Orbiters and Clouds: Bringing Mars Reconnaissance Orbiter Data to the People

Mike Fienen[1], Sam Congdon[2], Jay Laura[3], Brendan Wakefield[2]

[1] USGS Upper Midwest Water Science Center, Madison, WI

[2]USGS Cloud Hosting Solutions, Bozeman,MT

[3] USGS Astrogeology Science Center, Flagstaff, AZ

USGS
science for a changing world

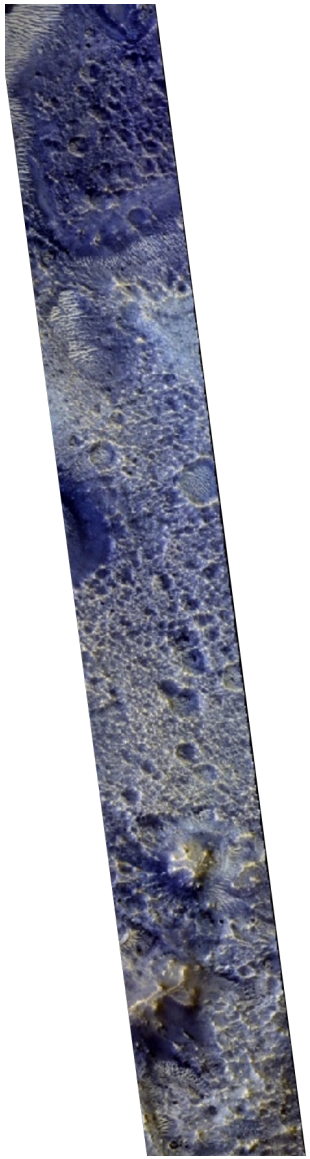

Overview

Building on the legacy of `condor_annex`, we created an open-source, pool-on-demand AWS-based platform for high-throughput computing with HTCondor

We illustrate with an example of making Mars Reconnaissance Orbiter HiRISE imagery data more open and accessible for science

USGS
science for a changing world

On Premises Resource Seems Ideal –
we have been running a large cluster for years

0. Can be difficult to buy and maintain on-premises computational resources
1. Consider the full costs on premises
   - buying equipment
   - the human cost of maintenance
   - power/cooling
   - carbon footprint of lacking efficiency
2. Who gets priority?

USGS
science for a changing world

What about the cloud?

The Good Stuff:

0. Can scale big
1. Can configure things pretty much how you want
2. “Bring the compute to the data”
3. Everyone gets their own cluster

The Bad Stuff:

0. IT nerdery overload to get going 🤓
1. Ubiquity of error has a side effect

USGS
science for a changing world

cloud-ht2c
Alpha Version 0.1

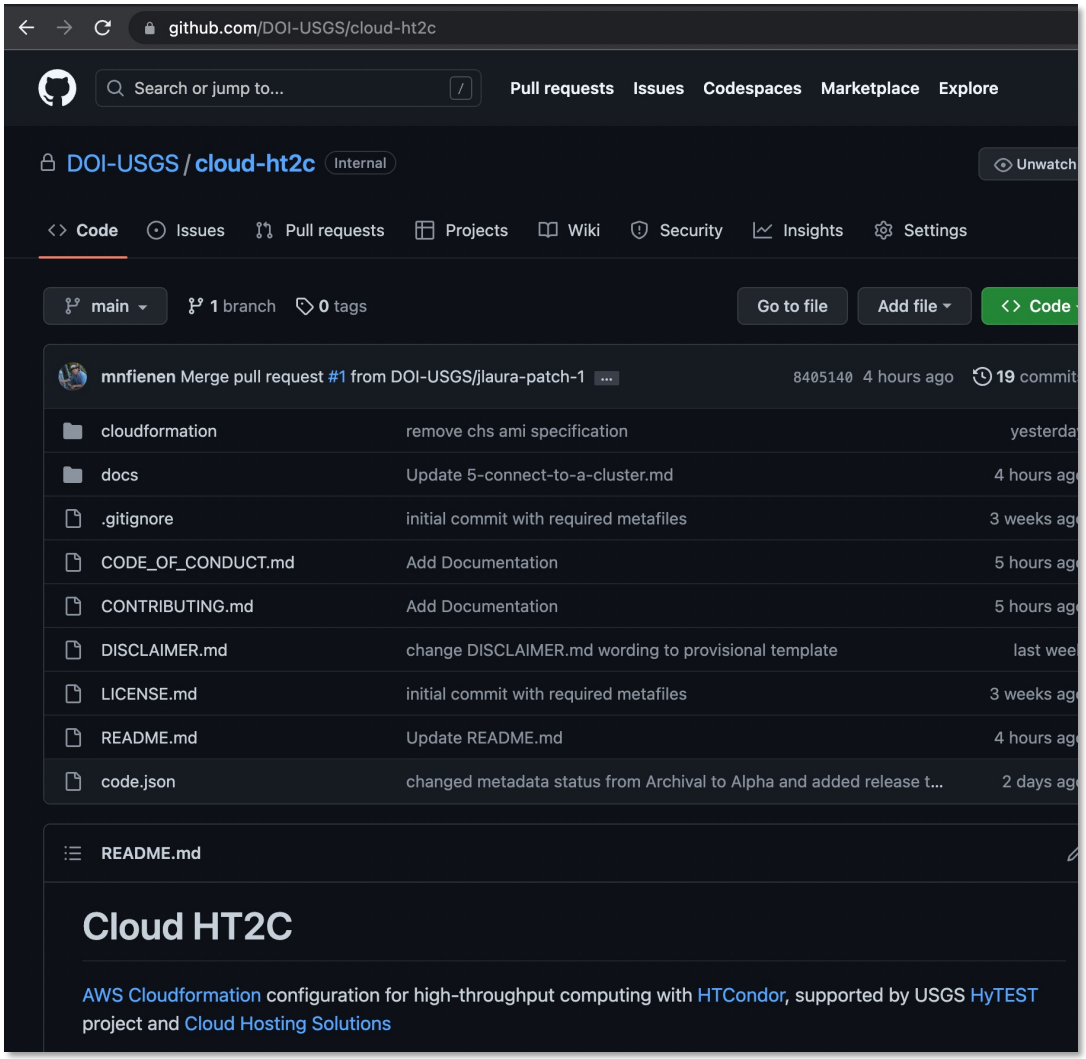

USGS
science for a changing world

Architecture of `cloud-ht2c`

0. Configure an AMI (Amazon Machine Image)
   - AMI=blueprint.
   - Single running instance=one building from the blueprint
   - There are a few prereqs for the AMI
1. Spin up a cluster on demand using cloudformation yml configuration
   - decide which AMI and architecture to base the control node and each worker on
   - decide how many workers you want
   - choose optional spot option
   - assign timeout options
2. Cloudformation coordinates a flexible setup with a single central manager and workers can come and go
   - HTCondor uses TCP/IP so port 9618 is open
   - Also ports 9700-9710 are available for PEST/PEST++

