

Network Working Group
Request for Comments: 1679
Category: Informational

D. Green
P. IreY
D. Marlow
K. O'Donoghue
NSWC-DD
August 1994

HPN Working Group Input to the IPng Requirements Solicitation

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document was submitted to the IETF IPng area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng area of any ideas expressed within. Comments should be submitted to the big-internet@munnari.oz.au mailing list.

Executive Summary

The Navy's High Performance Network (HPN) working group has studied the requirements of mission critical applications on Navy platforms. Based on this study, three basic categories of issues for IPng have been identified. The assumptions identified include accommodation of current functionality, commercial viability, and transitioning. The general requirements identified include addressing, integrated services architecture, mobility, multicast, and rapid route reconfiguration. Finally, the additional considerations identified include fault tolerance, policy based routing, security, and time synchronization. The HPN working group is interested in participating with the IETF in the development of standards which would apply to mission critical systems. In particular, the HPN working group is interested in the development of multicast functionality, an integrated services architecture, and support for high performance subnetworks.

1. Introduction

The HPN working group has been established to study future network architectures for mission critical applications aboard Navy platforms. As a result, the HPN working group is interested in the results of the IPng selection and development process. This document is a product of discussions within the HPN working group.

The purpose of this document is to provide what the HPN working group perceives as requirements for an IPng protocol set. Many of the necessary capabilities exist in current Internet and ISO network protocols; however, the HPN working group has identified needed capabilities that are beyond the existing standards.

The HPN working group has identified three categories of topics for discussion in this document. The first category is assumptions or those topics that the HPN working group believes the IPng process will solve satisfactorily without specific Navy input. The second category is general requirements. These are capabilities that are felt to be insufficiently addressed in existing network protocols and of key importance to Navy mission critical applications. Finally, a set of additional considerations has been identified. These are also issues of importance to the HPN working group. However, no guidance or specific requests can be provided at this time.

2. Background

The US Navy has set up a program through the Space and Naval Warfare Systems Command called the Next Generation Computer Resources (NGCR) Program. The purpose of this program is to identify the evolving needs for information system technology in Navy mission critical systems. The NGCR High Performance Network (HPN) working group was recently established by the NGCR program to examine high performance networks for use on future Navy platforms (aircraft, surface ships, submarines, and certain shore-based applications). This working group is currently reviewing Navy needs. The requirements provided below are based on the HPN working group's current understanding of these Navy application areas. The application areas of interest are further examined below. The time frame for design, development, and deployment of HPN based systems and subsystems is 1996 into the twenty first century.

Three general problem domains have been identified by the HPN working group. These are the particular problem domains within a mission critical environment that the HPN working group is targeting. The first is a distributed combat system environment. This problem domain is analogous to a collection of workstations involved in many varied applications involving multiple sources and types of information. Analog, audio, digital, discrete, graphic, textual, video, and voice information must be coordinated in order to present a single concise view to a commander, operator, or any end user. The second problem area highlights the general internetworking environment. The task of moving information to many heterogeneous systems over various subnetworks is addressed. Finally, the problem of providing a high speed interconnect for devices such as sensors and signal processors is identified. [1]

2.1 Application Area

The application area of HPN is the communication network which is a component of the mission critical systems of Navy platforms. The expected end points or users of the HPN include humans, computers, and the many devices (cameras, etc.) found on such platforms. The function of these end points includes sensor input, signal processors, operator consoles, navigation systems, etc. The endpoints are typically grouped into systems both on platforms and at shore-based sites. These systems perform functions including long range planning, analysis of sensor information, and machinery control in real-time.

Information types that have been identified as required by the HPN working group include voice, live and pre-recorded audio ranging from voice to CD quality (e.g., from sensors), video (1 to 30 frames per second in both monochrome and color), image data (static or from real-time sensors), reliable and connectionless data transfer, and very high-bandwidth (gigabits per second) unprocessed sensor data.

2.2 Services

Another way of categorizing the HPN application area is by considering the user services that need to be supported. Some of these services are the following:

1. process to process message passing
2. distributed file and database manipulation
3. e-mail (both within the platform and off the platform)
4. teleconferencing (with the platform, between platforms, and across the Internet)
5. video monitoring of various physical environments
6. voice distribution (as a minimum between computer processes and people)
7. image services
8. time synchronization
9. name or directory services
10. network and system management

11. security services (support of multilevel data security, privacy and protection)

3. Assumptions

The assumptions documented below are concerns that the HPN working group presumes will be accommodated in the IPng process. However, they are of enough importance to this working group to merit identification.

3.1 Accommodation of Current Functionality

The IPng protocols need to provide for at least the existing functionality. In particular, the following issues have been identified.

