WIND™

Building New Military Infrastructure with Open Virtualization Platforms and Cloud Technologies

By Chip Downing, Senior Director, Aerospace & Defense, Wind River

EXECUTIVE SUMMARY

How can commercial Internet of Things (IoT) and cloud technologies be integrated into next-generation military communications systems? How can vast sensor systems and intelligence systems be built with low-cost hardware and software, and with secure access to cloud analytics platforms with autonomous real-time response controls?

The rapid expansion of commercial IoT offerings is changing the business models of military information systems by creating the opportunity to reduce costs while vastly improving military communications and sensing systems. These new IoT offerings were built with a connected world in mind and have substantial hardware and software security capabilities designed in.

This paper describes how new IoT products can create a secure sensor-to-cloud environment built from commercial off-the-shelf (COTS) hardware and software components. These components range from small, low-power microcontrollers using low-cost software platforms to high-end data analytics platforms using dynamic application insertion strategies like Network Functions Virtualization (NFV). Strategies for provisioning, managing, and decommissioning cloud-connected devices will also be discussed.

TABLE OF CONTENTS

IoT is today's commercial initiative to interconnect a wide variety of technical and commercial information-generating components and to enable new business opportunities based on device and system intelligence. This initiative is the large-scale commercialization of technology that has been developed and proven by the U.S. Department of Defense and Intelligence Community over the past two decades. In much the same way that NASA and the early space program in the 1960s spurred innovations in chip technology, automation, propulsion, and miniaturization, solutions developed for network-centric operations (NCO) in the 1990s and sensor/reconnaissance systems in the 2000s translate directly to the foundations of today's commercial IoT.

Given that IoT concepts originated in the military-intelligence sector, can the commercialization of IoT provide new opportunities back to this sector? If so, how can vendors exploit these opportunities using COTS technologies based on open architectures and cloud technologies from leading companies such as Intel® and Wind River®?

THE NEXT-GENERATION BATTLEFIELD

The armed conflicts that the U.S. has been engaged in for more than 15 years have been bring-your-own-device (BYOD) wars where usable, in situ sensors, embedded devices, and consumer devices were nearly nonexistent—thus offering scant IoT device sensing support in battle zones. Without this sensor support, tens of thousands of discrete platforms—including manned and unmanned aircraft, satellites, and robotic vehicles—were needed to provide situational awareness in these relatively nonindustrialized battlefields.

To facilitate the collection of data, an extensive NCO communications environment was assembled, creating a vast network infrastructure consisting of ground, air, and space assets. This information was delivered back to centralized military information centers for multisource integration, analysis, and subsequent dissemination and distribution, creating an actionable common operating picture (COP). This communication distribution infrastructure equipment required extensive funding. Additionally, this "send it back" system architecture exposed design and security flaws and information-flow choke points that were exploited by potential adversaries (e.g., the compromise of drone video data).

Future conflicts will be different. A major distinguishing feature will be that millions of deployed data-sensing and collection devices will already exist and be in place, none of which will have been purchased nor delivered by the U.S. military or by coalition partners. Thanks to the exponential proliferation of smartphones, personal fitness devices, security cameras, tracking devices, and IoT devices located in almost every vehicle, building, home, and street corner, future war zones will be rich in sensing devices and real-time sensing data. Granted, there may still be areas with poor IoT infrastructure, but the trend is definitively moving toward a rich IoT sensing and data collection environment.

The U.S. military will not have to invest in the underlying IoT infrastructure, as it will have been funded by the local population, businesses, and government. This class of sensors will, in effect, be free. However, getting IoT data from these devices will not be free—it will need to be purchased from the IoT data aggregators or secretly obtained using electronic or cyber warfare techniques. Regardless of the methods utilized to procure the IoT data, investment will be needed. The transmission, analysis, and dissemination of this data will continue to have an associated cost.

FUTURE SYSTEMS WILL BE OPEN AND ADAPTABLE

To operate successfully in these new IoT-rich battlespaces, systems will need to be highly adaptable by design. The ever-changing, ever-evolving IoT landscape cannot be leveraged by static mission systems. Successful systems must be able to accommodate new devices, new technologies, and new security capabilities on demand—in many cases while already on a deployed mission. They must be able to respond to real-time events and to dynamically insert new technologies into nonstop mission and support platforms.

