18th ICCRTS

C2 in Underdeveloped, Denied, and Degraded Operational Environments

Paper 088

Maritime Operations in Disconnected, Intermittent, and Low-Bandwidth Environments

Topics: Networks and Networking, Architectures, Technology, and Tools, Autonomy

Stephan Lapic
Spawar Systems Center Pacific
San Diego, CA
Email: stephan.lapic@navy.mil

Chris Meagher
Spawar Systems Center Pacific
San Diego, CA
Email: christopher.meagher@navy.mil

Daryl Ching
Spawar Systems Center Pacific
Pearl City, HI
Email: daryl.ching@navy.mil

Lester Chong
Spawar Systems Center Pacific
Pearl City, HI
Email: lester.chong@navy.mil

Isabelle Labbé
Communications Research Centre
Ottawa, ON
Email: Isabelle.labbe@crc.gc.ca

Sarah Dumoulin
Communications Research Centre
Ottawa, ON
Email: sarah.dumoulin@crc.gc.ca

Martin Jordan
Spawar Systems Command
San Diego, CA
Email: martin.jordan@navy.mil

Rob Thompson
Computer Science Corporation
Honolulu, HI
Email: rthompson@ati-security.net

Abstract

The navies of the US and its allies and coalition partners have come to be increasingly reliant on the availability of stable, high-bandwidth ship-to-shore satellite communications (SATCOM) to deliver network and application services. Even before recent recognition that satellite dependence presents tactical and operational vulnerabilities, work within the five eyes AUSCANNZUKUS alliance and Multinational Maritime Information Systems Interoperability (M2I2) coalition communities have progressed architectures designed for Disconnected, Intermittent, and Low-Bandwidth (DIL) environments. Motivations for these efforts have derived from not only concerns about the risk that satellites can be jammed or even shot down during hostilities, but also concerns about the cost and availability of satellites world-wide and to all partners in potential coalitions even in peacetime. In this paper, we report on solution sets developed for maritime operations in DIL environments, focusing on five areas:

- Enhancing line of sight communications paths,
- Improving multi-bearer routing,
- Implementing a distributed security architecture,
- Developing applications for distributed operations, and
- Developing a dynamic distributed database to support operations in DIL environments

Introduction

When the People's Republic of China shot down an ageing weather satellite in January 2007, it did not come as a complete surprise to defense watchers. Indeed, the United States Office of the Secretary of Defense in 2003 accurately predicted that China was "conducting research and development on a direct ascent anti-satellite system that could be fielded in the 2005-2010 timeframe."[1] This event, however, more than any other, highlights defense dependence on satellite communications (SATCOM). Not only has it been demonstrated that can they be shot out of the sky, satellites can be disabled by electro-magnetic pulse (EMP) and are vulnerable to jamming. Between the time of Operation Desert Storm in 1990-1995 and Operation Iraqi Freedom in 2001-2003, the US and its major allies moved from operations supported by satellites to ones which depended on full satellite integration as an inherent and essential component of Command and Control (C2), Intelligence, Surveillance and Reconnaissance (ISR), and Positioning, Navigation, and Timing (PNT) systems.[2,3] The success of future network-centric operations and warfare require that we overcome this dependence.

Among allied and coalition Navies, satellite dependence manifests itself in architectures which rely on reach-back to the national Network Operations Center (NOC) ashore. The

[1] Office of the Secretary of Defense, FY04 Report to Congress on PRC Military Power Pursuant to the FY2000 National Defense Authorization Act, 28 July 2003
[2] Matthew E. Grant, Space Dependence – A Critical Vulnerability of the Net-Centric Operational Commander, Naval War College, 17 May 2005
[3] J. Wilson, The Ultimate High Ground, Armed Forces Journal, January 2004

NOC is the gateway between the satellite land-earth stations and terrestrial operational and strategic networks. The NOC also provides network services and hosts application servers which are accessed by deployed ships. In NOC-centric architectures, when connectivity to the NOC is lost, communication is impossible and applications will not work.

