A Data Model for Presence

Status of This Memo

Copyright Notice

Abstract

   This document defines the underlying presence data model used by
   Session Initiation Protocol (SIP) for Instant Messaging and Presence
   Leveraging Extensions (SIMPLE) presence agents. The data model
   provides guidance on how to map various communications systems into
   presence documents in a consistent fashion.

Table of Contents

1. Introduction

   Presence conveys the ability and willingness of a user to communicate
   across a set of devices. RFC 2778 [10] defines a model and
   terminology for describing systems that provide presence information.
   RFC 3863 [1] defines an XML [5] [6] [7] document format for
   representing presence information. In these specifications, presence
   information is modeled as a series of tuples, each of which contains
   a status, communications address, and other markup. However, neither
   specification gives guidance on exactly what a tuple is meant to
   model, or how to map real-world communications systems (and in

particular, those built around the Session Initiation Protocol (SIP)
[11]) into a presence document.

In particular, several important concepts are not clearly modeled or
well delineated by RFCs 2778 and 3863. These are the following:

Service: A communications service, such as instant messaging (IM) or
   telephony, is a system for interaction between users that provides
   certain modalities or content.

Device: A communications device is a physical component that a user
   interacts with in order to make or receive communications.
   Examples are a phone, PDA, or PC.

Person: A person is the end user, and for the purposes of presence,
   is characterized by states, such as "busy" or "sad", that impact
   their ability and willingness to communicate.

This specification defines these concepts more fully by means of a
presence data model, and concretely defines how to take real-world
systems and map them into presence documents using that model. This
data model is defined in terms of an extension to RFC 3863.

2. Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in RFC 2119 [9].

This document makes use of many additional terms beyond those defined
in RFC 2778 and RFC 3863. These new terms are as follows:

Device: A device models the physical environment in which services
   manifest themselves for users. Devices have characteristics that
   are useful in allowing a user to make a choice about which
   communications service to use.

Service: A service models a form of communication that can be used
   to interact with the user.

Person: A person models the human user and their states that are
   relevant to presence systems.

Occurrence: A single description of a particular service, a
   particular device, or a person. There may be multiple occurrences
   for a particular service or device, or multiple person occurrences
   in a Presence Information Data Format (PIDF) document, in cases
   where there is ambiguity that is best resolved by the watcher.

Presentity: A presentity combines devices, services, and person
   information for a complete picture of a user's presence status on
   the network.

Presentity URI: A URI that acts as a unique identifier for a
   presentity and provides a handle for obtaining presence
   information about that presentity.

Data Component: One of the device, service, or person parts of a
   presence document.

Status: Presence information about a service, person, or device that
   typically changes over time, in contrast to characteristics, which
   are generally static.

Characteristics: Presence information about a service, person, or
   device that is usually fixed over time, and descriptive in nature.
   Characteristics are useful in providing context that identifies
   the service or device as different from another service or device.

Attribute: A status or characteristic. It represents a single piece
   of presence information.

Presence Attribute: A synonym for attribute.

Composition: The act of combining a set of presence and event data
   about a presentity into a coherent picture of the state of that
   presentity.

