INSTITUTE FOR DEFENSE ANALYSES

Software Container Architecture for Distributed Mission Processes

Robert M. Rolfe, *Project Leader*

Francisco L. Loaiza-Lemos
Kevin E. Foltz
Russell J. Smith
Jagdeep Shah

October 2016

Approved for public release;
distribution is unlimited

IDA
NS D-8204

Log: H 2016-001126

Copy

INSTITUTE FOR DEFENSE
ANALYSES
4850 Mark Center Drive
Alexandria, Virginia 22311-1882

About This Publication

This work was conducted by the Institute for Defense Analyses (IDA) under Central Research Project C 5173. The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

For more information:

Robert M. Rolfe, Project Leader
Rolfe@ida.org, 703-845-6665

Margaret E. Myers, Director, Information Technology and Systems Division
mmyers@ida.org, 703-578-2782

Software Container Architecture for Distributed Mission Processes

Francisco L. Loaiza-Lemos
Kevin E. Foltz
Russell J. Smith
Jagdeep Shah

Robert Rolfe, *Project Leader*

October 26, 2016

INSTITUTE FOR DEFENSE ANALYSES

Abstract

The challenges posed by big data, both at the global and local levels, are viewed as critical challenges for the Department of Defense's (DoD) ability to accomplish its future missions. In addition to recent advances in the area of photonic communication technologies that may provide efficient means for rapid processing of large data sets by facilitating efficient data movement, we argue that leveraging the developments in the area of containerization and virtual machine technologies must also form part of the solution to big data analytics. This is so not only because virtual machines and containers offer an elegant way for implementing large-scale parallel processing but also because their use offers substantive cybersecurity advantages. As we recommended in our initial white paper,[1] the full advantages of these technologies will require that the solution architectures be optimized by considering all aspects of the system, including software, i.e., that they be holistically optimized. We also mention in this document areas where this type of holistically optimized combination of containerization and photonics could be a game changer for DoD.

[1] *Real-Time Information Extraction from Big Data,* R. Rolfe, J. Shah, F. Loaiza, IDA Non-Standard Document NS D-5618, October 2015.

Executive Summary

The Department of Defense (DoD) has the potential to leverage the next wave of innovation expected from developments such as the Internet of Things (IoT), which will make new and dramatic demands on the Department's ability to process massive quantities of data efficiently and effectively. This is because *big data* from the IoT will feed all the applications that are expected to provide quick and timely decision information to DoD decision makers. One approach to managing real-time information extraction from big data is using technologies that offer a realistic path toward achieving high levels of parallel distributed computing in combination with hardware solutions that are reaching a sufficient level of maturity. Specifically, developers' rapid progress and high level of interest in the use of virtual machines (VMs) and containers in combination with the recent developments in photonic technology makes the design of solution architectures that use both very appealing.

We argue that it is essential to holistically optimize both the VMs and containers, as well as the photonics interconnected processors that will compose the solution architectures. Marrying the excellent potential for scaling up that containers offer with the next-generation photonics processors capable of ultrafast in-memory graph data structure traversals, e.g., query executions, could potentially provide the required real-time information extraction capability for big data problems of interest to DoD. This approach would provide more than two orders of magnitude improvement in Performance/Watt for geographically localized nodes, as well as the corresponding cost benefits.[2] A subset of big data applications can be addressed by parallelization, and they are the focus of this paper. We mention areas where this holistically optimized combination of containerization and photonics could be a game changer for DoD.

1. Introduction

In a recent white paper,[3] the authors described the potential for using Silicon Photonics technology, developed by Defense Advanced Research Projects Agency's (DARPA) Photonically Optimized Embedded Microprocessors (POEM) Program. As noted in the white paper, IBM's Throughput Optimized POEM System (TOPS) program also investigated in detail the design and performance of a peta-scale, rack-size graph processor computer system that is holistically optimized for addressing big data problems by using underlying silicon photonic technologies.

[2] *See Real-Time Information Extraction from Big Data,* IDA Non-Standard Document NS D-5618, October 2015 for a description of cost and power savings related to a photonically interconnected graph processor.

[3] *Real-Time Information Extraction from Big Data*, R. Rolfe, J. Shah, F. Loaiza-Lemos, IDA Non-Standard Document NS D-5618, October 2015.

Also recently, one of the authors performed a short analysis and assessment of the advantages and potential uses of containers, as well as VMs, in the context of cybersecurity and as a means to reduce the time needed to obtain authorization to operate (ATO) for new software applications that will be hosted in DoD networks.[4] As noted therein, the extensive use of containers by some of the major commercial companies, such as Google, is a clear indication that containers, as well as VM technology, have reached a very high degree of maturity and have decisively improved the model of data center utilization and software reuse across systems. Given the above, the authors propose that container-based solution architectures merit consideration as a viable option for implementing scalable parallel distributed computing solutions.[5]

It should also be noted that interest in the development and maturation of graph processor technologies continues apace, as shown by the Broad Agency Announcement (DARPA-BAA-16-52) from the DARPA Microsystems Technology Office HIVE (Hierarchical Identify Verify Exploit) Program, issued on August 2, 2016, which solicits "*research proposals for the development of a generic and scalable graph processor specializing in processing sparse graph primitives*."[6] The desire to make progress in this area is based in part on the fact that many of the big data problems can be formulated in terms of Graph Problems,[7, 8] where data is represented as the vertices and connecting edges of an abstract graph. Ultrafast, single-purpose graph processors could, therefore, offer an elegant solution to many of the big data challenges.

