SYM-AM-18-085

PROCEEDINGS
OF THE
FIFTEENTH ANNUAL ACQUISITION RESEARCH SYMPOSIUM

THURSDAY SESSIONS
VOLUME II

Acquisition Research:
Creating Synergy for Informed Change

May 9–10, 2018

March 30, 2018

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Applying a DEVelopment OPerationS (DevOps) Reference Architecture to Accelerate Delivery of Emerging Technologies in Data Analytics, Deep Learning, and Artificial Intelligence to the Afloat U.S. Navy

CAPT Kurt Rothenhaus, USN—serves as the Program Manager of the PMW 160 Tactical Networks Program Office at PEO C4I in San Diego, CA. CAPT Rothenhaus holds a PhD in software engineering from the Naval Postgraduate School and has served as an Engineering Duty Officer since 2003. Previous assignments include service as Commanding Officer of SPAWAR System Center Pacific, Combat Systems/C5I Officer on USS *Harry. S. Truman* (CVN 75), and Deputy Program Manager for PMW/A 170 (Navy Communications and GPS Program Office). [kurt.rothenhaus@navy.mil]

CDR Kristine "Kris" De Soto, USN—currently serves as the Assistant Program Manager (APM; CANES Development) in PMW160 CANES Development. CDR De Soto was commissioned in 1999, following graduation from The Pennsylvania State University with a BS degree in computer science. CDR De Soto earned an MS in systems engineering with an emphasis in net-centric systems from the Naval Postgraduate School in 2009. Previous assignments include Combat Systems Officer on USS *John C. Stennis* (CVN 74), Protected Communications Manager Navy Multiband Terminal APM in PMW/A 170 Communications Program Office, and Littoral Combat Ship Integration Platform Manager in PMW 760. [kristine.m.desoto@navy.mil]

Emily Nguyen—serves as the Assistant Program Manager for PMW 160 Agile Core Services. Nguyen received her Bachelor of Science in computer science engineering from the University of Colorado and subsequently earned a Master of Business Administration at San Diego State University. Nguyen started as a software developer for Science and Technology projects sponsored by the Office of Naval Research and later transitioned one of the efforts to Marine Corps Systems Command which resulted in her role as the Project Manager at the Rapid Response Integration project providing agile application development of command and control technologies. [emily.nguyen@navy.mil]

Jeffrey Millard—serves as program and technical support for the PMW-160 CANES program and Agile Core Services subsystem. He holds a BA in chemistry from DePauw University and a PhD in physical chemistry from Indiana University. Prior to providing support to PMW-160, Dr. Millard developed mathematical/statistic models of the currency derivatives markets, shock physics material response models for satellite systems, electronics reliability models for missile navigation systems, and optical analyses laser-hardened aircraft canopies. He has designed electronic trading exchanges for non-standard financial assets, and coordinated nuclear weapons effects testing for Air Force and Navy programs. [jeffrey.millard.ctr@navy.mil]

Abstract

Twenty years ago, the Navy began expanding the use of commercial industry information technology (IT) to employ Internet Protocol (IP)–based client server and web-based technologies to improve software effectiveness and affordability on ships and submarines. Coupled with wideband satellite capabilities, these systems increased the Navy's ability to plan, communicate, command and control, and execute increasingly complex missions. With a sound foundation in commercial IT installed in the Fleet, the Navy is looking today to improve warfighting by leveraging emerging technologies in Data Analytics, Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL). These technologies have the potential to change how the Navy fights and will drive changes to the Fleet's Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) architecture and processes. This paper proposes a reference architecture, new processes, and tools to meet the dynamic nature of these emerging

technologies, to include employing the commercial DEVelopment and OPerationS (DevOps) construct. The reference architecture and processes have the potential not only to accelerate the modernization of the afloat Navy networking WAN/LAN infrastructure, but also to deliver important warfighting capabilities to support Command and Control, Intelligence, and Logistics software applications.

Introduction

In 2017, the Space and Naval Warfare System Center Pacific (SSC Pacific) presented senior Navy leadership a technical vision for the future of the afloat Navy in which AI technologies, Cyber, and increased cohesion in mission planning played a central role in warfighting success. In the video scenario presented, a Carrier Strike Group (CSG) in 2035 was deployed and tasked to conduct a humanitarian relief operation in a contested battlespace. Examining many of the concepts highlighted in Brynjolfsson & McAfee (2016), *Second Machine Age*, and the concept of improved human-machine teaming, the video focused on the interaction of the CSG staff as they wrestled with the development of courses of action (COAs) to balance rules of engagement, asset limitations, and force protection. During the planning session on the flagship USS *John F. Kennedy* (CVN 79), both CSG and ship leadership engaged with the ship's computing stack, nicknamed "Kennedy" via a natural language processing interface. Kennedy was able to not only understand the crew's instructions, but also to access tactical and logistic information resident aboard the ship and, using reach-back, interface with ashore command and control nodes. The information Kennedy processed spanned the tactical, operational, and strategic levels, providing the system requisite context across all domains.

