Network Working Group   N. Greene
Request for Comments: 2805   Nortel Networks
Category: Informational   M. Ramalho
                                                          Cisco Systems
                                                               B. Rosen
                                                                Marconi
                                                             April 2000

          Media Gateway Control Protocol Architecture and Requirements

Status of this Memo

Copyright Notice

Abstract

   This document describes protocol requirements for the Media Gateway
   Control Protocol between a Media Gateway Controller and a Media
   Gateway.

Table of Contents

1. Introduction

   This document describes requirements to be placed on the Media
   Gateway Control Protocol. When the word protocol is used on its own
   in this document it implicitly means the Media Gateway Control
   Protocol.

2. Terminology

   In this document, the key words "MUST", "MUST NOT", "REQUIRED",
   "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY",
   and "OPTIONAL" are to be interpreted as described in RFC 2119 [1] and
   indicate requirement levels for the protocol.

3. Definitions

   *   Connection

   Under the control of a Media Gateway Controller (MGC), the Media
   Gateway (MG) realizes connections. In this document, connections are
   associations of resources hosted by the MG. They typically involve
   two terminations, but may involve more.

   *   Line or Loop

   An analogue or digital access connection from a user terminal which
   carries user media content and telephony access signalling (DP, DTMF,
   BRI, proprietary business set).

   *   Media Gateway (MG) function

   A Media Gateway (MG) function provides the media mapping and/or
   transcoding functions between potentially dissimilar networks, one of
   which is presumed to be a packet, frame or cell network. For
   example, an MG might terminate switched circuit network (SCN)
   facilities (trunks, loops), packetize the media stream, if it is not
   already packetized, and deliver packetized traffic to a packet
   network. It would perform these functions in the reverse order for
   media streams flowing from the packet network to the SCN.

   Media Gateways are not limited to SCN <-> packet/frame/cell
   functions: A conference bridge with all packet interfaces could be an
   MG, as well as an (IVR) interactive voice recognition unit, an audio
   resource function, or a voice recognition system with a cell
   interface.

*   Media Gateway unit (MG-unit)

An MG-unit is a physical entity that contains an MG function and may
also contain other functions, e.g. an SG function.

*   Media Gateway Controller (MGC) function

A Media Gateway Controller (MGC) function controls a MG.

*   Media Resource

Examples of media resources are codecs, announcements, tones, and
modems, interactive voice response (IVR) units, bridges, etc.

*   Signaling Gateway (SG) function

An SG function receives/sends SCN native signalling at the edge of a
data network. For example the SG function may relay, translate or
terminate SS7 signaling in an SS7-Internet Gateway. The SG function
may also be co-resident with the MG function to process SCN
signalling associated with line or trunk terminations controlled by
the MG, such as the "D" channel of an ISDN PRI trunk.

*   Termination

A termination is a point of entry and/or exit of media flows relative
to the MG. When an MG is asked to connect two or more terminations,
it understands how the flows entering and leaving each termination
are related to each other.

Terminations are, for instance, DS0's, ATM VCs and RTP ports. Another
word for this is bearer point.

*   Trunk

An analog or digital connection from a circuit switch which carries
user media content and may carry telephony signalling (MF, R2, etc.).
Digital trunks may be transported and may appear at the Media Gateway
as channels within a framed bit stream, or as an ATM cell stream.
Trunks are typically provisioned in groups, each member of which
provides equivalent routing and service.

*   Type of Bearer

A Type of Bearer definition provides the detailed requirements for
its particular application/bearer type. A particular class of Media
Gateway, for example, would support a particular set of Bearer types.

4. Specific functions assumed within the MG

   This section provides an environment for the definition of the
   general Media Gateway Control Protocol requirements.

   MGs can be architected in many different ways depending where the
   media conversions and transcoding (if required) are performed, the
   level of programmability of resources, how conferences are supported,
   and how associated signalling is treated. The functions assumed to be
   within the MG must not be biased towards a particular architecture.

   For instance, announcements in a MG could be provided by media
   resources or by the bearer point resource or termination itself.
   Further, this difference must not be visible to MGC: The MGC must be
   able to issue the identical request to two different implementations
   and achieve the identical functionality.

   Depending on the application of the MG (e.g., trunking, residential),
   some functions listed below will be more prominent than others, and
   in some cases, functions may even disappear.

   Although media adaptation is the essence of the MG, it is not
   necessary for it to be involved every time. An MG may join two
   terminations/resources of the same type (i.e., the MG behaves as a
   switch). The required media conversion depends on the media type
   supported by the resources being joined together.

   In addition to media adaptation function, resources have a number of
   unique properties, for instance:

   *   certain types of resources have associated signalling
        capabilities (e.g., PRI signalling, DTMF),

   *   some resources perform maintenance functions (e.g., continuity
        tests),

   *   the MGC needs to know the state changes of resources (e.g., a
        trunk group going out of service),

   *   the MG retains some control over the allocation and control of
        some resources (e.g., resource name space: RTP port numbers).

   Therefore, an MG realizes point-to-point connections and conferences,
   and supports several resource functions. These functions include
   media conversion, resource allocation and management, and event
   notifications. Handling termination associated signalling is either
   done using event notifications, or is handled by the signalling
   backhaul part of a MG-unit (i.e. NOT directly handled by the MG).

MGs must also support some level of system related functions, such as establishing and maintaining some kind of MG-MGC association. This is essential for MGC redundancy, fail-over and resource sharing.

Therefore, an MG is assumed to contain these functions:

*   Reservation and release, of resources

*   Ability to provide state of resources

*   Maintenance of resources - It must be possible to make maintenance operations independent of other termination functions, for instance, some maintenance states should not affect the resources associated with that resource . Examples of maintenance functions are loopbacks and continuity tests.

*   Connection management, including connection state.

*   Media processing, using media resources: these provide services such as transcoding, conferencing, interactive voice recognition units, audio resource function units. Media resources may or may not be directly part of other resources.

*   Incoming digit analysis for terminations, interpretation of scripts for terminations

*   Event detection and signal insertion for per-channel signalling

*   Ability to configure signalling backhauls (for example, a Sigtran backhaul)

*   Management of the association between the MGC and MG, or between the MGC and MG resources.

5. Per-Call Requirements

5.1. Resource Reservation

The protocol must:

a.   Support reservation of bearer terminations and media resources for use by a particular call and support their subsequent release (which may be implicit or explicit).

b.   Allow release in a single exchange of messages, of all resources associated with a particular set of connectivity and/or associations between a given number terminations.

   c.   The MG is not required (or allowed) by the protocol to maintain
        a sense of future time: a reservation remains in effect until
        explicitly released by the MGC.

5.2. Connection Requirements

   The protocol must:

   a.   Support connections involving packet and circuit bearer
        terminations in any combination, including "hairpin" connections
        (connections between two circuit connections within the same
        MG).

   b.   Support connections involving TDM, Analogue, ATM, IP or FR
        transport in any combination.

   c.   Allow the specification of bearer plane (e.g. Frame Relay, IP,
        etc.) on a call by call basis.

   d.   Support unidirectional, symmetric bi-directional, and asymmetric
        bi-directional flows of media.

   e.   Support multiple media types (e.g. audio, text, video, T.120).

   f.   Support point-to-point and point-to-multipoint connections.

   g.   Support creation and modification of more complex flow
        topologies e.g. conference bridge capabilities. Be able to add
        or delete media streams during a call or session, and be able to
        add or subtract participants to/from a call or session.

   h.   Support inclusion of media resources into call or session as
        required. Depending on the protocol and resource type, media
        resources may be implicitly included, class-assigned, or
        individually assigned.

   i.   Provide unambiguous specification of which media flows pass
        through a point and which are blocked at a given point in time,
        if the protocol permits multiple flows to pass through the same
        point.

   j.   Allow modifications of an existing termination, for example, use
        of higher compression to compensate for insufficient bandwidth
        or changing transport network connections.

   k.   Allow the MGC to specify that a given connection has higher
        priority than other connections.

   l.   Allow a reference to a port/termination on the MG to be a
         logical identifier,

         with a one-to-one mapping between a logical identifier and a
         physical port.

   m.   Allow the MG to report events such as resource reservation and
         connection completion.