2. Research Questions

In exploring the challenges and opportunities that a combination of containerization technology and graph processors using photonic communication presents, we asked the following questions, which will be addressed in the indicated sections:

- How can container input/output (I/O) take advantage of optically optimized processors? (Section 3)
- What major issues exist when considering (1) intra-processor containers, (2) processor-to-processor data exchanges using locally optimized photonics communications, and (3) processor-to-processor exchanges using globally optimized photonics communications? (section 3)

[4] *Rapid Application Deployment with Secure Hosting,* L. Odell, R. Rolfe, IDA Document in preparation.

[5] "Borg, Omega, and Kubernetes," B. Burns, B. Grant, et al., ACMQUEUE Volume 14, issue 1. http://queue.acm.org/issuedetail.cfm? issue=2898442

[6] https://www.fbo.gov/index?s= opportunity&mode= form&id=daa4d6dbee8741f56d837c404eac726d &tab =core&_cview=1

[7] https://en.wikipedia.org/wiki/Graph_database

[8] http://hama.apache.org/

- Which future mission applications could be expected to benefit the most from parallel distributed computing and ought to be considered first for multiple-node container technology? (Section 4)

Among the mission areas that are likely to be challenged by big data processing demands, we consider the following:

- Command, control, communications, computers, intelligence, surveillance and reconnaissance (C4ISR) net assessment and analyses for command and control (C2);
- Decision support (e.g., Global Tag Track Locate for weapons of mass destruction (WMD) for homeland defense; C2 for global drone operations; sensor-based physical security, including national border security);
- Cyberspace command and control decision support (e.g., real-time network defense, real-time network offense);
- Cognitive big data and cloud robotics and autonomy (e.g., intelligence collection and exploitation applications);
- Organization knowledge management and retention (e.g., searches against large graph database repositories);
- Databases with homomorphic encryption (e.g., encrypted searches and manipulation of encrypted data).

3. Container Technology Coupled with Photonics Communications and Graph Processors – A Proposed Solution Architecture

We address the first question posed in the preceding section, namely, *How can container I/O take advantage of optically optimized processors?* by proposing a solution architecture, whose main characteristics are depicted in Figure 1.

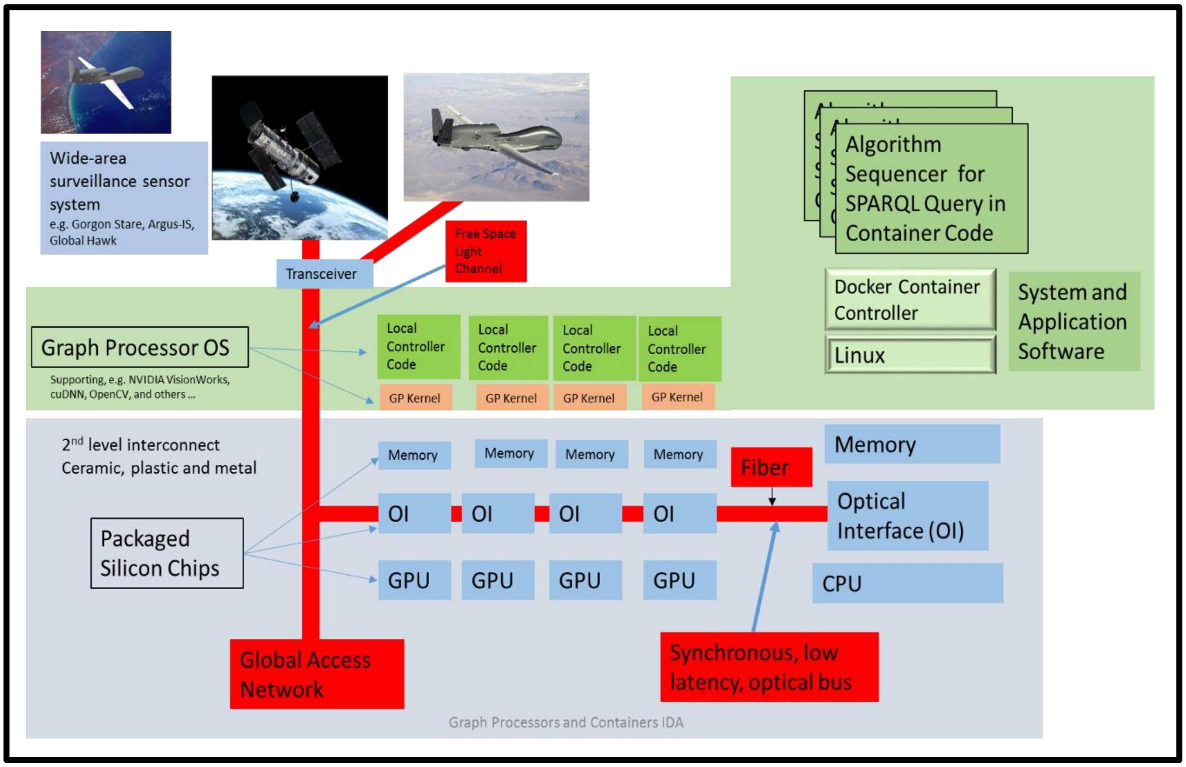
Wide-area surveillance sensor system
e.g. Gorgon Stare, Argus-IS, Global Hawk
Transceiver
Free Space Light Channel
Algorithm Sequencer for SPARQL Query in Container Code
Docker Container Controller
Linux
System and Application Software
Graph Processor OS
Supporting, e.g. NVIDIA VisionWorks, cuDNN, OpenCV, and others ...
Local Controller Code
Local Controller Code
Local Controller Code
Local Controller Code
GP Kernel
GP Kernel
GP Kernel
GP Kernel
2nd level interconnect
Ceramic, plastic and metal
Memory
Memory
Memory
Memory
Fiber
Memory
Packaged Silicon Chips
OI
OI
OI
OI
Optical Interface (OI)
GPU
GPU
GPU
GPU
CPU
Global Access Network
Synchronous, low latency, optical bus
Graph Processors and Containers IDA

Figure 1. Basic Architecture for Containers Deployed in Nodes Hosting Photonics Communications and Graph Processors

The figure shows schematically how a forward-deployed node hosting an array of ultrafast graph processors (GPUs), optically interconnected via a photonics bus and having access to large capacity memory banks (lower portion of Figure 1), could provide the large-scale abstract graph-processing capability required to traverse terabits of data per second,[9] which in some scenarios would be dynamically updated through data feeds linked to sensor platforms, e.g., Gorgon Stare, Argus-IS or Global Hawk.[10] These processors would have a second-level interconnect (ceramic, plastic, and metal) to support the required component connectivities.

