TESTING THE TPB'S TRAVEL MODELS AND MOVES MODELS ON AWS CLOUD SERVERS

Ray Ngo and Jinchul (JC) Park
TPB Transportation Engineers

TPB Travel Forecasting Subcommittee
November 18, 2022

Agenda Item 2

Metropolitan Washington
Council of Governments

COG's Migration from On-Premises to Cloud

MOVES
Gen2 Travel Model

Gen2 Travel Model
Gen3 Travel Model
GIS server
SAS server

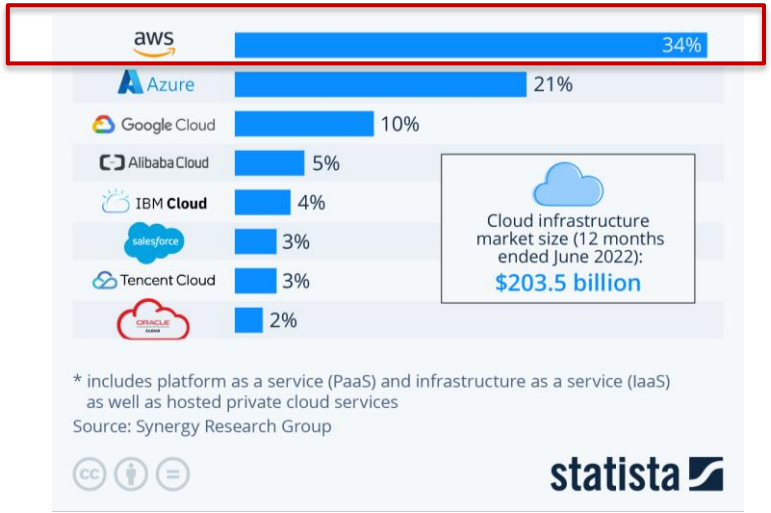

On-Premises and AWS Comparison

ON-PREM		AWS
More prone to reboots and shutdown		Fewer interruptions
Fixed and predictable costs		Potentially variable costs (depending on server strategy)
Slower and harder to acquire and set up a new or a replacement server		Fast to create multiple servers tailored to different needs
Higher cost for the third-party backup services		Lower backup cost and better disaster recoverability
COG IT staff to maintain data center		Amazon to mange data center and COG IT to provide support

Server Migration Considerations

- Existing and future needs
- Data security and compliance
- Cost analysis
- Staff expertise
- **Testing**

COG Server Migration Status

- Completed:
  - SAS, SQL servers
  - Some Commuter Connections servers
  - **MOVES servers**

- On going:
  - GIS server
  - Some Commuter Connections servers
  - **Travel Model Servers**

Testing MOVES Models on AWS

MOVES Modeling in AWS

- MOtor Vehicle Emission Simulator (MOVES) is software developed by the EPA for estimating mobile-source air pollution, both criteria pollutants & greenhouse gas (GHG) emissions
- COG/TPB staff developed a process, called the Emission Modeling Process (EMP), to create MOVES inputs, execute MOVES, and summarize outputs.
- In batch processing, EMP Version 2.0.1 coupled with MOVES3.0.4 can estimate nitrogen oxide (NOx) and volatile organic compounds (VOCs) emissions for the ozone non-attainment area and GHG emissions for the TPB Planning Area using an automated process.
  - 7-hour computer run time is necessary for estimating one year of NOx, VOC and GHG emissions
- Reasons for using AWS:
  - Interruptions resulting from power outages and software updates can more easily be avoided to ensure a complete run.
  - To save time and effort, having multiple computers is a must for any analysis when multiple scenarios are evaluated.
- Five "instances" are set up as MOVES dedicated computers in AWS: 2 for MOVES2014B, 2 for MOVES3.0.4 and 1 for MOVES3.0.3

MOVES Modeling in AWS (Cont'd)

- Performance: Estimating NOx, VOC and GHG emissions for one analysis year with MOVES3.0.4

Computer	Total Hours	Cost ($/hr)
Desktop (I7 CPU; 3.6 GHz; 8GB RAM; SSD)	7	0
Amazon AWS (m5zn.xlarge; 4 virtual CPUs [2 CPUs], 16 GB RAM, >4.5 GHz clock speed)	3.5	$0.51/hr
Run-Time Difference	~ 50% reduction	

- Results: Computers in AWS produce identical model estimates, but with 50% shorter runtime
- Evaluations:
  - AWS provides a stable environment and faster model results
  - User costs are subject to change, so there could be budgetary concerns in the future
  - Same-day analysis is possible and multiple simultaneous runs are executable → so helpful for production work or scenario testing

MOVES Modeling in AWS (Cont'd)

- Comments from User's Experience:
  - Sharing/Isolating Disk and Software: Users need to set up each instance correctly when cloning one instance to another. Each instance was isolated as each MOVES model uses a specific version of JAVA, Python, SQL software.
  - Software Version Control: To control software version, automatic software updates were disabled for all instances.
  - Simultaneous Access: Simultaneous access from multiple users should be available.
  - AWS User Cost: It is expected to increase based on usage time and disk storage space. Users need to devise a better file management system and set up a notification when a model run is over.
  - Handbook: Having a handbook about AWS will be helpful, which includes general use, budgets, troubleshooting, etc.

