

COTS Enables Low Cost Military C2 and IT Systems

Ron Busch and James A Lohr, *Member IEEE*, both of
DRS C3 and Aviation Company
400 Professional Drive
Gaithersburg, MD 20879 USA
busch@drs-c3a.com; jalohr@drs-c3a.com

Abstract - This paper discusses how Commercial-Off-The-Shelf (COTS) equipment and software, and strict compliance with Open Architecture principles, can be used to satisfy these unique military needs of associated C2 and IT systems. This discussion includes how COTS meets this goal and how it can be configured to also meet the ruggedization demands for operation in typical military environments, including at sea; and, satisfy the requirement to be battle ready and survivable allowing our military to be functional and be ready for action when called upon. The paper also discusses the selection of an Open Architecture (OA) framework for implementation, selection of COTS equipment that fits within the OA framework, interchangeability, and methods of cost reduction and life extension. Critical to success is the need for modularity such that updating one element (memory, processing, networking, storage, communications, etc.) will not affect or ripple into other elements. Compliance with common and popular industry standards and practices is discussed and its importance to a successful fielding is explained. A case study is presented of one Shipboard qualified system built from COTS elements.

I. INTRODUCTION

Development timelines are typically challenging. When military procurement time, testing times, qualification time, and then deployment times are added ... this development and fielding cycle can potentially take years. Tack onto this a typical military systems expected life (or at least hoped for life) of one or two decades. We want our military to have a competitive advantage. They need their systems to work when called upon. How do we maintain our systems at a state-of-the-art that is better than our adversaries, save the lives of our soldiers and sailors, and protect our freedom and way of life? How can these objectives be accomplished in the face of shrinking capital budgets; faster product evolution in IT and C2, faster pace obsolescence; and decades of system life; while maintaining a technological edge over those who oppose us? COTS provides a feasible solution.

Commercial-Off-The-Shelf (COTS) equipment gives us the tools and products we need to meet military objectives within the constraints identified. COTS supports the rapid change that can keep our military systems performing at a peak levels. COTS development funded by the commercial sector with performance driven by consumer applications, expectations, and market demand, augments our shrinking military budgets. COTS supports meeting the evolving military system demands.

While it is easy to conclude that COTS must be leveraged, how can COTS be applied to the needs of military procurement, qualification, training, and support? This paper provides some insights into how to make COTS usable in the military product life cycle; how to enable COTS to survive in the environments in which the military must operate; and how to architect COTS systems to meet the long supportable system life required of military systems.

II. MANAGING COTS

To apply COTS to military systems and to survive the rapid change and growth of technology associated with COTS requires several different techniques:

1. The principals of Modular Open Systems Architecture (MOSA) must be applied to the system. These principals are outlined in Fig. 1.
2. Applications of Standards at all levels of the design. One of the key factors in long term system success is compliance with standards. Standards exist at all levels of the system and product life cycle. Fig. 2 summarizes some of the standards employed at different levels of system design and development.

Plan for Change. It is not a matter of IF COTS will change, but WHEN. Planning for change is a basic requirement associated with COTS based systems and must be addressed during selection of components and supply chain management.




Fig. 1. MOSA Principals guide the application of COTS




Fig. 2. Application of OA, DOD, and Industry Standards at the component, protocol and Middleware levels enables COTS based solutions that abide by the MOSA principles

A. COTS Selection Criteria

Trade studies guide the selection of a specific COTS solution. One of the key attributes in a trade study is identification of the “sweet spot” of a COTS element on the product maturity curve.

Fig. 3 portrays that the optimal affordability strikes a balance between the availability and maturity of the COTS product. Emerging technologies have a high purchase cost and minimal support. Products in the declining stage have a high maintenance cost and lag in performance. Products in the optimal range offer a reasonable purchase cost associated with mainstream demand and are complemented by readily available support and training. [1]

A trade study will also help identify the best form factor that is applicable to a system element. For computing systems, this may be an ATX single board system, a VME rack based system, a 1U based system, or a Blade based systems. The trade study analysis can be a simple multi column and row table, or a complex multi-weighted formal Kepner-Tregoe. The analysis output will help system designers define the best solution for a given application. Once the basic architecture and form factor is determined, other trade studies can be performed to select the lower level elements, such as the best CPU, best graphics, best storage solution, etc. One benefit of trade studies is that they provide defense and quantification of the design selections and typically a wider range of possible options are considered when trade studies are employed.