[9] An area not covered in this paper but relevant to any future implementation of the proposed architecture is the specification of a system architecture that shows how the terabits of data are divided into different memory segments (see Figure 1), how each memory segment is optically accessed, and the pros and cons of the corresponding GPU requiring access from a given memory segment versus all GPUs requiring access to the total terabyte of memory. This will most likely be application-specific. Once such an architecture is developed, it would be interesting to analyze whether there are performance benefits to such an approach in addition to the expected software reusability, security, and cost.

[10] Another interesting question to investigate in a future phase of this study is how much of the big data analysis can be performed on a platform such a Global Hawk and how much needs to be performed at a large data center on the ground. The latter would be the case if comparison with large sets of historical data, spread over petabytes of memory, is required. Once again, this will be application-specific.

In addition, each graph processors would have its own GP kernel and local controller code, which, depending on the requirements, could be made very lightweight. Riding atop this configuration of special-purpose hardware and software would be a full operating system (OS); in our case, we chose Linux due to the robust support it offers for containerization and VM solutions, (shown within the green layer in Figure 1).

Within the Linux OS would be any number of containers, which again could be as lightweight as needed. In the proposed solution architecture, these containers are designed to execute single-purpose SPARQL queries. The queries could be pertinent to target tracking or any of the many tasks associated with wide area detection and surveillance.

Nodes configured with photonics communications supporting highly efficient graph processors, which can be tasked by software encapsulated in containers, offer a very attractive possibility for *ad hoc* tasking and reconfiguration of sensor utilization in support of highly dynamic operations.[11]

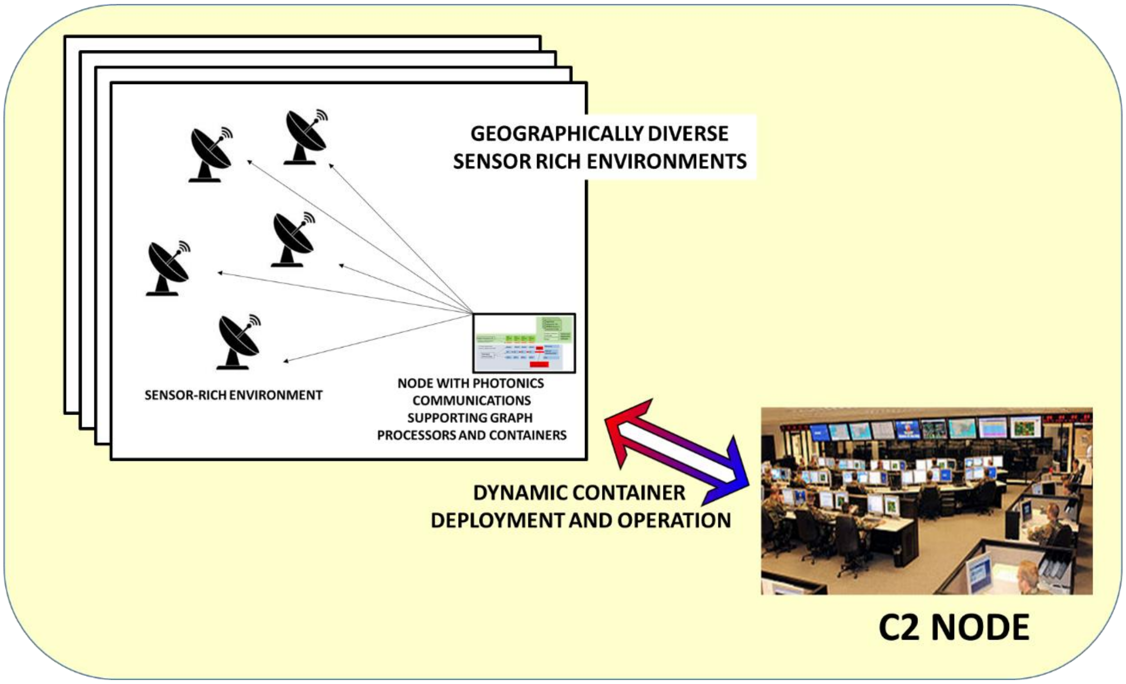
GEOGRAPHICALLY DIVERSE
SENSOR RICH ENVIRONMENTS
SENSOR-RICH ENVIRONMENT
NODE WITH PHOTONICS COMMUNICATIONS SUPPORTING GRAPH PROCESSORS AND CONTAINERS
DYNAMIC CONTAINER DEPLOYMENT AND OPERATION
C2 NODE

Figure 2. Dynamic Container Deployment and Operation

As shown schematically in Figure 2, one could easily envision scenarios in which geographically diverse, sensor-rich environments provide their data feeds directly to forward-deployed nodes hosting arrays of ultrafast graph processors (GPUs) supported by

[11] A future task, analyzing and quantifying the benefits of this approach for applications of interest to DoD, should be considered prior to implementation of the proposed architecture.

photonics communications – rather than centralized processing centers.[12] This architecture would allow the nodes to be dynamically tasked via containers by out-of-theater C2 nodes.

