Unidata Community Equipment Awards Cover Sheet

Proposal Title: Cloud-ready Processing and Dissemination of GOES-16 Geostationary Lighting Mapper Gridded Imagery

Date: 16 March 2018

Principal Investigator Name: Eric Bruning
Title: Associate Professor
Institution: Texas Tech University
Telephone number: 806-834-3120
FAX number: 806- 742-1738
Street Address: Texas Tech University, Box 41053, Lubbock, TX 79409-1053 USA

Email address: eric.bruning@ttu.edu

Signature of PI:____

Name of Institution Official (usually dept. chair): Amy Cook, J.D.
Title: Managing Director, Office of Research Services
Telephone number: 806-742-3884
FAX number: 806-742-3892

Email address: ors@ttu.edu

Signature of University Official:__________________________________

Project Summary

Texas Tech University is funded by the GOES-R program to develop gridded Geostationary Lightning Mapper products for NOAA/National Weather Service, with the real-time processing tailored to the requirements of the NWS operational environment. The processing method is also already suitable for research of case studies, where real-time processing is less of a concern. The final missing piece would be to develop a real-time infrastructure suitable for real-time processing of GLM data to serve the education and research community. By bundling GLM data processing into using modern cloud infrastructure (Docker) it will be transferable to a variety of cloud computing hosts, such as Amazon Web Services, Google Cloud, Microsoft Azure, and NSF XSEDE/Jetstream. Products developed in the GLM processing environment will then be disseminated by LDM and THREDDS, leveraging the Unidata efforts to Dockerize this community infrastructure.

Project Description

GOES-16 and GOES-17 are part of NOAA's next generation GOES-R series of geostationary weather satellites. The Advanced Baseline Imager instrument on each platform provides sixteen bands of imagery at up to 30 s and 500 m resolution. Such imagery matches the fluid timescale of convection in the atmosphere while also improving multi-band techniques for phenomenological detection (Schmit et al. 2017). Products from the ABI imager are by now well known, in part through prior support by Unidata for community-focused data dissemination efforts with a web presence. For instance, there are web viewers at CIRA, CIMSS, and College of DuPage.

The Geostationary Lightning Mapper (GLM) is a new instrument on the GOES-R series, and provides continuous, full-disk detection of the optical pulses produced by lightning at 500 frames per second, both day and night with better than 70% detection efficiency. The GLM data are provided in NetCDF format as a three-level point dataset: flashes, groups, and events (Carlomusto 2017). Events are individual pixel triggers on the GLM CCD focal plane and groups are spatially adjacent triggers in a 2 ms frame (with the location a radiance-weighted centroid). Groups within 16.5 km and 330 ms of one another are clustered into flashes, which are also reported as a radiance weighted centroid (Goodman et al. 2010).

The web presence for GLM data has largely been confined to point displays or point density imagery (such as the SSEC RealEarth GLM group density overlay). Such displays are readily made from the raw GLM point data. However, testing of GLM proxy data in the GOES-R Proving Ground (Goodman et al. 2012; Calhoun et al. 2014, 2015) demonstrated that a flash extent density product, which captures the event-level spatial footprint of each flash, is a preferred visualization method with a simple interpretation: it counts how many flashes illuminated each location. Pragmatically, though, the calculation of flash extent density is challenging, since it requires traversal of two levels of parent-child foreign key data in the dataset to associate flashes with events (Bruning 2017, 2018). Such automated traversal is not part of standard tools or meteorological dataset standards, even though the GLM Data are CF-compliant. Furthermore, initial operational trials (as reported by K. Calhoun at the 2017 GLM Science Team Meeting) revealed that even simple point density products are compromised. A variable detector pixel size over the field of view means that no one regular target grid can correct for both double-counting of events or false gaps between events that comprise a flash. Such visual artifacts compromise the reliability of the display. Further complicating matters, the GLM L2 data contain neither the pixel size nor the ID for each event, making it impossible to properly represent each flash's spatial footprint as actually observed.

So that GLM might join as a peer dataset for utility in the meteorological community, a logical first step would be to correct for the above effects and produce imagery for overlay on ABI so that the raw data can be routinely seen at community-focused web portals. For science users, it would also be helpful to disseminate, via LDM and/or THREDDS, GLM imagery grids in a CF-compliant NetCDF format that can be used for science purposes. Advanced users might also wish to develop new lighting post-processing techniques. Because the GOES-R datasets are big, it would be beneficial to conduct those analyses in a data-proximate computation environment.

