

A New Era of Afloat IP Services

Christopher B. Landis

[October-December 2016](#)

<https://www.doncio.navy.mil/CHIPS/ArticleDetails.aspx?ID=8349>

TAGS: [CTO: Networks](#), [Cybersecurity](#), [NGEN](#), [NMCI](#), [NNE](#), [Spectrum](#), [Telecommunications](#), [Wireless](#), [Workforce](#)

Cybersecurity is often considered in three pillars: confidentiality, integrity, and availability (CIA). In conjunction with the theme of this edition of *CHIPS*, “strengthening cybersecurity,” this article examines how the newest version of the Automated Digital Network System (ADNS) improves availability when employed to its greatest extent.

Many seagoing Navy communicators have experienced internet protocol (IP) and supporting radio frequency (RF) service shifts that have lasted several hours, leaving the ship without connectivity. The data path complexity inherent in ADNS Increment II and the supporting satellite and terrestrial transport lend themselves to such outages when shifting services across different satellite ground stations and IP providers.

Nonetheless, over the past few years, *a new era of afloat IP services* has dawned on the Navy. ADNS Inc. III can mitigate many of these prolonged outages and provide greater availability and resiliency to the fleet through cipher-text (CT) transport, removing inter-enclave dependencies, improved load balancing, and simplified troubleshooting.

To take advantage of the capabilities that ADNS Inc. III brings to the fleet, one needs two things: First is to understand how the ADNS Inc. III architecture improves upon ADNS Inc. II. Second is how fleet commanders should modify their tactics, techniques, and procedures regarding the selection of Standardized Tactical Entry Point (STEP) sites and their associated Network Operations Centers (NOCs) to provision RF and IP services, respectively. Some of the recommendations in this article will require a paradigm shift for career Navy communicators.

The Tactical Networks Program Office’s (PMW 160) official description of ADNS is as follows: “ADNS supplies the tactical wide area network (WAN) component of the naval communications system, providing surface ship, submarine, airborne, tactical-shore and shore-based WAN gateway services management. Since its initial deployment in 1997, ADNS has evolved through multiple increments. Increment III increases network capacity by adding the ability to utilize higher bandwidth available with newer satellite communication (SATCOM) systems. Efficiency is increased by converging to all-IP voice/video/data and implementing dynamic load distribution across SATCOM links.”

When a ship connects to a satellite to get IP services, there is a lot that happens on shore to actually provision IP services. Each step in the process must be in place before the next step will work. The fundamental end-to-end path is as follows: Starting with RF, cryptographic devices need to synchronize and terrestrial paths need to be in place before the IP/network layer begins to establish neighbor state. Once all this is in place, actual payload data for end users starts flowing.

There are some attributes that do not change when comparing ADNS Inc. II and Inc. III.

The RF portion is the most fundamental portion of the connection and, practically, does not change with the introduction of ADNS Inc. III. Note that Navy communications commonly use the term “RF”

to refer to more than the electromagnetic spectrum to include the associated baseband equipment. The Standardized Tactical Entry Point site establishes RF connectivity with the same satellite as the ship, which will lead to modem lock.

Modem lock is defined as the synchronizing of the ship's modem with the Standardized Tactical Entry Point site's modem. Modem lock depends on suitable look angles for the ship (satellite is above the horizon and not in a block zone for the ship's antenna), the right RF spectrum settings (beacon and carrier frequencies), and appropriate transmit power levels on the ship, at the STEP site, and on the satellite. The regional Wideband Satellite Operations Center (WSOC) controls satellite power levels.

Next, the transmission security (TRANSEC) at the STEP site will synchronize with the TRANSEC on the ship. Transmission security synchronization depends on modem lock, the right cryptographic device settings, and the right keying material (KEYMAT). At this point (synchronized TRANSEC), a ship is considered to be "good on RF" and after an hour of good RF, the ship is considered to be "stable on RF." These baseline conditions for RF are the same for ADNS Inc. II and III.

How did ADNS Inc. II work?

To understand how ADNS Inc. III improves upon ADNS Inc. II, it is necessary to have a basic understanding of ADNS Inc. II. Table 1 highlights some of these improvements and Figure 1 is a diagram of the following explanations, comparing ADNS Inc. II and ADNS Inc. III.