- 1) The IPng protocols need to provide for the basic connectionless transfer of information from one end-point to another.
- 2) The IPng protocols need to support multiple subnetwork technologies. This includes but is not limited to Ethernet, FDDI, Asynchronous Transfer Mode (ATM), Fiber Channel, and Scalable Coherent Interface (SCI). These are the subnetwork technologies that are of particular interest to the HPN working group. Ideally, IPng protocols should be subnetwork independent.
- 3) The IPng protocols need to support hosts that may be multihomed. Multihomed in this context implies that a single host may support multiple different subnetwork technologies. Multihomed hosts must have the capability to steer the traffic to selected subnetworks.
- 4) The IPng process needs to recognize that IPng may be only one of several network protocols that a host utilizes.
- 5) The IPng process needs to provide for appropriate network management in the finished product. Network management is of vital importance to the applications of interest to the HPN working group.

3.2 Commercial Viability

As is the case in the commercial world, the HPN working group feels strongly that the IPng protocols must be commercially viable. This includes but is not limited to the following issues:

- 1) The IPng protocols must function correctly. The Navy cannot afford to have network protocol problems in mission critical systems. There must be a high degree of confidence that the protocols are technically sound and multi-vendor interoperability is achievable.
- 2) The IPng protocols must have the support of the commercial/industrial community. This may first be demonstrated by a strong consensus within the IETF community.

3.3 Transition Plan

The Navy has a large number of existing networks including both Internet and ISO protocols as well as a number of proprietary systems. As a minimum, the IPng effort must address how to transition from existing IP based networks. Additionally, it would be desirable to have some guidance for transitioning from other network protocols including, but not limited to, CLNP and other commonly used network protocols. The transition plan for IPng needs to recognize the large existing infrastructure and the lack of funds for a full scale immediate transition. There will, in all likelihood, be a long period of co-existence that should be addressed.

4. General Requirements

The general requirements documented below are topics that the HPN working group considers to be of vital importance in a network protocol solution. It is hoped that the IPng solution will address all of these issues.

4.1 Addressing

The HPN working group has identified initial addressing requirements. First, a large number of addresses are required. In particular, the number of addressable entities on a single platform will range from the 100's to 100,000. The number of large platforms (ships, submarines, shore based sites) will range from a few hundred to several thousand. In addition, there will be 500 to 1000 or more small platforms, primarily aircraft. Since it is expected that in the future many of these platforms will be connected to global networks, the addresses must be globally unique.

The second requirement identified is for some form of addressing structure. It is felt that this structure should be flexible enough to allow for logical structures (not necessarily geographical) to be applied. It is also felt that this is important for the implementation of efficient routing solutions. In addition, the addressing structure must support multicast group addressing. At a

minimum 2**16 globally unique multicast groups must be distinguishable per platform.

4.2 Integrated Services Architecture

An important goal of the HPN working group is to identify existing and emerging technologies which provide mechanisms for integrating the services required by mission critical Navy systems. The HPN working group has identified two classes of problems under the general category of integrated services. The first is to provide for the multiple types of services identified in section 2.1. It is required to support these services in an integrated fashion in order to be able to correlate (in time) related streams of information.

The second class of problems relates to the predictable management of the various traffic flows associated with the above identified services. While many of these services require the delivery of a PDU within a specified time window, the applications in a mission critical environment can demand more stringent requirements. In areas where real-time systems are in use, such as machinery control, narrower and/or more predictable delivery windows may be required than in the case of the delivery of audio or video streams. The mission critical environment also requires the ability to assign end-to-end importance to instances of communications (i.e., invocations of a particular service). For example, an ongoing video stream may need to yield to machinery control commands to ensure that the commands are received before their deadline. The expense of this action is to degrade temporarily the video stream quality.

The HPN working group is looking for mechanisms in the IPng protocols to provide for both of these classes of problems in an integrated fashion. An integrated services architecture reduces design and integration complexities by providing a uniform set of tools for use by the mission critical system designer and application developer. Finally, the integrated services architecture must be flexible and scalable so that new services can be added in the future with minimum impact on systems using it. The HPN working group has intentionally avoided mentioning particular mechanisms that can be used to solve some of these problems in order to avoid requiring a particular solution.

4.3 Mobility

The HPN working group has identified two classes of mobility for the Navy mission critical environment. First, most platforms are themselves mobile. As these platforms move from port to port or from flight deck to flight deck, it is important that they are able to communicate with a number of defense installations via a general

infrastructure. Additionally, it is feasible that systems within a single platform may be mobile. Maintenance and damage assessment requires large amounts of information at numerous locations on a platform. This information could possibly be made available through mobile terminals.