B. Management of Change

The management of COTS is the management of change. COTS will change. Therefore, accepting change as inevitable and planning for it is essential to success. But as highlighted in Fig. 4, different technologies change at different rates. It is important to evaluate the technologies in your system, and map the associated cycles. This will assist in the planning of technology insertion and technology refresh cycles that provide long term support of COTS based solutions.




Fig. 3. Finding the Sweet Spot of COTS products.

III. NETCENTRIC DESIGN

A typical Command and Control (C2) system is comprised of four subsystems (see Fig. 5), which are presented in the following sections:

- A. Network Fabric
- B. Computing Platforms associated with hosting C2 applications and distribution of tactical data
- C. Human Machine Interface – which supports the operator interaction with the C2 applications through visualization and operator input (keyboard, trackball, touch input devices)
- D. Sensor Acquisition and Distribution

With the current advances in COTS network architectures and switching technology, C2 systems have been transformed from very platform specific implementations to net centric environments that support compliance to the Open Architectures. Attributes of this transformation are highlighted in Fig. 6.

The common advantages supported by these COTS technologies include:

- Reduction in the number of components and configuration items, thereby yielding enhanced reliability, maintainability, and system availability (Ao)
- Reduction in Size, Weight and Power (SWaP) of the C2 system
- Increase system capability




Fig. 4. Mapping the component change cycles facilitates planning for technology insertion and technology refresh events.




Fig. 5. Components of a typical Command and Control System




Fig. 6. Transformation of C2 System from specific Application to OACE Compliance

A. Network Fabric

Network Switch technology and Data Transport Services provide the network fabric that supports connectivity between the components of the C2 System. A typical architecture of the network fabric employs a distributed network that leverages Core, Distribution, and Edge Switches.

The Core switch provides a central high speed switching infrastructure based on standard IEEE 802.1 standards for Bridging, Virtual LANs (VLANs), Trunking, Port Aggregation, and Topology Control. With proper design the Core is easily distributed physically to different parts of the system, e.g. a ship platform, while logically remaining one entity.

The Distribution Layer provides the high speed routing infrastructure that supports the Data Transport Services and provides the connections to off-ship or out-of-network services. The distribution level also provides the control for wireless access points, when used in areas that need temporary network services or where wireless networking is the best networking option. Current COTS network switches now readily support 10 Gbps transfer rates and Port Aggregation. This allows the network architect to leverage 10 Gbps, and higher, network back bones to eliminate oversubscription issues amongst the Core, Distribution and Edge switches while providing full speed connections between compartments of a distributed network. A significant reduction in cabling and associated maintenance is realized with the application of high bandwidth network backbones.

Edge switches support the scalable connectivity to the end users. Higher port densities of current edge switches support redundant connections to both the central Core switches and user nodes that require fail-over redundancy. Edge switches also offer support of Power over Ethernet (PoE) which passes electrical power safely, along with data, on Ethernet cabling. PoE is especially useful for powering remote Ethernet switches, embedded computers, thin clients, wireless LAN access points, and cameras with pan tilt and zoom (PTZ). PoE provides additional flexibility to configure user ports to optimize power budget and specific node requirements thereby optimizing the cost of the network architecture.

Using this distributed model, proper architecture, and judicious implementation, it is possible to build a network that has the characteristic of Self-Similarity. For example, the network could be subdivided several times, as a result of battle damage, and could still survive as several full, although disconnected, networks. One of the key advantages is that the ship would still be able to fight as a unit or several disconnected units, and would be able to fully track, identify and prosecute targets assuming that a sensor, a console, a C2 processor suite and a weapon survive in one or more of the Self-Similar networks.

