Docker Containers

Container NWP Tutorial

https://dtcenter.org/met/docker-nwp/tutorial/container_nwp_tutorial/index.php

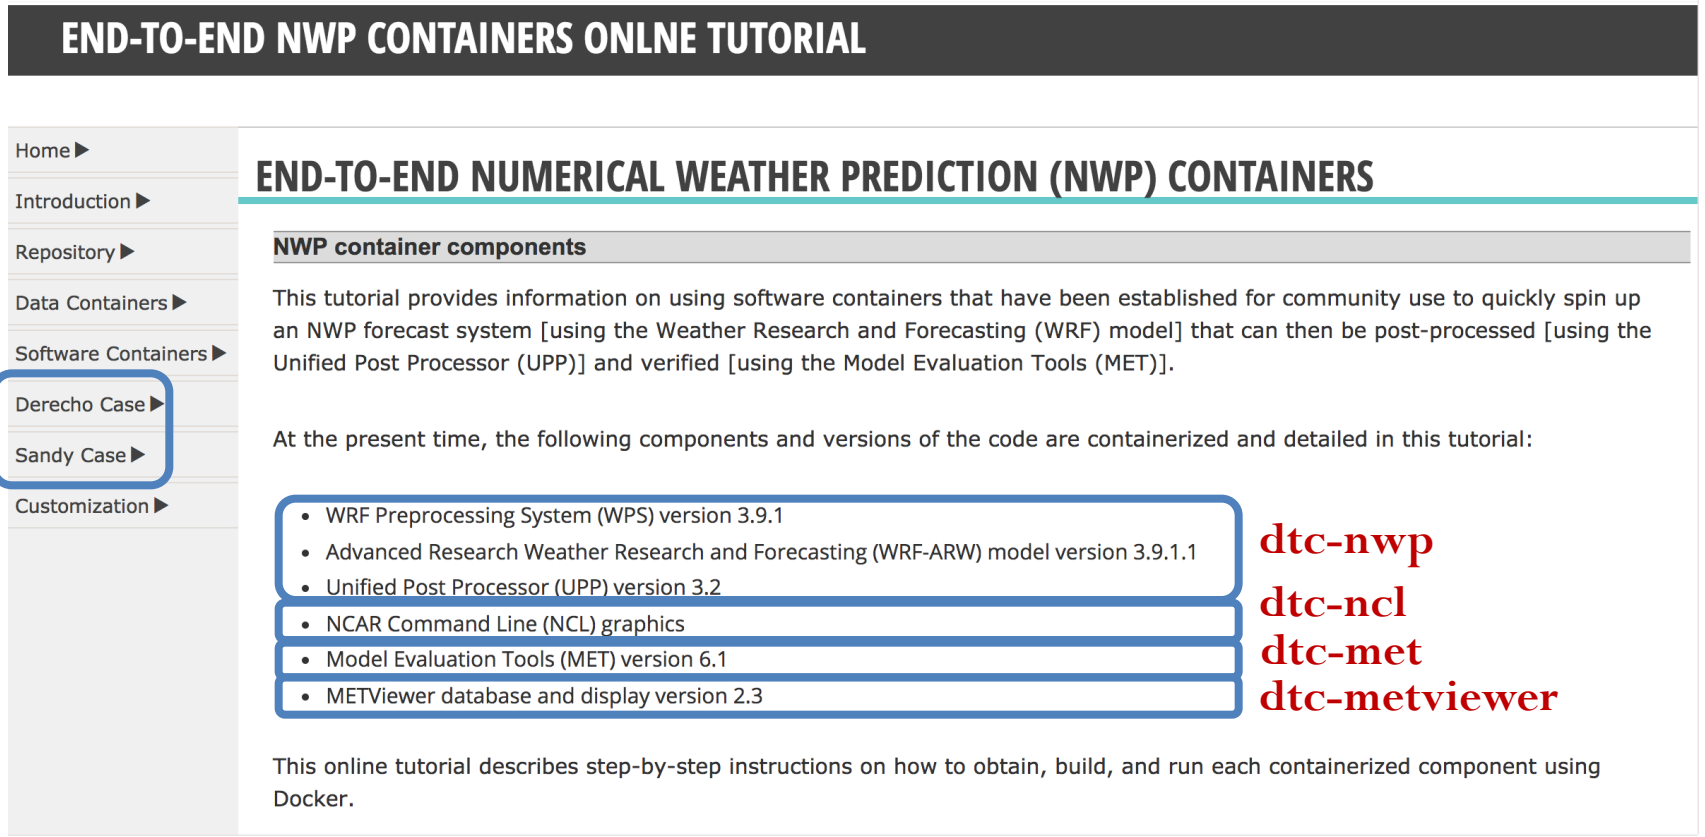

Why use containers?

- Software systems require substantial set-up to get all the necessary code, including external libraries, compiled on a specific platform → ***Containers help solve this problem!***
  - Efficient, lightweight, secure, and self-contained (including operating system, libraries, code, and executables) systems
  - Everything required is packaged into isolated components, ready for development, shipment, and deployment directly to users
  - Software should always run the same, regardless of where it is deployed
  - Eliminates possible frustrations with up-front system setup

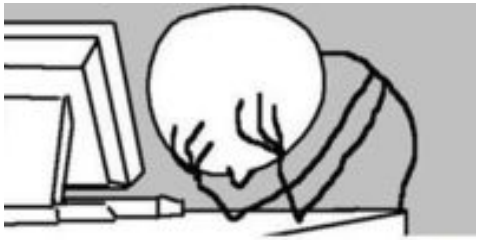

Someone trying to compile WRF

What is a Docker container?

