

APPLYING OPEN STANDARDS ELECTRONICS ARCHITECTURES FOR GROUND VEHICLES

**David Jedynak, Jacob Sealander, Jason DeChiaro, Charlie Kawasaki, Tim McElligott,
John Ormsby**

Curtiss-Wright Defense Solutions, Ashburn, VA

ABSTRACT

Open Standard are useful for designing and instantiating specific electronics architectures on vehicles. Successfully designing them requires understanding all the factors that impact their usefulness. These factors and associated trade-offs for intended vehicle types include quantitative factors such as operating environments, thermal management techniques, size, weight, and power, and acquisition cost. Additionally, integration challenges, acquisition models, and industrial base collaboration add additional layers of complexity. All of these need to be considered for successful application for ground vehicles.

1. INTRODUCTION

The last ten years in the defense market has seen significant growth and adoption of industry-driven or industry-supported open standards, in some cases with direct participation or leadership of the government organizations (e.g. CMOSS, VICTORY, Modular Active Protection). While the technical standards have advanced and matured, some challenges remain, most prominent of which are the decisions about to best apply and optimize the various standards for different applications and vehicles. There is not always a one-size-fits-all approach. A spectrum of options are available, driven by multiple considerations. Various modular chassis + line replaceable module (LRM) and standalone line replaceable unit (LRU) electronics standards can seem in conflict; however, they are really at different points of a spectrum.

The existing acquisition approach for platform technology is well understood: a singular focus on

providing a specific capability (e.g., battle command software running on a physical bolt-on appliqué). This single purpose approach typically provides a self-contained materiel solution consisting of a LRU, platform Installation Kit (IK), training, spares, etc. These recurring lifecycle costs are relatively fixed at the LRU-level and generally well-understood. In some cases, the IK costs as much or more than the LRU itself. The combination of the LRU and IK results in size, weight, and power plus cost (SWaP-C) allocated to the platform. An open standard approach can provide a lot more flexibility in the acquisition model.

Open Standards provide ways to understand and appropriately design architectures that fit the specific vehicle needs as well as the acquisition model needs. Furthermore, open standards encourage a stronger industrial base by establishing clear rules of engagement where more insular business approaches existed previously.

Clean open standard lines between who and what enable trades and optimization.

Building up the overall architecture with open standard approaches helps break through both technical and acquisition challenges because more design options become available for trade. With that added flexibility is the risk of misapplication. The following sections provide guidance to mitigate those risks and to realize the benefits.

2. UNDERSTANDING THE ENVIRONMENT

Environmental standards and challenges drive design choices. A thorough understanding of what is needed, what the fundamental limitations are, and how to best optimize are another important driver for determining the best approach. This includes understanding standardized A-Kit vehicle envelopes, such as the typical radio shelves in use across the fleet.

Ground Combat Vehicle environments are well defined in the US Army CCDC GVSC (previously TARDEC) *Automotive Tank Purchase Description: Interface Standard – Environmental Conditions for the Heavy Brigade Combat Team Tracked Vehicle Standard (ATPD-2404 21-OCT-2011 – Distribution Statement A: Approved for Public Release)* [1] and newer revisions. This relatively straightforward and short document pulls together a number of standards (e.g. MIL-STD-810) by reference to provide operational environment constraints that must be considered for any materiel solutions intended for use within ground combat platforms. In short, any equipment within such a platform will suffer wide temperature ranges, high shock, lots of vibration, and a dirty environment, among the many other environmental conditions never considered for electronic systems – most specifically computing systems – originally designed for operation in commercial benign environments (e.g. home, office, server room, data center). An example of a typical commercial rack mount server is shown in Figure 1. A unit like this is not designed for the

ground vehicle market, and as a result, will not be a good fit for the environment.




Figure 1: Typical Commercial rack-mount servers are not designed for ground combat environments

The Mean-Time-Between-Failures (MTBF), the Mean-Time-To-Repair (MTTR), and the Failure Modes, Effects, and Criticality Analysis (FMECA) all drive the overall estimation of Operational Availability (OA). The OA is a percentage of time the system will be available, and is based on the simple math considering how often it fails (MTBF) and how long it takes to repair once failed (MTTR). This is further reduced by schedule maintenance activities (downtime) – a primary reason for the continual push for condition-based maintenance and failure prediction / prognostics models. The FMECA guides what is actually a critical failure that prevents the most critical use of the system. In the

case of ground vehicles, the most important operational tasks are to Move, Shoot, and Communicate. With that operational context, the most critical environmental constraints for deployed computing systems are summarized in Table 1 along with assessment versus commercial (benign environment) computing equipment, and a general best case mitigation action. When considering open standards, or even de facto industry standards, it is critical to fully understand the intended environment, as driven by the intended market, for that open or industry standard. As the following table shows, highly reliable commercial servers powering datacenters throughout the world with “five nines” reliability (OA of 99.999% or better) are absolutely mismatched for the specific environmental requirements of a combat vehicle.