Systems must be based on open architectures to support this rapid change and evolution. Open architectures will enable the U.S. defense industry to request new capabilities, purchase these new capabilities quickly, and insert these capabilities into deployed platforms without disrupting ongoing operations. There are many examples of open common operating environments (COEs) in the military, in addition to open standards for a wide range of next-generation systems. One of the best examples is the Future Airborne Capability Environment (FACE™) approach that couples an open technical specification with an open business architecture.

The FACE technical team created a technical standard-of-standards utilizing more than 60 existing technical standards already proven in commercial and military aviation. They chose the ARINC 653 integrated modular avionics (IMA) standard—proven on multiple Boeing commercial aircraft and more than 80 global commercial and military aircraft—to be the open platform where an ecosystem of suppliers could supply new avionics technologies asynchronously on a shared compute platform by using robust ARINC 653 partitioning. They also selected subsets of the POSIX API standard—proven by global UNIX and communications platforms—to augment this multi-supplier partitioning strategy for rapid insertion of mission and sensor/IoT data applications.

From a business standpoint, the FACE designers quickly recognized that the traditional procurement model, where a prime contractor and its close ecosystem of suppliers maintain a platform for life, created an environment where innovation and rapid insertion of new technology could be slowed by government procurement processes. The FACE business team decided to create a role-based environment where any supplier could assume a particular role in new capability development. Similar to the roles defined in RTCA DO-297 *IMA Development Guidance and Certification Considerations*, which is used in commercial aviation IMA projects, this business model allows for the most efficient supply chain to be selected for any aircraft design or equipment upgrade.

Early in the standard's design, FACE architects thought that the ideal platform would be a platform-and-applications architecture similar to the Apple® iPhone® or Google Android environment. But further review of these popular architectures revealed that both had proprietary layers in their solution stack; therefore, a more open design was chosen.

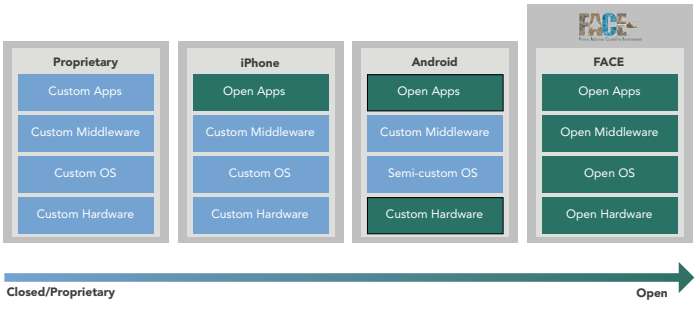

Figure 1. The evolution of closed to open architecture designs with the FACE approach

The best part of using an open standard and open architecture such as FACE is that the architectural design is free. Program managers do not need to invent yet another layered architecture model for their programs. Like many open standards that encourage the use of the standard with minimal startup costs, the FACE standards documents are available for free from the FACE website: www.opengroup.org/face. Additionally, there are no royalty or other charges for using FACE technology in military designs.

OPEN VIRTUALIZATION

Another must-have capability in next-generation, software-based military systems is virtualization. Virtualization enables the use of multiple application and operating system environments on a shared compute platform by abstracting away the exact computer architecture from the applications, thus removing underlying hardware and software dependencies of both new and legacy applications. Virtualization enables the use of a single compute platform to be used by multiple applications from different domains and different suppliers. This allows a new application to use an existing open virtualization platform at no cost.

This commoditization of computing hardware decouples legacy hardware/software dependencies and allows for rapidly repurposing mission platforms. It also enables systems at the edge to have dynamic "on-the-fly" information technology (IT)/operations technology (OT) functions—ideal in situations where there is a rapidly changing security or warfare environment, or where there is a high failure or degradation rate of capabilities.

For a virtualization platform to have the highest utility, it must allow applications using a broad range of commercial and proprietary guest operating systems to run without penalty. In addition, virtualized systems should enable the continued use of legacy software applications while combining them with new capabilities in new operating environments.

NETWORK FUNCTIONS VIRTUALIZATION (NFV) AND SOFTWARE DEFINED EVERYTHING

Systems having open architectures, robust virtualization, and system partitioning open the doors for new software paradigms for network and application management. Key benefits of NFV for the telecommunications industry include faster infrastructure

deployment, accelerated provisioning of new services, rapid scaling of resources, and reductions in both CAPEX and OPEX.