3. The Model

```
+---------------------------------------------------------------+
|                                                               |
|                    +---------------+                          |
|                    +---------------+|                         |
|                    |               ||                         |
|                    |    Person     ||                         |
|                  /|               ||\                         |
|                 / +---------------+| \                        |
|                /  |               |+  \                       |
|               /   |               |    \                      |
|              /    |               |     \                     |
|             V     V               V      V                    |
|   +---------------+  +---------------+  +---------------+      |
|   +---------------+| +---------------+| +---------------+|     |
|   |               || |               || |               ||    |
|   |    Service    || |    Service    || |    Service    ||    |
|   |               || |               || |               ||    |
|   |               |+ |               |+ |               |+    |
|   +---------------+  +---------------+  +---------------+      |
|           \              /    \                               |
|            \            /      \                              |
|             V          V        V                            |
|      +---------------+  +---------------+                     |
|      +---------------+| +---------------+|                    |
|      |               || |               ||                   |
|      |    Device     || |    Device     ||                   |
|      |               || |               ||                   |
|      |               |+ |               |+                   |
|      +---------------+  +---------------+                     |
|                                                               |
| Presentity (URI)                                              |
+---------------------------------------------------------------+
```
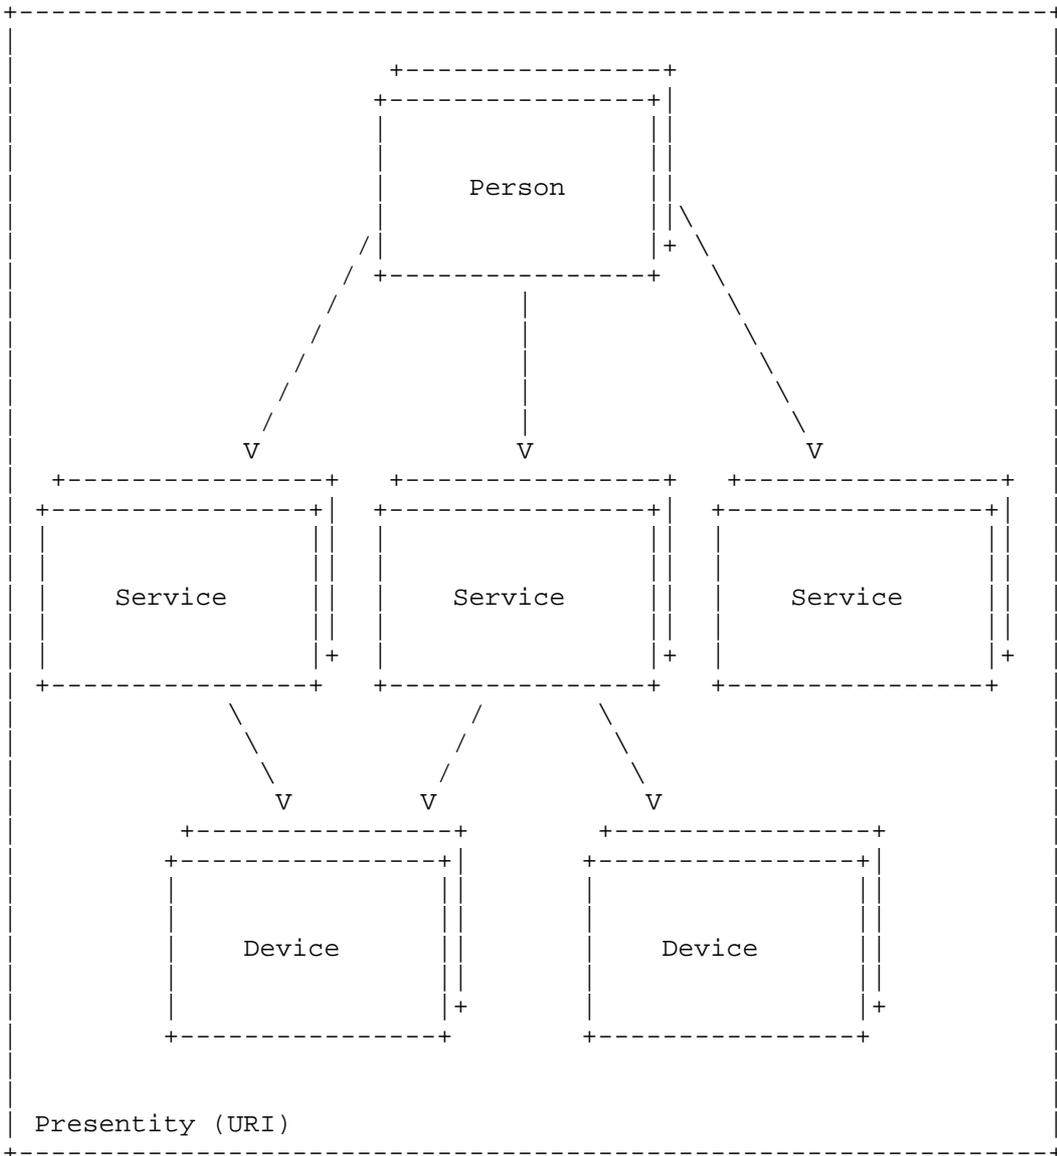
                             Figure 1

The data model for presence is shown in Figure 1. The model seeks to
describe the presentity, identified by a presentity URI. There are
three components in the model: the person, the service, and the
device. These three data components contain information (called
attributes) that provide a description of some aspect of the service,
person, or device. It is central to this model that each attribute
is affiliated with the service, person, or device because they
describe that service, presentity, or device. This is in contrast to
a model whereby the attributes are associated with the service,
presentity, or device because they were reported by that service,
presentity, or device. As an example, if a cell phone reports that a
user is in a meeting, this would be done by including an attribute as
part of the person information, indicating a status of
"in-a-meeting". The presence information may also include
information on the cell phone as a device. However, even though it
is the device that is reporting that the user is in a meeting, "in a
meeting" is a fact that describes the human user, not their physical
device. Consequently, this attribute is placed in the person
component of the document.

3.1. Presentity URI

The identifier for the presentity is a URI. For each unique
presentity in the network, there is one or more presentity URIs. A
presentity may have multiple URIs because they are identified by both
a URI from the Presence (pres) scheme [12] and a protocol-specific
URI, such as a SIP URI [11] or an Extensible Messaging and Presence
Protocol Internationalized Resource Identifier (XMPP IRI) [13]. Or,
it can be because a user has several aliases in a domain, all of
which are equivalent identifiers for the presentity.

When a document is constructed, the presentity URI is ideally set to
the identifier used to request the document in the first place. For
example, if a document was requested through a SIP SUBSCRIBE request,
the presentity URI would match the Request URI of the SUBSCRIBE
request. This follows the principle of least surprise, since the
entity requesting the document may not be aware of the other
identifiers for the presentity.

Irrespective of the scheme from which the URI is taken, the
presentity URI is independent of any of the services or devices that
the presentity possesses. However, the URI is not just a name - it
represents a resource that can be subscribed to, in order to find out
the status of the user. When the URI is a SIP URI, it will often be
the Address of Record for the user, to which SIP calls can be
directed. This equivalence is not mandated by this specification,
but is a recommended configuration for easing the burden of
remembering and storing identifiers for users.