5.3. Media Transformations

   The Protocol must:

   a.   Support mediation/adaptation of flows between different types of
         transport

   b.   Support invocation of additional processing such as echo
         cancellation.

   c.   Support mediation of flows between different content encoding
         (codecs, encryption/decryption)

   d.   Allow the MGC to specify whether text telephony/FAX/data modem
         traffic is to be terminated at the MG, modulated/demodulated,
         and converted to packets or forwarded by the MG in the media
         flow as voice band traffic.

   e.   Allow the MGC to specify that Dual-Tone MultiFrequency (DTMF)
         digits or other line and trunk signals and general Multi-
         Frequency (MF) tones are to be processed in the MG and how these
         digits/signals/tones are to be handled. The MGC must be able to
         specify any of the following handling of such
         digits/signals/tones:

   1.   The digits/signals/tones are to be encoded normally in the audio
         RTP stream (e.g., no analysis of the digits/signals/tones).

   2.   Analyzed and sent to the MGC.

   3.   Received from the MGC and inserted in the line-side audio
         stream.

   4.   Analyzed and sent as part of a separate RTP stream (e.g., DTMF
         digits sent via a RTP payload separate from the audio RTP
         stream).

   5.   Taken from a separate RTP stream and inserted in the line-side
         audio stream.

   6.   Handled according to a script of instructions. For all but the
        first case, an option to mute the digits/signals/tones with
        silence, comfort noise, or other means (e.g., notch filtering of
        some telephony tones) must be provided. As detection of these
        events may take up to tens of milliseconds, the first few
        milliseconds of such digit/signal/tone may be encoded and sent
        in the audio RTP stream before the digit/signal/tone can be
        verified. Therefore muting of such digits/signals/tones in the
        audio RTP stream with silence or comfort noise is understood to
        occur at the earliest opportunity after the digit/signal/tone is
        verified.

   f.   Allow the MGC to specify signalled flow characteristics on
        circuit as well as on packet bearer connections, e.g. u-law/a-
        law.

   g.   Allow for packet/cell transport adaptation only (no media
        adaptation) e.g. mid-stream (packet-to-packet)
        transpacketization/transcoding, or ATM AAL5 to and from ATM AAL2
        adaptation.

   h.   Allow the transport of audio normalization levels as a setup
        parameter, e.g., for conference bridging.

   i.   Allow conversion to take place between media types e.g., text to
        speech and speech to text.

5.4. Signal/Event Processing and Scripting

   The Protocol must:

   a.   Allow the MGC to enable/disable monitoring for specific
        supervision events at specific circuit terminations

   b.   Allow the MGC to enable/disable monitoring for specific events
        within specified media streams

   c.   Allow reporting of detected events on the MG to the MGC. The
        protocol should provide the means to minimize the messaging
        required to report commonly-occurring event sequences.

   d.   Allow the MGC to specify other actions (besides reporting) that
        the MG should take upon detection of specified events.

   e.   Allow the MGC to enable and/or mask events.

   f.   Provide a way for MGC to positively acknowledge event
        notification.

   g.   Allow the MGC to specify signals (e.g., supervision, ringing) to
         be applied at circuit terminations.

   h.   Allow the MGC to specify content of extended duration
         (announcements, continuous tones) to be inserted into specified
         media flows.

   i.   Allow the MGC to specify alternative conditions (detection of
         specific events, timeouts) under which the insertion of
         extended-duration signals should cease.

   j.   Allow the MGC to download, and specify a script to be invoked on
         the occurrence of an event.

   k.   Specify common events and signals to maximize MG/MGC
         interworking.

   l.   Provide an extension mechanism for implementation defined events
         and signals with, for example, IANA registration procedures. It
         may be useful to have an Organizational Identifier (i.e. ITU,
         ETSI, ANSI,) as part of the registration mechanism.

   m.   The protocol shall allow the MGC to request the arming of a
         mid-call trigger even after the call has been set up.

5.5. QoS/CoS

   The Protocol must:

   a.   Support the establishment of a bearer channel with a specified
         QoS/CoS.

   b.   Support the ability to specify QoS for the connection between
         MGs, and by direction.

   c.   Support a means to change QoS during a connection, as a whole
         and by direction.

   d.   Allow the MGC to set QOS thresholds and receive notification
         when such thresholds cannot be maintained.

   e.   Allow the jitter buffer parameters on RTP channels to be
         specified at connection setup.

5.6. Test Support

   The protocol must:

   a.   Support of the different types of PSTN Continuity Testing (COT)
         for both the originating and terminating ends of the circuit
         connection (2-wire and 4- wire).

   b.   Specifically support test line operation (e.g. 103, 105, 108).

5.7. Accounting

   The protocol must:

   a.   Support a common identifier to mark resources related to one
         connection.

   b.   Support collection of specified accounting information from MGs.

   c.   Provide the mechanism for the MGC to specify that the MG report
         accounting information automatically at end of call, in mid-call
         upon request, at specific time intervals as specified by the MGC
         and at unit usage thresholds as specified by the MGC.

   d.   Specifically support collection of:

   *   start and stop time, by media flow,

   *   volume of content carried (e.g. number of packets/cells
         transmitted, number received with and without error, inter-
         arrival jitter), by media flow,

   *   QOS statistics, by media flow.

   e.   Allow the MGC to have some control over which statistics are
         reported, to enable it to manage the amount of information
         transferred.

5.8. Signalling Control

   Establishment and provisioning of signalling backhaul channels (via
   SIGTRAN for example) is out of scope. However, the MG must be
   capable of supporting detection of events, and application of signals
   associated with basic analogue line, and CAS type signalling. The
   protocol must:

   a.   Support the signalling requirements of analogue lines and
         Channel Associated Signaling (CAS).

   b.   Support national variations of such signalling.

   c.   Provide mechanisms to support signalling without requiring MG-
         MGC timing constraints beyond that specified in this document.

   d.   Must not create a situation where the MGC and the MG must be
         homologated together as a mandatory requirement of using the
         protocol;

         i.e. it must be possible to optionally conceal signaling type
         variation from the MGC.

6. Resource Control

6.1. Resource Status Management

   The protocol must:

   a.   Allow the MG to report changes in status of physical entities
         supporting bearer terminations, media resources, and facility-
         associated signalling channels, due to failures, recovery, or
         administrative action. It must be able to report whether a
         termination is in service or out of service.

   b.   Support administrative blocking and release of TDM circuit
         terminations.

   Note: as the above point only relates to ISUP-controlled circuits, it
   may be unnecessary to require this since the MGC controls their use.
   However, it may be meaningful for MF and R2-signalled trunks, where
   supervisory states are set to make the trunks unavailable at the far
   end.

   c.   Provide a method for the MGC to request that the MG release all
         resources under the control of a particular MGC currently in
         use, or reserved, for any or all connections.

   d.   Provide an MG Resource Discovery mechanism which must allow an
         MGC to discover what resources the MG has. Expressing resources
         can be an arbitrarily difficult problem and the initial release
         of the protocol may have a simplistic view of resource
         discovery.