Table 1: Assessment of Commercial Equipment versus ATPD requirements

Typical ATPD Requirement		
Commercial Server	Suitability Assessment	Possible Mitigation
Operational Temperature: -56°C to +52°C, with warm-up kits allowed below -32°C. Induced up Temperature of 71°C for 6 hours.		
10°C to 35°C	Will not survive	Re-design for temperature ranges or create environmentally controlled enclosure
Humidity: Up to 100% (non-condensing)		
20% to 80% (non-condensing)	Will not survive	Ruggedize design for humidity or create environmentally sealed enclosure
Sand: 10.6 to 17.7 g/m³ of 0.01 to 1mm diameter with velocity at least 1.5 m/s		
No specification	Will not survive	Enclose in sealed case
Dust: 0.006 g/m³ of 0.0001 to 0.0 mm diameter with velocity of 1.5 m/s		
No specification	Will not survive	Enclose in sealed case
Vibration: 4.7 G_{rms} 10-500Hz continuous		
0.26 G _{rms} 5-350Hz for 15 mins	Will not survive	Vibration isolation mounting
Shock, Basic: 50G for 20ms; 50G for 6ms; 34G for 2.5ms; 32G for 1ms; All Axes		
31G for 2.6 ms in Z-axis only	Will not survive	Shock isolation mounting
Leakage (Immersion): 1 meter for 2 hours		
Not sealed	Will not survive	Enclose in sealed case
Steam & Water Jet cleaning: 172.2 to 241.3 kPa (25-35 psi)		
Not sealed	Will not survive	Enclose in sealed case
Noise Level: 85 dBA without protection		
49 dBA	Within requirements	

Ignition Protection: No ignition sources		
Not expected, low risk	Not certain	Enclose in sealed case
Rapid Decompression: 15,000 ft equivalent to 40,000 ft < 15 seconds		
Low Risk (normal shipping requirement)	Within requirements	
Contamination by Fluids (e.g. oil, fuel, cleaning): No performance or physical degradation		
Not designed for	Will not survive	Enclose in sealed case
Nuclear Hardness: Per USANCA criteria (SECRET)		
Not designed for	Will not survive	Re-design or enclose in radiation shielded case per USANCA
CBRN: Exposure and Decontamination		
Not designed for	Will not survive	Enclose in sealed case
Input voltage: MIL-STD-1275 (28VDC)		
50/60Hz 110-220VAC	Not compatible	Voltage Conversion

This mismatch is a clear demonstration of why understanding the end use environment is so critical to the design approach.

Theoretically, commercial equipment could be enclosed in a radiation shielded, sealed, and environmentally controlled case with shock and vibration isolation; however, this is impractical and prohibitive from a SWaP-C perspective. The enclosure would need to be too large and inefficient with regard to space, and would require significant provisions for thermal control. An aggressive estimate for an approach that can hold up to three servers, along with various additional equipment for thermal control is described in Table 2. Note that this barely leaves room for environmental control systems (heaters / chillers).

Table 2: Space-claim comparison to standard mount for commercial server solutions

MT-6352 Mount (for reference): ~15.9” W x 12.2” D, 8” H (nominally)				
AN/PRC-160(V)HF Manpack for comparison	7.9” wide x 9.2” deep x 3.3” high	240 in ³	9.1 lbs	Fits MT-6352 Mount
Server	17” wide x 24” deep x 1.75” high	714 in ³	33 lbs	Exceeds MT-6352 Mount
3x 19-inch rack mount case	21” wide x 30” deep x 7” high	4400 in ³ > 2.5 ft ³	>130 lbs	Exceeds MT-6352 Mount

Clearly, this approach would be far from optimal in a combat vehicle’s constrained environment.

3. THERMAL MANAGEMENT TECHNIQUES

Different techniques for thermal management have benefits and drawbacks in specific operational environments, including reliability, performance, maintenance concepts, and other constraints (e.g. mounting, noise, etc.). A summary of those benefits and drawbacks are listed below in Table 3. Consider system level impacts when evaluating each for suitability, especially regarding OA. It is important to understand these different techniques at the LRU-level, especially in the context of using LRMs and open standard LRM thermal management techniques described in section 3.5.

Table 3: LRU Thermal Management Techniques Summary

<i>Technique and Required Equipment</i>			
Thermal Performance	ATPD Alignment	Integration Complexity	MTBF / Maintenance
<i>Liquid Cooled with liquid loop through LRU out to radiator, plus pump.</i>			
Highest	Low	Very High	Low / High
<i>Forced Air with Fans mounted in LRU</i>			
Medium	Low	Medium	Med
<i>Cold Base Plate with chassis mounted to hull / large block as heat sink</i>			
Low-Medium	Good	High	High / Low
<i>Natural Convection requiring maximized surface area</i>			
Low	Best	Low	High / Low

3.1. Liquid

Fundamentally, liquid loops can provide the absolute highest thermal management performance; however, they present significant challenges with regard to platform integration, maintenance, reliability, and overall integration requirements (tubing, pumps, radiators, reservoirs, etc.). The number of elements that can fail and the difficulty of repair (especially field repair) can be prohibitive for overall Operational Availability goals.