3.2. Person

   The person data component models information about the user whom the
   presence data is trying to describe. This information consists of
   characteristics of the user, and their status.

   Characteristics of a person are the static information about a user
   that does not change under normal circumstances. Such information
   might include physical characteristics, such as age and height.
   Another example of a person characteristic is an alias. An alias is
   a URI that identities the same user, but with a different presentity
   URI. For example, a presentity "sip:bob@example.com" might have a
   presence document with a person component that indicates an alias of
   "sip:robert@example.com" and "sip:r.smith@example.com".

   Status information about a presentity represents the dynamic
   information about a user. This typically consists of things the
   user is doing, places the *user* is at, feelings the *user* has,
   and so on. Examples of typical person status are "in a meeting", "on
   the phone", "out to lunch", "happy", and "writing Internet Drafts".
   The line between static status information and dynamic status
   information is fuzzy, and it is not important that a line be drawn.
   The model does not differentiate in a syntactically or semantically
   meaningful way between these two types of attributes.

   In the model, there can be only one person component per presentity.
   In other words, the person component models a single human being, and
   includes characteristics and statuses that are related to the
   communication states for a single human being. Of course, the system
   has no way to verify that the human described by the person component
   is actually a single human being, as opposed to a group of users, or
   even a dog for that matter. As the saying goes, "on the Internet, no
   one knows you are a dog", and the same is true here. The person
   component is a facade for a single person; anything that can be made
   to look like a single person can be modeled with that facade.

   As an example, consider the task of using a presence document to
   describe a customer support help desk. The person component can be
   considered to be "busy" if none of the support staff are available,
   and "at lunch" if the help desk department has a group lunch
   together. The watcher that receives the document will consider the
   help desk to be a single person; nothing in the document (except
   perhaps the note element, should its value be "help desk" or
   something similar) conveys information that would indicate that the
   person in question is actually a help desk.

However, there can be multiple occurrences of the person component.
This happens in cases where the state of the person component is
ambiguous, as discussed in Section 3.5.

3.3. Service

Each presentity has access to a number of services. Each of these
represents a point of reachability for communications that can be
used to interact with the user. Examples of services are telephony
(that is, traditional circuit-based telephone service), push-to-talk,
instant messaging, Short Message Service (SMS), and Multimedia
Message Service (MMS).

It is difficult to give a precise definition for service. One
reasonable approach is to model each software or hardware agent in
the system as a service. If a user starts a softphone application on
their PC, then that represents a service. If a user has a videophone
device, then that represents another service. This is effectively a
physical view of services. This definition, however, starts to fall
apart when a service is spread across multiple software agents or
devices. For example, a SIP URI representing an address-of-record
can be routed to a softphone or a videophone, or both. In that case,
one might attempt instead to define a service based on its address on
the network. This definition also falls apart when modeling devices
or applications that receive calls and dispatch them to different
"helpers" based on potentially complex logic. For example, a
cellular telephone might house multiple SIP applications, each of
which can "register" different handlers based on the method or even
body type of the request. Each of those applications or handlers can
rightfully be considered a service, but it doesn't have an address on
the network distinct from the others.

Because of this inherent difficulty in precisely defining a service,
the data model doesn't try to constrain what can be considered a
service. Rather, anything can be considered a service so long as it
exhibits a set of key properties defined by this model. In
particular, each service is associated with characteristics that
identify the nature and capabilities of that service, with reach
information that indicates how to connect to the service, with status
information representing the state of that service, and relative
information that describes the ways in which that service relates to
others associated with the presentity.

As a consequence, in this model, services are not explicitly
enumerated. There is no central registry where one finds identifiers
for each service. Consequently, each service does not have a single
"service" attribute with values such as "ptt" or "telephony". That
doesn't mean that these consolidated monikers aren't useful; indeed,

they represent an essential summary of what the service is. Such
summarization is useful in creating icons that allow a user to choose
one service over another. A watcher is free to create such
summarization information from any of the information associated with
a service. The reach information often provides valuable information
for creating such a summarization. Oftentimes, the scheme of the URI
is synonymous with the view of what a service is. An "sms" URI [14]
clearly indicates SMS, for example. For some URIs, there may be many
services available, for example, SIP or tel [15], in which case the
scheme is less meaningful as a way of creating a summary. The reach
information could also indicate that certain application software has
to be invoked (such as a videogame), in which case that aspect of the
reach information would be useful for generating an iconic
representation of the game.

3.3.1. Characteristics

Each service is adorned with characteristics that describe the nature
and capabilities of the service that will be experienced when a
watcher invokes that URI. The nature of a service is a set of
properties that are relatively static across communication sessions
established to that service. The nature of a service tends to be
descriptive. Examples of the nature of a service are that it
represents an interactive voice response or voicemail server, that it
is an automaton, or that it is a telephony service used for the
purposes of work. Capabilities, on the other hand, represent
properties that might be exhibited, and whether they are exhibited
depends on negotiation and other dynamic functions that take place
during session establishment. Examples of such capabilities are the
type of media that might be used, the directionality of
communications that are permitted, the SIP extensions supported, and
so on. Capabilities can be very complex; for example, RFC 2533 [16]
describes a model for representing capabilities through N-ary boolean
functions. It is difficult to differentiate a capability with one
modality (e.g., this service only does voice) from a characteristic
that represents the nature of a service. However, it is not
important to do so.