         At a minimum, resource discovery must enumerate the names of
         available circuit terminations and the allowed values for
         parameters supported by terminations.

The protocol should be defined so that simple gateways could
respond with a relatively short, pre-stored response to the
discovery request mechanism. In general, if the protocol defines
a mechanism that allows the MGC to specify a setting or
parameter for a resource or connection in the MG, and MGs are
not required to support all possible values for that setting or
parameter, then the discovery mechanism should provide the MGC
with a method to determine what possible values such settings or
parameters are supported in a particular MG.

e.   Provide a mechanism to discover the current available resources
in the MG, where resources are dynamically consumed by
connections and the MGC cannot reasonably or reliably track the
consumption of such resources. It should also be possible to
discover resources currently in use, in order to reconcile
inconsistencies between the MGC and the MG.

f.   Not require an MGC to implement an SNMP manager function in
order to discover capabilities of an MG that may be specified
during context establishment.

6.2. Resource Assignment

The protocol must:

a.   Provide a way for the MG to indicate that it was unable to
perform a requested action because of resource exhaustion, or
because of temporary resource unavailability.

b.   Provide an ability for the MGC to indicate to an MG the resource
to use for a call (e.g. DS0) exactly, or indicate a set of
resources (e.g. pick a DS0 on a T1 line or a list of codec
types) via a "wild card" mechanism from which the MG can select
a specific resource for a call (e.g. the 16th timeslot, or
G.723).

c.   Allow the use of DNS names and IP addresses to identify MGs and
MGCs. This shall not preclude using other identifiers for MGs or
MGCs when other non IP transport technologies for the protocol
are used.

7. Operational/Management Requirements

7.1. Assurance of Control/Connectivity

To provide assurance of control and connectivity, the protocol must
provide the means to minimize duration of loss of control due to loss
of contact, or state mismatches.

The protocol must:

a.   Support detection and recovery from loss of contact due to
     failure/congestion of communication links or due to MG or MGC
     failure.

     Note that failover arrangements are one of the mechanisms which
     could be used to meet this requirement.

b.   Support detection and recovery from loss of synchronized view of
     resource and connection states between MGCs and MGs. (e.g.
     through the use of audits).

c.   Provide a means for MGC and MG to provide each other with
     booting and reboot indications, and what the MG's configuration
     is.

d.   Permit more than one backup MGC and provide an orderly way for
     the MG to contact one of its backups.

e.   Provide for an orderly switchback to the primary MGC after it
     recovers. How MGCs coordinate resources between themselves is
     outside the scope of the protocol.

f.   Provide a mechanism so that when an MGC fails, connections
     already established can be maintained. The protocol does not
     have to provide a capability to maintain connections in the
     process of being connected, but not actually connected when the
     failure occurs.

g.   The Protocol must allow the recovery or redistribution of
     traffic without call loss.

7.2. Error Control

The protocol must:

a.   Allow for the MG to report reasons for abnormal failure of lower
     layer connections e.g. TDM circuit failure, ATM VCC failure.

b.   Allow for the MG to report Usage Parameter Control (UPC) events.

c.   Provide means to ameliorate potential synchronization or focused
     overload of supervisory/signaling events that can be detrimental
     to either MG or MGC operation. Power restoration or signaling
     transport re-establishment are typical sources of potentially
     detrimental signaling showers from MG to MGC or vice-versa.

    d.   Allow the MG to notify the MGC that a termination was terminated
         and communicate a reason when a terminations is taken out-of-
         service unilaterally by the MG due to abnormal events.

    e.   Allow the MGC to acknowledge that a termination has been taken
         out-of-service.

    f.   Allow the MG to request the MGC to release a termination and
         communicate a reason.

    g.   Allow the MGC to specify, as a result of such a request its
         decision to take termination down, leave it as is or modify it.

7.3. MIB Requirements

    The Protocol must define a common MG MIB, which must be extensible,
    but must:

    a.   Provide information on:

    *   mapping between resources and supporting physical entities.

    *   statistics on quality of service on the control and signalling
         backhaul interfaces.

    *   statistics required for traffic engineering within the MG.

    b.   The protocol must allow the MG to provide to the MGC all
         information the MGC needs to provide in its MIB.

    c.   MG MIB must support implementation of H.341 by either the MG,
         MGC, or both acting together.

8. General Protocol Requirements

    The protocol must:

    a.   Support multiple operations to be invoked in one message and
         treated as a single transaction.

    b.   Be both modular and extensible. Not all implementations may wish
         to support all of the possible extensions for the protocol. This
         will permit lightweight implementations for specialized tasks
         where processing resources are constrained. This could be
         accomplished by defining particular profiles for particular uses
         of the protocol.

   c.   Be flexible in allocation of intelligence between MG and MGC.
         For example, an MGC may want to allow the MG to assign
         particular MG resources in some implementations, while in
         others, the MGC may want to be the one to assign MG resources
         for use.

   d.   Support scalability from very small to very large MGs: The
         protocol must support MGs with capacities ranging from one to
         millions of terminations.

   e.   Support scalability from very small to very large MGC span of
         control: The protocol should support MGCs that control from one
         MG to a few tens of thousands of MGs.

   f.   Support the needs of a residential gateway that supports one to
         a few lines, and the needs of a large PSTN gateway supporting
         tens of thousands of lines. Protocol mechanisms favoring one
         extreme or the other should be minimized in favor of more
         general purpose mechanism applicable to a wide range of MGs.
         Where special purpose mechanisms are proposed to optimize a
         subset of implementations, such mechanisms should be defined as
         optional, and should have minimal impact on the rest of the
         protocol.

   g.   Facilitate MG and MGC version upgrades independently of one
         another. The protocol must include a version identifier in the
         initial message exchange.

   h.   Facilitate the discovery of the protocol capabilities of the one
         entity to the other.

   i.   Specify commands as optional (they can be ignored) or mandatory
         (the command must be rejected), and within a command, to specify
         parameters as optional (they can be ignored) or mandatory (the
         command must be rejected).

8.1. MG-MGC Association Requirements

   The Protocol must:

   a.   Support the establishment of a control relationship between an
         MGC and an MG.

   b.   Allow multiple MGCs to send control messages to an MG. Thus, the
         protocol must allow control messages from multiple signalling
         addresses to a single MG.

    c.   Provide a method for the MG to tell an MGC that the MG received
        a command for a resource that is under the control of a
        different MGC.

    d.   Support a method for the MG to control the rate of requests it
        receives from the MGC (e.g. windowing techniques, exponential
        back-off).

    e.   Support a method for the MG to tell an MGC that it cannot handle
        any more requests.

8.2. Performance Requirements

    The protocol must:

    a.   Minimize message exchanges between MG and MGC, for example
        during boot/reboot, and during continuity tests.

    b.   Support Continuity test constraints which are a maximum of 200ms
        cross-MGC IAM (IAM is the name given to an SS7 connection setup
        msg) propagation delay, and a maximum of 200ms from end of
        dialing to IAM emission.

    c.   Make efficient use of the underlying transport mechanism. For
        example, protocol PDU sizes vs. transport MTU sizes needs to be
        considered in designing the protocol.

    d.   Not contain inherent architectural or signaling constraints that
        would prohibit peak calling rates on the order of 140
        calls/second on a moderately loaded network.

    e.   Allow for default/provisioned settings so that commands need
        only contain non-default parameters.

9. Transport

9.1. Assumptions made for underlying network

    The protocol must assume that the underlying network:

    a.   May be over large shared networks: proximity assumptions are not
        allowed.

    b.   Does not assure reliable delivery of messages.

    c.   Does not guarantee ordering of messages: Sequenced delivery of
        messages associated with the same source of events is not
        assumed.

   d.   Does not prevent duplicate transmissions.

