

# Implications of Precedence and Preemption Requirements on Packet Based Transport Architectures

R.G. Cole and B.S. Farroha  
JHU Applied Physics Laboratory  
{robert.cole,sam.farroha}@jhuapl.edu

**Abstract**—Over the last several years there have been various attempts to extend traditional Precedence enabled transport services to support all Command and Control (C2) applications. Traditional experience with Precedence based transport services support either circuit-switched, voice-based transport, e.g., the Defense Switched Network (DSN), or message-switched transport, e.g., the Defense Message System (DMS) and the Automated Message Handling System (AMHS). We believe these attempts to extend and define new Precedence enabled transport services have failed because of a lack of well defined requirements and an understanding of their implications.

In this paper, we offer a core set of ten requirements for Precedence and Preemption enabled transport services which aim to support all C2 applications. We make no claim as to the originality of these requirements; others have proposed subsets of these in the past. Based upon these ten requirements, we then discuss and identify their implications on network architectures for packet-based transport services. In the process we hope to better clarify the implications of the requirements and the network mechanisms to be designed and developed for future military Precedence enabled communications networks. We conclude by identifying areas for future research and development in order to bring packet-based Precedence-enabled transport services to all C2 applications.

## I. INTRODUCTION

The U.S. Army's Future Combat System (FCS) is reliant upon the development of a reliable, resilient communications capability under harsh, battlefield environments. During periods of crisis, the communications infrastructure must be capable of providing preferential delivery of information based upon the Future Force Warrior's indication of the importance of the information. In current circuit-switched networks, this is indicated by the message Precedence Level (PL). The Joint Chiefs of Staff (JCS) are developing instructions for the Global Information Grid (GIG), and hence FCS communications, for support of Precedence and Preemption (P&P) capabilities [1]. These current instructions only extend the traditional telephony Multi-Level P&P (MLPP) services to equivalent Voice over IP (VoIP) services. To date, little or no work has been performed to extend P&P requirements and resulting network architectures to all Command and Control (C2) applications to be supported in the GIG all-Internet Protocol (IP) packet-based transport network. In this paper we propose a small set of reasonable P&P requirements which encompass all application

types. We then discuss the packet-based transport architectural implications of these requirements.

Experience in providing P&P capabilities in communications services fall into two camps, i.e., traditional telephony services, e.g., the Defense Switched Network (DSN), and message handling services, e.g., the Defense Message System (DMS) and the Automated Message Handling System (AMHS). Naively mapping these onto an all-IP, packet-based transport network like the GIG is problematic. The DSN handled high PL traffic through signaling to indicate the PL and resource reservation for assured call set-up. Message handling systems provide preferential queuing and scheduling to high PL messages. The differences in supporting P&P in IP networks are numerous. The GIG all-IP network is required to provide transport services for all applications, ranging from high bandwidth, long duration, real time video to short, transactional, applications. The GIG's in-band signaling architecture relies upon common transport infrastructure for its signaling, infrastructure services, e.g., Domain Name Services (DNS), and data transport. Further the GIG will be comprised of high bandwidth fixed optical communications, satellite communications, wired digital communications, and dynamic mobile ad-hoc wireless communications. This diversity of applications, communications and architectures requires a total rethinking of approaches to developing a P&P enabled transport service for the GIG.

In this paper, we first propose a small set of critical transport requirements for this new P&P enabled packet transport service. We then explore the packet network architectural consequences of these requirements. In the process we hope to point out novel areas for new research and engineering. We further hope to stimulate a broader discussion of these concepts in the open literature.

We first present our core requirements for P&P packet transport services in the next section. We then discuss the architectural consequences of these requirements in the context of a packet based transport network. We then point out areas where we believe further research and development are required towards achieving a true Precedence enabled transport service capable of supporting all C2 applications. The last section contains conclusions and final thoughts.

## II. REQUIREMENTS

In this section we propose ten fundamental requirements for Precedence-enabled transport services. We do not suggest that these are all inclusive or by any means final. They will require much improvement and modification through open discussion. We do not suggest that we originated these requirements. We propose them here based upon discussions with numerous folks and based upon reading initial attempts by others in writing P&P requirements of packet-based transport services. We only offer these as our initial proposed set of requirements for a starting point for discussion in the rest of this paper.