The communications bandwidth requirements associated with large numbers of sensors may be kept small by preprocessing sensor data in real time. The same software deployed in containers in ISR processing at C2 node centers can be deployed on the sensor platform if the containers are designed to support narrowly scoped reusable processing tasks – such as a single SPARQL query on a data stream in the form of an abstract graph. We could envision the exploitation of other preprocessing software deployed at the sensors themselves, so that instead of pushing terabits per second of raw imagery data through the network pipes, the preprocessors would extract the features that are most pertinent for a given task and generate corresponding abstract graph snippets that can be stored and processed at the downstream nodes that host large general-purpose graph processors and the pertinent analysis containers.

4. High Pay-Off Mission Area Examples Amenable to Parallelization

4.1. Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) Net Assessment and Analyses for Command and Control (C2)

It is clear that future decision-support processes must be able to process the ever-increasing volume of data streaming from the totality of sources, both sensors and intelligence, and convert it into actionable decision-support information. Architectures along the lines of what has been suggested in Figure 2 have the potential for satisfying the scalability problem by partitioning the geographic space and reducing the load of centrally configured processing architectures. In addition, containerization may pave the way to exploiting many of the benefits that derive from a much higher degree of modularization in the software application development cycle.

One should also consider the inherent advantages of having narrowly scoped software applications, which can be made much harder to penetrate via malware and other hacking techniques, as an essential design consideration for future implementations of C2.

4.2. Decision Support (e.g., Global Tag Track Locate for WMD for homeland defense, C2 for global drone operations, sensor-based physical security, including national border security)

Tagging, tracking, and locating (TTL) capabilities typically provide the close-range, persistent surveillance necessary to: (1) collect information, (2) interdict or capture persons or things, and (3) apply precision force at the tactical level. These capabilities have been used extensively in operations across the Special Operations Forces (SOF) community for

[12] The communications requirements of such an approach need to be analyzed in a future task.

many years. Recently, TTL capabilities have begun to be applied aggressively and globally in the commercial sector, primarily for supply chain management, but more importantly for *asset tracking* (“things” – cars, containers, trucks, baggage, planes, ships, etc.), *people tracking and locating*, and *activity monitoring* (consumer shopping). DoD logistics and supply organizations have also embraced RF identification technologies and systems – there are future plans to “tag” all soldiers and their equipment as well.

Figure 3 illustrates the various resource layers that provide data streams for the TTL knowledge base. When the mission is situational understanding and alerting, the knowledge base will contain WMD life-cycle activities, products, and materials; context provided by the WMD activities of rogues, criminal groups, and nation-states; and the activities needed to detect, classify, identify, tag, track, locate, and target solids, fluids, gases, biological materials, components, and industrial measurement and production machines on an on-going basis.

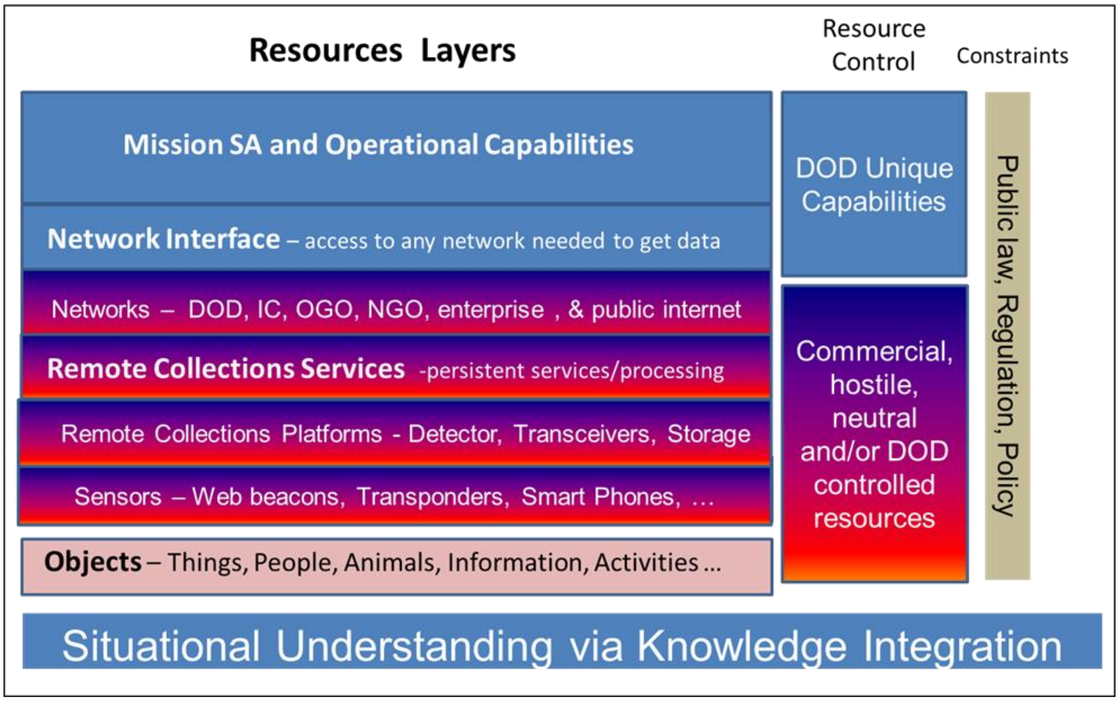
Resources Layers
Resource Control
Constraints
Mission SA and Operational Capabilities
Network Interface – access to any network needed to get data
Networks – DOD, IC, OGO, NGO, enterprise , & public internet
Remote Collections Services -persistent services/processing
Remote Collections Platforms - Detector, Transceivers, Storage
Sensors – Web beacons, Transponders, Smart Phones, …
Objects – Things, People, Animals, Information, Activities …
DOD Unique Capabilities
Commercial, hostile, neutral and/or DOD controlled resources
Public law, Regulation, Policy
Situational Understanding via Knowledge Integration

Figure 3. TTL Resource Layers

There is a steady need to provide globally integrated context across maritime, land, air, space, and cyberspace and to avoid stove-piping knowledge resources. Knowledge resources should be hosted in a dynamic, continuously growing knowledge base that can be processed by graph processors that leverage a library of container-based applications, and that can continuously execute the necessary data extraction needed to perform predictive analytics (i.e., connect the dots) and preemptively detect problems as they develop, thus anticipating how the dots are likely to be connected at any time in the future.