- Docker is one of the leading software containerization platforms
  - Home page: **https://www.docker.com**
  - Documentation: **https://docs.docker.com**
- A Docker container
  - is open source
  - is an easy way to build a development environment
  - can hold applications "inside the container"
  - is portable across Linux, Mac, and Windows machines
  - is much smaller than a virtual machine
  - sets up a user-defined partition between the host machine and "container land"
  - allows "root" inside the container, but does not alter permissions on the host machine
  - ***requires root access to install Docker***
  - ***runs on a single node***

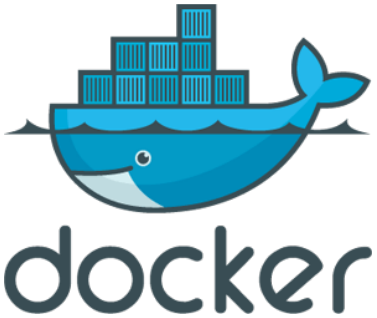

Virtual machine vs. containers

- Containers vs. virtual machines: VMs bundle a full operating system, whereas containers only contain necessary libraries and dependencies

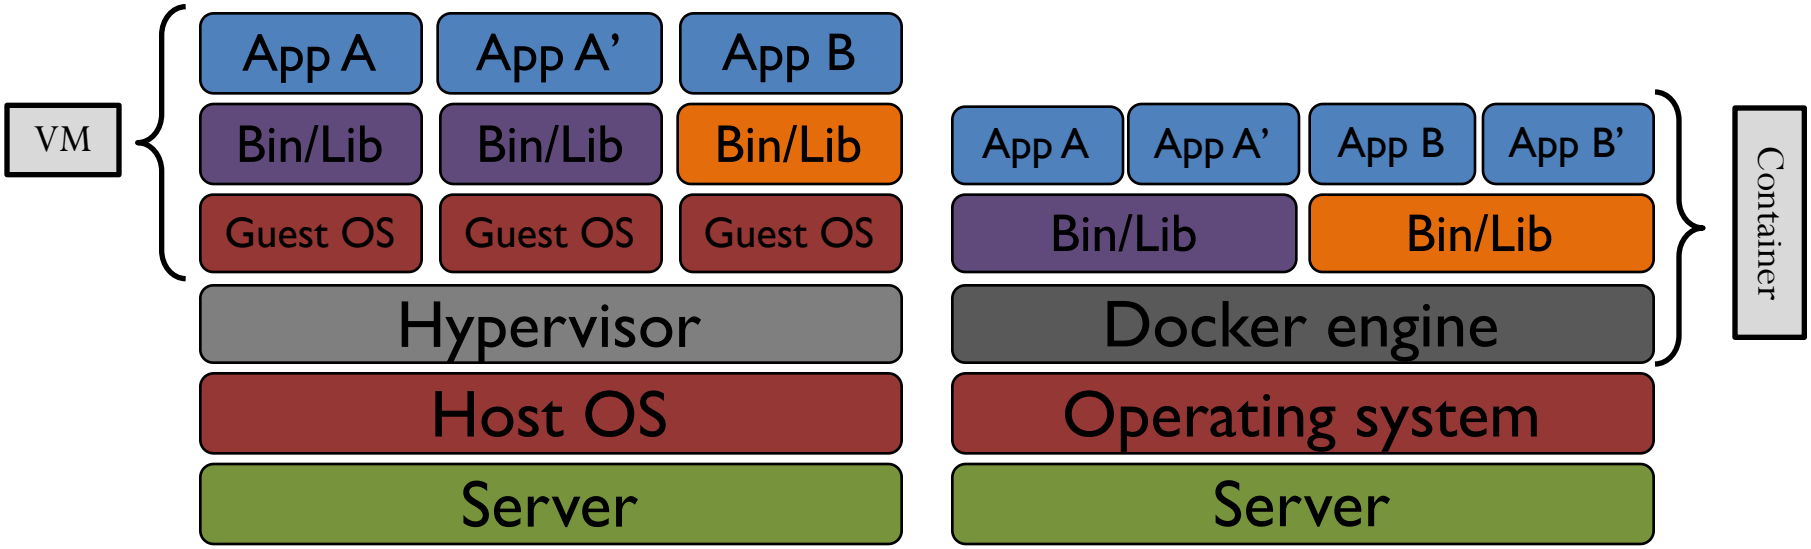

Why use containers for NWP?

Advantages:

- ✓ Reduces spin-up time to build necessary code components
- ✓ Highly portable
- ✓ Use in cloud computing
- ✓ Easily sharable with other collaborators
- ✓ Easy to replicate procedures and results

Who can benefit from containers?

- ✓ Graduate and undergrad students
- ✓ University faculty
- ✓ Researchers
- ✓ Tutorial participants

Ultimately, containers will substantially reduce the spin-up time with setting up software systems, which promotes greater efficiency in producing model and statistical output!

Images and Containers

- Images can be…
  - Built from scratch with a *Dockerfile (slower, but offers customization!)*
    - ***docker build –t dtc-met .***
  - Saved to a tar file, which can then be loaded for faster deployment
    - ***docker load –i dtc-met.tar.gz***

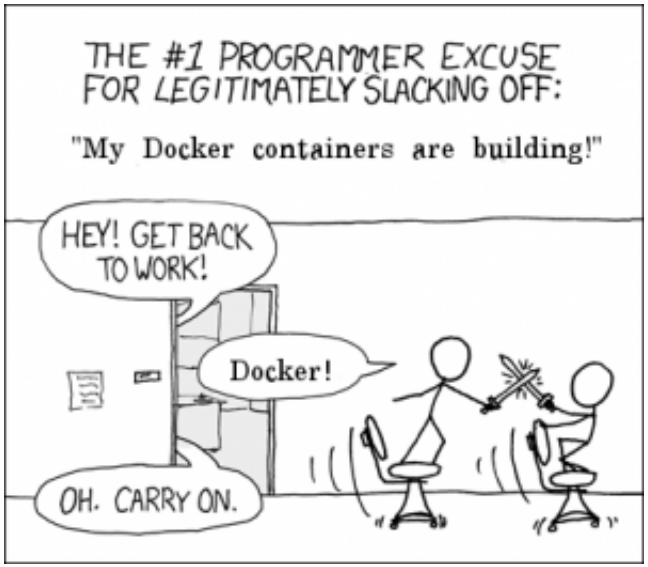

- Containers are…
  - An instance of an image
  - The result of docker run
    - ***docker run -it dtc-met /bin/bash***
  - The result of docker create
    - ***docker create --name wps_geog dtc-nwp-wps_geog***