Here we simply state the ten requirements. Note that they are expressed in terms of a generalized “message”, representing the natural information unit handled by the application in question. In the following sections, we discuss the implications of these requirements on packet-based transport architectures.

**Requirement 1 - User Selection:** Users should be able to indicate (to the network) the importance of the information through the message Precedence Level (PL) for all Command and Control (C2) applications.

**Requirement 2 - Content Importance:** Precedence Levels are associated with the importance of the information content of the message and not related to the application type.

**Requirement 3 - Rank Ordering:** Multiple Precedence Levels are to be supported and their handling strictly rank ordered.

**Requirement 4 - Level Authentication:** The network must be able to authenticate the users access to their given Precedence Level transport service.

**Requirement 5 - Preemption:** Under times of resource overload, the network can Preempt lower Precedence Level messages access to transport services in order to preserve or to provide transport service to higher Precedence Level messages.

**Requirement 6 - QoS:** The transport service (for a given Precedence Level handling) must support the necessary Quality of Service (QoS) to ensure proper application performance upon message receipt.

**Requirement 7 - Maximum Goodput:** The network transport service should maximize the Precedence Level goodput under all scenarios, anticipated or not anticipated.

**Requirement 8 - Precedence Authorization:** User initiated access to a given Precedence Level is specified by a higher authority, while a recipient is bestowed a given Precedence Level during the receipt of the message.

**Requirement 9 - Commanders Intent:** Under situations of autonomous network operations, Commander’s Intent should allow for local modification of policies related to user Precedence Level authentication and the nature of the Precedence Level transport service if doing so improves the mission effectiveness.

**Requirement 10 - Service Accounting:** The Precedence Level transport service should be accountable, traceable

and robust under all conditions.

Based upon these requirements, we represent their architectural consequences in terms of a Reference Architectural Model (RAM). In addition to supporting the requirements as defined above, we hold our architectural decisions to several goals. Specifically, the architecture for P&P services should be robust to situations that may arise, both expected and unexpected. Therefore, mechanisms which rely on a prediction of expected traffic volumes and volumes per Precedence-Level should be avoided. The network architecture should also be unaffected by end user behavior to the maximum extent possible. For example, the architecture should avoid allowing end users to submit unsolicited traffic at a Precedence-Level higher than its authorization. Basically, when discussing and deciding upon mechanisms for inclusion in the P&P RAM, we should continually ask ourselves what would happen if a) end users misbehave, b) what if our traffic assumptions are wrong, and c) what happens if large portions of the network infrastructure disappear, etc.

## III. ARCHITECTURAL IMPLICATIONS

In this section, we provide a discussion of the architectural implications of the requirements of the previous section in the context of our Reference Architectural Model (RAM) for P&P services. The RAM is based upon the supposition that the network should provide a core, packet-based capability for P&P, and that other services, e.g., Voice over IP (VoIP) services, can be supported as an overlay on top of this core, P&P enabled transport.

### A. Context

We refer to VoIP and Video as overlay services due to the heavy reliance of their service architectures upon overlay boxes, e.g., VoIP Gatekeepers and Gateways or Session Initiation Protocol (SIP) [11] Gateways and Proxy Servers, etc. These overlay services rely on the notion of an application defined session, and use overlay boxes to handle session initiation and negotiation, bandwidth reservations (where deemed necessary), accounting and provide other services related to the session. Clearly, the mechanisms used act on timescales related to the expected session duration and are not meant to act on the time scales of individual packet transmission times. This notion of transport and overlay is illustrated in Figure 1. Of course, it is important to develop the appropriate interface between these different service architectures.

In Figure 1, the architecture is divided into a core networking component and an overlay component. The core network component is cognizant of packet level functions and requirements, while the overlay component is cognizant of additional requirements related to the session nature of the applications, which are not simply available in the network core transport component. This is clearly not a strict division, but is meant to be “more sort of guidelines”. The purpose of this separation is to aid in making progress in planning, design and deployment. However, there obviously needs to be some level of understanding and communications between




Fig. 1. The relationship between network and overlay services.

the two groups designing the network core and the overlay components. Several obvious examples where this distinction is somewhat vague are the network core functions a) selective discard mechanisms [13] which rely on the loose notion of sessions to be effective, b) Flow Labels in IPv6 headers [12], c) the interaction between the session initiation signaling and the underlying transport network reservation mechanisms and d) the appropriate Precedence-Level mapping of the overlay control traffic when transported over the P&P enabled transport network. Others may exist.