Characteristics are important when multiple services are indicated.
That is because the purpose of listing multiple services in a
presence document is to give the watcher a *choice*. That is, the
presentity is explicitly offering the watcher an opportunity to
contact them using a multiplicity of different services. To help the
watcher make a decision, the presence document includes
characteristics of each service that help differentiate the services
from each other and give the watcher the context in which to make a
choice.

Because their purpose is primarily to facilitate choice, capabilities
do not impose a requirement on the way in which a user reaches that
service. For example, if a presence document includes two services,
and one supports audio only while the other supports only video, this
does not mean that, when contacting the first service, a user has to
offer only an audio stream, or when contacting the second service, a
user has to offer only a video stream. A user can use local policy
at its discretion in determining what capabilities or communications
modalities are offered when they choose to connect with a service.
It is not necessary for a watcher to add SIP caller preferences [2]
to request routing of the request to a service with the
characteristics described in the presence document.

If, in order to reach a service, the user agent must generate a
request that exhibits a particular capability or contains a specific
header, then this is indicated separately in the reach information,
described below.

One important characteristic of each service is the list of devices
on which that service executes. Each device is identified uniquely
by a device ID. As such, the service characteristics can include a
list of device IDs. A presence document might also contain
information on each device, but this is a separate part of the
document. Indeed, the information on each device might not even be
present in the document. In that case, the device IDs listed for
each service are nothing more than correlation identifiers, useful
for determining when two services run on the same device. The
benefit of this model is that information on the devices can be
filtered out of a presence document, yet the service information,
which includes the device IDs, remains useful and meaningful.

It is perfectly valid for a presence document to contain just a
single service. This is permitted even if the presentity actually
has multiple services at their disposal. The lack of multiple
services in the document merely means that the presentity is not
offering a choice to the watcher. In such a case, the service
characteristics are less important, but may be helpful in allowing a
watcher to decide if they wish to communicate at all.

3.3.2. Reach Information

The reach information for a service provides the instructions for the
recipient of a document on how to correctly contact that service.

When a service is accessible over a communications network, reach
information includes a URI that can be "hit" to access the service.
This URI is called the service URI. However, some services are not

accessible over a communications network (such as in-person
communications or a written letter), and as such, may not utilize a
URI.

Even for services reachable over a communications network, the URI
alone may not be sufficient. For example, two applications may be
running within a cellular telephone, both of which are reachable
through the user's SIP Address of Record. However, one application
is launched when the INVITE request contains a body of a particular
type, and the other is launched for other body types. As another
example, a service may provide complex application logic that
operates correctly only when contacted from matching application
software. In such a case, even though the communications between
instances utilizes a standard protocol (such as SIP), the user
experience will not be correct unless the applications are matched.

When the URI is not sufficient, additional attributes of the service
can be present that define the instructions on how the service is to
be reached. These attributes must be understood for the service to
be utilized. If a watcher receives a presence document containing
reach information it does not understand, it should discard the
service information.

The reach information is an important part of the service. When the
watcher makes a decision about which service of the presentity they
wish to access, the watcher utilizes the reach information for that
service. For this reason, each service has to have a unique set of
reach information. If this was not the case, the user would have no
way to choose between the services. This means that the reach
information represents a unique identifier for the service. However,
a presence document can contain multiple occurrences of a particular
service, each of which contains the same reach information, but
differs in its occurrence identifier. Multiple occurrences of a
service exist in a document when the state of the service is
ambiguous, as discussed in Section 3.5.

Because the reach information serves as an identifier for a service,
it also serves as a way to figure out whether a communications
capability should be represented as one service or more. Something
cannot be a service unless there is a way to reach it separately from
another service. As an example, consider a softphone application
that is capable of audio and video. It is not possible to describe
this softphone as two services - one capable of just audio, and one
capable of just video. That's because there is no way to reach the
video-only service; for example, sending a SIP INVITE with just a
video stream doesn't suffice, since one can always add the audio
stream later and it will work. Video and audio, in this case,
represent capabilities for a single service.

The reach information represents a weak form of contract; the
presentity tells the watcher that, if the watcher utilizes the reach
information included in the presence document, the watcher might be
connected to a service described by the characteristics included in
the presence document. It is important to stress that this is not a
guarantee in any way. It cannot be a guarantee for two reasons.
First, the service in the document might actually be modelling a
number of actual services used by the user, and it may not be
possible to connect the watcher to a service with all of the
characteristics described in the presence document. Second, the
preferences of the presentity always take precedence. The caller
might ask to be connected to the video service, but it is permissible
to connect them to a different service if that is the wish of the
presentity.

This loose contract also provides some guidance on the type of URI
that is most ideally suited for the service URI. A URN [3] can be
used as the service URI. However, since a URN could be resolved to
potentially any number of different URIs, the characteristics,
status, and relative information need to be sensible for all of the
URIs that can be resolved from the URN. As the URN becomes
increasingly "vague" in terms of the service it identifies, the
number of presence attributes that can be included decreases
correspondingly.

