

ARL-TR-8930 • MAR 2020



Leveraging Software Container Technologies to Enable Data-Centric Applications on High Performance Computing (HPC) Resources

by Matthew Dwyer and Mark Mittrick

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



Leveraging Software Container Technologies to Enable Data-Centric Applications on High Performance Computing (HPC) Resources

Matthew Dwyer and Mark Mittrick

Computational and Information Sciences Directorate, CCDC Army Research Laboratory

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) March 2020		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To) October 1, 2018 – September 30,2019	
4. TITLE AND SUBTITLE Leveraging Software Container Technologies to Enable Data-Centric Applications on High Performance Computing (HPC) Resources				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Matthew Dwyer and Mark Mittrick				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) CCDC Army Research Laboratory ATTN: FCDD-RLC-S Aberdeen Proving Ground, MD 21005				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-8930	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT For an intelligence analyst, sifting through large quantities of unstructured text trying to comprehend the data can be a manual and highly time-consuming process. Decisive Analytics Corporation (DAC) developed a suite of visual tools and capabilities in order to help address the Army's problem. This report documents the process of instantiating and containerizing the DAC software onto the US Army Combat Capabilities Development Command Army Research Laboratory's High Performance Computing environment.					
15. SUBJECT TERMS High Performance Computing, HPC, DOD Supercomputing Resource Center, DSRC, natural language processing, entity disambiguation					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 17	19a. NAME OF RESPONSIBLE PERSON Matthew Dwyer
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) (410) 278-6819

Contents

List of Figures	iv
List of Tables	iv
1. Introduction	1
2. Background	2
2.1 ARL DSRC and HPC Modernization Program	3
2.2 Software Virtualization (Container) Technology	4
2.3 Decisive Analytics Corporation (DAC): SBIR Program	5
2.4 DAC Resolution and Disambiguation (DRADIS)	6
3. Methods	6
3.1 Staging	6
3.2 Deployment	7
4. Conclusion and Future Work	7
5. References	9
List of Symbols, Abbreviations, and Acronyms	10
Distribution List	11

List of Figures

Fig. 1	The Army Cloud computing platform separated by ingest, storage, machine-learning, and interface components	1
Fig. 2	A side-by-side comparison of container virtualization technology (left) and hypervisor-based virtualization technology (right) ³	4
Fig. 3	A high-level overview of the DRADIS capability workflow	6

List of Tables

Table 1	Summary of unclassified HPC systems sorted by site ²	3
---------	---	---

1. Introduction

In order to provide an array of functionality and to minimize future maintenance efforts, software applications are now being divided into smaller subcomponents. Each subcomponent requires a unique set of dependencies and an environment configured specifically for the software. Data-centric applications leverage this paradigm to separate functionality between different components of the data science lifecycle including data ingest, data storage, data processing, and data visualization. Figure 1 shows this separation of functionality as it applies to the Army Cloud computing platform.

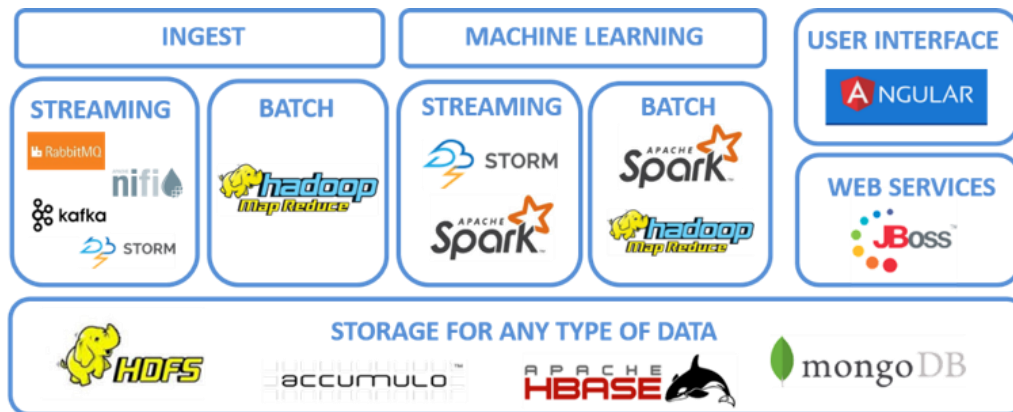


Fig. 1 The Army Cloud computing platform separated by ingest, storage, machine-learning, and interface components

