Next-Generation Packet Network Architectures with Decoupled Service Plane and Transport Plane

(Invited Paper)

Indra Widjaja
Bell Laboratories, Lucent Technologies
Murray Hill, NJ 07974, USA

Abstract— **The emerging packet network architectures with decoupled service plane and transport plane may play a prominent role in the future. Traditional packet network architectures are primarily concerned with providing quality "pipes" (better QoS, high bandwidth, more efficient and reliable transport, etc.). The emerging architectures shift the focus towards networks that readily enable new advanced services and allow operators to have deep control over their networks. We describe related architectures with decoupled service and transport planes that are currently being addressed by various standards development organizations such as 3GPP, 3GPP2, ETSI TISPAN, ITU-T NGN, Packet Cable and MSF. We explore various issues that need to be considered with these network architectures with a focus on resource management.**

I. Introduction

Packet network architectures have been and will be constantly evolving to adapt to new applications, services and requirements. In the past, newly developed packet network architectures have been mainly focussed on the *transport plane*, whose purpose is to deliver user packets[1] end-to-end and manage transport resources such as bandwidth, buffers, packet processing elements, etc. QoS information needed for dynamic resource management may be propagated at the transport plane, typically via a transport-plane signaling protocol (e.g., RSVP [1]). Examples of such architectures include ATM and MPLS.

We believe we are now entering an era where a new form of packet network architectures will receive increased attention. End user devices are becoming more intelligent and many are now capable of initiating innovative services. One example comes from devices that are capable of exchanging Session-Initiation-Protocol (SIP) signaling messages to initiate services [2]. These signaling messages are forwarded by SIP proxies over a *service plane*, whose purpose is to transfer application-level messages and enable a rich set of services such as IPTV, push-to-talk, instant messaging, multimedia, on-line gaming, etc. When negotiation at the service plane is successful, the end user devices often exchange additional users packets over the transport plane. The crucial point in these architectures is that *the service plane and the transport plane are decoupled*: (1) the service plane is aware of the service-related information pertaining to user requests but is ignorant of the transport resource information, and (2) the transport plane is aware of its own transport resource information but is ignorant of the service-related information.

There are other desirable characteristics to these architectures besides facilitating service creation. The architectures allow a single service protocol at the service plane to work with heterogeneous technologies at the transport plane (e.g., cable network, DSL access network, EPON, etc.), or multiple service protocols at the service plane to work on a single IP-based network (e.g., MPLS) at the transport plane. The architectures also allow some functions in one plane to be modified without impacting the other plane. However, decoupled network architectures introduce new challenges - the most outstanding being the coordination problem. For instance, while end users negotiate services and specify the required QoS level at the service plane, the actual user packets are transferred between the users at the transport plane. If there is no coordination between the two planes, it is possible that the transport plane may not be able to satisfy the requested QoS level. To solve this coordination problem, various standards development organizations (SDOs) have proposed a functional entity that has a northbound interface to the service plane and a southbound interface to the transport plane [3]-[11].

Fig. 1 illustrates an example of how the coordination works for the network architecture with decoupled service plane and transport plane. Before service can be delivered, both end users exchange application signaling messages over the service plane. These signaling messages are forwarded by Application Functions (AFs), which could be SIP proxies if the signaling is based on SIP. For the coordination with the transport plane, an AF in the service plane would contact a Resource Control (RC) to request an authorization for the service. The request contains relevant *service information* needed, such as resource require-

[1] In this paper, user packet is equivalent to bearer, media or data packet.

1-4244-0425-8/06/$20.00 ©2006 IEEE

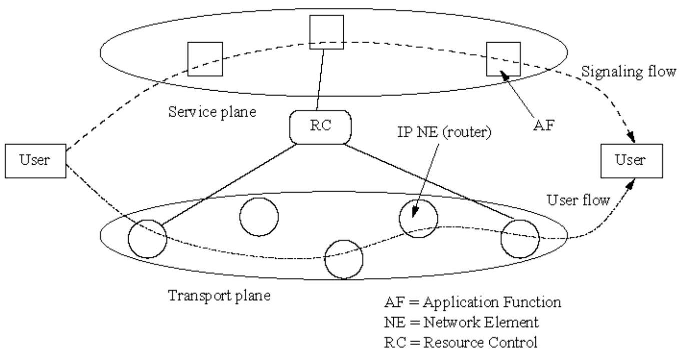

Fig. 1. Architecture with decoupled service and transport planes and the resource control that provides the linkage between the two planes.

ment, for authorization. If the authorization is granted, the RC will provide some instructions to the network elements (NEs) at the transport plane so that certain resources can be configured, reserved and enforced for the user packet flow(s) associated with the service. The contacted NEs are typically those residing at the network edge. If authorization is successful, the end users will eventually be notified and transfer of user packets with guaranteed QoS over the transport plane can begin. Although any technology at the transport plane can be used, in this paper we assume that the transport plane is IP-based.

In Sec. II, we review various standards that deal with the coordination problem between the service and transport planes. We discuss issues that arise with the architectures in Sec. III. We point out that different phases and models for resource control present some interesting challenges. Because QoS templates at the service and transport planes are incongruent, a mapping of these templates is needed. To further clarify the operation of a network with decoupled service and transport plane, Sec. IV provides examples of how control messages are being exchanged horizontally within the service plane and vertically between the two planes via the RC. We conclude in Sec. V and contemplate future research topics in QoS control for networks with decoupled service and transport planes.