USGS
science for a changing world

Cloud-based hybrid of HPC and HTC

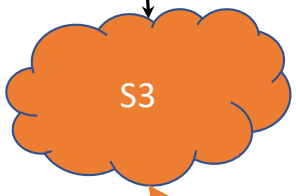

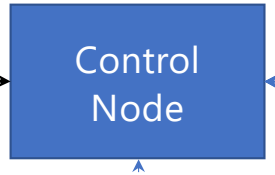

/mnt/condor_working

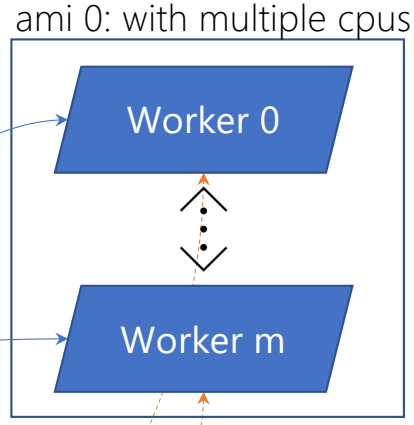

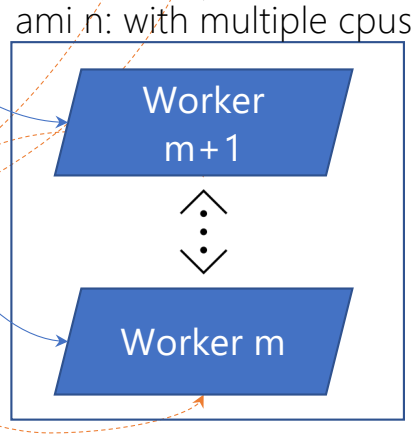

Configuring a cloud-based compute cluster

1.

Parameters

Network Configuration

Name Prefix
Enter a unique prefix to help identify your cluster's specific resources
mnf_demo
Malformed input-Parameter pNamePrefix must match pattern [\.\-_/#A-Za-z0-9]+

VPC
VPC Id for deployment
vpc-0af42fd592a1efc5b - csr-vpc-ephemeralDev

Subnets (Select 3)
Subnets for deployment
Choose options
subnet-06e8cb2055f9477b3 ✕ subnet-0f29464029b7f677c ✕
subnet-0bcd68692d4d2279d ✕

CIDR Range with access to the Control Node
Input the CIDR range of the selected VPC.
0.0.0.0/0
Default value can be overridden.

2.

Controller Configuration

Control Node Instance Type
Instance type for the Control Node
c5.xlarge

Control Node Volume Size
The EBS volume size of the Control Node. NOTE: AWS Elastic File System storage is comes configured on the Control Node.
256

Control Node AMI ID
The AMI to run the HTCondor Controller node on. Default is the chs-centos-7-awstools AMI.
ami-02b40f4187a3c2fe4
Default value can be overridden.

3.

Worker Configuration

Number of Worker Nodes to launch
The number of Worker Nodes to spin up. Update this parameter to alter available compute power or save on costs.
2000

Worker Node Instance Type
Instance type for Worker Nodes
c5.large

Worker Node Volume Size
The EBS volume size of the Worker Nodes.
64
Default value can be overridden.

Worker Node AMI ID
The AMI to run the HTCondor Worker nodes on. Default is the chs-amazon-linux-2-htcondo
ami-02b40f4187a3c2fe4
Default value can be overridden.

Worker Node Timeout
The Idle Time in seconds a Worker Node is allowed to exist before it is terminated. Set value to 0 to disable auto termination of idle Worker Nodes.
600
Default value can be overridden.

Worker Node Kill Switch Timestamp
If specified, executes a hard scale-in of all worker nodes in the AutoScaling Group to zero instances at the designated timestamp (UTC, 24-hr time notation). Enter "NULL" if you do not wish to enable a hard termination timestamp for your Worker Nodes. Timestamp must be formatted as "YYYY-MM-ddTHH:mm:ssZ", e.g. "2023-01-01T22:00:00Z"
2023-07-10T22:00:00Z

Worker Node Timeout
The Idle Time in seconds a Worker Nodes.
600
Default value can be overridden.

Worker Node Kill Switch Timestamp
If specified, executes a hard scale-in of all v
hr time notation). Enter "NULL" if you do no
formatted as "YYYY-MM-ddTHH:mm:ssZ", e
2023-07-10T22:00:00Z

4.

Spot Instance Configuration

Use Spot Instances
Whether to use Spot Instance for the worker nodes
true

(Optional) Spot Instance Type 1
An additional instance type to use in the Spot Pool for Worker Instances
c5.xlarge

(Optional) Spot Instance Type 2
An additional instance type to use in the Spot Pool for Worker Instances
c5.2xlarge

(Optional) Spot Instance Type 3
An additional instance type to use in the Spot Pool for Worker Instances
c5.4xlarge

Thanks to Todd Miller for the lowdown!