When Preemption is discussed within the above context, two forms of Preemption are considered. These two forms arise due to the fact that the transport network no longer treats, in some sense, the application data as an atomic unit as in the case of message switching as in AMHS or as voice circuits as in DSN. Instead, the packet transport switches on packets, multiples of which comprise the application-level message. The first form, often referred to as *Hard Preemption*, acts on the session-level by accepting, blocking or denying further access. Hard Preemption is often handled and communicated through a session-level signaling and reservation protocols. The other form of Preemption is referred to as *Soft Preemption*. This is often implemented within the transport network as increased packet discard probability. When the network is performing this form of Preemption, notification of the Soft Preemption may be given toward the end user. However, the actions taken by the end users based upon the Soft Preemption events are entirely up to the end user or another intervening device. Soft Preemption may or may not result in a Hard Preemption event.

### B. Network Architecture

In this section, we present our proposed RAM for supporting P&P capabilities within the packet transport service. We believe that this proposed network architecture meets the explicit P&P requirements defined above and that the implicit requirements for VoIP and Video services in [6] and [7] can be met by an overlay architecture in conjunction with the appropriate underlying packet transport service. Further, the overlay services will require that certain capabilities be provided by the core transport service. So, we expect that there will be some additional transport requirements from [6] and

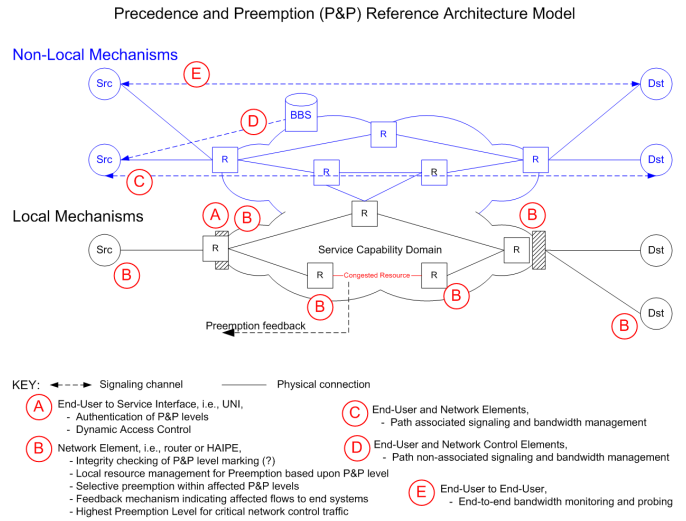


Fig. 2. Reference Architecture Model

[7] even though we categorized these documents as overlay requirements.

Our network architecture for P&P capabilities is illustrated in Figure 2. We have chosen a tiered reference model, comprising a *Local Tier* and a *Non-Local Tier*. The Local Tier comprises those functions and mechanisms that take actions based upon information and state which is local to the mechanisms. The Non-Local Tier comprises mechanisms and functions which require coordination and information dissemination between distant components in the network. For example, a Local Tier mechanism is the service discipline implemented at the access to a network resource. A Non-Local Tier mechanism is the transmission of a resource reservation request packet from the application to the destination which establishes state at each resource along the data path and which returns information regarding the success or failure of the path to provide the requested resources.