The tel URI [11] shares similar properties with a URN, and the same
considerations apply. If, for example, the telephone number exists
in ENUM [18] and multiple ENUM services are defined, including voice
and messaging, it is likely that very little characteristic
information can be included in that service. If, however, a tel URI
has only a single ENUM service defined, and it refers to a telephone
service on the Public Switched Telephone Network (PSTN), more can be
said about its characteristics, status, and relative priority.

It is important to point out that there can be a many-to-one mapping
of reach information to a service. That is, a particular service can
potentially be reachable through an infinite number of reach
information sets. This is true even if the reach information is just
the service URI; it is permissible for multiple service URIs to reach
the same service. Within any particular document, for a particular
service, there will be a single service URI. However, it is allowed
and even valuable to provide different service URIs to different
watchers, or to change the service URIs provided to a particular
watcher over time. Doing so affords many benefits, in fact. It can
allow the recipient of a communications attempt to determine the
context for that attempt - that the attempt was made as a result of

trying to reach a particular service in a particular presence
document. This can be used as a technique for preventing
communications spam, for example [19].

It is also possible for a presence document to contain a service that
has no reach information at all. In such a case, the presentity is
indicating that the service exists, but is electing not to offer the
watcher the opportunity to connect to it. One such example would be
to let a watcher know that a user has a telephony service, and that
they are busy, but in order to avoid receipt of a call, no reach
information is provided.

In an ideal system, the URI alone would represent sufficient reach
information for each service. A URI is supposed to provide
sufficient context for reaching the resource associated with the URI,
and thus in theory there is no need for additional context. However,
sometimes, additional information is needed. Since the reach
information has to be understood in order for the service to be
utilized, reach information beyond the URI should be defined and used
sparingly. Extensions to PIDF that define attributes that are reach
information should clearly call those attributes out as such.

3.3.3. Relative Information

Each service is also associated with a priority, which represents the
preference that the user has for usage of one service over another.
This does not mean that, when a watcher wishes to communicate with
the presentity, that they should always use the service with the
highest priority. If that were the case, there would be no point in
including multiple services in the presence document. Rather, the
priority says, "If you, the watcher, cannot decide which of these to
use, or if it is not important to you, this is the order in which I
would like you to contact me. However, I am giving you a choice."
The priorities are relative to each other, and have no meaning as
absolute numbers. If there are two services, and they have
priorities of 1 and .5, respectively, this is identical to giving
them priorities of .2 and .1, respectively.

3.3.4. Status

Each service also has a status. Status represents generally dynamic
information about the availability of communications using that
service. This is in contrast to characteristics, which describe
fairly static properties of the various services. The simplest form
of status is the basic status, which is a binary indicator of
availability for communications using that service. It can have
values of either "closed" or "open". "Closed" means that
communication to the service will, in all likelihood, fail, will not

reach the intended party, or will not result in communications as
described by the characteristics of the service. As an example, if a
call is forwarded to voicemail if the user is busy or unavailable,
the service is marked as "closed". Similarly, a presentity may
include a hotel phone number as a service URI. After checkout, the
phone number will still ring, but reach the chambermaid or the next
guest. Thus, it would be declared "closed" by that presentity. As
another example, if a user has a SIP URI as their service URI that
points to a SIP softphone application, and the PC shuts down, calls
to that SIP URI will return a 480 response code. This service would
also be declared "closed". "Open" implies the opposite - that
communications to this service will likely succeed and reach the
desired target.

It is also possible to have status information that is dependent on
the characteristics of the communications session that eventually
gets set up. For example, a status attribute can be defined that
indicates that a softphone service is available if instant messaging
is used, but unavailable if audio is used.

Other status information might indicate more details on why the
service is available or unavailable. For example, a telephony
service might have additional status to indicate that the user is on
the phone, or that the user is handling 3 calls for that service.

Services inherently have a lot of dynamic state associated with them.
For example, consider a wireless telephony service (i.e., a cell
phone). There are many dynamic statuses of this service - whether or
not the phone is registered, whether or not it is roaming, which
provider it has roamed into, its signal strength, how many calls it
has, what the state of those calls are, how long the user has been in
a call, and so on. As another example, consider an IM service. The
statuses in this service include whether the user is registered, how
long they have been registered, whether they have an IM conversation
in progress, how many IM conversations are in progress, whether the
user is typing, to whom they are typing, and so on.

However, not all of this dynamic state is appropriate to include
within a service data component of a presence document. Information
is included only when it has a bearing on helping the watcher decide
whether to initiate communications with that service, or helping the
watcher decide when to initiate it, if not now. As an example,
whether a cell phone has strong signal strength or just good signal
strength does not pass the litmus test. Knowing this is not likely
to have an impact on a decision to use this service.

3.4. Device

   Devices model the physical operating environment in which services
   execute. Examples of devices include cell phones, PCs, laptops,
   PDAs, consumer telephones, enterprise PBX extensions, and operator
   dispatch consoles.