Scripts and Configuration

- Shell scripts required to run each NWP component
- Necessary namelist and configuration files
  - Vtable.GFS
  - namelist.wps and namelist.input
  - MET configuration files
- Case-specific data
  - GFS files for ICs/LBCs
  - Observation data for gridded (Stage II) and point (NDAS prepbufr) verification
  - Sample METViewer plot xml

End-to-End NWP Workflow

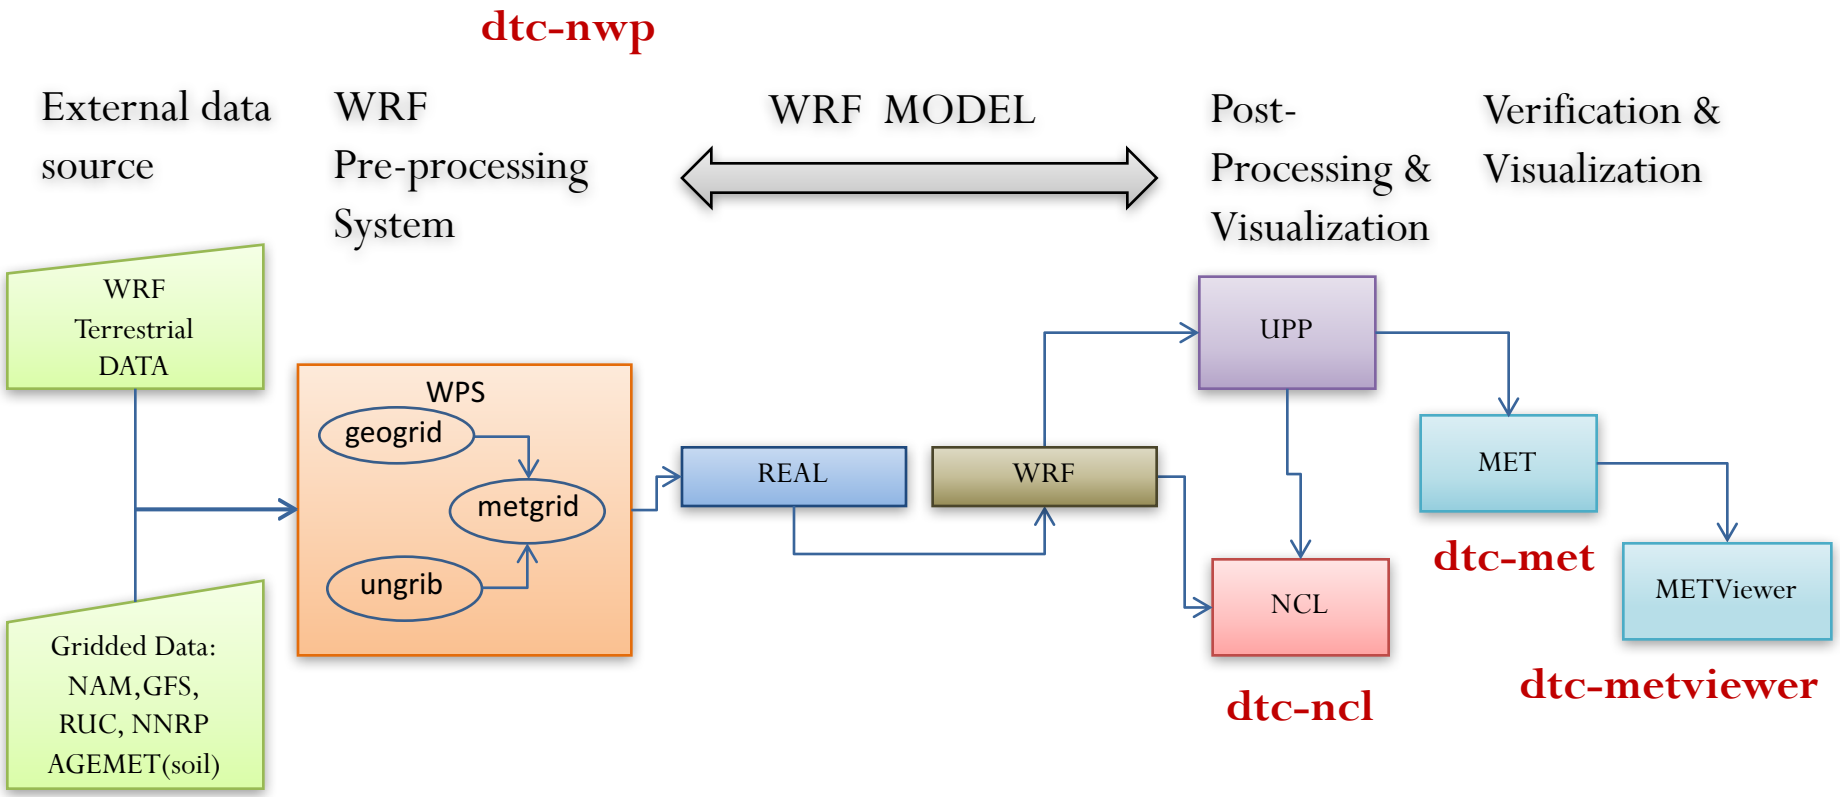

What is in the DTC containers? (1/2)

- DTC containers package everything that is needed to build and run the model and produce verification, including code and data
  - Uses gfortran; can be run serially or with distributed memory
- Two cases with full datasets are provided in current inventory
  - Hurricane Sandy (Initialized on 27 Oct. 2012)
    - 40-km domain centered over East Coast (6-h forecast)
  - Derecho event over the Eastern CONUS (Initialized on 29 June 2012)
    - 12-km parent domain with 3-km nest over southern Ohio (24-h forecast)