We present the RAM in terms of this tiered model, i.e., local versus non-local, for various reasons. First, the local tier of the architectural model comprises those mechanisms that are minimally implemented to form a base P&P enabled transport service. We suspect that there will be programs that are satisfied in implementing only the local tier architectural components and that these mechanisms offer the desired functionality for those programs. We elaborate on the benefits of the local versus non-local tier components throughout our analysis. Further, some programs may be incapable of implementing Non-Local Tier mechanisms due to highly variable channel and link conditions in their communications systems, e.g., tactical mobile ad-hoc networks. However, we know that in order to meet all requirements, programs will have to implement some non-local mechanisms in addition to the local tier mechanisms. For example, suppose that a high PL high bandwidth video stream is to be established across the network. Without non-local capabilities to probe the network for sufficient resources, the stream may be established when a

network resource is incapable of supporting it. This may result in local Preemption events which remove all access to the congested resource for lower PL traffic while simultaneously not meeting the QoS requirements of the high PL video traffic. It is up to the specific programs to determine the desire and the expense in order to decide whether the non-local mechanisms need to be included into their networks. In our analysis, we assume that a non-local mechanism is desired for bandwidth reservation and we use a model of end-to-end, in-path, bandwidth reservation signaling, such as that afforded by the Resource Reservation Protocol (RSVP) from the IETF.

So, the network architecture we propose, minimally consists of the following Local Tier components:

- *Authentication and Access Control* - P&P capabilities allow for higher precedence messages to access the network resources in favor of lower precedence messages<sup>1</sup>. Hence it is imperative that effective authentication mechanisms are part of the architecture. Authentication functions must be performed at the network service interface, i.e., the User to Network Interface (UNI), as indicated by the label Circle-A in the figure. These mechanisms address Requirements 4 (Authentication) and 8 (Authorization) above. Specifically, we assume that the network interface implements some form of access control which is minimally configured into the interface by a network management and provisioning function. This indicates the highest Precedence-Level of the packets that are allowed to enter the network over that interface, unless generated in response to outbound packets of a higher Precedence-Level. Hence, the access control is required to have a dynamic component, which correlates outbound Precedence-Level indications to allowed inbound Precedence-Level indications. The dynamic component is necessary due to Requirement 8 (Authorization) above.
- *Packet Marking* - fundamentally, the GIG is a packet transport and the base transport network takes actions, e.g., forwarding, on individual packets. Hence, it is necessary to indicate on the packet the P&P Network Precedence-Level associated with the end user or source's message. This is required because not all communications (or messages) are session-based; often an application may generate one or a few packets. Thus all packets are required to carry the Network Precedence-Level of the message. This also requires that packet marking be performed by the end user elements connected to the network service. Packet marking is necessary to meet Requirements 1 (User Selection) and 2 (Content Importance), as users must be able to indicate PL for all application types, i.e., signaling does not suffice. We realize that this opens a host of Information Assurance (IA) issues that need to be addressed, see, e.g., [4].
- *Network Control Level* - DoD style P&P is often referred

<sup>1</sup>The [1] defines the terms Precedence and Preemption. We use the term Precedence-Level as the precedence level of the message as indicated by the end user or source of the message. We will refer to the Network Precedence-Level of a packet as indicated by packet marking.

to as ruthless. There are no guarantees afforded to the transport of lower Network Precedence-Level packets in the presence of higher Network Precedence-Level traffic. Thus, it is necessary that a highest Network Precedence-Level be identified, above all end user indications of Precedence-Level, for the sole purpose of carrying only network critical control traffic. This is traffic which supports the basic connectivity requirements of the network. An analysis of what control traffic is considered appropriate for inclusion into this Critical Network Control (CNC) Precedence-Level. The objective here is to keep the network functioning while not overloading this CNC Precedence-Level with unnecessary traffic. Further, no end user traffic should be capable of accessing the CNC Precedence-Level. This implies restrictions on the Authentication and Access Control functions at the network interface, i.e., end users should not be receiving CNC traffic so the dynamic access control should never allow an end user to send CNC traffic into the network. However, questions remain, such as what if the end user is really a router and required to run a dynamic routing protocol over the interface, etc. More work and discussion is required on this issue. A highest Network PL is necessary to address Requirements 3 (Rank Ordering) and 10 (Service Accounting). Requirement 3 (Rank Ordering) mandates a strict ordering based upon PL. This may result in starvation to lower PL traffic. Requirement 10 (Service Accounting) requires robust network operations under all conditions. To prevent starvation of critical network control traffic (causing failure to the network), we propose the Critical Network Control (CNC) PL.