In one vignette, SSC Pacific highlighted the potential use case for human-machine learning in a scene featuring a conversation between the staff Operations Officer and the Battle Group Commander regarding COAs. Having monitored the conversation and without prompting, Kennedy interjected with an independent COA not considered by the staff. In the past, human-machine teaming of this level was the stuff of science fiction. Today, however, advances across multiple technologies are making these capabilities a potential reality—albeit in incremental steps.

The Consolidated Afloat Network Enterprise Services (CANES) provides the computing infrastructure and the foundation of the Navy's Information Warfare Platform afloat (NEJ, 2018). To deliver a more mission-effective, cyber-secure and affordable afloat Information Warfare Platform, CANES is modernizing its software application hosting and application integration processes and tools to implement a cloud-enabled DevOps framework. Agile Core Services (ACS) is a critical sub-system of CANES that has two broad categories of capabilities—core services and data analytics. Both support data sharing and analytics for CANES' hosted applications. ACS also provides a platform for the rapid insertion and management of software and facilitates applications' access to commercial analytics tools optimized to support the maritime battlespace. As shown in Figure 1, ACS sits on the stack as part of the CANES system between software application logic and the computing system.

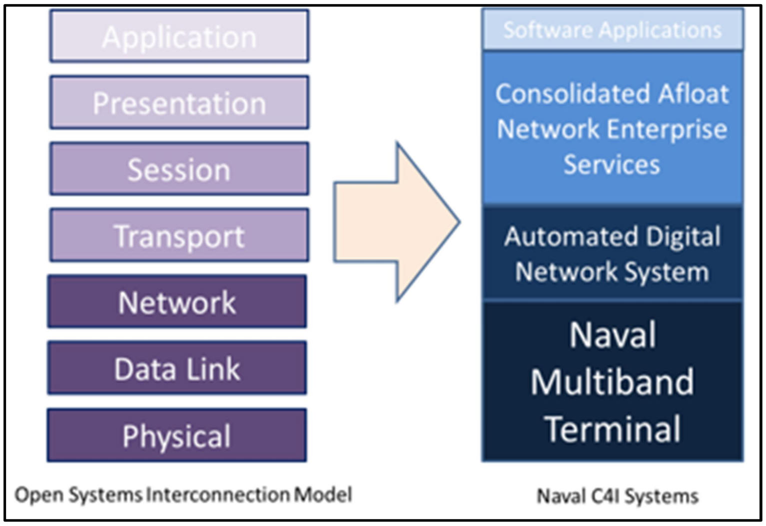

Figure 1. Open System Interconnection Model Mapping to Present Naval C4I Systems

An examination of the emergence of AI and ML technologies indicates that ACS also provides a framework to deliver these emerging technologies at the maritime computing edge afloat. AI technologies applied to command and control, intelligence, and logistics missions have the potential to provide new insights and speed of command. In this paper, we examine how the Navy can expand CANES to support AI and ML. In the first section of the paper, we examine ACS and its present capabilities. After that, we discuss how the current DevOps processes can support AI and ML. Next, we explore some of the commercial AI and ML technologies that exhibit initial potential to support Navy missions afloat. Finally, we cover potential Fleet use of the technologies and discuss implementation at sea.

Agile Core Services (ACS) Capabilities

Agile Core Services (ACS) provides a core set of software services that collectively create a shared framework for applications to build, deploy, and operate mission threads. ACS provides enterprise computing with integrated solutions of commercial off-the-shelf products creating an underlying service infrastructure to support the modernization of applications. Two primary services are provided to promote connection, collaboration, and communication between applications. These are Data Analytics and Common Services. Figure 2 provides an overall system resource flow description.