Total precip for Hurricane Sandy model forecast

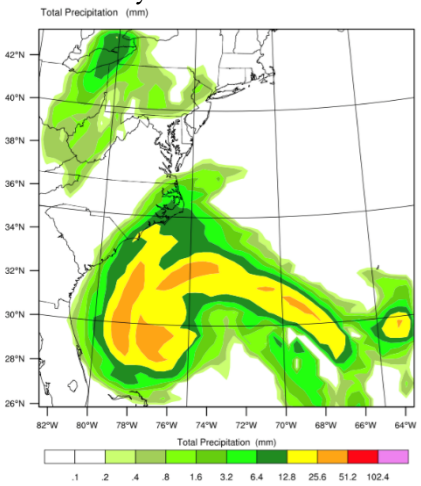

Derecho case domain configuration

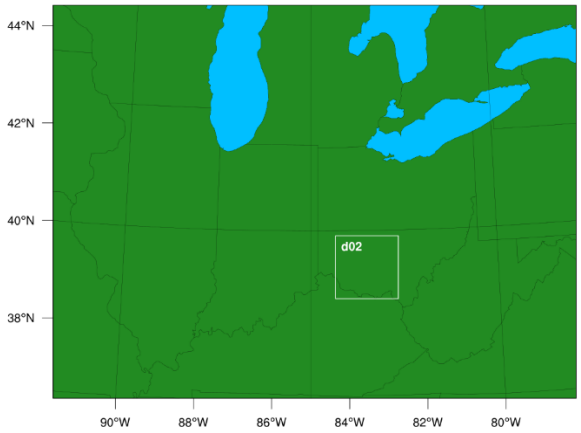

What is in the DTC containers? (2/2)

- README files with explicit instructions for building and running WPS, WRF, UPP, NCL, MET, and METViewer
- Necessary namelist and configuration files
  - Vtable.GFS
  - namelist.wps and namelist.input
  - MET configuration files
- Case-specific data
  - GFS files for ICs/LBCs
  - Observation data for gridded (Stage II) and point (NDAS prepbufr) verification
  - Sample METViewer XML plots

MET Container

- Dockerfile defines all steps necessary to compile MET from OS base image ~ ***docker build***

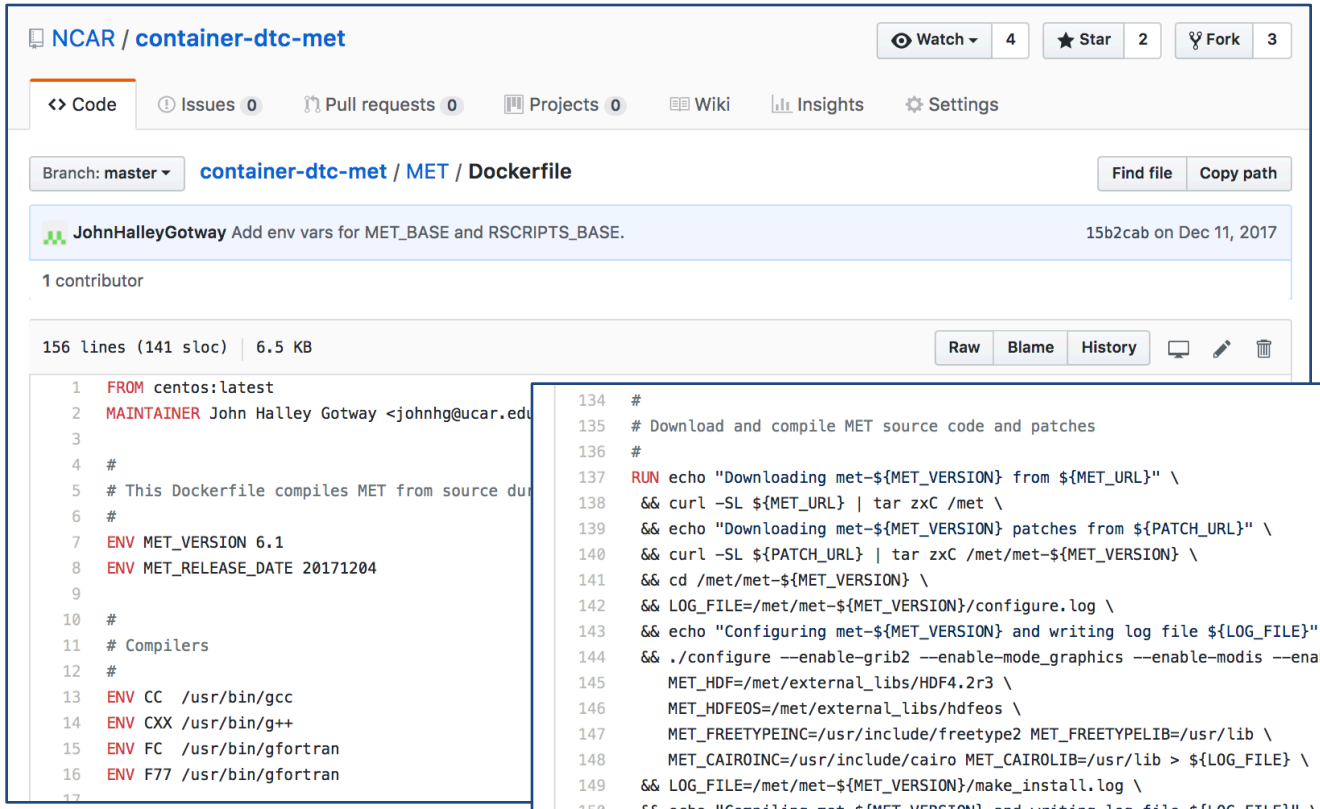

NCAR / container-dtc-met

Watch 4 | Star 2 | Fork 3

Code | Issues 0 | Pull requests 0 | Projects 0 | Wiki | Insights | Settings

Branch: master ▾ container-dtc-met / MET / Dockerfile — Find file | Copy path

JohnHalleyGotway Add env vars for MET_BASE and RSCRIPTS_BASE. 15b2cab on Dec 11, 2017