Bruning (2018) reported on an open-source toolkit, *glmtools* (https://github.com/deeplycloudy/glmtools/), that automates the traversal of GLM data files and that can produce gridded imagery. It is written in Python, following standard packaging conventions, and relies heavily on *xarray* (including our patch for dealing with the odd way event lat/lon are reported as unsigned integers by the GLM ground system). This spring we consulted with the GOES-R program, NESDIS, and Lockheed Martin to allow for navigation of the GLM data to the ABI fixed grid (resulting in perfect ABI-aligned overlays, Figure 1, left and bottom) and reconstruction of the GLM CCD pixel geometry to allow for proper representation of the flash footprint on an arbitrary target grid (Figure 1, top right). The following suite of GLM derived imagery are produced are at 1 min intervals, oversampled at 2 km resolution over the full disk:

- Flash and group extent density
- Flash and group centroid density (not using event-level spatial footprint)
- Average flash and group area (area-weighted extent density)
- Total radiant energy (energy-weighted event density)

The polygon-based description of the GLM pixels prevents use of ordinary image resampling approaches. Instead, in our approach, the event, group, and flash polygons are carefully remapped and area-weighted onto the target grid of interest. The computational geometry required to do this correctly and quickly is

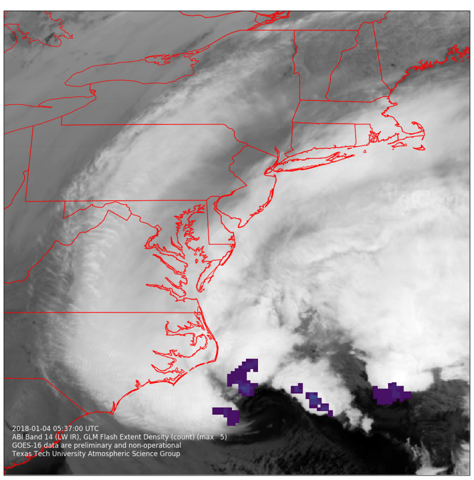

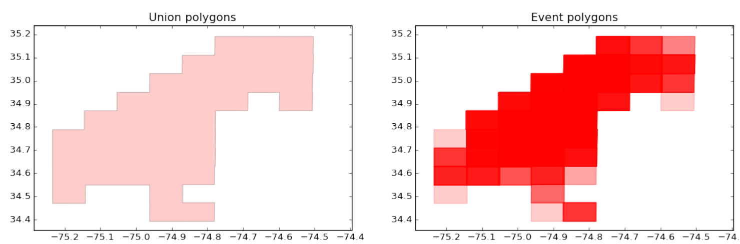

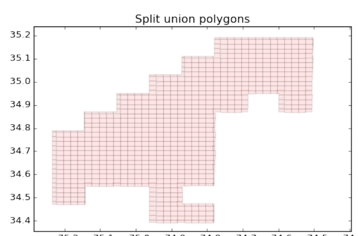

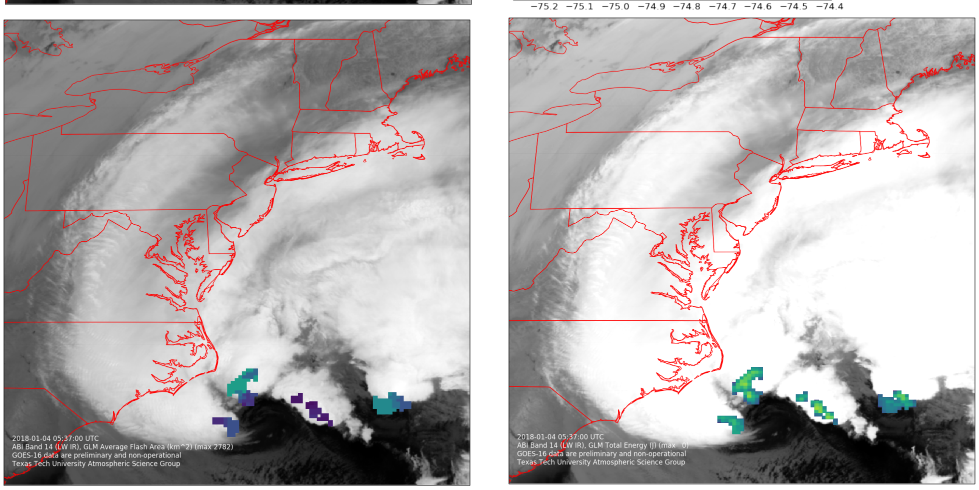

Figure 1. (Left and bottom) GLM Flash extent density, average flash area, and total radiant energy overlaid on ABI band 14 infrared imagery from a mesoscale sector on 4 January 2018. (Top right) One flash in the northwestern-most storm cell is broken into its constituent events (darker shading is more overlapping events), the union of those event polygons to make a flash polygon, and the flash polygon as sliced into its constituent polygons on the target 2 km grid. Coordinates are longitude and latitude, though the overload data were gridded on the ABI fixed grid.

not part of any standard meteorological packages, and would be a significant effort to re-implement by individual research groups.