In NFV, software-based virtual network functions (VNFs) run on one or more virtual machines (VMs) and are chained together to create communications services. NFV solutions consist of four core components:

1. A VNF layer executing virtualized communications services
2. An NFV infrastructure (NFVI) layer that includes the virtualization foundation (hypervisor), high availability (HA) plugins, virtualized infrastructure management (VIM), telemetry, and middleware
3. A hardware layer that consists of server-class processors, such as Intel Xeon® platforms, along with network and storage hardware components
4. An NFV service orchestration capability

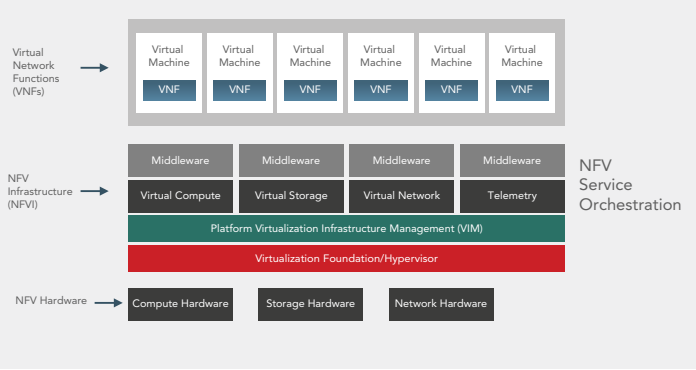

Figure 2. Generic architecture of an NFV system

The same principles that underpin NFV also enable the expansion of this virtualization strategy into other use cases, like Software Defined Everything (SDE or SDx) technologies, an umbrella term that includes Software Defined Networking (SDN), Software Defined Computing (SDC), and Software Defined Storage (SDS). This same IT concept can then be expanded to include operational technology (OT) environments that also consolidate legacy platforms onto a shared compute platform. In all SDE cases, the computing infrastructure is virtualized and delivered as a service, and management and control are automated by software rather than by the hardware components of the infrastructure.

Implementation of NFV in network elements requires a range of critical elements, including:

- Hardened Linux (or other server operating system)
- High-performance virtualization
- Data plane acceleration, including high-performance virtual switching
- Integration with OpenStack or other open VM orchestration tools
- High availability on the order of six-nines (99.9999%) uptime, with at most 31.5 seconds downtime per year

Availability is a key component of a platform virtualization strategy. High levels of availability (i.e., above five nines) cannot be achieved by singular NFV platform elements. It can only be achieved with robust interoperability between the platform hardware (including hardware virtualization IP), dynamic software virtualization layers, operating systems, failover software, function services, and corresponding processes. These high levels of availability, with only seconds of outage per year, define and differentiate today's highly reliable, dynamic virtualization solutions.

Table 1. Maximum downtime per year to achieve a given level of availability

Availability Chart	
Availability %	**Downtime per Year**
90% (one nine)	36.5 days
99.0% (two nines)	3.65 days
99.9% (three nines)	8.76 hours
99.99% (four nines)	52.56 minutes
99.999% (five nines)	5.26 minutes
99.9999% (six nines)	31.5 seconds
99.99999% (seven nines)	3.15 seconds
99.999999% (eight nines)	315.569 milliseconds
99.9999999% (nine nines)	31.5569 milliseconds

Virtualization abstracts hardware that allows multiple workloads to share a common set of compute, network, and storage resources. In NFV environments, a variety of network function workloads can colocate on shared hardware while maintaining full isolation from each other. In addition, the VNFs can migrate across infrastructures and scale as required. Virtualization also improves server utilization and consolidation and makes on-demand self-provisioning of services and software-defined orchestration of resources possible.

Open virtualization platforms must include hardware assist virtualization technology to support unmodified guest OS execution in VMs. This hardware virtualization assist IP—such as Intel virtualization technology (VT)—complements virtualization software/hypervisors and reduces size, cost, and complexity while improving security with hardware-enforced separation.

Fully utilizing hardware virtualization assist capabilities fundamentally changes computing at the edge and affords many benefits:

- Hardware virtualization assist enables the support of unmodified guest OSes, including IT/enterprise operating systems such as Windows and Linux.
- Hardware virtualization assist moves the responsibility of partitioning and application separation from traditional software-plus-memory-management-unit (MMU) solutions to a hardware virtualization layer that manages the processor core.
- Hardware virtualization assist on multi-core processors enables hardware-controlled separation of hardware resources (e.g., cores, memory, devices, IP) for increased safety and security.
- Unmodified guest OS support removes both cost and risk from porting legacy platforms to new execution environments.
- Unmodified guest OS support enables the rapid, dynamic insertion of new capabilities into multiple platforms.