A counter argument can be made that ground vehicles already have one liquid system – fuel. Some still have hydraulic systems (e.g. turret drives); however, the move is away from those to electric turret drives for many reasons, including reliability and complexity. While it is conceivable that a fuel or hydraulic system could also be used to provide some level of liquid cooling loops for

the electronics (route fuel lines through electronics enclosures to carry away heat), the cascading impact from a single leak could impact not just the ability to move and/or shoot, but then also communicate as the electronic systems lack proper cooling. By using those existing loops to also provide electronics cooling, the sheer number of parts (and thus potential failures) will need to increase. From a system perspective, it is a difficult MTBF, MTTR, and FMECA challenge, and not a strong approach to support OA goals.

3.2. Forced Air

Forced air cooling with fans provides a lot of thermal management capability. In the right environment, fans are a good solution; however, fans introduce numerous problems versus the ground combat environment, precluding compliance to the following ATPD-2404 requirements:

- Sand, Dust, & Fluid contamination
- Immersion & Wash-down
- CBRN decontamination
- Requiring scheduled maintenance to meet OA
- Noise
- Filter maintenance

For example, one type of fan frequently used on military systems is designed to mitigate issues such as sand and dust, but its noise level is 82dBA at a moderate static pressure. Two of those fans would result in 85dBA – the limit in ATPD-2404. In addition, the MTBF of a fan like this is about 40,000 hours in a ground benign environment. Using standard MIL-HDBK-217 environmental conversion factors [2] (in this case, ground mobile MTBF is 20% of ground benign), a fan like this would have an MTBF of about 8,000 hours. With a vehicle-level operational duty cycle of about 50% (~4000 hours a year), the expectation is that the fan would need replacement every 2 years. If more than one fan is used on the platform this gets

worse – the MTBFs compound, e.g. 8 fans as a group would have a MTBF of about 1000 hours, requiring a change to a fan in the system roughly every 3 months. Furthermore, a single fan failure would result in degraded performance (if two are used) or a fault / system shut-down (if only one is used). That single failure when two are required will cause a temperature rise, which then stresses the remaining fan. The MTBF of the remaining fan will decrease as the it heats (roughly a factor of two with a 10° to 15°C increase). The fans are a challenge to meeting OA goals.

Forced air cooling is a actually a very good example of platform suitability considerations. As problematic as it is on ground vehicles, it is very well suited to aircraft, especially fast jets. The contrast between the two is instructive. The overall mission length (hours, not days) and maintenance concept (every flight) is very amenable to frequent servicing. Having said that, the very environment that the fans operate in is far more controlled and never anywhere near the temperatures of a ground vehicle. At the same temperature, the MTBF versus ground benign is only 10% per MIL-HDBK-217 (versus 20% for ground mobile) equating to about 4,000 hours; however the maximum temperature experienced is up to 30°C lower than the initial ground benign temperature estimate, resulting in about 4x improvement from that 10% baseline, thus about 16,000 hours. Furthermore, the fans are used in very different environment: centrally provided filtered airflow inlet plenums, no concerns about immersion or wash down, no noise concerns, and not the same sort of exposures to things like CBRN contamination inside the avionics bays. Furthermore, there is no chance that inlets or outlets will be blocked by other equipment when deployed in such well-defined and constrained environments. Understanding and considering these differences, in contrast to ground vehicles, can be very useful when considering the trade-offs for thermal management.

3.3. Cold Base Plate

This cooling approach requires the unit to be bolted directly to a large (relatively) cold mass, e.g., directly to the hull with a thermal interface material between the electronics chassis and the vehicle hull. Although this may initially appear to be a good approach to the electronics designer (no moving parts, thus higher MTBF), this method poses significant challenges in integration and requires the following considerations:

- Is there a suitable surface for the base plate to interface to, especially given equipment rack layouts, human factors, etc.?
- What is the temperature range of that surface and would it actually provide the appropriate ΔT needed to be useful across temperature ranges?
- Does that surface actually always stay cooler than the electronics, or could it get hotter? (e.g. solar load, engine heat, etc.)
- What impact at the platform level is there if heat is rejecting through that surface and mass? Could it impact the overall IR signature of the platform?

Cold base plate cooled enclosures do have the advantage of being sealed enclosures (versus forced air), but preclude any sort of shock / vibration mounting if needed – although those should not be needed with appropriately rugged open standard electronics. As sealed enclosures without any sort of moving parts, they will have much better alignment with the ground vehicle environment.

3.4. Natural Convection

Natural convection is generally the easiest for integration and reliability (no moving parts), but typically also provides the lowest performance. It relies on radiating heat surfaces, usually with fins to increase surface area, through which heat transfers to the (assumed static) surrounding ambient air. Fundamentally, this approach

conducts heat through a mass out to a surface, and the limiting factor is heat / area, which is then scales with the ΔT between the surface and the air. For general order-of-magnitude estimates, the upper bound of performance is generally 1 Watt per square inch, thus a 100 Watt unit needs 100 square inches of radiating surface (e.g. 10 x 10 inch). Proper design of fins and surfaces ensure the correct boundary conditions to maximize thermal transfer.