These products will be produced, using *glmtools*, on an NWS operations prototyping system this spring. They will be the first GLM data available in AWIPS throughout late spring and early summer (S. Rudlosky, NOAA/NESDIS, personal communication). However, these products are for internal NWS use and are not expected to leave the NOAA firewall. Therefore, there is a gap in the dissemination of research best practice for basic GLM data processing to the Unidata user community.

Furthermore, within TTU the routine production of GLM imagery assists in education and research. We have operated real-time processing of the West Texas Lightning Mapping Array for six years, and it allows us to understand and mentally synthesize lightning behavior alongside other real-time satellite and radar data feeds. Such display makes it possible to identify research-relevant case studies as they occur. Furthermore, a calendar-based archive of imagery makes it easy to quickly browse past cases as necessary and bring new graduate students up to speed. This trail of imagery reduces impedance to science, and allows students and researchers to focus their efforts on analyses that lead to research results, instead of basic data visualization and case selection. We would anticipate using the capability regularly as part of our funded work under GOES-16 and GOES-17 GLM validation, as part of VORTEX-SE, and in upcoming NSF field campaigns that have been proposed. In fact, we have already promised to generate and provide NetCDF grids to support the NSF/DOE RELAMPAGO campaign in Argentina (S. Nesbitt, Twitter communication, 21 Feb. 2018).

The availability of routine GLM imagery makes possible another application. In the PI's experience, it is common to receive reports on Twitter of surprise lightning flashes (for instance, during winter storms) or even meteors (which GLM has been proven to detect). In more intense summer storms, one might notice a particularly long, rumbling lightning flash produced by large thunderstorm complex (Lang et al. 2017). Such reports often lack temporal precision (a 5-30 min window is not uncommon), but gridded imagery can make it possible, with low data transit volume, to spot the likeliest 1 min interval for the event of interest, at which point only a few raw GLM files must be investigated to find the flash of interest. A prototype notebook has been developed (glmtools/examples/basic_read_plot.ipynb) for interactively browsing through GLM data from a time period of interest, including retrieval of GLM data from a THREDDS server using Unidata's Siphon Python package.

Therefore, the purpose of this proposal is to:

- acquire hardware to enable routine GLM processing on TTU systems for creation and dissemination of GLM gridded imagery to the community,
- develop a basic website at TTU for display of these images, and
- use the above experience to develop cloud-ready packages for data-proximate replication of our processing and visualization toolchain, and transition TTU processing to the cloud.

Specifically, we propose to:

- Pull GLM data over LDM to a new storage pool attached to an existing processing server at TTU
- Create a Dockerized GLM processing environment to produce GLM NetCDF grids
    - Use Unidata's already-Dockerized LDM and THREDDS instances where appropriate
- Make GLM grids available to the community over LDM from an existing server
- Replicate the above processing with cloud computing resources:
    - On the XSEDE Jetstream environment, where Unidata has a presence and experience
        - Jupyter notebook environment for the "find a flash" application
    - On Amazon / AWS, where there is already an S3 bucket of GLM data
    - Possibly on other cloud providers (Google Cloud, Microsoft Azure)
- Develop and expand glmtools documentation, including how to run glmtools in the cloud

Resource considerations

The shared computing resources in the Texas Tech Atmospheric Science Group are as follows. This discussion excludes resources in the TTU High Performance Computing Center (which are part of University-wide shared infrastructure) and those resources managed for the operational needs of the West Texas Mesonet and West Texas Lightning Mapping Array (WTLMA). Most analyses are carried out on individual laptops or desktop workstations, and because we do not have an undergraduate program those analyses have a heavy emphasis on research and graduate education. There is one common computer lab (four older Linux machines, primarily used for the occasional class project requiring the solo3 radar data editing software). Faculty and student teaching and research webpages are hosted on www.atmo.ttu.edu, which is an old Dell desktop computer running an outdated version of Linux on an in-office network connection. That server also hosts a site for browsing the twice-daily output of the TTU Ensemble Prediction System (http://www.atmo.ttu.edu/bancell/real_time_ENS/ttuenshome.php), as a part of Big Weather Web and NOAA CSTAR projects. The only other server-style, shared computing resource in TTU ATMO is a 10-core server (graupel.ttu.edu) with 10 TB of storage. It is linked to another dedicated WTLMA processing server, which runs a project-specific LDM instance. Therefore, this proposal will modernize and augment close to 100% of the shared resources for larger, non-HPC analyses at TTU. According to the Unidata website, TTU has never received an equipment award.