II. Relevant Standards Activities

In this section, we provide an overview of the relevant standards that deal with the function that provides the coordination between the service plane and the transport plane. As all the standard activities are still ongoing, it is likely that some functions in the architectures will evolve in the future.

A. 3GPP PCRF

The third Generation Partnership Project (3GPP) consists of partners (national standard bodies), whose main goal is to produce technical specifications for 3rd generation mobile systems based on evolved GSM. The 3GPP output is massive and diverse. In this paper, we will only focus on the relevant output pertaining to the RC.

The work in 3GPP has been evolving with the most recent release (Release 7) defining the functional entity between the two planes called Policy and Charging Rule Function (PCRF) [3]. The PCRF has an interface to the AF residing at the service plane, and an interface to the wireless access gateway residing at the transport plane. The gateway is called Gateway GPRS Support Node (GGSN). The main objective of the PCRF is to interpret the request from the AF containing the service information, and use this information to apply policy and charging control (PCC) to the gateway. The policy control contains gate control (allowing or blocking certain user flows) and QoS control (e.g., packet marking and resource enforcement). The charging control instructs the gateway on how certain user flows are to be metered (volume based or time based) and charged (offline or online). The work in 3GPP is limited to a specific wireless access network application.

B. 3GPP2 SBBC

3GPP2 is similar to 3GPP, except that it is focussed on CDMA wireless systems. The work in 3GPP2 pertaining to the RC tracks the corresponding 3GPP work very closely. The differences mainly come from the terminology. For example, the capabilities of the entity providing the coordination is called Service Based Bearer Control (SBBC) in 3GPP2 [4], instead of PCC in 3GPP. Also the PCRF in 3GPP2 is composed of two entities called Policy Decision Function (PDF) and Charging Rule Function (CRF). However, the overall functions and protocols are almost identical to those defined in 3GPP, including the PCC and SBBC. The gateway that is controlled by the PCRF is called Packet Data Serving Node (PSDN).

C. ITU-T RACF

In the ITU-T Next-Generation Network (NGN) architecture, the RC is called the Resource and Admission Control Function (RACF), which is envisioned to be an arbitrator for QoS between Service Control Functions (SCFs) residing at the service plane and Transport Functions (TFs) residing at the transport plane [5]. Unlike the limited applicability of the RC equivalence defined in 3GPP or 3GPP2, the TFs may reside in access or core networks and deal with wireless or wireline applications. Fig. 2 shows the functional architecture of RACF.

RACF contains two functional entities: Policy Decision Functional Entity (PD-FE) and Transport Resource Control Functional Entity (TRC-FE). The TRC-FE is an entity

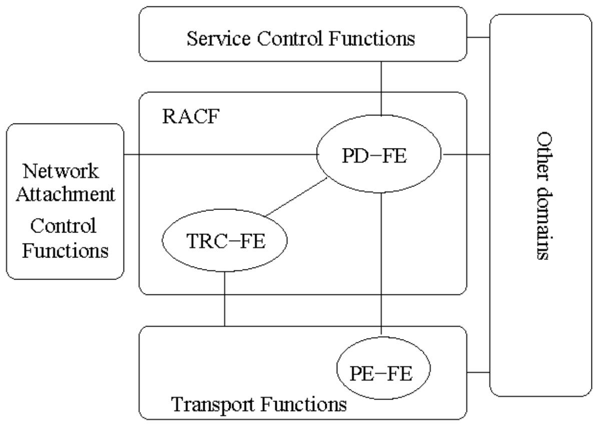

Fig. 2. RACF functional architecture.

that does not exist in 3GPP or 3GPP2. Its main purpose is to collect network resource information and possibly apply admission control on the resources. The TRC-FE allows for the "domain-wide" resources (as opposed to "gateway-specific" resources in 3GPP or 3GPP2) to be monitored and applied for the purpose of admission control, thus enabling RACF to also work within a core network. Based on the service information from the SCF, resource information from the TRC-FE, local policy rules and subscription information, the PD-FE performs the final policy and admission control decision. RACF also defines a variety of features that it can install at the Policy Enforcement Functional Entities (PE-FEs), including gating, packet marking, policing, firewall selection, and Network Address and Port Translation (NAPT). However, the PD-FE at this point has not incorporated a charging function yet.

D. ETSI RACS

Resource and Admission Control Subsystem (RACS) is an ETSI standard that provides control services within the access networks and at interconnection points between core networks [6]. Fig. 2 shows the functional architecture of RACS with service plane consisting of Application Functions (AFs) and the transport plane consisting of Resource Control Enforcement Functions (RCEFs) and Border Gateway Functions (BGFs).