Natural convection has the advantages of being fully sealed and allowing for shock / vibration mounting if required. Aside from the temperature challenges, this is the easiest approach for ATPD-2404 compliance.

It is important to note that natural convection does require some amount of space around the unit. The space required is generally a clearance of a few inches (e.g. 2 or 3) from the thermal surfaces, allowing diffusion and natural air currents to form. This “chimney” effect creates updraft currents as hot air rises. With appropriate air space at the base of the unit and above the unit, air will flow, carrying the heat upwards and away.

For illustration sake, the absurd counter-example would be putting a natural convection cooled chassis in a sealed box only slightly larger than the chassis itself. Another caution is to ensure the chassis does not become a “shelf” for things to be placed (e.g. a sack). These sorts of precautions would apply to most electronic systems, not just those which are natural convection.

Because of how typical environmental conditions are stated (e.g. ATPD-2404), no consideration is given in the design to the actual realistic environmental conditions. Designs focus on two major highly conservative constraints: static air and steady-state temperature extremes. In the former, the assumption is made that the air in the environment of the unit is completely static. This means no gently circulating airflow due to any vehicle level air handlers, nor air flow due to open hatch air currents. Either one of these can provide just enough airflow to improve the

thermal performance over the static air assumption, often significantly (e.g. max allowable temperature). Compounding this is the assumption that a temperature is fixed and steady state, e.g. 71°C for a long number of hours while the system is under full load. In actuality, the specified temperature extremes are often required to ensure hot soak turn on operation (e.g. power up of unoccupied and sealed vehicle in afternoon sun), and continued operation through a period of time while other thermal management systems (crew focused air handlers, if present, or simply opening a hatches during ingress) reduce the maximum specified temperature. Understanding these differences between the specified maximums and actual operational profiles provides a significant amount of risk mitigation to natural convection approaches.

3.5. OpenVPX and LRM Cooling

The above sections explain the various approaches for LRU-level thermal management. On the surface, this may seem tangential to the application of open standards. On the contrary, the many LRU-level thermal management techniques demonstrate the advantage of open standard interfaces for use in modular open standard LRM approaches, specifically that of OpenVPX-specified thermal management, specifically conduction cooling mechanisms. A 3U OpenVPX LRM is shown in Figure 2.




Figure 2: 3U OpenVPX LRM highlighting wedge locks for thermal interface

Within the specification for OpenVPX are requirements and referenced specifications for conduction cooling of boards, with the primary

standardized thermal interface (wedge locks) at the boundary between the LRM and the inner surface of the LRU (inner side-walls), as shown in Figure 3. While there are other methods within the OpenVPX specification (various air and liquid methods), conduction cooling is the most common, and given the constraints of ATPD, the only method really suitable for ground vehicles.




Figure 3: OpenVPX Conduction Cooling to LRU inner side wall

This figure depicts the various thermal interfaces, and temperature deltas, across the thermal path. The hot parts (e.g. processor IC) conduct to a heat frame which is part of the LRM. This path results in a T_{edge} temperature specification. Typical for OpenVPX modules is a card edge temperature requirement of 85°C . The design of the open standard module is such that if the module's card edge temperature is kept at 85°C or below, the parts will stay within their maximum junction temperatures, and the module will continue to operate normally. The interface between the card edge and the inner side wall of the chassis has its own thermal drop through the standard OpenVPX specified wedge locks. These are designed to provide a certain amount of normal force between the card edge's thermal surface and the chassis inner side wall surface such that a well-controlled thermal drop will occur. This is typically 0.2°C per Watt. For example, a 50 Watt LRM, appropriately designed to maintain proper operation when both card edges are held at 85°C , will carry 25 Watts to each edge. When installed in the chassis with wedge locks appropriately torqued, the thermal drop from card edge to chassis side wall will be 25 Watts times

the $0.2^{\circ}\text{C} / \text{Watt}$ drop, or 5°C . This means that the LRU must be designed (using whatever appropriate LRU level thermal management approach) such that in the maximum temperature environment (e.g. 71°C), the chassis side wall will stay at or below 85°C card edge – 5°C wedge-lock drop = 80°C .

This open standard thermal and mechanical interface is incredibly important for understanding how to apply open standards for ground vehicles. Regardless of how the LRU is cooled, the open standard LRM interface is fixed. As long as an LRM meets those standards, it will drop in to the slot and thermal management will function per design. From this perspective, the LRM is viewed as a heat source, regardless of what is generating that – a CPU, a switch, an FPGA, a radio, a power amplifier, or a heater. The implications of this are far ranging, impacting both technical and acquisition approaches.