One of the most pronounced aspects of ADNS Inc. II is that it employs Secure Internet Protocol Router Network (SIPRNet) for its network layer transport. All other enclaves, such as Non-Secure Internet Protocol Router Network (NIPRNet), Joint Worldwide Intelligence Communications System (JWICS), Combined Enterprise Regional Information Exchange System (CENTRIXS), etc., tunnel through SIPRNet between ship and shore.

At the datalink layer, ADNS Inc. II employs Automated Digital Multiplexing System (ADMS) (Timeplex/ST-1000) or High-Speed Global Ring (HSGR)/Circuit-to-Packet (CTP), depending on the variant of ADNS Inc. II. Only about a dozen Timeplex ships remain at the time of this publishing.

In anticipation of establishing good RF, the STEP site and technical control facility (TCF) at the communications station (COMSTA), with the servicing NOC (where the ship will receive IP services), will have worked together to establish a terrestrial path between each other for the ship. (See list of NOCs in Table 2.)

For a Timeplex ship, the STEP site will patch its own Timeplex with the ship's Timeplex and build a path through the shore ST-1000s to reach the NOC. For a non-Timeplex ship, the technical control facility will build a path through circuit-to-packet. The point is that the path needs to be built manually between the Standardized Tactical Entry Point site and NOC.

Further Explanation: Previous to mid-2016, the path between a STEP site and a NOC for ADNS Inc. II was built through the High-Speed Global Ring via a TNX-1100 Asynchronous Transfer Mode (ATM) network. In an effort to support the Defense Information Systems Agency's (DISA's) Time-Division Multiplexing (TDM) drawdown, HSGR has been replaced by circuit-to-packet, finishing in mid-2016.

In the first year of its existence, CTP has already proven to be more user-friendly and reliable than HSGR — with features like a built-in bit error-rate test (BERT) — and highly versatile, being convenient for many more payloads than just ADNS. Timeplex and ST-1000 use CTP for transport. Circuit-to-packet consists

of Juniper circuit emulation (CEM) routers so that the serial connectivity used over the RF path can travel as IP packets between the STEP site and NOC.

For the IP path portion of circuit-to-packet, CTP employs the Navy Tactical Wide-Area Network (NTWAN), a Navy cipher-text partial-mesh network formerly employed only by ADNS Inc. III and NOC-to-NOC connectivity. Whether Timeplex, HSGR, or circuit-to-packet, the path still needs to be built manually between the Standardized Tactical Entry Point site and Network Operations Center.

Now that the ADNS Inc. II ship's connection is "good on RF" and its terrestrial path, the technical control facility at the COMSTA with the servicing NOC patches the path to a SIPRNet communications security (COMSEC) device. This will synchronize with the SIPRNet COMSEC device on the ship. Communications security synchronization depends on a reliable path between the two COMSEC devices (RF and terrestrial), the right cryptographic device settings, and the right KEYMAT. With COMSEC synchronization, now IP services (finally) come into the discussion.

The first part of IP services for an ADNS Inc. II ship is establishing an Open Shortest Path First (OSPF) neighbor state on SIPRNet between the ship's ADNS router and the shore's frequency router.

Neighbor state can be defined as two or more routers having exchanged route information, contributing to each other's routing tables. Routing protocols, like OSPF, define the algorithms by which routers exchange route information. Once an OSPF neighbor state is established, the routers on the ship and shore will have routes in their routing tables that will allow traffic flow between the ship and shore.

Part of that traffic flow will include the other enclaves. For each of the other enclaves, there will be a separate communications security device that synchronizes with the ship's COMSEC device for that enclave. As each of those enclaves' COMSEC devices synchronizes, their respective routers will also need to establish a neighbor state with the shore routers. Most other enclaves use Open Shortest Path First, including NIPRNet, but some use Enhanced Interior Gateway Routing Protocol (EIGRP), a Cisco-proprietary routing protocol.