USGS
science for a changing world

How to deploy dependencies and model files

Amazon EFS

`/mnt/condor_working`

wine

Amazon Machine Image (AMI)

"AMI is the OG Docker"

conda/**conda-pack**

Package conda environments for redistribution

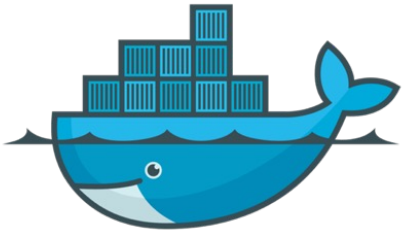

docker

USGS science for a changing world

Increasing complexity

Mars Reconnaissance Orbiter and Analysis Ready Data

Different missions have different specs/needs.

Big issues are image format (JPEG 2000 are not registered in space so harder to find).
Cloud-optimized geoTIFF is explorable via browser.

Common projections enhance interoperability

Processing and serving the entire dataset empowers scientists to focus on the science rather than endless fussing with a complex image/geo-processing workflow

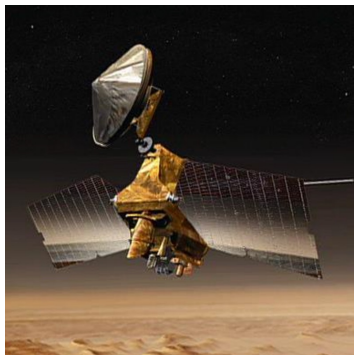

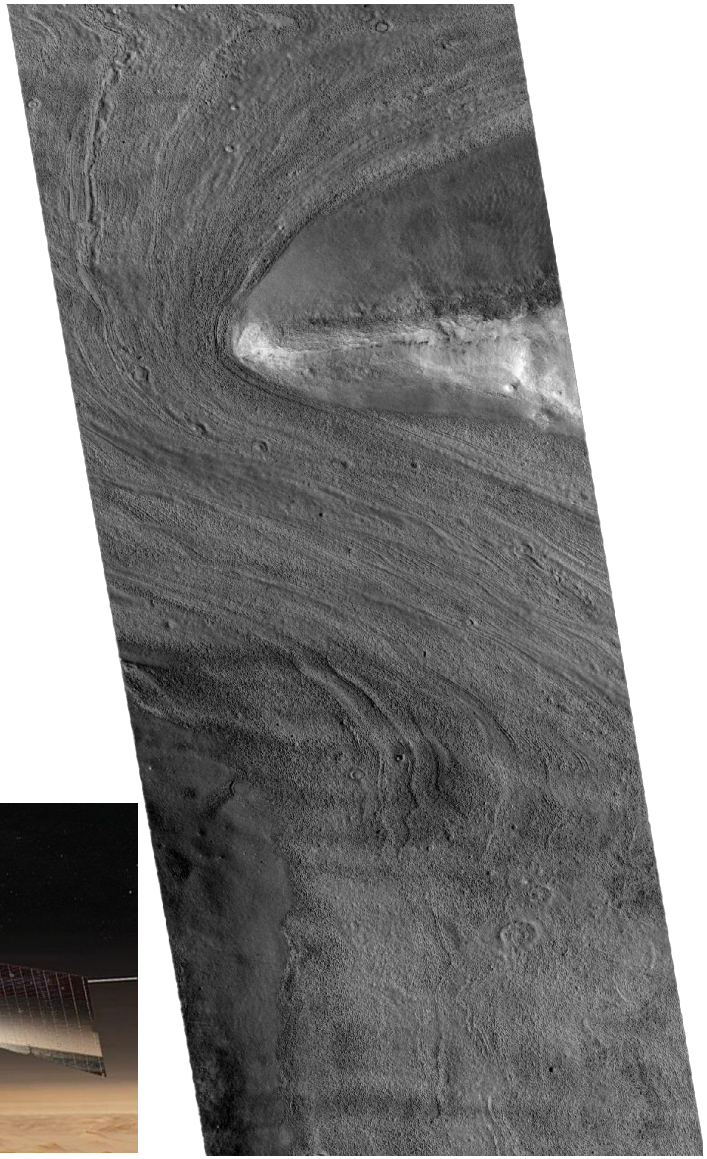

USGS
science for a changing world

Mars Reconnaissance Orbiter and Analysis Ready Data by the Numbers

155,317 total images to process

Input from USGS/NASA S3 bucket: 100 TB

Output to NASA/Amazon Registry of Open Data S3: 114 TB

2,000 AMIs supporting 4,664 concurrent jobs

Total wall time: about 4 hours
(not quite enough time to watch *The Martian* twice!)

Cost: about 10% of planned (#spotInstnaces)

