Subsecond vrrp timers

16 July 2003

Karen O'Donoghue
Naval Surface Warfare Center Dahlgren Division
OdonoghueKF@nswc.navy.mil
(540) 653 – 1567

Bob Hott
Naval Surface Warfare Center Dahlgren Division
HottRW@nswc.navy.mil
(540) 653 – 1497

Background

- The shipboard environment has many of the same communications requirements that face commercial organizations
  - Support for voice, data, video and imaging, etc…
  - Security, Performance, Fault Tolerance (survivability)
- But, there are a number of requirements beyond typical shore based networks…
- Historically, each shipboard system would address its communications requirements in its own way, typically with built-to-spec systems and point-to-point links
  - Advantages: met initial requirements
  - Disadvantages: Scalability, technology refresh, cost ?

Architecture Evolution
(Equipment)

Federated

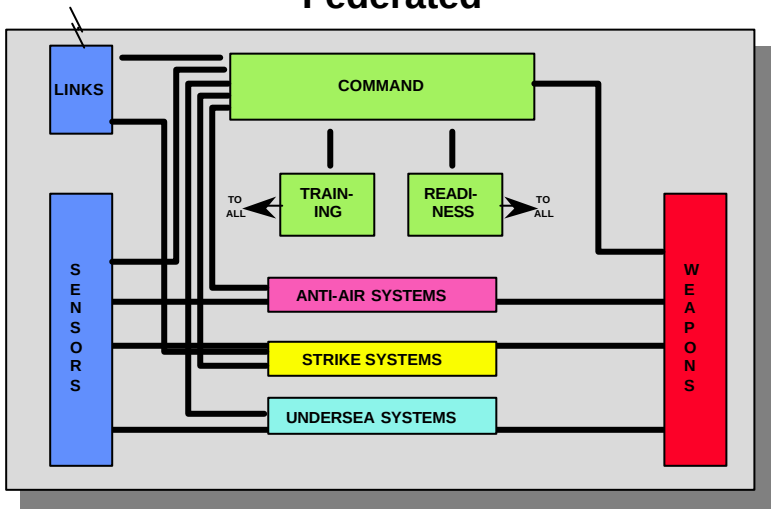

Network-of-Networks

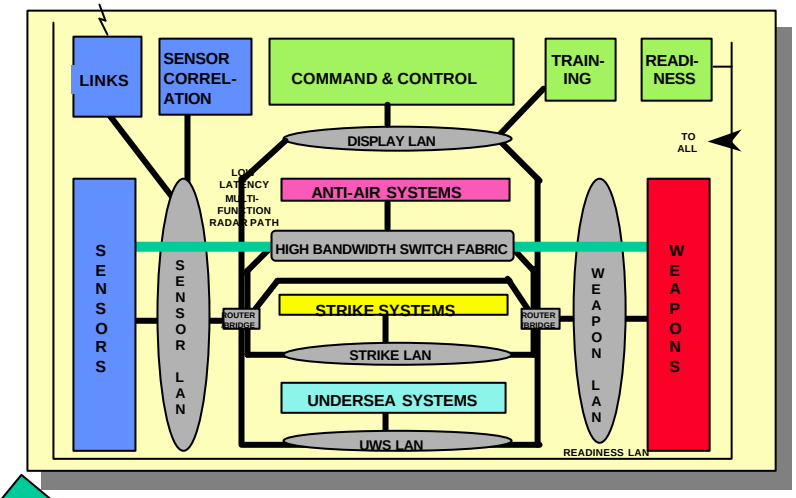

Total Ship Computing

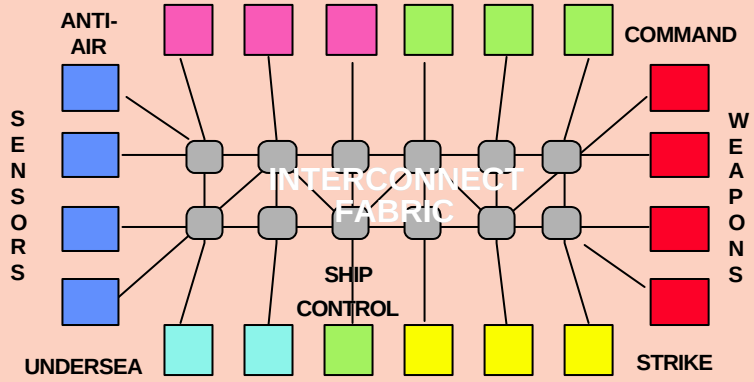

Technical characteristics

- COTS-based open equipment
- Fault tolerant, scaleable
- Instrumental and verifiable
- Uniform high throughput, low latency
- Guaranteed quality of service
- System-wide resource sharing
- Dynamic resource allocation

Requirements Summary

- The U.S. Navy needs to build shipboard networks that can provide:
  - End-to-end latencies of less than 200 microseconds
  - Scalable end-to-end throughput greater than 100 Mbps
  - End-to-end outage of less than 1 second due to a network fault (i.e., reconfiguration time includes the detection and recovery from the network fault)
  - The ability to interface with other open, commercial (COTS-based) networking solutions (e.g., routers, switches, encryption devices)

Current Navy Direction

- Current direction for meeting our survivability requirements:
  - Dual homed end stations (Ramix proprietary solution)
  - HSRP with subsecond timers
- We have done performance and functionality testing on HSRP/ESRP/VRRP.
- We understand there are negative implications for subsecond timers (including instability and scalability), but we still feel these are useful in a certain subset of applications with an appropriate engineered network.

Possible vrrp approaches

- We need the ability to use subsecond timers, and we would like this to be supported in a standard manner.
- Possible approaches within vrrp specification
  - Add a flag (option) to indicate advertisement interval is in milliseconds instead of seconds
    - Carried in the vrrp frame
    - Configuration option (state machine clarification)
  - Modify the units of advertisement interval

Approach

- Don’t have proposed text today (… my apologies)
- What needs to be done (if you want to proceed)…
  - Decide on basic approach
    - Configuration option/state machine clarification (?)
  - Specification changes (preliminary impact review)
    - Advertisement_Interval
    - Master_Down_Timer
    - Adver_Timer
  - Add caveat/warning text about usage of this type of option – maybe followup with an information RFC with guidance for usage.
- We are willing to provide this text in the near future if necessary.

- Thoughts…