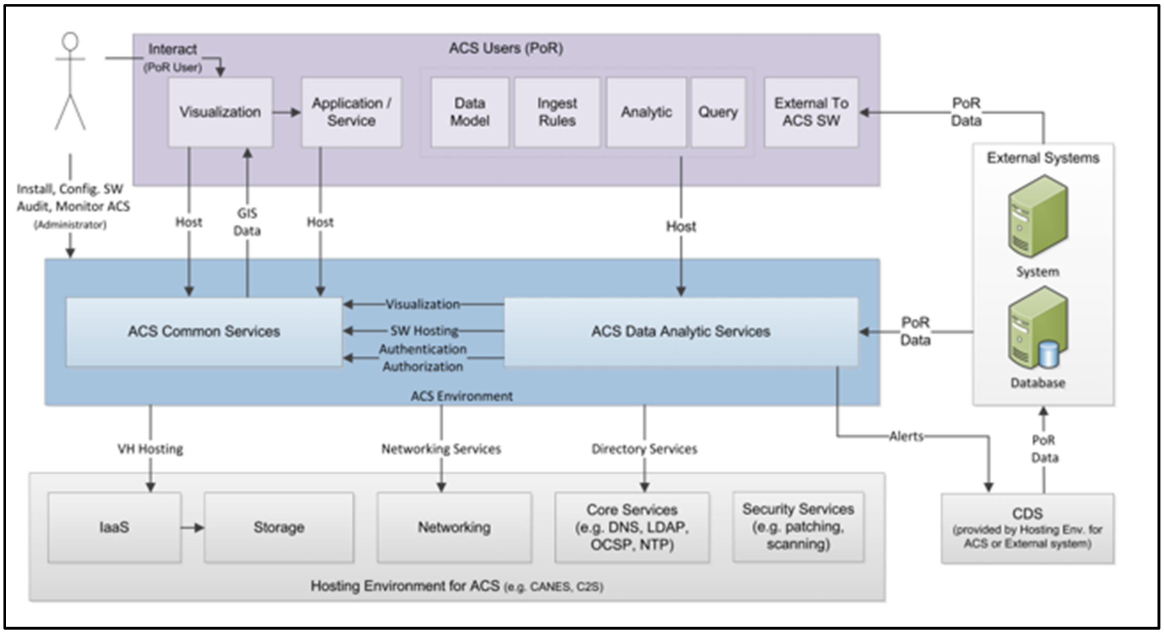

Figure 2. Agile Core Services Functional Diagram

ACS Data Analytics Services

The abundance of data within the Navy has driven the demand for ACS to provide a common data analytic architecture. By leveraging new industry technologies, ACS is able to provide a data analytic architecture that can process far greater volumes and variety of data at unprecedented velocities. The Data Processing and Analytics Framework utilizes clustering for the streaming, batch, and data processing to enable parallel processing for improved performance, robustness, and scalability.

Today, ACS leverages Apache's Spark and Storm, two common open source distributed compute engines. Both perform analytics and distributed compute tasks, but with distinctive implementations and focuses. Apache Spark provides fast cluster computing as a general-purpose distributed computing platform. It does provide limited stream processing; however, Apache Storm is specialized in reliably processing unbounded streams of data. Apache Storm provides real-time analytics, online ML, continuous computation, distributed Remote Procedure Call (RPC), and Extract, Transform, Load (ETL). A Storm topology consumes streams of data and processes those streams in arbitrarily complex ways, repartitioning the streams between each stage of the computation however needed.

The two engines described above are only a subset of the ACS data analytic framework. The following is a synopsis of all of the capabilities:

- Streaming Processing Framework for non-interactive manipulation and analysis of data moving at high velocity
- Batch Processing Framework for efficient, non-interactive analysis of big data stored at rest
- Shared Semantics Framework enabling a common vocabulary across heterogeneous data sources
- Application Programmer Interfaces (APIs) for data ingest, normalization, enrichment, and fusion

- Pre-Defined Query Processor to support interactive queries across heterogeneous data (structured, unstructured, and semantic)
- Alert Processor to generate alerts by comparing data against alert criteria

Identifying trends, finding patterns and relationships, and drawing conclusions are all reasons why data analytics has become vitally important. Given the variety of data available in the maritime environment, including operational data, content data (e.g., documentation, videos, and imagery), authoritative data (e.g., sensor data), and system-generated data (e.g., system logs), the data analytic solution has to be flexible enough to handle structured and unstructured data. With the diversity in data there are great opportunities to develop new decision aids that can help assist the Fleet fight and win.

Until now, many applications have leveraged data to do a specific task. As the Navy moves into the future, data will drive decisions by providing additional solutions that may have been overlooked due to the previous inability to analyze and provide relationships between overwhelming quantities of data. This is where a data analytic engine can facilitate advanced analytics and can accelerate the delivery of emerging technologies in DL and AI. Analytics include ML, data mining, and statistical analysis. When applied in real time and presented in an operational context, analytics can enhance the warfighters' ability to complete complex missions afloat. This unlocks many possibilities for the advancement in applications developed for the warfighter.

ACS Common Services

To support the agility required in present-day software development and to keep pace with the demand for new updates and patches, ACS Common Services provides a suite of services to aid in the modernization of applications. Each service is decomposed below.

Geospatial Data Access Layer and Geographic Information System (GIS) Services

A prevailing set of data that is common to many Navy applications is geospatial data. Geographic data is large in size and requires ample storage in the realm of terabytes. The Geospatial Data Access Layer and GIS Services provides a full suite of geographic data persistence, analytic, query, and mapping capabilities based around standard Open Geospatial Consortium (OGC) interfaces. By providing a common service for applications to retrieve geospatial data, this reduces the burden of hosting multiple map servers, each of which require maintenance while consuming terabytes of storage space reducing cost, man-hours, and storage space.

Mediation

The mediation service provides middleware platforms for hosting and integrating modular Java software components and sharing data via machine to machine messaging using topics and queues. Together these components support implementation of common enterprise integration patterns.

Visualization

Learning new applications can be incredibly challenging for Fleet Sailors due to the current portfolio of diverse user facing software. The visualization service provides a common user interface framework to help support applications with similar mission threads providing consistency to the look and feel of the application. A cohesive user interface presenting different types of data processed in various ways in a common structure enables warfighters to focus on decision making with the presented data instead of struggling with the complexities of using an application. Combining natural language with user interfaces will be the next step in user application interaction.

Platform as a Service (PaaS)

PaaS enables application services to change quickly, innovate easily, and remain competitive by supporting a services/microservices architecture reducing the complexity of building and maintaining the infrastructure typically associated with developing an application. A growing trend is the demand for cloud computing with microservices that can be scaled and deployed separately, enabling shorter release cycles. A PaaS service enables easy deployment of microservices by standardizing how services are executed, maintained, and orchestrated with a standard framework using containers to isolate elements of application deployment which reduces integration risks often seen in monolithic software architectures. As a result, service providers can more easily implement continuous delivery of updates to their services incorporating key practices of DevOps discussed in more detail in the next section.