1 contributor

156 lines (141 sloc) | 6.5 KB — Raw | Blame | History

```
1  FROM centos:latest
2  MAINTAINER John Halley Gotway <johnhg@ucar.edu
3
4  #
5  # This Dockerfile compiles MET from source du
6  #
7  ENV MET_VERSION 6.1
8  ENV MET_RELEASE_DATE 20171204
9
10 #
11 # Compilers
12 #
13 ENV CC  /usr/bin/gcc
14 ENV CXX /usr/bin/g++
15 ENV FC  /usr/bin/gfortran
16 ENV F77 /usr/bin/gfortran
```

```
134 #
135 # Download and compile MET source code and patches
136 #
137 RUN echo "Downloading met-${MET_VERSION} from ${MET_URL}" \
138  && curl -SL ${MET_URL} | tar zxC /met \
139  && echo "Downloading met-${MET_VERSION} patches from ${PATCH_URL}" \
140  && curl -SL ${PATCH_URL} | tar zxC /met/met-${MET_VERSION} \
141  && cd /met/met-${MET_VERSION} \
142  && LOG_FILE=/met/met-${MET_VERSION}/configure.log \
143  && echo "Configuring met-${MET_VERSION} and writing log file ${LOG_FILE}" \
144  && ./configure --enable-grib2 --enable-mode_graphics --enable-modis --enable-lidar2nc \
145     MET_HDF=/met/external_libs/HDF4.2r3 \
146     MET_HDFEOS=/met/external_libs/hdfeos \
147     MET_FREETYPEINC=/usr/include/freetype2 MET_FREETYPELIB=/usr/lib \
148     MET_CAIROINC=/usr/include/cairo MET_CAIROLIB=/usr/lib > ${LOG_FILE} \
149  && LOG_FILE=/met/met-${MET_VERSION}/make_install.log \
150  && echo "Compiling met-${MET_VERSION} and writing log file ${LOG_FILE}" \
151  && make install > ${LOG_FILE} \
152  && LOG_FILE=/met/met-${MET_VERSION}/make_test.log \
153  && echo "Testing met-${MET_VERSION} and writing log file ${LOG_FILE}" \
154  && make test > ${LOG_FILE} 2>&1
```

METViewer Container

- METViewer only available publicly as a container.
- Dockerfile defines all steps necessary to compile METViewer from OS base image ~ ***docker build***

NCAR / container-dtc-metviewer

Watch 4 Star 0

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights Settings

Branch: master container-dtc-metviewer / METViewer / Dockerfile

TatianaBurek METViewer v2.3 e46c9f1

1 contributor

70 lines (52 sloc) 1.61 KB Raw Blame History

```
1  FROM centos:latest
2
3  MAINTAINER Tatiana Burek <tatiana@ucar.edu>
51 #
52 # Download and Deploy METViewer load/batch
53 #
54
55 RUN  curl -L ${METVIEWER_URL} | tar xzC /tmp \
56   && mv  /tmp/METViewer/dist/*.war $CATALINA_HOME/webapps/ \
57   && mv /tmp/METViewer /METViewer
58
59 #
60 # Set working directory
61 #
62 WORKDIR /METViewer
63
64
65 #
66 # Start Tomcat
67 #
68 ENTRYPOINT ${CATALINA_HOME}/bin/startup.sh && /bin/bash
69 CMD ["true"]
```

NCAR / container-dtc-metviewer

Watch 4 Star 0 Fork 0

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights Settings

Branch: master container-dtc-metviewer / docker-compose.yml Find file Copy path

TatianaBurek use downloaded mysql image fe805e5 on Nov 10, 2017

2 contributors

29 lines (26 sloc) 756 Bytes Raw Blame History

```
1  db:
2      image: mysql:5.7
3      container_name: mysql_mv
4      ports:
5          - "6603:3306"
6      volumes:
7          - ${MYSQL_DIR}:/var/lib/mysql
8
9      environment:
10         MYSQL_ROOT_PASSWORD: mvuser
11     tty: true
12 #
13 metviewer:
14     image: metviewer
15     container_name: metviewer_1
16     ports:
17     - "8080:8080"
18     volumes:
19       - ${METVIEWER_DATA}:/data
20       - ${METVIEWER_DIR}/metviewer_output/xml:/opt/tomcat/webapps/metviewer_output/xml
21       - ${METVIEWER_DIR}/metviewer_output/plots:/opt/tomcat/webapps/metviewer_output/plots
22       - ${METVIEWER_DIR}/metviewer_output/data:/opt/tomcat/webapps/metviewer_output/data
23       - ${METVIEWER_DIR}/metviewer_output/scripts:/opt/tomcat/webapps/metviewer_output/scripts
24     links:
25       - db:mysql_mv
26     tty: true
27
28 #
```

- Launch METViewer using METViewer and MySQL images ~ ***docker-compose***
- http://localhost:8080/metviewer/metviewer1.jsp

Helpful Docker commands

- docker build -t my-name . : builds image
- docker images : see what images are built
- docker rmi : remove image
- docker save my-name > my-name.tar.gz : save an image
- docker load < my-name.tar.gz : load a saved image

- docker run --rm -it --volumes-from -v --name : run a command in a new container
  - --rm: Automatically remove the container when it exits
  - -it: create an interactive bash shell in the container
  - --volumes-from: Mount volumes from the specified container(s)
  - -v: Bind mount a volume
  - --name: Assign a name to the container
- docker ps -a : see what containers are running; obtain container ID
- docker rm : remove container using ID

What does this look like for the end-to-end DTC container?