Aware of this vulnerability, work within the five eyes AUSCANNZUKUS maritime alliance and the M2I2 coalition community have progressed architectures to address satellite-denied environments and where the available communications paths are intermittent and low-bandwidth. Pursuit of NOC-less architectures was motivated not only be awareness of the threats during wartime operations but also by the recognition that, even in peacetime, satellites may not be available to all potential partners and the availability of alternatives to SATCOM can reduce operating costs. Barring operations in a conflict where satellites are denied, there is nothing that points out our reliance more than engaging in operations with a partner that does not have satellite connectivity or cannot afford continuous satellite coverage.

Much of this work has been conducted under the aegis of the Trident Warrior (TW) exercises, the premier C41 sea trials for the US Navy. Trident Warrior was first conducted in 2003 and allies and coalition partners have acted as observers and full participants since 2004. The AUSCANNZUKUS nations and other partners have sponsored technologies, provided operational and engineering stuff to support the trials, and have embedded their own national platforms within the US-led TW Task Group.

A significant step in migrating from a focus on the NOC ashore to the Task Group afloat has been the deployment of alternatives to satellite communications: line of sight (LOS) and extended line of sight (ELOS) communications. To provide this capability for coalition operations, Ultra-High Frequency (UHF) Subnet Relay (SNR) and High Frequency (HF) Internet Protocol (IP) are installed between the existing coalition Combined Enterprise Regional Information Exchange System (CENTRIXS) router and the shipboard radio and cryptographic equipment. These two systems provide mobile ad hoc self-organizing, self-healing communications capability.

Installation of the LOS/ELOS bearers, however, is only the first step. To make effective use of such a network requires the design and deployment of services and applications that can effectively operate using it. The end goal is self-organizing, self-healing *applications* that are robust to denied, intermittent, and low-bandwidth (DIL) communications links. To make this vision a reality also required a routing redesign, a reconsideration of the security architecture, and an ability to dynamically reconfigure shipboard applications to support distributed operations via a dynamic distributed database.

The next sections provide details of the LOS and ELOS bearers, the routing architecture, the security architecture, the mechanisms developed to implement dynamic, distributed applications, and the dynamic distributed database. A final section provides conclusions.

Line of Sight Communications

Figure 1 depicts the shipboard equipment configuration employed on shipboard networks when LOS and ELOS bearers are added. On US and allied ships, coalition enclave traffic is typically IP-encrypted and tunneled over the ship's gateway router, via SATCOM, to shore. In the case of the US, the gateway is the Automated Digital Network System (ADNS) router. This router multiplexes the IP-encrypted traffic from multiple enclaves (only the coalition enclave is depicted in the figure) and may divide the traffic among multiple SATCOM bearers if more than one is available. In the event of degraded or unavailable SATCOM services, the ship would be effectively disconnected from the rest of the task group.

To support operations, LOS and ELOS mobile ad hoc bearers are connected directly to the coalition router to provide ship-to-ship connectivity. While connecting these bearers through the gateway router might appear to be more consistent with the existing architecture, connecting the gateway routers between ships of different nations using LOS/ELOS connections is not currently possible because they are operated at different classification levels. Both UHF SNR and HF IP have a controller and an external modem, but use existing shipboard cryptographic equipment, as well as radio, coupler, and antenna infrastructure. Re-use of as much existing infrastructure as possible is a driving force in any practical implementation of capability evolution.