4. UNDERSTANDING SWAP TRADE-OFF

Size, Weight, and Power (SWaP) are a primary consideration when considering vehicle electronics, especially suitability for ground vehicles. Different approaches are optimized differently, leading to breakpoints between one or another. The simplest of these details are size, weight, and quantity of IKSs, as shown in Table 4, which compares multiple LRUs versus a single chassis LRU containing multiple LRMs. For this analysis, assume the LRM-based single chassis is intended to be some sort of common mounted chassis suitable for use on multiple platforms, sized to fit on a relatively standard radio equipment shelf typical in ground vehicles as noted in previous sections (MT-6352, $15.9'' \times 12.2'' \times \sim 8''$). Assume a single capability LRU is something like a modern processing unit used for battle management applications, e.g. $13'' \times 10'' \times 3'' = 0.22 \text{ ft}^3$ and about 10 lbs.

Table 4: Size, Weight, and IK comparison for single LRU vs LRM approaches

	Size (ft ³)	Weight (lbs)	Quantity of IKs (includes harnesses & mounts)
Single Capability LRU	0.22	10	1
Total for 8 Capabilities	1.76	80	8
Common mounted chassis LRU (size of shelf) for 8 LRMs	0.9	30	1
LRM per Capability	fits inside	1	None / interfaces to slot
Total for 8 Capabilities	0.9	38	1
Size, Weight, and IKs Versus 8x LRUs	52%	48%	13%

From just these three parameters, the trade-off of a common mounted chassis with LRMs versus a single capability LRU should be clear: when needing to optimize multiple systems at once, significant size, weight, and IK reductions can be realized with the LRM approach. The elimination of individual duplicative physical parts (housings, rugged connectors, thermal management, power supplies, etc.) and IKs for each capability drives a significant return of size and weight back to the platform. Reduction of IKs also results in simplification of the platform wire harnessing and commensurate reduction in associated size, cable runs, and weight.

Further efficiencies on the order of 10-20% are gained with regard to power via consolidation of power supplies. This is mainly due to elimination of losses through various power carrying and conditioning electronics (e.g. wire, diodes, inductors) that are common in power front-ends as well as overall reductions in load from LRMs and the elimination of circuitry normally required at the LRU level (e.g. power-hungry Ethernet PHYs for cable transmission between LRUs versus simple SERDES-level board-to-board communications on a backplane). These savings are not insignificant, but not as clearly straightforward to estimate as Size, Weight, and the quantity of IKs.

The important conclusion here is that the overall platform end-use and electronics requirements

need to be considered for the trade-off. Primarily, that means understanding if the platform needs just a single capability (e.g. a simple vehicle computer) or if it needs a more extensible and scalable architecture for many different applications. Alternatively, it may be that the platform will have a simple vehicle level need (a single vehicle control computer) as a stand-alone LRU, and then needs provisions for a more generic common mounted chassis for hosting multiple different add-on capabilities.

5. UNDERSTANDING COST

Cost is a critical parameter for savings. If each IK is estimated at an average of 25% the cost of a capability – a lower estimate given that some IKs are 200% or more the cost of the LRU – then an interesting model can be constructed. Assume the common mounted chassis LRU plus IK cost is anywhere from 4x to 6x the cost of an average IK (25% of LRU). Assume also that each LRM cost is about 75% to 80% of an equivalent LRU due to the elimination of LRU-level connectors, housing, and discrete power supplies. The graph below shows the overall benefit to the acquisition enterprise in the context of recurring cost. The results are compelling. With a single filled 8-slot common mounted chassis, up to 30% aggregate recurring cost is saved.

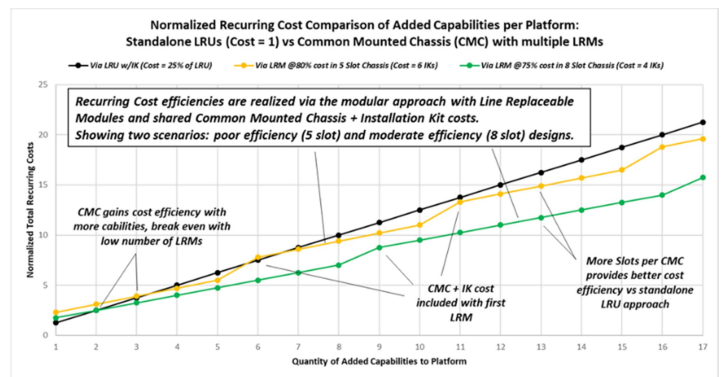


Figure 4: LRU vs LRM cost model

The focus above is on recurring cost. Development non-recurring engineering cost can be reduced overall with the LRM approach since certain aspects of environmental qualification will not need to be performed for each LRM. For example, chassis level items such as MIL-STD-1275 power testing, various wash down, immersion, and thermal tests are all performed only at the LRU level. Given the nature of the open standard modules, each will already be qualified by vendors.