9.2. Transport Requirements

   The protocol must:

   a.   Provide the ability to abort delivery of obsolete messages at
         the sending end if their transmission has not been successfully
         completed. For example, aborting a command that has been
         overtaken by events.

   b.   Support priority messages: The protocol shall allow a command
         precedence to allow priority messages to supercede non-priority
         messages.

   c.   Support of large fan-out at the MGC.

   d.   Provide a way for one entity to correlate commands and responses
         with the other entity.

   e.   Provide a reason for any command failure.

   f.   Provide that loss of a packet not stall messages not related to
         the message(s) contained in the packet lost.

   Note that there may be enough protocol reliability requirements here
   to warrant a separate reliable transport layer be written apart from
   the Media Gateway Control Protocol. Also need to compare Megaco
   reliable transport requirements with similar Sigtran requirements.

10. Security Requirements

   Security mechanisms may be specified as provided in underlying
   transport mechanisms, such as IPSEC. The protocol, or such
   mechanisms, must:

   a.   Allow for mutual authentication at the start of an MGC-MG
         association

   b.   Allow for preservation of the of control messages once the
         association has been established.

   c.   Allow for optional confidentiality protection of control
         messages. The mechanism should allow a choice in the algorithm
         to be used.

   d.   Operate across untrusted domains in a secure fashion.

   e.   Support non-repudiation for a customer-located MG talking to a
         network operator's MGC.

   f.   Define mechanisms to mitigate denial of service attacks

   Note: the protocol document will need to include an extended
   discussion of security requirements, offering more precision on each
   threat and giving a complete picture of the defense including non-
   protocol measures such as configuration.

   g.   It would be desirable for the protocol to be able to pass
         through commonly-used firewalls.

11. Requirements specific to particular bearer types

   The bearer types listed in Table 1 can be packaged into different
   types of MGs. Examples are listed in the following sections. How
   they are packaged is outside the scope of the general Media Gateway
   control protocol. The protocol must support all types of bearer types
   listed in Table 1.

Table 1: Bearer Types and Applications

Bearer Type	Applications	Transit Network
=============	==============	=================
Trunk+ISUP	trunking/access Voice,Fax,NAS, Multimedia	IP, ATM, FR
Trunk+MF	trunking/access Voice,Fax,NAS, Multimedia	IP, ATM, FR
ISDN	trunking/access Voice,Fax,NAS, Multimedia	IP, ATM, FR
Analogue	Voice,Fax, Text Telephony	IP, ATM, FR
Termination in a Restricted Capability Gateway	Voice,Fax, Text Telephony	IP, ATM, FR
Application Termination	IVR,ARF, Announcement Server, Voice Recognition Server,...	
Multimedia H.323	H.323 Multimedia Gateway and MCU	IP, ATM, FR
Multimedia H.320	H.323 GW and MCU	ISDN, IP, ATM, FR

11.1. Media-specific Bearer Types

   This section describes requirements for handling terminations
   attached to specific types of networks.

11.1.1. Requirements for TDM PSTN (Circuit)

   This bearer type is applicable to a Trunking GW, Access GW, ...

   The protocol must allow:

   a.   the MGC to specify the encoding to use on the attached circuit.

   b.   In general, if something is set by a global signalling protocol
        (e.g. ISUP allows mu-Law or A-Law to be signaled using ISUP)
        then it must be settable by the protocol.

   c.   TDM attributes:

   *   Echo cancellation,

   *   PCM encoding or other voice compression (e.g. mu-law or A-law),

   *   encryption,

   *   rate adaptation (e.g. V.110, or V.120).

   d.   for incoming calls, identification of a specific TDM circuit
        (timeslot and facility).

   e.   for calls outgoing to the circuit network, identification of a
        specific circuit or identification of a circuit group with the
        indication that the MG must select and return the identification
        of an available member of that group.

   f.   specification of the default encoding of content passing to and
        from a given circuit, possibly on a logical or physical circuit
        group basis.

   g.   specification at any point during the life of a connection of
        variable aspects of the content encoding, particularly including
        channel information capacity.

   h.   specification at any point during the life of a connection of
        loss padding to be applied to incoming and outgoing media
        streams at the circuit termination.

   i.   specification at any point during the life of a connection of
        the applicability of echo cancellation to the outgoing media
        stream.

   j.   Multi-rate calls to/from the SCN.

   k.   H-channel (n x 64K) calls to/from the SCN.

   l.   B channel aggregation protocols for creating high speed channels
        for multimedia over the SCN.

   m.   Modem terminations and negotiations.

   The protocol may also allow:

   n.   specification of sub-channel media streams,

   o.   specification of multi-channel media streams.

11.1.2. Packet Bearer Type

   The protocol must be able to specify:

   a.   ingress and egress coding (i.e. the way packets coming in and
         out are encoded) (including encryption).

   b.   near and far-end ports and other session parameters for RTP and
         RTCP.

   The protocol must support reporting of:

   c.   re-negotiation of codec for cause - for further study

   d.   on Trunking and Access Gateways, resources capable of more than
         one active connection at a time must also be capable of mixing
         and packet duplication.

   The protocol must allow:

   e.   specification of parameters for outgoing and incoming packet
         flows at separate points in the life of the connection (because
         far-end port addresses are typically obtained through a separate
         signalling exchange before or after the near-end port addresses
         are assigned).

   f.   the possibility for each Media Gateway to allocate the ports on
         which it will receive packet flows (including RTCP as well as
         media streams) and report its allocations to the Media Gateway
         Controller for signalling to the far end. Note that support of
         different IP backbone providers on a per call basis would
         require that the ports on which packets flow be selected by the
         MGC. (but only if the IP address of the MG is different for each
         backbone provider).

   g.   the specification at any point during the life of a connection
         of RTP payload type and RTP session number for each RTP-
         encapsulated media flow.

   h.   the ability to specify whether outgoing flows are to be uni-cast
         or multi-cast. Note that on an IP network this information is
         implicit in the destination address, but in other networks this
         is a connection parameter.

   i.   invoking of encryption/decryption on media flows and
         specification of the associated algorithm and key.

The protocol should also allow:

j.   the MGC to configure non-RTP (proprietary or other) encapsulated
      packet flows.

11.1.3. Bearer type requirements for ATM

This bearer type is applicable to Trunking GW, Access GW, ...

11.1.3.1. Addressing

a.   The protocol must be able to specify the following termination
      attributes:

*   VC identifier,

*   VC identifier plus AAL2 slot, and variant of these allowing the
      gateway to choose (part of) the identifier,

*   remote termination network address, remote MG name.

b.   Allow specification of an ATM termination which is to be
      assigned to an MG connection as a VC identifier, a VC identifier
      plus AAL2 slot, a wild-carded variant of either of these. A
      remote termination network address, or a remote MG name could
      also be used when the MG can select the VC and change the VC
      during the life of the connection by using ATM signalling.

c.   Provide an indication by the MG of the VC identifier and
      possibly AAL2 slot of the termination actually assigned to a
      connection.

d.   Provide a means to refer subsequently to that termination.

e.   Refer to an existing VCC as the physical interface + Virtual
      Path Identifier (VPI) + Virtual Circuit Identifier (VCI).

f.   Where the VCC is locally established (SVCs signalled by the
      Gateway through UNI or PNNI signalling or similar), the VCC must
      be indirectly referred to in terms which are of significance to
      both ends of the VCC. For example, a global name or the ATM
      address of the ATM devices at each end of the VCC. However, it
      is possible/probable that there may be several VCCs between a
      given pair of ATM devices. Therefore the ATM address pair must
      be further resolved by a VCC identifier unambiguous within the
      context of the ATM address pair.

g.   refer to a VCC as the Remote GW ATM End System Address + VCCI.

    h.   allow the VCCI to be selected by the MG or imposed on the MG.

    i.   support all ATM addressing variants (e.g. ATM End System Address
          (AESA) and E.164).

