



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2022

Tactical ISR/C2 Integration with AI/ML Augmentation

Maule, Randy W.

Monterey, California: Naval Postgraduate School

<https://hdl.handle.net/10945/71895>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>



NAVAL RESEARCH PROGRAM

NAVAL POSTGRADUATE SCHOOL



TACTICAL ISR/C2 INTEGRATION WITH AI/ML AUGMENTATION

NPS-22-N215-A

Dr. Randy W. Maule
Department of Information Sciences

December 2022

Prepared for:
Commander, Naval Surface Forces (CNSF)

Distribution Statement E: Distribution authorized to DoD components only (Operational Use)
(December 2022). Other requests for this document to President, Code 261,
Naval Postgraduate School, Monterey, CA 93943-5000.



Research Topic



- NAVPLAN 2021 and 2022 specify Distributed Maritime Operations (DMO) with a tactical grid to connect distributed nodes for processing at the tactical edge with Artificial Intelligence/Machine Learning (AI/ML) to support:
 - Expeditionary Advanced Base Operations (EABO)
 - Littoral Operations in a Contested Environment (LOCE)
 - Joint All-Domain Command and Control (JADC2)

- Intelligence, Surveillance and Reconnaissance (ISR) and Command and Control (C2) hardware and software have yet to be fully integrated and configurations tested.



Objectives



- Evaluate options for ISR and C2 integration into a universal Common Operational Picture (COP)
 - Hardware Infrastructure: Tactical cloud hardware and deployment options
 - Software Infrastructure: Tactical cloud software and deployment options
 - Application Services: C2/ISR integrated solutions sufficient to support a universal COP from HQ to tactical commands to warfighters at the far edge on mobile devices
- AI/ML for decision and automation support



Questions



- Which cloud hardware infrastructure configurations best support tactical operations for mobile warfighters?
- Which distributed hybrid cloud software architectures best support a universal COP in mobile and disconnected operations?
- Can C2 and ISR software be integrated to provide a universal COP on hybrid tactical cloud architecture?
- How can AI/ML be integrated into hybrid cloud and COP operations to enhance decision support?



Method



- Data Requirements / Systems Review: Tactical cloud hardware, hybrid cloud software, C2 and ISR systems
- Request/Obtain Data from Topic Sponsors: Collect resources to support topic sponsor requirements for a universal COP
- Collect and Analyze Data: Evaluate distributed and converged data processing technologies suitable for DMO/EABO tactical edge nodes
- Provide Empirical Evidence: Determine specifications for a universal COP and determine feasibility for tactical cloud deployment
- Evaluate tactical edge software and signal processing options for D-DIL, EMS, GPS, and cyber challenged operations
- Determine sustainment options for self-contained tactical edge equipment, peer and reachback services, with AI/ML decision support.
- Final Report/ Final Presentation



Guidance

Supported Initiatives



Digital Modernization



Electromagnetic Spectrum Operations

- *Resilient, secure, and adaptive tactical*
- *Contested, congested, and operationally limited EMS environment*

Hyper-Converged Infrastructure (HCI)

- *Tightly-integrated compute, storage, networking, and virtualization*

DoD CIO Priorities

- Cybersecurity
- Artificial Intelligence (AI)
- Cloud
- Command, Control and Communications (C3)



Information Superiority



Department of the Navy Information Superiority Vision



February 2020


Infrastructure – modernize from the current state of fragmented, non-performant, outdated, and indefensible architectures to a unified, logical modern infrastructure capable of delivering information advantage.

Naval Mesh – leverage the existing Naval Tactical Grid (NTG) to create a Naval Mesh Network that extends the DON network and in D-DIL environments can operate cut off from the DON network until connections are reestablished.

Cloud – provide a performant, defensible cloud-enabled network with unified shore and tactical edge processing, storage and networks with identity management across the grid so that tactical edge networks operate as one logical construct.



Maritime Superiority



- Distributed Maritime Operations (DMO)
- Expeditionary Advanced Base Operations (EABO)
- Littoral Operations in a Contested Environment (LOCE)

Accelerate *Ready, Relevant Learning* (RRL).
To retain our competitive advantage.

Instill *continuous learning* behaviors to
broaden and deepen warfighting knowledge

Enable adaptation, improvement, and
strengthen mission command to out-think
and outfight any adversary.



Expeditionary Advanced Based Operations



- We do not collect the data we need systematically, *we lack the processes and technology to make sense of the data we do collect*, and we do not leverage the data we have to identify the decision space in manning, training, and equipping the force.
- We will make strategic investments in *data science, machine learning, and artificial intelligence*. Initial investments will be focused on challenges we are confronting in talent management, predictive maintenance, logistics, intelligence, and training.
- We will explore investments in *decision support tools* that leverage data science and artificial intelligence for the *tactical commander*. Success is defined in terms of finding the smallest, lowest signature options that yield the maximum operational utility .





Tri-Service Maritime Strategy



- The Naval Service will accelerate delivery of the next-generation *Naval Operational Architecture*, composed of the *Naval Tactical Grid*, battle management aids, data structures and infrastructure that underpin distributed operations.
- This network will be fully interoperable with *Joint All-Domain Command and Control* systems and will combine inputs into an actionable common operational picture.
- Leveraging *artificial intelligence* and *machine learning*, we will give our warfighters enhanced situational awareness and facilitate decision making at tactically relevant speeds.



Joint All-Domain Command and Control



Congressional Research Service
 Improving the legislative process since 1949

IN FOCUS

Updated March 18, 2021

Joint All-Domain Command and Control (JADC2)

What is JADC2?
 Joint All-Domain Command and Control (JADC2) is the Department of Defense's (DOD's) concept to connect sensors from all of the military services—Air Force, Army, Marine Corps, Navy, and Space Force—into a single network. Traditionally, each of the military services developed its own tactical network that was incompatible with those of other services (i.e., Army networks were unable to interface with Navy or Air Force networks). DOD officials have argued that future conflicts may require decisions to be made within hours, minutes, or potentially seconds compared with the current military process to analyze the operating environment and issue commands. They have also stated that the Department's existing command and control architecture is insufficient to meet the demands of the National Defense Strategy (NDS). Congress may be interested in the concept because it is being used to develop many high-profile procurement programs.

DOD uses ride-sharing service Uber as an analogy to describe its desired endstate for JADC2. Uber combines two different apps—one for riders and a second for drivers. Using the respective users' positions, the Uber algorithm determines the optimal match based on distance, travel time, and passenger (rating) and other variables. The application then seamlessly provides directions for the driver to follow, delivering the passenger to their destination. Uber relies on cellular and Wi-Fi networks to transmit data to match riders and provide driving instructions.

Figure 1. Visualization of JADC2 Vision

Source: <https://www.seniorconservative.com/defense/174-us-aid-usaf-parties-to-jadc2.html>.

JADC2 envisions providing a cloud-like environment for the Joint force to share intelligence, surveillance, and reconnaissance data, transmitting across many communication networks, to enable faster decisionmaking (see Figure 1). JADC2 intends to enable commanders to make better decisions by collecting data from numerous sensors, processing the data using artificial intelligence

algorithms to identify targets, then recommending the optimal weapon—both kinetic and nonkinetic (e.g., cyber or electronic weapons)—to engage the target.

Some analysts take a more skeptical approach to JADC2. They raise questions about its technical maturity and affordability, and whether it is even possible to field a network that can securely and reliably connect sensors to shooters and support command and control in a lethal, electronic warfare-rich environment. Analysts also ask who would have decisionmaking authority across domains, given that, traditionally, command authorities are delegated to each domain rather than to a human or machine-empowered perspective. Some also question how much a human will be needed for JADC2 to make decisions in real time, and whether it is appropriate to reduce the amount of human involvement in military-related decisions.

Why Change Current C2 Structures?
 The future operating environment articulated by the NDS, the NDS Commission, and other sources describes how potential adversaries have developed sophisticated anti-access/area-denial (A2/AD) capabilities (see Figure 2). These capabilities include electronic warfare, cyber weapons, long-range missiles, and advanced air defenses. U.S. competitors have pursued A2/AD capabilities as a means of countering traditional U.S. military advantages—such as the ability to project power—and improving their ability to win quick, decisive engagements.

Figure 2. A2/AD Environment

Source: <https://www.japcc.org/electronic-warfare-the-future-of-battle/>.

Senior DOD leaders have stated that access to information will be critical in the future operating environment. In addition, these leaders have stated that to challenge potential peer adversaries, a multi-domain approach is required (where U.S. forces would use ground, air, naval,

<https://crsreports.congress.gov>


- Joint All-Domain Command and Control (JADC2) is the Department of Defense's (DOD's) concept to *connect sensors from all of the military services*—Air Force, Army, Marine Corps, Navy, and Space Force—into a single network.
- JADC2 will enable commanders to *make better decisions by collecting data from numerous sensors, processing the data using artificial intelligence algorithms to identify targets*, then recommending the optimal weapon—both kinetic and non-kinetic (e.g., cyber or electronic weapons)—to engage the target.



CNO NAVPLAN



- The Navigation Plan charts the course to execute the *Tri-Service Maritime Strategy* using DMO, LOCE, and EABO for sea and shore-based fires from distributed platforms.
- The *Naval Operational Architecture* provides counter-C5ISR capabilities; weapons of increasing range and speed; and directed-energy systems capable of defeating anti-ship cruise missiles.
- All connected in the *Naval Operational Architecture* (NOA) that integrates with JADC2; the NOA collection of networks, infrastructure, data, and analytic tools will provide decision advantage.

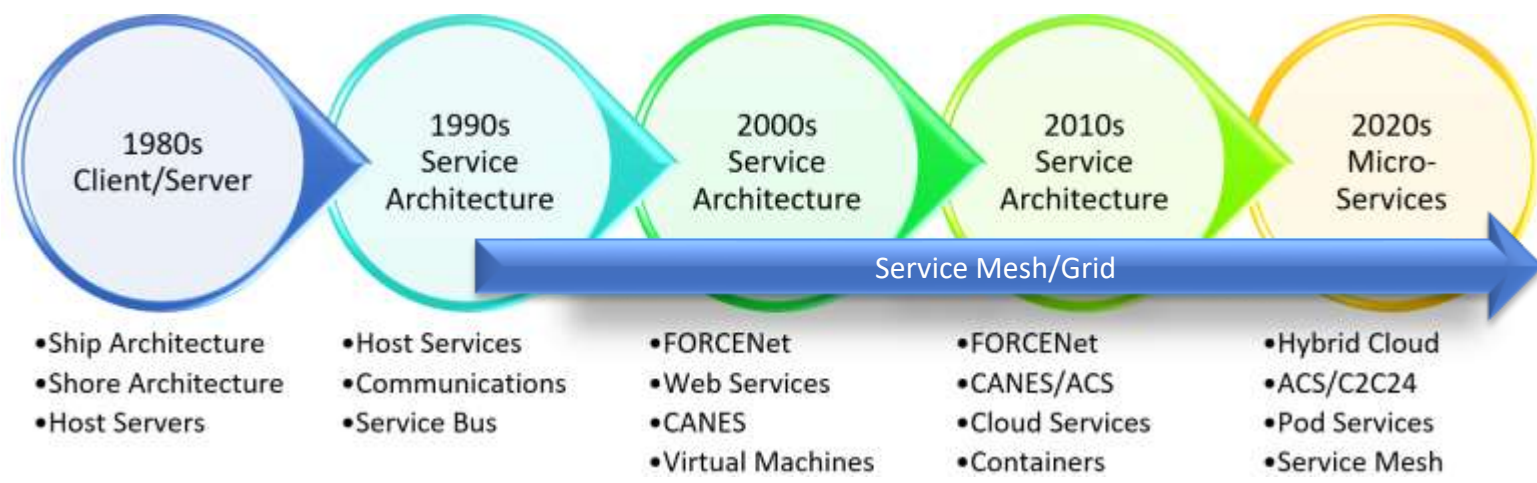




Tactical Infrastructure

HQ – Tactical Command – Far Edge

Tactical Cloud Architecture



AI /ML at the tactical edge (source: IBM)

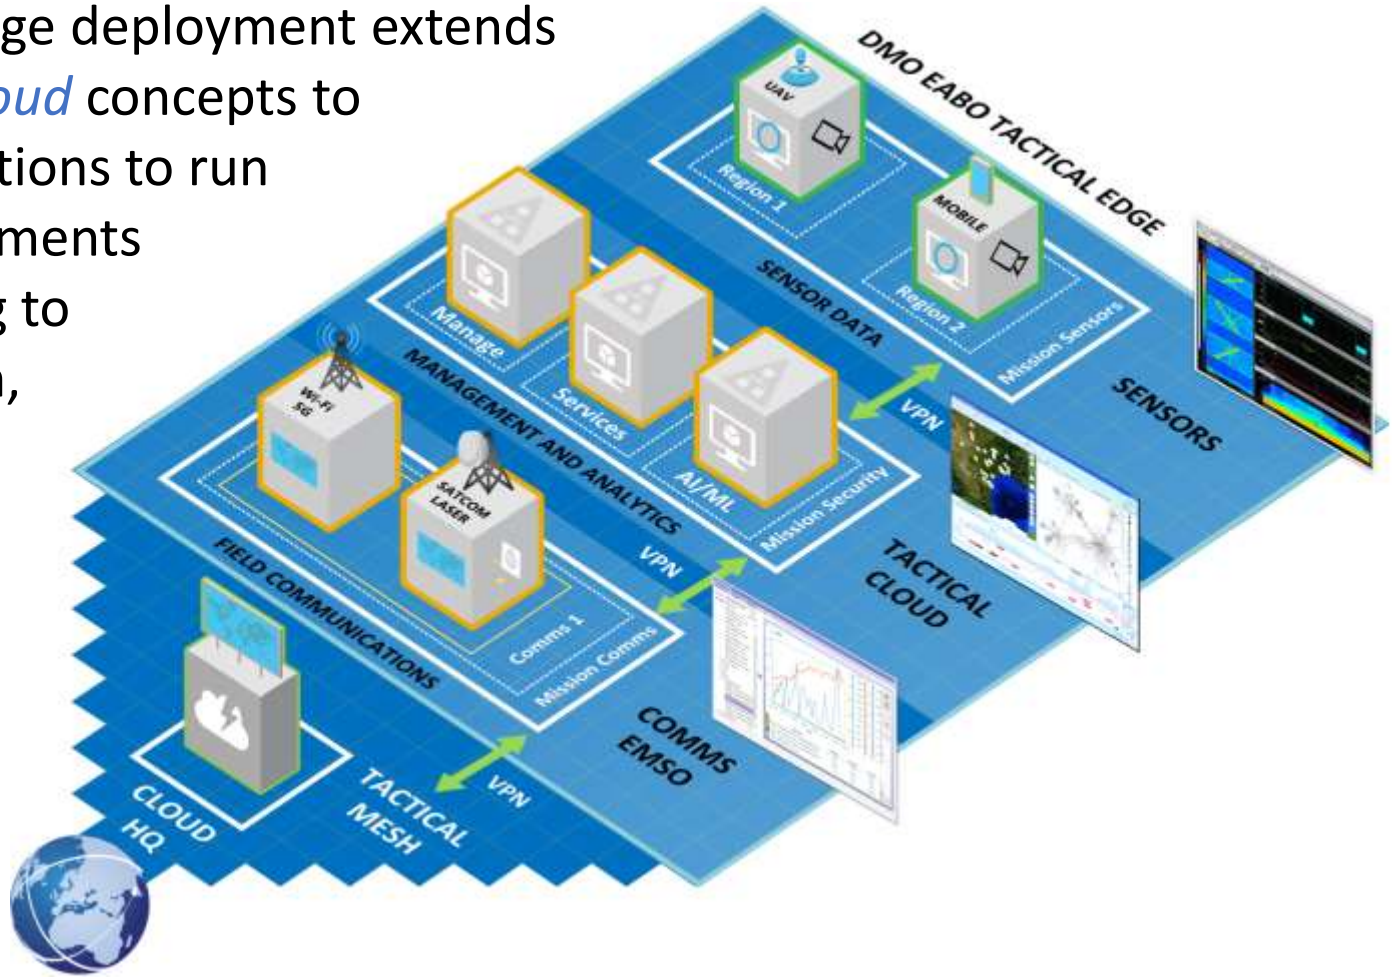
- *Distributed platforms* (ship/shore)
- *Converged hardware/software*
- *Hybrid cloud edge* services
- **DMO tactical grid:**
 - *EABO*
 - *LOCE*



Tactical Edge Node



The Tactical Edge deployment extends *open hybrid cloud* concepts to enable applications to run across environments without having to rebuild, retrain, or maintain separate systems.






Edge Computing



- Consistent deployment model from the core to the edge
- Dynamic and static caching for lower response time
- Automated provisioning, updating, maintenance
- Flexible connectivity and management options
- On-site aggregation and big data analytics
- Higher resiliency and lower costs
- Highly available applications
- Real-time monitoring
- Local data security
- Low latency



Low-latency edge computing (source: HPE)