- *Local Resource Management* - some of the most effective schemes for resource management, e.g., P&P, are local management capabilities resident at the congested resource. Hence, our architecture provides a local resource management function. Specifically, these functions include service discipline and active queue management capabilities within the transport architecture. We will refer to these mechanisms collectively as the Per Hop Behavior (PHB) mechanisms. These mechanisms should be supported in all network elements performing a statistical multiplexing function, as indicated by the label Circle-B in Figure 2. Alternative proposals for implementing these functions within the context of a P&P enabled transport network are possible. We do not intend to specify a given set of Local Resource Management implementation. Instead we intend to analyze alternatives and identify their respective trade offs [2] [3]. Local P&P processing is necessary due to Requirement 1 (User Selection). Requirement 3 (Rank Ordering) mandates a treatment ordering and Requirement 5 (Preemption) provides for Preemption to achieve this ordering. Also, Requirements 6 (QoS) and 7 (Maximum Goodput) address issues of scheduling for joint QoS and P&P treatment.
- *Selective Discard Mechanisms* - when packets are discarded during a soft preemption event, the network el-

ements should do so in the most unobtrusive way as possible. Hence the RAM includes a Selective Discard capability. This should be defined so as to minimize the impact of preemption on the affected P&P Level's data. This should be implemented in a fashion which reasonably considers the capabilities of routers. The specific mechanisms for doing this are going to be dependent upon the specific network technology, i.e., IPv4 versus IPv6. More work and discussion is required with respect to the mechanisms and the policies for selective discarding. No guidance is given in the current DoD requirements, e.g., [1]. Requirement 7 (Maximum Goodput) mandates maximizing each PL's goodput. It does not define the goodput, which is left for further discussion. However, we suggest, independent of a strict goodput definition, the only way to maximize a goodput metric with dependence upon PL is to provide some form of Selective Discard.

- *Explicit Notification* - it is useful that feedback be provided upstream to voice and video applications, end-users or middle boxes in the event that flows are being preempted. This capability is not required for non-session based applications. A form of explicit notification is included at all network elements performing a statistical multiplexing function, as indicated by the label Circle-B in the figure. Here, the selective discard mechanism identifies a subset of the flows on which to implement preemption, i.e., packet discard for soft preemption. It must maintain a memory related to the selected flows and send notifications back up stream toward the source of the packet flows suffering preemption. It is not necessary that notifications are sent for every dropped packet, but only occasionally for sessions being preempted. These soft preemption notifications will carry the P&P Level markings and be sent to the address of the application source. The exact algorithm for tagging packets with an explicit notification is currently undetermined (see the section below on challenges and future work items). However, the contents and form of the notification messages need to be analyzed within the context of various analysis/deployment models for IPv4 technologies, for IPv6 technologies and for MPLS technologies. It is left to the application, middle boxes, or other elements in the overall service architecture to take actions based upon these notifications. We consider this aspect outside the scope of this paper. For hard preemption events, the notification messages should be sent using the associated signaling and reservation protocol. These capabilities are necessary due to Requirement 10 (Service Accounting) above.

The network architecture we propose, additionally consists of Non-Local Tier components. Not all non-local components are necessary, but some form of non-local mechanisms are required to fully meet the requirements and goals discussed earlier. Possible non-local mechanisms include:

- *Application Level Signaling* - it is desirable for the application to be able to indicate, through some form of signaling, their minimum desired QoS requirements and associated PL. This allows the transport service to be able to more intelligently coordinate and determine when Preemption conditions exist and how to coordinate Preemption to optimize network-wide performance. How the network uses this information and communicates this information to network components is a separate issue. (see the following three bullets). Requirements 6 (QoS) and 7 (Maximum Goodput) mandate these capabilities as discussed below.

The application signaling could have a local interaction, termed *User to Edge Signaling* through an application to network interface signaling method. The application may choose to signal the network indicating its minimum expected performance levels. The network may use this information to configure rate filters at the UNI for that flow. Note, currently this approach addresses only rate controls. It is not clear how to extend this concept to other metric indicators. This approach has the benefit of leaving the network elements relatively simple, while placing the burden on the network interface elements.