11.1.3.2. Connection related requirements

    The protocol must:

    a.   Allow for the de-coupling of creation/deletion of the narrow-
          band connection from the creation/deletion of the underlying
          VCC.

    b.   Allow for efficient disconnection of all connections associated
          with a physical port or VCC. As an example, this could aggregate
          disconnections across a broadband circuit which experienced a
          physical error.

    c.   Allow the connection established using this protocol to be
          carried over a VCC, which may be a:

    *   PVC or SPVC,

    *   an SVC established on demand, either by the MGC itself or by a
          broker acting on its behalf or,

    *   an SVC originated as required by the local MG, or by the remote
          end to the local MG through UNI or PNNI signalling.

    d.   Allow ATM transport parameters and QoS parameters to be passed
          to the MG.

    e.   Allow blocking and unblocking of a physical interface, a VCC or
          an AAL1/AAL2 channel.

    The protocol should:

    f.   Where a VCC is required to be established on a per narrow-band
          call basis, allow all necessary information to be passed in one
          message.

11.1.3.3. Media adaptation

    The protocol must:

    a.   Allow AAL parameters to be passed to the MG.

   b.   Allow AAL1/AAL2 multiple narrow-band calls to be mapped to a
        single VCC. For AAL2, these calls are differentiated within each
        VCC by a AAL2 channel identifier. An AAL2 connection may span
        more than 1 VCC and transit AAL2 switching devices.   ITU
        Q.2630.1 [2] defines an end-to-end identifier called the Served
        User Generated Reference (SUGR). It carries information from the
        originating user of the AAL2 signalling protocol to the
        terminating user transparently and unmodified.

   c.   Allow unambiguous binding of a narrow band call to an AAL2
        connection identifier, or AAL1 channel, within the specified
        VCC.

   d.   Allow the AAL2 connection identifier, or AAL1 channel, to be
        selected by the MG or imposed on the MG.

   e.   Allow the use of the AAL2 channel identifier (cid) instead of
        the AAL2 connection identifier.

   f.   Allow the AAL2 voice profile to be imposed or negotiated before
        the start of the connection. AAL2 allows for variable length
        packets and varying packet rates, with multiple codecs possible
        within a given profile. Thus a given call may upgrade or
        downgrade the codec within the lifetime of the call. Idle
        channels may generate zero bandwidth. Thus an AAL2 VCC may vary
        in bandwidth and possibly exceed its contract. Congestion
        controls within a gateway may react to congestion by modifying
        codec rates/types.

   g.   Allow the MGC to instruct the MG on how individual narrow-band
        calls behave under congestion.

   h.   Allow for the MGC to specify an AAL5 bearer, with the following
        choices:

   *   Per ATM Forum standard AF-VTOA-0083 [4],

   *   RTP with IP/UDP,

   *   RTP without IP/IDP per H.323v2 Annex C [5],

   *   Compressed RTP per ATM Forum AF-SAA-0124.000 [6].

   i.   Allow unambiguous binding of a narrow band call to an AAL1
        channel within the specified VCC. (In AAL1, multiple narrow-band
        calls may be mapped to a single VCC.)

11.1.3.4. Reporting requirements

   The protocol should:

   a.   Allow any end-of-call statistics to show loss/restoration of
        underlying VCC within the calls duration, together with duration
        of loss.

   b.   Allow notification, as requested by MGC, of any congestion
        avoidance actions taken by the MG.

   The protocol must:

   c.   Allow for ATM VCCs or AAL2 channels to be audited by the MGC.

   d.   Allow changes in status of ATM VCCs or AAL2 channels to be
        notified as requested by the MGC.

   e.   Allow the MGC to query the resource and endpoint availability.
        Resources may include VCCs, and DSPs. VCCs may be up or down.
        End-points may be connection-free, connected or unavailable.

11.1.3.5. Functional requirements

   The protocol must:

   a.   Allow an MGC to reserve a bearer, and specify a route for it
        through the network.

11.2. Application-Specific Requirements

11.2.1. Trunking Gateway

   A Trunking Gateway is an interface between SCN networks and Voice
   over IP or Voice over ATM networks. Such gateways typically
   interface to SS7 or other NNI signalling on the SCN and manage a
   large number of digital circuits.

   The protocol must:

   a.   Provide circuit and packet-side loopback.

   b.   Provide circuit-side n x 64kbs connections.

   c.   Provide subrate and multirate connections for further study.

     d.   Provide the capability to support Reporting/generation of
          per-trunk CAS signalling (DP, DTMF, MF, R2, J2, and national
          variants).

     e.   Provide the capability to support reporting of detected DTMF
          events either digit-by-digit, as a sequence of detected digits
          with a flexible mechanism For the MG to determine the likely end
          of dial string, or in a separate RTP stream.

     f.   Provide the capability to support ANI and DNIS generation and
          reception.

11.2.2. Access Gateway

   An Access Gateway connects UNI interfaces like ISDN (PRI and BRI) or
   traditional analog voice terminal interfaces, to a Voice over IP or
   Voice over ATM network, or Voice over Frame Relay network.

   The Protocol must:

     a.   Support detection and generation of analog line signaling
          (hook-state, ring generation).

     b.   Provide the capability to support reporting of detected DTMF
          events either digit-by-digit, as a sequence of detected digits
          with a flexible mechanism For the MG to determine the likely end
          of dial string, or in a separate RTP stream.

     c.   Not require scripting mechanisms, event buffering, digit map
          storage when implementing restricted function (1-2 line)
          gateways with very limited capabilities.

     d.   Provide the capability to support CallerID generation and
          reception.

   Proxying of the protocol is for further study.

11.2.3. Trunking/Access Gateway with fax ports

     a.   the protocol must be able to indicate detection of fax media.

     b.   the protocol must be able to specify T.38 for the transport of
          the fax.

     c.   the protocol must be able to specify G.711 encoding for
          transport of fax tones across a packet network.

11.2.4. Trunking/Access Gateway with text telephone access ports

   An access gateway with ports capable of text telephone communication,
   must provide communication between text telephones in the SCN and
   text conversation channels in the packet network.

   Text telephone capability of ports is assumed to be possible to
   combine with other options for calls as described in section 11.2.6
   (e.) on "Adaptable NASes".

   The port is assumed to adjust for the differences in the supported
   text telephone protocols, so that the text media stream can be
   communicated T.140 coded in the packet network without further
   transcoding [7].

   The protocol must be capable of reporting the type of text telephone
   that is connected to the SCN port. The foreseen types are the same as
   the ones supported by ITU-T V.18: DTMF, EDT, Baudot-45, Baudot-50,
   Bell, V.21, Minitel and V.18. It should be possible to control which
   protocols are supported. The SCN port is assumed to contain ITU-T
   V.18 functionality [8].

   The protocol must be able to control the following functionality
   levels of text telephone support:

   a.   Simple text-only support: The call is set into text mode from
        the beginning of the call, in order to conduct a text-only
        conversation.

   b.   Alternating text-voice support: The call may begin in voice mode
        or text mode and, at any moment during the call, change mode on
        request by the SCN user. On the packet side, the two media
        streams for voice and text must be opened, and it must be
        possible to control the feeding of each stream by the protocol.

   c.   Simultaneous text and voice support: The call is performed in a
        mode when simultaneous text and voice streams are supported. The
        call may start in voice mode and during the call change state to
        a text-and-voice call.