Wireless networks offer an alternative to the prohibitive costs associated with cable plant upgrades to remote locations, especially where only low bandwidth connections are required. Wireless Access Points combined with Wireless

Intrusion Detection in a combined appliance provides an economical node that supports Information Assurance requirements.

IEEE 802.1 Networking standards support creation of Virtual LANs (VLANs) which can help to provide a logical separation by function and risk within the network fabric. VLANs enhance the IA posture by creating logical isolation

within a distributed Layer Two network fabric. Reconfigurations of specific user LANs can be applied to the network fabric with the need for additional hardware investment. Subnets for different types and classes of user nodes can be created, controlled, and modified as the needs for the network fabric change over time and application.

B. Computing Platform

Aided by the advances in processing technology driven by the telecommunication market, the emergence of 1U servers provided a cost effective alternate to computing platforms based on VME. As computing requirements for the telecom industry and typical C2 systems have grown, the 1U server form factor has reached limits associated with the required computing and power densities.

From the COTS domain, the Bladecenter architecture has evolved to meet the restrictions imposed by the 1U server form factor. While IBM™ was a fore runner for the Bladecenter, other commercial vendors such as Dell™ and Cisco™ have brought similar systems to market. In addition, the ATCA standard has also addressed part of the market space.

The main advantage to the Bladecenter architecture is the reduction in the infrastructure complexity of the computing platform. Individual processing elements in the form of blades replace individual 1U servers. LAN connectivity is provided within the Bladecenter chassis thereby removing the need for large number of cables for network connectivity. Network switch modules within the Bladecenter also remove some of the requirements for separate physical network switches. Control of the independent processing elements is supported by an integrated management module which removes the need for separate KVM appliances for user control and multiple server management tools. Along with increased compute densities, the Bladecenter offers efficient power distribution. Fig. 7 provides an example that highlights the efficiencies gained with the Bladecenter.

The resulting conservation of energy, heat and physical space are key attributes for implementation of military systems that are uniquely constrained in these domains.

Since the Bladecenter provides a consolidated physical system it is well suited for a virtualized environment that aggregates the physical system resources into shared pools from which virtual machines access virtual resources. Dynamic resource management, fail over support, and redundancy features provided by the virtual machine environment and enabled by the Bladecenter architecture assist in support of a self-repairing C2 infrastructure.




Fig. 7. Efficiencies gained with a BladeCenter approach as compared to individual 1U Servers

Distributed storage is another key technology for a self-repairing C2 infrastructure. Some amount of data, even if it is just virtual images of key C2 software must be stored across the network to be able to resume functionality after battle damage. Two of the key concepts are data durability and consistency; the ability to keep the data in many places (durability) and the same in all places (consistency). Modern storage systems have greatly reduced the cost of mass storage, while making the storage devices more robust by adding advanced anti-shock mechanisms or removing mechanical systems altogether. As with the virtual machines, a control plane would be the mechanism to provide the “virtual” data to requesting clients.

C. Human Machine Interface

Traditional implementations for the Human Machine Interface (HMI) for C2 systems require a computing platform at each user node for visualization and interaction with the associated C2 applications. While the server technology allows for a fairly efficient local computing platform in terms of physical size and power, the ruggedization requirements demanded by the military domain incurs a cost for hosting a computing platform at each user node.

In the COTS domain, virtual machine environments have been implemented to address large number of user interfaces typically associated with data processing centers.

The PC-Over-IP[®] (PCoIP[®]) protocol was developed by the Teradici Corporation to provide an IP addressable interface between the computing platforms hosting applications and the end user. The PCoIP protocol compresses, encrypts and encodes the entire computing experience at the host and

transmits it 'pixels only' across a standard IP network to stateless PCoIP-enabled desktop devices. It is frequently referred to as a “zero client” because it requires no application OS, no drivers, no local storage, and no need for antivirus/spyware. [2]

PCoIP can be implemented in either software or hardware. Though the graphics rendering requirements typical of a C2 application may limit the implementation to hardware PCoIP encode and decode.