Executing these applications on separate environments is a cumbersome task, but is often necessary to achieve scaling requirements. Every environment has differences in its configuration and users do not always have the ability to make the necessary changes to the execution environment. This is the case for the High Performance Computing (HPC) resources at the ARL DOD Supercomputing Resource Center (DSRC), where users do not have administrator privileges. Traditionally, applications executed on HPC resources need to be built from source and configured with the specific HPC architecture in mind. Because of the unique configuration of the HPC resources, this transition is nontrivial and is, often times, deemed to be more trouble than it is worth. Software virtualization, referred to as containers, provides a mechanism to ease this transition. Through leveraging this container technology, application dependencies and environment configurations are packaged with the software application. The resulting package, called a container image, can be launched on systems that are equipped with a container technology. The ARL DSRC is equipped with this container runtime technology. This report documents the efforts to use container technology to transition a data-centric, natural language processing (NLP) software application called DRADIS

developed by Decisive Analytics Corporation (DAC) to HPC resources across the DOD.

2. Background

For an intelligence analyst, sifting through large quantities of unstructured text trying to comprehend the data can be a manual and highly time-consuming process—one that needs to be improved upon.

To address the Army’s problem, DAC developed a suite of visual tools and capabilities that extracts entities, relationships, topics, and semantic concepts from unstructured text and multi-source datasets in order to automatically disambiguate entities and build relationship-rich entity network graphs through their NLP pipeline.

Meaning from text is extracted, indexed, and made searchable in order to help analysts quickly accomplish their mission. The relationships and concepts in which an entity takes part as well as the entity’s surrounding context are also used as features within the subsequent entity disambiguation. A topic modeling component organizes content and entities into automatically learned categories, which make evaluating whether two textual mentions refer to the same real-world entity a tractable problem for our machine-learning algorithms. Once entity disambiguation has clustered and resolved the entities, the system presents entity-resolution recommendations to the analyst instead of the analyst having to manually tag and resolve each mention of an entity.

With these tools, Army analysts are now able to see automatically generated network graphs of disambiguated entities. This means textual mentions of relationships for an entity that occur within separate documents but still refer to the same real-world person, place, or location can be resolved to the same entity node in the graph even if they use different spellings or titles for the entity. The resultant network graphs can be used by the analyst to vet the algorithm’s results, allowing the analyst to quickly refer to the source material for each entity mention and relationship to ensure accuracy before approving and publishing the results into their targeting tool.¹

2.1 ARL DSRC and HPC Modernization Program

Established by Congress in 1992, the High Performance Computing Modernization Program (HPCMP) has four major areas of focus including the modernization and support of the DOD Research and Engineering Network (DREN), HPC systems, software and training, and the Data Analysis and Assessment Center (DAAC). The ARL DSRC, located at Aberdeen Proving Ground, Maryland, is one of five DSRCs that provide HPC resources to the DOD community under the management of the HPCMP. DSRCs are located across the United States and associated with the different branches of the DOD. Each DSRC provides a spectrum of support for its users, but each has separate areas of focus.

The ARL DSRC is focused on providing classified cycles where 90% of the classified cycles the HPCMP offers are supported at an ARL facility. The US Air Force Research Laboratory (AFRL) DSRC provides development and support of user services and HPCMP-wide software management. The Army Corps of Engineers Engineer Research and Development Center (ERDC) DSRC is home to HPCMP management, provides the largest amount of unclassified computing hours, and is home to the DAAC, which focuses on large-scale visual analysis of modeling and simulations executed on the HPCMP resources. Navy DSRC provides HPC user support, long-term data archival, and expertise in Advanced Computational Environment and Climate Weather Ocean Modeling. All four of the aforementioned DSRCs provide production HPC cycles to the DOD community. Together these 4 DSRCs provide over 650,000 production compute cores to the unclassified community. Table 1 provides details on each system including the number of cores, memory per node, and expected decommission date. The Maui High Performance Computing Center (MHPCC) has a focus on emerging technology, high productivity computing, and focused HPC-backed solutions. To support these focus areas, the MHPCC resources typically have a high turnover rate and are not included in the production resources offered by the HPCMP.