   A port may implement only level a, or any level combination of a, b
   and c, always including level a.

   The protocol must support:

   d.   A text based alternative to the interactive voice response, or
        audio resource functionality of the gateway when the port is
        used in text telephone mode.

    e.   Selection of what national translation table to be used between
          the Unicode based T.140 and the 5-7 bit based text telephone
          protocols.

    f.   Control of the V.18 probe message to be used on incoming calls.

11.2.5. Network Access Server

    A NAS is an access gateway, or Media Gateway (MG), which terminates
    modem signals or synchronous HDLC connections from a network (e.g.
    SCN or xDSL network) and provides data access to the packet network.
    Only those requirements specific to a NAS are described here.

    Figure 1 provides a reference architecture for a Network Access
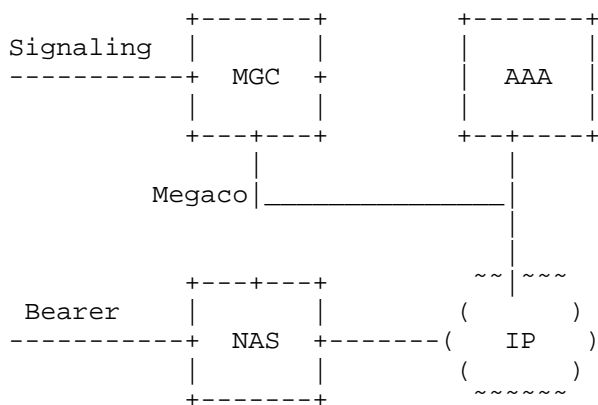   Server (NAS). Signaling comes into the MGC and the MGC controls the
    NAS.

```
                    +-------+          +-------+
        Signaling   |       |          |       |
        -----------+  MGC   +          |  AAA  |
                    |       |          |       |
                    +---+---+          +--+----+
                        |                 |
            Megaco|_____|
                        |
                        |
                    +---+---+          ~~ |~~~
        Bearer      |       |          (       )
        -----------+  NAS   +-------(   IP   )
                    |       |          (       )
                    +-------+          ~~~~~~
```

                  Figure 1: NAS reference architecture

    The Protocol must support:

    a.   Callback capabilities:

    *   Callback

    b.   Modem calls. The protocol must be able to specify the modem
          type(s) to be used for the call.

    c.   Carriage of bearer information. The protocol must be able to
          specify the data rate of the TDM connection (e.g., 64 kbit/s, 56
          kbit/s, 384 kbit/s), if this is available from the SCN.

   d.   Rate Adaptation: The protocol must be able to specify the type
        of rate adaptation to be used for the call including indicating
        the subrate, if this is available from the SCN (e.g. 56K, or
        V.110 signaled in Bearer capabilities with subrate connection of
        19.2kbit/s).

   e.   Adaptable NASes: The protocol must be able to support multiple
        options for an incoming call to allow the NAS to dynamically
        select the proper type of call. For example, an incoming ISDN
        call coded for "Speech" Bearer Capability could actually be a
        voice, modem, fax, text telephone, or 56 kbit/s synchronous
        call. The protocol should allow the NAS to report back to the
        MGC the actual type of call once it is detected.

   The 4 basic types of bearer for a NAS are:

   1.   Circuit Mode, 64-kbps, 8-khz structured, Speech

   2.   Circuit Mode, 64-kbps, 8-khz structured, 3.1-khz, Audio

   3.   Circuit Mode, 64-kbps, 8-khz structured, Unrestricted Digital
        Transmission-Rate Adapted from 56-kbps

   4.   Circuit Mode, 64-kbps, 8-khz structure, Unrestricted Digital
        Transmission

   f.   Passage of Called and Calling Party Number information to the
        NAS from the MGC. Also, passage of Charge Number/Billing Number,
        Redirecting Number, and Original Call Number, if known, to the
        NAS from the MGC. If there are other Q.931 fields that need to
        be passed from the MGC to the MG, then it should be possible to
        pass them [9].

   g.   Ability for the MGC to direct the NAS to connect to a specific
        tunnel, for example to an LNS, or to an AAA server.

   h.   When asked by the MGC, be able to report capability information,
        for example, connection types (V.34/V90/Synch ISDN..), AAA
        mechanism (RADIUS/DIAMETER/..), access type (PPP/SLIP/..) after
        restart or upgrade.

11.2.6. Restricted Capability Gateway

   The requirements here may also be applied to small analog gateways,
   and to cable/xDSL modems. See also the section on access gateways.

The Protocol must support:

a.   The ability to provide a scaled down version of the protocol.
     When features of the protocol are not supported, an appropriate
     error message must be sent. Appropriate default action must be
     defined. Where this is defined may be outside the scope of the
     protocol.

b.   The ability to provide device capability information to the MGC
     with respect to the use of the protocol.

11.2.7. Multimedia Gateway

The protocol must have sufficient capability to support a multimedia
gateway. H.320 and H.324 are characterized by a single data stream
with multiple media streams multiplexed on it.

If the mapping is from H.320 or H.324 on the circuit side, and H.323
on the packet side, it is assumed that the MG knows how to map
respective subchannels from H.320/H.324 side to streams on packet
side. If extra information is required when connecting two
terminations, then it must be supplied so that the connections are
not ambiguous.

The Multimedia Gateway:

1)   should support Bonding Bearer channel aggregation,

2)   must support 2xB (and possibly higher rates) aggregation via
     H.221,

3)   must be able to dynamically change the size of audio, video and
     data channels within the h.320 multiplex,

4)   must react to changes in the H.320 multiplex on 20 msec
     boundaries,

5)   must support TCS4/IIS BAS commands,

6)   must support detection and creation of DTMF tones,

7)   should support SNMP MIBS as specified in H.341 [3]

a.   If some of the above cannot be handled by the MGC to MG protocol
     due to timing constraints, then it is likely that the H.245 to
     H.242 processing must take place in the MG. Otherwise, support
     for this functionality in the multimedia gateway are protocol
     requirements.

   b.   It must be possible on a call by call basis for the protocol to
        specify different applications. Thus, one call might be PSTN to
        PSTN under SS7 control, while the next might be ISDN/H.320 under
        SS7 control to H.323. This is only one example; the key
        requirement is that the protocol not prevent such applications.

11.2.8. Audio Resource Function

   An Audio Resource Function (ARF) consists of one or more functional
   modules which can be deployed on an stand alone media gateway server
   IVR, Intelligent Peripheral, speech/speaker recognition unit, etc. or
   a traditional media gateway. Such a media gateway is known as an
   Audio Enabled Gateway (AEG) if it performs tasks defined in one or
   more of the following ARF functional modules:

                 Play Audio,
                 DTMF Collect,
                 Record Audio,
                 Speech Recognition,
                 Speaker Verification/Identification,
                 Auditory Feature Extraction/Recognition, or
                 Audio Conferencing.

   Additional ARF function modules that support human to machine
   communications through the use of telephony tones (e.g., DTMF) or
   auditory means (e.g. speech) may be appended to the AEG definition
   in future versions of these requirements.

   Generic scripting packages for any module must support all the
   requirements for that module. Any package extension for a given
   module must include, by inheritance or explicit reference, the
   requirements for that given module.

   The protocol requirements for each of the ARF modules are provided in
   the following subsections.