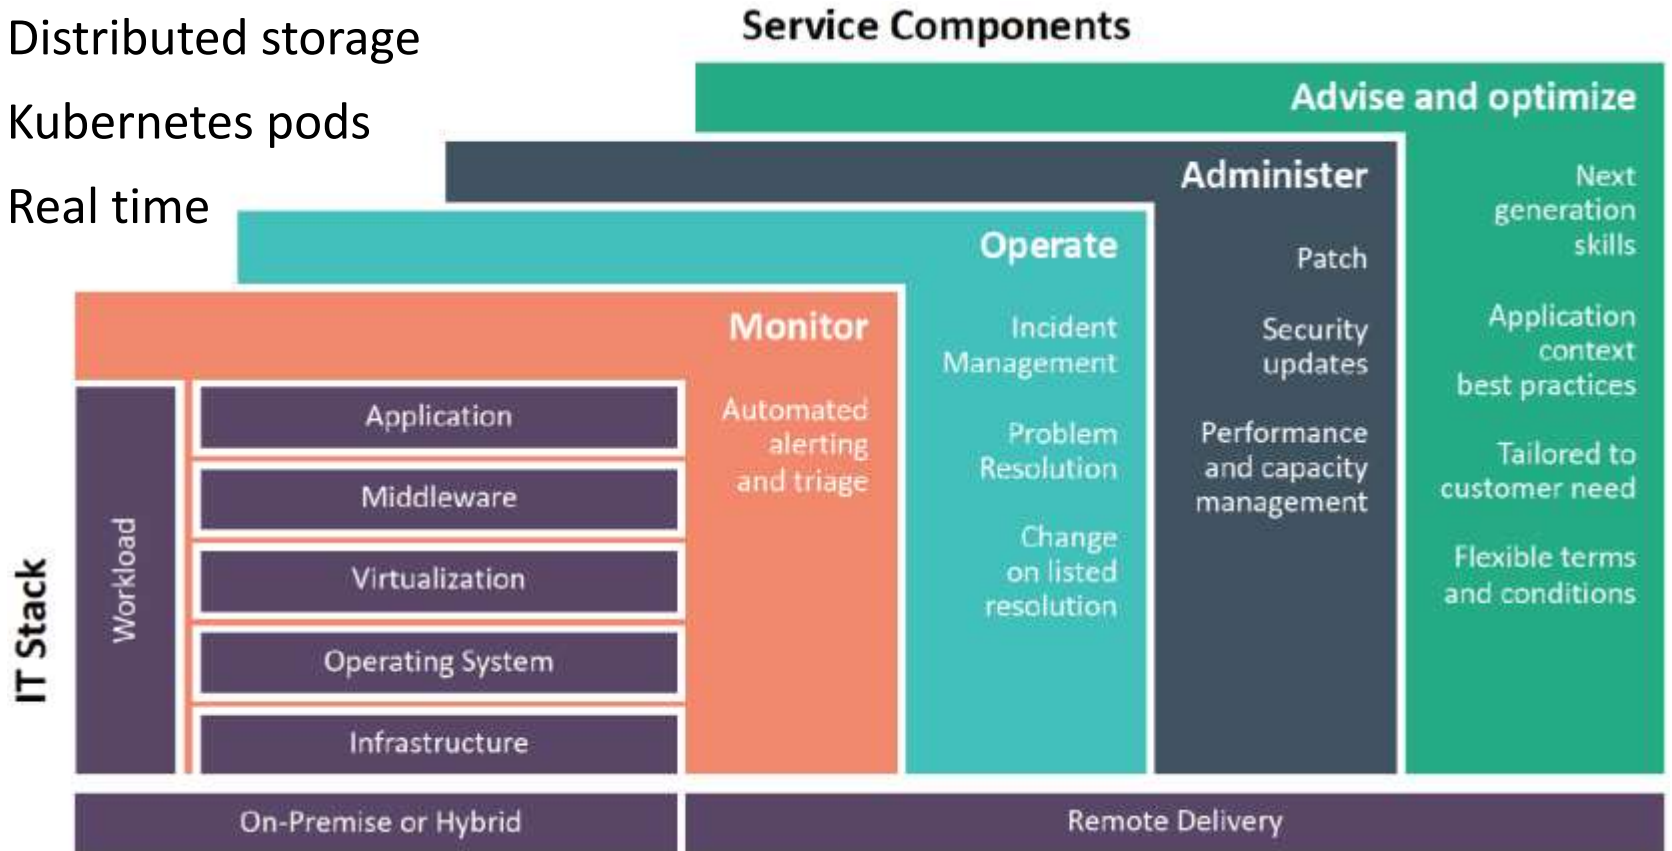


Hyper-Converged Infrastructure



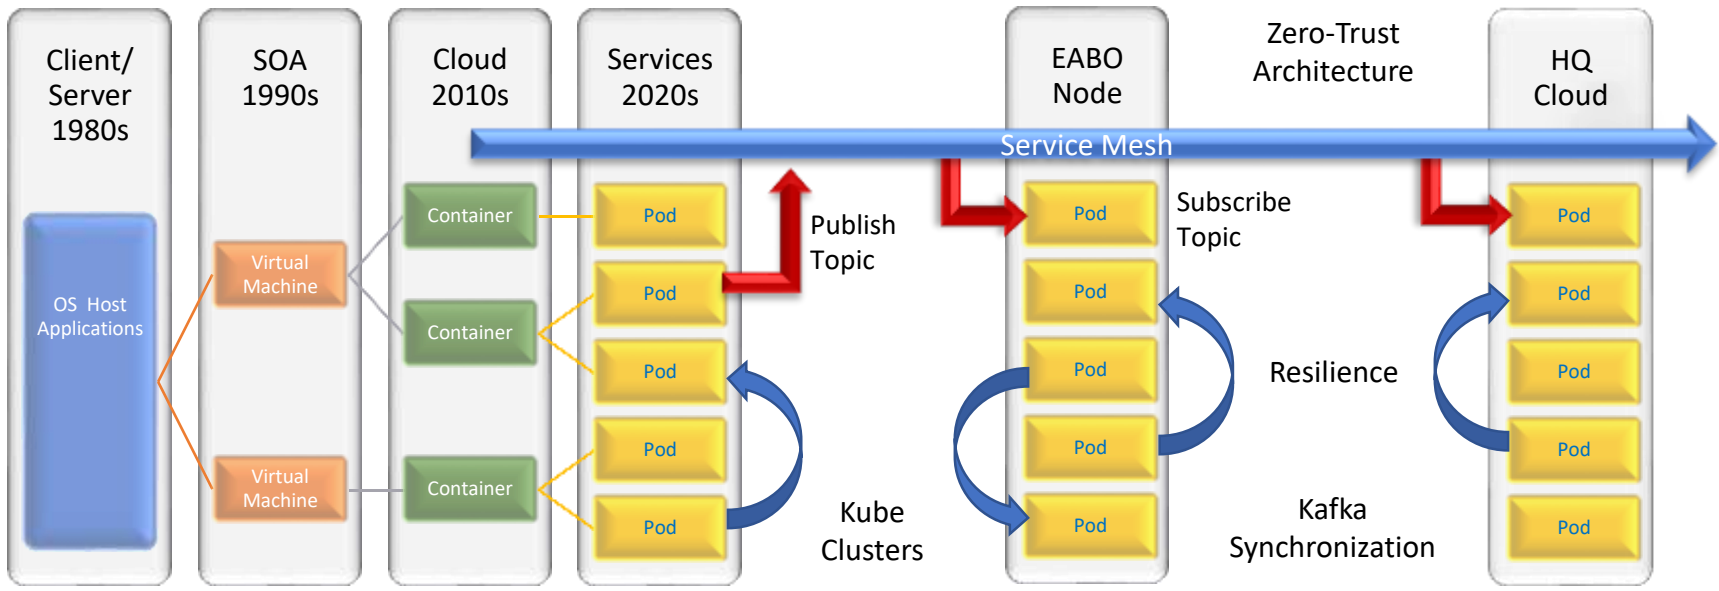
- High availability to operate, regardless of WAN connection state
- Low Space, Weight and Power, Low Cost (SWaP-C)
- Fully autonomous clusters
- Distributed storage
- Kubernetes pods
- Real time

Source: HPE





Service Evolution Zero Trust



- Event-driven services better mirror real-world events to improve SA
- Capabilities can be added to services without reprogramming
- Architecture is better able to manage topology changes, systems failure
- Operations in D-DIL, EMS, and cyber challenged environments



Cloud Software

DMO/EABO Tactical Grid



DMO/EABO Tactical Edge



- Hybrid cloud distributed services and storage, data center capabilities at the far tactical edge
- Tolerant of geographically distributed data sources and high-latency/low-bandwidth interconnects
- Limited physical space, restrictive power, heat generation, vibration and shock, restricted and intermittent connections
- Offline operations with regional nodes on the tactical grid, with central synchronization when communications are available
- Scale pods, nodes and clusters with remote out-of-band management agents to synchronize across tactical and regional/HQ clouds
- Integrated artificial intelligence with machine and deep learning for decision support at the tactical edge



Containers



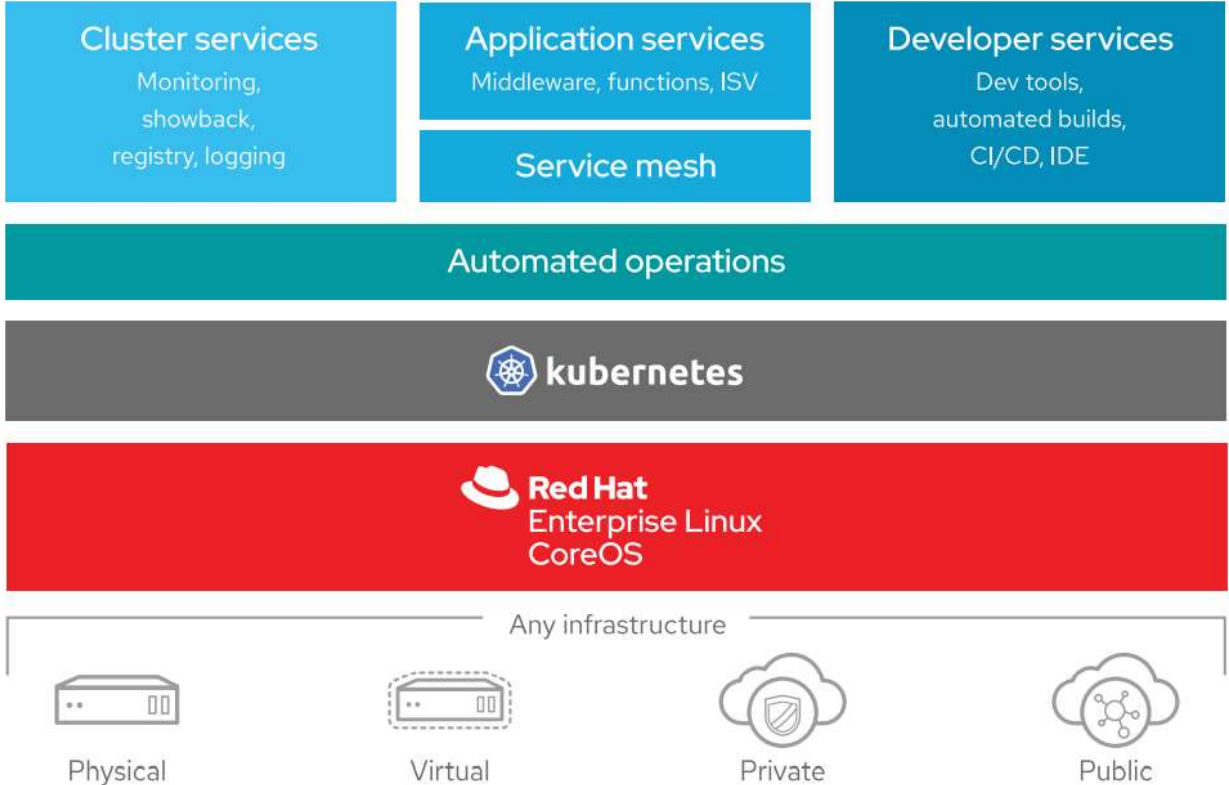
- Extend virtualization to containers to speed app development and deployment, agility, and portability.
- Package and isolate applications that include the entire runtime environment to eliminate physical machines and operating systems.
- Deploy without a kernel with embedded dependencies to enable conflicting software on the same host.
- Eliminate competition for resources such as networking and storage.
- Automatically transition apps between host environments and across systems and geographic areas.
- Micro-services scale to meet demand, scaling only the services, not the entire application.
- Security in the container pipeline to make the containers scalable and trusted.



Orchestration



- Kubernetes de-facto standard for open source container orchestration to automate deployment, scaling, management of containerized applications.
- Containers wrapped into pods with metadata as single deployment entity.
- Controllers determine the number of pods for the workload.
- Devices w/o IPs.
- Secure access.
- Architecture:
 - Micro-services
 - Serverless



Source: IBM/Red Hat

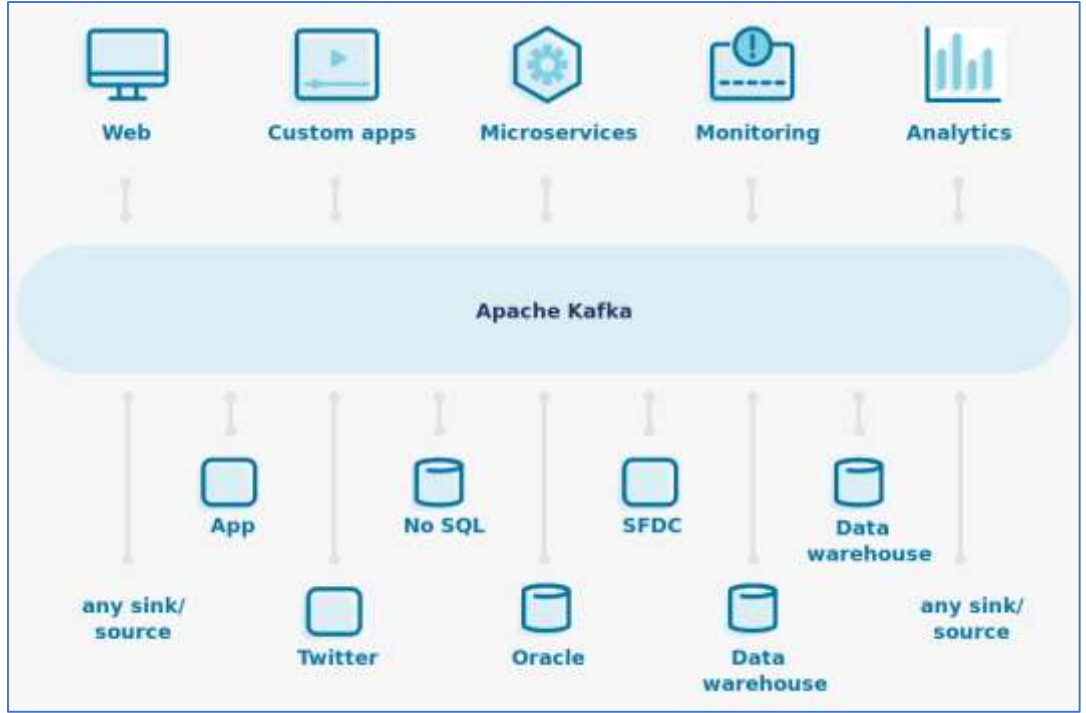


Real Time Streaming Architecture



Apache Kafka:

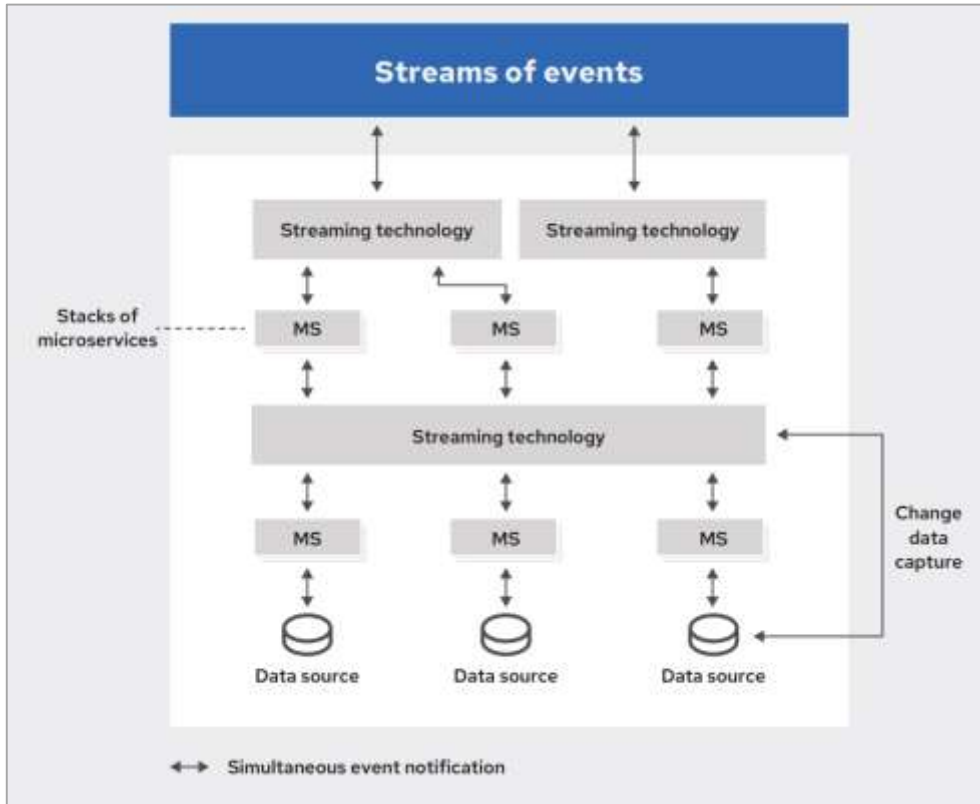
- Distributed system designed for streaming data and media management.
- De-facto communication bus for event-driven and real time architecture.
- Highly resilient, horizontally-scalable, and fault-tolerant.
- High-performance data pipelines, analytics, EDA.
- Topics published as stream of events consumed by subscribers.
- Stream can be consumed within applications and micro-services.
- Optimized for C2/ISR sensor integration



Source: Confluent



Kafka Streams



Source: IBM/Red Hat

Event-Driven Architecture (EDA)

- Asynchronous communication
- Cloud and container development
- Agile, flexible, and scalable
- Distributed microservices
- Stateless microservices
- Loosely coupled
- Sender/publisher objects
- Receiver/subscriber objects
- Users receive events simultaneously
- Low latency, high throughput
- Real time event reaction
- Improved situational awareness

- Kafka EDA is optimized for data in motion in real-time streams. Streams can be captured and replayed, or transformed into new streams and published to subscribers.
- Non-blocking communication releases resources without a response.



Test Cases

Microsoft Azure Stack Hub/HCI

IBM/Red Hat OpenShift

- Microsoft makes two versions of Azure Stack. Both are premise, hybrid cloud versions of Azure:
 1. *Azure Stack Hub* is a private, autonomous cloud that provides connected or disconnected cloud-native apps for Azure services in premise deployments.
 2. *Azure Stack HCI* is a virtualization host that uses a hybrid solution that integrates with Azure public cloud to provide scalable virtualization and storage for high-performance workloads in edge deployments. Does not currently support disconnected operations.

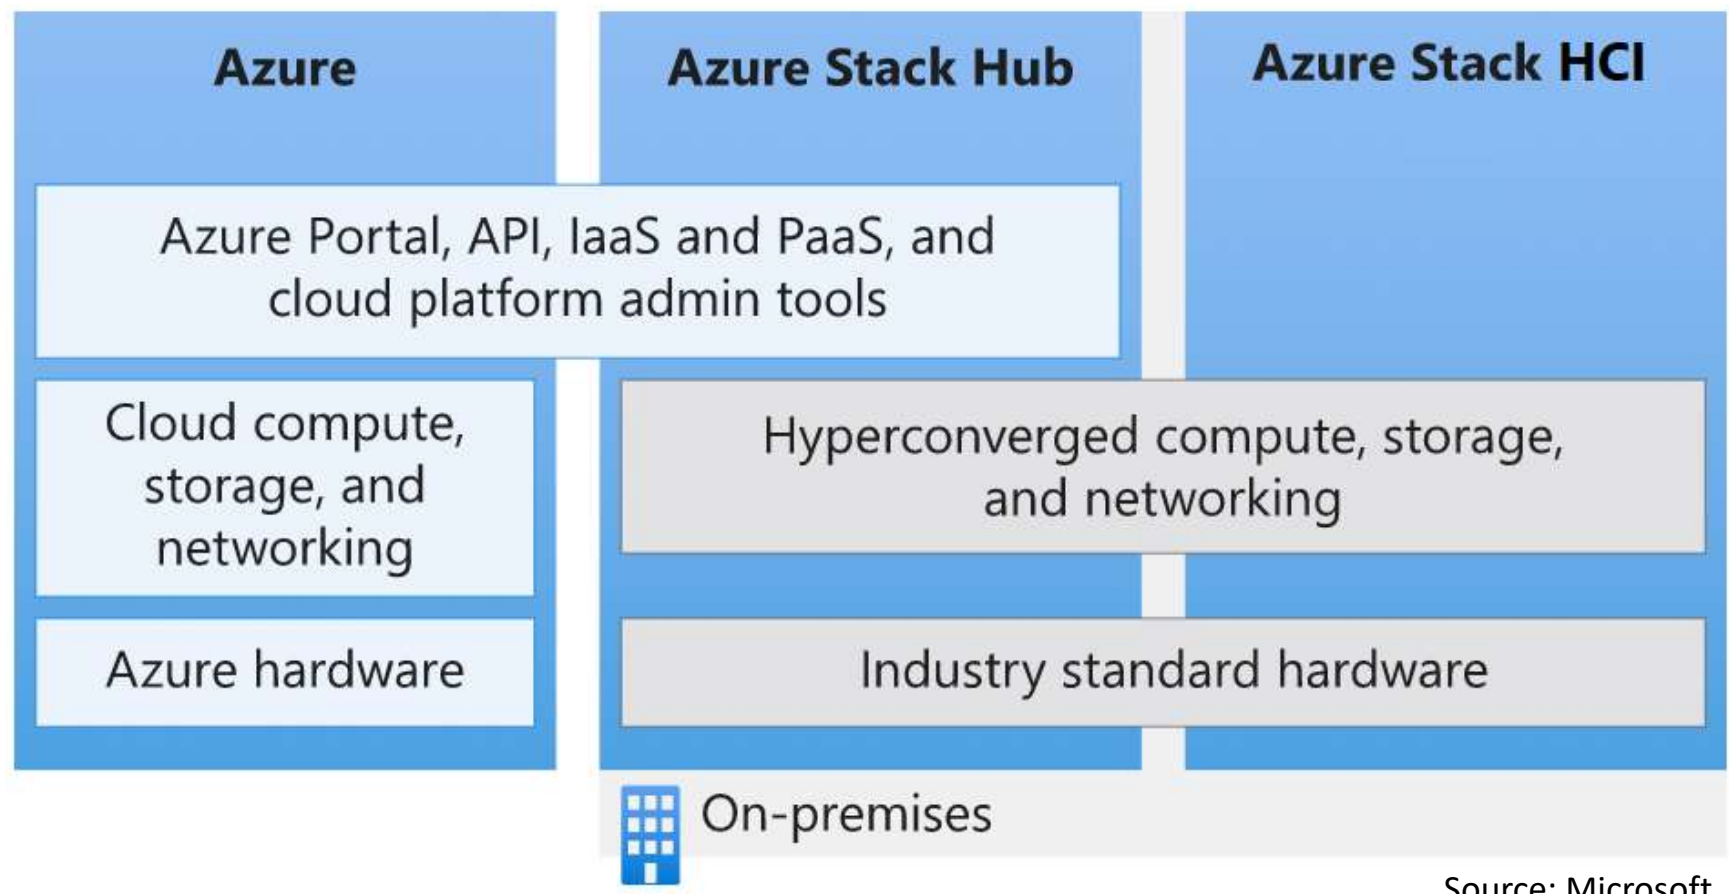
Microsoft makes an Azure edge appliance that is a subscription service from Azure that was not tested for this project since it cannot be independently deployed.