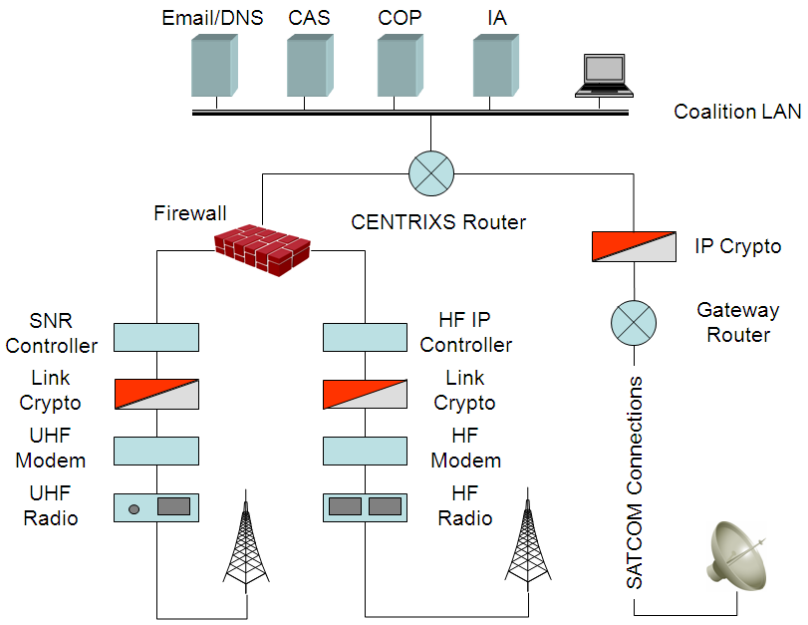

Figure 1: Shipboard Equipment Configuration

UHF SNR provides a synchronous time-division media access (TDMA) service and layer-2 relay in multi-hop topologies[4]. HF IP provides a wireless token bus media access service and relays at layer 3[5]. Lacking crypto bypass, both systems are configured to transmit at a particular burst rate, set based on the anticipated conditions. Narrowband SNR is typically set to burst at 51.2-64 kbps in a 25 kHz channel. Narrowband HF IP is typically configured to burst at 6.4 kbps or 12.8 kbps in either a 3 kHz or 6 kHz channel, depending on whether the HF equipment is operated in single side band (SSB) or independent side band (ISB) mode. These channel allocations – 25 kHz for UHF and 3 or 6 kHz for HF – are standard for tactical voice.

In Trident Warriors 10, 11, and 12, successful demonstrations higher capacity LOS and ELOS bearers were conducted. With an upgrade to the UHF modem in figure 1 and a new radio when necessary, high data rate (HDR) SNR was operated at 384 kbps in a 100 kHz channel and up to 1.92 Mbps in a 500 kHz channel. By upgrading the HF modem to one which implemented MIL-STD-188-110C appendix D[6] and installing a wideband HF radio, data rates of up to 96 kbps in a 24 kHz channel were achieved. This capacity is shared between the network participants. For comparison, a single International Maritime Satellite (INMARSAT) Fleet Broadband connection provides up to 432 kbps, which is shared between serial data and multiple IP enclaves, typically resulting in 32-64 kbps dedicated for coalition IP data.

Besides bandwidth, delay is another important element of bearer performance. During TW trial events, HDR SNR network latencies were measured using IP ping and analysis of TCP handshakes at approximately 1.5 seconds while WB HF latencies were approximately 2 seconds. Several elements in the LOS/ELOS systems conspire to add latency including the legacy cryptos and long interleaver times to combat multi-path. The half-duplex nature of the LOS/ELOS channels further exacerbate problems with delay so that WAN optimizers and special TCP modifications are often deployed to improve performance. These latencies, however, are comparable with those experienced over SATCOM which typically have ping times exceeding 2 seconds themselves. Whether SATCOM-centric or over LOS/ELOS bearers, tactical maritime communications for coalitions are characterized by lower bandwidth and higher latency than terrestrial networks. Maritime forces must – and do – operate in this challenging communications environment.

A critical consideration for the employment of LOS/ELOS links is their connectivity range. UHF connections are limited to the RF horizon, which is approximately 20 nautical miles (NM) with the antenna heights typically employed on ships. HDR SNR has been demonstrated to the RF horizon in TW 11. HF connections, on the other hand, extend past LOS due to atmospheric refraction. HF ground wave connections can extend

[4] STANAG 4691, "Mobile ad hoc relay line of sight networking," edition 1 draft 2, NATO standardization agreement, 2010
[5] STANAG 5066, "Profile for HF Data Communications," edition 2, draft 2, NATO standardization agreement, 2006
[6] MIL-STD-188-110C, "Interoperability and Performance Standards for Data Modems," Department of Defense Interface standard, 2011

to approximately 150 NM over sea water. WB HF was demonstrated out to 75 NM in TW 12, that being the furthest connection that was attempted.

Figure 2: WB HF IP TW 11 Test Locations

HF IP was designed to operate over ground wave so that is the primary focus of our investigations. However, its possibility as a long-haul strategic back-haul alternative to SATCOM is intriguing. During TW 11, wideband HF IP trials were conducted over sky wave between four locations in the US and Canada. These are illustrated in figure 2. Operation of HF IP using the MIL-STD-188-110C Appendix D waveform was successful in this setting, proving to be robust to both ultra-short and short interleavers. The WB HF IP sky wave network also supported the critical CENTRIXS collaboration applications – Sametime and Domino – in this setting. Future testing of WB HF IP ship-to-shore is planned for the future.