4.4 Multicast

Multicast transfer is a very critical IPng requirement for the Navy's mission critical systems. Aboard a Naval platform there are many hosts (e.g., workstations) connected via numerous subnetworks. These hosts are all working different aspects of the problem of keeping the platform operational to perform its mission. In support of this environment, multicast transfer is needed to share data that is needed by multiple hosts. For example, aboard a ship platform, environmental data (roll, pitch, heading...) is needed by almost all systems. Video conferencing may be used for communication among operational personnel at multiple places aboard this ship. Video conferencing could also be used for communicating with personnel on other platforms or at shore facilities. Both of these examples, in addition to a number of DoD and NATO studies, have highlighted the need for multicast functionality in mission critical systems.

One of the limiting factors with the present IP version 4 multicast is the optional nature of this multicast, particularly with respect to routers. The use of tunnels, while enabling the initial deployment of multicast in the Internet, appears to limit its potential. The HPN working group believes that the best approach to provision of multicast functionality is to consider it as a basic functionality to be provided by IPng. In addition, sensible mechanisms are needed to control multicast traffic (i.e., scope control). Finally, support is required to enable multicast functionality in IPng in areas such as group addressing and scalable multicast routing.

4.5 Rapid Route Reconfiguration

The HPN project will be using very high bandwidth subnetwork technology. In the mission critical environment one very important problem is placing a very low bound on the time it takes to identify a subnetwork problem and to complete the necessary route reconfigurations. The Navy's mission critical environment needs to be able to trade-off bandwidth to enable a short detection/reconfiguration time on subnetwork faults. A maximum bound on this time is felt to be less than 1 second.

5. Additional considerations

This section represents additional concerns of the mission critical environment which may impact IPng. The HPN working group felt that these issues are important for the mission critical environment; however, it was not clear how or whether it is necessary to accommodate them in IPng solutions. It may suffice that designers of IPng are aware of these issues and therefore do not preclude reasonable solutions to these problems.

5.1 Fault Tolerance

The mission critical environment is particularly sensitive to the area of fault tolerance. Any mechanisms that can be accommodated within the IPng protocol set, including routing and management, to support various levels of fault tolerance are desirable. In particular, the following features should be supported: error detection, error reporting, traffic analysis, and status reporting.

5.2 Policy Based Routing

The HPN working group feels that there may be some uses for policy based routing within the Navy's mission critical systems. The primary interest is in support of a very capable security facility. Other uses discussed are as a means for keeping certain types of data on certain subnetworks (for multiply homed hosts) and providing for automatic reconfiguration in the event of particular subnetwork failures.

5.3 Security

Security is an important requirement for most Navy applications and thus the ability for the network functions to be designed to support security services are essential. The following are several security services in particular that the HPN working group believes the network function should be able to support: rule based access control, labeling, authentication, audit, connection oriented and connectionless confidentiality, selective routing, traffic flow confidentiality, connection oriented and connectionless integrity, denial of service protection, continuity of operations, and precedence/preemption. In addition to these services, the network function should also support the security management of these security services. In particular, key management is of importance.

Currently, the IPSEC of the IETF has several draft memos being considered to incorporate various security services in the network functions. It is of concern to the HPN working group that the IPng be able to support the concepts currently being developed by the IPSEC

and also provide the ability for the addition of future security services.

5.4 Time Synchronization

Time synchronization among the various components of mission critical systems is of vital importance to the Navy. It is desirable to be able to synchronize systems on multiple subnetworks via a network layer infrastructure. Some hooks for time synchronization can be envisioned in the network layer. However, the HPN working group feels that, as a minimum, efficient time synchronization algorithms must be able to function above an IPng infrastructure. For HPN systems, it is desirable that a time-of-day synchronization capability be supported of at least an accuracy of one microsecond among all hosts in a platform or campus network. The IPng protocols should not arbitrarily prevent this type of synchronization capability.

6. Conclusions

A number of concerns specific to mission critical systems targeted by the HPN working group have been identified. The HPN working group is interested in participating with the IETF in the development of standards which would apply to mission critical systems. In particular, the HPN working group is interested in the development of multicast functionality, an integrated services architecture, and support for high performance subnetworks.

7. References

- [1] HPN Planning Group, "Concepts and Guidance for High Performance Network (HPN)", Work in Progress, May 17, 1993.

8. Security Considerations

Security issues are discussed in Section 5.3.

9. Authors' Addresses

Dan Green
NSWC-DD
Code B35 NSWCDD
Dahlgren, VA 22448

Phone: (703) 663-1571
EMail: dtgreen@relay.nswc.navy.mil

Phil Ireys
NSWC-DD
Code B35 NSWCDD
Dahlgren, VA 22448

Phone: (703) 663-1571
EMail: pirey@relay.nswc.navy.mil

Dave Marlow
NSWC-DD
Code B35 NSWCDD
Dahlgren, VA 22448

Phone: (703) 663-1571
EMail: dmarlow@relay.nswc.navy.mil

Karen O'Donoghue
NSWC-DD
Code B35 NSWCDD
Dahlgren, VA 22448

Phone: (703) 663-1571
EMail: kodonog@relay.nswc.navy.mil

