

**AWS PROCESSING PIPELINE FOR MAVEN IUVS.** J. P. Elliott<sup>1</sup>, R. Meisner<sup>1</sup>, B. McClellan<sup>1</sup>, E. Jacobshagen<sup>1</sup>, S. Jain, N. Schneider<sup>1</sup>. <sup>1</sup>Laboratory for Atmospheric and Space Physics, University of Colorado, 3665 Discovery Dr, Boulder, CO 80303 (joshua.elliott@lasp.colorado.edu)

**Introduction:** The Imaging Ultraviolet Spectrograph (IUVS) [1] aboard the Mars Atmosphere and Volatile Evolution Mission (MAVEN) spacecraft is designed to measure the characteristics of the atmosphere and ionosphere of Mars, as well as atmospheric dayglow, nightglow and aurorae in the Far-UV and Mid-UV portions of the spectrum (collectively 110nm-340nm).

As presented by Elliott et al. 2019 [2], the increasing data volume over the continuing mission has prompted the need for a higher degree of data processing parallelization in order to meet mission deliverable requirements including our regular PDS release schedule and full-mission reprocessing. This led us to explore cloud-based options for our data processing needs. Our new AWS based processing pipeline gives us the flexibility and scalability to meet those needs.

**AWS Pipeline:** In recent months we have finished the migration of our Level 1 data product pipeline to the AWS cloud, allowing us to scale our processing infrastructure to meet mission processing demands. For our regular three-month PDS releases this means a quicker turn around (on the order of 10 hours rather than a week or more) for our science team so that they have more time to perform data analyses and generate the higher level products that they are responsible for.

There is also the periodic need to reprocess the entire mission dataset. This arises when a significant change to the L1 pipeline code or data product format occurs. Since Mars orbit insertion on 21 September 2014, nearly seven years of data have accumulated. Our new AWS pipeline has the ability to complete a full mission reprocess in just a matter of a week rather than many weeks or even months on our older in-house processing system.

Our pipeline can also be configured to perform our regular weekly staging processing, which provides our science team with quick access to newly downlinked data from the spacecraft for analysis.

In each of these three processing scenarios EC2 compute resources are spun up only when needed and scaled to meet the demands of each type of processing job.

**Benchmarks:** We present a comparison of the architecture of our old and new processing pipelines. Advantages and disadvantages of each system are presented along with a yearly cost breakdown.

Computation times and benchmarks will also be presented.

**References:**

[1] McClintock W. E. et al. (2015) *Space Science Reviews*, 195, 75-124.

[2] Elliott J. P. et al. (2019) *4<sup>th</sup> Planetary Data Workshop*, 2151, id.7029.