Development Operations (DevOps) Processes, Capability Delivery

One of CANES' objectives is to lower the barrier of entry for applications that are hosted on or connect to CANES. Serving more than 100 applications from across the Department of Defense (DoD) and the Navy, CANES can drive affordability and increased speed to capability by providing application developers a common development environment, in some cases a cloud-based capability, along with the governance and processes to rapidly progress to the testing, information assurance, and fielding pipeline. To increase their ability to release capabilities, features, and patches out faster than ever before, industry established a software development culture and practice to bridge the gap between development and operations, which is known as DevOps.

Riding on the CANES platform, ACS provides an operational environment where it enables many of the tenets of the DevOps movement. As discussed previously, ACS provides a PaaS, in which automated development, deployment, provisioning, security, and other application lifecycle management tools are supported. In DevOps, PaaS enables applications to have a representative operational environment during development so that each step of taking a capability from development to operations can be automated. Automation accelerates the development cycle, increasing deployment frequency, while maintaining stability.

Beyond providing the technology and tools for application to leverage for DevOps, CANES/ACS has strategically aligned the people, processes, and tools to create the DevOps objective end state, which is depicted in Figure 3. It highlights several organizational and cultural changes that are described in detail in the next section to how CANES/ACS is developed and deployed operationally.

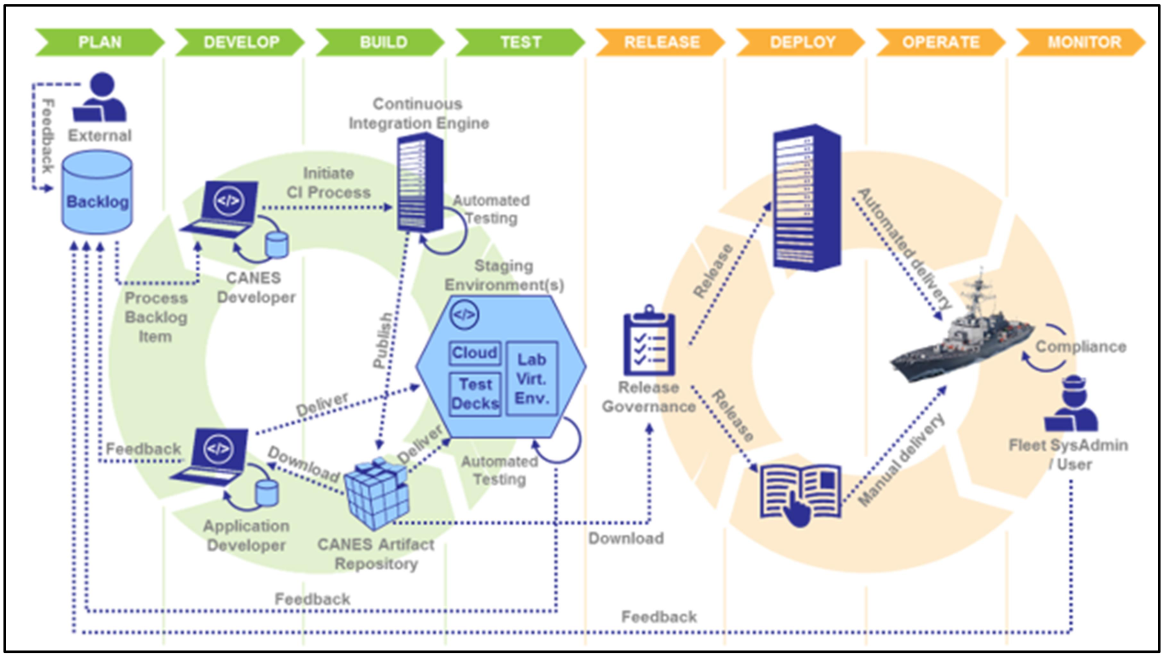

Figure 3. Afloat DEVelopment and OPerationS Framework

Planning

Each cycle starts with planning following the Agile Scrum methodologies. Work is described as product backlog items where it is prioritized based upon criticality and need. An important change that is depicted in the Figure 3 diagram is the feedback loop that comes from operational users and/or automated monitors afloat in the CANES/ACS platform that can provide suggestions and/or issues back to developers. As DevOps matures, this cycle between receiving feedback to deployment of changes to the system becomes shorter and shorter, transitioning from a lifecycle of years to possibly days.

Development/Build

Continuous Integration is a foundational part of DevOps that occurs during development in which developers regularly merge their code changes into a central repository (i.e., software version control), and automated builds and tests are executed. Both CANES and application stakeholders that deploy applications on CANES will drive affordability and improved interoperability by adopting this development practice to minimize integration efforts throughout the development and deployment process and to be able to continually deliver new capability as they become readily available.