Contains necessary build and run commands to run end-to-end NWP workflow

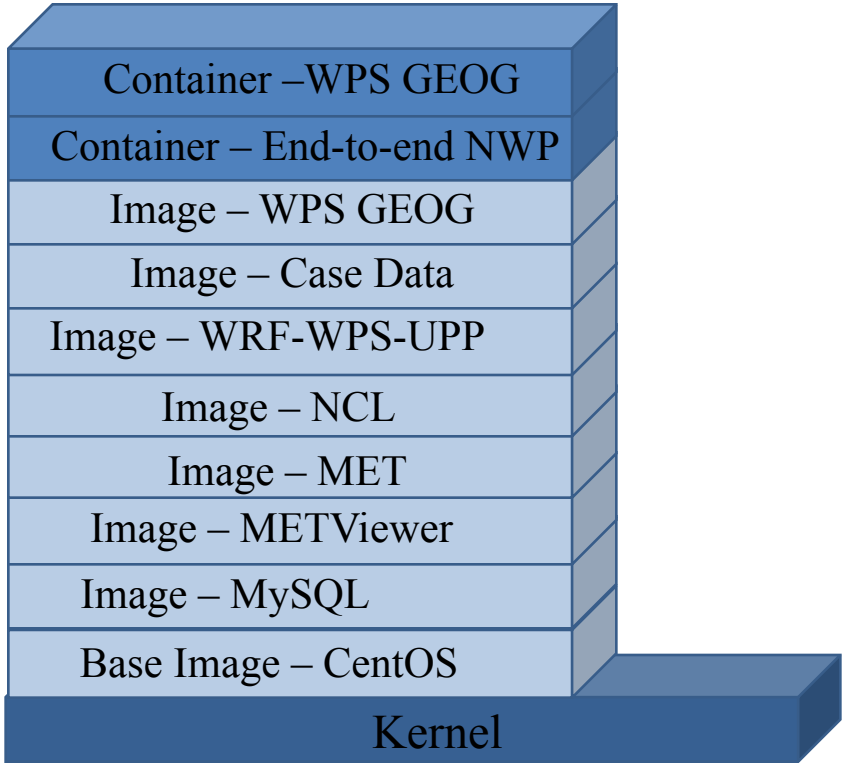


```
[nitro:~/container-dtc-nwp/components] jwolff% docker images
REPOSITORY          TAG                 IMAGE ID            CREATED             SIZE
[nitro:~] jwolff% docker ps -a
CONTAINER ID        IMAGE               COMMAND             CREATED             STATUS              PORTS               NAMES
4b96a3a0d0b7        dtc-nwp-derecho     "true"              2 minutes ago       Created                                 derecho
856064629ea5        dtc-nwp-wps_geog    "true"              3 minutes ago       Created                                 wps_geog
dtc-nwp             latest              fe1145f1002e        25 minutes ago      2.02GB
dtc-nwp-derecho     latest              8f7ff6e1e5ca        About an hour ago   773MB
dtc-nwp-wps_geog    latest              869b67961d1f        About an hour ago   16.5GB
debian              jessie              25fc9eb3417f        4 days ago          123MB
centos              latest              d123f4e55e12        4 days ago          197MB
```