Routing Architecture

When LOS/ELOS links are available, the objectives of the routing design are that ship-to-ship traffic would be preferentially routed using the direct connections, even when SATCOM is also available, and that shore-to-ship traffic be sent via direct satellite whenever possible and only routed via another ship (to be relayed via LOS) when direct SATCOM is down. Furthermore, traffic should be routed via a same nation ship if possible. Routing should be via another nation's ship as a last resort. When more than one LOS/ELOS bearer is available, for example UHF SNR and HF IP, the higher capacity bearer (UHF SNR) would be used first.

The routing architecture in the coalition setting is complicated because, while Open Shortest Path First (OSPF) is the routing protocol used between ship and shore, the expectation is that Border Gateway Protocol (BGP) be employed between different national NOCs. Since routes learned by interior gateway protocols such as OSPF are

typically preferred to those learned from exterior gateway protocols such as BGP, a naïve design would have, say, traffic from the US to a Canadian ship sent first to a US ship and then relayed over LOS to the destination, if the LOS connection was available, rather than sending the traffic to the Canadian NOC to have it delivered via SATCOM, clearly the preferable alternative.

The multi-bearer routing architectures developed to support LOS networking use OSPF tags and BGP community strings to communicate additional routing information. The shipboard routers mark routes that they generate, i.e. the shipboard local area network (LAN) and those they receive through LOS/ELOS connections with different OSPF tags so that different routes can be distinguished at the NOC. The ship has the best visibility into which routes are locally originated, versus those it receives from SNR/HFIP, and can pass that information up to the NOC.

The NOC meanwhile preserves the OSPF tags passed up from the ship and converts them to BGP community strings so the routes can remain distinguishable. Within BGP, autonomous system (AS) prepend is used to add additional AS paths to routes learned from the LOS/ELOS connections. This insures that each ship's own national NOC will remain the preferred path from shore when SATCOM connectivity is available and other national NOCs will serve as backups only when desired. At the same time, each national NOC uses BGP route maps to assign a BGP weight of 40000 to other country's ship routes when advertised from their respective NOCs. This is done so that the BGP path will be preferred over the OSPF routes passed from the ships. Adjusting the BGP weight prevents the NOC from attempting to use the LOS to reach ships while their SATCOM, with higher bandwidth, is available.

Further, specific subnet and summarizes networks to leverage routing's ability to use prefix length as an additional decision point for routing. Ships advertise their specific LAN subnet (/26) up to their local NOC while advertising a summarized route into the OSPF process running over the LOS links for other NOCs to advertise out as a backup path. Out on the wide area network (WAN), two routes are then advertised out: the specific subnet and the summarized class C. The router treats the two, each with different subnets (/25 and /24), as two distinct different routes and allows both to propagate. Routers prefer the more specific route (e.g. /26 over /25, and /25 over /24). The local NOC that advertises the shorter prefix route becomes the preferred path to the ship. The NOC advertising the summarized route with the longer prefix becomes the backup path to the ship.