With regard to load balancing across multiple paths, the best that ADNS Inc. II can achieve is on ships that have both super-high frequency (SHF) and the commercial broadband satellite program (CBSP). NIPRNet will be routed across one path and SIPRNet via the other. Through a manual process of cross-connecting the SHF and CBSP paths, it is possible to glean additional throughput according to a ship's demand, according to Cmdr. Eric Johnson, in his CHIPS article, "[USS Enterprise SHF Cross-Connect Configuration Increases Allocated Bandwidth by 100 Percent.](#)" [1]

That sounds complicated. How does ADNS Inc. III improve upon this?

The ADNS Inc. III improvements on ADNS Inc. II are game-changing. Refer to Figure 1 again to see the differences at a glance. As mentioned above, the RF portion of the ship-to-shore connection remains the same as with ADNS Inc. II. But instead of using SIPRNet for its network layer transport, ADNS Inc. III uses its own cipher-text enclave for transport. From the TRANSEC at the Standardized Tactical Entry Point site, the STEP site connects the path directly to an ADNS Inc. III border router. The ship's ADNS router (which is a CT router instead of a secret router) will establish an Open Shortest Path First neighbor state with this border router. The routing between the border router at the STEP site and the policy router at the NOC is automatic so that each ship doesn't need a separate, manually built path. All of the ADNS Inc. III missions for a Standardized Tactical Entry Point site are aggregated across parallel/redundant IP paths between the STEP site and NOC.

Further Explanation: Formerly, the Navy Tactical Wide-Area Network has been employed for this IP path but the ADNS Inc. III Service Pack (SP) 3 upgrade removes the added complexity of the NTWAN, a partial-mesh network of leased circuits riding over DISA's cipher-text partial-mesh transport network, and sends the circuits directly over DISA's cipher-text network. At the time of publication, the ADNS Inc. III SP3 shore upgrade is already at Pacific Region Network Operations Center (PRNOC) and European-Central Region Network Operations Center (ECRNOC), in progress at Indian Ocean Region Network Operations Center (IORNOC), and forthcoming at Unified Atlantic Region Network Operations Center (UARNOC), likely to be completed in early-2017.

Since ADNS Inc. III uses a cipher-text transport that is independent of any of the enclaves installed on the ship, each of the network enclaves, such as SIPRNet, NIPRNet, JWICS, CENTRIXS, etc., are independent of each other. This is in direct contrast to ADNS Inc. II.

Like ADNS Inc. II, however, the communications security device on the ship for each enclave will synchronize with the matching shore COMSEC device and the enclaves' respective routers will also establish a neighbor state with the shore routers.

So far, all improvements described have been architectural. A functional improvement is the load balancing capability of ADNS Inc. III.

To enable the load balancing across all paths, first consider that the traffic is bidirectional and each direction is load balanced at its source. All traffic leaving the ship is load balanced by the cipher-text ADNS router. Since all of the data paths for the ship converge at the CT ADNS router, it is relatively easy for it to load balance traffic that it is transmitting across each of those paths.

On the shore, all traffic sent to the ship is load balanced by the ADNS CT border routers, which are co-located with the RF terminations, with assistance from the ADNS performance routing (PfR) master control (MC) routers.

Further Explanation: The ADNS performance routing master control routers actually determine how shore-to-ship traffic is load balanced, establishing and modifying policies in real-time, according to each ship's traffic demands. Traffic is categorized according to protocol and enclave. For example, web traffic on NIPRNet is lower priority than video teleconference (VTC) traffic on SIPRNet. Since load balancing occurs in the cipher-text enclave, the only way that the routers know which protocol is encrypted in each packet is by the flag that each enclave's packet shaper is configured to write into each packet's IP header. The tactical local-area network encryption (TACLANE) do not encrypt this flag field so that it is readable in the CT enclave.

Note that as the links reach saturation, performance routing will start dropping the lowest traffic classes altogether (in order of precedence) to ensure that the higher priority traffic can get through. In other words, while you might not be able to browse social media sites on the web, SIPRNet VTC users will not even notice any degradation caused by the traffic congestion.

Also note that performance routing may cause asymmetric routing within the ADNS cipher-text core: The ship routes a specific traffic class over the commercial broadband satellite program and the shore sends the same traffic class back over SHF. While this by itself is not a problem, complications in troubleshooting can arise if the RF quality of a path begins to degrade, especially when performance routing starts moving traffic classes around based on the ship's demand. Performance routing does not consider path quality or reliability in its algorithm. There are still plenty of other details not described here surrounding the operation of PfR that could readily serve as the subject of another article.