Test

Since CANES provides a fully integrated computing infrastructure comprised of hardware and software, the staging environment supports various test environments. Applications at the software level can be tested against the latest version of CANES in a commercial cloud environment, and applications that are tested at a hardware level can test in a lab environment with representative hardware. The key benefit of the staging environment is to provide integration early on and throughout the development of CANES and integrated applications. The combination of a staging environment and automated testing is an indispensable part of quality assurance. Adopting this new test strategy will free up time that used to be spent on manual tests that were too slow to keep up with the rapid development that typically occurs and help teams focus on quality enhancements that until now have been addressed at the tail end of development.

Release/Deploy

Release will be a governance process to be defined based on existing processes and new options now available as automation is introduced. In addition, technologies provided by ACS such as PaaS help simplifies the deployment and management of modern web applications.

Operations/Monitoring

Operations is the last stage in DevOps where CANES is deployed and is operating on ship. As monitoring tools continually advance providing platforms that aggregate data and perform cross analysis, commercial tools are improving the user experience with easy-to-use UIs and system alerts to show correlations between events providing performance degradation or user experience issues quickly back to the developers is the crucial data in shortening the cycle between identified issues and deployment of fixes. When examining the Operations side of DevOps, it's important to highlight that unlike our commercial counterparts, our afloat networks are mobile, global, and engineered to be shot at both kinetically and are prime nation-state cyber targets. Furthermore, our systems are maintained by Sailors with at times limited access to shore support. These attributes drive an even more critical need to get the right system data from the platform and to use every interchange with the afloat platform right.

Commercial Technologies

To envision the CANES/ACS DevOps and tactical Data Analytics platform for the near future, it is instructive to look at some of the current and emerging trends in this space:

- Advanced DevOps/microservices capabilities
- Self-healing and self-protecting systems
- Serverless platforms and function as a service (FaaS)
- ML, especially deep machine learning
- Augmented analytics
- Game theory and ML/DL
- Advanced analytics processors

Advanced DevOps/Microservices Capabilities

As discussed earlier, ACS provides an operational environment that is integral to DevOps. In particular, the PaaS capability provides containers, where applications can build in an environment that encapsulates the necessary software dependencies and enables robust, fault-tolerant remote installations. While we have yet to fully exploit all of the benefits, the capability exists today (Farcic, 2017) to perform zero-downtime software upgrades, immediate rollback to prior versions if problems occur, and phased rollouts ("blue/green deployments"). Of particular interest to data analytics is the ability to perform "A/B testing." In this case, two (or more) different versions of a data analytic could run simultaneously and, based on the results, the more successful version would be retained. Blue/green and A/B deployments allow new analytics to be pushed to a ship and tested, evolving the analytics on a tactical platform without disrupting current missions.

Self-Healing and Self-Protecting Systems

The PaaS + Microservices architecture supports self-healing software, a design pattern for high availability systems that is actually fairly common in distributed data analytics platforms such as Hadoop. These self-healing systems are as follows (Bonér et al., 2014):

- Reactive—maintaining rapid and consistent response times
- Resilient to failure—through replication, containerization, isolation, and delegation
- Elastic—rapidly scaling under varying workloads
- Asynchronous—implementing loose coupling between components, so failures are isolated

By using ML algorithms for anomaly detection and classification, it will soon be possible to build systems that are self-protecting at the architecture, application, and OS levels (Yuan, 2016). This adaptive, intelligent defense, coupled with traditional security protections, will be very important in improving our defenses against the growing frequency, complexity, and sophistication of cyberattacks. In fact, the trend is toward self-managing, self-healing, self-configuring, and self-protecting software systems enabled by ML.

Serverless Computing/Function as a Service

Serverless computing allows highly efficient sharing and utilization of compute resources (Roberts, 2016). The term serverless is generally acknowledged as a misnomer—there are definitely still servers in the system—but the details of the infrastructure, servers, operating systems, and runtime environments have been abstracted away from the application developer, so they need not be concerned with the implementation or management of the underlying system. The serverless platform provider manages all the details of the environment and the dynamic allocation of compute resources. All the application provider does is deliver the code. This code is an event-driven, functional program (as opposed to object-oriented or procedural code), hence the alternative term *function as a service* (FaaS). Amazon Web Services Lambda, Azure Functions, and Google Cloud Functions are all examples of cloud-based FaaS. Apache OpenWhisk, originally developed by IBM, is an open source serverless platform that can be deployed on-premises.

In FaaS, an event (such as incoming data) triggers execution of the function, which can scale horizontally with varying load. The code is automatically containerized, deployed, scaled, and executed. The serverless architecture has built-in high availability and fault tolerance, and allows for very efficient utilization of compute resources. This resource efficiency is very appealing in the fixed-footprint compute environment of a tactical platform. Because FaaS compute is event driven and transient, there is a reduced attack surface and it is more resilient.

Not all applications are suited for FaaS. The best applications are for short computations, triggered by very short bursts of activity. Long-running processes, or ones with consistent loads, do not benefit much from FaaS. Nevertheless, FaaS is excellent in periods of rapid activity and is highly optimized for performance in such situations.