As in RACF, RACS also contains two functional entities: Service-based Policy Decision Function (SPDF) and Access-Resource and Admission Control Function (A-RACF). The SPDF performs similar functions as the PD-FE in RACF and may interact with the A-RACF and/or BGF depending on the service information received from the AF and its local policy rules. The A-RACF is different from TRC-FE. The A-RACF receives requests from the SPDF and performs policy and admission control for an access network. If control is passed, the A-RACF will

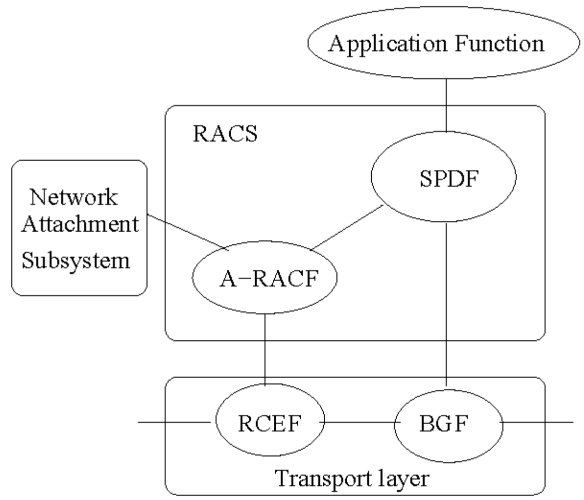

Fig. 3. RACS functional architecture.

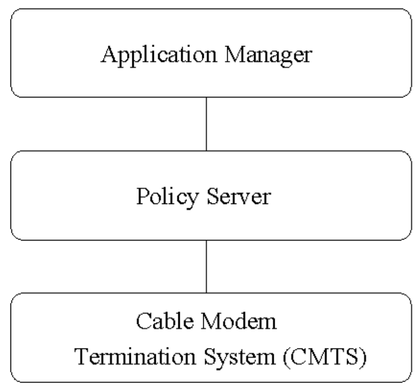

Fig. 4. PacketCable functional architecture.

install the decision to the RCEF. There could be different instances of A-RACF, each dealing with a different access network type.

The BGF provides an interface between two network domains (access-core or core-core), and supports the following functions: gating, NAPT, NAT traversal, packet marking, resource allocation, policing and usage metering. The RCEF is a specific type of BGF that sits in an access network or at one of its edges. Functions supported by the RCEF include gating, packet marking and policing.

E. PacketCable PS

PacketCable is specified by CableLabs for providing multimedia services over DOCSIS 1.1 or greater access network [7] [8]. PacketCable Release 2 relies on 3GPP (Release 6) for the service plane implementing IP Multimedia Subsystem (IMS) [12]. Fig. 4 depicts the simplified functional architecture of PacketCable Release 2 without the service plane, which is provided by IMS [8].

The Application Manager (AM) is mainly responsible for determining the QoS resources needed for a session

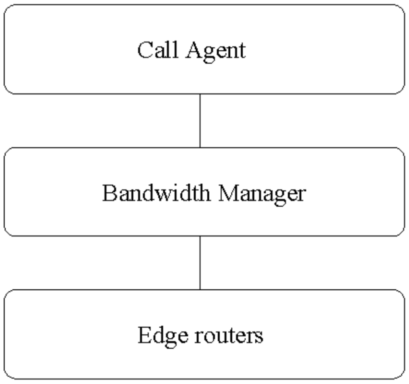

Fig. 5. MSF functional architecture.

based on the information it receives from an IMS element (not shown in the figure) called Proxy Call Session Control Function (P-CSCF). The Policy Server (PS) is equivalent to the RC. It receives requests from the AM and apply policy decisions to the CMTS, which resides at the transport plane. The interface between the PS and the AM (or CMTS) is specified in [9].

F. MSF BM

The Multi Service Forum (MSF) is an organization consisting of service providers and vendors. Its primary focus is on developing implementation agreements that complement standards specified by other SDOs. Early work was targeted to architectures with separate control and switching functions such as a separate media gateway controller that can control one or more media gateways [10]. Recent work is aligned to the architecture with decoupled service and transport planes [11]. Fig. 2 shows the simplified functional architecture of MSF Release 2. Other functions such as application servers and service broker are not shown for clarity.

The Call Agent provides call-control functions such as session setup, session tear-down and session control. It interfaces with the service broker for value-added services and interfaces with the Bandwidth Manager (BM) for bandwidth allocation. The BM is the entity that provides the coordination and assures QoS for each given service.

III. Issues and Challenges

As evidenced by the considerable standard activities, the preceding section underscores the significance of the emerging architectures with decoupled service and transport planes. In this section, we present important research issues pertaining to the role played by the RC.

A. Transport Plane Resources

Because different standards focus on different NE types or network technologies, they typically define different resources that need to be controlled by the RC and enforced at the transport plane. Nevertheless, there are some resources that are common.

One concept that is common at the transport plane is that of gating. A **gate** is a data-path construct residing in an NE at the transport plane that contains a variety of resources that are to be controlled by the RC. For example, the QoS resource may contain packet marking, policing, shaping and other traffic management elements. The charging resource may contain various types of usage-based meters. The Network-Address-Translation resource may contain various types of internal-external mapping of addresses and ports. A gate typically also contains packet classifiers and operates a binary switch that can be opened to allow packets matching the classifying rules to pass through or closed to block the packets.[2] Note that the NE may contain other data-path resources that are *not* controlled by the RC; for example, forwarding information base. Such resources will not be covered in our discussion.