Table 1 Summary of unclassified HPC systems sorted by site²

System	Site	No. of cores	No. of std. nodes	Mem per node (GB)	Decommission date
Mustang	AFRL	56448	1128	192 768	12/2023
Thunder	AFRL	125888	3216	128 768	03/2020
Centennial	ARL	73920	1784	128 512	07/2021
Excalibur	ARL	101184	3098	128 512	02/2020
Onyx	ERDC	214568	4810	128	05/2021

Table 1 Summary of unclassified HPC systems sorted by site² (continued)

System	Site	No. of cores	No. of std. nodes	Mem per node (GB)	Decommission date
Copper	ERDC	14720	460	64	01/2020
Conrad	NAVY	51008	1523	128 512	09/2020
Gaffney	NAVY	35328	704	192 768	01/2024
Gordon	NAVY	51008	1523	128 512	09/2020
Koehr	NAVY	35328	704	192 768	01/2024

2.2 Software Virtualization (Container) Technology

Container technology is a lightweight software virtualization technology that is enabled through leveraging the kernel virtualization functionality. Unlike virtual machines (VMs), containers are built and run using the host machine’s kernel. When executed, this allows containers to bypass a layer of translation and results in a negligible performance overhead as opposed to VMs that use a dedicated kernel and accompanying layer of translation (see Fig. 2). The performance achieved in a container makes the technology more suitable for environments where major variances in kernel functionality are not of concern.

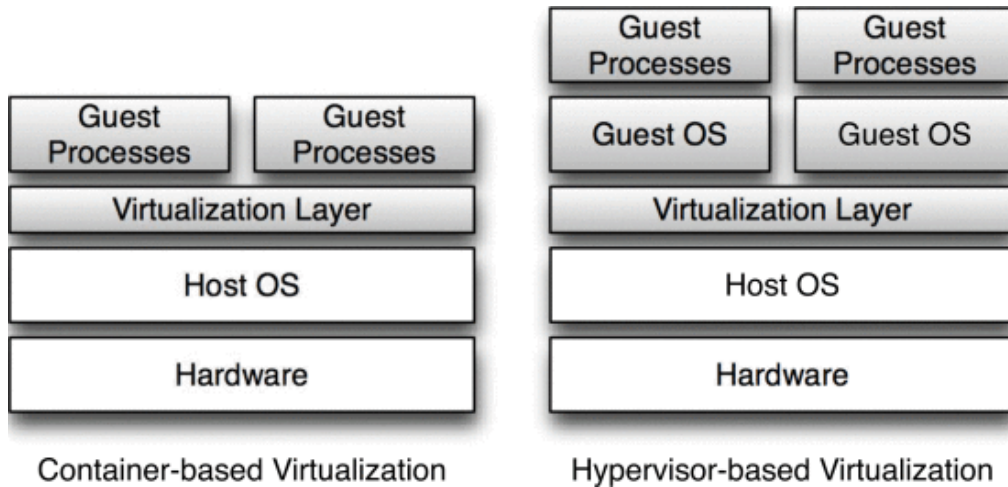


Fig. 2 A side-by-side comparison of container virtualization technology (left) and hypervisor-based virtualization technology (right)³

Though container technology is built into the kernel and is capable of being leveraged by any user, software technologies that provide mechanisms to facilitate the creation, launching, and management of containers burst onto the scene. In early 2013, Docker was publicly released as open source. Since then, Docker has evolved into the primary container ecosystem supported largely by enterprise data center applications. Docker Hub, which is a web repository for container images, currently

hosts over 500,000 images that can be leveraged to create secure private container repositories.⁴

Developed by a team of researchers out of Lawrence Berkeley National Laboratory, the Singularity software application was launched in 2015 as an open-source project aimed at extending the benefits of container technology to the HPC community.⁵ In 2017, the Singularity runtime software was chosen as the supported container runtime technology by the HPCMP and is now installed as a user-level application on all of the HPCMP's production systems.