Deep Machine Learning

It is now commonplace to encounter ML applications in our day-to-day lives. Digital assistants such as Alexa, Siri, and Cortana are generally able to respond to questions and voice commands with reasonable accuracy. ML speech-to-text, natural language processing, and text-to-speech algorithms are all employed by these digital assistants. Photo library software on our personal computers can do facial recognition. Using clustering algorithms, recommender systems from Netflix or Amazon suggest what to watch or buy. DL is a subset of ML that brings us closer to AI. Whereas ML can perform natural language processing (NLP), finding common words, n-grams, and phrases, DL is more capable of natural language understanding (NLU). NLU not only identifies common words and phrases, it can analyze sentences and groups of sentences to discern topics and context and thus it comes closer to understanding conversation. Furthermore, speech-to-text ML can couple with audio ML analysis (e.g., speech-to-emotion) or image recognition (body language), providing further understanding beyond word recognition. Even tougher NLU problems such as identifying sarcasm and irony, which are difficult with purely speech-to-text analysis, are amenable to these hybrid analyses. It is feasible to use these technologies to construct a model of the social hierarchy of a group (think of a ship's Flag Command Center or a boardroom in a business). The implications of these technologies to assist decision-makers during periods of stress, for example, to determine if group dynamics are negatively impacting the decision-making process are tremendous. Similar to how better cockpit resource management made aviation safer, perhaps in the above scenario, these technologies could help decision-makers in real time make better strategic decision by better understanding the human teaming in the room.

Augmented Analytics

In the SSC Pacific vision video, when "Kennedy" suggested a COA not considered by the staff, this represented Augmented Analytics (AA; Su, 2017). AA uses ML, NLP, and other tools to perform data source selection, preprocessing (cleaning), analysis, insight generation, and presentation. Indeed, behind the scenes, Kennedy was presumably continuously selecting and analyzing data to generate actionable insights, and then communicated those insights to the ship's staff without the benefit of a human data scientist in the decision loop. No one commanded Kennedy to conduct an analysis or build a DL model, but rather it apparently was simply churning out insights on its own, waiting for an applicable situation to arise.

Today, AA maturity as a near-term capability is unlikely. If performed as an exhaustive survey of a vast parameter space, AA is very resource intensive, roughly akin to categorizing everything in a haystack, finding multiple "non-hay" objects in the haystack, recognizing one of these non-hay objects is a needle, and generating the insight that this sharp object can be used to puncture a balloon, and that having the insight that puncturing a balloon would be a useful tactical objective. Even so, AA has near-term potential in streamlining the tedious parts of analytics: data collection, cleaning, labeling, classification, and preliminary analysis, setting the stage for people to make focused inquiries and interpreting and generating insights from the results.

Game Theory and ML

The discussion of AA brings us to the intersection of Game Theory (GT) and ML/DL (Perez, 2016). If we (oversimplify) GT as modeling cooperation and competition (decision-making) strategies based on deterministic game rules, we can see the role of DL as discovering rules based on imperfect knowledge of the outcome of games. What does that mean? To better understand this, consider the use of adversarial neural networks (ANNs), one of many possible approaches in DL. With ANNs, we pit one neural network against another in a game. The classic example in image recognition is that we have one neural network (NN) trying to distinguish between real and fake images. At the same time, a second NN is trying to generate convincing fake images. As time goes on, the first NN is trained to distinguish real versus fake images, while the second NN learns how to make better counterfeits. In the end, both NNs "learn" what makes an image "real" and what makes it "fake." The result of the game is that both NNs learn from imperfect data to become better players.

These technologies have potential to apply to tactical missions as they mature to a broader range of human machine interaction. For example, developing a DL algorithm about how to avoid detection of a ship's electromagnetic signature could be modelled using ANNs under various environmental/weather conditions. These are the type of "games" Kennedy could play against itself to improve its algorithms. In Figure 4, we present a simplified NN that ingests data from various domains into the hidden layer and ultimately presents a probability of success of a specific mission course of action (COA).