Fig. 8 depicts an implementation of the PCoIP in conjunction with a consolidated BladeCenter computing platform. Each user application and associated graphics rendering are performed within an individual blade server. The “desktop” associated with each user interface is encoded by the PCoIP interface and connected to the user via the network. A decoder at each user interface provides connectivity to the typical HMI devices for video, audio and USB. The net result is that at each user node the required compute platform is reduced to a small form factor decoder with a typical power requirement of 25W.

With the user application efficiently implemented in the BladeCenter, the local requirements for SWaP at the “stateless” user node are dramatically reduced. In addition, the security posture at the user node is greatly enhanced in that only HMI data, and not application data, is accessible. USB devices can be configured for authentication, and restricted HMI devices can be locked down.

D. Sensor Acquisition and Distribution

The distribution of sensor data can leverage the netcentric environment if the sensor acquisition can be made IP compatible. A radar sensor is typically integrated in a C2




Fig. 8. HMI supported by PCoIP connection to the host computing platform

system and may require distribution of both radar video and processed radar data, i.e. track data.

In legacy systems the radar video was distributed to the user node as analog data via a point – point connection. At the user node the analog radar video was digitized and scan converted for display as a graphical rendering.

To leverage the netcentric environment, the radar video must be digitized, typically near the actual sensor, and

published across the network fabric. Client applications subscribe to the published digital sensor data and perform the scan conversion, i.e. the conversion from polar to raster format. Typically the Graphical Processing Unit (GPU) on the standard graphics card is harnessed to perform the scan conversion, thereby eliminating the need for any specialized hardware. A block diagram of a sensor distribution is shown in Fig. 9. Sensor distribution via the network typically requires




Fig. 9. Graphical Processing Unit (GPU) is harnessed to perform the scan conversion

modest bandwidth for real time full fidelity radar video and eliminates dedicated, therefore expensive, distribution cabling. Use of publish/subscribe data distribution mechanisms allows a wide range of access to sensor data, thereby increasing system capability.

IV. MILITARIZATION

In order to survive in naval environments the COTS equipment must be protected from:

- Shock and Vibration,
- Rain/drip and Salt Spray,
- EMI emissions, both radiated and conducted
- Ships power

A. Shock and Vibration

Military Specification MIL-S-901D is a military specification for High Impact mechanical shock which applies to equipment mounted on ships. Two levels of shock survivability apply:

- Grade A items are items which are essential to the safety and continued combat capability of the ship;
- Grade B items are items whose operation is not essential to the safety and combat capability of the ship may not become a hazard to personnel, to Grade A items, or to the ship as a whole as a result of exposure to this high level of shock.

A Barge Test is qualification testing normally performed on the subject device situated on a floating barge where explosive charges are detonated at various distances and depths to impart high shock to the equipment under test. This shock is designed to simulate a torpedo, missile, or mine striking the ship. COTS equipment must be protected with a shock isolation system that ideally attenuates a MIL-S-901D shock event to 15 Gs. Most COTS products can withstand shock between 15 and 30 Gs. This isolation design should also eliminate ship vibrations from passing to the equipment. If not properly designed, the shock isolation mechanism may actually amplify the vibrations. Hence it is important to perform modeling of the shock isolation system, the ships resident deck frequency, and the physical characteristics of the equipment.

There are three principal methods of shock isolation:

- Coiled wires
- Elastomer mounts
- Piston based shock isolators

A coiled wire mount is inexpensive and reliable if properly sized to match the weight of the equipment. It is, however, subject to “crimping” which would require it to be replaced after a shock event. Replacement is permitted by the MIL Specification, but requires additional stock of isolators on-

board. Also, the equipment may not be protected in the event of subsequent shock events.

Another popular shock isolator is an elastomer mount. This is just as effective as the coiled wire and it does not crimp, so it is reusable after a shock event. Its disadvantages are limited sources of supply, long delivery lead time, and moderate cost.

