

Email Queues Can Be a False Indicator

Christopher B. Landis

[January-March 2015](#)

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

TAGS: [CTO: Networks](#), [InfoSharing](#), [NNE](#), [Spectrum](#), [Telecommunications](#), [Wireless](#)

Navy communicators must consider multiple factors to correctly assess the status of afloat email systems

At 0900, 25 emails are on queue at the European Central Region Network Operations Center (ECRNOC) for USS Neptune. At 1000, 25 emails are still on queue at ECRNOC. Has Neptune lost her connection to the Department of Defense Information Networks (DoDIN)? Is the Exchange Server onboard down? Were no emails delivered to Neptune between 0900 and 1000? Or, was she connected throughout and received 100 emails? Navy communicators must consider multiple factors to correctly assess the status of afloat email systems.

Email queues are FIFO (first-in, first-out). The first email received at the NOC would be the first one delivered to the ship. Between 0900 and 1000, if ECRNOC received 100 new emails to deliver to the Neptune, the 25 emails that were queued at 0900 would be delivered first and then the first 75 of the 100 new emails. By 1000, Neptune's email queue at ECRNOC would consist of the last 25 of the 100 new emails.

There could be multiple email queues. There are four fleet NOCs: ECR, Unified Atlantic Region (UAR), Indian Ocean Region (IOR), and Pacific Region (PR). When a ship shifts its IP service provider to another NOC, either UAR or PRNOC will "shift" that ship's email Domain Name Service (DNS) record, or MX record, to the gaining NOC. In reality, the MX record does not shift (or move) anywhere. It stays on the DNS server and replicates to other DNS servers for redundancy.

An MX record shift simply consists of changing the MX record's IP address attribute to designate a different NOC's email server to receive email on behalf of the ship. Based on this standard operating procedure, a communicator might incorrectly assume that the designated NOC is the only location that queues emails bound for this ship. In reality, the MX record only affects emails coming from locations other than the four fleet NOCs. For example, an email coming from a ship receiving IP services from IORNOC bound for the Neptune (receiving services from ECRNOC) would queue on IORNOC's email server instead of the ECRNOC's. This is because NOC email servers are configured to send mail to ships, regardless of where they are receiving their IP services. In other words, Neptune may have emails queued at the other NOCs in addition to the 25 emails queued at ECRNOC in the example above.

Why do we shift MX DNS records? On a system as critical as DNS — the Internet would almost cease to exist as something useful without DNS — it is generally better not to make changes where changes are not needed. One line of thought is that we should set fleet MX DNS records once and not change them. Each time a NOC operator changes the DNS configuration, there is an opportunity for error. But the status quo argument in favor of shifting MX DNS records is flawed: "There will be too much latency if we do not change a ship's MX DNS record to point to the NOC from which it receives IP services." The counterargument: "Who do you expect to experience any latency

whatsoever when an email travels between a NOC's email server and a ship's email server?" The answer is no one.

Furthermore, typical NOC-to-NOC latency is significantly less than the latency experienced when communicating with a satellite in geosynchronous orbit, such as the Defense Satellite Communications System (DSCS) and Wideband Global Satellite (WGS). The real reason that we shift MX DNS records is to maximize fleet commanders' lines of communication to their ships through localization of resources. MX DNS record shifts enable one additional line of communications in a communications-degraded environment, as in an area of responsibility (AOR) isolation scenario. Even though email is non-real-time and best effort (not guaranteed) delivery, it may become the most effective path for a commander to direct operations because of other communications degradations, and it needs to be able to work without resources from other AORs.

Afloat units use virtual email routing. All ships are configured to send their emails to the same IP address. Each NOC has this virtual IP address in its local network. It is not the IP address of any server, per se, but a load balancer that simply forwards email on to be virus scanned and queued on one of its NOC's many email servers. This is how email for a ship receiving IP services from ECRNOC can have email on queue at IORNOC. By the way, a NOC's email queues are aggregately counted across all of its email servers. In summary, to get an accurate count of the number of emails on queue for a ship, one must consider the queues at each of the four NOCs.

SPAWAR Systems Center Atlantic knows where the email is located. NOCs monitor their email queues on an out-of-band network, meaning that there is no fleet IP traffic, including email, on that network because it is used for network and security management, and it does not connect to the "production" network that carries fleet traffic. It does; however, connect to a NOC laboratory located in Charleston, South Carolina, at Space and Naval Warfare Systems Center (SPAWARSYSCEN) Atlantic. With each of the out-of-band NOC networks connecting to SPAWARSYSCEN Atlantic, SPAWARSYSCEN Atlantic automatically becomes the ideal location to host a server that aggregates each of the NOCs' email queues. This information could then be made available to fleet communicators as a more accurate count of emails on queue for their ships. But this still does not solve the problem of consistent queue sizes, highlighted in the opening example of the Neptune.

Show me the data. The missing bit of information (pun intended) to develop an accurate picture of a ship's connectivity to the DoDIN based on her email queue is the measurement of data transferred to the ship from that queue over a certain time period. With this metric and queue information from SPAWARSYSCEN Atlantic, communicators can better assess a ship's status: "Neptune had 25 emails on queue across all four NOCs at 0900, the same number at 1000, and 42 megabytes of emails were transferred from NOC email queues to the Neptune between these two times." Now the common misconception of no connectivity and zero emails transferred based on consistent email queues is immediately ruled-out. To make the consolidated email queues at SPAWARSYSCEN Atlantic worthwhile, the amount of data transferred with respect to the time period must be part of the information.

Disclaimer. The Navy should not create another real-time or near real-time sensor to determine when ships lose or regain their connections. This is why we have WhatsUp Gold or similar network management software. However, fleet communicators must understand the capabilities and limitations of using fleet email queues as an indicator of DoDIN connectivity.

Can Navy communicators rely on email queues alone to determine the status of a ship's connectivity to the DoDIN? Currently, we cannot. Nevertheless, we may be able to improve our understanding of

afloat email systems and ship connectivity with consolidated queue information and throughput data from the NOC laboratory at SPAWARSYSCEN Atlantic. Hopefully, this article raises awareness of our fleet email systems for fleet communicators and corrects some common misconceptions.