Although all paths converge at the ADNS cipher-text policy router on shore, the paths at the ADNS cipher-text policy router are shared among all of the units provisioned in that area of responsibility (AOR). To see the individual paths to each ship, the ADNS CT policy router needs to extend its visibility to the ADNS border routers with which the ships' ADNS CT router establishes an OSPF neighbor state.

Further Explanation: The ADNS performance routing master control routers push the load balancing policies through the ADNS cipher-text policy router to its border routers for execution. As such, each of the four fleet NOCs has an ADNS CT policy router that connects to its AOR's border routers via tunnels (like a virtual private network). If, for example, a ship is transiting between AORs and has one path connected to a border router for UARNOC and another path connected to a border router for ECRNOC, the load balancing from shore to ship will not work across those two paths. Figure 2 depicts each ADNS CT policy router's connectivity with each other and the ADNS border routers.

These architectural and functional improvements just described are great for resiliency. What's the paradigm shift mentioned in the introduction?

Perhaps the greatest paradigm shift that ADNS Inc. III brings for career fleet communicators is that it is possible for a ship to change NOCs without any shore assistance. It is as easy as pointing a TACLANE at a different IP address. In other words, a ship receiving services from PRNOC could shift its own IP services to IORNOC without any assistance from shore. Furthermore, since each network enclave is independent of the others, a ship could get NIPRNet IP services from UARNOC and SIPRNet IP services from ECRNOC. Interestingly, neither the losing nor gaining NOC may notice the ship's shift right away because all of the routing is automatic. There are some questions and concerns that come with this paradigm shift:

--What about the ship's mail exchanger (MX) record? Don't we have to point it at the gaining NOC for the ship to receive IP services from that NOC? No, read my CHIPS article, "[Email Queues Can be a False Indicator](#)," [2] which explains why. In short, it's a best practice to have the MX record aligned with the servicing NOC to localize services to the greatest extent possible in one AOR.

--Aren't ships just going to shift to another NOC to avoid troubleshooting a problem that they are having? Maybe so. Nevertheless, if other units serviced by the same NOC are not having the same problem, then the problem will most likely follow the ship when it shifts to other NOCs so the ship will end up coming back to the original NOC and completing the troubleshooting to fix the problem anyway. It just costs the ship time with no connectivity for that enclave.

--Is the IP connection really the problem? Perhaps surprisingly, I have experienced that many outages reported as IP service outages actually have nothing wrong with the IP services. Instead, a loss of IP services is just a symptom and the true problem lies in the quality of the RF path. In other words, if your IP connectivity is suffering, consider whether it is one or more enclaves. If it's multiple enclaves, then most likely, you'll need to troubleshoot with the Standardized Tactical Entry Point site instead of the NOC. Perhaps the ship has moved enough that you just need to coordinate power balancing the RF with the Wideband Satellite Operations Center, as described above.

--What about configuration management and troubleshooting support? If ships are out of baseline configuration (TACLANEs or routers), their clandestine IP service shift may make the ship's connectivity worse. Furthermore, the NOC that didn't realize that it just gained another ship will not necessarily be prepared to assist in troubleshooting with that ship.

--*What about the inconsistencies between NOCs?* Note that not all four NOCs are exactly the same. If a ship has never been to a specific NOC before with its current configuration (including ADNS and Common Personal Computer (PC) Operating System Environment (COMPOSE) or Consolidated Afloat Networks and Enterprise Services (CANES)), the ship may find other routing issues when it shifts to that NOC for the first time. For example, the ship might not be listed in that NOC's firewall's trusted host table or the NOC may not be configured to advertise a route to DISA for that specific ship so no traffic from outside of the NOC will be able to find its way back to the ship.

--*Isn't there a formal request for shifting IP services between NOCs?* Even in emergent shifts, ships are required by Global Communications Information Bulletin (GCIB) 3A to submit an IP Service Request to establish services officially. GCIB-3A IP Service Requests are not just a formality. There is actually information conveyed in the IP Service Request that is necessary for the servicing fleet NOC to have.