6. INTEGRATION CHALLENGES

Often integration is cited as a high-risk area, depending on the complexity of interfaces. What is important to understand is that the various open standards provide methods of reducing risk and can enable and even acceleration the migration from one set of electronic hardware to another, representing a high return-on-invest from integration activities. However, open-standards allow for increased technology insertion, potentially resulting in added complexity. Use of an open-standards based software management solution that provides an integrated, single-pane of glass graphical user interface (GUI) that is hardware agnostic reduces complexity, improves visualization across all vehicle systems, and reduces training burden. This reduces training requirements for integrators, operators, and maintainers regardless of which platforms they use – and also insulates them from changes to underlying electronics from vehicle to vehicle. Furthermore, it enables the integrate-then-migrate cycle for multiple hardware types.

To address the complexity and training burdens introduced by evolving technologies, software tools which offer an intuitive user interface to simplify component provisioning, integration, and maintaining consistency throughout any vehicle architecture can be used. Providing a single interface for both on-platform, in-vehicle network requirements, as well as remote, off-platform interfaces can facilitate greater situational

awareness from higher headquarters. It is additionally beneficial to further extend configuration management by comparing configuration differences and imposing necessary changes onto local and remote vehicles from upper tiers while preserving the change records. These tools can simplify the setup, configuration, and management of the underlying equipment used in an open standard architecture, (e.g. VICTORY). environment.

Open standards make implementation and application of tools such as these more straightforward because the underlying interfaces are standardized. Even in cases where non-standard or proprietary interfaces exist, the appropriate application of management software can encapsulate and abstract away the non-standard interfaces. All of this reduces the integration challenge because the variables of configuration can be both captured and controlled, with clear migration paths forward.

7. NAVIGATING ACQUISITION MODELS

The traditional acquisition approach (separate systems, no interaction) is not necessarily optimal for platforms; however, from a program complexity and scope standpoint, it can be considered efficient. This model has worked in the past to bring relatively small sets of new or upgraded capabilities to existing vehicles without much integration complexity (bolt-on). Clean lines of separation and limited interaction between each capability are unintended consequences of separate and uncoordinated materiel acquisition solutions. However, with the drive toward accelerated technology refresh and the convergence of enterprise-wide multi-domain services and systems, it is now critical to streamline how we bring new capabilities to the fight to achieve overmatch. The question today should be how to make the overall delivery of new capabilities to platforms more efficient, and what adjustments to the acquisition approach are needed

to ensure the right enabling infrastructure is prepositioned throughout the enterprise.

Bringing technology to the battlefield is not just a technical challenge, but an acquisition and industrial-base challenge. Competition for funding and competition for business often discourage collaborative efforts. Open standards – used strategically – can break down those barriers and enabled better use of funding and faster acquisition cycles.

A key challenge in doing this is the reallocation of certain acquisition and performance responsibilities. In a single function / single system model (proprietary or otherwise), the responsibility and funding for the application is well defined and mostly self-contained. In an open standard modular approach, the standards offload some of the responsibility and shift cost (both NRE and recurring) to the mating sides of an open standard interface. For a new system to take advantage of an open standard architecture (e.g. an OpenVPX LRM software defined radio hosting a waveform), other elements of the open standard architecture need to be provided ahead of time, or in parallel. For example, a chassis with appropriate backplane is needed for the LRM to install in to, and the LRM itself needs software to run.

The acquisition model needs to shift to consider the open standard interfaces, with specific attention to the following:

- Are the physical standards applied to multiple acquisition programs in synchronization and coordination?
- Are the software standards applied to multiple acquisition programs in synchronization and coordination?
- Is there a top down directive driving the adoption of the open standard approach across multiple acquisition programs?

In the case of ground vehicles, the CMOSS standards, coupled with inclusion of threshold

requirements in multiple acquisition programs, provides the path to application of open standards. Timing of acquisition and deployment is the next challenge – the various corresponding open standard infrastructure elements need to be available for elements like LRMs to be integrated.

From a higher level, this synchronization and coordination means that the entire acquisition approach needs to be aligned with the use of open standards, ranging from POM cycles to the many various materiel solution development paths (BAA, SBIR, OTA, JUONS, and standard Programs of Record), whether they are S&TCD funds, RDT&E funds, or even O&M funds. Furthermore, it means that the long term sustainment model needs to be considered, as the entire life-cycle needs to be managed at the open standard level.

It is noteworthy to realize the long term impact of the proper application of open standards. A single LRM can be used for many different platforms. In addition, that same LRM can be used for many different purposes with different software loads. From a supply chain and sparing standpoint, this means the sustainment model for multiple different capabilities can merge in to a common sustainment model. Furthermore, the open standard nature of the LRMs means that the supply chain can actually be fed by multiple interchangeable parts, and even drop in technology refresh parts which maintain backward compatibility. By no means is this sort of acquisition transformation simple, but the end benefit can be immense.

8. INDUSTRIAL BASE COLLABORATION

Just as importantly, the proper application of open standards can be used to drive collaboration between industry vendors, replacing competition with collaboration, assuring that each company can bring what each does best in cooperation with the unique value brought by other vendors. With open standard interfaces, and overall capabilities built from multiple open standard building blocks,

and acquired through an open standard approach, companies can more easily specialize in certain technology offerings which co-exist within the application space.