Source: Microsoft



Azure Stack Comparison



Source: Microsoft



Azure Stack Comparison



Characteristic	Azure Stack Hub	Azure Stack HCI
Number of nodes	4-16	2-16
Hardware	OEM	OEM
Support disconnected scenarios	Yes	No
Modernize aging storage	No	Yes
Cloud billing for on-prem data workloads	Yes	Yes
Provide Azure Consistent IaaS and PaaS	Yes	No
Build modern apps across cloud and on-premises using Azure services	Yes	No
Small-footprint branch office scenarios	No	Yes
Ruggedized form-factors in harsh or remote environments	No	No
Support for repurposed hardware	No	Yes
Trusted enterprise virtualization	No	Yes
High availability for virtual machines	Yes	Yes
Built-in disaster recovery capabilities	No	Yes

Azure Stack Hub allows a restricted set of administrative tasks via well-defined, constrained interfaces but is able to run disconnected from the network and Azure cloud.

Azure Stack HCI provides full, direct access to the underlying hardware and the operating system running on cluster nodes but cannot run disconnected from Azure cloud.





Azure Stack Hub



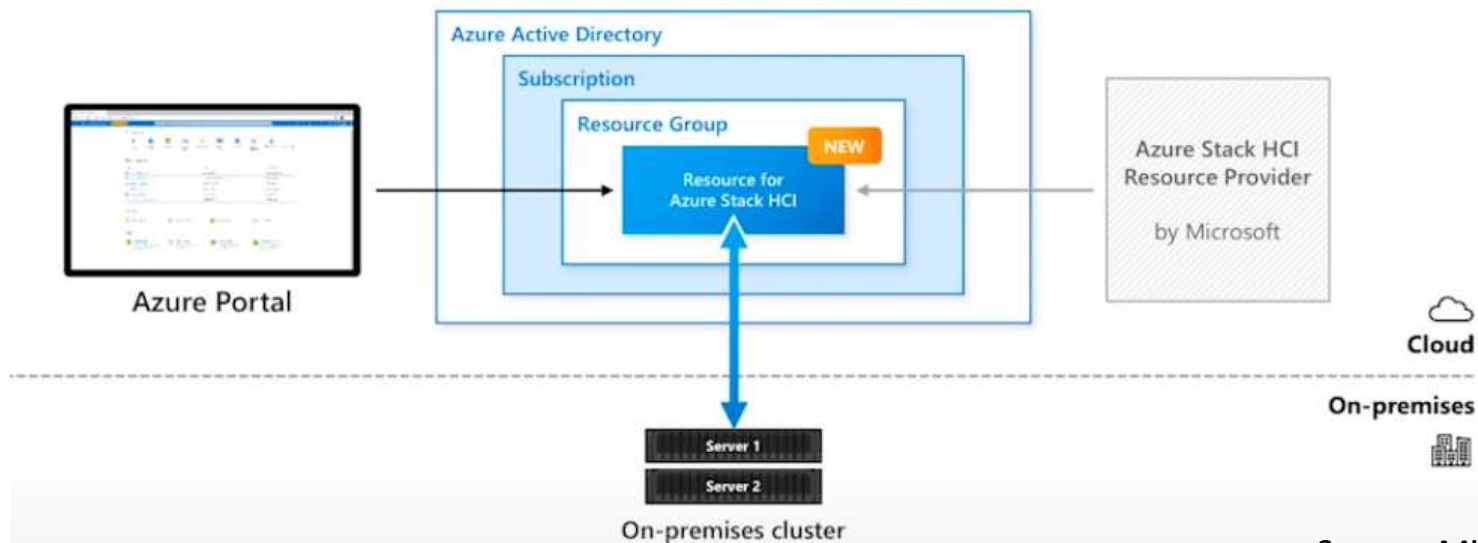
- Azure Stack Hub is purchased pre-configured from an industry partner (e.g., IBM, Dell, HPE)
- Azure Stack Hub is an extension of Azure for on premise cloud computing.
- Azure Stack Hub can provide Azure services either connected to the internet (and Azure) or in disconnected environments with no internet connectivity.
- Azure Stack Hub uses the same underlying technologies as public Azure, which includes Infrastructure-as-a-Service (IaaS), Software-as-a-Service (SaaS), and optional Platform-as-a-Service (PaaS) capabilities.
- Azure Stack Hub operators can offer PaaS services to users including:
 - Service Fabric
 - Kubernetes Container Service
 - Ethereum Blockchain
 - Cloud Foundry



Azure Stack HCI



- Azure Stack Hyper-Converged infrastructure (HCI) operating system is delivered as an Azure service.
- Azure Stack HCI can be deployed and run Windows and Linux virtual machines (VMs) in premise datacenters or at the tactical edge.
- Azure Stack HCI can integrate back to Azure public cloud for backup, monitoring, and to use Azure Security Center.
- Azure Stack HCI support for disconnected operations TBD.



Source: Microsoft



OpenShift



- OpenShift is a platform for developing and running containerized applications that can scale from a few machines and applications to thousands.
- OpenShift Container Platform (OCP) incorporates a Kubernetes foundation to extend containerized applications from a single cloud to multi-cloud environments.
- Our default cluster will consist of four bare metal machines: a bastion node and 3 cluster nodes. Virtualization is available through the Kubernetes API.
- The bastion node will host all of the infrastructure and services required for the cluster to operate using a layer 2 switch to connect all four nodes.
- For added security, the bastion node is the only node to connect to the network and all communications are through the bastion.
- Each node has an out-of-band management interface on the same layer 2 network.
- The bastion node runs Red Hat Enterprise Linux and the nodes run Red Hat CoreOS (RHCOS) that includes:
 - The CRI-O Kubernetes native container runtime that integrates with the OS for running, stopping, and restarting containers.
 - The Kubelet node agent for Kubernetes for launching and monitoring containers.



OpenShift Hybrid Multi-Cloud



THE FORRESTER WAVE™
Multicloud Container Development Platforms
Q3 2020



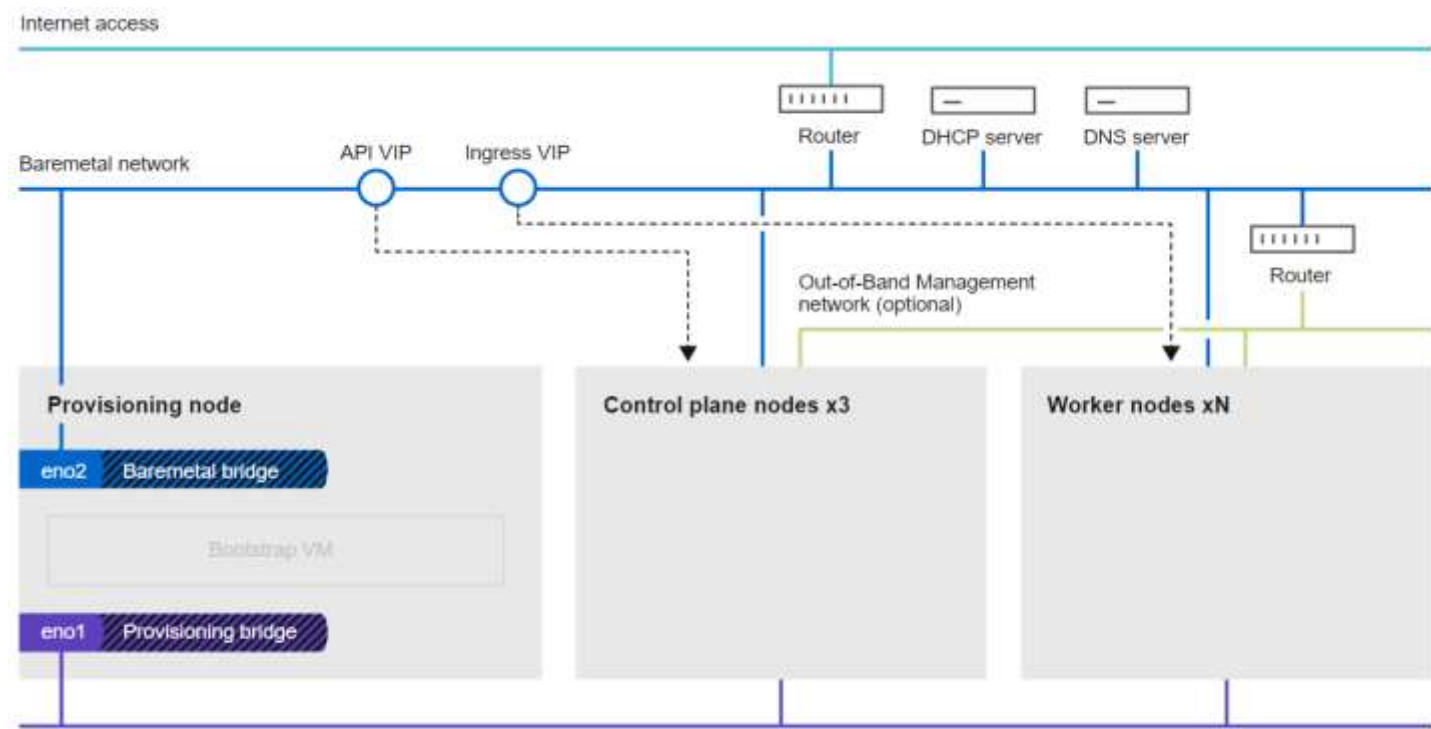
- Most widely deployed multi-cloud container platform
- Integrated development and unified operations across public and on-premises platforms
- Build once, deploy anywhere
- Micro-services application development
- Seamless integration with public and private cloud services.



OpenShift Control Plane



- The control plane is composed of the master machines that manage the cluster and workloads on the compute/worker machines.
- The cluster manages upgrades to the machines by the actions of the cluster version, machine configuration, and individual operators.



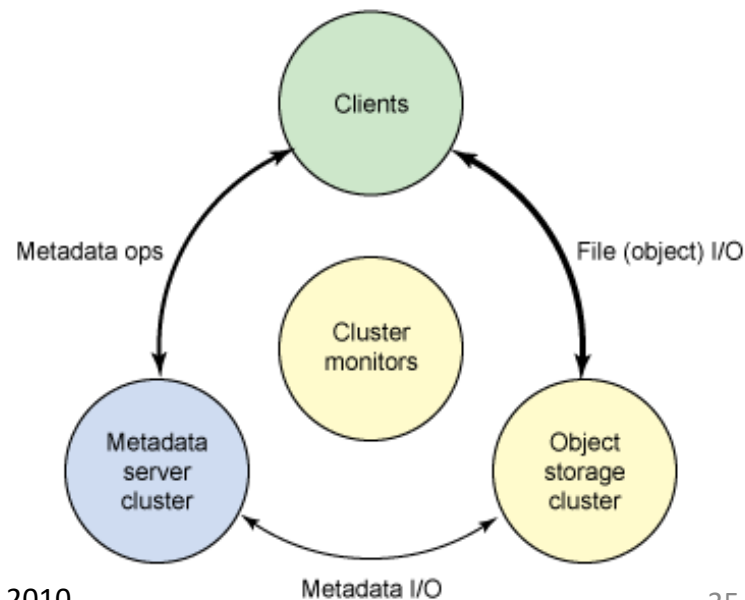
Source: IBM/Red Hat



OpenShift Storage



- IBM/Red Hat OpenShift Container Platform (OCP), renamed OpenShift Data Foundation (ODF), provides software-defined storage for containers that support Kubernetes private, hybrid, and multi-cloud deployments.
- Multi-cloud gateways abstract storage infrastructure so data can be stored in many different places but seen as one persistent store.
- Data can be formatted as files, blocks, or objects to support different Kubernetes workloads and help developers deploy applications across multiple tactical clouds.
- ODF is based on the Ceph open source storage standard for unified storage across single and distributed clusters:
 - Distributed operations
 - No single point of failure
 - Scalable to the exabyte level
 - Replicates data for fault-tolerance
 - Self-healing and self-managing





Selected Tactical Cloud Hardware

NPS Laboratory Tests






Configuration



- Tactical cloud node configurations from Hewlett Packard Enterprise, IBM, and Dell Computer were evaluated.
- At the time of this writing only the HPE EL8000 provided the required specifications for low SWaP-C, ruggedized, high performance computing for tactical edge deployment.
- “Ruggedized” in this instance refers to water, shock, and vibration per MIL-STD 810G tests.



DMO/EABO Tactical Cloud Node



[Video Link](#)




HPE EL8000 HCI



Hyper-Converged Infrastructure (HCI):

➤ Compute and storage components located in the same cluster:

- 112 Xeon cores (CPU)
- 6TB memory (RAM)
- 122TB storage (NVMe)
- NVIDIA AI/ML (GPU)



Source: HPE

Multi-access Edge Compute (MEC) for C5ISRT, AR/VR, video analytics, AI/ML, AI/DNN



EL8000: Specifications



EL8000 HCI tactical cloud edge node:

- Multi-access Edge Compute (MEC), 5G optimization
- Optimized for IoT sensor processing
- Real-time data acquisition and analytics
- Remote management
- Rugged, compact, energy-efficient
- Scalable and modular for real-time AI workloads
- Components can be combined, scaled and hot-swapped
- Intel and Xilinx FPGAs, Intel or Mellanox NICs
- 2TB NVMe internal storage drives per slot
- 6TB memory and 122TB storage per chassis
- 4 PCIe slots per CPU socket, NVIDIA Tesla AI/ML GPUs
- Bare metal to virtualized AI workloads
- Units can be combined for global data center workloads





Power



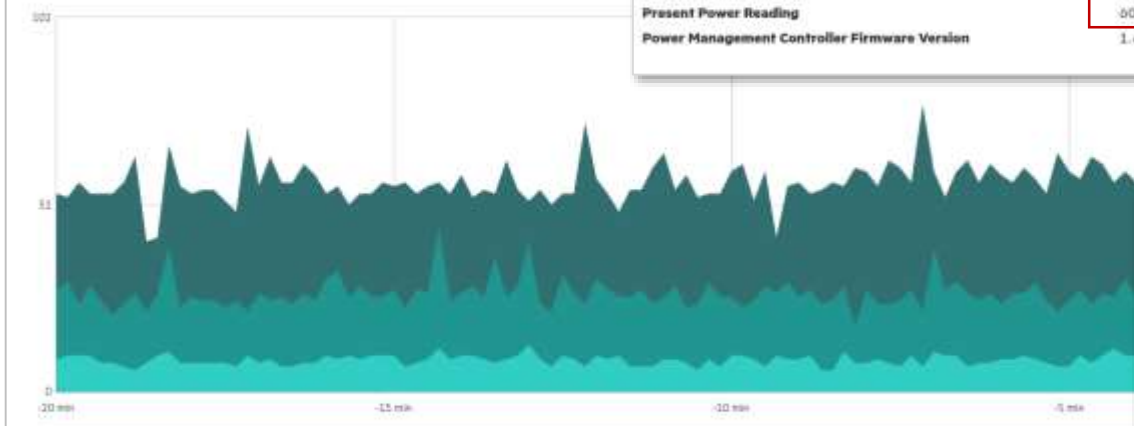
Power & Thermal - Power Meter

Server Power | **Power Meter** | Power Settings | Power | Fans | Temperatures

Select an interval: **20 min** | 24 hr | 1 week

20-Minute History Graph

Power consumption over the past 20 minutes at two-second intervals



Power & Thermal - Power Information

Server Power | Power Meter | Power Settings | **Power** | Fans | Temperatures

Power Summary

Present Power Reading 60 Watts

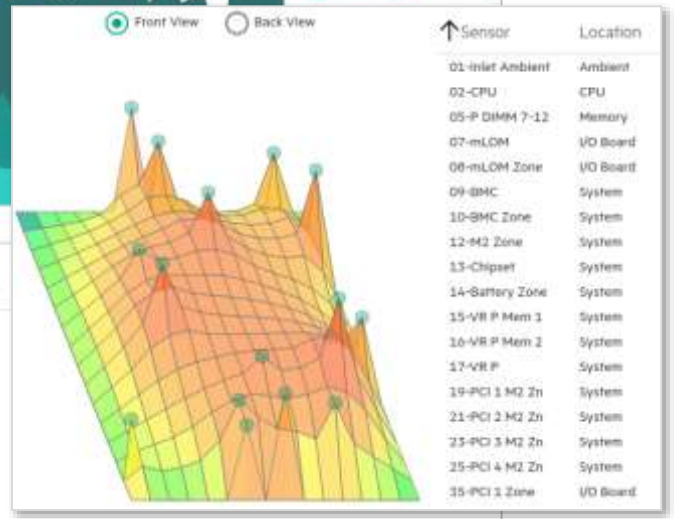
Power Management Controller Firmware Version 1.4.28

- Maximum 64
- Average 64
- Total CPU 24
- Total DIMM 6

2021-07-14 17:44:53

Watts | BTU/hr

Power Status	Now	Power History	5 min
Present Power Reading	64	Maximum Power	75
Power Input Voltage	114	Average Power	54
Power Supply Capacity	1500	Minimum Power	44
Peak Measured Power	91		
Power Regulator Mode	Min (static low)		



NEBS/ASHRAE

- Thermal
- Shock
- Vibration




Agentless Management



- Agentless Management uses out-of-band communication for increased security and stability.
- Health monitoring and alerting is built into the system and begins when power is connected.
- Runs on the iLO hardware, independent of the operating system and processor.
- The management network provides access to the servers in the event of failure in the production network.
- The management network cannot be accessed from production.