The third class of shock isolators are piston based, similar to those found on automobiles. They are very effective, can be used repeatedly without damage, but require unique mechanical designs of the equipment and of the shock pistons themselves, which must be matched to the equipment weight and shock levels expected. This makes piston based isolators very expensive from both a Non Recurring Engineering (NRE) and Recurring cost points of view. Their advantage: they save space, do not fail after a single shock event, and need less “sway space” allowance around the equipment than the other types of shock isolators.

B. Rain, Drip and Salt Spray

Rain/Drip and Salt Spray can be protected against by enclosing the entire equipment bay inside a sealed compartment, and using gasketing on all access doors and panels.

To allow air to enter/exit the enclosure wave guide beyond cutoff honeycombs are used in all openings allowing air to flow as freely as possible, while attenuating EMI radiation. When these honeycombs are set at an angle of about 45 degrees, drip and rain are prevented from entering the enclosure. Primer and paint of the right quality protect the exposed finishes from Salt spray.

All openings and seams of the external enclosure must be gasketed to prevent rain and other moisture from entering the cabinet. This gasket can serve the dual purpose of making the enclosure watertight, and EMI tight preventing electromagnetic radiation from escaping the chassis. This applies to all connectors and switches that go through the external wall of the cabinet.

C. EMI Protection

All signals, including power lines, Ethernet lines, USB interfaces and any other signal lines must be filtered to prevent EMI conducted radiation. All seams and openings, areas around doors and access panels must be used with compressed gaskets to prevent EMI radiated leakage.

D. Ships Power

MIL-STD 1399-300 defines ships power requirements. Power conditioning is usually required for operation with COTS products as well as (for mission critical applications) an uninterruptable power source (UPS). These devices are often packaged together. Available ship platform power is typically 440V AC, three- (3) phase or 115V AC single- (1) phase.

These input variants combined with requirements of MIL-STD-1399-300 typically requires power subsystems specifically developed for military use.

Key requirements of MIL-STD-1399, Section 300 are listed in Table 1. [3]

Table 1. Key areas of MIL-STD-1399, Section 300

Nominal Line Voltage	+/-5%
Nominal Line Freq (60Hz)	+/-3%
Transient Line Voltage (2 sec)	+/-20%
TransientLine Freq (2 sec)	+/-4%
EMERGENCY Line Voltage (2 min)	-100% to +35%
EMERGENCY Line Freq (2 min)	-100% to +12%
Total Harmonic Voltage	5%
Max Single Harmonic Voltage/Current	3%
Power Factor	.80 lag to .95 lead
Voltage Spike (50us)	2500
Inrush /Nom Curent Ratio	10:1
Line Current Unbalance (max)	5%

Notice that the transients are extreme and power conditioning will be required in front of essentially all commercial equipments.

V. EXAMPLES OF ARCHITECTURE (C2 AND IT)

USN’s Repair Station Console (RSC) program is an example of COTS applied to a military application. After the attack on the USS Cole, the Navy upgraded their damage

control repair stations with a modern, COTS based system. The RSC is designed to survive at sea. It is cocooned in a shock isolated rack, that also protects the equipment from power variations, EMI threats, and vibrations found onboard DDG-51 class destroyers. The general architecture of the six-compartment RCS is shown in Fig. 10.

The RSC is built with COTS products, following well accepted standards supported by multiple vendors. The CPU chip and the graphics module are two of the most volatile technologies used in this system with product life times of about 2 years, but only a 1 year production sweet spot. It was therefore important to base the design on COTS standards that will keep up with this rapid product cycle. The ATX motherboard form factor was selected, along with PCIe bus plug-in module. In addition to graphics and Ethernet controllers, the ATX standards define standard power voltages and connectors, plug-in memory modules, BIOS, disk interfaces, and USB interfaces. This openness will allow us to support the RCS for years to come. The Intel i7 quad core CPU provides the needed computational power.

In order to cocoon this system a two level isolation technique was used. The ATX based compute platform was housed in a slide mounted chassis (DRS’ Genesis chassis) with shock isolators under the mounting points of the ATX card. Expanded cooling fans were utilized to overcome restrictions induced by EMI honeycomb in the main shock isolation rack. The packaging concept is shown in Fig. 11.