B. Resource Control Phase

The preceding section describes different types of resources (e.g., QoS, charging and security) that the RC needs to control. In this paper, we mainly focus on QoS resource. QoS resource control performed by the RC for a given session can generally be categorized into three phases [5][13]:

- **Authorization**: The phase when the amount of QoS resource is authorized based on policy rules and subscription information. Service-based admission control may play a role in this phase.
- **Reservation**: The phase when the amount of QoS resource is reserved based on the authorized resource, the requested QoS resource and current resource availability. Resource-based admission control plays an important role in this phase. If admission control is successful, the requested resource will be reserved, but the associated gate remains closed. The resource that has been reserved but not committed may be used for best-effort traffic, but may not be used for higher quality traffic as it may be preempted when the resource is committed.
- **Commitment**: The phase when the amount of QoS resource is committed (and enforced) for the user flows when the associated gate is open. This is when the packets asso-

[2]Alternatively, packets may be passed through without giving them with high-quality service; e.g., they may only receive best-effort service. Opening a gate is sometimes referred to as setting up a pinhole.

ciated with the gate are allowed to pass through. Normally, usage-based meters for billing start recording only when the gate is open.

From the above definition, the amount of the reserved resource clearly should be less than or equal to the amount of the authorized resource. For example, an initial request from the AF for a given subscriber may be authorized by the RC with a maximum of 10 Mbits/sec. When a subsequent request indicates that the requested resource is 2 Mbits/sec, this amount will be checked by admission control against current resource availability. If admitted, the amount of the reserved resource will be 2 Mbits/sec. In general, the amount of the committed resource is equal to the amount of the reserved resource. In some cases, the amount of the committed resource can be less than that of the reserved resource. For example, this may happen when reservation may have to deal with a list of codecs while only one codec is known to be active at the commitment phase.

Although there are three phases for QoS resource control, they need not be executed separately in practice. For example, for services where the delay between request and service delivery should be minimal, authorization, reservation and commitment of resources may need to be combined and executed in a *single phase*. For some services such as interactive voice calls, a *two-phase* execution process may be necessary. For example, authorization and reservation may be combined and made before the callee is alerted in order to prevent "ghost ring" [14]. Resource commitment is made only after the callee "goes off-hook". Finally, authorization, reservation and commitment of resources may also be executed in three different phases when transport resource is scarce and there is a significant delay between authorization and reservation.

It is sometimes mentioned that the two-phase process (reservation phase followed by commitment phase) is less efficient in terms of bandwidth than the single-phase authorize/reserve/commit process (e.g., see [15]). Although this may impact efficiency measurably when the delay between reservation and commitment is substantial, another more significant problem is that the single-phase process applies charging before the callee goes off-hook if commitment is also made in the first phase. This is clearly unacceptable if the service is never provided because the callee does not answer.

Note that admission at the authorization phase does not guarantee admission at the reservation phase. On the other hand, reservation generally leads to successful commitment. Because resources are held but not used at reservation, it appears possible to mitigate this resulting inefficiency by permitting, with small probability, reservation to lead to failed commitment later.

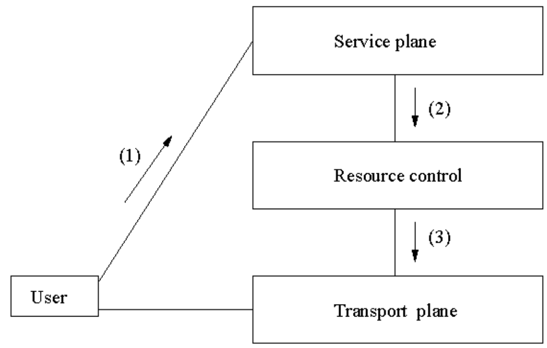

Fig. 6. Control flows in the push model.

C. Resource Control Model

Resource control model describes the flow of control messages between various components from the time a user initiates a request to the service plane until the service is enabled at the transport plane so that the corresponding user packets receive appropriate treatments. There are two models that have been identified [5][8].

The first is the **push model**, as shown in Fig. 6. Here, the user first sends a service request message (e.g., SIP invite) to the service plane (step 1). Normally, one of the AFs at the service plane derives the service information in the request message and sends a request for authorization and reservation to the RC (step 2). The RC performs the necessary admission control. If the admission control passes, the RC *pushes* the information needed to install the gate in the NE(s) and commit the resource at the transport plane (step 3). Acknowledgements (not shown in the figure) are returned from the NE to the RC, from the RC to the AF and eventually from the AF to the user, which can then start transferring its user packets over the transport plane. Note that the push model can generally work in any number of phases.

Fig. 7 shows the alternative **pull model**. As in the push model, the user first initiates the request to the AF at the service plane (step 1), which then sends a request for authorization to the RC (step 2). If the request is authorized, the RC may return an authorization token in the acknowledgement (step 3), which will be forwarded to the user. An example of an authorization token is described in [16]. Upon authorization, the user sends a transport-plane signaling message (e.g., via RSVP with the authorization token carried in the policy-data object of an RSVP message [17]) to perform reservation (step 4). The signaling message triggers one or more NEs at the transport plane to *pull* the information needed to install the gate by sending a request message to the RC (step 5). The NEs may

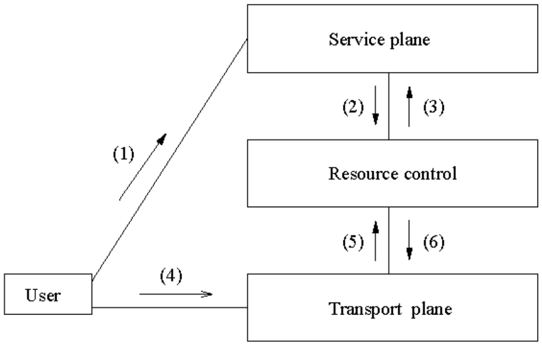

Fig. 7. Control flows in the pull model.

also perform admission control to determine the amount of resources for the reservation. After determining that the authorization token received from the NE correlates with the one authorized and issued previously, the RC returns the gate information to the NE (6). Finally, the NE can acknowledge the signaling request (e.g., via RSVP Resv message) back to the user, indicating that the reservation is successful. Note that the above pull model requires at least two phases.