2.3 Decisive Analytics Corporation (DAC): SBIR Program

DAC is an employee-owned technology company headquartered in Arlington, Virginia, with offices in San Diego, California; Huntsville, Alabama; Colorado Springs, Colorado; and Dayton, Ohio. Founded in 1996, DAC provides systems engineering at the Office of the Secretary of Defense level, as well as develops cutting edge technologies for some of the country's most advanced laboratories such as ONR, ARL, AFRL, and DARPA. Under these efforts DAC provides Fusion Algorithms, Software, and Services that automate the human decision-making process. DAC's focus is on military and intelligence applications that are characterized by data-rich and information-poor environments, uncertain outcomes, and real-time constraints. DAC develops these algorithms, software, and services through three broad areas. First, they develop innovative computational probability techniques enabling dynamic inference under uncertainty. Second, they build machine-learning capabilities that perform real-time data mining for complex, time-constrained problems. Finally, they provide these capabilities at scale in cloud-computing environments that solve real-world problems. These capabilities have been transitioned from the lab to operational environments at the Navy, Army, SOCOM, DTRA, FBI, SMDC, and the USMC. DAC's capabilities have been developed through Small Business Innovation Research (SBIR) and BAA efforts and are delivered with Government-Purpose Rights and therefore no license fees; DAC software engineers support integration, test, and deployment into operational systems.⁶

The SBIR program is a highly competitive awards-based program that encourages small businesses to partner with the federal government in order to further scientific research, stimulate technological innovation, and fill capability gaps, with the incentive to potentially profit from its commercialization.⁷

Over the past 6 years, DAC was awarded over \$7,000,000 in funding to develop, test, and integrate NLP capabilities. This funding was awarded through multiple contracts with the following programs: ARL SBIR program (Phase II contract

W911QX-05-C-0009), Army INSCOM Futures program (SBIR Phase III contracts W911W4-13-C-0008, W911W4-13-C-0007), and the DOD Rapid Innovation Fund (W911QX-15-C-0031).

2.4 DAC Resolution and Disambiguation (DRADIS)

The DRADIS capability leverages various technologies and methods to provide a comprehensive solution to processing and analyzing unstructured text. At the highest level, the DRADIS solution includes the following processes: NLP entity extraction, statistical topic modeling (STM), semantic role labeling (SRL), and both intra-document and inter-document resolutions. The DRADIS capability workflow combines both streaming and batch data ingest, data storage, data processing, and data visualization into one solution. Unstructured text is ingested into the data storage mechanism, processed, and visualized. Figure 3 illustrates how the DRADIS solution uses NLP to extract entities and refines those extractions with STM and SRL to provide semantic concepts, topics, and relationships. The processed information is then resolved using an entity disambiguation algorithm before it is placed in a repository where analysts can view and interact with the results.

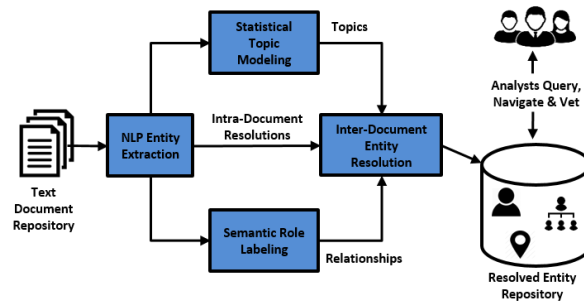


Fig. 3 A high-level overview of the DRADIS capability workflow

3. Methods

The rapid deployment of the DRADIS software on HPC resources involved two major phases: staging and deployment.

