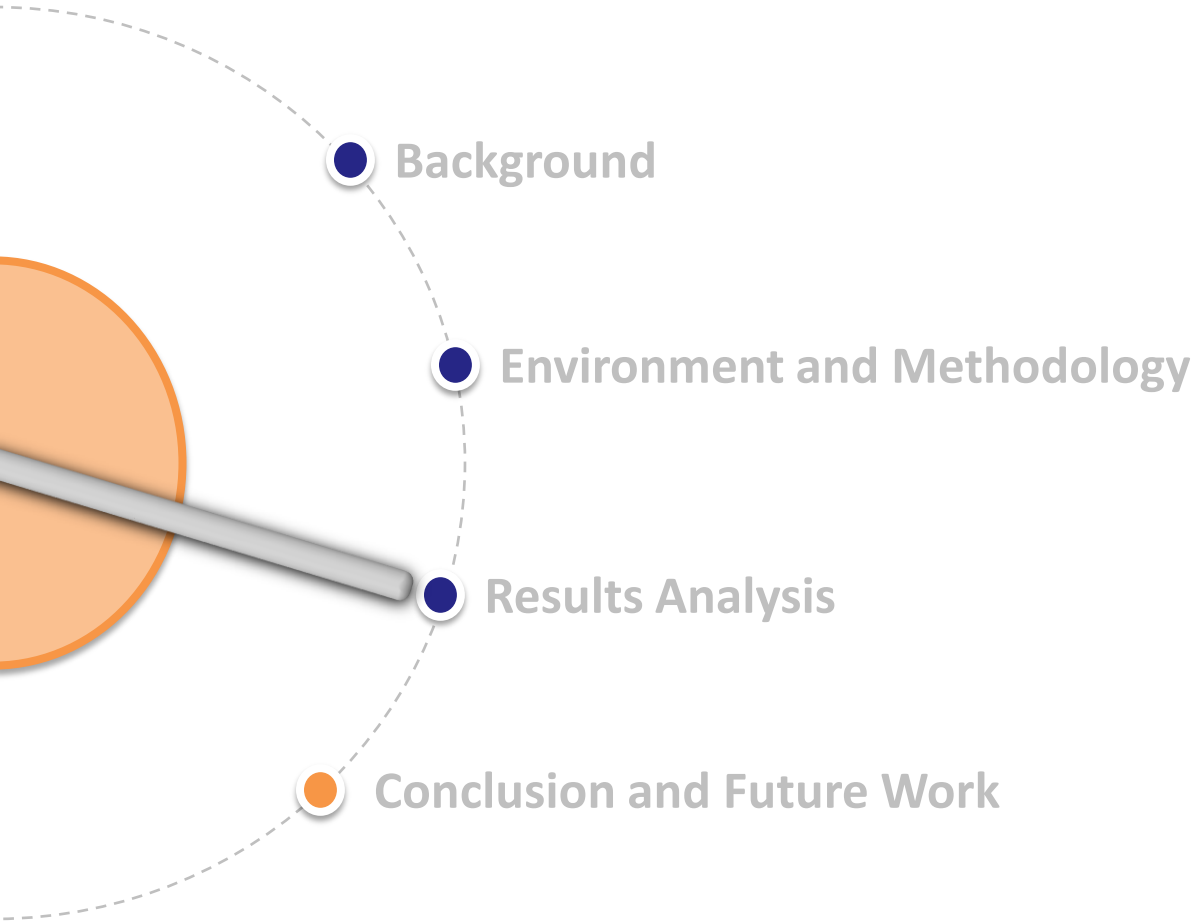
Background
Environment and Methodology
Results Analysis
Conclusion and Future Work

Conclusion

- Cloud infrastructure offers possibilities to move HPC applications from legacy code reducing.
- Variability in performance was low in the cloud, an important characteristic for HPC.
- Scaling BRAMS up to 168 cores, beyond 88 cores <2% exec time reduction
- Scaling model resolution from 50km to 5km in the cloud deployment.

Future Work

- Obtaining results from other cloud providers as Amazon EC2 and Google Compute.
- Capturing more metrics to cover all the aspects of the execution and try to improve the performance of this HPC application in a cloud environment.
- Parallelize some of the sequential steps in BRAMS workflow if possible.
- Test new instances like Linux VMs with Infiniband.

Performance Analysis of a Numerical Weather Prediction Application in Microsoft Azure

Emmanuell D Carreño, Eduardo Roloff, Jimmy V. Sanchez, and Philippe O. A. Navaux

WSPPD 2015 - XIII Workshop de Processamento Paralelo e Distribuído

UFRGS
UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

GPPD
Grupo de Processamento Paralelo e Distribuído

CNPq
Conselho Nacional de Desenvolvimento Científico e Tecnológico

.inf
INSTITUTO DE INFORMÁTICA
UFRGS

THANK YOU!