Derecho Case

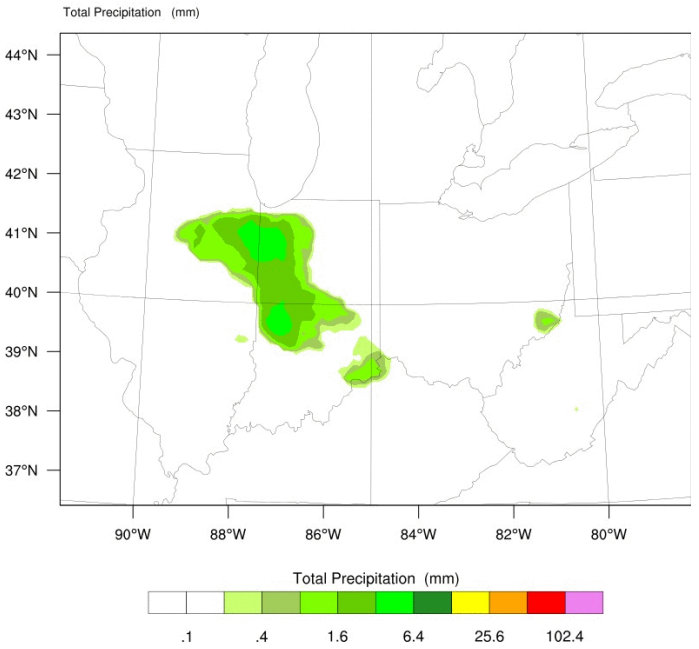

WPS Domain Configuration

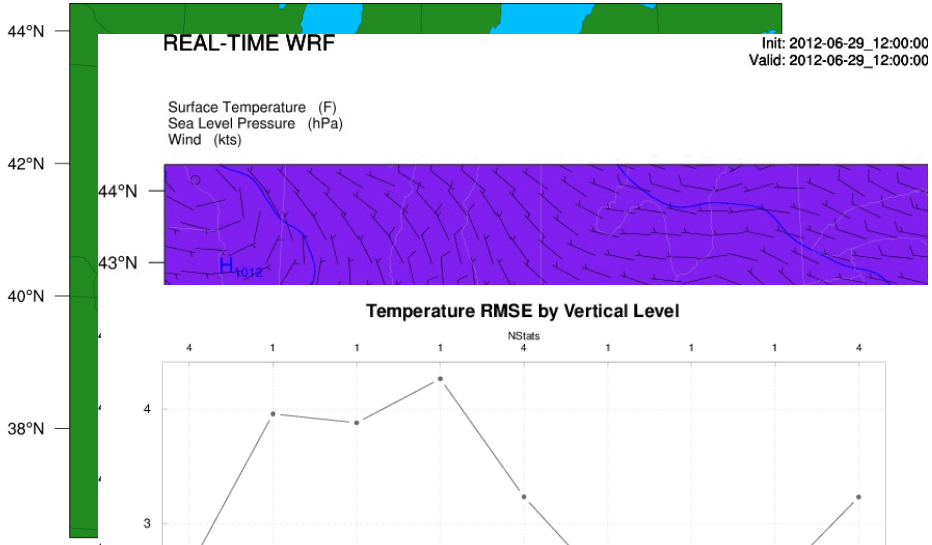

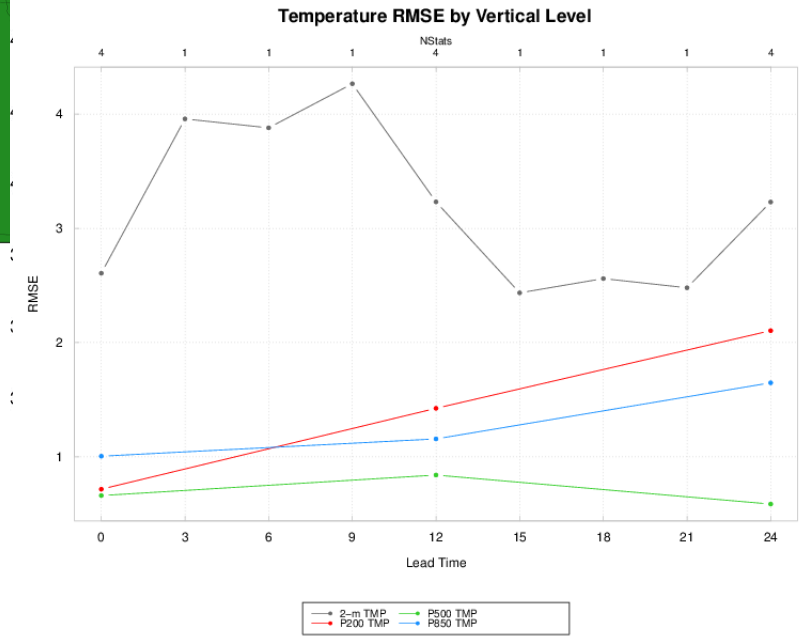

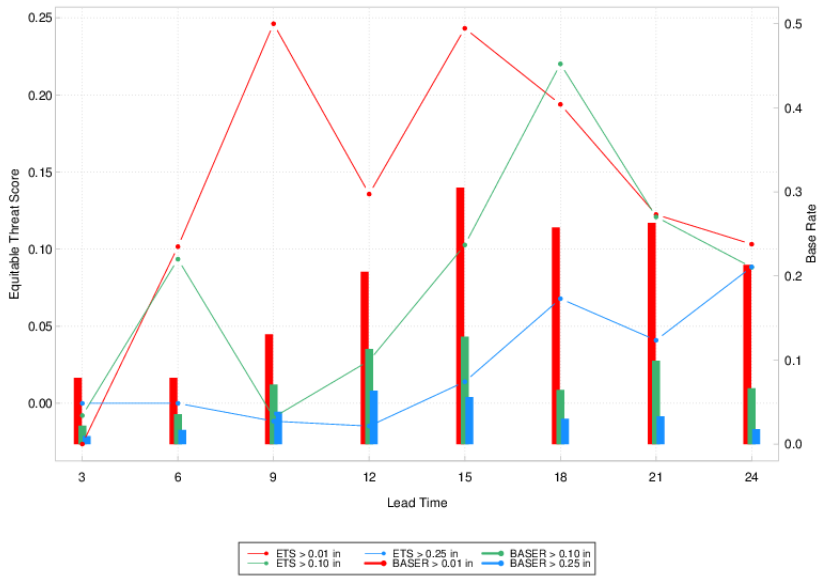

Sandy Case

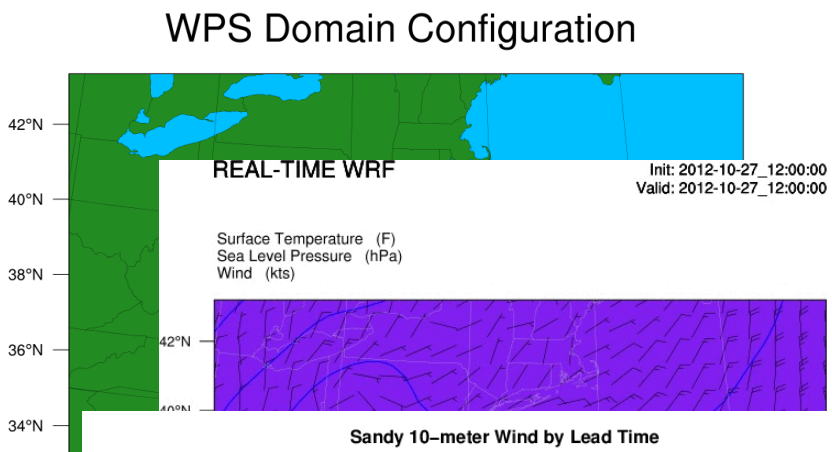
WPS Domain Configuration
REAL-TIME WRF
Init: 2012-10-27_12:00:00
Valid: 2012-10-27_12:00:00
Surface Temperature (F)
Sea Level Pressure (hPa)
Wind (kts)
42°N
40°N
38°N
36°N
34°N
32°N
30°N
28°N
26°N