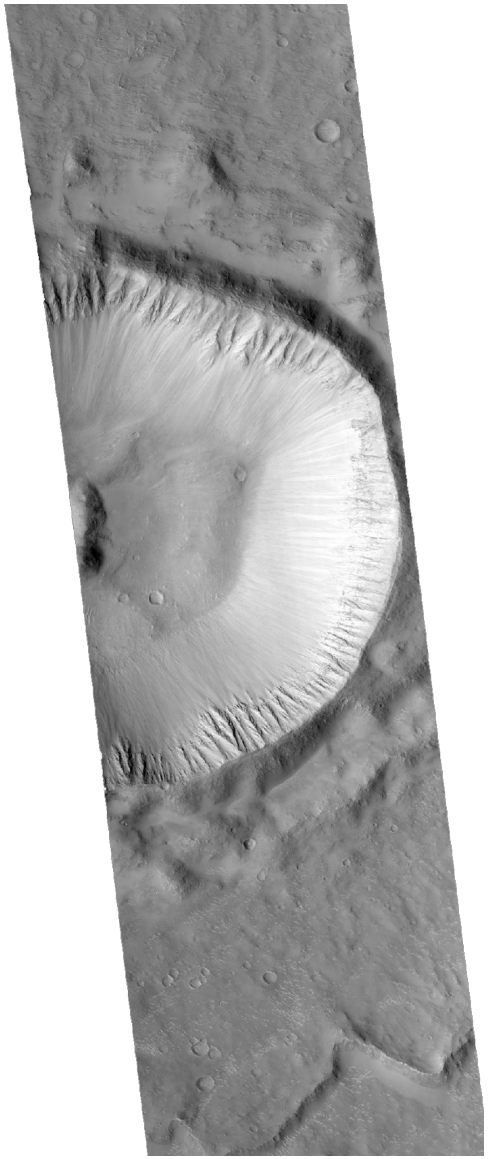

USGS
science for a changing world

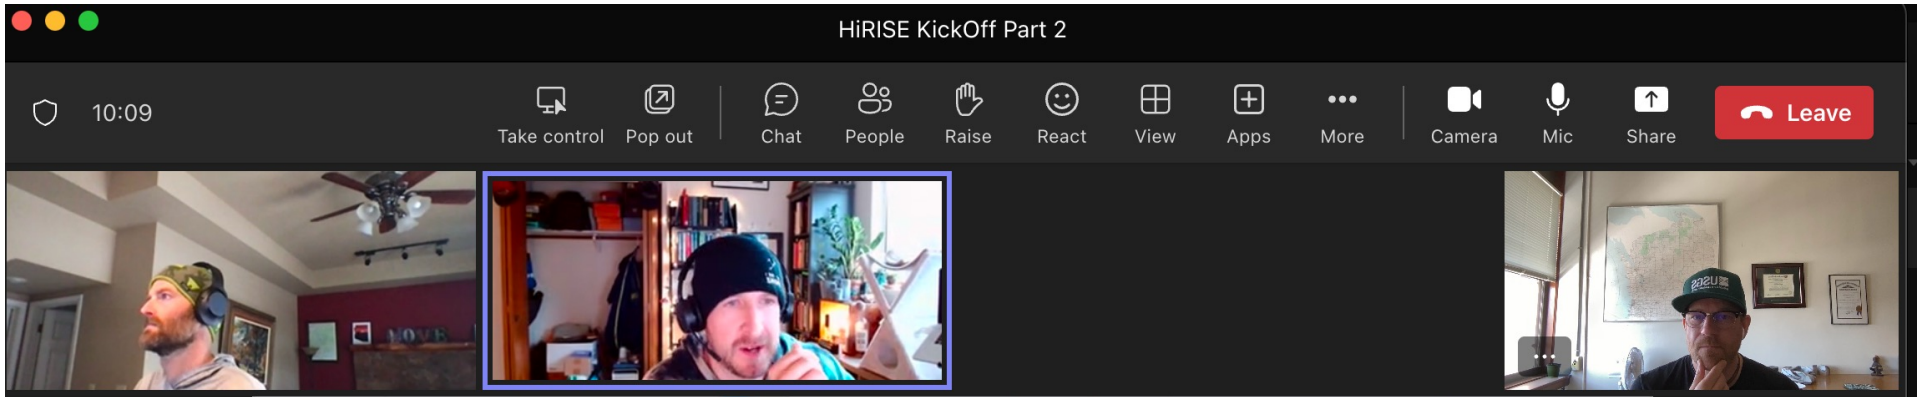

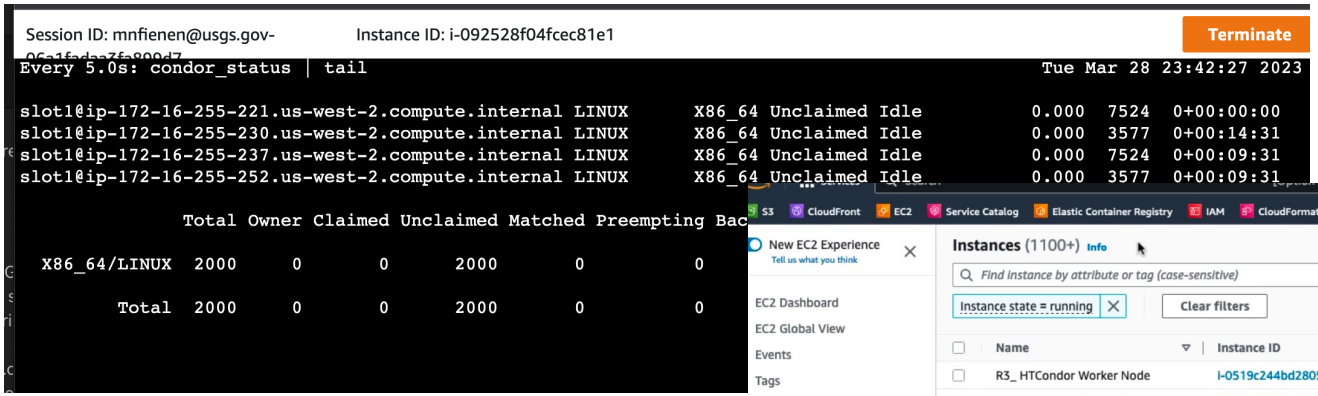