   The mapping of services to devices are many to many. A single
   service can execute in multiple devices. Consider a SIP telephony
   service. Two SIP phones can register against a single Address of
   Record for this service. As a result, the SIP service is associated
   with two devices. Similarly, a single device can support a
   multiplicity of services. A cell phone can support a SIP telephony
   service, an SMS service, and an MMS service. Similarly, a PC can
   support a SIP telephony service and a SIP videophone service.

   Furthermore, a single device can support no services. In such a
   case, the device has no useful presence information by itself.
   However, when composed with other documents that describe this same
   device in relation to a service, a richer presence document can be
   created. For example, consider a Radio Frequency ID (RFID) tag as a
   device. This device does not execute any services. However, as a
   device, it has properties, such as location, and it may have network
   connectivity with which it can report its status and characteristics.
   If a video telephone were to report that it was running a video
   service, and one of its properties was that it was tagged with that
   RFID, a compositor could combine the two documents together, and use
   the location of the RFID to say something about the location of the
   video telephony device.

   Devices are identified with a device ID. A device ID is a URI that
   is a globally and temporally unique identifier for the device. In
   particular, a device ID is a URN. The URN has to be unique across
   all other devices for a particular presentity. However, it is also
   highly desirable that it be persistent across time, globally unique,
   and computable in a fashion so that different systems are likely to
   refer to the device using the same ID. With these properties,
   differing sources of presence information based on device status can
   be combined. The last of these three properties - readily computable
   - is particularly useful. It allows for a compositor to combine
   disparate sources of information about a device, all linked by a
   common device ID that each source has independently used to identify
   the device in question.

   Unfortunately, due to the variety of different devices in existence,
   it is difficult for a single URN scheme to be used that will have
   these properties. It is anticipated that multiple schemes will be
   defined, with different ones appropriate for different types of

devices. For cellular telephones, the Electronic Serial Number
(ESN), for example, is a good identifier. For IP devices, the MAC
address is another good one. The MAC address has the property of
being readily computable, but lacks persistence across time (it would
change if the interface card on a device were to change). In any
case, neither of these are associated with URN schemes at this time.
In the interim, the Universally Unique IDentifier (UUID) URN [20] can
be used. For devices with a MAC address, version 1 UUIDs are
RECOMMENDED, as they result in a time-based identifier that makes use
of the MAC address. For devices without a MAC, a version 4 UUID is
RECOMMENDED. This is a purely random identifier, providing
uniqueness. The UUID for a device would typically be chosen at the
time of fabrication in the device, and then persisted in the device
within flash or some other kind of non-volatile storage. The UUID
URN has the properties of being globally and temporally unique, but
because of its random component, it is not at all readily computable,
and therefore useless as a correlation ID with other presence sources
on a network. It is anticipated that future specifications will be
developed that provide additional, superior device IDs.

Though each device is identified by a unique device ID, there can be
multiple occurrences of a particular device represented in a
document. Each one will share the same device ID, but differ in its
occurrence identifier. Multiple occurrences of a device exist in a
document when the state of the device is ambiguous, as discussed in
Section 3.5.

Though this document does not mandate a particular implementation
approach, the device ID is most useful when all of the services on
the device have a way to obtain the device ID and get the same value
for it. This would argue for its placement as an operating system
feature. Operating system developers interested in implementing this
specification are encouraged to provide APIs that allow applications
to obtain the device ID. Absent such APIs, applications that report
presence information about their devices will have to generate their
own device IDs. This leads to the possibility that the applications
may choose different device IDs, using different algorithms or data.
In the worst case, these may mean that two services that run on the
same device, do not appear to.

Like services and person data components, device data components have
generally static characteristics and generally dynamic status.
Characteristics of a device include its physical dimensions and
capabilities – the size of its display, the speed of its CPU, and the
amount of memory. Status information includes dynamic information
about the device. This includes whether the device is powered on or
off, the amount of battery power that remains in the device, the
geographic location of the device, and so on.

The characteristics and status information reported about a device
are for the purposes of choice - to allow the user to choose the
service based on knowledge of what the device is. The device
characteristics and status cannot, in any reliable way, be used to
extract information about the nature of the service that will be
received on the device. For example, if the device characteristics
include the speed of the CPU, and the speed is sufficient to support
high-quality video compression, this cannot be interpreted to mean
that video quality would be good for a video service on that device.
Other constraints on the system may reduce the amount of CPU
available to that service. If there is a desire to indicate that
higher-quality video is available on a device, that should be done by
including service characteristics that say just that. The speed of
the CPU might be useful in helping the watcher differentiate between
a device that is a PC and one that is a cell phone, in the case where
the watcher wishes to call the user's cell phone.

Similarly, if there is dynamic device status (such as whether the
device is on or off), and this state impacts the state of the
service, this is represented by adjusting the state of the service.
Unless a consumer of a presence document has a priori knowledge
indicating otherwise (note that presence agents often do), the state
of a device has no bearing on the state of the service.

Just like services, there is no enumeration of device types - PCs,
PDAs, cell phones, etc. Rather, the device is defined by its
characteristics, from which a watcher can extrapolate whether the
device is a PDA, cell phone, or what have you.