The initial implementation of GLM processing will be on graupel.ttu.edu, where it has already been tested but storage is limited. We will receive the GLM data via IDD over LDM. The parallelized code can process a full disk GLM grid just faster than real time on 4 CPU cores and with < 8 GB memory. The GLM grids are quite sparse, so a full disk compresses to a few 10s of KB per minute. Each day of raw GLM files from GOES-16 is approximately 2 GB; with the addition of GOES-17, this is about 1.5 TB/yr. The Lubbock NWS office, with whom we work closely, has expressed interest in having the ability to access archived data, including from ABI. At least one MS student in the PI's group, with others to come, would benefit from the ready access to local data, with further data processing carried out on local resources. While saving all ABI data from both satellites is prohibitive, maintenance of a rolling archive gives us the opportunity to select and retain time periods of interest. Therefore, this proposal requests support for an initial 10 TB archive at TTU, which doubles our current storage capacity.

The WTLMA server's LDM will be used disseminate the GLM grids. We also will produce routine GLM imagery and disseminate it on www.atmo.ttu.edu, much as we have for the WLTMA data. We propose to move this server to a more sustainable VM environment operated by TTU. The move to a VM will also stabilize the TTU Ensemble Prediction System viewer, as well as other, smaller projects such as the real-time StickNet display built for VORTEX-SE.

For resources on XSEDE/Jetstream, our GLM processing is a good match for an m1.medium VM (6 vCPUs, 16 GB RAM), which consumes 6 SU per hour. We anticipate applying for a Jetstream startup allocation, which is typically 50,000 SU per XSEDE documentation. That allocation would cover nearly a full year of GLM processing. We may also divide this allocation to support LDM and THREDDS instances later in the project. On the current XSEDE prototype individual 6 core / 16 GB RAM instances are used to run both LDM and THREDDS, though our needs will likely be less given the small GLM grids to be disseminated. We have already demonstrated the ability to run the GLM processing on XSEDE in the Unidata Science Gateway (J. Chastang, personal communication), and we will gladly collaborate further to ensure our efforts integrate with Unidata's funded work in this area.

Eventually, our TTU processing server will need to be used for other local processing needs, and so we propose to transition our routine processing to the cloud. Our initial target for this processing will be on Amazon Web Services (AWS), making use of lessons learned from the local and XSEDE environments, while moving the processing environment to be proximate to the AWS S3 bucket hosting a real-time feed of the GOES-16 dataset. Other cloud providers (Google Earth, Microsoft Azure) also host GOES data, and if time permits we will also attempt to acquire free compute hours to demonstrate GLM processing on those hosts.

The following personnel will accomplish the work. Research associate John Geesling is already supported by GOES-R GLM Validation activities, and by some College-level funds for general research computing support. Infrastructure purchased under this award was designed in consultation with Mr. Geesling, and will be brought under his administrative supervision in support of our GLM activities. Resources made available with this award will be the deployment targets favored by PI Bruning as he implements the GLM work to which he is already committed, and will ensure that the work being done for NOAA impacts and is accessible to the wider research community. One or more MS students (for instance Cameron Nixon, already funded as a Teaching Assistant and working on analyses of tornadic storms and storm environments with GLM) will also be instructed in the use of the new infrastructure.

In summary, our equipment request meets two of the four areas of special consideration for 2018 awards: data proximate analysis of large datasets, and student use and access to GOES-R data. Furthermore, the use of xarray in our data processing supports the field testing of crucial technologies that are being used to facilitate the application of machine learning techniques and data analytics in atmospheric science (a third area of emphasis). Finally, the TTU ATMO web server also disseminates ensemble weather prediction imagery, the fourth area of emphasis. More broadly, this proposal contributes to the Unidata community capability and broadens its scope to include state-of-the art GLM processing, which will be disseminated by participation in the IDD. Students benefit from using the same processing technique used by the NWS, furthering their education and research. The Texas Tech Atmospheric Science Group will benefit from worked experience with modern VM and cloud-based data processing workflows that integrate with atmospheric science data infrastructure.

Budget

Item	**Description**	**Cost**
Dell storage server	1 CPU, 6 cores; 4 GB memory, 10 TB + 24 spare bays storage; RAID.	$6,988.00
TOSM Virtual Machine	2 vCPUs, 4 GB RAM, 500 GB storage	$775.50
XSEDE Startup Allocation	50,000 SU	$0.00
AWS EC2 Compute	c5.2xlarge, 3 year convertible term, reserved instance, 36 mo	$3,259.00
AWS EBS Storage	Throughput optimized HDD, 500 GB, 36 mo	$810.00
University Overhead (F&A)	49% on non-capital purchases	$2374.00
Total		**$14,207.00**

The budget for this proposal includes components for initial processing of data at TTU. An existing 10-core machine is already being used for lightning data processing at TTU. To support the volume of GLM data to be processed we propose a companion storage server, to provide 10 TB of initial storage, with an additional 24 drive bays available for later expansion. We will pair this purchase with funds from our other current awards to enable additional storage purchase (at least $4K is available) as stated in their data management plans. Cost is that quoted under the TTU Dell contract rates. We also propose resources in the TTU data center to provide a public web presence for the GLM imagery. Using published rates at the TTU Technology and Operations System Management datacenter, we request funds for a virtual machine with 2 CPUs, 4 GB RAM, and 500 GB storage.