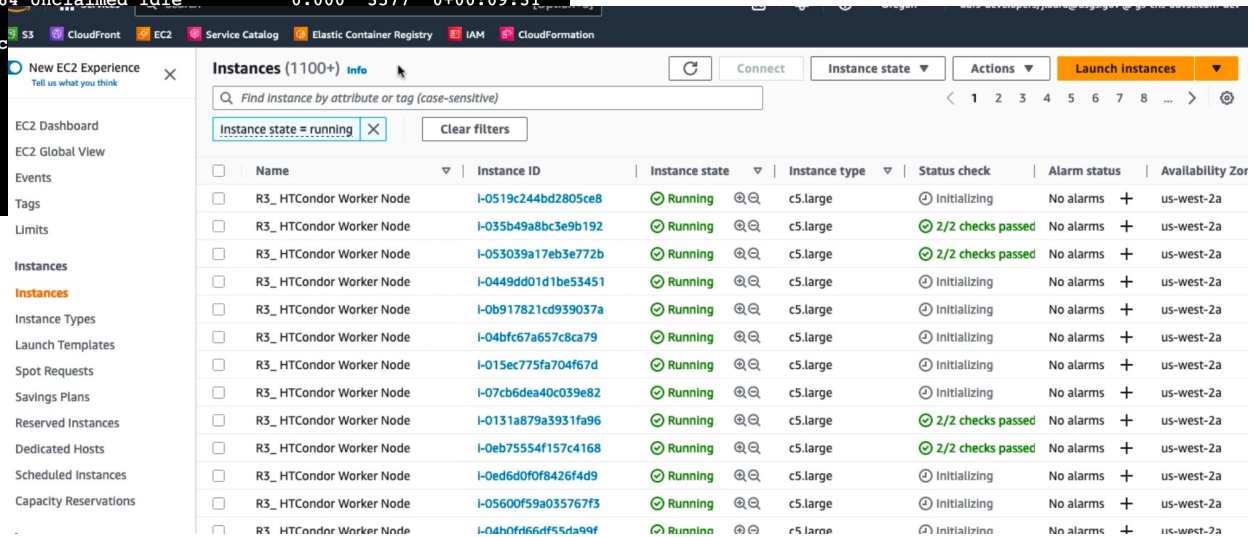

USGS
science for a changing world

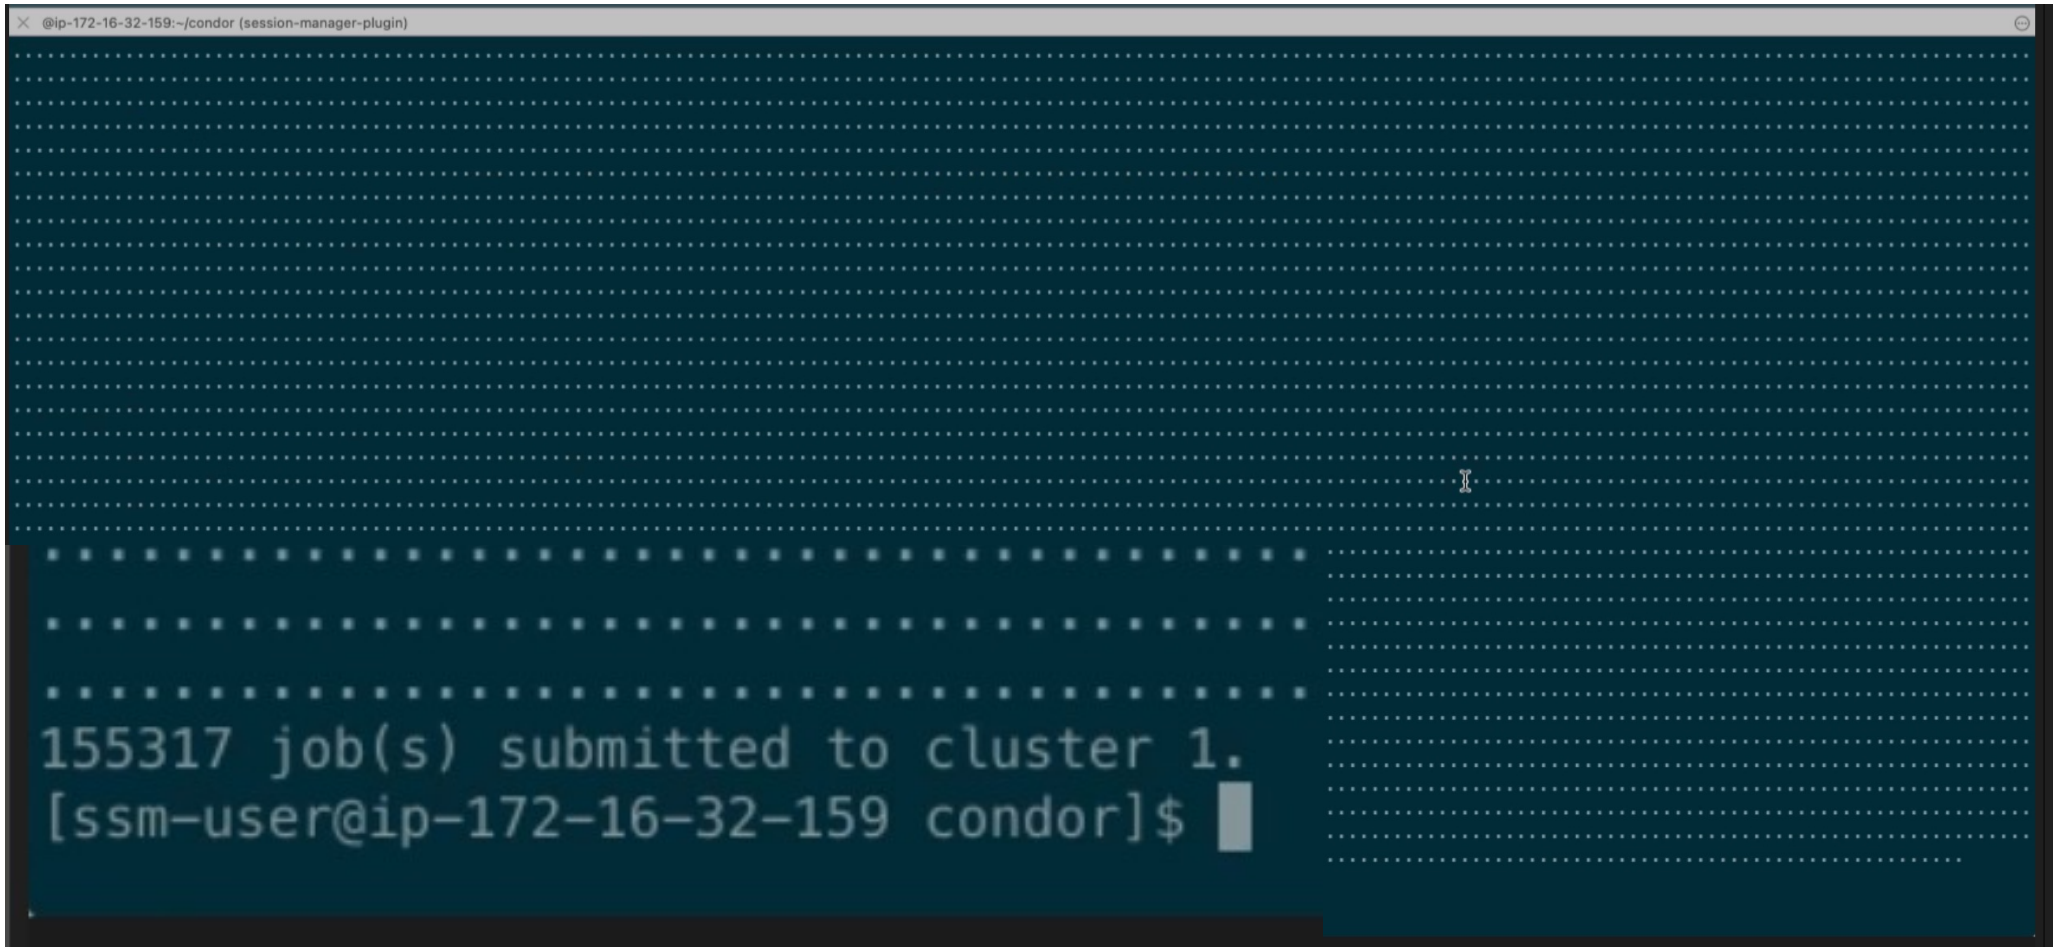

That's a lot of dots!