It is important to point out that the device is a *model* of the
underlying physical systems in which services execute. There is
nothing that says that this model cannot be used to talk about
systems where services run in virtualized systems, rather than real
ones. For example, if a PC is executing a virtual machine and
running services within that virtual machine, it is perfectly
acceptable to use this model to talk about that PC as being composed
of two separate devices.

3.5. Modeling Ambiguity

Ambiguity is a reality of a presence system, and it is explicitly
modeled by this specification. Ambiguity exists when there are
multiple pieces of information about a person, a particular device,
or a particular service. This ambiguity naturally arises when
multiple elements publish information about the person, a particular
service, or a particular device. In some cases, a compositor can
resolve the ambiguity in an automated way, and combine the data about
the person, device, or service into a single coherent description.

In other cases, it cannot, perhaps because the compositor lacks the
ability to do so.

However, in many cases, the resolution of this ambiguity is best left
to the watcher that consumes the document. This consumer could be an
application with more information than the compositor, and thus be
able to do a better job of resolving the ambiguity. Or, it may be
presented to the human user, and the human can often resolve the
ambiguity. Unsurprisingly, a human can often do this far better than
an automaton can.

To model ambiguity, the model allows each service, each device, or
the person component to contain multiple occurrences. Each
occurrence has a unique identifier, called the occurrence identifier.
This identifier is unique across all other occurrence identifiers for
any service, device, or person. That is, its uniqueness is scoped
within all of the services, devices, and person elements for a
particular presentity. The identifier ideally persists over time,
since it serves as a valuable handle for setting composition and
authorization policies. Even if there is a single occurrence for a
particular device, service, or person, the occurrence has an
occurrence identifier.

The occurrence identifier is not to be confused with the instance ID
defined in the SIP Outbound specification [27]. A user agent
instance is best modeled as a service, and indeed, a Globally
Routable User Agent URI (GRUU) [22], which is derived from the
instance ID, represents a reasonable choice for a service URI.
However, if the status of such a UA instance could not be determined
unambiguously, a presence document could include two or more
occurrences of the service modeling that UA instance. In such a
case, each occurrence has a unique occurrence ID, but they share the
same service URI, and consequently, the same instance ID.

When multiple occurrences exist in a document, it is important that
some of the attributes of the device, service, or person help the
recipient resolve the ambiguity. For humans, the note field and
timestamp serve as valuable tools. For an automaton, nearly any
attribute of the device, service, or person can be used to resolve
the ambiguity. The timestamp in particular is very useful for both
humans and automatons. As described in RFC 3863 [1], the timestamp
provides the time of most recent change for the tuple. This
specification defines the timestamp for person and device components
as well, with the same meaning. Absent other information, the
person, device, or service that most recently changed can be used as
the more reliable source of data. However, such a resolution
algorithm is not normatively required in any way.

3.6. The Meaning of Nothing

   It is clear that the existence of a presence attribute in a document
   tells something to a watcher about the value of that presence
   attribute. However, what does the absence of a presence attribute
   say? This data model follows the lead of RFC 3840 [17], which is
   used to define capabilities for SIP user agents. In that
   specification, if a capability declaration omits a particular feature
   tag, it means that the agent is making no definitive statement either
   way about whether this feature tag is supported. The same is true
   here - the absence of a presence attribute from a document means that
   a watcher cannot make any definitive statement about the value for
   that presence attribute. It may be absent because it is being
   withheld from the watcher, or it may be absent because that attribute
   is not supported by the presentity's software. Neither conclusion
   can be drawn.

   Because the absence of a presence attribute conveys no information
   whatsoever, presence documents achieve their maximum value when they
   have as many presence attributes as possible. As such, it is
   RECOMMENDED that a presence document contain as many presence
   attributes as the presentity is willing to and able to provide to a
   watcher.

3.7. Status vs. Characteristics

   The data model tries to separate status information from
   characteristics, generally by defining status as a relatively dynamic
   state about a person, device, or service, whereas a characteristic is
   relatively static. However, this distinction is often artificial.
   Almost any characteristic can change over time, and sometimes
   characteristics can change relatively quickly. As a result, the
   distinction between status and characteristics is merely a conceptual
   one to facilitate understanding about the different types of presence
   information. Nothing in a presence document indicates whether an
   element is a characteristic vs. a status, and when a presence
   attribute is defined, there is no need for it to be declared one or
   the other. Presence documents allow any presence attribute, whether
   it can be thought of as a characteristic or a status, to change at
   any time.

   Unfortunately, the original PIDF specification did have a separate
   part of a tuple for describing status, and the basic status was
   defined to exist within that part of the tuple. This specification
   does not change PIDF; however, all future presence attributes MUST be
   defined as children of the <tuple> and not the <status> element.
   Furthermore, the schemas defined here do not contain a <status>
   element for either the <person> or <device> elements.

3.8. Presence Document Properties

   The overall presence document has several important properties that
   are essential to this model.

   First, a presence document has a concrete meaning independent of how
   it is transported or where it is found. The semantics of a document
   are the same regardless of whether a document is published by a
   presence user agent to its compositor, or whether it is distributed
   from a presence agent to watchers. There are no required or implied
   behaviors for a recipient of a document. Rather, there are well-
   defined semantics for the document itself, and a recipient of a
   document can take whatever actions it chooses based on those
   semantics.