Based on our expected load for the core GLM processing task, and because of our lack of familiarity with actual resource utilization on Amazon, we request an c5.2xlarge (8 vCPUs, 16 GB memory), compute-optimized instance. We have budgeted a 3-year, convertible fixed term so that the resources might be used for extended real-time processing or split to purchase other instances (for data dissemination). We also include 3 years of storage charges for 500 GB of data. Total cost is $3259+$710 according to the AWS pricing chart. While the award is for one year, we will purchase all AWS credits immediately and use them over three years to sustain GLM processing, thereby supporting the community over the long term.

No costs are expected for our work using the XSEDE infrastructure.

F&A is charged at the approved rate of 49% of modified total direct costs (MTDC). MTDC excludes equipment costs (items costing $5,000 or more); graduate student tuition and fee remission; participant support costs, and the amount of each subaward over $25,000.

Project Milestones

Date	Milestone
1 May 2018	Award Starts
	Request Jetstream XSEDE allocation
	Purchase local storage server and virtual machine
30 June 2018	Deploy local storage server
30 July 2018	Complete local Docker processing environment
	Begin disseminating data from TTU
1 August 2018	RELÁMPAGO field campaign begins
30 August 2018	Initial tests with Docker on Jetstream, AWS
30 December 2018	Stable processing and dissemination on Jetstream and AWS
January–March 2019	Performance tuning and consolidation
1 April 2019	Write final report
30 April 2019	Award Ends

References

Bruning, E. C., 2018: Processing and quality-controlled visualization of GOES-16 GLM data. Eighth Symposium on Advances in Modeling and Analysis Using Python, Austin, TX, USA, American Meteorological Society. P. 7.4.

Bruning, E. C., 2017: From the field to geostationary orbit: Mapping lightning with Python. Conference on Scientific Computing with Python (SciPy). URL https://github.com/deeplycloudy/scipy2017/tree/master

Calhoun, K. M., G. T. Stano, E. C. Bruning, S. D. Rudlosky, and D. M. Kingfield, 2014: The Hazardous Weather Testbed: Moving total lightning data from a research tool into forecast and warning operations. Intl. Conf. Atmos. Elec., Norman, OK, Paper O–04–02.

Calhoun, K. M., 2015: Forecaster use of total lightning data for short-term forecasts and warnings in the Hazardous Weather Testbed. Seventh Conference on the Meteorological Applications of Lightning Data, Paper 3.5, AMS Annual Meeting, Phoenix, AZ, USA, invited.

Carlomusto, M., 2017: GOES-R Series Product Definition and User's Guide (PUG), Volume 5: Level 2+ Products. Technical Report v1.1, NOAA/Harris Corp.

Goodman, S. J., R. J. Blakeslee, W. J. Koshak, D. Mach, J. Bailey, D. Buechler, L. Carey, C. Schultz, M. Bateman, E. McCaul, and G. Stano, 2013: The GOES-R Geostationary Lightning Mapper (GLM). Atmospheric Research, 125-126, 34–49, doi:10.1016/j.atmosres.2013.01.006.

Goodman, S. J., J. Gurka, M. DeMaria, T. J. Schmit, A. Mostek, G. Jedlovec, C. Siewert, W. Feltz, J. Gerth, R. Brummer, and et al., 2012: The GOES-R Proving Ground: Accelerating user readiness for the next-generation geostationary environmental satellite system. Bull. Amer. Meteor. Soc., 93, 1029–1040, doi:10.1175/bams-d-11-00175.1.

Goodman, S. J., D. M. Mach, W. J. Koshak, and R. J. Blakeslee, 2010: Algorithm Theoretical Basis Document (ATBD) for the GLM Lightning Cluster-Filter Algorithm. Technical Report v2.0, NOAA/NESDIS Center for Satellite Applications and Research.

Lang, T. J., S. Pedeboy, W. Rison, R. S. Cerveny, J. Montany`a, S. Chauzy, D. R. MacGorman, R. L. Holle, E. E. A´vila, Y. Zhang, and et al., 2017: WMO world record lightning extremes: Longest reported flash distance and longest reported flash duration. Bulletin of the American Meteorological Society, 98, 1153–1168, doi:10.1175/bams-d-16-0061.1.

