2026 NDIA MICHIGAN CHAPTER
GROUND VEHICLE SYSTEMS ENGINEERING
AND TECHNOLOGY SYMPOSIUM
MODULAR OPEN SYSTEMS APPROACH (MOSA) TECHNICAL SESSION
AUGUST 11-13, 2026 - NOVI, MICHIGAN

OPTIMIZING DATA DELIVERY FOR GROUND VEHICLE TEAMS

Jafar Al-Gharaibeh, PhD[1], Jordan Bonney[1]

[1]Architecture Technology Corp. (ATCorp), Eden Prairie, MN

ABSTRACT

The radio-based wide area network (WAN) that forms the command-and-control backbone for deployments of multiple ground vehicles is a classic DIL (disconnected, intermittent, limited) communications infrastructure for bandwidth-intensive services like video streams, situational awareness feeds, and command-and-control messages.

This paper describes an inter-vehicle network architecture that leverages the radio-aware routing features of an existing vehicle Ethernet switch/router to provide mission-optimized, fault-tolerant data delivery while minimizing both network overhead and vehicle SWaP requirements.

Citation: Al-Gharaibeh, Bonney, "*Optimizing Data Delivery for Ground Vehicle Teams*," In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 11-13, 2026.

1. INTRODUCTION

Weapons replaceable units (WRUs) commonly use one or more Ethernet switches/routers as the interconnect between various onboard sensors, processors, and actuators. These onboard computing elements have unique internet protocol (IP) addresses within the same IP subnet range. To simplify manufacturing and deployment, the same IP subnet is used for every instance of the WRU that is built.

Consider the network topology depicted in Figure 1. Each WRU (A, B, C) has the same internal local area network (LAN) subnet of 192.168.1.0/24. The WRUs are connected to each other by a wide area network (WAN) that uses subnet 10.1.0.0/16.

If IP communication is to occur between any specific element of a WRU and an entity that resides outside of the WRU's LAN, a Network Address Translation (NAT) [1,2] function must be provided at the edge of the WRU to disambiguate the common addresses used by all WRUs.

For example, assume all three WRUs in the diagram have a video camera that, by convention, is always assigned an IP address of 192.168.1.3. If a command-and-control (C2) node located on the WAN wants to access WRU A's camera, the C2 node can't directly specify 192.168.1.3 as the target address because i.) the 192.168.1.0 LAN subnet is outside of the WAN subnet's address range (10.1.0.0/16), and ii.) 192.168.1.3 is an IP address that's been assigned to three different physical cameras (one on each WRU).

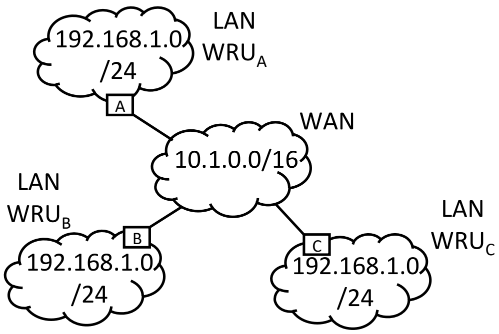

Figure 1. Replicated LAN Subnets Connected by a Common WAN.

Specific configuration of the NAT function on the switch/routers that are located at the edge of each WRU subnet (shown as boxes A, B, and C in Figure 1) can be handled in a number of ways. One option is to configure a single, unique WAN address in the 10.1.0.0/16 subnet for each edge switch/router and also configure port-forwarding rules for specific IP hosts on the corresponding WRU's LAN subnet to accommodate inbound (WAN-to-LAN) connection requests – a classic one-to-many NAT. An alternative configuration is a one-to-one (static) NAT strategy where addresses from the (larger) WAN subnet pool are mapped directly to corresponding host IP addresses on the (smaller) LAN subnets.

Whether one-to-many or many-to-one NAT rules (or a combination of the two) are configured, external access to the LAN elements of a specific WRU is accomplished by the external entity specifying the WRU's unique WAN address(es). For example, assume the following one-to-one NAT mapping has been configured for each WRU's camera:

WRU	Camera LAN IP	Camera WAN IP
A	192.168.1.3	10.1.1.3
B	192.168.1.3	10.1.2.3
C	192.168.1.3	10.1.3.3

With this one-to-one NAT mapping for each WRU, accessing any particular camera is straightforward: external entities initiate data requests to the corresponding WAN address (10.1.1.3, 2.3, or 3.3) of a specific WRU's camera.

The WAN itself may be comprised of redundant, diverse radio links. These radio links vary in terms of connectivity and throughput over time (a disconnected, intermittent, and limited, or "DIL," underlay), which is discussed later; however the limited throughput of the radio links makes multicast support imperative (e.g., video streaming to multiple consumers across a bandwidth-limited WAN [3]).

In our experience, organizations deploying networks like that described above (DIL underlay WAN with NAT'ing to connect multiple LANs with common IP subnets) want to avoid the additional complexity of implementing protocol independent multicast (PIM) or Layer 3 dynamic routing protocols like OSPF [4], preferring instead to rely on the internet group management protocol (IGMP) for multicast [5] and the address resolution protocol (ARP [6]) to resolve next-hop addressing on the WAN.