USGS
science for a changing world

```
slot1@ip-172-16-255-175.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000 1124  0+00:26:49
slot1_1@ip-172-16-255-175.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:26
slot1_2@ip-172-16-255-175.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:04
slot1@ip-172-16-255-176.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000 2620  0+00:23:38
slot1_1@ip-172-16-255-176.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:00
slot1_2@ip-172-16-255-176.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:00
slot1_3@ip-172-16-255-176.us-west-2.compute.internal LINUX      X86_64 Claimed   Idle      0.000 3200  0+00:00:00
slot1_4@ip-172-16-255-176.us-west-2.compute.internal LINUX      X86_64 Claimed   Idle      0.000 3200  0+00:00:00
slot1@ip-172-16-255-188.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000 1124  0+00:20:49
slot1_1@ip-172-16-255-188.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:28
slot1_2@ip-172-16-255-188.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:05
slot1@ip-172-16-255-221.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000 1124  0+00:19:29
slot1_1@ip-172-16-255-221.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:00
slot1_2@ip-172-16-255-221.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:00
slot1@ip-172-16-255-230.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000  377  0+00:32:38
slot1_1@ip-172-16-255-230.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:06
slot1@ip-172-16-255-237.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000 1124  0+00:29:01
slot1_1@ip-172-16-255-237.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:26
slot1_2@ip-172-16-255-237.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:07
slot1@ip-172-16-255-252.us-west-2.compute.internal   LINUX      X86_64 Unclaimed Idle      0.000  377  0+00:28:50
slot1_1@ip-172-16-255-252.us-west-2.compute.internal LINUX      X86_64 Claimed   Busy      0.000 3200  0+00:00:10

               Total Owner Claimed Unclaimed Matched Preempting Backfill  Drain

  X86_64/LINUX  6664     0    4664      2000       0          0        0      0

         Total  6664     0    4664      2000       0          0        0      0
[ssm-user@ip-172-16-32-159 ~]$
```

```
-- Schedd: ip-172-16-32-159.us-west-2.compute.internal : <172.16.32.159:9618?... @ 03/29/23 12:15:49
OWNER    BATCH_NAME    SUBMITTED   DONE   RUN    IDLE   HOLD  TOTAL JOB_IDS
ssm-user ID: 1        3/28 23:58 155293      _      _     24 155317 1.1422-89355

Total for query: 24 jobs; 0 completed, 0 removed, 0 idle, 0 running, 24 held, 0 suspended
Total for ssm-user: 24 jobs; 0 completed, 0 removed, 0 idle, 0 running, 24 held, 0 suspended
Total for all users: 24 jobs; 0 completed, 0 removed, 0 idle, 0 running, 24 held, 0 suspended
```