Schmit, T. J., P. Griffith, M. M. Gunshor, J. M. Daniels, S. J. Goodman, and W. J. Lebair, 2017: A closer look at the ABI on the GOES-R series. Bulletin of the American Meteorological Society, 98, 681–698, doi:10.1175/bams-d- 15-00230.1.

PowerEdge T630 - Fully configurable

Price: $6,987.96

Ships in 3 - 5 business days

PowerEdge T630 - Fully configurable

Unit Price: $6,987.96
Order Code: PE_T630_1023

Option	Selection	SKU/Product Code	Quantity
PowerEdge T630	PowerEdge T630 Server	210-ACWJ, 329-BCZH	1
Trusted Platform Module (TPM)	No Trusted Platform Module	461-AADZ	1
Chassis Configuration	Chassis with up to 32, 2.5" Hard Drives, Rack Configuration	321-BBKD	1
Shipping	PowerEdge T630 Shipping	340-AKLY	1
Processor	Intel® Xeon® E5-2609 v3 1.9GHz,15M Cache,6.40GT/s QPI,No Turbo,No HT,6C/6T (85W) Max Mem 1600MHz	338-BFFT	1
Additional Processor	No Additional Processor	374-BBBX	1
Processor Thermal Configuration	1 CPU up to 105W	412-AADW	1
Memory DIMM Type and Speed	2400MT/s RDIMMs	370-ACPH	1
Memory Configuration Type	Performance Optimized	370-AAIP	1
Memory Capacity	4GB RDIMM, 2400MT/s, Single Rank, x8 Data Width	370-ACOG	1
RAID Configuration	RAID 6 for H730/H730P (4-32 HDDs or SSDs)	780-BBHY	1
RAID Controller	PERC H730 RAID Controller, 1GB NV Cache	405-AADT	1
Hard Drives	2TB 7.2K RPM NLSAS 12Gbps 512n 2.5in Hot-plug Hard Drive	400-AMUC	8
Additional Network Cards	On-Board Intel i350-AM2 Dual Port 1Gb LOM	542-BBBP	1
Embedded Systems Management	iDRAC8 Express, integrated Dell Remote Access Controller, Express	385-BBHN	1
Internal Optical Drive	No Internal Optical Drive for x16+x16 and x32 Chassis	429-AAOJ	1
Rack Rails	No Rack Rails, No Cable Management Arm, No Casters	770-BBCR	1
Bezel	No Bezel	350-BBBW	1
Power Management BIOS Settings	Performance BIOS Setting	384-BBBL	1
Power Cords	NEMA 5-15P to C13 Wall Plug, 125 Volt, 15 AMP, 10 Feet (3m), Power Cord, North America	450-AALV	1
Power Supply	Single, Hot-plug Power Supply (1+0), 495W	450-ADWP	1
System Documentation	No Systems Documentation, No OpenManage DVD Kit	631-AACK	1
Operating System	No Operating System	619-ABVR	1
OS Media Kits	No Media Required	421-5736	1
Service	5 Year Basic Hardware Warranty Repair, 5X10 HW-Only, 5x10 NBD On-site	976-8234, 976-8285, 996-8029	1
Deployment Services	No Installation	900-9997	1
Remote Consulting Services	Declined Remote Consulting Service	973-2426	1

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

Same day shipment subject to order size limitations, Dell standard shipping methods and payment via credit card, gift card or Dell Business Credit. Notification will be provided if there are payment delays which could impact shipping date. Electronics and accessories may ship separately.

Smart Selection. Limited quantities. Only available for orders placed by 5:59 p.m. CT Mon.–Thurs. Systems shipped the next business day after an order is placed. Subject to order approval. Software and accessories not part of the configuration will be shipped separately and may arrive after your system. Please note that Smart Selection Configuration pricing cannot be combined with other pricing offers or discounts provided or agreed to by Dell. ** Orders with Custom Factory Integration might require additional processing time.

^Dell Business Credit: Offered to business customers by WebBank, Member FDIC, who determines qualifications for and terms of credit. Taxes, shipping and other charges are extra and vary. Minimum monthly payments are the greater of $15 or 3% of the new balance shown on the monthly billing statement. Dell and the Dell logo are trademarks of Dell Inc.

**Payment solutions provided and serviced by Dell Financial Services L.L.C. or its affiliate or designee ("DFS") for qualified customers. Offers may not be available or may vary in certain countries. Where available,

offers may be changed without notice and are subject to product availability, credit approval, execution of documentation provided by and acceptable to DFS, and may be subject to minimum transaction size. Offers not available for personal, family or household use. Dell and the Dell logo are trademarks of Dell Inc. Restrictions and additional requirements may apply to transactions with governmental or public entities.

¹Subject to applicable law and regulations.

ALL ORDERS ARE SUBJECT TO APPROVAL AND ACCEPTANCE BY DELL. Pricing, availability and other terms of offer may be withdrawn or changed without notice. Dell cannot be held responsible for errors in typography or photography.