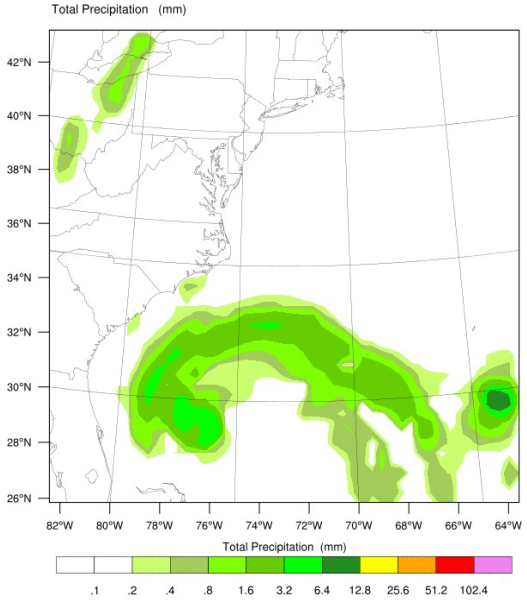
REAL-TIME WRF
Init: 2012-10-27_12:00:00
Valid: 2012-10-27_13:00:00
Total Precipitation (mm)
42°N
40°N
38°N
36°N
34°N
32°N
30°N
28°N
26°N
82°W
80°W
78°W
76°W
74°W
72°W
70°W
68°W
66°W
64°W
Total Precipitation (mm)
.1
.2
.4
.8
1.6
3.2
6.4
12.8
25.6
51.2
102.4

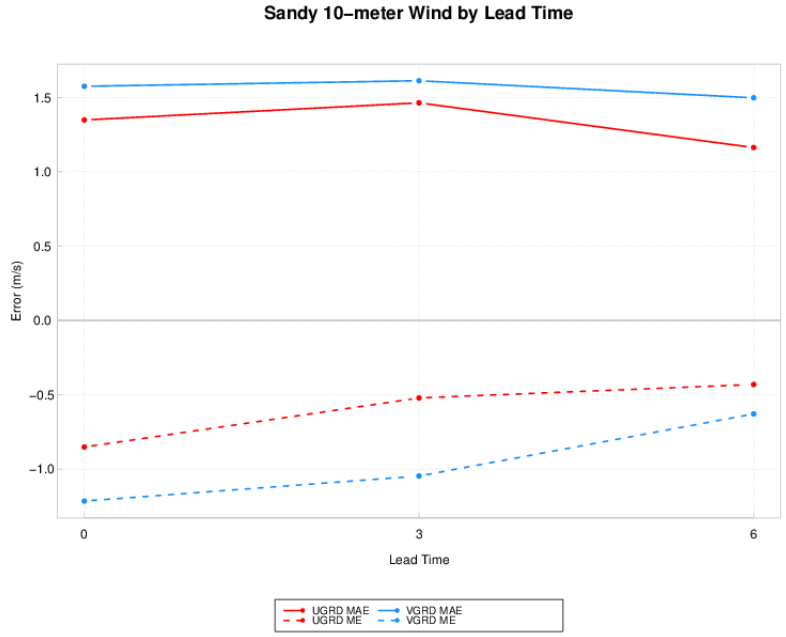
Sandy 10–meter Wind by Lead Time
Error (m/s)
1.5
1.0
0.5
0.0
−0.5
−1.0
0
3
6
Lead Time
UGRD MAE
UGRD ME
VGRD MAE
VGRD ME

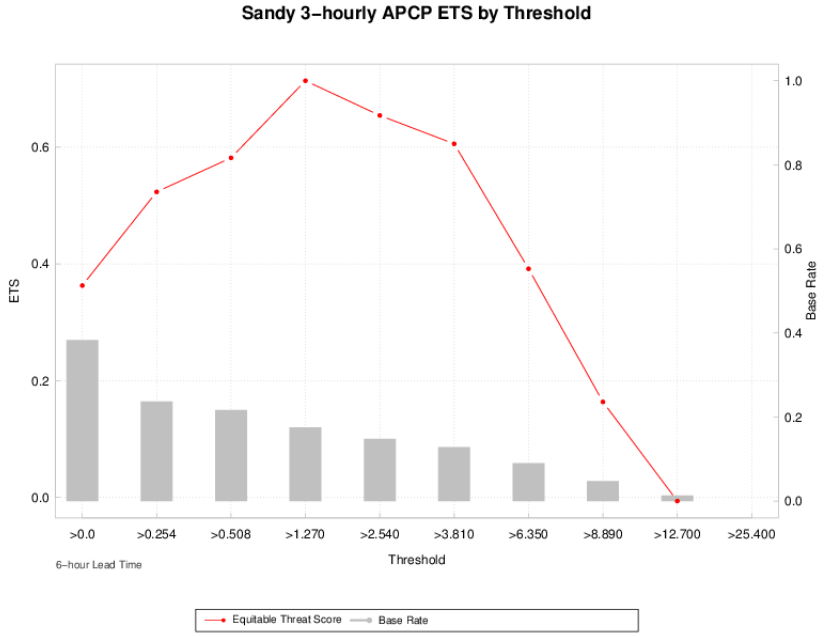
Sandy 3–hourly APCP ETS by Threshold
ETS
Base Rate
0.6
0.4
0.2
0.0
1.0
0.8
0.6
0.4
0.2
0.0
>0.0
>0.254
>0.508
>1.270
>2.540
>3.810
>6.350
>8.890
>12.700
>25.400
Threshold
6–hour Lead Time
Equitable Threat Score
Base Rate

DTC Container Links

- **WRF, WPS, NCL, and UPP**
  - https://github.com/NCAR/container-dtc-nwp/
- **MET**
  - https://github.com/NCAR/container-dtc-met/
  - https://dtcenter.org/met/users/downloads/docker_container.php
    (Instructions for installing and running pre-built container)
- **METViewer (Containers for MySQL and METViewer)**
  - https://github.com/NCAR/container-dtc-metviewer/
- **End-to-end NWP container online tutorial**
  - https://dtcenter.org/met/docker-nwp/tutorial/container_nwp_tutorial/index.php

Future Work

- Expanded use for in-person tutorials
- Containerize Gridpoint Statistical Interpolation (GSI) data assimilation code
- Seek projects for using containers in cloud computing
- Deploy MET and METViewer containers to DockerHub
- Explore alternatives to Docker
  - Root access requirement is limiting
  - Issues mapping directories in Windows
  - Need to run with multiple nodes on HPC