USGS
science for a changing world

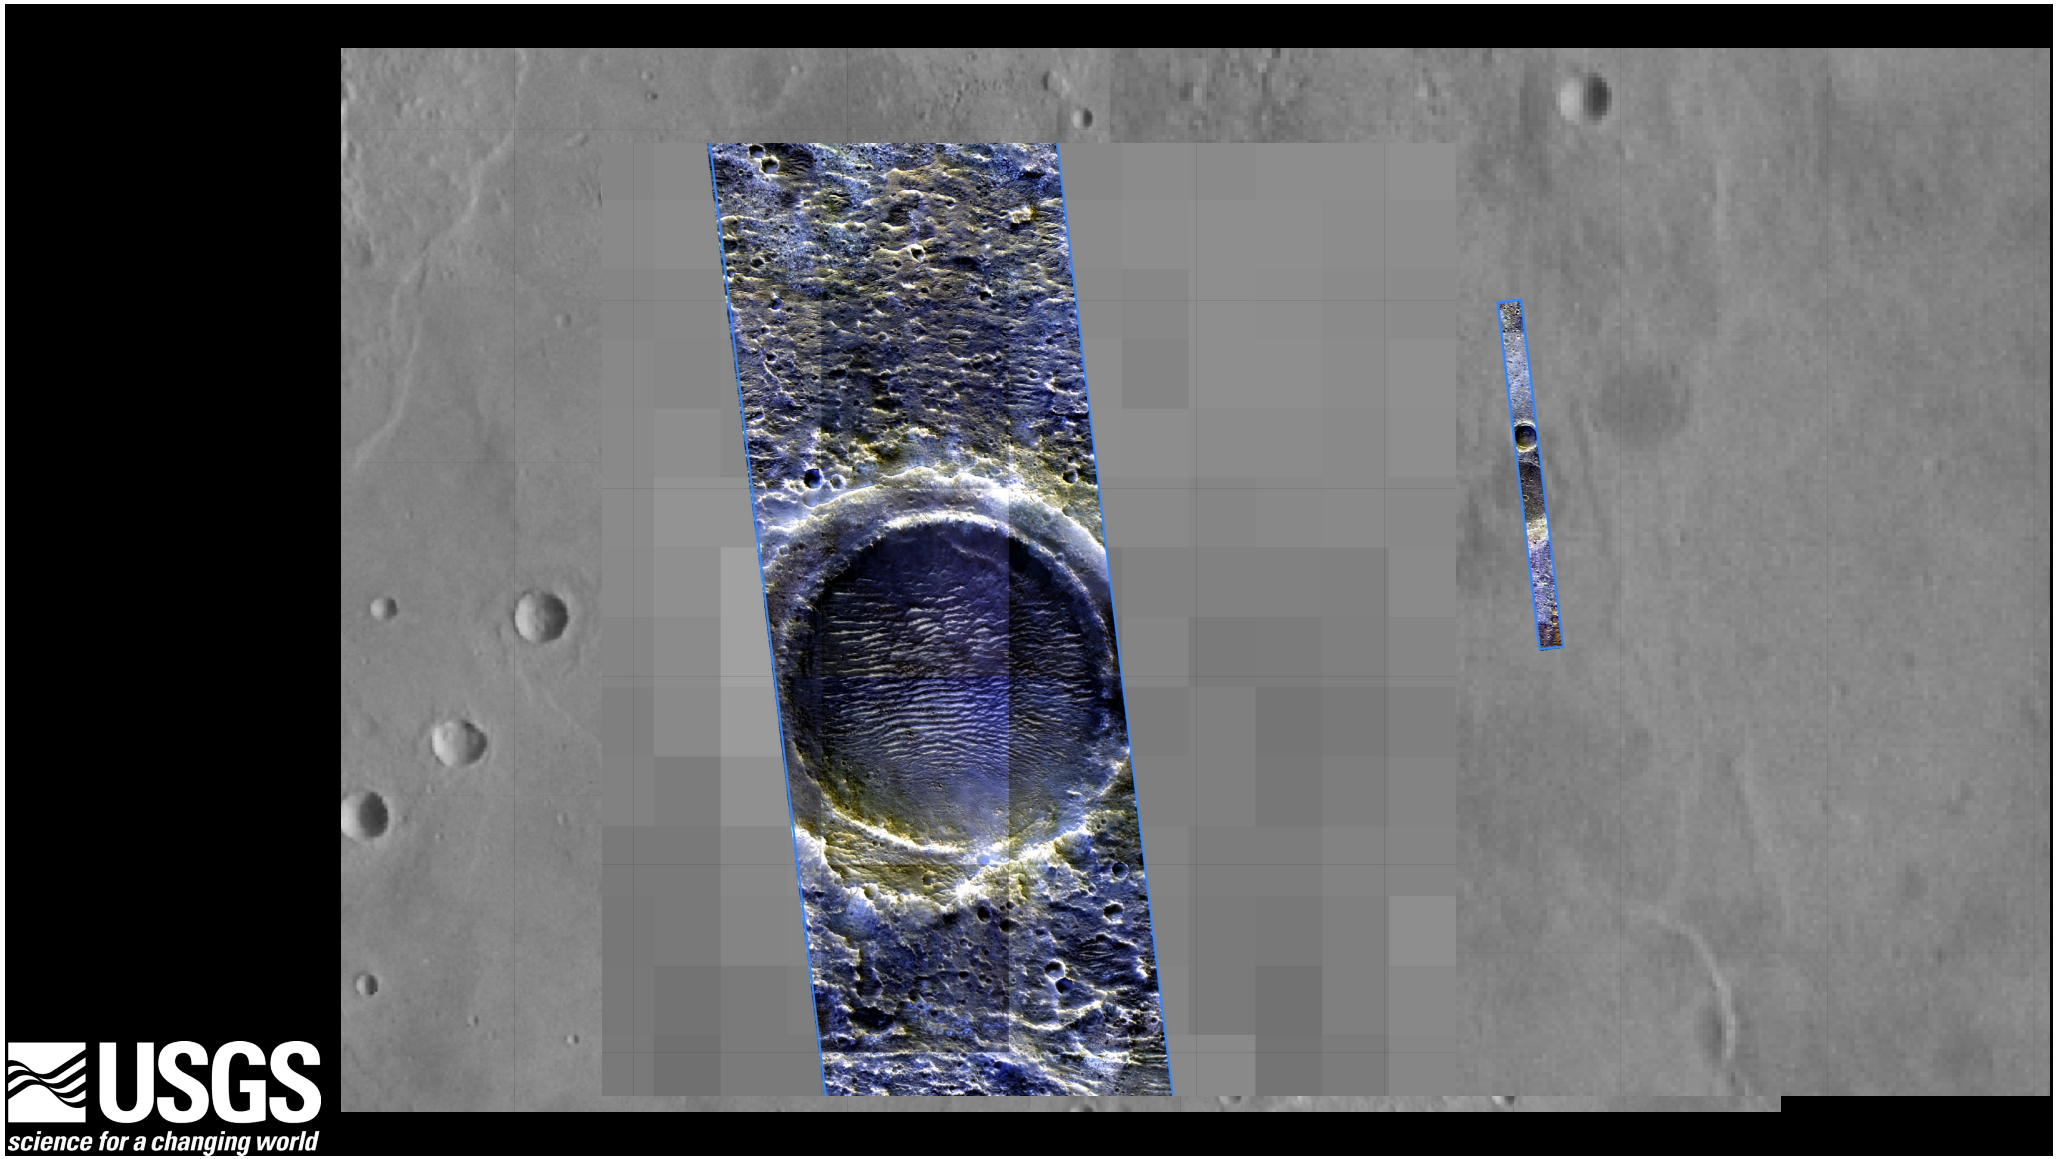

USGS
science for a changing world

Let's talk about MONEY

On-demand pricing

Instance name	On-Demand hourly rate	vCPU	Memory	Storage	Netwo per
c5.large	$0.085	2	4 GiB	EBS Only	
c5.xlarge	$0.17	4	8 GiB	E	
c5.2xlarge	$0.34	8	16 GiB	E	
c5.4xlarge	$0.68	16	32 GiB		
c5.9xlarge	$1.53	36	72 GiB		
c5.12xlarge	$2.04	48	96 GiB		
c5.18xlarge	$3.06	72	144 GiB	E	
c5.24xlarge	$4.08	96	192 GiB	E	
c5.metal	$4.08	96	192 GiB	E	
c5a.large	$0.077	2	4 GiB	E	
c5a.xlarge	$0.154	4	8 GiB	EBS Only	Up t
c5a.2xlarge	$0.308	8	16 GiB	EBS Only	Up to 10 G