It is worth noting that not all users can freely select between the two models. For example, user devices that can only negotiate QoS information at the service plane (e.g., via SDP [18]) can only support the push model. User devices that can negotiate QoS information at the transport plane but not at the service plane can only support the pull model.

In a practical deployment, it is likely that a mixed environment contains sessions operating in both push and pull models. Thus, the challenge is to design an effective system that can coordinate admission control initiated by two different planes.

D. QoS Templates at Service and Transport Planes

It is clear from the preceding section that QoS template can be described at the service and transport planes. The main purpose of this section is to highlight the differences that QoS templates are described at the two planes. The objective of the RC is to reconcile the differences when performing admission control and providing gate information.

Currently, Session Description Protocol (SDP) is the most dominant standard that can be used to describe QoS template at the service plane. SDP session description is typically conveyed in certain SIP messages. The session description consists of a session-level description and one or more media-level descriptions. The media-level descriptions, which describe the QoS information of each media flow belonging to a session, are mostly pertinent to the QoS template at the session level. We will only give a brief overview of the media-level descriptions through an example (see [18] for detailed information).

Suppose a user (user 1) wants to establish a session with another user (user 2) that consists of three types of media (voice, video and data), user 1 may describe its session in an SDP offer [19] as follows (part of the session-level description is not shown for brevity):

```
c=IN IP4 192.128.1.10
t=0 0
m=audio 49100 RTP/AVP 0
a=sendrecv
m=video 49536 RTP/AVP 31
a=sendonly
m=application 36700 udp wb
a=sendrecv
```

The message describes the media flows with the corresponding user 1's receiving port information, media types and attributes. For example, the third line indicates audio media with receiving port 49100 and RTP/AVP transport using audio/video profile carried over UDP. The digit '0' indicates that the payload type is μ-law PCM sampled at 8000 Hz [20]. The attribute in the 'a=' lines indicates whether the user would like to send, receive, or send and receive. The attribute can also describe other information such as proposed bandwidth, packetization time, etc.

Note that at this point only the information from user 1 is known. When user 2 receives this information, it may return an SDP answer as follows:

```
c=IN IP4 204.150.20.68
t=0 0
m=audio 48710 RTP/AVP 0
a=sendrecv
m=video 48010 RTP/AVP 31
a=recvonly
m=application 31535 udp wb
a=sendrecv
```

At this point, the AF has the complete QoS template for the session, which is needed by the RC to perform admission control.

In other cases, user 1 may indicate a list of codecs for a given media flow in its offer; for example, with codec list $\{1, 2, 3, 4, 5\}$ such as:

```
m=audio 3587 RTP/AVP 18 96 2 15 0
a=rtpmap:96 i:LBC/8000
a=sendrecv
```

The list indicates the order of preference with the first one being preferred. The list also allows a user to switch to a different codec in the list in the middle of a session. When user 2 receives the offer, it may respond with its own answer; for example, with codec list $\{1, 2, 5\}$. In this case, admission control must be able to take into account the possible codec $\{1, 2, 5\}$ being used during the session.

At the transport plane, the QoS template generally depends on the specific network technology. In general, IP QoS supports *absolute* QoS where certain performance objectives such as delay, jitter and packet loss can be achieved, and *relative* QoS where only relative services are provided. Absolute QoS relies on admission control and policing to guarantee the performance objectives. In such a case, the QoS template may include marking information, policing information, bandwidth information and performance metrics. Relative QoS may only rely on packet marking with appropriate scheduling, so marking information may be sufficient in the QoS template.

E. QoS Mapping

The preceding section illustrates different QoS templates at the service and transport planes, and consequently mapping of QoS templates is necessary.

As SDP is one of the possible templates that can be used at the service plane, it may be argued that mapping is needed to arrive at common *service information* that describes the QoS template at the interface between the AF and the RC. To take to another level, we may also argue that common *transport information* is needed to describe the QoS template at the interface between the RC and the NE. The transport information may be further mapped to another template specific to each technology at the NE (e.g, MPLS QoS, Ethernet QoS, etc.).

There are issues in derivation of service information. On the one hand, minimal information that is common to most QoS templates at the service plane may be preferred to simplify the interface protocol. On the other hand, since admission control may work poorly if crucial information is omitted, it may be useful to provide all the information contained in the original session description.