4.3. Cyberspace Command and Control Decision Support (e.g., real-time network defense, real-time network offense)

An intuitively straightforward approach to defending against hostile cyber operations is to have perfect visibility of the data that flows through the network infrastructure that all our key software and hardware applications run on top of. The nature and variety of the data is schematically depicted in Figure 4.

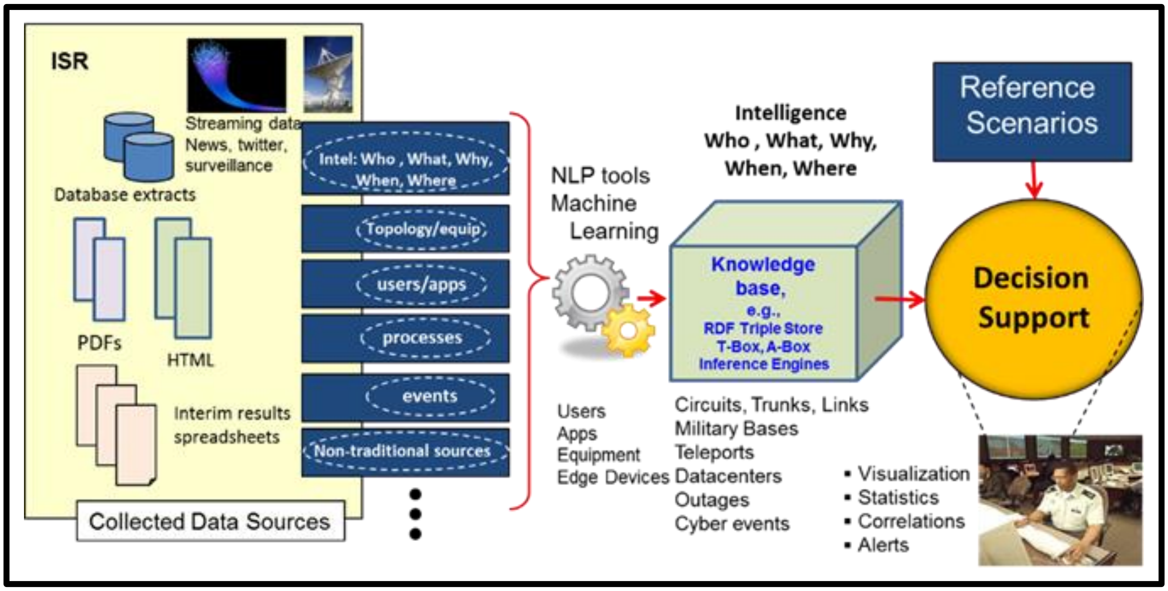

Figure 4. The Data Processing Workload Associated with Network Defense

The left portion of the figure displays key types of sources, such as database extracts, unstructured data in the form of PDF and HTML documents, semi-structured data such as spreadsheets, and other tabular data. These sources would need to be mined via artificial intelligence techniques, such as supervised and unsupervised machine learning, to generate essential information that can be stored in graph databases and processed in near real time to support the decision-making process.

This task requires massive computational capacity, not just to sample the traffic, but also to process the associated metadata, to monitor and update the continuously evolving profiles of the activities our adversaries are engaged in, to identify the types of traces those activities leave—the so-called digital exhaust, and to adjust the response times according to the sensitivity of the various types of information assets that reside at the government and contractor facilities.

Containers, coupled with graph processors supported by photonics communications, arguably could satisfy many of the above-mentioned requirements.[13] In addition, the solution architectures could take advantage of the decentralized processing that containers and graph processors supported by photonics communications offer, which not only scales well to handle the loads in parallel but also make the management and oversight of applications simpler and more effective.

4.4. Cognitive Big Data and Cloud Robotics and Autonomy (e.g., Intelligence Collection and Exploitation Applications)

Semi-autonomous robotic platforms, such as those that have achieved a fairly advanced degree of operational maturity (see Figure 5) will be expected to execute the exceedingly complex operations associated with detection, classification, identification, and recognition of objects of interest.

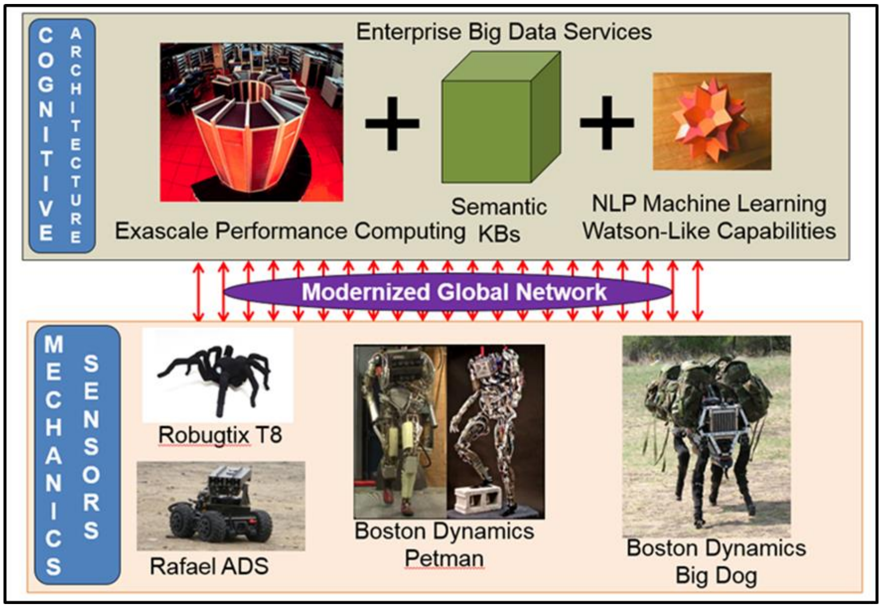
COGNITIVE
ARCHITECTURE
Enterprise Big Data Services
+
+
Exascale Performance Computing
Semantic KBs
NLP Machine Learning Watson-Like Capabilities
Modernized Global Network
MECHANICS
SENSORS
Robugtix T8
Rafael ADS
Boston Dynamics Petman
Boston Dynamics Big Dog