3.1 Staging

During the staging phase, a Docker container image containing the DRADIS software stack was provided by DRADIS developers. The image was converted to a suitable format, Singularity image, to run on the HPC machines. This was accomplished by using the conversion mechanisms built into the Singularity software. Though the successful execution of converted container images can be

achieved in a one-step process, the particular configuration of the provided Docker container required modification to execute the stack in a container image using Singularity runtime software. Once successfully modified, the stack was tested and verified to run using the Singularity runtime software via the following steps:

- 1) Convert Docker container image to a singularity container image
 - a. `singularity build docker://<image-name>:<image-tag>`
- 2) Resolve conversion conflicts
 - a. Permissions + read write access
- 3) Verify container functionality in development mode
 - a. Navigate to the DRADIS interface from the browser

3.2 Deployment

Once the DRADIS stack was verified to execute successfully in a container in the staging area we transitioned to the deployment phase. In this phase we performed the following tasks:

- 1) Transfer container to HPC resources
 - a. `scp <container-name>.simg centennial:~/`
- 2) Create batch script to run container on compute resources
- 3) Launch batch script
 - a. `qsub run-dradis.pbs`
- 4) Verify that the container is running
 - a. `qstat -u <username>`
- 5) Connect to the external facing interface
 - a. Port forward to the external facing interface on the node
 - b. Navigate to the DRADIS interface from the browser

4. Conclusion and Future Work

The methods for staging and deploying the DRADIS stack resulted in a successful deployment of the DRADIS software on HPC machines across the DSRC. Users were able to follow deployment steps to launch batch jobs and connect to the running jobs.

Moving forward we expect efforts to have the following areas of focus:

- 1) Ingestion of data sets with
 - a. Scaling quantities of data
 - b. Streaming data
 - c. Disparate data
- 2) Scale the software stack both vertically and horizontally to include maximizing the efficiency of ingestion and analysis processes
- 3) Follow-on software (TechTrakr)
 - a. Separate components of the stack into separate containers

5. References

1. Decisive Analytics Corporation. ARL SBIR success – decisive NLP. Arlington (VA): Decisive Analytics Corporation; 2015.
2. DoD High Performance Computing Modernization Program. Systems Summary. Aberdeen Proving Ground (MD): DoD High Performance Computing Modernization Program; 2019 July [accessed 2019 Sep 3]. <https://centers.hpc.mil/systems/summary.html>.
3. Xavier MG, Neves MV, Rossi FD, Ferreto TC, Lange T, De Rose CAF. Performance evaluation of container-based virtualization for high performance computing environments. Proceedings of the 2013 21st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing; 2013; Belfast, UK. Piscataway (NJ): IEEE.
4. Zhao N, Tarasov V, Albahar H, Anwar A, Rupprecht L, Skourtis D, Warke AS, Mohamed M, Butt AR. Large-scale analysis of the docker hub dataset. Proceedings of the 2019 IEEE International Conference on Cluster Computing (CLUSTER); 2019 Sep 23–26; Albuquerque, NM. Piscataway (NJ): IEEE.
5. Kurtzer GM, Sochat V, Bauer MW. Singularity: scientific containers for mobility of compute. PLoS One. 2017;12(5):e0177459.
6. About Us. Arlington (VA): Decisive Analytics Corporation; n.d. [accessed 2019 Sep 3]. <http://www.dac.us/aboutus>.
7. About SBIR. The Small Business Innovation Research. [accessed 2019 Sep 3]. <https://www.sbir.gov/about/about-sbir>.

List of Symbols, Abbreviations, and Acronyms

AFRL	US Air Force Research Laboratory
ARL	Army Research Laboratory
DAAC	Data Analysis and Assessment Center
DAC	Decisive Analytics Corporation
DARPA	Defense Advanced Research Projects Agency
DOD	Department of Defense
DRADIS	DAC resolution and disambiguation
DREN	DOD research and engineering network
DSRC	DOD Supercomputing Resource Center
DTRA	Defense Threat Reduction Agency
ERDC	Engineer Research and Development Center
FBI	Federal Bureau of Investigation
HPC	High Performance Computing
HPCMP	High Performance Computing Modernization Program
MHPCC	Maui High Performance Computing Center
NLP	natural language processing
ONR	Office of Naval Research
SBIR	Small Business Innovation Research
SMDC	Space and Missile Defense Command
SOCOM	US Special Operations Command
SRL	semantic role labeling
STM	statistical topic modeling
USMC	United States Marine Corps
VM	virtual machine

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 CCDC ARL
(PDF) FCDD RLD CL
TECH LIB

2 CCDC ARL
(PDF) FCDD RLC S
M DWYER
FCDD RLC IT
M MITTRICK