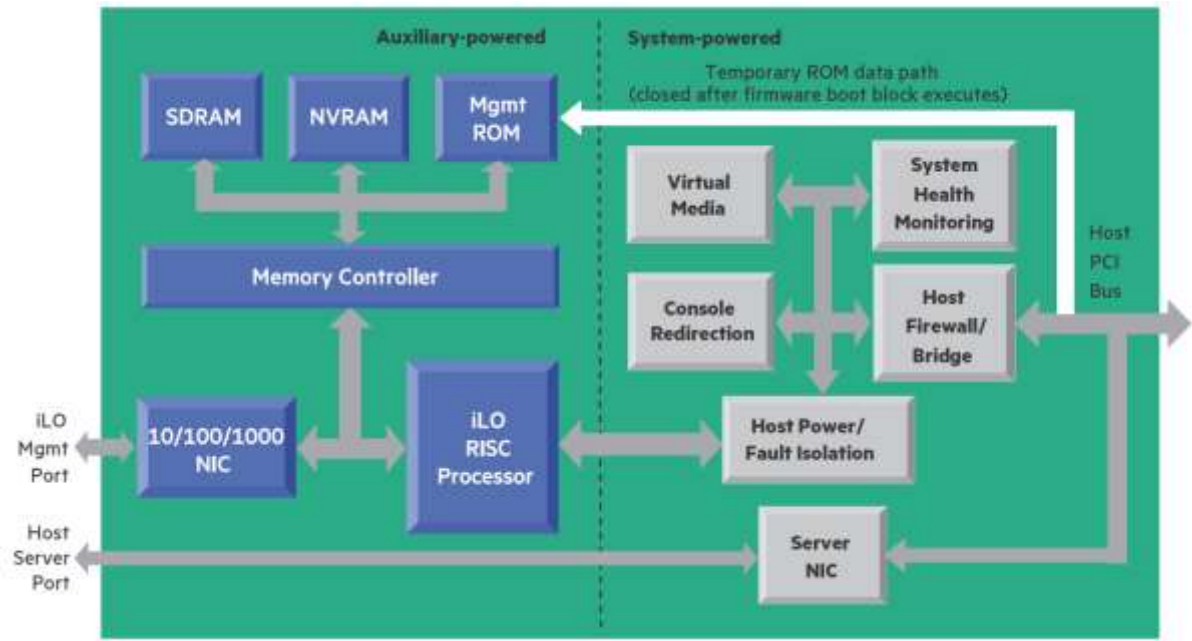
Component	Agentless Management without AMS	Additional information provided when AMS is installed
Server health	<ul style="list-style-type: none"> • Fans • Temperatures • Power supplies • Memory • CPU • NVDIMM 	N/A
Storage	<ul style="list-style-type: none"> • Smart Array • SMART Drive Monitoring (connected to Smart Array) • Internal and external drives connected to Smart Array • Smart Storage Energy Pack monitoring (supported servers only) 	<ul style="list-style-type: none"> • SMART Drive Monitoring (connected to Smart Array, Smart HBA, and AHCI) • iSCSI (Windows) • NVMe drives
Network	<ul style="list-style-type: none"> • MAC addresses for embedded NICs that support NC-SI over MCTP • Physical link connectivity and link up/link down traps for NICs that support NC-SI over MCTP • Fibre Channel adapters that support Hewlett Packard Enterprise vendor-defined MCTP commands 	<ul style="list-style-type: none"> • MAC and IP address for standup and embedded NICs • Link up/link down traps • NIC teaming and bridging information (Windows and Linux) • Supported Fibre Channel adapters • VLAN information (Windows and Linux)
	<ul style="list-style-type: none"> • iLO data • Firmware inventory • Device inventory 	<ul style="list-style-type: none"> • OS information (host SNMP MIB) • Driver/service inventory • Logging events to OS logs ^{1, 2}
	<ul style="list-style-type: none"> • Memory • Drives (physical and logical) 	N/A





Integrated Lights-Out (iLO)

- Protected PCI bus – iLO shields keys and data stored in memory and firmware, and does not allow direct access to keys via the PCI bus.
- Network and management ports – iLO’s firewall and bridge logic prevent any connection between the iLO management port and the server Ethernet port so attacks on the server network cannot compromise iLO and vice-versa.
- Services include:
 - Two-factor authentication
 - One-button secure erase NIST 800-88r1
 - Intelligent OS provisioning
 - Automatic backup, restore, reimage
 - System diagnostics
 - Remote repair
 - Server power
 - Thermal control
 - RESTful API
 - Browser
 - Mobile app
 - SSH client
 - BIOS access



Source: HPE



System Monitor



System Information - NIC Information

Summary Processors Memory **Network** Device Inventory

Collapse All

Physical Network Adapters

Adapter 1 - Intel(R) I210 Gigabit Backplane Connection

Location: Embedded LOM
Firmware: N/A
Status: OK

Network Ports

Port	MAC Address	IPv4 Address	IPv6 Address
1	b4:7a:f1:35:a1:3d	192.168.0.159	2600:bc52:b200:2600:bc52:b200:fe80:ed525cc00f

Adapter 2 - HPE e910 2p 1GbE RJ45 Mod

Location: Embedded mLOM
Firmware: 1.2074.0
Status: Degraded

Network Ports

Port	MAC Address	IPv4 Address	IPv6 Address
1	b4:7a:f1:35:a1:3e	N/A	N/A
2	b4:7a:f1:35:a1:3f	N/A	N/A
3	b4:7a:f1:35:a1:40	N/A	N/A
4	b4:7a:f1:35:a1:41	N/A	N/A

Information - Dashboard

Dashboard Session List Logs

Blades Health State

3 Blades New 3 Blades New 0 Blades Low 0 Blades Low

Blades: 2 Blades Good Health: 2 Blades Warning: 0 Blades Critical: 0 Blades

Topology

SW1 SW2 SW3 SW4 SW5 SW6 SW7 SW8 SW9 SW10 SW11 SW12 SW13 SW14 SW15 SW16 SW17 SW18 SW19 SW20 SW21 SW22 SW23 SW24 SW25 SW26 SW27 SW28 SW29 SW30 SW31 SW32 SW33 SW34 SW35 SW36 SW37 SW38 SW39 SW40 SW41 SW42 SW43 SW44 SW45 SW46 SW47 SW48 SW49 SW50 SW51 SW52 SW53 SW54 SW55 SW56 SW57 SW58 SW59 SW60 SW61 SW62 SW63 SW64 SW65 SW66 SW67 SW68 SW69 SW70 SW71 SW72 SW73 SW74 SW75 SW76 SW77 SW78 SW79 SW80 SW81 SW82 SW83 SW84 SW85 SW86 SW87 SW88 SW89 SW90 SW91 SW92 SW93 SW94 SW95 SW96 SW97 SW98 SW99 SW100

FW1 FW2 FW3 FW4 FW5 FW6 FW7 FW8 FW9 FW10 FW11 FW12 FW13 FW14 FW15 FW16 FW17 FW18 FW19 FW20 FW21 FW22 FW23 FW24 FW25 FW26 FW27 FW28 FW29 FW30 FW31 FW32 FW33 FW34 FW35 FW36 FW37 FW38 FW39 FW40 FW41 FW42 FW43 FW44 FW45 FW46 FW47 FW48 FW49 FW50 FW51 FW52 FW53 FW54 FW55 FW56 FW57 FW58 FW59 FW60 FW61 FW62 FW63 FW64 FW65 FW66 FW67 FW68 FW69 FW70 FW71 FW72 FW73 FW74 FW75 FW76 FW77 FW78 FW79 FW80 FW81 FW82 FW83 FW84 FW85 FW86 FW87 FW88 FW89 FW90 FW91 FW92 FW93 FW94 FW95 FW96 FW97 FW98 FW99 FW100

EL8000CH

Chassis Information

Product Name: EL8000 S11 Blade Infrastructure
Manufacturer: Hewlett Packard Enterprise
Serial Number: MN03460W25
Part Number: P13305-001
Asset Tag: P13305-001
EL8000CH Firmware Version: 3.4-b291
IPv4: 192.168.0.149
IPv6: 2600:bc52:b200:0944:0c9f:80b:10
Host Name: EL8000M440G0460W25
CPLD Version: Synopsys EL8000 v01.00.06 PCA_REV_1D5
EL8000CH Date & Time: 2021-03-07 11:21:53
Chassis Health: OK
Chassis UID: LOD OFF

- Endpoint tamper detection
- 2- factor authentication CAC/PIV
- Global directory/Kerberos authentication
- iLO federation across servers with HTML 5 interface

Source: SEA Laboratory



Remote Provisioning



Intelligent Provisioning (remote server management) – All needed firmware, drivers, and tools are available on the system so the server is immediately ready for provisioning:

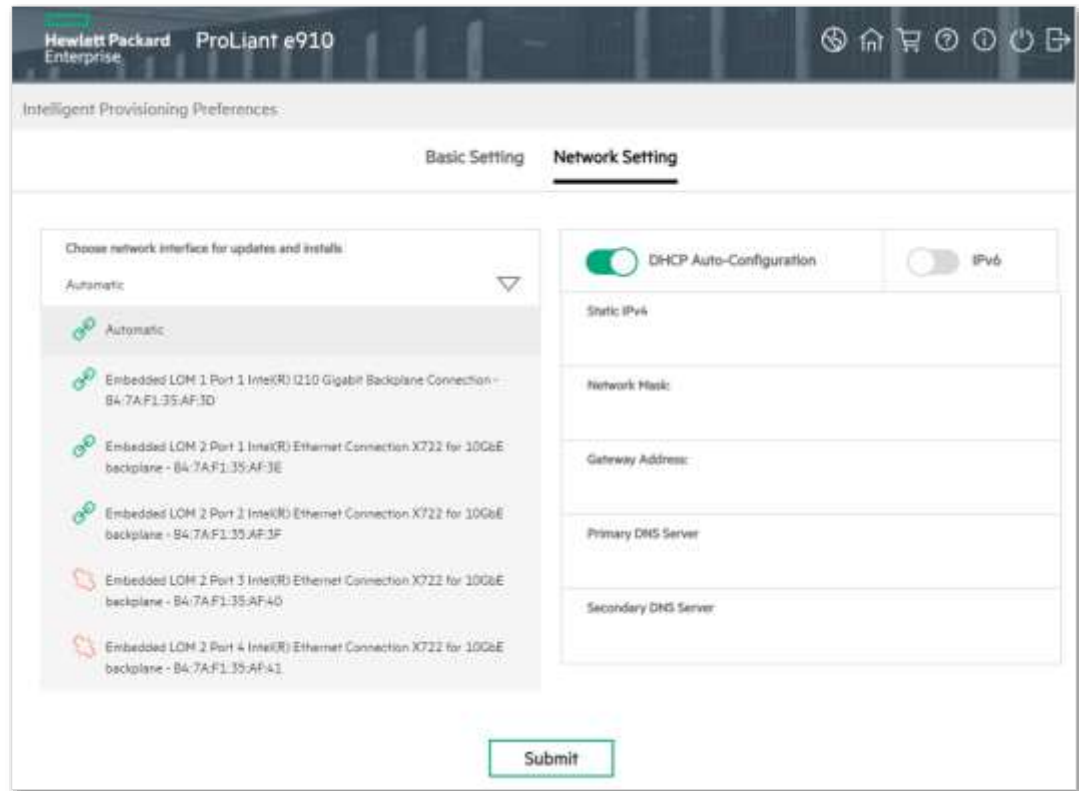
- Perform functions when the server is OFF
- Perform tasks while running an operating system without powering OFF the server.

iLO-dedicated management:

- Secure management firewall
- Out-of-band communication
- Increased security and stability

Monitor internal subsystems:

- Thermal
- Power
- Memory
- Storage
- Machine learning
- Predictive analytics
- Problem recommendations





Erase and Reset



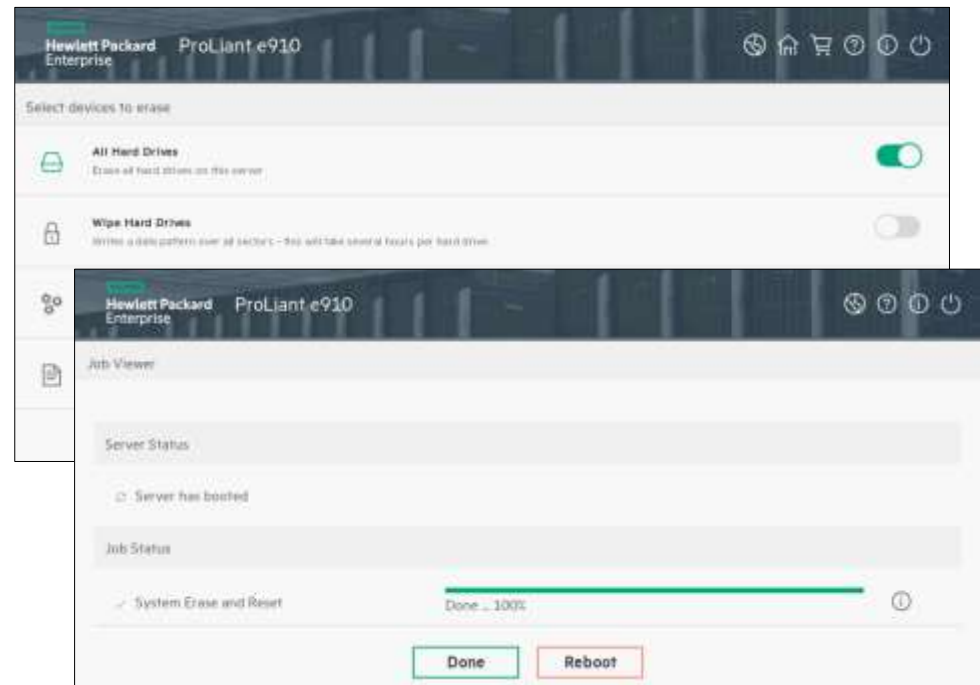
Intelligent Provisioning provides two methods to remotely decommission or repurpose a server :

➤ One-button Secure Erase

- Automatically returns the server and supported components to the default state following NIST Special Publication 800-88 Revision 1, Guidelines for Media Sanitization.

➤ System Erase and Reset

- The System Erase and Reset function overwrites data on drives by using the guidelines from DoD 5220.22-M.
- Software overwrites all block devices attached to the system by applying random patterns.



Source: SEA Laboratory



Remote Log Analysis



Active Health System Viewer

➤ Parse File

- Process and Collate Data
- Analyze for Faults
- Generate Output

➤ Load database

- Configuration data
- Build information
- Fault Analysis data
- PCI Information
- USB Information
- Link Module Information
- Event data

The screenshot displays the HPE InfoSight 'Analyze Logs' interface. At the top, it shows the log file 'Log HPE_MXQ0460W2H_20210318.ahs'. Below this, there are three main sections: SUMMARY, WARRANTY/SUPPORT STATUS, and HARDWARE STATUS. The SUMMARY section lists Product Name (ProLiant #R10), Product ID (P13381-021), and Serial Number (MXQ0460W2H). The WARRANTY/SUPPORT STATUS section shows a status of 'Unknown' and an end date. The HARDWARE STATUS section provides a breakdown of components like CPU, Fan, PSU, Memory, Storage, iFCs, and Other.

Below these sections is a navigation bar with tabs for Overview, Health Alerts, Event Logs, Server Hardware, Firmware & Software, and Customer Advisory. The main content area is divided into several hardware configuration panels:

- System Board:** Serial Part#: MXQ0460W2H, Product ID: P13381-021, PCA Part Number: P0904-001, Build Date: 11/13/2020 10:07:18 AM.
- iLO:** Firmware: iLO 5 v1.31sp4 built on Feb 18 2021, License: iLO Advanced, Hardware: Not Available, IP Address: Not Available.
- Processors:** Processor 0: Intel(R) Xeon(R) Gold 6330H CPU @ 2.30GHz, Stepping: Cascade Lake SP B1.
- Memory:** Installed: 2 x 64 GB 2933 MHz.
- Power Supply:** PSU15: @ 1500 Watts, Firmware: 01.
- Storage:** Slot undefined: HVM Express Controller, Not Available, Device: 1 drives, 1 x 2.05 TB total.
- Networking:** NIC 0: Intel(R) E10 Gigabit Backplane Connection, Firmware: N/A, NIC 2: HPE #910 2p 30Gb RJ45 Mod, Firmware: 1.2874.0.
- BIOS:** BIOS Version: HPE, BIOS Version: H58 v1.82, BIOS Date: 06/21/2020.



Selected Tactical Cloud Software

NPS Laboratory Tests





About



- Azure Stack Hub required 5 servers while the EL8000 provides 4. The HPE team dedicated to an Azure 4-server conversion for the EL8000 conversion for this project was discontinued by HPE. Microsoft did not participate. Azure Stack HCI is an extension of Windows Server Data Center that can run on two servers but at the time of our tests could not operate offline, requiring connection to the Azure public cloud. Tests were discontinued.
- OpenShift also normally requires 5 servers, but an IBM/Red Hat team worked with HPE to make the conversion to 4 servers for the EL8000 configuration used in this project. OpenShift is the software used in Fleet CANES/ACS and C2C24/Project Overmatch.



Configuration

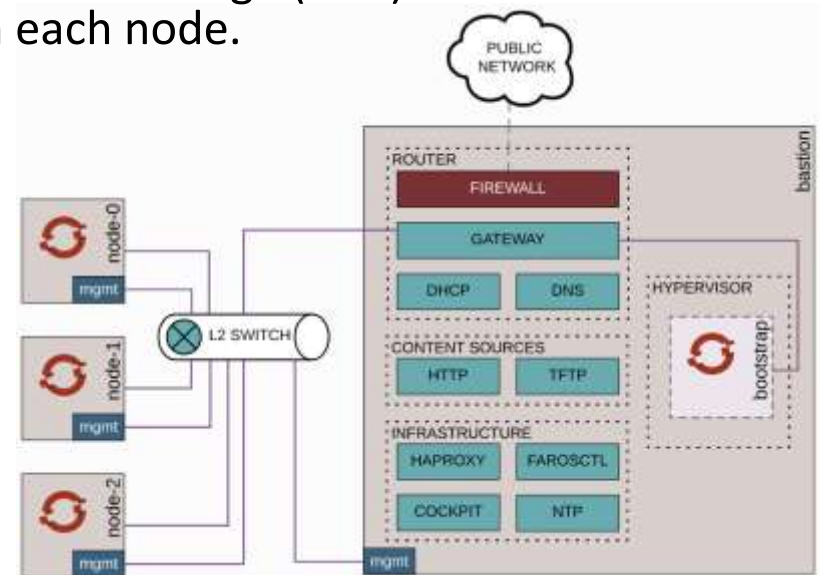


This project tested 4 independent compute nodes/servers in a 5U chassis on small footprint, bare-metal clusters:

- Node 1 is the Bastion/Management Node: The bastion node runs Red Hat Enterprise Linux and hosts the scripts, files, and tools to provision the bootstrap, control-plane, and compute nodes. After deployment the bastion node serves as the administrative node for the cluster.
- Nodes 2, 3 and 4 are the OpenShift Cluster Nodes: The “worker-nodes” run OpenShift Kubernetes and OpenShift Container Storage (OCS) across nodes with management agents on each node.

The architecture is specifically designed for far edge environments with minimal power draw and minimal heat production – to be self-sustaining without dependence on other infrastructure.

The Red Hat OpenShift Container Storage applies Federal Information Processing Standard (FIPS) 140-2 (FIPS-140-2) security requirements for cryptographic modules.



Source: HPE and RedHat



Cluster



The screenshot displays the Red Hat OpenShift console interface. The main content area is titled "Overview" and "Cluster".

Control Plane status popup:

Component	Response rate
API Servers	100%
Controller Managers	100%
Schedulers	100%
API Request Success Rate	100%

Cluster Details:

- Cluster API address: <https://api.edge.seacyber.lah/ov9442>
- Cluster ID: k8s22024-2545-416f-4a1b-6dc0dc06008a
- Provider: None
- OpenShift version: 4.7.0
- Update channel: stable-4.7

Status: Cluster, Control Plane, Operators, Insights (all green)

Alert: Alerts are not configured to be sent to a notification system, meaning that you may not be notified in a timely fashion when important failures occur. Check the OpenShift documentation to learn how to configure notifications with Alertmanager.

Cluster utilization:

Resource	Usage	4:30 PM	4:45 PM	4:30 PM	4:45 PM
CPU 107 available of 100	4.17	[Graph showing low CPU usage]			
Memory 50 GB available of 276.1 GB	50.7% GB	[Graph showing memory usage]			
Filesystem 5.3 TB available of 5.00 TB	71.1 GB	[Graph showing filesystem usage]			
Network transfer	2.93 Mi/s in 4.33 Mi/s out	[Graph showing network transfer]			
Pod count	100	[Graph showing pod count]			

Cluster inventory:

- 1 Nodes
- 208 Pods
- 0 StorageClasses
- 0 PVCs

Quick Starts: Get started with Spring, Write your sample application, Get started with Quarkus using a Helm Chart.

Activity: Recent events include "All pending requests processed", "Received signal to terminate, becoming...", "The minimal shutdown duration of pods...", "All pre-shutdown hooks have been finish...", "Server has stopped listening", etc.