As an example, 3GPP uses relatively simple mapping rules from SDP to service information [21]. For each media component (the 'm=' line) in the session, the AF roughly perform the following mapping to the service information:

- The direction attribute is mapped to "uplink", "downlink" or "both".
- The bandwidth attribute is mapped to Max-Requested-Bandwidth in units of bits/sec, for each direction.
- The media types are mapped to different QoS classes. For example, audio or video media type is generally mapped to class-A QoS, except for uni-directional flow which is mapped to class-B QoS. Control is mapped to class-C QoS and data is mapped to class-E QoS.

We see that the simplicity of this mapping has some drawbacks. First, since traffic parameters that are needed for policing are not mapped, absolute QoS cannot be precisely provided. Second, the mapping does not take into account the choice of multiple codecs.

Another contrasting example comes from PacketCable where it only specifies the mapping from SDP to RSVP QoS parameters rather than to service information [13]. Specifically, the mapping derives: bucket depth (b), bucket rate (r), peak rate (p), minimum policed unit (m) and maximum packet size (M) for the TSpec, and the reserved bandwidth (R) for the RSpec. For example, for G711 codec with packetization time of 20 ms, the corresponding values for the TSpec are $b = m = M = 200$ bytes and $r = p = 10,000$ bytes/sec. The reserved bandwidth in the RSpec is $R \geq r$. When multiple codecs are involved, the mapping uses the "Least-Upper-Bound" method to guarantee the QoS of any active codec. Fig. 8 shows the comparative performance of two admission control approaches. In the first case, the bandwidth reserved for each session is based on the Least Upper Bound (LUB) of all codecs that can be used during the session. The second case with statistical multiplexing allows for the total bandwidth requirement to exceed the total reserved bandwidth by a small probability (10^{-6}). The total number of sessions is assumed to be 100. The results show that substantial savings in the amount of bandwidth that needs to be reserved can be gained by using a better representation for the service information and by designing an appropriate admission control technique.

In summary, a systematic methodology is needed to find mapping procedures that are flexible for different types of service and transport planes, and that incur little or no loss of information so that resources can be efficiently reserved.

F. Admission Control

Admission control is one of the most important functions in the ITU-T RACF and ETSI RACS standards. Although admission control can principally be applied at any element in the architectures with decoupled service plane and transport plane, the RC is normally the main entity that is involved in admission control. In many cases, the NE at the transport plane may also participate in admission control.

In one approach, the RC may perform *service-based admission control* based on service information from the AF, subscription information and local policy rules. Observe

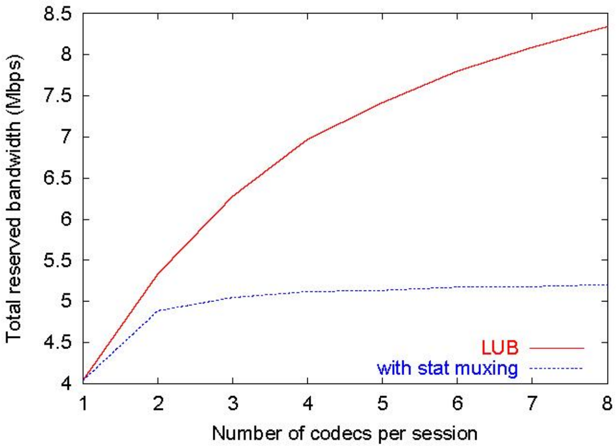

Fig. 8. Bandwidth reservation with Least Upper Bound (LUB) and with statistical multiplexing for different numbers of codecs per session.

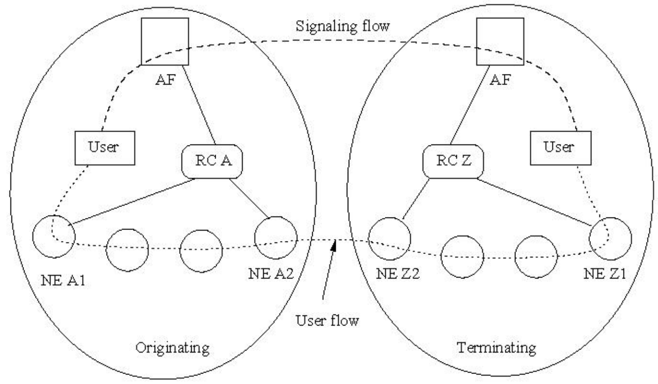

Fig. 9. End-to-end QoS model

that service-based admission control can only be used to perform the authorization but not reservation phase. If admission fails, the RC will directly reject the AF request. If admission passes, a check against resource availability at the transport plane is needed in the reservation phase. If the RC does not have access to the resource availability information, the RC may send a request for admission control to another entity; for example, the TRC-FE in RACF or the NE at the transport plane. This entity then performs *resource-based admission control* based on transport information received from the RC and resource availability information that the entity gathers using its own means. It is to be noted that although multiple entities may perform admission control, the RC is the entity that makes final decision.

Clearly, prior work on admission control, including accounting-based admission control and measurement-based admission control, may be applicable to the present context. However, there are distinct differences in the way the QoS request is defined, including the richness of a session description with multiple media types, the ability of describing a codec list and the policy rules that play a role in admission. It is for further research to investigate if multiple levels of admission controls are needed and how they should be designed.

G. End-to-End QoS

The discussions so far were focused on resource management vertically for the AF-RC pair and the RC-NE pair. In this section, we discuss the end-to-end characteristic at the transport plane between two end users.