For example, to deploy an advanced electronic warfare system, three companies could collaborate to bring forward a materiel solution which meets the need. Assuming an open standard chassis already exists on the platform, along with all the open standard interfaces required for system level interconnects, the three separate companies can provide the necessary open standard building blocks most appropriate to their technology focus.

One company can provide high performance RF devices (tuners) designed to the physical standards (OpenVPX) with software interfaces designed to the system and software standards (e.g. VICTORY, MORA, VITA 49). Another company can provide a high performance processing engine (e.g. CPU + GPGPU + FPGAs) designed to the physical standards (OpenVPX) with corresponding interfaces to ingest data from the RF tuner. A third company can provide sophisticated analytics algorithms (e.g. machine learning) which run on the processing engine to perform the analysis of the data from the RF tuner. In an open standard architecture, all three companies can collaborate as a team to provide the solution, or they can each independently provide their elements to a final integrator, such as the government, without ever working together at all. Without the benefit of open standard modularization of the greater standard, it would be more difficult for the three companies to work together, requiring engineering rework to align to each other's interfaces, architectural models, and business goals. Often this level of friction against coordination would drive the teams to a resource protecting no-bid or an attempt to go it alone with less than best-in-class elements outside their main technology specialty.

This approach brings the added benefit of keeping any concerns of intellectual property conflicts or competition from crossing company

boundaries. This also provides cleaner lines for companies to receive revenue from separate funding lines (e.g. one for RF tuner, one for processing, and one for algorithms). The topic of revenue and profit is even more important when considering the traditional multi-tier supplier structure. If three expert providers of complementary technologies can provide a materiel solution directly to the government for open standard integration, then the necessary tier / sub-tier margin stacking critical to the health of the industrial base can be avoided or diminished.

One other benefit of applying the open standards is the benefit to Foreign Military Sales (FMS) and the various requirements around that. Quite often, in FMS arrangements, some technologies are substituted before or after sale. In some cases, this is to protect specific capabilities from being exported. An open standard building-block (e.g. an LRM) which is sensitive can be replaced with an ITAR-free open standard module before export. In other cases, the foreign government may have in-country manufacturing requirements, often set at a certain percentage of the system cost. With open standard approaches, companies within the foreign nation can manufacturer locally sourced modules without needing complex technical data packages to replicate key functionality.

In all the various examples above, it is critical to understand the industrial base benefit. A healthy and robust industrial base is critical for driving technology investment, and with it, technology to the battlefield. Open standards significantly reduce the various friction points within a free-society (vs command driven) industrial base to collaborate and accelerate technology advancement.

9. EXAMPLE APPLICATION

Returning to the standard commercial rack-mount server example, a natural convection approach meets the ATPD-2404 requirements as shown in Table 5 and within the Size and Weight baseline of typical equipment shown in Table 6.

Table 5: Performance of Natural Convection OpenVPX LRU versus ATPD Requirements

Typical ATPD Requirement	
Natural Convection OpenVPX LRU	Suitability Assessment
Operational Temperature: -56°C to +52°C, with warm-up kits allowed below -32°C. Induced up Temperature of 71°C for 6 hours.	
Nominal -40° to 71°C, extension down to -56°C for cold-soak turn-on with extended boot time.	Meets Requirements
Humidity: Up to 100% (non-condensing)	
0-100% (non-condensing)	Meets Requirements
Sand: 10.6 to 17.7 g/m³ of 0.01 to 1mm diameter with velocity at least 1.5 m/s	
Sealed, tested per MIL-STD-810	Meets Requirements
Dust: 0.006 g/m³ of 0.0001 to 0.0 mm diameter with velocity of 1.5 m/s	
Sealed (including dust caps), tested per MIL-STD-810	Meets Requirements
Vibration: 4.7 G_{rms} 10-500Hz, continuous	
10 Grms→ 5-2000 Hz continuous	Meets Requirements
Shock, Basic: 50G for 20ms; 50G for 6ms; 34G for 2.5ms; 32G for 1ms; All Axes	
40G 11ms (standard, tested beyond as needed); All Axes	Meets Requirements
Leakage (Immersion): 1 meter for 2 hours	
Sealed, tested per MIL-STD-810	Meets Requirements
Steam & Water Jet cleaning: 172.2 to 241.3 kPa (25-35 psi)	
Sealed, tested per MIL-STD-810	Meets Requirements
Noise Level: 85 dBA without protection	
None	Meets Requirements
Ignition Protection: No ignition sources	
No ignition sources	Meets Requirements
Rapid Decompression: 15,000 ft equivalent to 40,000 ft < 15 seconds	
Relief vents for sealed chassis, tested per MIL-STD-810	Meets requirements
Contamination by Fluids (e.g. oil, fuel, cleaning): No performance or physical degradation	
Sealed, CARC paint or anodized, tested per MIL-STD-810	Meets Requirements
Nuclear Hardness: Per USANCA criteria (SECRET)	
Nuclear Event Detector (NED) with electrical provisions, tested by appropriate lab	Meets Requirements
CBRN: Exposure and Decontamination	
Sealed, CARC paint or anodized, tested per MIL-STD-810	Meets Requirements
Input voltage: MIL-STD-1275 (28VDC)	
Designed and tested to MIL-STD-1275 (28VDC)	Meets Requirements