The application signaling could have a more global interaction, termed *Path Reservation Signaling*, where the network passes the application level signaling information along the desired router path across the network and to allow the components along the path to build state to enhance overall network performance and handling of Preemption events. An example of this type of signaling would be end-to-end, path associated RSVP signaling.

The application signaling could interact with a *Bandwidth Broker Reservation System* where the network may choose to pass application level signaling information to a centralized Bandwidth Broker Service (BBS) which would be responsible for reserving the appropriate state information along the desired routed path to enhance overall network performance and handling of Preemption events. The details of this type of network signaling is out of the scope of this document.

- *Network Usage Policy* - it is clearly desirable to understand the limitations of the deployed architecture for P&P services and to employ forms of Correct Usage Policy to ensure correct functioning of the network. This is also an example of a Non-Local Tier function that may be employed within the RAM.
- *Policy Based Networking* - policy based network management capabilities are necessary in order to reconfigure the P&P mechanisms based upon Requirement 9 (Commanders Intent). Further, the implication of specific policy decisions must be measurable, hence Requirement 10 (Service Accounting) is suggested.

We suspect that it is desirable to implement all of the Local Tier mechanisms in a multi-service P&P transport network. However, due to the somewhat overlapping nature of the

Non-Local Tier mechanisms we have identified, not all Non-Local Tier mechanisms would be necessary within an actual network deployment. It is clear, however, that these Non-Local Tier mechanisms are a valuable and useful component or any network deployment of P&P transport services. In order to achieve end-to-end support for P&P enabled transport across a multi-domain network like the GIG, it is necessary to pass a consistent set of information across domain boundaries related to the above discussed non-local components. This aspect of the RAM has not been fully analyzed. For example, some domains may implement an end-to-end, path associated signaling and reservation mechanism, while other domains may implement a Bandwidth Broker Service mechanisms. In these cases, the information passed between the domains should be independent of the actual internal domain mechanisms and architectures. Internal to a domain, this architecture is not mandated. Even how to implement preemption notification may be dependent upon the domain internal architectures. This level of flexibility must be kept in mind when design a specific implementation and deployment of the P&P RAM.

Additional architecture enhancements are possible and should be considered. We list these here so they do not get dropped from future consideration. These may include:

- *Preemption Initiation of Alternate Routing* - it may be desirable that the network attempt to mitigate congestion at a limiting resource by initiating actions, e.g., alternate routing. Other actions may be considered as well.
- *Middleboxes for Performance Control* - middleboxes for performance enhancements are desirable in situations where an interior subnet is not fully compliant to the Reference Architectural Model for P&P. Although we do not view these as integral to the RAM, they clearly are advantageous in certain migration and deployment situations. Therefore, analysis of these types of middleboxes should be considered. See [10] for a discussion of middleboxes.

P&P and QoS architectures are intertwined concepts and need to be considered as such. Critical areas where the two concepts merge include:

- *Identification of Preemption Conditions* - the decision to preempt should be cognizant of the QoS requirements of the applications currently inputting traffic into the congested resource. For example, if an elastic data flow at a relatively high Precedence-Level can be throttled back somewhat (and still meet its minimum requested QoS) instead of preempting a lower Precedence-Level flow, then this is what the architecture should strive to do.
- *Per Hop Behavior* - packet scheduling and active queue management define the per packet behavior managing a specific resource, e.g., a trunk interface. The PHB must support both the QoS requirements of the applications as well as the P&P requirements as indicated by the end users. Thus, the PHB must be cognizant of both the QoS and the P&P requirements. For example, Figure 3 shows an expansion of a networking element's interface

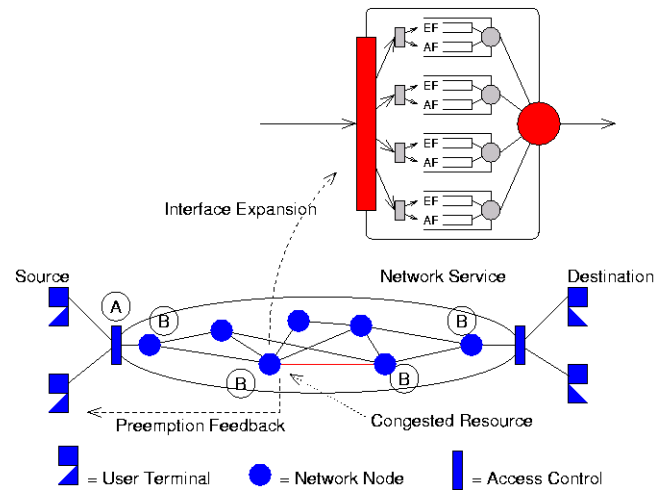


Fig. 3. An expansion of a local resource management capability at an interface card.