TOSM Server and Storage Services

Pricing & Rates

Last Updated: August 2016

Datacenter Infrastructure	Customer Costs
Rack space for rack-mountable servers/storage	$0
Power (redundant UPS, PDU and generator)	$0
Cooling and humidity control	$0
Network connections for physical servers and devices	$225 per gigabit connection, $2525 per 10gbE connection
Server monitoring	$0; only eligible on systems with active maintenance contracts.

Virtual Servers	Customer Costs
General Purpose VM - Base Server (includes Windows Server 2012 Standard or Oracle Linux 6/7, 2GB RAM, 1 vCPU, 1gb shared network connection, storage prices are below)	$400 per year
Database VM - Base Server (includes Windows Server 2012 Standard or Oracle Linux 6/7, 4GB RAM, 2 vCPUs, 1gb shared network connection, storage prices are below)	$650 per year
Virtual Server Add-Ons	
Additional RAM (unit is 1GB, up to 16GB total)	$25/GB per year
General purpose storage (unit is 50GB)	$0.25/GB per year ($12.50/50GB per year)
Database storage (unit is 50GB)	$1/GB per year ($50/50GB per year)
Additional vCPU (unit is 1 vCPU, up to 4 total)	$100/vCPU per year
Microsoft SQL Server 2014 Standard license (VMs only)	$200 per year

TOSM Server and Storage Services

Pricing & Rates

Last Updated: August 2016

Storage Services	Customer Costs
STORAGE01 - ONSITE WITH OFFSITE BACKUP (Also referred to as TechSHARE; Each TTU/TTUS department receives 100GB free)	**$.20/GB per year***
STORAGE02 – ONSITE WITH NO BACKUP	**$.10/GB per year***
STORAGE03 – ALTERNATE SITE REPLICA (source data must be on one of the other TOSM Storage Services)	**$.10/GB per year****
STORAGE04 – OFFSITE WITH NO BACKUP	**$.10/GB per year***

* **Storage available via UNC path; the customer manages access using eRaider accounts and groups**

** **Storage replication will be configured and maintained by TOSM personnel**

Backup Services	Customer Costs
TOSM Virtual Server Backups (must reside on TOSM's VMware cluster)	**2 TB of data is free; $.15/GB per year for additional data****
Physical Server Backups for servers hosted in the data center	**1 TB of data is free; $.15/GB per year for additional data****
Remote Physical Server Backups (servers NOT hosted in the data center; must have a 1gb network connection; must be able to complete a full backup cycle within a normal nightly backup window)	**250GB of data is free; $.15/GB per year for additional data****

** **Additional data beyond the threshold amount will be billed at $.15/GB per year, based on the maximum amount of a single backup at any point during the fiscal year and will be billed in arrears at the beginning of the next fiscal year.**

Get Started for Free

Create Free Account

Amazon EBS Pricing

With Amazon EBS, you only pay for what you use. The pricing for Amazon EBS volumes is listed below.

Region:

US East (N. Virginia)

Amazon EBS General Purpose SSD (gp2) volumes

- $0.10 per GB-month of provisioned storage

Amazon EBS Provisioned IOPS SSD (io1) volumes

- $0.125 per GB-month of provisioned storage
- $0.065 per provisioned IOPS-month

Amazon EBS Throughput Optimized HDD (st1) volumes

- $0.045 per GB-month of provisioned storage

Amazon EBS Cold HDD (sc1) volumes

- $0.025 per GB-month of provisioned storage

Amazon EBS Snapshots to Amazon S3

- $0.05 per GB-month of data stored

Except as otherwise noted, our prices are exclusive of applicable taxes and duties, including VAT and applicable sales tax. For customers with a Japanese billing address, use of AWS is subject to Japanese Consumption Tax. Learn more.

Get Started with AWS for Free

Create a Free Account

AWS Free Tier includes 30GB of Storage, 2 million I/Os, and 1GB of snapshot storage with Amazon Elastic Block Store (EBS).

View AWS Free Tier Details »

Amazon EC2 Reserved Instances

Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them.

Get AW
Cr
AW includ Linux t2.m each year. the Fr EC2 Micro instances.
View AWS Free Tier Details »

Reserved Instances Purchase Options

Standard and Convertible Reserved Instances Pricing

Standard Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing, and can be purchased for a 1-year or 3-year term. Customers have the flexibility to change the Availability Zone, the instance size, and networking type of their Standard Reserved Instances.

Purchase Convertible Reserved Instances if you need additional flexibility, such as the ability to use different instance families, operating systems, or tenancies over the Reserved Instance term. Convertible Reserved Instances provide you

Documentation

Release Notes

Discussion Forum

Amazon EC2 SLA

Get Started for Free

Create Free Account

with a significant discount (up to 54%) compared to On-Demand Instances and can be purchased for a 1-year or 3-year term.