Table 6: Space-claim comparison to standard mount for Natural Convection OpenVPX 8 slot LRU

MT-6352 Mount (for reference): ~15.9" W x 12.2" D, 8" H (nominally)				
AN/PRC-160(V)HF Manpack for comparison	7.9" wide x 9.2" deep x 3.3" high	240 in ³	9.1 lbs	Fits MT-6352 Mount
Natural Convection OpenVPX LRU (8 slot)	15.9" wide x 12.2" deep x 8" high	1552 in ³	38 lbs	Designed to fit MT-6352 Mount envelope

Anticipating future requirements, the OpenVPX – as part of CMOSS – ecosystem has sets of well-defined module interface definitions called *profiles*, which make interchangeability and upgradeability straightforward, while simplifying drop-in replacement or technology refresh. The module profiles and corresponding backplane slot profiles within a common LRU are interconnected with well-defined backplane topologies and capabilities. A subset of these have been captured in the CMOSS (and broader SOSA) standards, providing even tighter interface definition for technology refresh and reconfiguration.

To highlight this, an example baseline 8-slot common mounted chassis with conceptual technology refresh configuration is shown in Table 7. Note how some capabilities have collapsed in to a single module, opening up slots for new capabilities.

Table 7: Example baseline 8-slot common mounted chassis

Slot	Slot Type	Current Configuration	Future Configuration
1	Central Switch	40 Gigabit Ethernet Switch	100 Gigabit Ethernet Switch
2	Central Timing	Assured PNT Module with M-Code GPS receiver	APNT Module with M-Code GPS receiver + additional signal receivers and algorithms
3	Processing & I/O	Processor with Mission Command Software and interfaces to platform displays	Next Generation Processor with enhanced AI Engines and Augmented Reality Graphics Processing and headset interfaces for next generation Mission Enhanced Situational Awareness Software
4	Payload	Processor running tactical intelligence software	Next Generation Processor running tactical intelligence and targeting software
5	Payload	Processor running targeting software	Next Generation AI Accelerator supporting slot 4
6	Payload	Graphics Processing Unit providing AI acceleration for slots 4 and 5	Mobile ad-hoc network (MANET) transceiver with built-in CSfC-based Data-in-Transit encryption
7	Payload	Software Defined Radio rehosting existing DoD waveforms	Software Defined Radio simultaneously rehosting existing DoD waveforms, 4G/5G, WiFi, and commercial SATCOM
8	Payload	Multi-channel SIGINT Receiver	Multi-channel SIGINT and Passive Radar receiver for Active Protection Systems

Illustrating the concepts discussed in previous sections, the modules, chassis, and software can all be different vendors. Of course risk can be reduced by having some elements pre-integrated where some technical features sit at the ragged edge of open standardization (e.g. Processing and I/O modules, typically connecting to some legacy systems with legacy I/O). In some cases, this will be preferable.

Importantly, the various portions of the system can be updated independently. Software can be updated on existing modules. Existing modules can be migrated to newer modules. As long as the modules and software continue to adhere to the open standards, the overall system integration – especially aided by software tools – can maintain stability. Other factors to consider when performing upgrades:

- Overall power envelope for a given module – this must stay within the power margins defined for the slot, both for consumption and thermal management constraints
- Module profile – this must stay the same as previous module and the backplane slot

This is simple and straightforward as long as the constraints of the open standards have been adhered too by all parties in the development of the systems.

10. SUMMARY

The application of open standard electronics architectures for ground vehicles has clear benefit. The technical, cost, and risk reduction benefits have been shown. The benefit to the industrial base has also been explained. The real challenge is that of acquisition model. The leading acquisition program for a single new capability will always be at some cost disadvantage if it is also required to deploy the open standard enterprise infrastructure (e.g., common mounted chassis) for the collective benefit it provides to other contemporary and emerging requirements. Top level coordination and synchronization is necessary to drive success; nevertheless, the strategic benefit of applying open standards as infrastructure is essential for swiftly deploying new technologies to the field. Investment for the greater good of the warfighting enterprise will enable the technology breakout and multi-domain convergence essential to increase Warfighter and Weapon System effectiveness for our collective national defense.

1. REFERENCES

- [1] *ATPD-2404 Environmental Conditions for the Heavy Brigade Combat Team Tracked Vehicle*. Warren, MI: US Army, 2011
- [2] *MIL-HDBK-217F Reliability Prediction of Electronic Equipment*. Washington, DC: US Department of Defense, 1995.