card which manages the limited resources associated with buffers and bandwidth. Here, we show combined mechanisms which are cognizant of both packet P-L and QoS markings. We are not promoting a specific PHB in the figure; it is merely intended for illustrative purposes. However, in [2] and [3], we propose a specific PHB where new active queue management functions are designed to address the P&P requirements while traditional QoS scheduling mechanisms, e.g., Expedited Forwarding (EF) [9] or Assured Forwarding (AF) [8], support the QoS requirements. In [2] and [3], we are investigating these PHBs through extensive simulations for the development of joint QoS and P&P enabled Per Hop Behaviors which are robust under a range of traffic conditions.

- *Selective Discard Decisions* - we have not yet fully addressed this issue of how to perform selective discard for soft preemption events in the RAM. This set of considerations also applies to decisions to hard preempt sessions, e.g., which sessions to preempt given a choice of multiple types of session applications. Various possible implementations have been identified, but more analysis is required. Alternatives proposed include a) randomly selecting flows to discard from, b) monitor and tagged SLA non-compliant flows and selective discard from them first, c) always favor voice applications first, and d) policy mechanisms pushed down from a policy manager based upon the commander's intent.

Up to this point, we have been careful not to associate the architecture components with specific protocol implementations, e.g., Differentiated Services (DiffSrv) Code Points (DSCPs), Explicit Congestion Notifications (ECNs), IPv6 Flow Labels, Weighted Random Early Discards, etc., because these are associated with a specific layer in the protocol stack. However, depending upon the specific transport technologies upon which our P&P architecture is to reside, other comparable protocol implementations must be considered. For example, in the case

of packet marking, it may be appropriate to rely on a subset of the DSCP space to indicate the P&P level of the message (although this is open to further analysis and discussion). If instead, the transport architecture is supporting Multi-Protocol Label Swapping (MPLS) transport, then we must find the appropriate mapping of the markings to the MPLS header or in the MPLS label space. The same consideration holds in broadcast, multi-access subnet such as Ethernet. Here, the P&P requirements dictate that the P&P levels should be carried on the Ethernet headers and that some form of priority back off in the event of collision detection occurring. Another example may be in the deployment of IP-based satellite modems which support dynamic bandwidth allocation. Therefore, it will be necessary to consider the P&P architecture within the specific set of networking technologies on which it is to be implemented and deployed.

#### IV. CHALLENGES AND FUTURE WORK

The known challenges we identify here. We point out the following research and analysis challenges associated with the P&P RAM:

- *Selective Discard Algorithms* - within an affected P&P Level at a congested resource, there may be messages from a complex mix of applications. Given that some, but not all, messages will require some form of preemption, the interesting question is how to choose the subset of application messages to preempt. This requires further study and guidance.
- *Preemption Notification Algorithms* - the exact nature of the algorithm to implement preemption notification is not yet defined. What are the algorithms to identify flows, to decide when to send notifications and what is required to communicate this information across a security tunneling device? These, and other, question needs to be addressed.
- *Authentication Mechanisms* - per packet marking is expected in this architecture and, in fact, per packet preemption is expected by the requirements. It is necessary to authenticate that the packet marking of P&P Levels are allowed when the packet enters the network. Typically, in data networks this is handled based upon the port associated with a single or group of end-users with common authentication privileges, or by association of source address with privilege level, etc. We are currently proposing a dynamic based filtering mechanisms at each network interface. This need further study.
- *High Speed Integrity Checking* - is it necessary for each statistical multiplexing network element to check the integrity of the packet markings prior to assigning to a P&P treatment? This would require some form of high speed integrity checking capability, or it may be possible to sample a subset of packets related into a given flow, or some other method. Does this offer any improvement in the overall security of the architecture? This requires further study.
- *High Speed Integrity Checking and Correction* - is it necessary for each statistical multiplexing network element

to check the integrity of the packet markings prior to assigning to a P&P treatment? If yes and the check fails, it may be desirable to have the capability to correct the modified P&P level to the correct level as specified by the application. This requires further study.

