

UNICLOUD

Docker Use at Unidata

Julien Chastang (UCAR, Boulder, CO USA), Ward Fisher,
Sean Arms

January 14, 2016

Outline

What We do at Unidata

Five-year (2013-2018) plan and the Cloud

What is Docker?

Microsoft Azure Grant

Goal of Unidata Server (Motherlode) in the Cloud

"Dockerizing" the LDM, TDS, RAMADDA

Lessons Learned, Other Efforts, Future Work

What we do at Unidata

- ▶ End-to-end geoscience data services
 - ▶ IDD (13 TB/day to 263 IDD sites)
- ▶ Tools for scientific data life-cycle
 - ▶ LDM, IDV, TDS, netCDF
 - ▶ RAMADDA, McIDAS/ADDE
- ▶ Community support help desk for all of the above
- ▶ Primarily serving the US academic community

Unidata 2018 Five-year Plan and the Cloud

- ▶ *Transition to [cloud computing for geoscience] is vital to our community's ability to pursue research and education in the 21st century.*
- ▶ Cloud shifts IT infrastructure (hardware/OS/software) from in-house departments to data centers managed by cloud providers yielding time, cost savings, and new opportunities.
- ▶ Advantageous for big data problems coupled with data proximate analysis
- ▶ For profit (Amazon, etc.) and non-profit clouds (NSF sponsored XSEDE for HPC and soon cloud computing)

What is Docker?

- ▶ Docker is an open source platform for uniformly building, deploying and running software in the cloud.
- ▶ Borrows from the notion of standard size shipping containers, but for software
- ▶ Linux only, so most applicable (but not limited) to server-side
- ▶ Barrier to entry is low and adoption is already wide-spread
- ▶ In last nine month, Unidata has made a significant foray into Docker.

Azure for Research Grant

- ▶ Unidata colleague Ward Fisher obtained "Azure for Research" Grant
- ▶ \$20,000 of cloud-computing resources
- ▶ Grant allows Unidata to experiment with cloud resources
- ▶ Ubuntu Linux VMs with out-of-box Docker support

What does it Mean to "Dockerize" Software

- ▶ "Dockerization" is the process of defining minimal OS environments that will allow your application to run
- ▶ These environment are codified in Dockerfiles
- ▶ Build and share these environments AKA images (e.g., via dockerhub)
- ▶ An instance of an image is called a "container"

Dockerizing the TDS/RAMADDA

- ▶ TDS and RAMADDA are Tomcat Java web applications
- ▶ For TDS and RAMADDA, use the canonical tomcat base image which buys us a lot
- ▶ For TDS, leveraged work of TDS Team
- ▶ For TDS, also must Dockerize TDM for creating index files and TDS notification
- ▶ For RAMADDA, simply tuned Tomcat, and defined data directory

Software and Data Configuration

- ▶ LDM: Configure what data you want via pqact configuration file
- ▶ Data: 13km RR, 1° GFS, GOES Satellite, NEXRAD Composites
- ▶ TDS: Configure catalogs, and where to find data on file system (reuse TdsConfig project).
- ▶ RAMADDA: Standard web configuration and server-side views

Other Efforts and Future Work

- ▶ IDV/cloud control, Python, ADDE, GEMPAK, AWIPSII
- ▶ Gradually scale up to more voluminous data feeds
- ▶ Make more robust and oriented towards operations
- ▶ Continue researching a viable business model (perhaps similar to HPC?)

Acknowledgments

- ▶ National Science Foundation (Grant NSF-1344155)
- ▶ Microsoft "Azure for Research" program
- ▶ Tom Yoksas for Unidata operations expertise