2. Radio-Aware Optimization

The radio WAN is commonly comprised of two or more different mesh-network radio waveforms (e.g., Streamcaster, Trellisware, WaveRelay, etc.) that provide redundant, wireless subchannels for the WAN. These redundant, but disparate, WAN subchannels individually have instantaneous bandwidth, loss, and latency that is greatly influenced by signal strength at their respective RF endpoints as well as at any intermediate relays. Receive signal strength (RSS) at any element in a subchannel is impacted by the physical distance to the closest neighbor, antenna design, transmitter power, terrain/obstructions, and RF spectrum noise (including jamming).

The challenge is determining which of the radio subchannels is "best" at a particular moment or if data should be load balanced across all available subchannels to provide greater aggregate throughput than any single subchannel can provide.

The OnTime Networks line of rugged switch/routers that are used within the ground vehicles community include a radio-aware tactical routing feature called "underlay mode" (or "underlay links") that can be used to dynamically ascertain which subchannel(s) to use and how to prioritize different data flows within the subchannel(s).

Underlay mode activates a neighbor-discovery mechanism on a switch/router's WAN interfaces, i.e., the switch/router's interface to a radio terminal (the "underlay links"). Underlay links perform realtime link-quality monitoring to determine which WAN subchannel is best.

Consider two OnTime switch/routers S/R_A and S/R_B, as depicted in Figure 2, below. Each has a direct interface to a corresponding radio terminal (rad_{Ax} and rad_{Ay} for S/R_A, and rad_{Bx} and rad_{By} for S/R_B), which in turn provides wireless connectivity to a corresponding WAN subchannel (WF_x and WF_y).

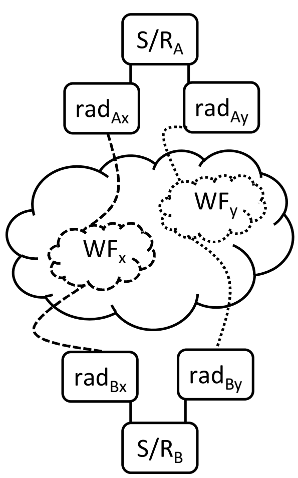

Figure 2. Switch/Router Interfaces to Radios that Form Discrete WAN Subchannels.

The maximum throughput of WAN subchannels WF_x and WF_y can vary over time because of changes in RSS at a specific radio terminal. WF_x might provide an order of magnitude better throughput than WF_y for a period of time while the opposite might be true five minutes later.

S/R_A and S/R_B continually assess the end-to-end throughput across WF_x and WF_y and adjust the data transmission accordingly. I.e., if one WAN subchannel is an 10x faster than the other, then 10x more data will be sent over the faster radio.

From an addressing perspective within the context of the network laydown presented in Figure 1, a /24 IP subnet is created on each switch/router; the logical addressing is provided in Figure 3. The subnet address of WF_x and WF_y are not shown in the diagram; they can be any unique (non-overlapping) range within the overall network addressing scheme.

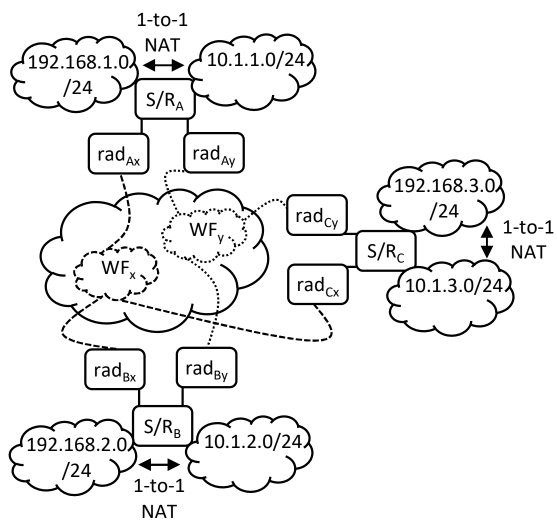

Figure 3. Network Architecture and Addressing Scheme for Radio-Aware Routing Evaluation

3. Configuration Impact

Additional configuration required to invoke underlay mode on an OnTime switch/router consists of changing the following clause of the ares.json file that corresponds to a radio-link (WAN) interface from the following form:

```
"Links" : [
    {
      "Name" : "link1",
      "Type" : "simple",
      "Interface" : {
        "Name" : "eth0",
        "Inet4Addresses" : ["10.0.1.1/16"],
      }
    }
  ]
```

To:

```
"Links" : [
    {
      "Name" : "link1",
      "Type" : "underlay",
      "Interface" : {
        "Name" : "eth0",
        "Inet4Addresses" : ["10.0.1.1/16"],
      }
    }
  ]
```

4. Network Overhead Impact

With the original architecture shown in Figure 1, attempts to access an individual element of a remote WRU (such as a camera) triggers an ARP request that is broadcast over the WAN. ARP requests are initiated periodically by each requesting host's network stack. Though ARP messages are small, their impact on aggregate WAN overhead is amplified by the fact that they must be broadcast to every endpoint of a given radio subchannel and that the ARP request is initiated for every remote WRU element that is being accessed.