Nodes



Red Hat OpenShift Container Platform

Nodes

Name	Status	Role	Pods	Memory	CPU	Filesystem	Created	Instance type
node-0.edge.seacyberlan	Ready	master, worker	40	15.54 GiB / 125.6 GiB	0.93 / 40 cores	24.13 GiB / 186 TiB	Jun 28, 2023, 3:42 PM	
node-1.edge.seacyberlan	Ready							
node-2.edge.seacyberlan	Ready							

node-0.edge.seacyberlan

Overview Details YAML Pods Events Terminal

Details

Node name: node-0.edge.seacyberlan

Role: master, worker

Instance type: m5.xlarge

Zone: us-east-1a

Node address: 192.168.1.100

Inventory

40 Pods

Storage

Status

Ready

Health checks: Unconfigured

Utilization

Resource	Usage	6:40 PM	7:00 PM	7:15 PM	7:30 PM
CPU	1%	[Graph showing low CPU usage]			
Memory	12.5 GiB / 125.6 GiB	[Graph showing low memory usage]			
Filesystem	24.13 GiB / 186 TiB	[Graph showing low filesystem usage]			
Network transfer	8.07 MiB/s in, 2.07 MiB/s out	[Graph showing network activity]			
Pod count	40	[Graph showing pod count]			

Activity

Creating

Recent events

- 6:13 PM Node node-0.edge.seacyberlan event 8
- 6:11 PM Node node-0.edge.seacyberlan event 8
- 6:04 PM Node node-0.edge.seacyberlan event 8
- 6:27 PM Node node-0.edge.seacyberlan event 8
- 6:24 PM Node node-0.edge.seacyberlan event 8
- 6:22 PM Node node-0.edge.seacyberlan event 8
- 6:22 PM Setting node node-0.edge.seacyberlan...
- 6:19 PM Node node-0.edge.seacyberlan event 8



Pods



The screenshot shows the Red Hat OpenShift console interface. On the left is a navigation sidebar with sections like Administrator, Home, Operators, and Workloads. The main area displays a list of pods under the heading 'Pods'. The table below shows the details of these pods.

Name	Namespace	Status
alermanager-main-0	openshift-monitoring	Running
alermanager-main-1	openshift-monitoring	Running
alermanager-main-2	openshift-monitoring	Running
apicenter-8c70512f-4nbf	openshift-casfi-apicenter	Running
apicenter-8c70512f-4y07	openshift-casfi-apicenter	Running
apicenter-8c70512f-4rhwk	openshift-casfi-apicenter	Running
apicenter-2-444716d7-3reqf	openshift-apicenter	Running
apicenter-644470d0-1f8eq	openshift-apicenter	Running
apicenter-644470d0-1790c	openshift-apicenter	Running
authentication-operator-685456d1-cq7fn	openshift-authentication-operator	Running
celery-operator-5d549b6d1-cfvz6	openshift-operator-licycle-manager	Running
certified-operators-gf5vz	openshift-workspace	Running
cloud-controller-operator-4460b4c9-7u2f	openshift-cloud-controller-operator	Running
cluster-autoscaler-operator-7d8ff0c6-bmxf	openshift-trafficops	Running

On the right, a detailed view of a pod is shown. It features a grid of circular indicators representing different components and their pod counts:

- prometheus-k8s: 2 Pods
- Scaled to 0
- node-exporter: 0 Pods
- openshift-etrics: 1 Pod
- alertrn_r-main: 3 Pods
- cluste_erator: 1 Pod
- thanos-querier: 2 Pods
- kube-s_etrics: 1 Pod
- grafana: 1 Pod
- promet_dapter: 2 Pods
- promet_erator: 1 Pod
- telemeter-client: 1 Pod



Auditors



The screenshot shows the Red Hat Insights dashboard for a single system. The interface includes a navigation sidebar on the left with options like Dashboard, Advisor, Drift, Inventory, Security insights, Vulnerability, Compliance, Policies, Patch, Subscriptions, Register Systems, Remediations, and Product Materials. The main content area is divided into several panels:

- System Overview:** Shows 1 system registered with Insights, 0 stale systems, and 0 systems to be removed. A "Register systems" button is present.
- Vulnerability:** A summary card stating "Red Hat recommends addressing these CVEs with high priority due to heightened risk associated with these security issues". It shows 0 CVEs with security rules impacting 1 or more systems and 0 CVEs with known exploits impacting 1 or more systems. Buttons for "View CVEs" and "View known exploits" are provided.
- CVEs by CVSS score:** A table showing the distribution of CVEs by severity score.

CVSS score	CVE totals	Known exploits
8.0 - 10	0	0
4.0 - 7.9	0	0
0.0 - 3.9	0	0
- Advisor recommendations:** A card indicating "1 incident detected" with the note "Problematic conditions that cause an issue have been actively detected on your systems" and a "View incidents" button.
- Recommendations by total risk:** A bar chart showing 0 Critical, 1 Important, 1 Moderate, and 0 Low risk recommendations.
- Recommendations by category:** A pie chart showing 1 Availability, 0 Stability, 1 Performance, and 0 Security recommendations.
- Compliance:** A card stating "No policies" and explaining that the compliance service uses OpenSCAP policies to track adherence to requirements.
- Remediations:** A card stating "You haven't created any remediation playbooks yet" and suggesting to "Create an Ansible Playbook to remediate or mitigate vulnerabilities or configuration issues".



Node Metrics



Project: openshift-monitoring

Monitoring

Dashboard Metrics Alerts Events

Rate of received packets [Show FromQL](#)

30m [Reset zoom](#) Stacked

pod 1

- thanos-gateway-5d85cc7b-7q4qg
- thanos-gateway-5d85cc7b-r2fwe
- telemeter-client-c094c3a-9u48
- node-exporter-cc84b
- cluster-monitoring-operator-bc755d796-6y8qp
- node-exporter-4d67j
- openshift-state-metrics-8255645c54-jqct8
- prometheus-k8s-1
- prometheus-operator-79568f5c5-2jpd
- prometheus-adapter-6470999c5d-8y6f9
- alertmanager-main-0
- kube-state-metrics-55a6f5cc7-1056e
- node-exporter-r2krm
- prometheus-k8s-0
- alertmanager-main-2
- alertmanager-main-1
- prometheus-adapter-6470999c5d-8k8d
- grafana-6d88678d5f-1scqj

Project: dmo

Pods: Pod details

postgresql-1-kc2dg Restart Actions

[Details](#) [YAML](#) [Environment](#) [Logs](#) [Events](#) [Terminal](#)

Pod details

Memory usage

CPU usage

Filesystem

Network in

Network out

Network out



Pod Metrics

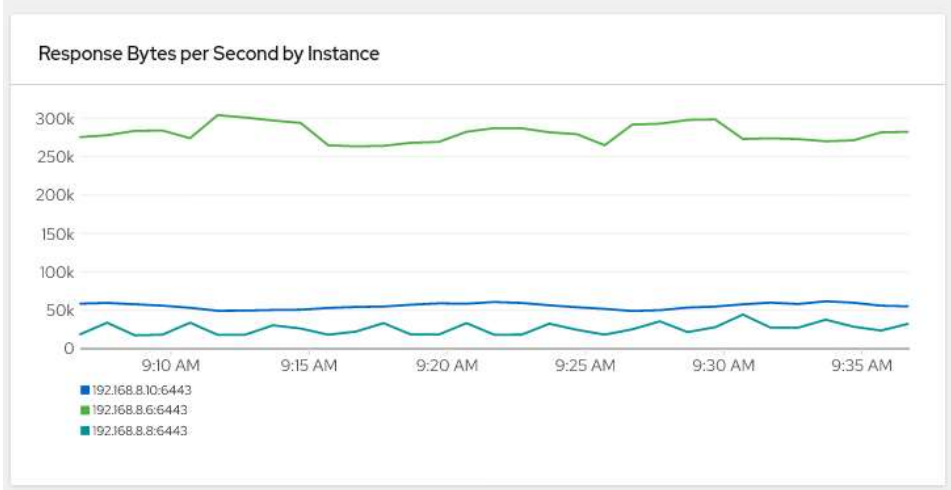

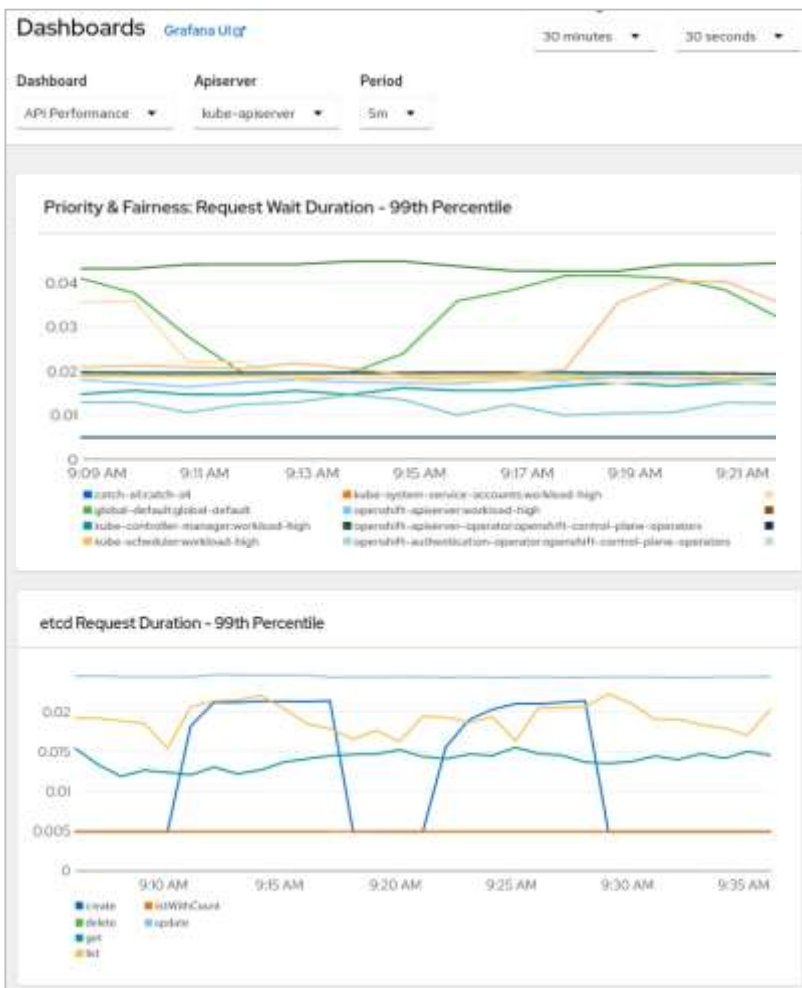


The screenshot displays the Red Hat OpenShift Container Platform monitoring interface. The main view shows the 'Monitoring' dashboard for the 'opensearch-apiserver' project. A 'Pod details' window is open for the pod 'kube-apiserver-node-0.edge.seacyber.lan', which is in a 'Running' state. This window contains two line graphs: 'Memory usage' and 'CPU usage', both plotted from 3:30 PM to 4:00 PM. The memory usage graph shows a peak near 2 GiB, and the CPU usage graph shows a peak near 200m. Below the graphs, a table lists the pod's resource usage and status.

Pod Name	Status	Replicas	Ready	Age	Node	Memory	CPU	Created
kube-controller-manager-node-0.edge.seacyber.lan	Running	4/4	0		node-0.edge.seacyber.lan	377.5 MiB	0.027 cores	Jun 28, 2021, 3:57 PM
kube-controller-manager-node-1.edge.seacyber.lan	Running	4/4	0		node-1.edge.seacyber.lan	250.3 MiB	0.010 cores	Jun 28, 2021, 3:56 PM
kube-controller-manager-node-2.edge.seacyber.lan	Running	4/4	0		node-2.edge.seacyber.lan	339.0 MiB	0.022 cores	Jun 28, 2021, 3:55 PM

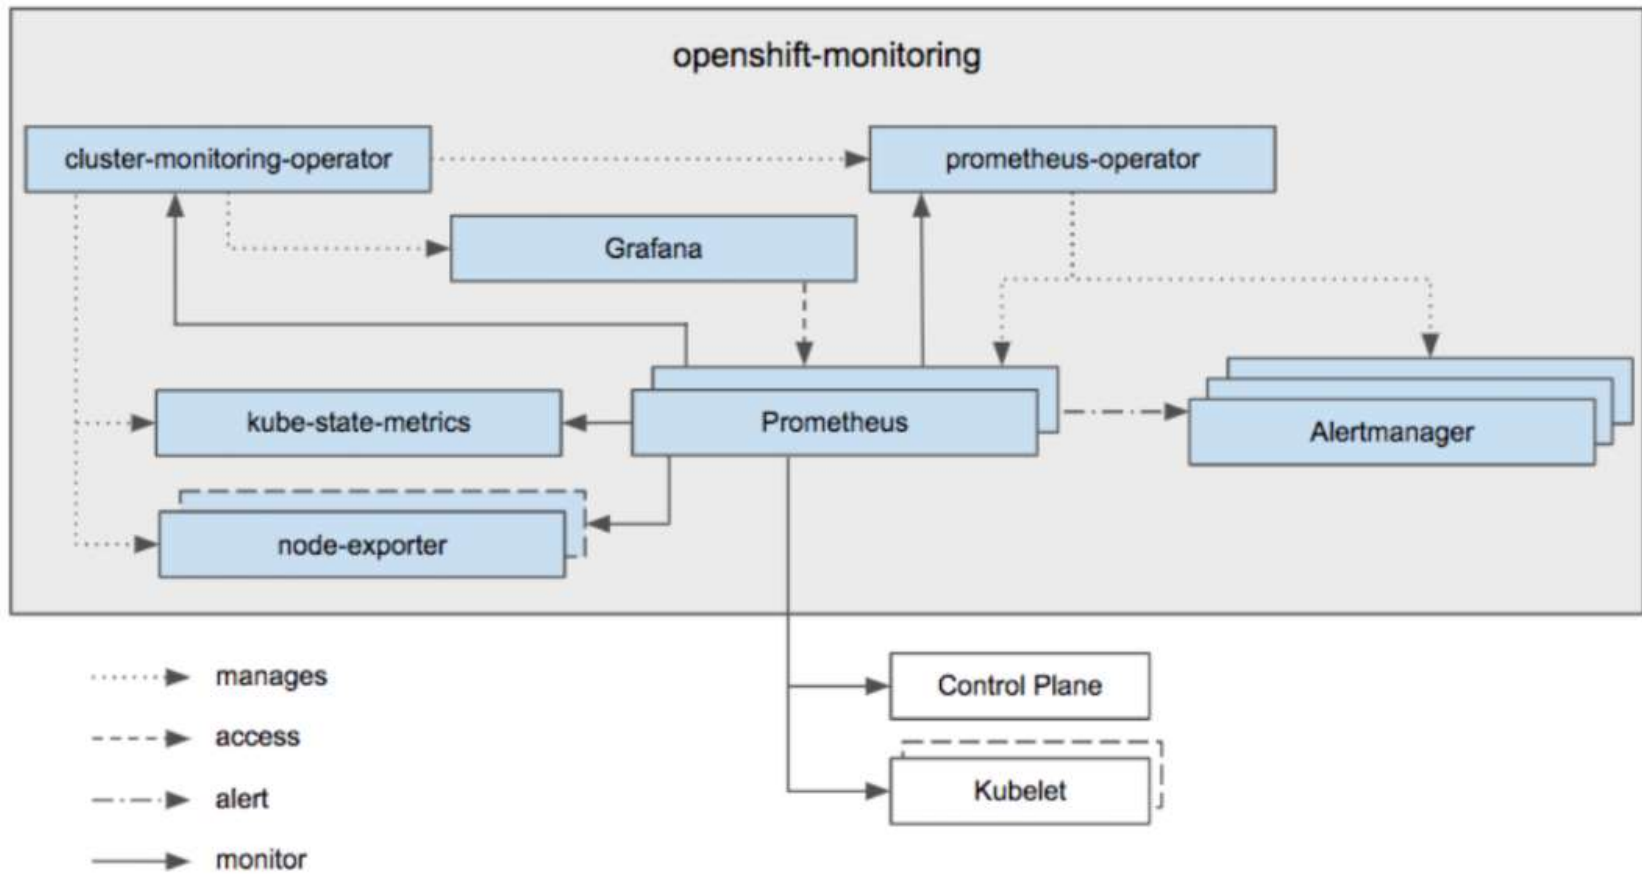


Grafana





Prometheus





Prometheus



Red Hat OpenShift Container Platform

Alerting **Alertmanager UI**

Alerts Silences Alerting rules

Filter Name Search by name

Name	Severity
etcdMembersDown etcd cluster 'etcd' members are down (0)	Critical
TargetDown 33.33% of the kube-controller-manager/kube-controller-manager targets in openshift-kube-	Warning
TargetDown 33.33% of the network-metrics-service/network-metrics-service targets in openshift-multus namespace...	Warning
TargetDown 33.33% of the scheduler/scheduler targets in openshift-kube-scheduler namespace are down.	Warning
TargetDown 33.33% of the kubefit/kubefit targets in kube-system namespace are down.	Warning
TargetDown 33.33% of the node-exporter/node-exporter targets in openshift-monitoring namespace are down.	Warning
TargetDown 33.33% of the controller-manager/controller-manager targets in openshift-controller-manager namespace in...	Warning
TargetDown 33.33% of the dns-default/dns-default targets in openshift-k8s namespace are down.	Warning
TargetDown 33.33% of the nvidia-admission-controller/nvidia-admission-controller targets in openshift-multu...	Warning

Metrics **Platform Prometheus UI**

Refresh off Actions

30m Reset zoom

Hide graph Stacked

Insert metric at cursor

Add query Run queries

```
sort_desc(sum(sum_over_time(ALERTS[alertstate="firing"])[24h])) by (alertname))
```

alertname	Value
Watchdog	1809
AlertmanagerReceiversNotConfigured	1786
KubePodNotReady	12
ClusterOperatorDown	7

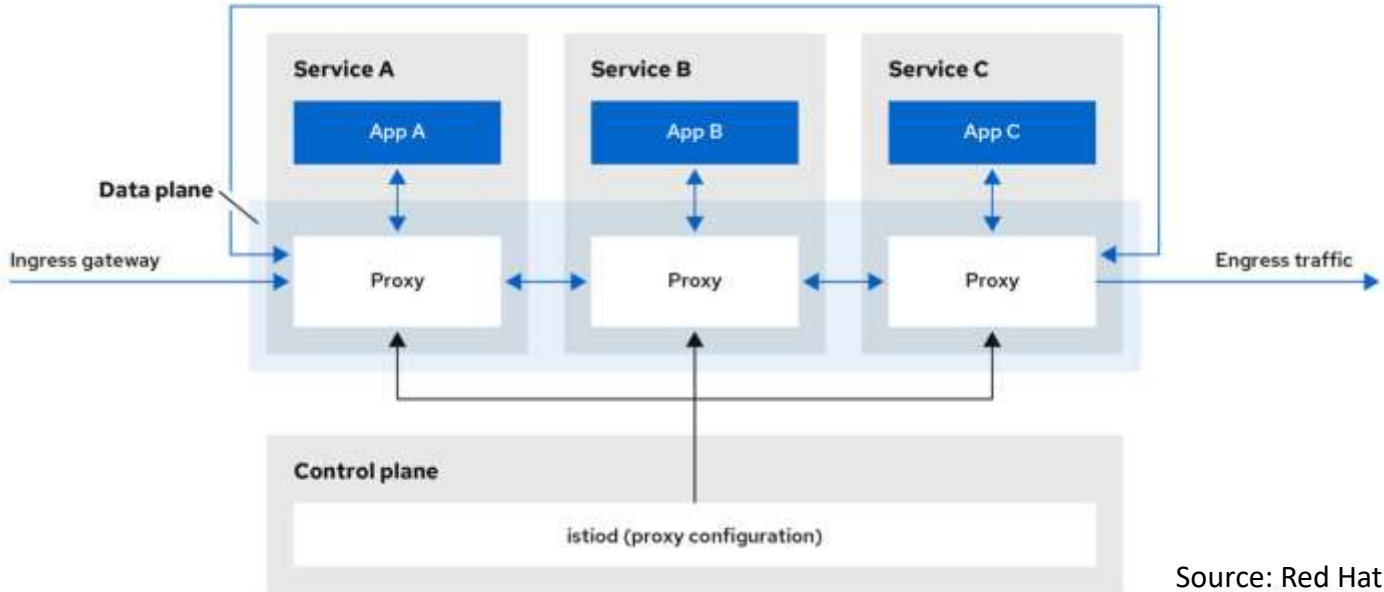
1 - 4 of 4 of 1



Service Mesh



- The service mesh (Istio open source) consists of a data plane and a control plane
- Intelligent proxies run alongside pod application containers to intercept, control and modify inbound and outbound communication
- Benefits include:
 - Centralized point of control in an application
 - Enterprise applications split into modular services to ease scaling and maintenance
 - Load balancing
 - Authentication
 - Failover
 - Monitoring
 - Rate limiting
 - Access control.