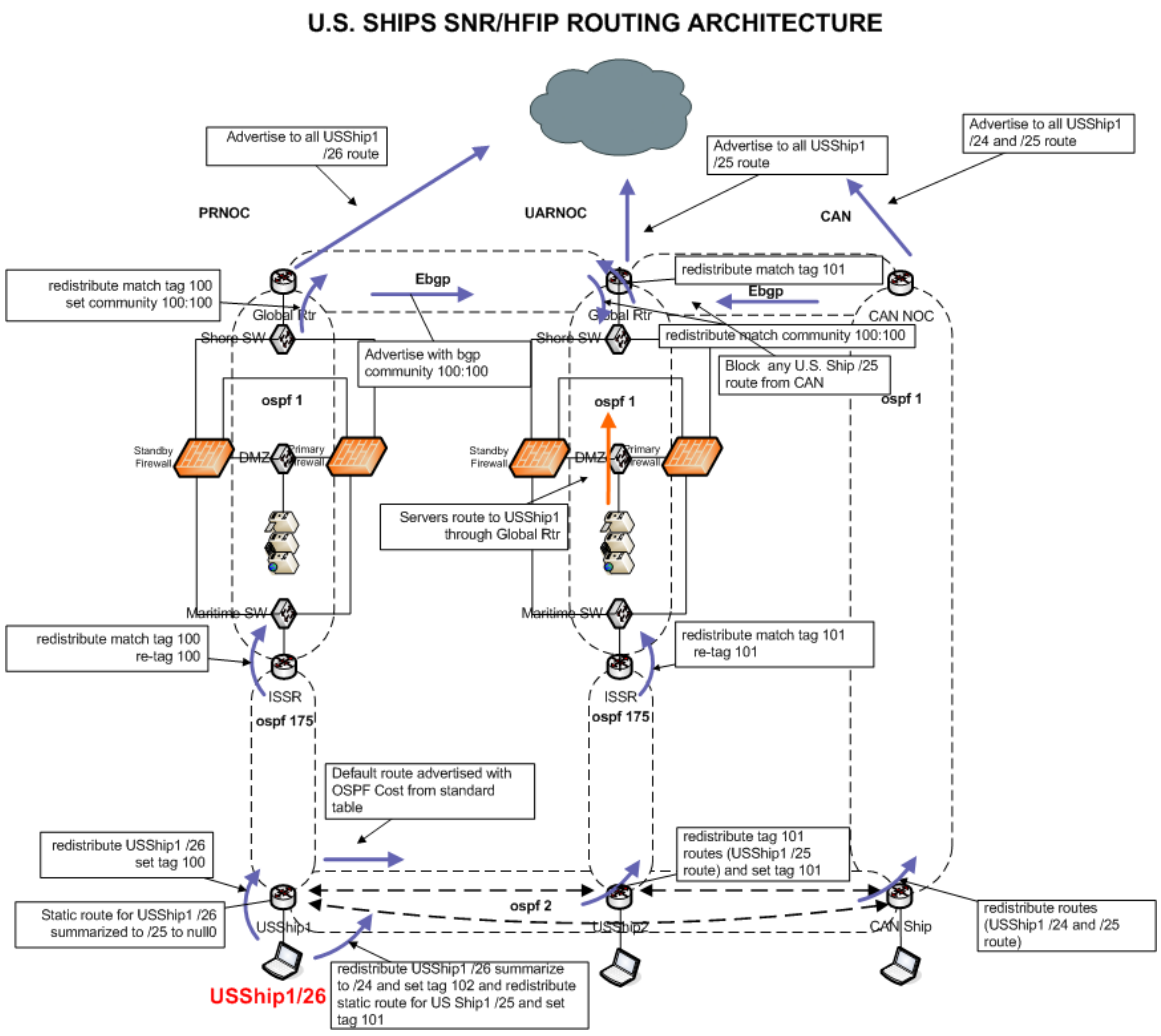

Figure 3: Routing design for inclusion of LOS/ELOS links

Figure 3 summarizes the routing architecture from the US point of view. A similar diagram can be constructed for each of the coalition NOCs. The full details are described in the M2I2 ad hoc routing standard.[7]

Although it was critical from a practical point of view that the addition of a ship into a LOS-equipped coalition task group require no additional configuration on the ship routers, the current design requires additional configuration of the conditional BGP forwarding at each of the national NOCs and coordination between them to complete this. Thus, it is a longer term objective to further simplify and automate this process.

[7] M2I2 Ad Hoc Networking Focus Group, "M2I2 ad hoc network parameter standards for LOS networking standards on CENTRIXS," version 17, 11 September 2012

Security Architecture

The introduction of direct ship-to-ship connections to the current SATCOM-centric architecture presents new security vulnerabilities. These include increased risk of network penetration and host compromise, as well as increased opportunity to compromise traffic delivery. The LOS/ELOS links employ a broadcast media which exposes vulnerability to eavesdropping and jamming. The security risks associated with LOS connections become even more pronounced when multiple nations are involved, as they are in a coalition network because the different national components are under separate administrative controls. Technologies and procedures must be deployed to mitigate these risks.

Table I indicates some of the most severe risks associated with LOS/ELOS connection and some of the mitigation techniques available. Multi-level defense in depth is called for to protect the networks, the hosts, and the data in transit. Network monitoring, management, and support for recovery must be part of the solution as well. Finally, it should be noted that security mechanisms deployed on ships must not create an undue administrative or resource burden for shipboard personnel; as a consequence shipboard security solutions may differ from their shore-based counterparts.