11.2.8.1. Play Audio Module

   a.   Be able to provide the following basic operation:

   -    request an ARF MG to play an announcement.

   b.   Be able to specify these play characteristics:

   -    Play volume

   -    Play speed

-     Play iterations

-     Interval between play iterations

-     Play duration

c.   Permit the specification of voice variables such as DN, number,
      date, time, etc. The protocol must allow specification of both
      the value (eg 234-3456), and well as the type (Directory
      number).

d.   Using the terminology that a segment is a unit of playable
      speech, or is an abstraction that is resolvable to a unit of
      playable speech, permit specification of the following segment
      types:

-     A provisioned recording.

-     A block of text to be converted to speech.

-     A block of text to be displayed on a device.

-     A length of silence qualified by duration.

-     An algorithmically generated tone.

-     A voice variable, specified by type and value. Given a variable
      type and value, the IVR/ARF unit would dynamically assemble the
      phrases required for its playback.

-     An abstraction that represents a sequence of audio segments.
      Nesting of these abstractions must also be permitted.

An example of this abstraction is a sequence of audio segments, the
first of which is a recording of the words "The number you have
dialed", followed by a Directory Number variable, followed by a
recording of the words "is no longer in service".

-     An abstraction that represents a set of audio segments and which
      is resolved to a single segment by a qualifier. Nesting of
      these abstractions must be permitted.

For example take a set of audio segments recorded in different
languages all of which express the semantic concept "The number you
have dialed is no longer in service". The set is resolved by a
language qualifier. If the qualifier is "French", the set resolves to
the French version of this announcement.

   In the case of a nested abstraction consisting of a set qualified by
   language at one level and and a set qualified by gender at another
   level, it would be possible to specify that an announcement be
   played in French and spoken by a female voice.

   e.   Provide two different methods of audio specification:

   -    Direct specification of the audio components to be played by
        specifying the sequence of segments in the command itself.

   -    Indirect specification of the audio components to be played by
        reference to a single identifier that resolves to a provisioned
        sequence of audio segments.

11.2.8.2. DTMF Collect Module

   The DTMF Collect Module must support all of the requirements in the
   Play Module in addition to the following requirements:

   a.   Be able to provide the following basic operation:

   -    request an AEG to play an announcement, which may optionally
        terminated by DTMF, and then collect DTMF

   b.   Be able to specify these event collection characteristics:

   -    The number of attempts to give the user to enter a valid DTMF
        pattern.

   c.   With respect to digit timers, allow the specification of:

   -    Time allowed to enter the first digit.

   -    Time allowed for user to enter each digit subsequent to the
        first digit.

   -    Time allowed for user to enter a digit once the maximum expected
        number of digits has been entered.

   d.   To be able to allow multiple prompt operations DTMF digit
        collection, voice recording (if supported), and/or speech
        recognition analysis (if supported) provide the following types
        of prompts:

   -    Initial Prompt

   -    Reprompt

-    Error prompt

-    Failure announcement

-    Success announcement.

e.   To allow digit pattern matching, allow the specification of:

-    maximum number of digits to collect.

-    minimum number of digits to collect.

-    a digit pattern using a regular expression.

f.   To allow digit buffer control, allow the specification of:

-    Ability to clear digit buffer prior to playing initial prompt
     (default is not to clear buffer).

-    Default clearing of buffer following playing of un-interruptible
     announcement segment.

-    Default clearing of buffer before playing a re-prompt in
     response to previous invalid input.

g.   Provide a method to specify DTMF interruptibility on a per audio
     segment basis.

h.   Allow the specification of definable key sequences for DTMF
     digit collection to:

-    Discard collected digits in progress, replay the prompt, and
     resume DTMF digit collection.

-    Discard collected digits in progress and resume DTMF digit
     collection.

-    Terminate the current operation and return the terminating key
     sequence to the MGC.

i.   Provide a way to ask the ARF MG to support the following
     definable keys for digit collection and recording. These keys
     would then be able to be acted upon by the ARF MG:

-    A key to terminate playing of an announcement in progress.

-    A set of one or more keys that can be accepted as the first
     digit to be collected.

-    A key that signals the end of user input. The key may or may
     not be returned to the MGC along with the input already
     collected.

-    Keys to stop playing the current announcement and resume playing
     at the beginning of the first segment of the announcement, last
     segment of the announcement, previous segment of the
     announcement, next segment of the announcement, or the current
     announcement segment.

11.2.8.3. Record Audio Module

   The Record Module must support all of the requirements in the Play
   Module as in addition to the following requirements:

   a.   Be able to provide the following basic operation:

   -    request an AEG to play an announcement and then record voice.

   b.   Be able to specify these event collection characteristics:

   -    The number of attempts to give the user to make a recording.

   c.   With respect to recording timers, allow the specification of:

   -    Time to wait for the user to initially speak.

   -    The amount of silence necessary following the last speech
        segment for the recording to be considered complete.

   -    The maximum allowable length of the recording (not including
        pre- and post- speech silence).

   d.   To be able to allow multiple prompt operations for DTMF digit
        collection (if supported), voice recording (if supported),
        speech recognition analysis (if supported) and/or speech
        verification/identification (if supported) and then to provide
        the following types of prompts:

   -    Initial Prompt

   -    Reprompt

   -    Error prompt

   -    Failure announcement

   -    Success announcement.

   e.   Allow the specification of definable key sequences for digit
        recording or speech recognition analysis (if supported) to:

   -    Discard recording in progress, replay the prompt, and resume
        recording.

   -    Discard recording in progress and resume recording.

   -    Terminate the current operation and return the terminating key
        sequence to the MGC.

   f.   Provide a way to ask the ARF MG to support the following
        definable keys for recording. These keys would then be able to
        be acted upon by the ARF MG:

   -    A key to terminate playing of an announcement in progress.

   -    A key that signals the end of user input. The key may or may
        not be returned to the MGC along with the input already
        collected.

   -    Keys to stop playing the current announcement and resume playing
        at the beginning of the first segment of the announcement, last
        segment of the announcement, previous segment of the
        announcement, next segment of the announcement, or the current
        announcement segment.

   g.   While audio prompts are usually provisioned in IVR/ARF MGs,
        support changing the provisioned prompts in a voice session
        rather than a data session. In particular, with respect to
        audio management:

   -    A method to replace provisioned audio with audio recorded during
        a call. The newly recorded audio must be accessible using the
        identifier of the audio it replaces.

   -    A method to revert from replaced audio to the original
        provisioned audio.

   -    A method to take audio recorded during a call and store it such
        that it is accessible to the current call only through its own
        newly created unique identifier.

   -    A method to take audio recorded during a call and store it such
        that it is accessible to any subsequent call through its own
        newly created identifier.

11.2.8.4. Speech Recognition Module

   The speech recognition module can be used for a number of speech
   recognition applications, such as:

   -    Limited Vocabulary Isolated Speech Recognition (e.g., "yes",
        "no", the number "four"),

   -    Limited Vocabulary Continuous Speech Feature Recognition (e.g.,
        the utterance "four hundred twenty-three dollars"),and/or

   -    Continuous Speech Recognition (e.g., unconstrained speech
        recognition tasks).

   The Speech Recognition Module must support all of the requirements in
   the Play Module as in addition to the following requirements:

   a.   Be able to provide the following basic operation: request an AEG
        to play an announcement and then perform speech recognition
        analysis.

   b.   Be able to specify these event collection characteristics:

   -    The number of attempts to give to perform speech recognition
        task.

   c.   With respect to speech recognition analysis timers, allow the
        specification of:

   -    Time to wait for the user to initially speak.

   -    The amount of silence necessary following the last speech
        segment for the speech recognition analysis segment to be
        considered complete.