Figure 5. A Cognitive Architecture to Support Semi-Autonomous Robotic Platforms

These cognitive operations will encompass activities such as finding force components under a forest canopy; recognizing a unit from separate observations of its components; drawing inferences; identifying relationships; observing force movements over time; understanding relationships and functionality of aggregates of components; learning, teaching, incorporating experience into a knowledge base; learning through discourse with man and machine; conceptualizing by generalization; predicting and analyzing future situations; evaluating alternative courses of action for different future

[13] Analysis of the photonic network architecture and the associated control plane required to effectively address these requirements is essential prior to any implementation of the proposed architecture.

scenarios via the application of rules and models; integrating heterogeneous models to handle team and group behavior (for groups of cognitive agents); developing team strategies; and coordinating operations of task-constituted reconnaissance teams. These are the most pressing ones in the context of military missions.

Also, as shown in Figure 5, the basic architecture needed to support these semi-autonomous robotic platforms comprises elements quite similar to those discussed in the preceding sections. Here too, there is a need to convert large quantities of data about the environment, the characteristics of objects of interest, their known and expected behavior, etc., into easily processable representations such as abstract graphs. When the abstract graphs are hosted in appropriate repositories, they can then be processed in real time to guide the reaction of the semi-autonomous cognitive agents, leveraging the inherent parallelizable nature of forward-deployed nodes hosting graph processors supported by photonics communications, that can run containers dynamically deployed and activated and specifically designed to retrieve the needed data from the stored abstract graphs.

As the technology matures for both ultrafast graph processors and the containerization of narrowly scoped software that can be dynamically reconfigured and orchestrated to address mission needs, it will be possible to support the computational demands of the rich, multi-dimensional components involved in cognition (see Figure 6).

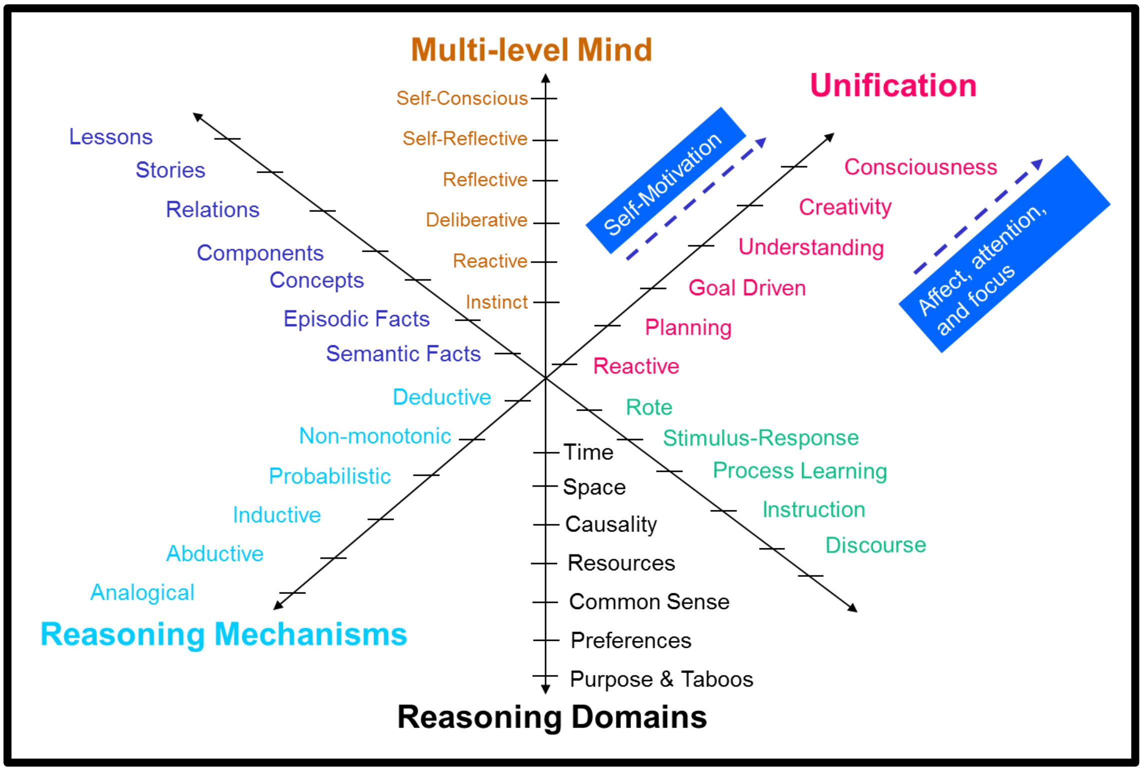

Figure 6. The Multi-Dimensional Components of Cognition

4.5. Organization Knowledge Management and Retention (e.g., searches against large graph database repositories)

The Federal Government has invested hundreds of billions of dollars in the creation of structured data repositories. These information resources tend to be hard to integrate when using traditional *extraction, transformation, and loading* (ETL) techniques based on mappings of the underlying relational database management system (RDBMS) schemata. In fact, one could claim that the traditional ETL approach to retaining and reusing valuable legacy data stores has no chance of success because of its very high cost and amount of time involved.