Table I: Security risks and mitigations for LOS/ELOS networks

Risk \ Mitigation	Traffic screening	Packet inspection	Anti-virus checking	Intrusion detection & prevention	Authentication	Access control	Monitoring	Backup & recovery	Encryption	Patch management
Network Penetration	✓	✓			✓	✓	✓			
Malicious logic		✓	✓	✓						✓
Data Integrity					✓	✓	✓			
Interference							✓			
Eavesdropping									✓	
Indirect exposure	✓									
Spoofing	✓	✓			✓	✓				
Human error					✓					
HW/SW error							✓	✓		
Cryptanalysis									✓	
Misuse						✓	✓			

With SATCOM only, the security perimeter is maintained at the NOC. Ship networks are protected by implementing access controls, firewalls, anti-virus (AV) scanning, intrusion detection and prevention systems, network monitors, and other mechanisms on shore. With LOS connections, the security perimeter must be moved to the ship itself.

To implement network defense in depth on the LOS-equipped ships, several mechanisms are implemented. These are illustrated in fig. 1. The type-1 cryptographic equipment used to provide data confidentiality, communications security (COMSEC), and transmission security (TRANSEC) is already there. A firewall appliance is placed between the LOS and ELOS bearers and the coalition router to perform traffic screening, deep packet inspection, and network intrusion detection. Additional screening is employed by implementing access control lists on the coalition router. On the network hosts and servers, malicious code detection and host-based intrusion detection and prevention are implemented. This is supported by a patch management system indicated by the IA server shown in fig. 1 which automates the downloading of Windows patches and Symantec AV definition files from shore and updating of deployed hosts and servers.

Distributed Applications

The most critical applications used for coordination within a maritime coalition are email, a web-based suite of tools known as Collaboration at Sea (CAS) based on IBM/Lotus Sametime and Domino[8], and Common Operational Picture (COP). The collaboration tools, first designed for the commercial, fixed infrastructure environment have been retrofitted into the maritime tactical space. The fit becomes particularly problematic when working in a satellite-denied scenario.

In a SATCOM-only environment, ships send Microsoft Exchange email directly to the shore and the home mail server for the ships is located at the NOC. For ships to be able to exchange email directly in a SATCOM-denied environment, each LOS-capable ship must have its Domain Name Service (DNS) zone files reconfigured to support conditional forwarding to the other LOS-capable ships in the Task Group. In initial trials of this capability, the reconfiguration was done manually. This solution, while effective, scales poorly if support of unplanned or large coalition is required. Recent efforts have automated this capability. Automation requires the distribution of a directory that contains all the information necessary to do this on each ship and the implementation of software scripts to make the changes.

Domino has been the biggest challenge by far. Domino provides replication services between servers to maintain consistent document databases that are web-enabled. Domino replication is controlled by connection documents. These are essentially static routes at layer 7 that define the Domino replication architecture. The connection documents can only be changed by system administrators, usually at the NOC. Access to them is tightly controlled.

The hub-and-spoke replication architecture used in CENTRIXS today breaks down when connectivity to the NOC is lost. Currently, ships point their Domino servers at the NOC and replicate only with the server ashore. A manual method of changing the connection documents shipboard which did not grant full administrator access was successfully trialed during TW 11. Engineers or operators on a small ship could "with the push of a

[8] Mark Lenci, "CS 101: Collaborating at Sea: A Domino Based Carrier Battle Group Solution," USN Technical Report, 23 February 2000