Moving data efficiently through systems is also a key capability. This integrated hardware/software solution needs to increase the performance of layer 2 through 4 packet processing while maintaining or increasing security. On multi-core Intel processors, this is accomplished by:

- Leveraging the Data Plane Development Kit (DPDK) to accelerate data plane traffic
- Utilizing Intel's high speed and scalable IPv4 and IPv6 forwarding that can support over 10 Mpps per core
- Using accelerated IPsec and IKE stacks that support over 200 Gbps throughput over tens of thousands of tunnels using GRE, GTP, L2TP, MPLS, PPP, VXLAN, and other network industry tunneling standards
- Implementing accelerated TCP/UDP stacks that can support over 100 million concurrent sessions with session setup rates of 5 million sessions per second

All of these NFV and SDE platforms can run software from leading network infrastructure, network function, and critical infrastructure suppliers using standard COTS enterprise platforms and ruggedized embedded hardware platforms. Today's rapid pace of silicon innovation coupled with support for efficient, open virtualization in the next-generation silicon is quickly driving compute systems designed for enterprise IT into smaller edge devices along with embedded systems and sensors—the OT domain. This enables the easy migration of enterprise-class server applications such as NFV and SDE to OT edge and embedded devices. This migration drives the feasibility of tactical clouds and real-time analytics on forward mission systems.

Using today's advanced multi-core processors also adds incremental hardware-based security in three significant ways:

1. The ability to place different application security domains on separate processor cores provides for immediate execution segregation.
2. The use of hardware virtualization assist creates robust, hardware-controlled virtual machine containers for each application/OS environment, ensuring that each domain is controlled by hardware separation capabilities instead of by the MMU and software separation kernels.
3. The use of hardware-based security creates incremental layered security strategies.

CLOUDS

Once the data center and embedded platforms have a foundation of virtualization established, cloud computing is the next logical step. Cloud computing is a deployment model for enabling on-demand network access to a shared pool of configurable computing resources—including networks, servers, storage, applications, and even deployed devices—that can be rapidly provisioned, utilized, and then released for other services, all on demand. The pool of configurable resources typically requires minimal management effort by the cloud service provider but can provide significant economies of scale. Operational and management costs of larger cloud systems can be lowered up to 60% over traditional IT deployments.

There are some obvious benefits of cloud computing:

- No need to spend time and money housing, powering, cooling, maintaining, and protecting the infrastructure
- Limited upfront capital costs by shifting infrastructure costs from CAPEX to OPEX
- Paying only for the resources you use
- Ability to scale resources up or down on demand
- Faster deployment with rapid provisioning by infrastructure experts and automated provisioning
- Ability to focus labor costs on increasing capability versus maintaining IT infrastructure

CLOUD SECURITY

The number one concern for commercial or military entities regarding cloud technology is security. When military organizations consider cloud computing, they initially gravitate toward on-premises clouds for reasons of security, control, and trust. But these fears may be ill-founded, based on the assumption that their IT organization has better security experience than large cloud service providers like Amazon, Microsoft®, IBM, and Google. In fact, with so much at stake for large commercial cloud service providers, their security has to be some of the best in the world. These providers can also be bound by precise service level agreements that include location and security.

There are some clear security advantages to cloud computing and storage, including storage, including the following:

- Systems and data are in virtualized environments, constantly moving, and rarely in the same execution environment, making them harder to target for attacks.
- All data-in-motion and data-at-rest have high levels of encryption.
- Cloud providers have redundancy capabilities in place.
- Cloud providers have disaster recovery plans in place.
- Cloud providers secure an entire facility, not just a room (e.g., USB drives and small disk drives are prohibited to prevent interference, infection, or the stealing of information).

CLOUD-CONNECTED DEVICE MANAGEMENT

With the proliferation of network-connected devices in the OT arena, the ability to deploy, monitor, manage, service, update, and decommission those devices becomes vital. Cloud-connected equipment and devices need to be actively managed to provision platform software; maintain security and install security patches; and collect, manage, and aggregate machine-generated data to maintain peak operational readiness and efficiency. Management systems like Wind River Helix™ Device Cloud can perform the following essential services:

- Easily collect, store, manage, and integrate data from disparate devices, machines, and systems
- Protect data-in-transit with a secure, scalable, and customizable on-demand encryption infrastructure
- Remotely configure, monitor, and update connected machines from a single management console
- Streamline the management of hundreds of thousands of IoT devices for provisioning, configuration, and decommissioning

These capabilities solve the problem of connecting and managing devices remotely, diagnosing situations, and forwarding data to enterprise systems for analysis and long-term storage. In addition, complex IoT and sensor systems can scale to mission demands.