Source: Red Hat



Selected C2/ISR Services

ESRI - Industry

NPS Laboratory Tests





Enterprise Suite



Asset		TCPED				AI/ML		Hybrid Cloud						
		Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka/Stream	Tactical Cloud
Enterprise	ArcGIS Enterprise	X	X	X	X	X	X	X	X	X	X	X	X	X
	ArcGIS Enterprise Portal		X	X	X	X			X	X	X			X
	ArcGIS Data Store									X	X	X		X
	ArcGIS Web Adaptor									X	X	X		X
	ArcGIS Server [Advanced]		X	X	X	X	X	X	X	X	X	X	X	X

- Enterprise Portal serves as the central hub and common user interface for C2/ISR services.
- Data Store is the data storage server; Web Adaptor to integrate with existing servers and security infrastructure.
- Server is the primary enterprise geodatabase with feature and geodata services, including advanced raster analysis and surface generation, and integrated analysis of raster and vector data.

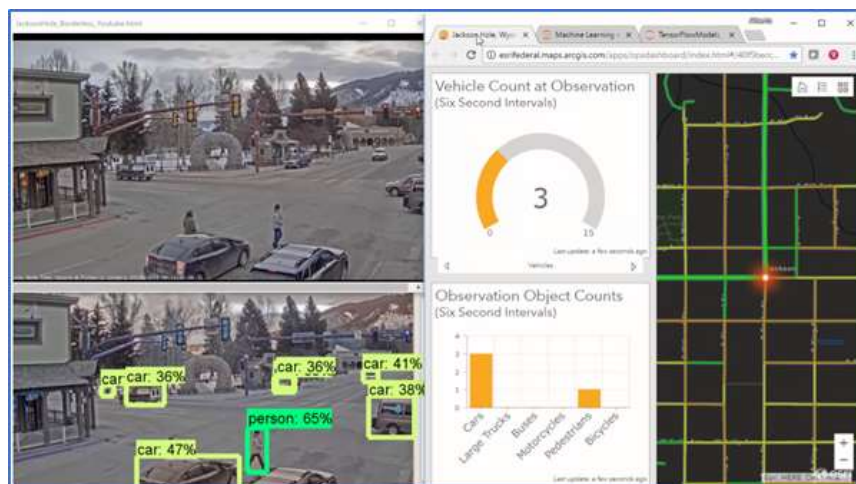


Enterprise Extensions



Asset	TCPED				AI/ML			Hybrid Cloud					
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka/Stream	Tactical Cloud
ArcGIS GeoAI Toolbox			X	X		X	X						X
ArcGIS GeoAnalytics Server			X	X		X	X						X
ArcGIS GeoEvent Server		X	X	X		X	X				X		X
ArcGIS Image Server		X	X	X	X	X	X	X					X
ArcGIS Knowledge Server		X	X	X	X	X	X	X					X
ArcGIS Mission Server	X	X	X	X	X			X				X	X
ArcGIS Notebook Server	X	X	X	X	X	X	X		X	X	X		X
ArcGIS Workflow Manager Server	X	X	X	X	X	X	X	X					X

- GeoAnalytics Server: Workflows.
- GeoAI Toolbox: Geospatial AI/ML.
- GeoEvent Server: Sensor streams.
- Mission Server: Tactical SA.
- Notebook Server: Mobile AI/ML.
- Workflow Manager: Scheduling.





Applications



Asset	TCPED				AI/ML		Hybrid Cloud						
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka/Stream	Tactical Cloud
ArcGIS ATAK/iTAK	X	X			X			X				X	X
ArcGIS Dashboards					X								X
ArcGIS Data Reviewer			X			X	X						
ArcGIS Defense Mapping		X	X	X									X
ArcGIS Drone2Map	X	X	X	X		X						X	X
ArcGIS Excalibur			X	X	X								X
ArcGIS Experience Builder			X	X	X	X	X						X
ArcGIS Insights						X	X						X
ArcGIS Intelligence Toolbox		X	X	X		X	X						X
ArcGIS LocateXT		X	X	X	X	X	X						X
ArcGIS Mission Manager	X	X	X	X	X			X				X	X
ArcGIS Pro Intelligence			X	X	X	X	X						X
ArcGIS Production Mapping			X	X	X								X
ArcGIS Publisher					X								X
ArcGIS Workflow Manager	X				X		X	X					X

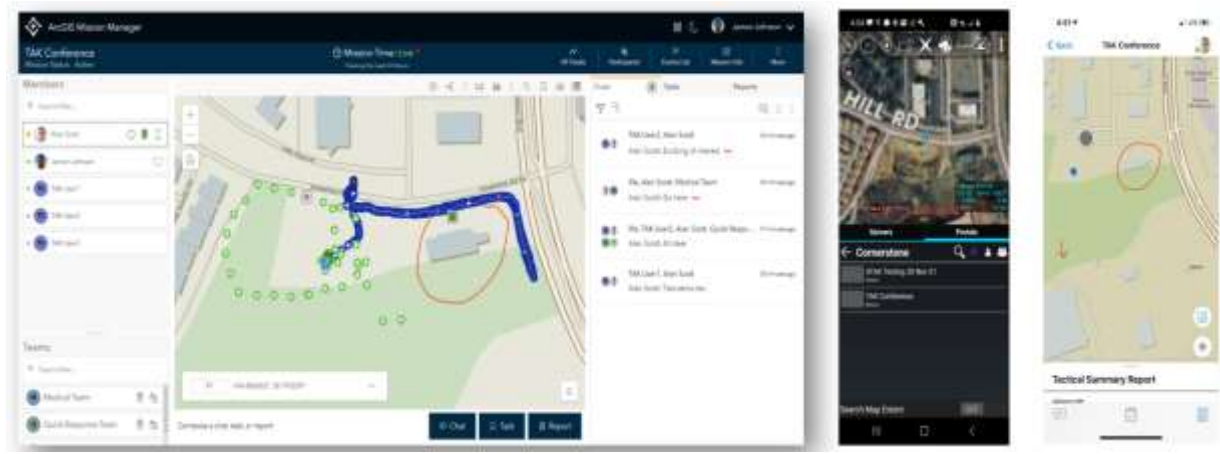
Applications



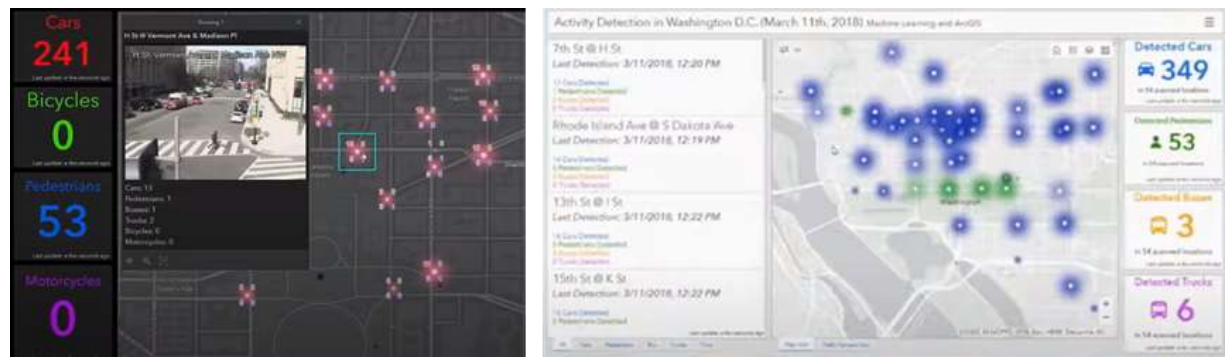
Applications



ArcGIS TAK Integration (Source: ESRI)

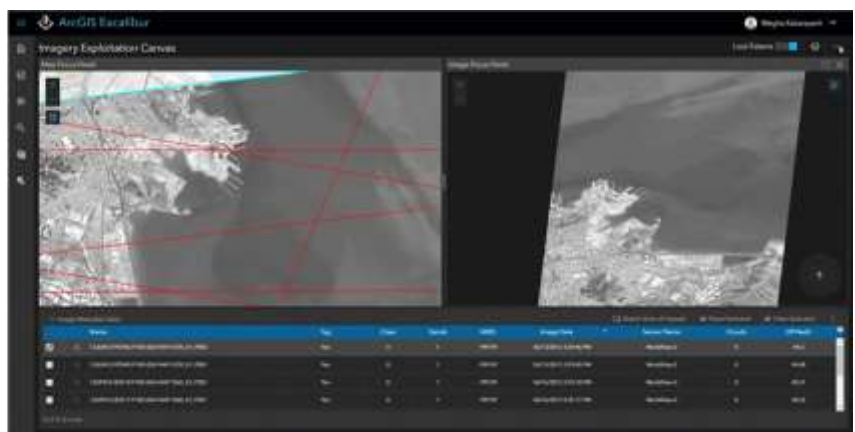
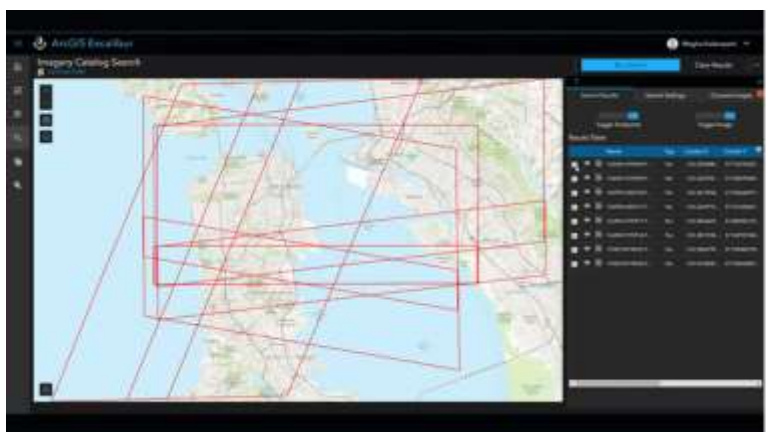


Command Operations Dashboards (Source: ESRI)





Applications



Excalibur Tactical Exploitation (Source: ESRI)



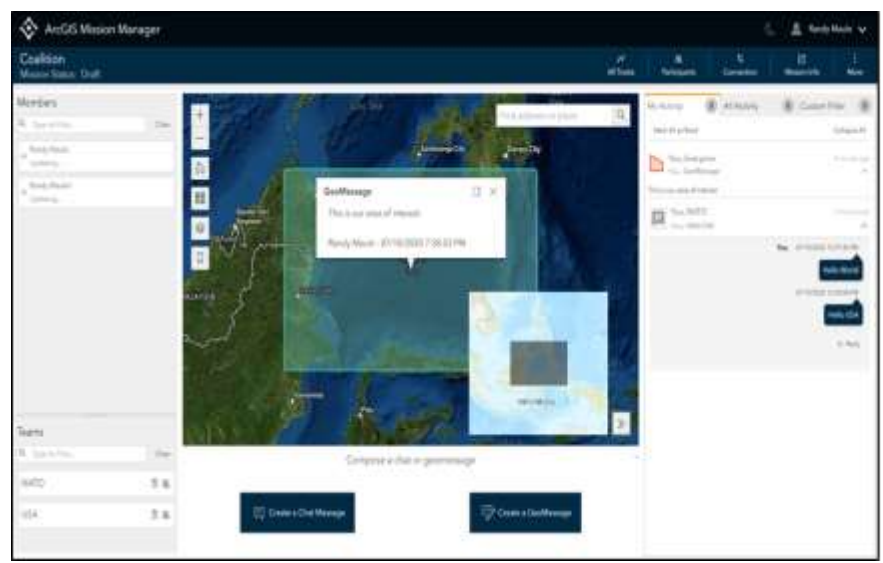
Mission Dashboard (Source: ESRI)



ArcGIS Drone2Map (Source: ESRI)

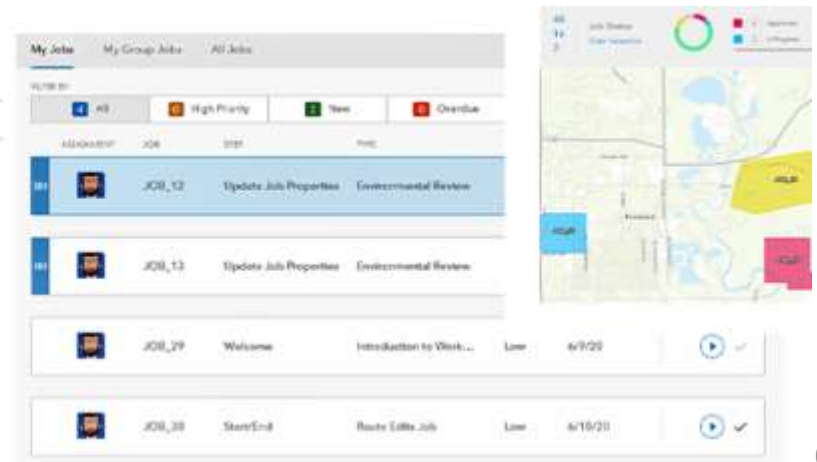
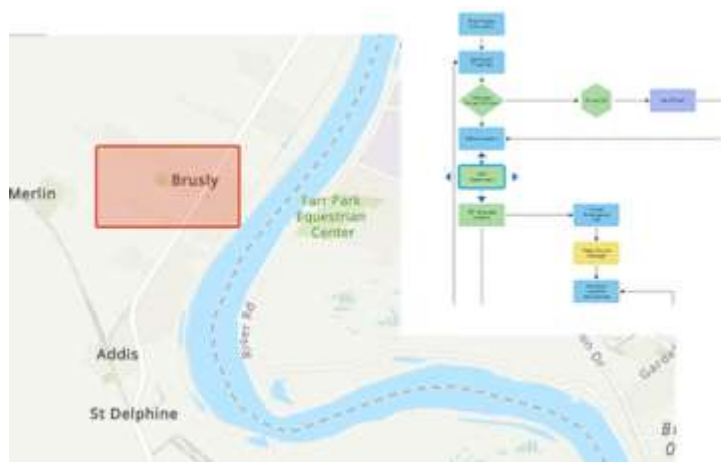


Applications



Mission Portal
(Source: ESRI)

Workflow
Manager
(Source: ESRI)





Field Apps



Asset	TCPED				AI/ML		Hybrid Cloud							
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka/Stream	Tactical Cloud	
ArcGIS Collector		X											X	
ArcGIS Field Maps		X	X	X	X		X	X					X	X
ArcGIS Mission Responder	X	X	X	X	X			X					X	X
ArcGIS Navigator			X		X			X					X	X
ArcGIS QuickCapture		X											X	X
ArcGIS Survey123		X	X	X	X		X	X					X	X
ArcGIS Workforce	X	X	X	X	X			X					X	X

- Collector: Mobile iOS/ Android.
- Navigator: Turn-by-turn mobile navigation.
- Survey123: Field collection audio, images, and questions.
- Workforce: CTP field to HQ.





Selected C2/ISR Services

TAK – Government

NPS Laboratory Tests





Tactical Assault Kit (TAK)



Asset	TCPED				AI/ML			Hybrid Cloud					
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
TAK Server	X	X	X	X	X			X	X	X	X	X	X
ATAK/iTAK	X	X						X					X
WebTAK	X				X								X
VTAK	X												X


Servers, containers, and mobile apps with versions for military forces, law enforcement, and emergency responders:

- ATAK/iTAK: Android/iPhone
- WebTAK: Browser
- VTAK: Virtual Reality

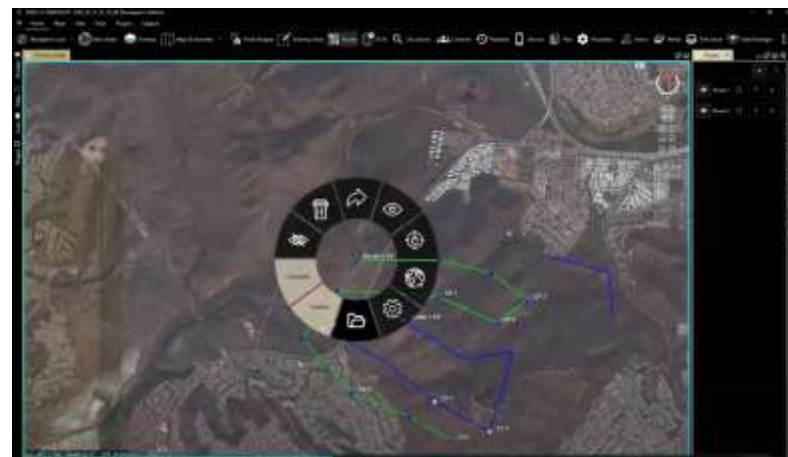




Tactical Assault Kit (TAK)



ATAK map interface (Source: TAK)



TACX map tools (Source: TAK)



VTAK TOC (Source: TAK)



TAK Data Plugins

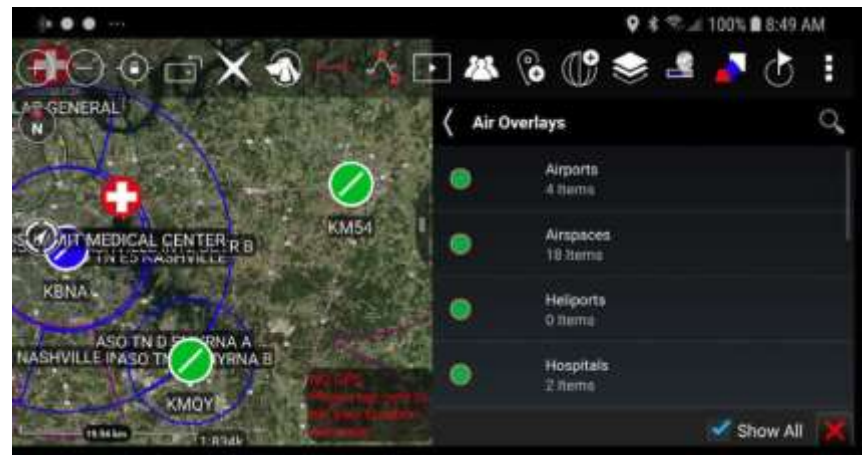


Asset		TCPED			AI/ML	Hybrid Cloud								
		Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
Data	AO Update		X	X	X									
	Air Overlays			X	X									
	ArcGIS		X											
	Arc4Recon	X	X	X										
	Chokepoint		X	X	X									
	EZAZ	X	X	X										
	Fire Area Survey	X	X	X										
	GeoTAKCam	X	X	X										
	GEEP		X	X										
	Talon Point		X	X										

- Add data and functions to the TAK family of devices.
- Open API and SDK facilitate plugin development to enhance the core mapping application with tools for mission requirements.