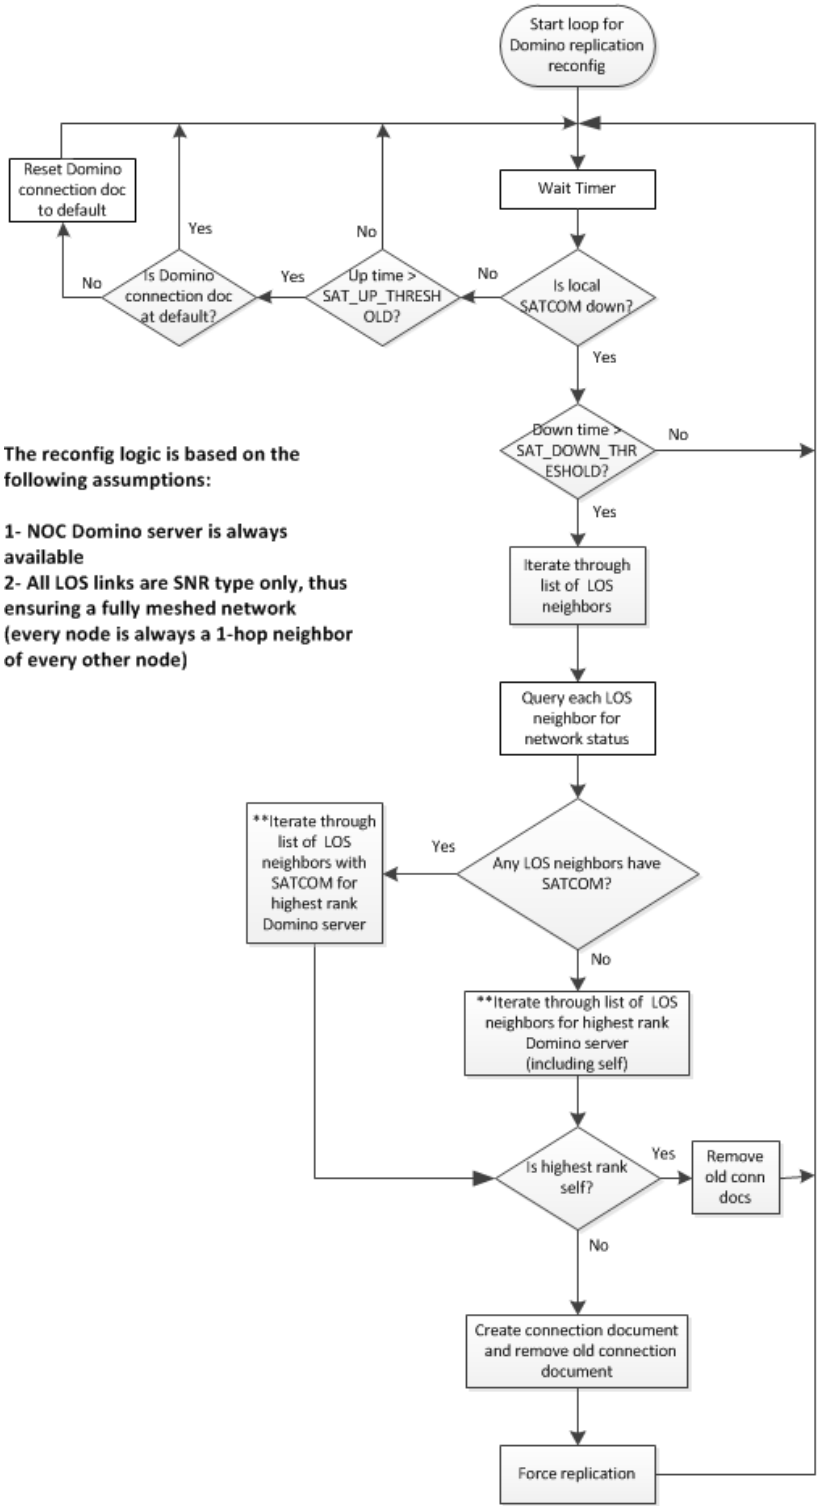

Figure 4: Domino Reconfiguration Logic

button" point Domino at either the shore or another ship in the task group acting as the hub server. The human still needed to decide where to point, based on the communications paths available at the time.

The next step is to remove the human from the loop. To implement Domino as a self-organizing, self-healing mobile ad hoc application requires replication logic be developed that decides where to point and when to do it. A software agent has been developed that runs on the Domino server that automates the process. Without human intervention, the agent determines when Domino should replicate with the shore, when it should replicate with another ship in the Task Group and if so who, or when it should act in the role of the large deck and serve as the home office for the Task Group afloat in the event of satellite-denial. Figure 4 shows the Domino reconfiguration logic. This decision is based on local information about connectivity and link characteristics and distributed data on the capabilities

Although specific to each application, similar logic is implemented for other critical applications.