MOVING FORWARD WITH TACTICAL CLOUDS

The most common view of cloud technology is that it is a part of an IT environment. The primary challenge for the military and intelligence community sector is delivering the ever-increasing volume of data generated by their field sensing systems into a usable cloud environment—typically in or near a military operations center. This intelligence is controlled by a tasking, collection, processing, exploitation, and dissemination (TCPED) process, incorporating a "send it back" data communications model where field operations data is sent back via beyond-line-of-sight (BLOS) satellite communications (SATCOM) channels to service and agency clouds. There has been large growth in the use and adoption of automated data processing and decision support tools to try to unblock the TCPED logjam, but the growth of sensor data-generating assets and the increasing demand for speed of action have indicated the current architecture may be losing the battle for efficient and reliable information management.

One of the core inefficiencies of the current sensor-to-cloud architecture is that about 75% of the data collected goes straight to archive, with no current operation using the data. Creating local combat/tactical clouds near the forward field of active missions that can deliver real-time intelligence directly to the operations personnel will potentially reduce this data communications bandwidth bottleneck as well as provide better real-time situational intelligence to operators in the field.

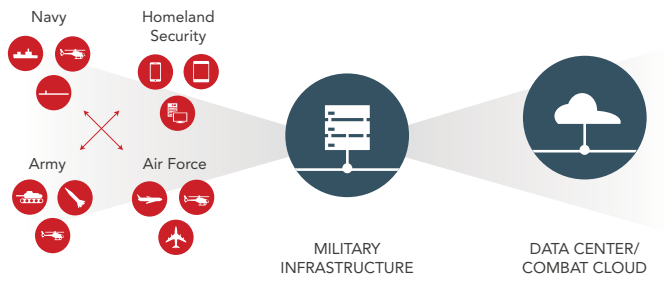

Figure 3. Single cloud view of data center and combat/OT clouds

Next-generation OT/tactical/combat clouds will have powerful analytics engines to aggregate, integrate, and refine sensor data and fuse them to other relevant mission data in real time. Combining this compute power with dynamic NFV and SDE capabilities will drive huge operational advantages, with a greater ability to export both data and assets in the field for joint operations, providing all connected entities with a real-time COP. This will also rapidly create an enhanced environment for dynamically enabling coalition partner operations.

NFV/SDE IT/OT edge devices will most often utilize line-of-sight (LOS) communications, which are more diverse and ubiquitous in forward-deployed systems. These LOS communications can handle higher volumes of data with a greater selection of data paths, and they are mostly immune to adversarial disruption with increasingly low probability of intercept/low probability of detection (LPI/LPD) characteristics. Using this local communications backbone, the future of collective military computing is enabled with locally provisioned, multilevel secure (MLS) data sharing, virtualized applications systems, and automated data analytics tools. Multiple devices can enter and exit the cloud, optimized for the collective capabilities available at any given moment to support missions on demand.

TACTICAL CLOUDS INCREASE SOLDIER SPEED AND AGILITY

Traditional sensor systems were highly dependent on a platform and a network that aggregated the sensor data and then sent it back to a rearward operations center for analysis (i.e., TCPED). With commercial IoT sensor packages now easily connecting with OT and tactical cloud systems, the architecture of data collection, aggregation, and intelligence dissemination can be moved closer to the field of operation, increasing the speed of situational awareness and reducing the demand for BLOS data communications. Most of these tactical/OT cloud systems can operate in a local field of vision with LOS communications, reducing the demand on bandwidth-constrained BLOS communications platforms and adding communications architecture robustness and continued operations in degraded environments.

CONCLUSION

In the IoT era, consumers are realizing the benefits—and businesses are monetizing the intelligence—gained from technologies tested and proven in the military and intelligence community over the last 20 years. This commercial investment in COTS IoT and cloud technologies can drive huge cost savings for next-generation military services and security agency systems. With commercial investments from industry leaders like Intel, Wind River, and other commercial IoT and cloud providers, the military communications and intelligence community can now reap the benefits of transforming its systems into the next generation of high-value, network-enabled, software-defined solutions increasing the knowledge, speed, and utility of future military communication systems.

WIND

Wind River is a global leader in delivering software for the Internet of Things. The company's technology is found in more than 2 billion devices, backed by world-class professional services and customer support. Wind River delivers the software and expertise that enable the innovation and deployment of safe, secure, and reliable intelligent systems.