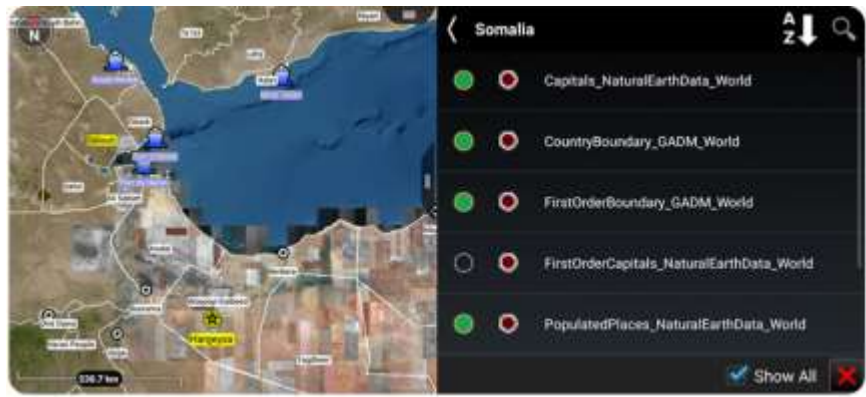
TAK Data Plugins



Air Overlay (Source: TAK)



GeoTAKCam (Source: TAK)



Google Streaming Services (Source: TAK)



Assault Zone DB (Source: TAK)



TAK Communication Plugins

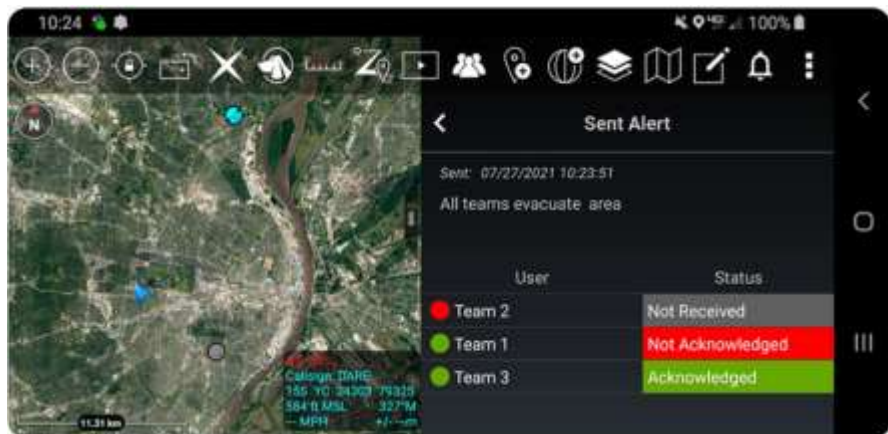


Asset	TCPED				AI/ML		Hybrid Cloud						
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
Alert		X		X			X						
ESChat							X						
Hammer							X						
HUD				X									
ICE Voice							X						
ICOM							X						
MobileJECL							X						
RF Propagation		X	X	X									
SPR							X						
SIP							X						
TAK Chat		X	X		X		X						
TAK ICU		X			X		X						
Wave							X						

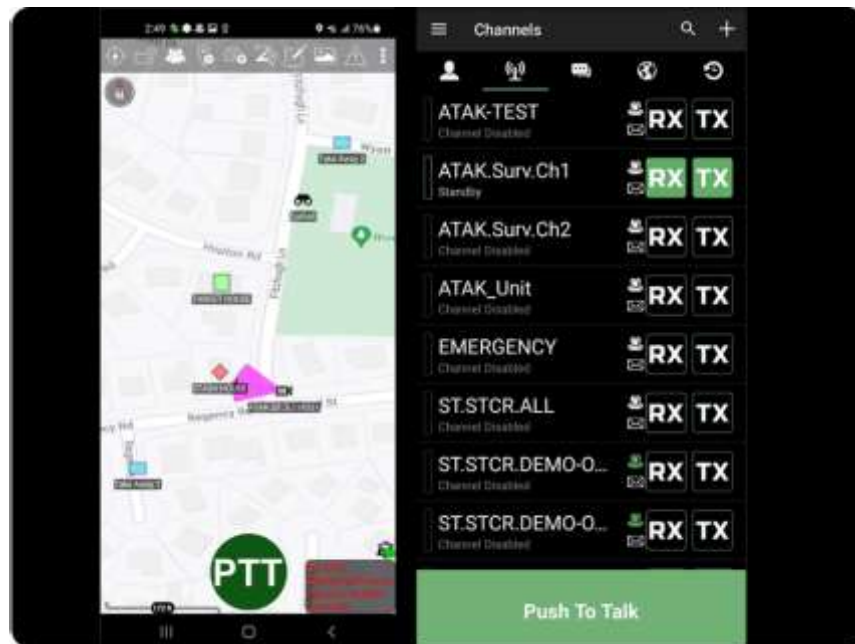
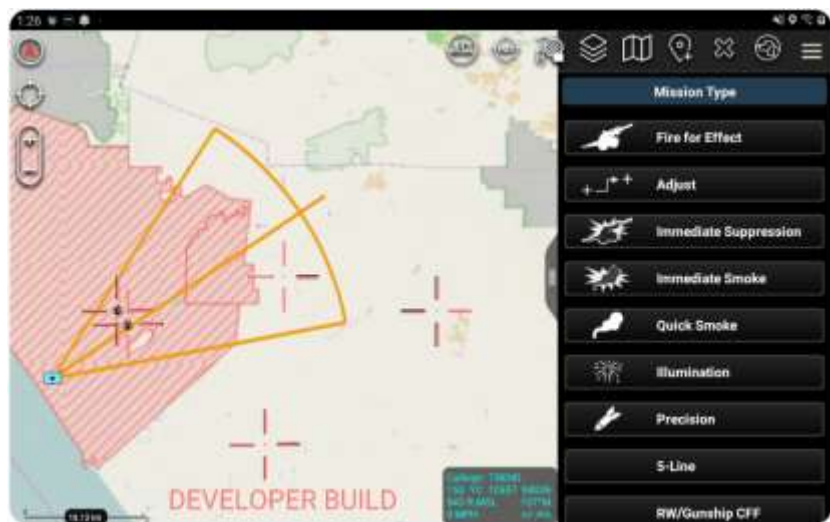
➤ Enhancements range from team communication, to external device integration, to alerts and notifications.



TAK Communication Plugins



Alert Status Messages (Source: TAK)



Push to Talk (Source: TAK)

Joint Effects Coordination Link (Source: TAK)



TAK GPS Plugins



Asset	TCPED				AI/ML		Hybrid Cloud						
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
AR Repeater		X	X										
ATOS		X	X										
AuSS		X	X	X									
Compass Nav			X										
DFT Sensors		X	X	X									
Drifter				X									
EMAPS		X	X	X									
Intercep			X	X									
Last Known Location		X	X										
Munter				X									
Neon		X	X	X									
Stack Manager		X	X	X									
Tetra		X	X	X	X								
VNS			X										

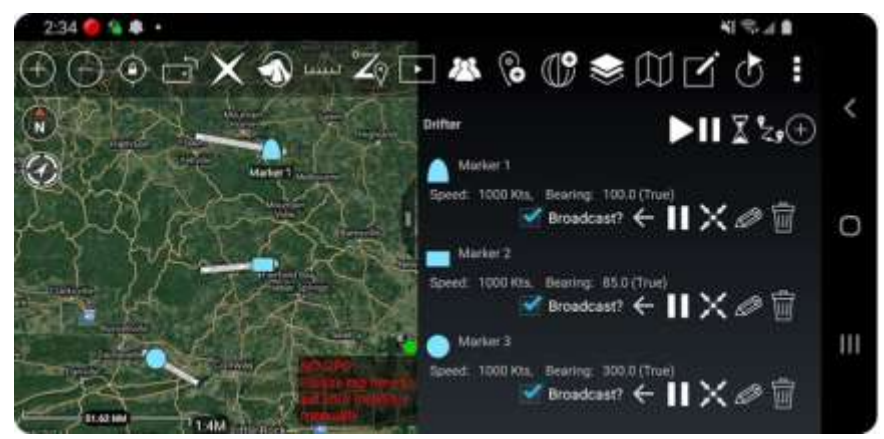
➤ Enhancements range from user navigation, to Augmented Reality (AR), to tracking and targeting.




TAK GPS Plugins



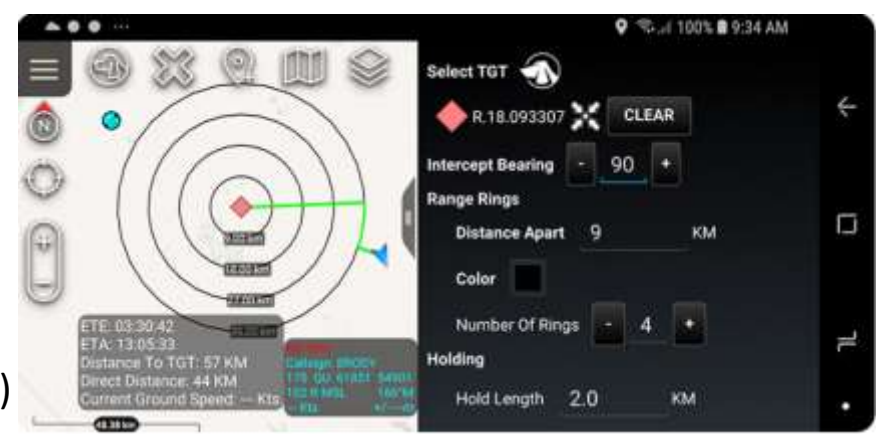
ATOS LOS Tracker (Source: TAK)



Drifter Dead Reckoning (Source: TAK)



GPS Denied Tracking (Source: TAK)



Intercept Bearing Calculation
(Source: TAK)

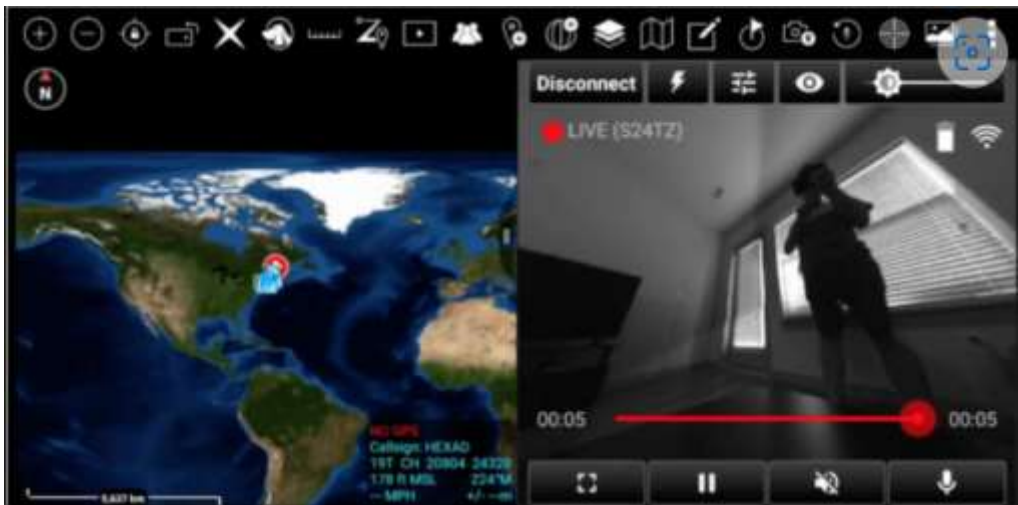


TAK Image Plugins



Asset	TCPED				AI/ML			Hybrid Cloud					
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
Bounce Viewer	X	X											
Checkpoints	X	X	X										
Milestone	X	X											
SSE Tool	X			X									
Vulcane	X												

Images

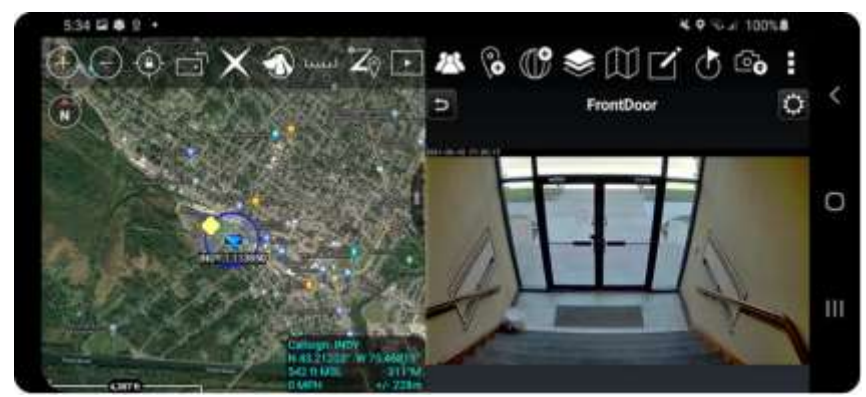


Bounce Omnidirectional Tactical Camera (Source: TAK)

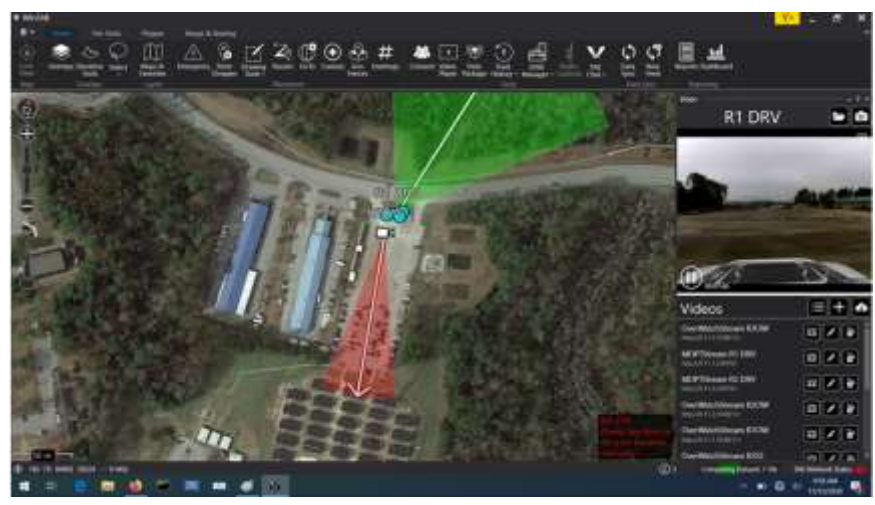


Checkpoints Video Detection and Alert
(Source: TAK)

Milestone Location Stream



(Source: TAK)



Vulcane Vehicle Camera System
(Source: TAK)



TAK Sensor Fusion Plugins



Asset	TCPED				AI/ML		Hybrid Cloud						
	Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
ADS-B	X	X	X										
AVO	X	X	X										
Building Manager	X	X	X										
CBRN	X	X	X										
Deep Purple	X	X	X										
Effects	X	X	X										
FoCUS	X												
Ninja	X	X	X										
Point Mensuration Tool	X	X	X										
Prowl	X	X	X										
Somewhere	X	X					X						
UAS Tool	X	X											
UGV Tool	X	X											
Wx Report	X	X	X										

➤ To increase the number of sensed entities available on the platform, the sophistication of the fusion, and the processing capabilities.



TAK Sensor Fusion Plugins



ADS-B Service Integration (Source: TAK)



CBRN Sensor Integration (Source: TAK)



Somewhere Satellite Hotspot

(Source: TAK)





TAK Operations Plugins



Asset		TCPED				AI/ML		Hybrid Cloud						
		Task	Collect	Process	Exploit	Disseminate	Data Science	Automation	Collaborate	Containers	Microservices	Kubernetes	Kafka	Tactical Cloud
Operations	CMP	X	X	X	X	X			X					
	Data Sync		X	X	X	X			X					
	ExCheck		X	X		X			X					
	Mission Workflow	X		X		X			X					
	Pager		X	X	X		X							
	Reports		X	X										
	TAK-ML		X	X	X		X	X						
	TAK Replay				X									
	TRAX		X	X	X	X								
	WASP		X	X	X	X			X					

Assessment, reporting and decision support functions:

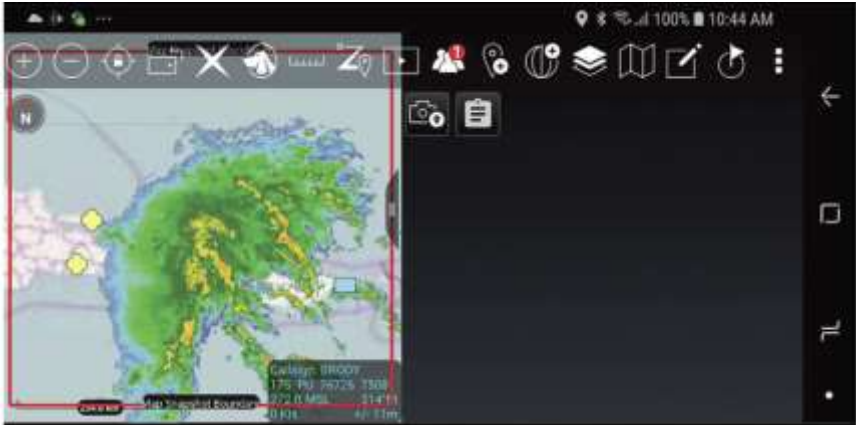
- Calculations for conflict management
- Tactical guidance
- Libraries for the development of AI/ML for decision support



TAK Operations Plugins



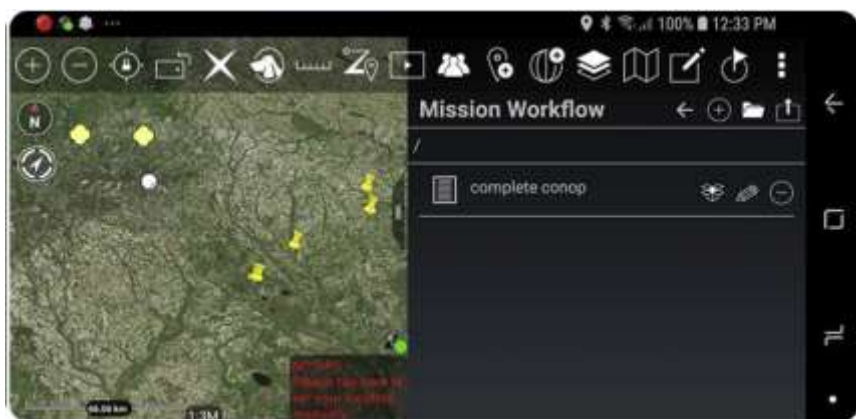
Mission Planner (Source: TAK)



Air-Maritime-Ground COP (Source: TAK)



Mission Workflow (Source: TAK)



Search and Rescue (Source: TAK)





Edge AI/ML

NPS Laboratory Tests





About



- This section discusses AI/ML for the EL8000 tactical node:
 - AI/ML is available for the EL8000 tactical node chassis and servers
 - AI/ML is available for the OpenShift tactical node operating software.
- Sensor Open Systems Architecture (SOSA) standards are implemented for hardware, software, functions and behaviors; and for electrical-mechanical interfaces for communications, EO/IR, EW, radar and SIGINT interoperability.
- AI/ML and AI deep neural network (AI/DNN) algorithms provide context-aware applications that can track and identify objects, analyze motion for events, and extract intelligence from analog or digital streams using an open, low-latency streaming web interface and control API to:
 - Learn the spectrum instantly and automatically with contextual analysis
 - Detect and classify RF emissions across bandwidths to report anomalies and threats in near real-time
 - Assess wide-band and narrow-band signals, analog single carrier modulations, multi-carrier modulation schemes, cellular and infrastructure signals, ISM-band signals (e.g., Wi-Fi, Bluetooth) and mobile radio services
 - Apply AI/DNN in real-time for signal identification and re-train the neural network as needed for new signals/anomalies.