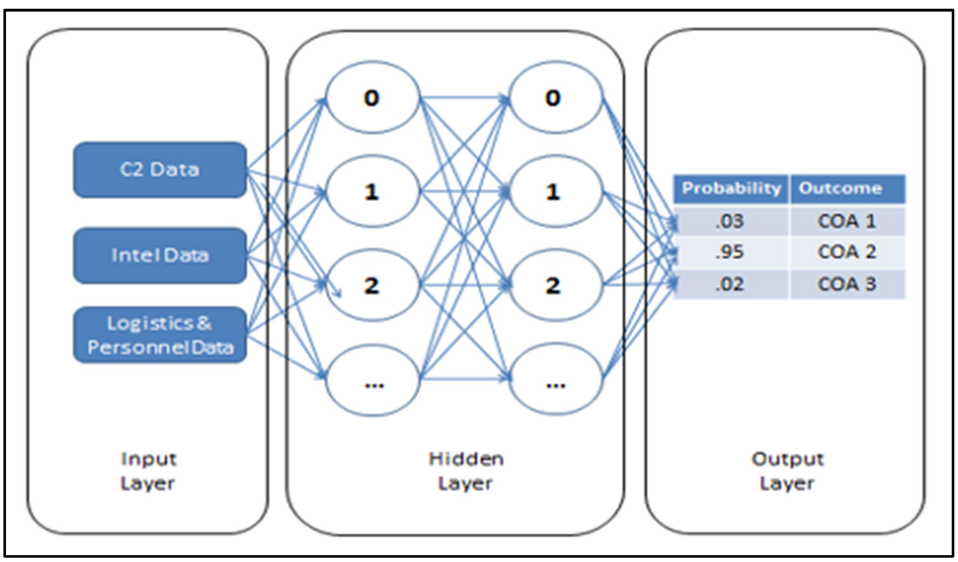

Figure 4. Simplified Example of Neural Network in a Maritime Tactical Network

Advanced Analytics Processors

Finally, all of these data analytics make high demands on the compute resources on ship. Part of the problem is that generic (e.g., x86) processors are optimized for general purpose computing and not data analytics. Graphical Processing Units (GPUs) and Field-Programmable Gate Arrays (FPGAs) are better suited for analytics processing, but Application-Specific Integrated Circuits (ASICs), purpose-built for NN (ML/DL) computations, such as Google Tensor Processing Unit, have the potential to greatly accelerate DL calculations. These have a much smaller footprint and lower power consumption (relative to NN compute capacity), so they are enablers of more data analytics capacity in a constrained footprint, like an afloat platform. If we expect our systems (Kennedy) to be performing AA and ANN behind the scenes, we will need to rely on purpose-built hardware to achieve the necessary throughput and allow the ship to remain tactically relevant even if disconnected from the shore infrastructure.

Overall, the near-term Data Analytics reference architecture will exploit DevOps capabilities to rapidly deploy new capabilities to the tactical edge. Looking toward the future, advances in technology will enable the platform to be self-healing and self-protecting, with highly efficient use of computing resources, including purpose-built processors. Potentially, we will leverage ML and DL for machine-human interaction and rely on AA to prepare and stage data for our analytics. Those analytics will be able to train themselves using game theoretic and DL technologies, and be ready to facilitate insights that will give our warfighters a critical advantage. As we engineer the CANES and ACS stack afloat to operate afloat in what is a very harsh maritime battlespace with intermittent and limited connectivity, we will need to examine how we integrate these emerging technologies afloat and allow them to be survivable and maintained by our Sailors afloat.

Fleet Employment and Management of AI, DL, and ML Technologies

In the 2018 National Defense Strategy, rapid technological change and challenges from our adversaries in every operating domain are identified as causes of our complex security environment. The strategic approach identifies several steps to build a more lethal force, including modernizing key capabilities in the C4ISR mission. The secretary of Defense states that investments will be prioritized for resilient and federated networks to assist in gaining information. AI, DL, and ML can support the objectives in the defense strategy as these technologies have the potential to give Navy warfighters the advantage in combat operations if they are deployed with the reliability and security that is needed for fleet missions.

In the book *The Master Algorithm*, Pedro Domingos (2018) discusses the increasing growth of ML technologies across a wide spectrum of activities. Domingos highlights the initial use in the 1980s of early ML in the financial sector, and the migration toward other commercial uses in the subsequent decades. Additionally, Domingos discusses the rapid rise and early employment of ML tools following the attacks of 9/11 by the DoD and other agencies, and how ML is not a single set of technologies or algorithms, but rather a multi-discipline body of knowledge or “tribes.” These “five tribes” that Domingos explores each approach the desired outcomes of ML from varied theoretical and applied approaches, suggesting that the state of practice today requires different technical approaches or in many cases a blending of technologies, based on the questions being asked. Hence, as we look toward ML solutions for the Navy, we will likely need to field a collection of capabilities that are optimized to our unique connectivity and maritime concept of operations.

The CSG organization provides an excellent use case of where these technologies can be applied to support Fleet operations. In the CSG organizational structure, the Battle

Group Commander's staff executes a series of functions required for the mission within each coded directorates. On the CSG flagship, the CVN or Nuclear Powered Aircraft Carrier, staff functions are supported by departments that exist within the shipboard organization. The Operations Department executes and supports current tasking and future missions while maintaining schedule. The Combat Systems Department is responsible for maintaining weapons and communications systems and ensuring the enterprise network is available. The Engineering and Reactor Departments are charged with ensuring safe and sustained propulsion, as well as maintaining all hull, machinery, and electrical equipment. The Air Department enables flight operations to occur hazard free and within environmental constraints of any geographical area. The Supply Department supports logistical needs in parts, materials, and consumable equipment. The Navigation Department is responsible for ensuring safe passage during open ocean transit, as well as in constrained and heavily trafficked waterways. The arrangement of departments onboard a CVN is similar to what may be found on smaller units, such as cruisers and destroyers—with the roles of each department differing slightly based on the specific mission of the ship. Weaving the activities of these departments, along with the Aircraft Carrier's mission to deliver air power via an embarked airwing and other ships in the strike group, provides a compelling use case for the potential capabilities ML can provide.