Reserved Instances Payment Options

You can choose between three payment options when you purchase a Standard or Convertible Reserved Instance. With the **All Upfront** option, you pay for the entire Reserved Instance term with one upfront payment. This option provides you with the largest discount compared to On-Demand instance pricing. With the **Partial Upfront** option, you make a low upfront payment and are then charged a discounted hourly rate for the instance for the duration of the Reserved Instance term. The **No Upfront** option does not require any upfront payment and provides a discounted hourly rate for the duration of the term.

Note: For T2 Unlimited instances, CPU Credits are charged at:

- $0.05 per vCPU-Hour for Linux, RHEL and SLES, and
- $0.096 per vCPU-Hour for Windows and Windows with SQL Web

The CPU Credit pricing is the same for all T2 instance sizes, for On-Demand and Reserved Instances, and across all T2 regions.

See T2 Unlimited documentation for details on when CPU Credits are charged.

Linux RHEL SLES Windows

Windows with SQL Standard Windows with SQL Web

Windows with SQL Enterprise

Region: US East (N. Virginia)

t2.nano

STANDARD 1-YEAR TERM			
Payment	Effective	Savings over On-	On-Demand

Upfront

c5.xlarge

STANDARD 1-YEAR TERM					
Payment Option	Upfront	Monthly*	Effective Hourly**	Savings over On-Demand	On-Demand Hourly
No Upfront	$0	$78.11	$0.107	37%	$0.17 per Hour
Partial Upfront	$447	$37.23	$0.102	40%	
All Upfront	$876	$0	$0.100	41%	

CONVERTIBLE 1-YEAR TERM					
Payment Option	Upfront	Monthly*	Effective Hourly**	Savings over On-Demand	On-Demand Hourly
No Upfront	$0	$89.79	$0.123	28%	$0.17 per Hour
Partial Upfront	$514	$43.07	$0.118	31%	
All Upfront	$1007	$0	$0.115	32%	

STANDARD 3-YEAR TERM	

Payment Option	Upfront	Monthly*	Effective Hourly**	Savings over On-Demand	On-Demand Hourly
No Upfront	$0	$51.83	$0.071	58%	$0.17 per Hour
Partial Upfront	$867	$24.09	$0.066	61%	
All Upfront	$1629	$0	$0.062	64%	

CONVERTIBLE 3-YEAR TERM					
Payment Option	Upfront	Monthly*	Effective Hourly**	Savings over On-Demand	On-Demand Hourly
No Upfront	$0	$59.86	$0.082	52%	$0.17 per Hour
Partial Upfront	$999	$27.74	$0.076	55%	
All Upfront	$1957	$0	$0.074	56%	

c5.2xlarge

STANDARD 1-YEAR TERM					
Payment Option	Upfront	Monthly*	Effective Hourly**	Savings over On-Demand	On-Demand Hourly

No Upfront	$0	$156.22	$0.214	37%	$0.34 per Hour
Partial Upfront	$894	$74.46	$0.204	40%	
All Upfront	$1751	$0	$0.200	41%	

CONVERTIBLE 1-YEAR TERM					
Payment Option	**Upfront**	**Monthly***	**Effective Hourly****	**Savings over On-Demand**	**On-Demand Hourly**
No Upfront	$0	$179.58	$0.246	28%	$0.34 per Hour
Partial Upfront	$1028	$85.41	$0.234	31%	
All Upfront	$2014	$0	$0.230	32%	

STANDARD 3-YEAR TERM					
Payment Option	**Upfront**	**Monthly***	**Effective Hourly****	**Savings over On-Demand**	**On-Demand Hourly**
No Upfront	$0	$103.66	$0.142	58%	$0.34 per Hour
Partial Upfront	$1733	$48.18	$0.132	61%	
All	$3259	$0	$0.124	64%	

Upfront					

CONVERTIBLE 3-YEAR TERM					
Payment Option	**Upfront**	**Monthly***	**Effective Hourly****	**Savings over On-Demand**	**On-Demand Hourly**
No Upfront	$0	$119.72	$0.164	52%	$0.34 per Hour
Partial Upfront	$1997	$55.48	$0.152	55%	
All Upfront	$3914	$0	$0.149	56%	

c5.4xlarge

STANDARD 1-YEAR TERM					
Payment Option	**Upfront**	**Monthly***	**Effective Hourly****	**Savings over On-Demand**	**On-Demand Hourly**
No Upfront	$0	$312.44	$0.428	37%	$0.68 per Hour
Partial Upfront	$1787	$148.92	$0.408	40%	
All Upfront	$3503	$0	$0.400	41%	

CONVERTIBLE 1-YEAR TERM	