On the other hand, the conversion of legacy RDBMS data to abstract graph representations is an easy and straightforward process. *Any* abstract graph composed of snippets having the form of **<node><edge><node>** can *always* be loaded into *any graph database application*. This means that after the conversion of the contents of any number of legacy RDBMSs, one can proceed without any technical impediment to collect all the generated graphs into an *abstract graph data lake* (see Figure 7). At this point, the solution architecture that was introduced at the beginning of this paper can be applied again. In this situation, nodes with photonics communications supporting ultrafast abstract graph processors can be deployed to ensure extremely low latency for processing the queries posed by the geographically distributed end users. Because the functionality expected by the end users is encapsulated in the applications that run inside the containers, and because these containers can be dynamically configured to match the needs of the end users, the architecture offers an elegant way to ensure that the investments made to build the information resources are not lost and the useful life of the resources themselves is extended with relatively low cost.

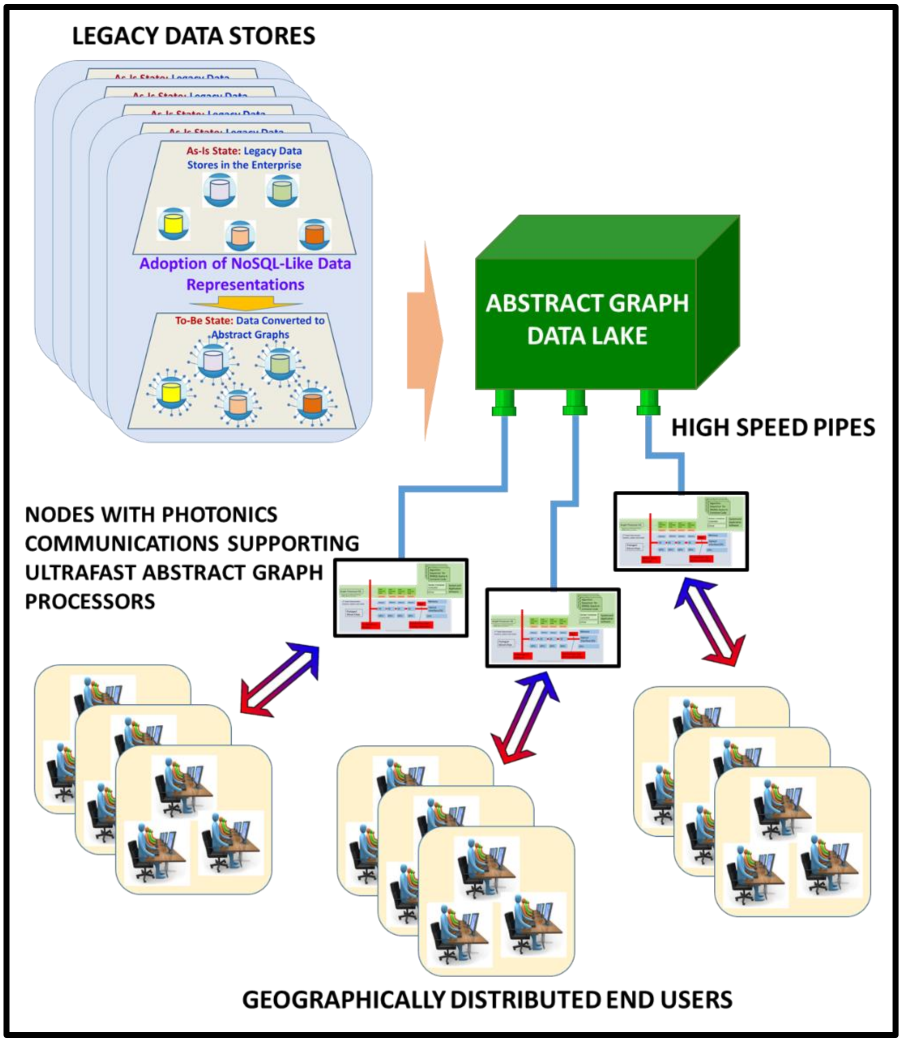
LEGACY DATA STORES
As-Is State: Legacy Data Stores in the Enterprise
Adoption of NoSQL-Like Data Representations
To-Be State: Data Converted to Abstract Graphs
ABSTRACT GRAPH DATA LAKE
HIGH SPEED PIPES
NODES WITH PHOTONICS COMMUNICATIONS SUPPORTING ULTRAFAST ABSTRACT GRAPH PROCESSORS
GEOGRAPHICALLY DISTRIBUTED END USERS

Figure 7. Leverage and Retention of Structured Data

4.6. Databases with Homomorphic Encryption (e.g., encrypted searches and manipulation of encrypted data)

Cybersecurity and cloud processing have inherent conflicts. Cybersecurity requires control over data and cryptographic keys; cloud processing involves offloading data and keys to third parties. The third parties have physical and logical access to all data and keys that they host. As a result, the secure solution is to host in a private cloud because the public cloud cannot be trusted.

Solutions for the public cloud do exist, but they lack security, performance, or both. Any solution that does not encrypt data in the cloud allows the provider to access the data. To alleviate this concern, data can be encrypted before it is sent to the cloud and decrypted when it is retrieved. However, this only works for well-defined blocks of data. Searches over knowledge bases (large connected graphs) and databases and manipulation of large data sets requires downloading large portions of the encrypted data store, decrypting, performing computations, encrypting, and uploading the results. It is often easier to simply host the unencrypted data locally.