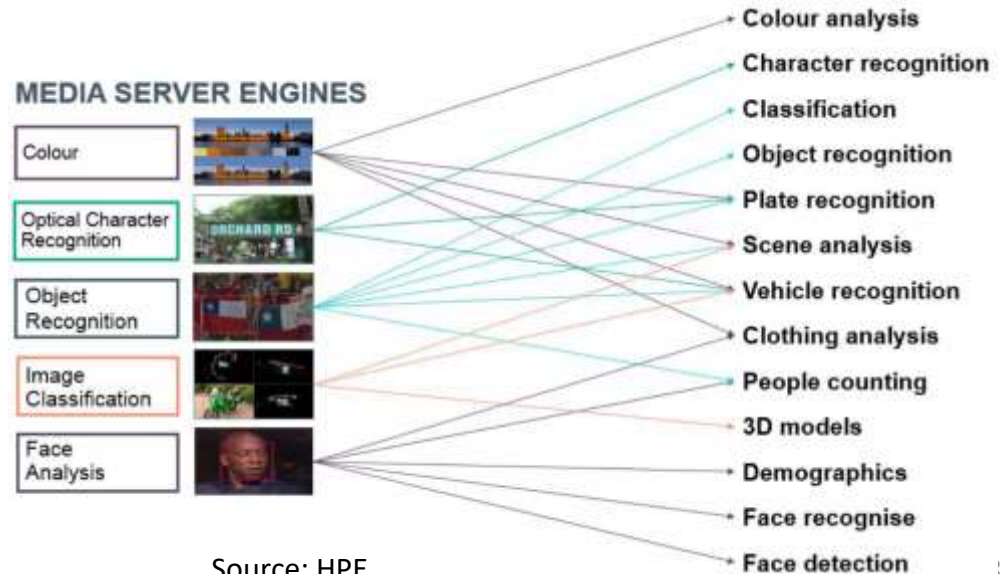
Embedded AI/ML



EL8000 Software Capabilities	EL8000 Hardware
AI/ML	NVIDIA Tesla T4
Anomaly Detection	NVIDIA Tesla T4
Body Recognition	NVIDIA Tesla T4
Facial Attributes (gender, age, etc.)	NVIDIA Tesla T4
Facial Recognition	NVIDIA Tesla T4
Facial Expression Analysis	NVIDIA Tesla T4
License Plate Recognition	NVIDIA Tesla T4
Object Detection and Classification	NVIDIA Tesla T4
Object Tracking and Pathing	NVIDIA Tesla T4

- EL8000 integrates hardware and software to optimize the platform for sensor collection and processing at the tactical edge
- Embedded AI/ML for situational awareness and GPU-accelerated data visualization for tactical decision support

- AI/ML enables the tactical node to learn from examples.
- AI deep neural network (AI/DNN) algorithms automate that training.



Source: HPE



AI as a Service (AaaS)

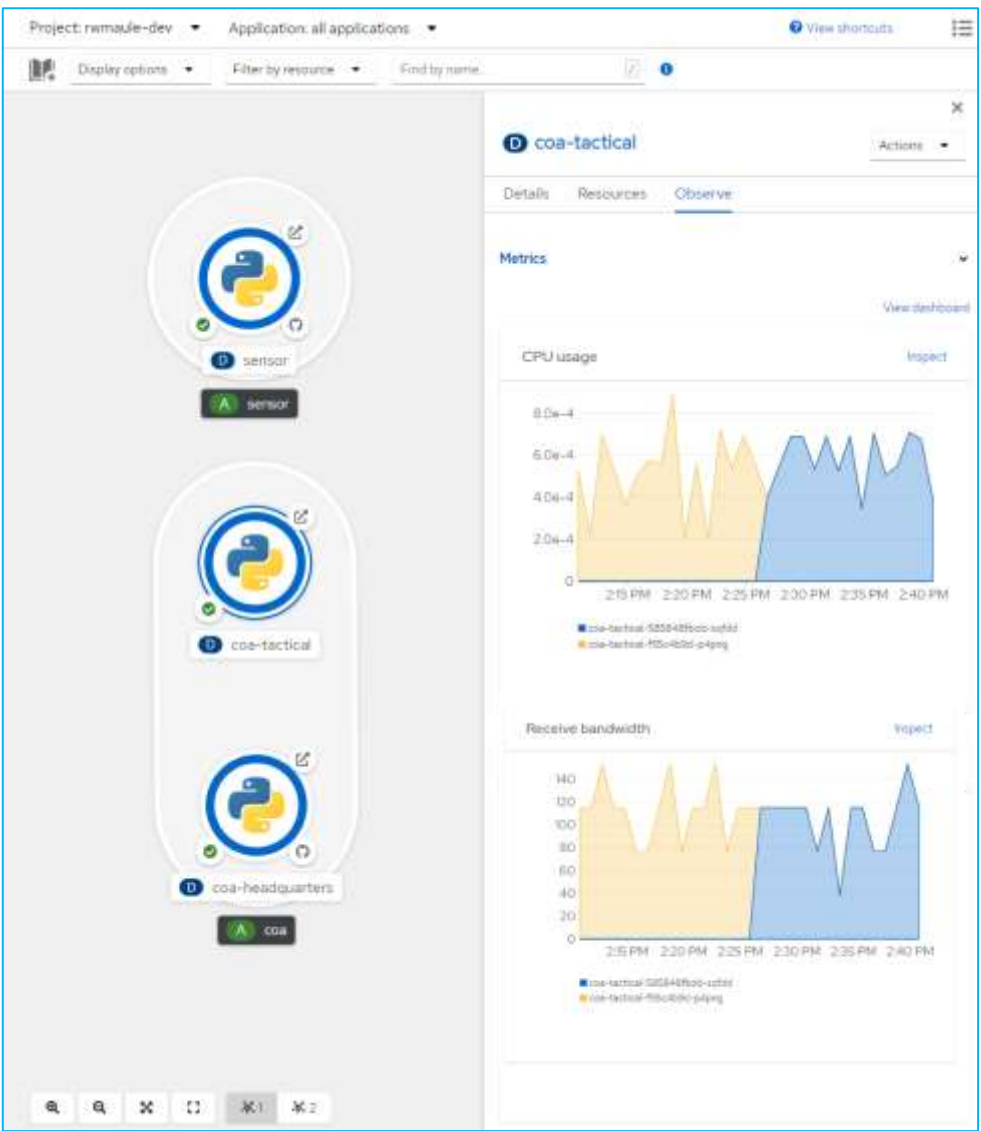


Open Data Hub for AI as a Service (AaaS) on Kubernetes:

- Ceph Object Storage for analytics at the tactical edge
- Inherits upstream from Kafka/Strimzi and Kubeflow
- Jupyter supports interactive data science and scientific computing
- Scikit-learn ML libraries for Python
- Numpy, Scipy, and Matplotlib for predictive analysis
- TensorFlow end-to-end AI/ML to build/deploy ML-powered apps
- PyTorch open source ML framework for computer vision and NLP
- Jupyter notebooks with integrated TensorFlow, PyTorch, and Apache Spark model development frameworks
- IBM Watson Studio for building and managing models at scale
- OpenVINO and oneAPI analytics toolkits for optimizing and tuning models
- Seldon for deploying, managing, and monitoring models
- Starburst Galaxy for data integration



Tactical Edge AaaS Topology



- The Topology view provides a visual representation of the applications within a project, their build status, and the components and services associated with them.
- Pods can be packaged as a container image to run AaaS that can be called from other applications.
- The number of pods for a service can be scaled up or down to increase or decrease the number of instances of the application.
- For serverless applications, the Pods are automatically scaled down to zero when idle and scaled up depending on the channel traffic.
- In this instance we have Sensor Pods, and Course of Action (COA) Pods for the tactical edge and Headquarters.



Node Replica Sets



```

1 kind: Pod
2 apiVersion: v1
3 metadata:
4   generateName: sensor-864646b7cd-
5   annotations:
6     k8s.v1.cni.cncf.io/network-status: |-
7       [
8         {
9           "name": "openshift-sdn",
10          "interface": "eth0",
11          "ips": [
12            "10.128.5.144"
13          ],
14          "default": true,
15          "dns": {}
16        }
17      ]
18     k8s.v1.cni.cncf.io/networks-status: |-
19       [
20         {
21           "name": "openshift-sdn",
22           "interface": "eth0",
23           "ips": [
24             "10.128.5.144"
25           ],
26           "default": true,
27           "dns": {}
28         }
29       ]
30   kubernetes.io/limit-ranger: |-
31     LimitRanger plugin set: cpu, memory request for c
32     memory limit for container sensor
33     openshift.io/scc: restricted
34     resourceVersion: '162672463'
35     name: sensor-864646b7cd-w7qj4
36     uid: f5637221-b4bd-4fa5-bc1a-985239f3a836
37     creationTimestamp: '2022-01-01T19:34:08Z'
38   managedFields: ---
39     namespace: rwmale-dev
40     ownerReferences:
41     - apiVersion: apps/v1
42       kind: ReplicaSet
43       name: sensor-864646b7cd
44       uid: 8a74badf-0ed3-4474-9891-6d2d383db09e
45       controller: true
46       blockOwnerDeletion: true
47     labels:
48       app: sensor
49       deploymentconfig: sensor
50       pod-template-hash: 864646b7cd

```

Streaming events...

- sensor-864646b7cd-w7qj4 rwmale-dev
 Generated from kubelet on ip-10-0-160-156.ec2.internal
 Created container sensor
- sensor-864646b7cd-w7qj4 rwmale-dev
 Generated from kubelet on ip-10-0-160-156.ec2.internal
 Started container sensor

Project: rwmale-dev

Deployments > Deployment details

sensor

Details Metrics YAML ReplicaSets Pods Environment Events

Name Search by name...

Name	Status	Labels	Owner
sensor-75f6667cf8	0 of 0 pods	app=sensor deploymentconfig=sensor pod-template-hash=75f6667cf8	sensor
sensor-864646b7cd	1 of 1 pods	app=sensor deploymentconfig=sensor pod-template-hash=864646b7cd	sensor



Far Edge Data Collection



In the Jupyter notebook we take a small data set of 52 samples and use the Markovify algorithm to simulate a data set of 1000 derived from the original 52 samples.

```
import pandas as pd

pd.set_option('display.max_colwidth', None)
df = pd.read_csv('dataset/sensor.csv')
df.sample(10)
```

	Timestamp	Sensor	Area	Source	Function	Issue	COA
0	2021/12/01 11:46:33 AM CST	HF-1	Shore	Blue	COMINT	Jam	Defend
26	2021/12/01 11:46:33 AM CST	HF-1	Shore	Red	COMINT	Detect	Jam
17	2021/12/01 11:46:33 AM CST	Radar-4	Ship	Blue	SIGINT	Latency	Maintain
12	2021/12/01 11:46:33 AM CST	SATCOM-1	Shore	Blue	COMINT	Jam	Defend

```
def train_markov_type(data, coa):
    return markovify.Text(data[data["coa"] == coa].issue, retain_original=False, state_size=2)

#Function takes one of the 'issue' models and creates a randomly-generated sentence of length
def make_sentence(model, length=100):
    return model.make_short_sentence(length, max_overlap_ratio = .7, max_overlap_total=15)

#built models
defend_model = train_markov_type(subset, "Defend")
jam_model = train_markov_type(subset, "Jam")
maintain_model = train_markov_type(subset, "Maintain")
attack_model = train_markov_type(subset, "Attack")
```

Latency	Maintain
Detect	Attack
Jam	Defend
Latency	Maintain



AI/ML Model



- Characterize sensor status from free text descriptions entered by users through Natural Language Processing (NLP).
- Package the code to create a service that can be queried from an application.
- Train the model on the simulated data, and once trained, enter sensor issues to see if the model has correctly categorized the status.
- Use the TensorFlow AI/ML libraries to run and share the code:
 - 80% training, 20% testing.
- Text entered by warfighters is converted into contextual vectors with numeric representations to form an index.
- Scikit-learn is used to convert label strings into a numbered index to enable the AI/ML algorithms to work with categorical data.

```
[0 0 0 1 1 0 1 1 1 1 0 2 1 2 1 1 0 0 1 0 0 1 1 1 0 2 0 0 0 0 1 2 1 1 2 0 0 2  
0 2 0 2 0 0 1 0 1 1 2 0 2 0 0 0 0 1 1 0 0 1 1 2 0 2 0 0 0 2 1 0 0 2 0 1 2  
1 2 2 1 2 0 1 2 0 2 0 0 1 0 0 2 0 2 2 2 0 2 0 0 1 2 2 1 2 1 1 0 2 0 0 0 0  
0 0 2 2 2 1 0 0 0 0 2 0 0 0 2 0 2 0 0 1 1 1 0 2 0 1 2 0 0 1 0 0 1 1 1 0 1  
0 1 0 1 0 2 0 0 2 2 0 2 0 2 0 1 0 0 0 0 0 0 0 0 1 0 0 0 0 2 0 0 0 1 0 0 0  
2 2 1 0 2 2 0 1 2 0 1 0 2 0 0 1 0 2 1 0 2 2 0 1 0 1 0 1 2 0 0 0 0 1 2 0 1  
1 2 1 0 0 2 1 1 0 2 0 1 1 1 0 2 0 1 2 0 1 0 1 0 2 0 1 1 0 0 0 0 1 0 0 2 2  
0 0 2 0 0 2 0 0 0 2 0 0 0 1 0 0 2 0 2 1 0 0 0 2 1 0 1 2 0 0 0 0 1 0 0 0 2  
0 2 1 2 2 1 0 0 1 2 1 1 1 2 0 0 0 0 2 2 2 1 0 2 1 2 0 0 0 1 1 0 1 0 0 2  
1 2 0 1 2 2 0 0 0 2 2 2 2 1 0 2 0 1 1 0 2 2 0 0 0 0 1 1 2 1 0 0 0 1 1 0  
1 2 0 0 0 0 0 1 0 0 0 2 0 1 0 2 1 0 2 2 1 2 2 1 0 2 0 1 2 2 0 0 0 1 0 0  
0 2 2 0 1 0 1 0 1 0 0 1 0 1 0 0 2 2 2 0 0 0 0 0 1 0 2 0 1 0 0 2 2 2 2 1 1  
1 1 2 1 0 1 0 0 2 0 1 1 1 0 1 2 0 0 1 1 0 0 1 2 0 0 2 0 1 0 0 0 0 0 1 0 1  
0 2 0 0 2 1 0 1 0 0 0 1 1 1 2 0 2 2 0 1 0 0 0 0 2 1 1 0 1 0 2 2 0 0 2 0 0  
0 0 2 0 1 0 0 0 0 0 1 0 0 0 0 2 2 2 0 1 1 0 2 2 0 1 0 1 0 1 2 0 0 2 2 0  
2 0 0 0 0 1 0 0 0 0 1 0 0 0 1 0 1 0 0 0 0 1 0 2 2 2 1 2 0 1 1 0 2 0 0 2  
0 1 1 1 0 1 1 1 1 0 1 2 0 2 0 2 0 2 0 0 2 0 1 1 0 0 0 2 0 1 2 0 1 2 0 1 2  
2 0 2 2 0 0 2 2 2 2 0 1 2 0 2 2 1 0 0 2 0 0 0 0 0 2 0 0 0 0 2 2 1 1 0 1 2  
2 0 2 0 2 0 0 0 2 2 0 0 0 0 0 2 0 0 0 1 0 0 0 2 1 2 2 2 1 2 0 2 1 0 1 2 2  
0 1 1 0 0 0 0 2 0 0 2 0 1 2 1 2 2 1 0 1 2 0 0 0 2 0 2 0 2 0 2 2 2 1 0 2 1  
2 2 0 0 2 1 0 0 0 0 2 0 1 0 0 0 2 2 0 0 2 2 1 0 1 2 2 1 0 1 0 1 0 0 0 1 0  
1 0 0 0 0 2 1 2 0 0 0 0 0 0 0 2 0 0 1 0 0 2]
```



Train and Test



- Softmax calculates probabilities for each category in each document.
- Epochs represent the number of times model calculations pass through the data.
- TensorFlow binary is optimized with oneAPI AI/DNN library.

```
def predict(single_test_text):  
  
    text_as_series = pd.Series(single_test_text) #do a data convers  
    single_x_test = tokenize.texts_to_matrix(text_as_series)  
    single_prediction = model.predict(np.array([single_x_test]))  
  
    single_predicted_label = text_labels[np.argmax(single_predictio  
-----  
    return {'prediction': single_predicted_label}  
  
#####  
#Run the firs time in order to save the model|  
#####  
single_test_text = 'HE-1 is being jammed'  
print(single_test_text)  
  
prediction = predict(single_test_text)  
print(prediction)---
```

```
[('COMINT LTE-2 Ship Blue Latency', 'maintain'),  
( 'IMINT UAV-1 Shore Blue Jam', 'defend'),  
( 'COMINT 5G-2 Ship Blue Latency', 'maintain'),  
( 'SIGINT Radar-4 Ship Blue Cyber', 'defend'),  
( 'COMINT WiFi-2 Ship Blue Latency', 'maintain'),  
( 'COMINT 5G-2 Ship Blue Latency', 'maintain'),  
( 'COMINT SATCOM-2 Ship Blue Latency', 'maintain'),  
( 'SIGINT Cyber-3 Ship Red Detect', 'jam'),  
( 'IMINT UAV-3 Shore Blue Jam', 'defend'),  
( 'COMINT LTE-1 Shore Red Detect', 'jam'),  
( 'IMINT UAV-3 Shore Blue Jam', 'defend'),  
( 'COMINT LAN-2 Ship Blue Latency', 'maintain'),  
( 'SIGINT Radar-3 Ship Red Detect', 'jam'),  
( 'COMINT SATCOM-1 Shore Blue Jam', 'defend'),  
( 'COMINT WAN-2 Ship Blue Latency', 'maintain'),  
( 'COMINT WAN-3 Ship Blue Cyber', 'defend'),  
( 'COMINT WAN-2 Ship Blue Latency', 'maintain'),  
( 'COMINT WAN-3 Ship Blue Cyber', 'defend'),  
( 'COMINT LTE-2 Ship Blue Latency', 'maintain'),
```



Recommend/Predict



Inline Curl recommendation

```
!curl -X POST -H "Content-Type: application/json" --data '{"data": "HF-1 blue is being
{
  "prediction": "defend"
}
```

Embedded Python recommendation

```
import requests
import json
response = requests.post('http://127.0.0.1:5000/prediction', '{"data":
response.json()
{'prediction': 'defend'}
```

HF-1 is being jammed
{'prediction': 'defend'}

latency issue on HF-2
{'prediction': 'maintain'}

COMINT latency on blue ship SATCOM-2
{'prediction': 'maintain'}

detect COMINT red ship HF-2
{'prediction': 'jam'}

red shore LAN-1 detect
{'prediction': 'jam'}

SIGINT jam radar-1 shore blue
{'prediction': 'defend'}

```
Epoch 1/2
23/23 [=====] - 0s 9ms/step - loss: 0.8811 - accuracy: 0.6428 - val_loss: 0.3439 - val_accuracy: 1.0000
Epoch 2/2
23/23 [=====] - 0s 4ms/step - loss: 0.2551 - accuracy: 1.0000 - val_loss: 0.0658 - val_accuracy: 1.0000
7/7 [=====] - 0s 1ms/step - loss: 0.0669 - accuracy: 1.0000
Test loss: 0.06687449663877487
Test accuracy: 1.0
```



Test Cases



Function	Component	Supported	Tested	Result
Orchestration	Kubernetes	Yes	Yes	Recommend
Streams	Kafka	Yes	Yes	Recommend
AI/ML	Jupyter	Yes	Yes	Recommend
AI/ML	TensorFlow	Yes	Yes	Recommend
AI/ML	NLP	Yes	Yes	Recommend
AI/ML	PyTorch	Yes	No	Viable
AI/ML	Spark	Yes	No	Viable
AI/ML	Watson	Yes	No	Viable
AI/ML	Scikit-learn	Yes	Yes	Recommend
AI/DNN	TensorFlow	Yes	No	Viable
AI/DNN	OpenVINO	Yes	No	TBD
AI/ML	Seldon	Yes	No	TBD



Conclusion

Summary and Recommendations



Conclusion



- This project informs DMO, EABO, LOCE, and JADC2 objectives with technical designs for hardware, software, processing, and AI/ML at the tactical edge.
- Hardware was selected to support tactical cloud edge nodes, and software to support hybrid multi-cloud distributed tactical computing in high security architecture suitable for forward deployed forces in D-DIL and challenged EMS and cyber environments.
- Best-of-class industry and government software offering the potential to support an integrated C2/ISR universal COP with legacy and next generation JADC2 sensors and services were evaluated.
- Micro-service mesh/grid, real time streaming architecture, and AI/ML were evaluated for integrated C2/ISR universal COP tactical edge decision support and process automation.
- Future research may continue to refine hardware/software for mobile tactical clouds for extreme edge deployments in challenged environments to support an integrated C2/ISR universal COP with AI/ML services including analytics for enhanced SA, automation, and prediction.