   -    The maximum allowable length of the speech recognition analysis
        (not including pre- and post- speech silence).

   d.   To be able to allow multiple prompt operations for DTMF digit
        collection (if supported), voice recording (if supported),
        and/or speech recognition analysis and then to provide the
        following types of prompts:

   -    Initial Prompt

   -    Reprompt

   -    Error prompt

          -   Failure announcement

          -   Success announcement.

     e.   Allow the specification of definable key sequences for digit
          recording (if supported) or speech recognition analysis to:

          -   Discard in process analysis, replay the prompt, and resume
              analysis.

          -   Discard recording in progress and resume analysis.

          -   Terminate the current operation and return the terminating key
              sequence to the MGC.

     f.   Provide a way to ask the ARF MG to support the following
          definable keys for speech recognition analysis. These keys would
          then be able to be acted upon by the ARF MG:

          -   A key to terminate playing of an announcement in progress.

          -   A key that signals the end of user input. The key may or may
              not be returned to the MGC along with the input already
              collected.

          -   Keys to stop playing the current announcement and resume playing
              at the beginning of the first segment of the announcement, last
              segment of the announcement, previous segment of the
              announcement, next segment of the announcement, or the current
              announcement segment.

11.2.8.5. Speaker Verification/Identification Module

   The speech verification/identification module returns parameters that
   indicate either the likelihood of the speaker to be the person that
   they claim to be (verification task) or the likelihood of the speaker
   being one of the persons contained in a set of previously
   characterized speakers (identification task).

   The Speaker Verification/Identification Module must support all of
   the requirements in the Play Module in addition to the following
   requirements:

   a.   Be able to download parameters, such as speaker templates
        (verification task) or sets of potential speaker templates
        (identification task), either prior to the session or in mid-
        session.

   b.   Be able to download application specific software to the ARF
         either prior to the session or in mid-session.

   c.   Be able to return parameters indicating either the likelihood of
         the speaker to be the person that they claim to be (verification
         task) or the likelihood of the speaker being one of the persons
         contained in a set of previously characterized speakers
         (identification task).

   d.   Be able to provide the following basic operation: request an AEG
         to play an announcement and then perform speech
         verification/identification analysis.

   e.   Be able to specify these event collection characteristics: The
         number of attempts to give to perform speech
         verification/identification task.

   f.   With respect to speech verification/identification analysis
         timers, allow the specification of:

   -     Time to wait for the user to initially speak.

   -     The amount of silence necessary following the last speech
         segment for the speech verification/identification analysis
         segment to be considered complete.

   -     The maximum allowable length of the speech
         verification/identification analysis (not including pre- and
         post- speech silence).

   g.   To be able to allow multiple prompt operations for DTMF digit
         collection (if supported), voice recording, (if supported),
         speech recognition analysis (if supported) and/or speech
         verification/identification and provide the following types of
         prompts:

   -     Initial Prompt

   -     Reprompt

   -     Error prompt

   -     Failure announcement

   -     Success announcement.

   h.   Allow the specification of definable key sequences for digit
         recording (if supported) or speech recognition (if supported) in
         the speech verification/identification analysis to:

   -     Discard speech verification/identification in analysis, replay
         the prompt, and resume analysis.

   -     Discard speech verification/identification analysis in progress
         and resume analysis.

   -     Terminate the current operation and return the terminating key
         sequence to the MGC.

   i.   Provide a way to ask the ARF MG to support the following
         definable keys for speech verification/identification analysis.
         These keys would then be able to be acted upon by the ARF MG:

   -     A key to terminate playing of an announcement in progress.

   -     A key that signals the end of user input. The key may or may
         not be returned to the MGC along with the input already
         collected.

   -     Keys to stop playing the current announcement and resume speech
         verification/identification at the beginning of the first
         segment of the announcement, last segment of the announcement,
         previous segment of the announcement, next segment of the
         announcement, or the current announcement segment.

11.2.8.6. Auditory Feature Extraction/Recognition Module

   The auditory feature extraction/recognition module is engineered to
   continuously monitor the auditory stream for the appearance of
   particular auditory signals or speech utterances of interest and to
   report these events (and optionally a signal feature representation
   of these events) to network servers or MGCs.

   The Auditory Feature Extraction/Recognition Module must support the
   following requirements:

   a.   Be able to download application specific software to the ARF
         either prior to the session or in mid-session.

   b.   Be able to download parameters, such as a representation of the
         auditory feature to extract/recognize, for prior to the session
         or in mid-session.

   c. Be able to return parameters indicating the auditory event found
       or a representation of the feature found (i.e., auditory
       feature).

11.2.8.7. Audio Conferencing Module

   The protocol must support:

   a. a mechanism to create multi-point conferences of audio only and
       multimedia conferences in the MG.

   b. audio mixing; mixing multiple audio streams into a new composite
       audio stream

   c. audio switching; selection of incoming audio stream to be sent
       out to all conference participants.

11.2.9. Multipoint Control Units

   The protocol must support:

   a. a mechanism to create multi-point conferences of audio only and
       multimedia conferences in the MG.

   b. audio mixing; mixing multiple audio streams into a new composite
       audio stream

   c. audio switching; selection of incoming audio stream to be sent
       out to all conference participants.

   d. video switching; selection of video stream to be sent out to all
       conference participants

   e. lecture video mode; a video selection option where on video
       source is sent out to all conference users

   f. multi-point of T.120 data conferencing.

   g. The ability for the MG to function as an H.323 MP, and for the
       MGC to function as an H.323 MC, connected by this protocol
       (MEGACOP/H.248). It should be possible for audio, data, and
       video MG/MPs to be physically separate while being under the
       control of a single MGC/H.323 MC.

12. References

   [1]   Bradner, S., "Key words for use in RFCs to Indicate Requirement
         Levels", BCP 14, RFC 2119, March 1997.

   [2]   ITU-T Recommendation Q.2630.1, AAL type 2 Signalling Protocol
         (Capability Set 1), December 1999.

   [3]   ITU-T Recommendation H.341, Line Transmission of Non-Telephone
         Signals, May 1999.

   [4]   ATM Forum Technical Committee, af-vtoa-0083.001, Voice and
         Telephony Over ATM to the Desktop Specification, March 1999.

   [5]   ITU-T Recommendation H.323v3, Packet-based Multimedia
         Communications Systems (includes Annex C - H.323 on ATM),
         September 1999.

   [6]   ATM Forum Technical Committee, af-saa-0124.000, Gateway for
         H.323 Media Transport Over ATM, May 1999.

   [7]   ITU-T Recommendation T.140, Protocol for Multimedia Application
         Text Conversation, February 1998.

   [8]   ITU-T Recommendation V.18, Operational and Interworking
         Requirements for DCEs Operating in Text Telephone Mode, February
         1998.

   [9]   ITU-T Recommendation Q.931, Digital Subscriber Signalling System
         No. 1 (DSS 1) - ISDN User - Network Interface Layer 3
         Specification for Basic Call Control, May 1998.

14. Acknowledgements

15. Authors' Addresses

Nancy Greene
Nortel Networks
P.O. Box 3511 Stn C
Ottawa, ON, Canada K1Y 4H7

Phone: (514) 271-7221
EMail: ngreene@nortelnetworks.com

Michael A. Ramalho
Cisco Systems
1802 Rue de la Port
Wall Township, NJ

Phone: +1.732.449.5762
EMail: mramalho@cisco.com

Brian Rosen
Marconi
1000 FORE Drive, Warrendale, PA 15086

Phone: (724) 742-6826
EMail: brosen@eng.fore.com

16. Full Copyright Statement

Acknowledgement