Testing Travel Models on AWS

On-Prem Travel Model Servers

Server	tms5	tms6	tms7	tms8
Processor name(s)	Intel Xeon X5690	Intel Xeon E5-2690	Intel Xeon E5-2687W V3	Intel Xeon Gold 6146
Clock speed of processor (GHz)	3.47	2.9	3.1	3.2
Total number of physical cores	12	16	20	24
Total Physical Memory (GB)	12	224	288	256

AWS Instance for Testing

Server	tms8	r6id.8xlarge	r6i.8xlarge	r6a.8xlarge
Processor name(s)	Intel Xeon	Intel Xeon	Intel Xeon	AMD
Clock speed of processor (GHz)	3.2	3.5	3.5	3.6
Total number of physical cores	24	16	16	16
Total Physical Memory (GB)	256	256	256	256
EBS bandwidth		Local + EBS drive (10 Gbps)	EBS only (up to 10 Gbps)	EBS only (6.6 Gbps)
On-demand cost ($/hour)		3.89	3.49	3.29

EBS: Elastic Block Storage

EBS volume type: gp3 (General Purpose SSD – high I/O per second and throughput)

Runtime Reduction on AWS (vs tms8): Simplified Gen3 (Phase 1) Model

How do the selected instances perform when compared with COG's tms8?