In examining each department's specific areas of responsibility, it is evident that there are myriad opportunities for AI, DL, and ML to contribute to the naval mission. Examining the use cases, we see applications across Command and Control, Intelligence, Cyber, Logistics, and Personnel, to use the data in these domains to improve decision making. For example, one subset within the Operations Department is Operations Information (OI). The Operations Specialists in OI Division comprise the watchteam that maintains a current operational picture (COP) that provides the mission commander situational awareness to make decisions for the entire CSG. Key assets in the mission are identified and tracked with tools that present the commander with a comprehensive view of the battlespace. Contacts on the display are identified as friendly, neutral, or foe. Since the environment is constantly changing, this mission area is one that has the potential to benefit from AI technologies and a DevOps framework to rapidly integrate and process new sensors, data feeds, and contingences.

The Navy has laid much of the groundwork associated with implementing AI technologies by employing data models and software patterns that leverage the commercial technologies that are many of the underpinning connectivity methods the newer AI technologies employ such as eXtensible markup language (XML) and other web-based connectivity tools (Rothenhaus IEEE, 304). To manage a growing demand for faster decision cycles and improved battlespace awareness, supporting systems such as Global Command and Control System Maritime (GCCS-M), as well as intelligence systems, the Navy can ensure our warfighters are not missing key elements of the operational picture due to over-complexity and human error by implementing intelligent systems able to process data automatically and present that correlation of data to Sailors and decision-makers. In the "Kennedy" vignette discussed earlier, tools are dynamic enough to adapt to emerging conditions, giving the warfighter the advantage of not having to focus on updating changing information, and instead allowing warfighters to focus precious attention on strategies to support mission success.

Further Research

Although AI is not a new field of study, new methods and approaches enabled by increased computing, storage, and data management techniques are driving renewed interest in the commercial sector and are fertile areas of future research in the Navy enterprise. The adoption of those technologies for the Navy presents additional areas for multi-disciplinary research in the areas of computer science, system of systems engineering, and acquisition with likely many more potential areas of study. Areas of study include questions about how to characterize the data the Fleet manages from the perspective of volume and variety. Many of the technologies highlighted in this paper rely on significant quantities of data to gain the insights they deliver. Although ships have large databases, they may not have the types of enterprise-level data that would permit meaningful analysis. Additionally, our ships operate in a unique connectivity construct ranging from full wideband connectivity to being completely disconnected. Future research in the areas of data strategies to support sharing of processing between on-board and off-board processing with a focus of graceful scaling from very elastic processing environments ashore to more limited processing on shipboard data centers could help future engineers design systems to support the Fleet's unique requirements.

Conclusion

Navy warfare increasingly depends on advanced software capabilities, deployed across numerous platforms and systems. As the amount of information grows with the increase in quality and quantity of sensors and new autonomous platforms, the Navy is looking to leverage commercial information technologies to improve warfighting outcomes and enhance mission execution. Leveraging AI technologies, Navy Program Offices are examining methods to provide a ready platform to rapidly and affordably integrate and test emerging AI capabilities. As we integrate these types of technologies into an already complex system of systems afloat, it is critical to manage the complexity to ensure our Sailors can employ and support them. In the areas of command and control, intelligence, and logistics, AI and ML have the potential to deliver to the nation that integrates them first a tactical advantage. For the Navy, CANES and ACS are the target platform on which to integrate those technologies to deliver important warfighting capabilities across Command and Control, Intelligence, and Logistics software applications.

References

A Summary of the 2018 National Defense Strategy of the United States of America. (2018). Retrieved from https://www.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf

Bonér, J., Farley, D., Kuhn, R., & Thompson, M. (2014). The reactive manifesto. Retrieved from https://www.reactivemanifesto.org/

Brynjolfsson, E., & McAfee, A. (2016). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies.* New York, NY: W.W. Norton.

Buduma, N., & Locasio, N. (2017). *Fundamentals of deep learning: Designing next-generation machine intelligence algorithms*. Sebastapol, CA: O'Reilly Books.

Domingos, P. (2018). *The master algorithm: How the quest for the ultimate learning machine will remake our world*. New York, NY: Basic Books.

Farcic, V. (2017). *The DevOps 2.0 toolkit: Automating the continuous deployment pipeline with containerized microservices*. Retrieved from http://www.amazon.com

Perez, C. (2016). Game theory reveals the future of deep learning. Retrieved from https://medium.com/intuitionmachine/game-theory-maps-the-future-of-deep-learning-21e193b0e33a

Roberts, M. (2016). Serverless architectures. Retrieved from https://martinfowler.com/articles/serverless.html

Rothenhaus, K., Bonwit, B., Galdorisi G., & Stang, A. (2018). Distributed lethality, command and control software engineering, and naval laboratories. *Naval Engineers Journal* (5).

Rothenhaus K., Michael, B., & Shing, M. (2008). Architectural patterns and auto-fusion process for automated multi-sensor fusion in SOA system of systems. *IEEE Systems Journal.*

Su, B. (2017). Augmented analytics demystified: What it means and why it is the future of data analytics. Retrieved from https://medium.com/analytics-for-humans/augmented-analytics-demystified-326e227ef68f

Yuan, E. (2016). *Architecture-based self-protecting software systems* (Doctoral dissertation, George Mason University). Retrieved from https://seal.ics.uci.edu/dissertations/eric.pdf

Disclaimer & Distribution Statement

ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL
555 DYER ROAD, INGERSOLL HALL
MONTEREY, CA 93943

www.acquisitionresearch.net