- *Critical Network Control Traffic* - it is necessary to identify a highest Precedence Level, called Critical Network Control (CNC). This level is to contain only that control traffic which is critical to the proper functioning of the network. It should not carry end user generated traffic. It should not carry all network management data. So, we must also carefully identify what control and management traffic is deemed critical and to be placed within this highest Precedence Level. This requires further study.
- *Joint QoS and P&P PHBs* - much work exists in the open literature in defining PHBs for QoS. There is little work on joint P&P and QoS PHBs. We have embarked on a set of studies in this area [2] [3]. But much more work is required.

#### V. CONCLUSIONS

We have proposed a small core set of requirements for P&P enabled network services to support all C2 applications. This set of requirements is neither complete nor well vetted. We discussed the architectural implications of these requirements on packet-based transport services like the DOD's GIG. We discuss architectures in the form of a Reference Architectural Model (RAM). In the process, we hope to have stimulated interest of researchers in investigating new mechanisms required for a complete end-to-end P&P architecture for packet-based network services.

#### ACKNOWLEDGMENTS

We wish to thank B. Doshi of JHU/APL for numerous and informative discussions of P&P enabled architectures. We are grateful for the numerous and enlightening discussions with and presentations from the members of the 2005 GIG QoS Working Group Precedence and Preemption sub-Team. Finally, we wish to thank the reviewers for their insightful comments and suggestions to improve this paper.

#### REFERENCES

- [1] Chairmen of the Joint Chiefs of Staff - Instructions (CJCSI), *Policy, Responsibilities, Processes, and Administration for the Department of Defense Global Information GRID Networks*, CJCSI 6215.03C, Draft, 31 July 2006.
- [2] Cole, R.G. and P.F. Chimento, *Modeling and Preliminary Simulation Studies for Packet-based Precedence and Preemption for FCS Communications*, Army Science Conference 2006, Orlando, FL, USA, December 2006.
- [3] Cole, R.G., *Impact of Precedence Enabled Per Hop Behaviors on TCP Flows*, submitted to IEEE MILCOM 2007, Orlando, FL, USA, October 2007.
- [4] Farroha, B.S., Cole, R.G., Farroha, D.L. and A. DeSimone, *An Investigative Analysis of Information Assurance Issues Associated With the GIG's P&P Architecture*, SPIE 2007, Orlando, FL, USA, April 2007.
- [5] reference to GIG MA-ICD. title, reference number, date.
- [6] IAW NCS, *Telecommunications Operations Government Emergency Telecommunication Services (GETS)*, IAW NCS Directive 3-10, 10 February 2001.

- [7] IAW NCS, *Telecommunications Service Priority (TSP) System for National Security and Emergency Preparedness (NS/EP)*, IAW NCS Directive 3-1, 10 August 2000.
- [8] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, *Assured Forwarding PHB Group*, Internet Engineering Task Force (IETF), RFC 2597, June 1999.
- [9] Davie, E., et.al., *An Expedited Forwarding PHB (Per-Hop Behavior)*, Internet Engineering Task Force (IETF), RFC 3246, March 2002.
- [10] Carpenter, B. and S. Brim, *Middleboxes: Taxonomy and Issues*, Internet Engineering Task Force (IETF) Request for Comment (RFC) 3234, February 2002.
- [11] Rosenberg, J., et.al., *SIP: Session Initiation Protocol*, Internet Engineering Task Force (IETF), RFC 3261, June 2002.
- [12] Deering, S. and R. Hinden, *Internet Protocol, Version 6 (IPv6)*, Internet Engineering Task Force (IETF), RFC 2460, December 1998.
- [13] Floyd, S. and V. Jacobson, *Random Early Detection Gateways for Congestion Avoidance*, IEEE/ACM Transactions on Networking, August 1993.