Fig. 10. RSC Functional Block Diagram




Fig. 11. Genesis rugged chassis with shock isolated CPU motherboard and expanded cooling fans

For the 20" monitor a COTS ruggedized unit that was already mil qualified was selected. This needed no additional protection. Because of the large glass area, the 40" monitors pose unique issues associated with the shock requirements. Since the 40" monitors were only available with a Grade B rating a secondary shock isolation mount was created to attenuate the shock input to Grade B levels.

Ethernet switches were mounted on frames behind the front panel and connected to the compute chassis, and the rear I/O panel on the RSC Rack. This EMI protected them, and shock isolated them. Fixturing for standard Retma or EIA-STD-310 mount, allows the switches to be upgraded with any 1U units meeting this widely available 19" rack mount standards.

All of the above elements are housed in a shock and vibration isolation rack, with EMI protection and filters. This rack is shown in Fig. 12.




Fig. 12. RSC Shock Isolated Equipment Enclosure housing COTS based components

VI. CONCLUSION

COTS provides unique advantages for application in military systems, the most prevalent is the ability to provide reduced component costs by leveraging the commercial industry investment in evolving technologies. To realize the benefits of COTS, the component selection, supply chain management, and tech refresh events must accommodate the product evolution cycles. Open Architecture and MOSA principles provide guidance and requirements that allow COTS product cycles to be effectively managed.

Military C2 systems can provide enhanced capability and performance by instituting net centric architectures that incorporate distributed network fabrics, centralized computing platforms, and Internet Protocol (IP) based HMI. Efficiencies gained by using COTS in the areas of performance and SWaP allow lower NRE and Recurring costs. Enclosure packaging and shock isolation concepts allow COTS products to be isolated from the military environment, thereby reducing the need for specialized hardware. Enclosures also provide a “hotel” for mounting space, power, and cooling that support multiple low cost tech upgrade and refresh cycles as COTS products evolve.

The attributes of COTS managed effectively during design, development, production and fielding offer a lower Total Ownership Cost (TOC) and higher performance and capability for military C2 and IT systems.

ACKNOWLEDGMENT

The authors acknowledge the assistance of the staff and managers of DRS C3 and Aviation Company whose expertise in application of COTS in military systems were very valuable in the creation of this paper. Many have contributed elements of this paper, as well as reviews that improved the content, and the authors gratefully acknowledge that assistance.

The authors specifically acknowledge Mr. Richard Cerpa, Director of Networks & Info Assurance at DRS C3 and Aviation Company, for his contributions to the networking section of this paper.

REFERENCES

- [1] Figure and description adapted from Open Architecture, The Critical Network- Centric Warfare Enabler, Presented by Captain Ric Rushton (N766-22) at the American Society of Naval Engineers Conference, Washington D.C. June 2004
- [2] Some data on the PC-over-IP protocol was obtained from the Teradici® website. See <https://www.teradici.com> for more information.
- [3] Interface Standard for Shipboard Systems, Electric Power, Alternating Current (Metric, MIL-STD-1399(NAVY)SECTION 300A

BIBLIOGRAPHIES

Ronald Busch is the Technical Director of Naval C4 Systems at DRS Technologies, C3 and Aviation Company. He has over 25 years of experience in development and integration of shipboard C4ISR components, subsystems, and systems at DRS Technologies. His naval experience includes project and systems engineering for programs such as Shipboard Display Emulator, the Common Display Systems, and the LHA Tech Refresh. Ron holds a B.S. degree in Electrical Engineering from the University of Maryland.

James A Lohr is a Staff Systems Engineer for DRS Technologies, C3 and Aviation Company. Mr. Lohr is the Principal Systems Engineer on the RSC program and the Common Display System (CDS) program (Consoles Phase 1, and Remotes Phase 1 and 2) and several other C4ISR systems. Mr. Lohr has over 30 years' experience in complex electronic systems, most in Command and Control and Communications applications.