Spot pricing

c5.large	$0.0349 per Hour	$0.1242 per Hour
xlarge	$0.0736 per Hour	$0.2484 per Hour
	$0.1356 per Hour	$0.4968 per Hour
	$0.2576 per Hour	$0.9936 per Hour
	$0.5795 per Hour	$2.2355 per Hour
	$0.7747 per Hour	$2.9807 per Hour
	$1.1683 per Hour	$4.471 per Hour
	$1.6719 per Hour	$5.9614 per Hour
	$1.5454 per Hour	$5.9614 per Hour
	$0.0329 per Hour	$0.1242 per Hour
	$0.0649 per Hour	$0.2484 per Hour
	$0.1288 per Hour	$0.4968 per Hour
	$0.2801 per Hour	$0.9936 per Hour
	$0.5887 per Hour	$2.2355 per Hour
12xlarge	$0.7727 per Hour	$2.9807 per Hour
c5d.18xlarge	$1.1671 per Hour	$4.471 per Hour
c5d.24xlarge	$1.5454 per Hour	$5.9614 per Hour

USGS
science for a changing world

Let's talk more about MONEY

Break it down:

each CPU on-demand is ~$0.043/hr

but likely need to reserve some (25%?) CPUs for OS, so **~$0.057/hr**

Spot market? Cost is often ~40%-50% so maybe **$0.02-$0.03/hr**?

BUT! If runs get dropped, can add up – think about forward runtimes

Spot is cheaper for AMIs with more CPUs, but that may hurt performance

USGS
science for a changing world

THANKS!

Go kick the tires on cloud-ht2c – just remember it's Alpha v 0.1 – let me know how it goes mnfienen@usgs.gov

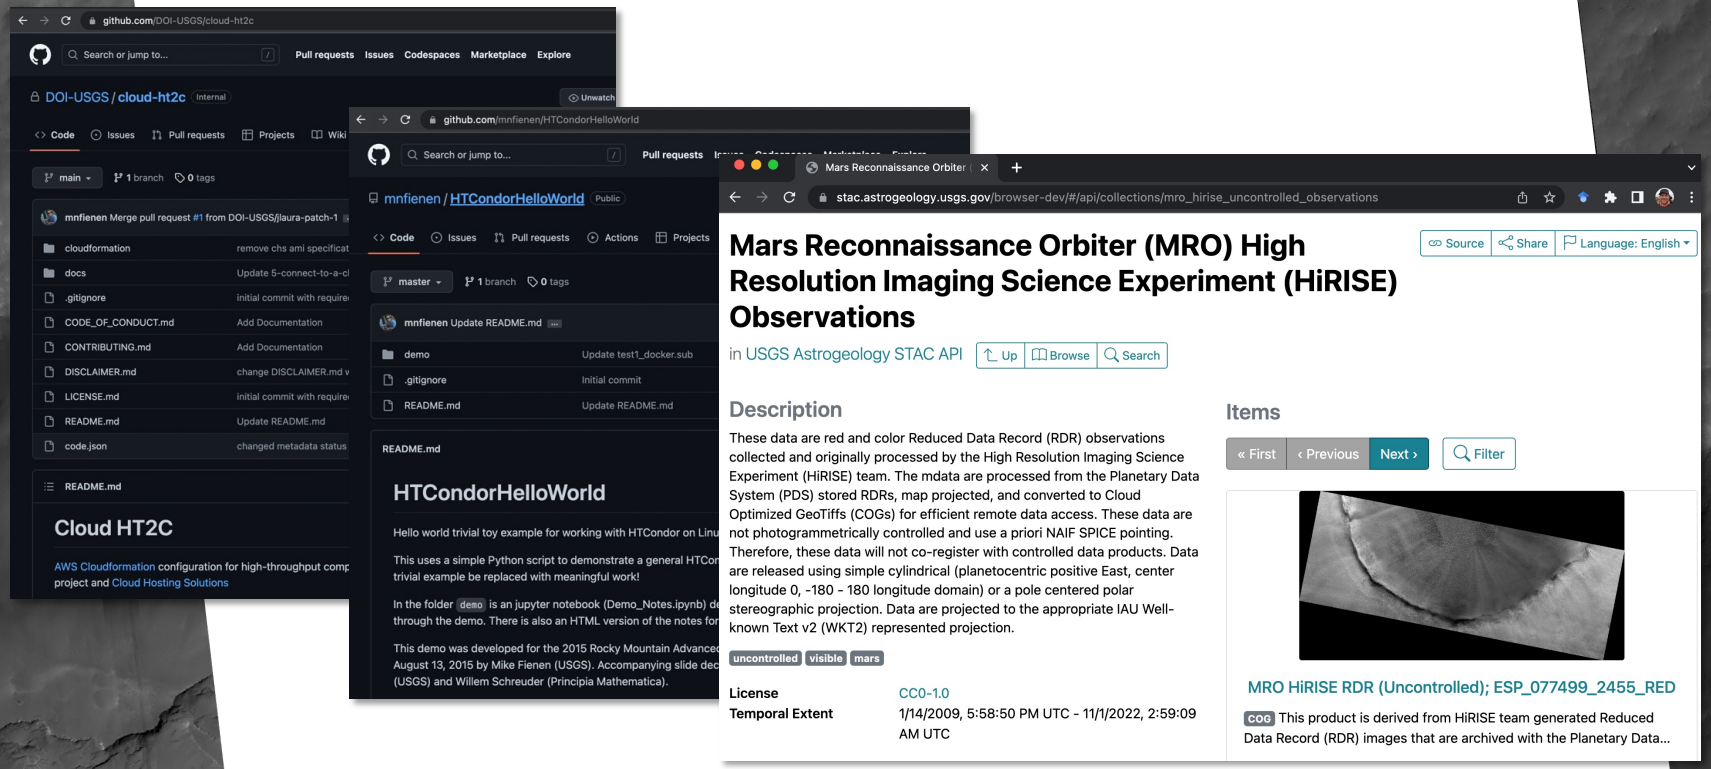

HyTEST

Cloud Hosting Solutions

USGS science for a changing world