Using underlay mode eliminates these ARP requests. Underlay links issue periodic neighbor-discovery messages to ensure reachability of remote switch/routers. The neighbor-discovery messages, however, are sent on a per-subnet basis. This should lead to a WAN overhead reduction that can be measured through packet-capture and analysis.

5. SWaP Impact

The underlay-mode feature already exists in the OnTime switch/router appliances that are used on a variety of ground-vehicle systems, which suggests there would be no SWaP increase for these vehicles, or even a SWaP decrease for those that use a separate load-balancing appliance for the radio-WAN subchannels.

6. Additional QoS Considerations

Underlay-mode switch/routers can also be used to ensure that a single video stream, whose bit rate exceeds that of the (currently) fastest WAN subchannel, is preempted by lower priority flows whose aggregate bit rate is less than the WAN subchannel's current capacity. Doing so ensures that lower-priority traffic can still pass if higher-priority traffic would otherwise saturate the available WAN subchannel.

This concept is illustrated in Figure 4, where the network load on the left side of the diagram exceeds the available capacity of the WAN subchannel. The underlay-mode switch/ router preempts the biggest flow (shown in blue), which is also designated as a high-priority flow. The remaining, smaller high-priority flow (green) as well as the low priority flow (purple) of approximately the same bit rate as the green high-priority flow *both* get forwarded through the WAN subchannel.

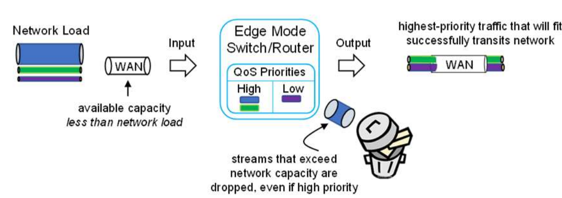

Figure 4. Advanced Prioritization.

7. Planned Experiments

In the coming months experiments will be developed that compare data-delivery performance in the ARP/IGMP approach depicted in Figure 1 with that obtained from the radio-aware approach of Figure 3.

The experiments will define performance profiles for the WAN subchannels. The WAN subchannel performance profiles will define epochs of different lengths during which each subchannel's bandwidth remains constant. For example, during epoch 0, WF_x and WF_y could both be configured to provide 10 Mbps. During epoch 1, WF_x might drop to 5 Mbps while WF_y maintains 10 Mbps. Longer epochs (10's of seconds) will be used to assess overall throughput, which shorter epochs (less than 1 second) will stress the edge routers' ability to detect changes in available subchannel throughput and adapt accordingly.

Transmitted data will emphasize connectionless, constant bit rate flows that are representative of live video streams.

Experiment scalability will be limited by hardware resources, but we anticipate tests with up to 32 WRU LANs each with up to 32 source/sink elements (hosts).

8. Conclusions

Network architectures that rely on one-to-one NAT'ing of WAN-side IP addresses to LAN-side address spaces that are common to every WRU are effective. However, their isolation from the details of the WAN's underlying radio networks likely necessitates the need for an external load balancing appliance to optimize data delivery across a range of operating conditions that impact the receive signal strength of different radios at different points in a mission (e.g., terrain, foliage, antenna placement, jamming).

The radio-aware routing features of the OnTime Networks rugged switch/router products need to be evaluated to determine if they can provide performance improvements, cost savings, and reduced SWaP for ground vehicle systems.

The first step in evaluating these radio-aware routing features should be experiments that are conducted using different numbers of WRU LANs, different numbers of LAN sources/sinks, and a variety of different WAN performance profiles.

9. REFERENCES

[1] P. Srisuresh, M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations," IETF RFC 2663, Aug. 1999.

[2] P. Srisuresh, K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)," IETF RFC 3022, Jan. 2001.

[3] N. Bahr, et al., "Efficient End-to-End IP Multicast in Crypto-Partitioned Networks," MILCOM 2015.

[4] J. Moy, "OSPF Version 2," IETF RFC 2328, Apr. 1998.

[5] B. Cain, et al., "Internet Group Management Protocol, Version 3." IETF RFC 3376. Oct. 2002.

[6] D. Plummer, "An Ethernet Address Resolution Protocol," IETF RFC 826. Nov. 1982.

10. CONTACT INFORMATION

Jafar Al-Gharaibeh, Research Scientist
Architecture Technology Corp.
9971 Valley View Rd.
Eden Prairie, MN
jafar@atcorp.com
952-829-5864 x172

Jordan Bonney, Chief Technology Officer
Architecture Technology Corp.
9971 Valley View Rd.
Eden Prairie, MN
jbonney@atcorp.com
952-829-5864 x104

11. ACKNOWLEDGMENT

ATCorp thanks their customers for their ongoing product feedback and radio-aware routing requirements.

12. ACRONYMS

ARP	Address Resolution Protocol
C2	Command and Control
IP	Internet Protocol
LAN	Local Area Network
NAT	Network Address Translation
PIM	Protocol Independent Multicast
WAN	Wide Area Network
WRU	Weapons Replaceable Unit