Homomorphic encryption allows encrypted data to be uploaded to the cloud and manipulated without decryption. Encrypted values can be compared, sorted, and added while remaining encrypted. This provides the benefits of encryption without the burden of transferring large sets of raw data. The queries are completed in the cloud on encrypted data, and only the encrypted results are returned.

The main challenge with homomorphic encryption is the computational burden it presents. In particular, loading a database into the cloud requires transforming each data element into an encrypted version that is capable of homomorphic encryption operations. This is computationally very expensive for large data sets but also highly parallelizable. A hardware architecture with high degrees of computational parallelization is well-suited to this problem.

5. Acquisition Considerations

Containers provide a tremendous improvement over the way DoD acquires information technology (IT) systems for C2 today in terms of distributed reuse and resiliency. IT system acquisition is the subject of numerous studies, programs, and initiatives, all geared toward improving the speed with which new capabilities are delivered to the warfighter. Improvements are often focused on the process—*how* we buy systems. With the technologies highlighted here, the Department could greatly improve the acquisition process because of *what* we buy.

Containers should become the essential building block for all IT acquisitions. Through abstraction of the computing infrastructure, an acquirer need only focus on the unique aspects of the system requirements; as long as the system can run on the DoD-approved virtualized infrastructure, only the specific capability or function must be acquired. Through abstraction of the IT infrastructure, costs are driven down by not including every component (computing, storage, software, hardware, security, network, etc.) as an integrated vertical solution to a mission need. Virtualization, through containers, greatly simplifies the acquisition and maintenance process to make what we buy easier, cheaper, and quicker to produce.

6. Next Steps

Both the container technology and the technology underlying nodes, with photonics communications supporting the ultrafast abstract graph processors in the proposed solution architectures discussed in the preceding sections, are evolving rapidly and are likely to exhibit substantive modifications in their final formulations. It is, therefore, necessary to continue exploring how the novel features that ultrafast graph processors and containers for software applications will incorporate can be exploited in the context of not only those mission areas mentioned in this document, but also emerging application areas such as homeland security and emergency response operations. With that in mind, we suggest the following next steps in the process:

- Explore the costs and benefits of implementing a highly dynamic deployment and utilization of containers across distributed geographic areas, as compared to centralized computing capabilities.
- Continue monitoring the evolution of container technology, specifically, in the context of additional players, e.g., Microsoft, with particular emphasis on modularization and OS-independence.
- Develop a system design for parallelizable applications of interest to DoD by holistically optimizing the entire system, including photonic communications network and system architectures, software, and photonic technologies. Minimizing latencies introduced by synchronization, optical-electrical-optical (OEO) conversions, etc., needs to be analyzed. Are photonic buffers/memories needed? If so, what are the required capacity and access times?
- Analyze and quantify the performance benefits of such holistically optimized big data analytics for DoD applications.
- Assess the maturity and applicability of results obtained within DARPA research initiatives in the areas of algorithms and hardware specifically targeted to graph processing.
- Reassess how emerging missions for the DoD will impact the use and applicability of the technologies explored in this paper.

7. Conclusions

We assert that containers can provide a very flexible and powerful approach for leveraging the robust processing capabilities that graph processors offer. This is particularly the case when the graph processors are linked via fast, synchronous photonics communications. Deployment and utilization of these containers, in both strategic centralized and tactical distributed modes, may bring DoD the kinds of benefits that commercial enterprises have already achieved: a higher degree of exploitation of available computational resources; more cost-effective software development cycles; and a better paradigm for overall management of the acquisition, operations, and sustainment of capabilities. Also, encapsulating mission-critical software in containers may add further protection against malware and other types of malicious hacking, making the adoption of container technologies not only desirable from a pure economic point of view, but sensible in an environment in which hostile cyber activities could have disastrous consequences if left unchallenged.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE (DD-MM-YY) 11-10-16 00-10-16	2. REPORT TYPE Non-Standard	3. DATES COVERED (From – To)

4. TITLE AND SUBTITLE

Software Container Architecture for Distributed Mission Processes

5a. CONTRACT NUMBER

HQ0034-14-D-0001

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBERS

6. AUTHOR(S)

Francisco L. Loaiza-Lemos, Kevin E. Foltz, Russell J. Smith, Jagdeep Shah

5d. PROJECT NUMBER

C 5173

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESSES

Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, VA 22311-1882

8. PERFORMING ORGANIZATION REPORT NUMBER

NS D-8204
H 2016-001126

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Institute for Defense Analyses

10.

IDA

11. SPONSOR'S / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

Project Leader: Robert M. Rolfe

14. ABSTRACT

The challenges posed by big data, both at the global and local levels, are viewed as critical challenges for the Department of Defense's (DoD) ability to accomplish its future missions. In addition to recent advances in the area of photonic communication technologies that may provide efficient means for rapid processing of large data sets by facilitating efficient data movement, we argue that leveraging the developments in the area of containerization and virtual machine technologies must also form part of the solution to big data analytics. This is so not only because virtual machines and containers offer an elegant way for implementing large-scale parallel processing, but also because their use offers substantive cyber security advantages. As we recommended in our initial white paper , the full advantages of these technologies will require that the solution architectures be optimized by considering all aspects of the system, including software, i.e., that they be holistically optimized. We also mention in this document areas where this type of holistically optimized combination of containerization and photonics could be a game changer for DoD.

15. SUBJECT TERMS

Holistic optimization; ultra-fast graph processors; photonics communications; software containers; big data; big data analytics; Internet of Things (IoT); Command and Control (C2); parallel processing; cyber security

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	Unlimited	15	19b. TELEPHONE NUMBER (Include Area Code)