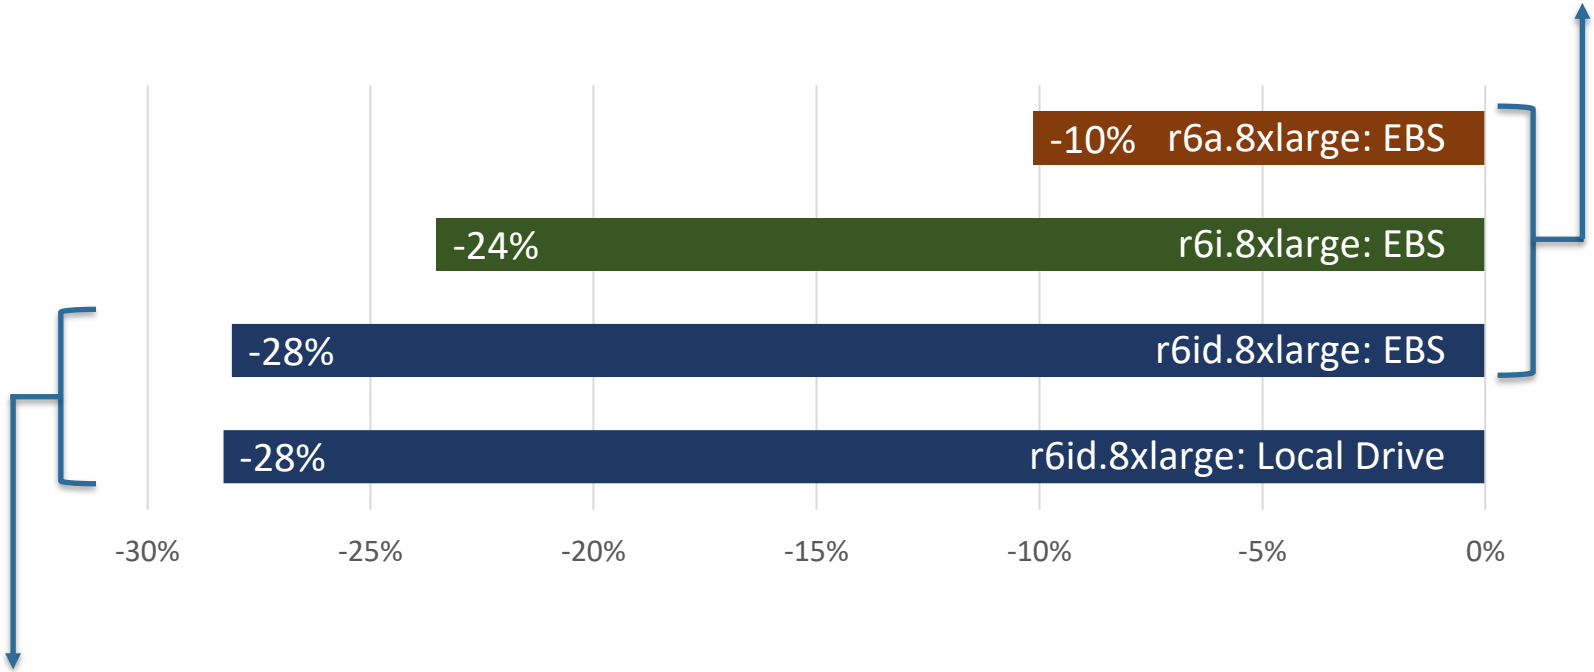

Should we conduct the model runs on an instance's local or EBS drives?

Runtime Reduction on AWS (vs tms7 and tms8): Full Gen2 and Gen3 (Phase 1) Models

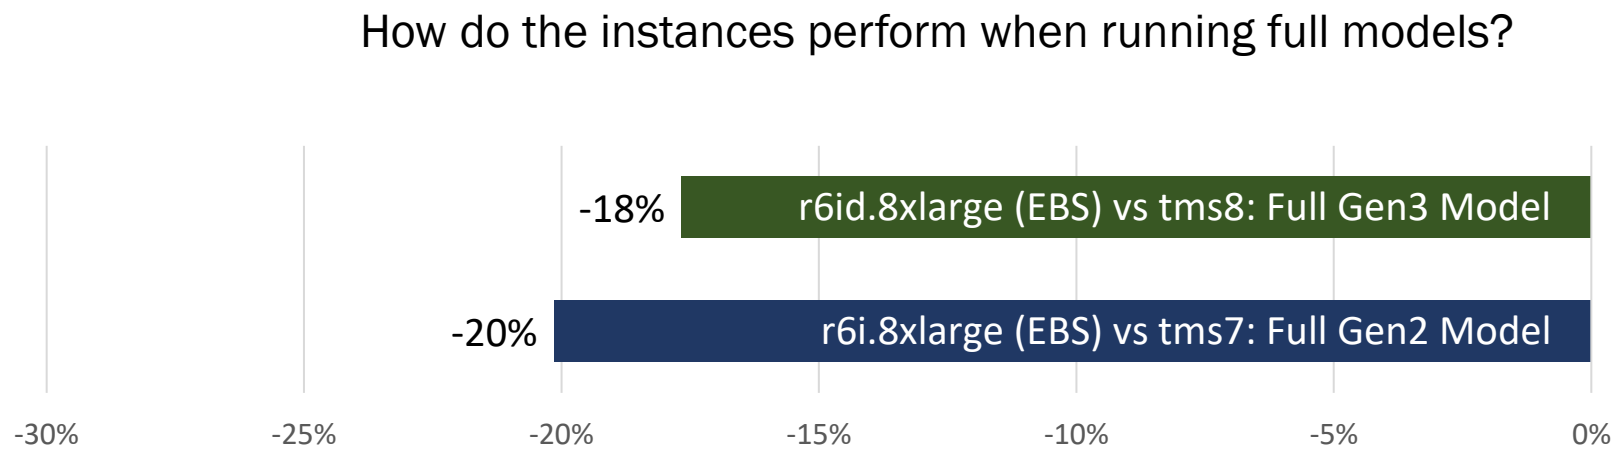

AWS Instance for Gen2 Model

Server	tms6	m6a.4xlarge
Processor name(s)	Intel Xeon E5-2690	AMD
Clock speed of processor (GHz)	2.9	3.6
Total number of physical cores	16	8
Total Physical Memory (GB)	224	64
EBS bandwidth		EBS only (up to 10 Gbps)
On-demand cost ($/hour)		1.43