Traditional packet networks such as ATM and RSVP-based networks generally involves checking and managing the QoS resources for a new request through admission control at each NE *hop-by-hop* along the path between the end users. Admission control passes only if each link along the path has enough resources to accommodate the new request. Another extreme is to check and manage the QoS between the NEs at *end-to-end* without explicitly managing the intermediate NEs. For example, the end NEs can perform online or offline measurement to determine whether a new request should be admitted.

In the context of the architectures in this paper, it is desirable for the RC to commit resources only at the edge NEs (NE $A1$, NE $A2$, NE $Z1$ and NE $Z2$) rather than at each NE along the user-flow path, as shown in Fig. 9. This allows the RC to discover and install decisions to fewer NEs. Obviously, the resources at other intermediate NEs must also be available if QoS is to be guaranteed. This can be done, for example, by provisioning MPLS LSPs with reserved bandwidth between the respective edge NEs. For the originating domain, RC A needs to determine the respective two edge NEs the path will traverse at the transport plane. Admission control passes if there is enough resource along the path between the two edge NEs. Similarly, RC Z in the terminating domain needs to determine its respective edge NEs. There may be a communication need between RC A and RC Z to ensure that both domains select compatible edge NEs $A2$ and $Z2$. This QoS model is sometimes called the *segmented model* as each segment (domain) performs its own admission control and resource management. Ref [22] also provides discussions on resource allocation across domains.

For cases where the path between the edge NEs cannot be pinned down (e.g., with connectionless IP), bandwidth reservation at the intermediate NEs may not be supported. A possible solution is to use Differentiated Services and some over provisioning between the edge NEs. In some other cases, there may be one or more intermediate domains between the originating and terminating domains.

How resources at the intermediate domain should be managed (e.g., at the per-flow or per-aggregate level) presents an interesting research topic.

IV. Call Flow Examples

This section provides examples of how resources at the traffic plane are authorized, reserved and committed during a session setup between two users, assuming that SIP is used over the service plane.

Fig. 10 shows the call flow that triggers resource authorization and reservation at the RC. To simplify the exposition, we omit other possible elements between the AF and user 2. Service initiation begins with the exchange of SIP invite request and 183 Progress. Upon receiving 183 Progress (step 2), the AF knows the service information from the SDP offer (attached in the Invite request) and SDP answer (attached in the 183 response). The AF then sends a request for resource authorization and reservation to the RC (step 3). After performing admission control and QoS mapping, the RC makes a resource reservation at the NE (step 4). The NE may perform its own resource-based admission control. If admission passes, the NE creates the gate with the necessary parameters specified by the RC, but keep the gate closed as user 2 has not been alerted yet at this point and thus resources have not been committed. Upon receiving the acknowledgement from the NE (step 5), the RC responds to the AF (step 6).

Fig. 11 shows the call flow that triggers resource commitment at the RC. After exchanging several more SIP signaling messages, as described in [14] for example, the response from user 2 finally triggers a 200 OK message (step 1). After receiving this message, the AF sends a request for resource commitment to the RC (step 2), which follows by committing the resource at the NE (step 3). The NE normally would simply open the gate and returns an acknowledgement indicating that resources have been committed (step 4). The RC then returns the commitment response to the AF (step 5), which proceeds by continuing the signaling message at the service layer (step 6). At this point, user 1 responds with an Ack, and both users can then exchange user packets over the transport plane.

V. Conclusion

Packet network architectures with decoupled service plane and transport plane are expected to become prominent as convergence of fixed and mobile systems with a common IP core moves closer towards a reality. In this paper, we have reviewed related architectures that are being worked out in different standards development organizations. We have identified various components and issues that are unique in the architectures. In addition, there are challenging combinations of admission control approaches at different levels coupled with decomposition of end-to-end QoS into multiple domains.

VI. Acknowledgement

The author thanks Kerry Fendick for initially suggesting him to look into various relevant standards. He also thanks Dong Sun and Hui-Lan Lu for the comments on an earlier version of the paper.

References

[1] R. Braden at al,, "Resource ReSerVation Protocol (RSVP) - Functional Specification", RFC 2205, Sep. 1997.

[2] J. Rosenberg et al., "SIP: Session Initiation Protocol," RFC 3261, Jun. 2002.

[3] 3GPP TS 23.203, "Policy and charging control architecture (Release 7)," V 1.0.0, May 2006.

[4] 3GPP2 X.S0012-012-0, "Service based bearer control - Stage 2," Draft Version, Feb. 2006.

[5] ITU-T Y.RACF, "Functional architecture and requirements for resource and admission control functions in next generation networks," Draft Version 8.1, Jan. 2006.

[6] ETSI-ES 282 003, "Resource and admission control subsystem (RACS)," Draft Version 1.1.1, Mar 2006.

[7] PacketCable, "Architecture framework technical report," PKT-TR-ARCH-FRM-v01-060406, Apr. 2006.

[8] PacketCable Technical Report, "Multimedia architecture framework," PKT-TR-MM-ARCH-005-FINAL, Dec. 2005.

[9] PacketCable Specification, "Multimedia Specification," PKT-SP-MM-I03-051221, Dec. 2005.

[10] MSF Implementation Agreement, "System architecture implementation agreement," MSF-ARCH-001.01-FINAL, May 2000.