Although not a paradigm shift, ADNS Inc. III affords greater flexibility to fleet commanders and Regional SATCOM Support Centers (RSSCs) and has proven itself more resilient than any other fleet IP services prior to ADNS Inc. III. If a ship has more than one SHF path, then each path should terminate at different STEP sites within the same AOR (so all of them are on the same ADNS CT policy router). Here are two of my anonymized experiences that demonstrate why path diversity within an AOR is the key to resiliency with ADNS Inc. III.

- The afloat unit, USS NEPTUNE, was up on both a 100 Mbps pier connection and up on an 8 Mbps SHF path simultaneously. The Standardized Tactical Entry Point site terminating NEPTUNE's SHF path suffered a Navy Tactical Wide-Area Network casualty so all ADNS Inc. III connections at that STEP site lost their connection to the NOC via the NTWAN. Since ADNS Inc. III employs an IP CT transport, OSPF on the border router having NEPTUNE's SHF path found a path to the NOC via NEPTUNE's pier connection. The other ships on that border router experienced a slightly slower connection because they were now sharing NEPTUNE's 8 Mbps SHF path. The other ships experienced greater latency because they made two round trips to satellites, their own and NEPTUNE's. But through all of this, they were still connected and able to continue their missions. Because NEPTUNE's 100 Mbps pier connection was nowhere near saturated the ship did not even realize it became part of the service provider's path.
- Several months later, the same standardized tactical entry point site mentioned above suffered a Defense Information System Network (DISN) casualty. Although USS NEPTUNE was not on her pier connection then, there were multiple ships with multiple, diverse SHF and commercial broadband satellite program paths in the AOR. Those ships that were terminated at this Standardized Tactical Entry Point site and terminated at another STEP site or via CBSP provided a path to the NOC for all of the other ships at the affected STEP site that did not have multiple or diverse paths. As before, although their connections were slightly slower and had greater latency, they were still connected and able to continue their missions because of the path diversity in place.

ADNS Increment III brings great, new potential to fleet communicators, if employed to its full set of capabilities. The advantages that ADNS Inc. III introduces with cipher-text transport enable greater improvements in automated routing, traffic shaping, load balancing, and the removal of inter-enclave dependencies. Its streamlined architecture simplifies troubleshooting and reduces the number of single points of vulnerability.

ADNS Inc. III empowers afloat units with less dependence on the shore for establishing connectivity and increased resiliency through path diversity within an AOR. When it comes to resilient networks,

availability is the key. ADNS Inc. III provides greater resiliency than any afloat IP provisioning system the U.S. Navy has seen before and it is up to us, as career fleet communicators, to use ADNS Inc. III to its greatest extent.

Author's Note: The descriptions and diagrams of the various communications paths and equipment strings are generalizations and do not include every piece of equipment in the path. The reason is two-fold: First, that level of detail is not necessary to convey the advantages of ADNS Inc. III over Inc. II or how to best employ ADNS Inc. III. Second, much of the equipment not included in the descriptions and diagrams varies from site to site in configuration.

Lt. Cmdr. Christopher B. Landis is approaching the end of his three-year tour at NCTS Naples, after which he will report to USS JOHN C STENNIS (CVN 74) as the Combat Systems Information Officer. Previously, he served on USS PAUL HAMILTON (DDG 60) as a Surface Warfare Officer in the Engineering Department and then in 2009 transferred to the Information Professional Officer Community to serve on USS INDEPENDENCE (LCS 2) Gold Crew as the Electronics Material Officer. He has a Master of Information Technology Strategy from Carnegie Mellon University.

The views expressed here are solely those of the author, and do not necessarily reflect those of the Department of the Navy, Department of Defense or the United States government.

References

[1] Eric Johnson. "USS Enterprise SHF Cross-Connect Configuration Increases Allocated Bandwidth by 100 Percent." *CHIPS*, April–June 2011.

<https://www.doncio.navy.mil/chips/ArticleDetails.aspx?id=2298>

[2] Christopher B. Landis. "Email Queues Can Be a False Indicator." *CHIPS*, January–March 2015.

<https://www.doncio.navy.mil/chips/ArticleDetails.aspx?ID=5930>