Runtime Reduction on AWS (vs tms6): Full Gen2

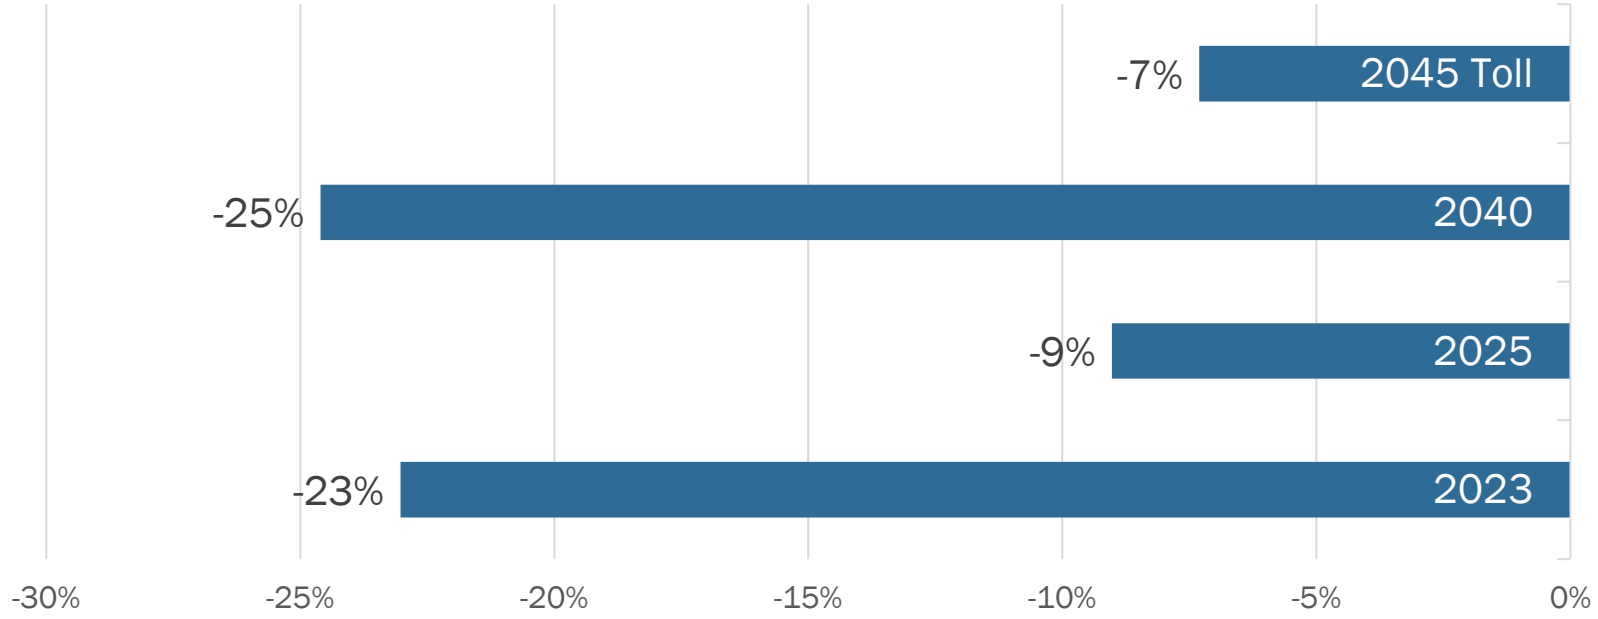

Final AWS Computer Specs

Computer Specs	Hi	Medium	Low
Usage	Gen3 Model	Gen2 Model	Other Tasks
Instance type	r6i.8xlarge	m6a.4xlarge	t3.2xlarge
Processor name(s)	Intel Xeon	AMD	Intel Xeon
Clock speed of processor (GHz)	3.5	3.6	3.1
Total number of physical cores	16	8	≤8
Total Physical Memory (GB)	256	64	32
EBS bandwidth	EBS only (up to 10 Gbps)	EBS only (up to 10 Gbps)	EBS only (up to 5 Gbps)
On-demand cost ($/hour)	3.49	1.43	0.48

Findings and Comments

- Staff chose three instance choices for each on-demand AWS server.
- Our models runs are faster on the selected servers than our on-prem.
- Two concurrent runs took 40% longer than one run on the medium-specs instance server.
- COG/TPB modelers have been shown the basics of running models on AWS servers.
- Migrating from on-prem servers to AWS requires cooperation of COG IT and COG DTP staff.
- On-demand instances are very quick to set up, but they could be very expensive. Auto shutdown feature helps reduce cost.
- The performance of an AWS instance depends on many factors, including chips and storage type. Gp3 EBS drive and S3 Glacier are recommended for storage.

Next Steps

- Set up six AWS on-demand servers, each with three instance types, for travel models.
- Transfer data on on-premises servers to AWS using Snowball by mid-December.

Acknowledgements

Mark Moran, Feng Xie, Daniel Son (Travel Forecasting and Emissions Analysis, COG/TPB)

Lionel Bernard, Robert Francois, Daniel Gobeze (Information Technology and Facilities Management, COG/TPB)

Jinchul (JC) Park (for MOVES Models)

Principle Transportation Engineer
(202) 962-3320
jpark@mwcog.org

Ray Ngo (for Travel Demand Models)

Senior Transportation Engineer
(202) 962-3231
rngo@mwcog.org

mwcog.org

777 North Capitol Street NE, Suite 300
Washington, DC 20002

Metropolitan Washington
Council of Governments