[11] MSF Implementation Agreement, "MSF Release 2 Architecture," MSF-ARCH-002.00-FINAL, Jan. 2005.

[12] 3GPP TS 23.228, "IP Multimedia Subsystem (IMS) - Stage 2 (Release 6)," V6.14.0, Jun. 2006.

[13] PacketCable Specifications, "Dynamic Quality-of-Service," PKT-SP-DQOS1.5-I02-050812, Aug. 2005.

[14] G. Camarillo, W. Marshall and J. Rosenberg, "Integration of resource management and session initiation protocol," RFC 3312, Oct. 2002.

[15] MSF Technical Report, "Bandwidth management in next generation packet networks," MSF-TR-ARCH-005-FINAL, Aug. 2005.

[16] W. Marshall, "Private Session Initiation Protocol (SIP) extensions for media authorization," RFC 3313, Jan. 2003.

[17] S. Herzog, "RSVP extensions for policy control," RFC 2750, Jan. 2000.

[18] M. Handley and V. Jacobson, "SDP: Session Description Protocol," RFC 2327, Apr. 1998.

[19] J. Rosenberg and H. Schulzrinne, "An offer/answer model with the session description protocol (SDP)," RFC 3264, Jun. 2002.

[20] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," RFC 1890, Jan. 1996.

[21] 3GPP TS 29.208, "End-to-end quality of service (QoS) signaling flows (Release 6)," Dec. 2005.

[22] T.W Anderson, I. Faynberg, H. Lu and D. Sun, "On the mechanisms for real-time application-driven resource management in next generation networks," ICIN 2006, Bordeaux, France, May 2006.

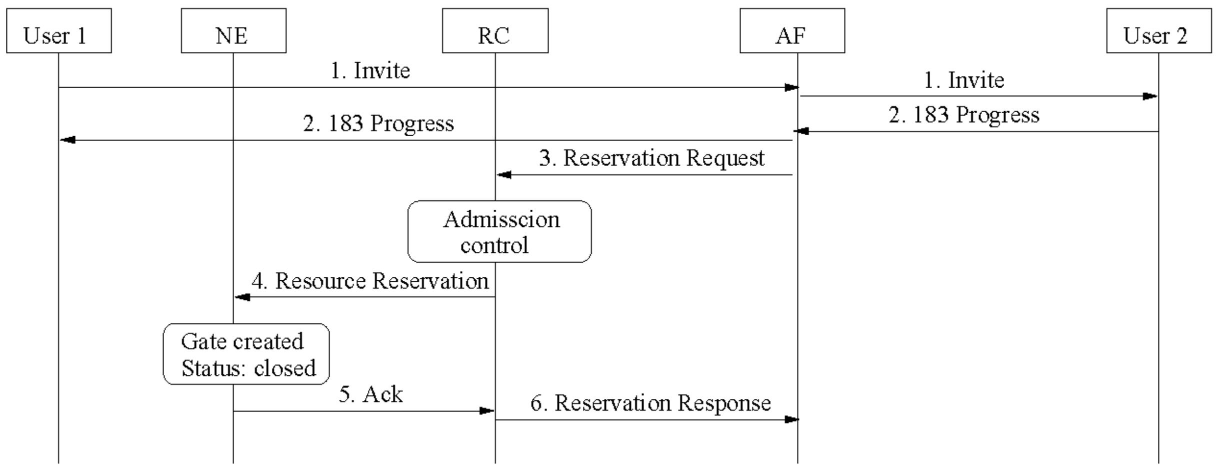

Fig. 10. Call flow example for reservation.

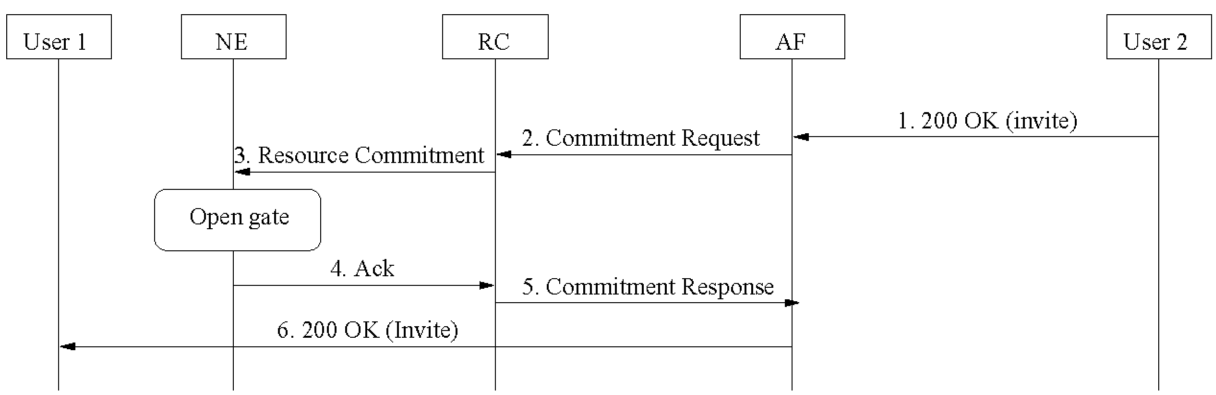

Fig. 11. Call flow example for commitment.