Dynamic Distributed Database

The development of a Dynamic Distributed Database (DDD) is a core element enabling the distributed operation of networks and applications, as described in this document. The DDD is a database containing all the relevant information required to reconfigure the applications, routing, and other network services within a Task Group when SATCOM is denied. Software agents that execute within IBM Domino server software obtain network status information from ship routers, populate the database, and use the information contained within it to automatically reconfigure DNS, Exchange, and Domino connection documents to support ship-to-ship networking in a satellite-denied environment. The entries in the database also include performance information to help optimize application configuration. Figure 5 gives a snapshot of entries in the DDD. In current testing, the DDD is replicated using Domino itself.

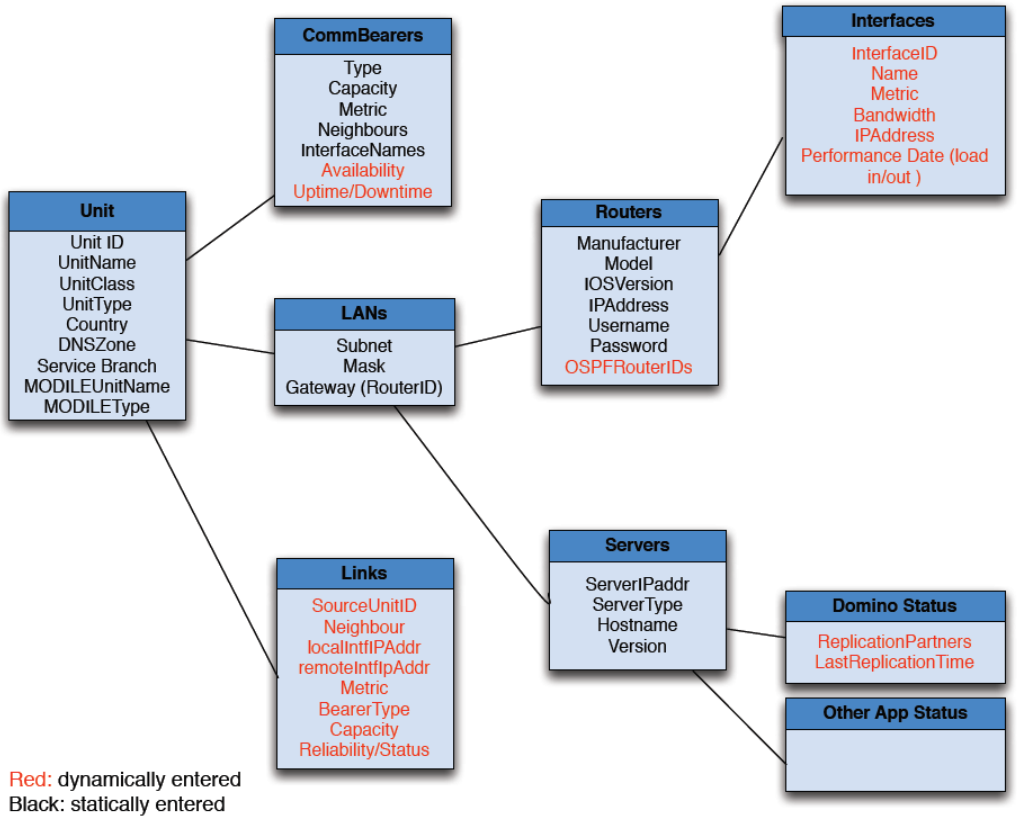

Figure 5: Dynamic Distributed Database for application and network reconfiguration

Conclusion

An end-to-end capability has been developed and demonstrated that permits coalition maritime operations to continue in the absence of satellite connectivity. Network designs that integrate LOS/ELOS links with SATCOM have been proven to improve capacity and eliminate single points of failure. Routing, security architectures, and other network services have been developed and tested in order to demonstrate the ability to successfully operate critical C4I applications in support of distributed and disconnected operations. This work continues; data, services, and technical capabilities must be continually re-examined and improved in order to ensure that the correct information is moved forward to the tactical edge in order to meet the operational requirements for coalition forces. Multi-national research and development groups, combined with realistic experimentation venues (i.e. Trident Warrior), provide us with the opportunities to examine and improve Coalition Maritime networking interoperability and information sharing now and in the future.