   A corollary of this property is that presence systems are infinitely
   composeable. A presence user agent can publish a document to its
   presence server. That presence server can compose it with other
   documents, and place the result in a notification to a watcher. That
   watcher can actually be another presence agent, combining that
   document with others it has received, and placing those results in
   yet another notify.

   Yet another corollary of this property is that implied behaviors in
   reaction to the document cannot ever be assumed. For example, just
   because a service indicates that it supports audio does not mean that
   a watcher will offer audio in a communications attempt to that
   service. If doing so is necessary to reach the service, this must be
   indicated explicitly through reach information.

   It is also important to understand that the role of the presence
   document is to help a user make a choice amongst a set of services,
   and furthermore, to know ahead of time with as much certainty as
   possible whether a communications attempt will succeed or fail.
   Success is a combination of many factors: Does the watcher understand
   the service URI? Can it act on all of the reach information? Does
   it support a subset of the capabilities associated with the service?
   Does the person information indicate that the user is likely to
   answer? All of these checks should ideally be made before attempting
   communication.

   Because the presence document serves to help a user to choose and
   establish communications, the presentity URI - as the index to that
   document - represents a form of "one-number" communications.
   Starting from this URI, all of the communications modalities and
   their URIs for a user can be discovered, and then used to invoke a
   particular communications service. Rather than having to give out a
   separate phone number, email address, IM address, Voice over Internet

Protocol (VoIP) address, and so on, the presentity URI can be
provided, and all of the others can be learned from there.

4. Motivation for the Model

Presence is defined in [21] as the ability, willingness, or desire to
communicate across a set of devices. The core of this definition is
the conveyance of information about the ability, willingness, or
desire for communications. Thus, the presence data model needs to be
tailored around conveying information that achieves this goal.

The person data component is targeted at conveying willingness and
desire for communications. It is used to represent information about
the users themselves that affects willingness and desire to
communicate. Whether I am in a meeting, whether I am on the phone -
each of these says something about my willingness to communicate, and
thus makes sense for inclusion in a presence document.

The service component of the data model aims to convey information on
the ability to communicate. The ability to communicate is defined by
the services by which a user is reachable. Thus, including them is
essential.

How do devices fit in? For many users, devices represent the ability
to communicate, not services. Frequently, users make statements
like, "Call me on my cell phone" or "I'm at my desk". These are
statements for preference for communications using a specific device,
as opposed to a service. Thus, it is our expectation that users will
want to represent devices as part of the presence data.

Furthermore, the concept of device adds the ability to correlate
services together. The device models the underlying platform that
supports all of the services on the phone. Its state therefore
impacts all services. For example, if a presence server can
determine that a cell phone is off, this says something about the
services that run on that device: they are all not available. Thus,
if services include indicators about the devices on which they run,
device state can be obtained and thus used to compute the state of
the services on the device.

The data model tries hard to separate device, service, and person as
different concepts. Part of this differentiation is that many
attributes will be applicable to some of these, but not others. For
example, geographic location is a meaningful attribute of the person
(the user has a location) and of a device (the device has a
location), but not of a service (services don't inherently have
locations). Based on this, geographic location information should
only appear as part of device or person, never service. Furthermore,

it is possible and meaningful for location information to be conveyed
for both device and person, and for these locations to be different.
The fact that the presence system might try to determine the location
of the person by extrapolation from the location of one of the
devices is irrelevant from a data modeling perspective. Person
location and device location are not the same thing.

[25] defines the <geopriv> XML element for conveying location
information, and indicates that it is carried as a child of the
<tuple> element in a PIDF document. [25] was developed prior to this
specification, and unfortunately, its recommendation to include
location objects underneath <tuple> runs contrary to the
recommendations here. As such, implementations based on this
specification SHOULD include <geopriv> location objects as part of
person and/or device components of the document, but SHOULD be
prepared to receive presence documents with that object as a child to
<tuple>. A <geopriv> location object would be included in a person
component when the document means to convey the location of the user,
and within a device component when it means to convey the location of
the device.

5. Encoding

   Information represented according to the data model described above
   needs to be mapped into an on-the-wire format for transport and
   storage. The Presence Information Data Format [1] is used for
   representation of presence data.

   The <presence> element contains the presence information for the
   presentity. The "entity" attribute of this element contains the
   presentity URI.

   The existing <tuple> element in the PIDF document is used to
   represent the service. This is consistent with the original intent
   of RFC 2778 and RFC 3863, and achieves backward compatibility with
   implementations developed before the model described here was
   complete. The <contact> element in the <tuple> element is used to
   encode the service URI. New presence attributes, whether they
   represent dynamic status or static characteristics, appear directly
   as children of <tuple>. However, attributes defined prior to
   publication of this specification that were defined as children of
   <status> (such as <basic>) remain as children of <status>, for
   purposes of backward compatibility. Consequently, a presence
   attribute describing a service could appear as either a child of
   <status> or directly as a child of <tuple>, but never both.

The "id" attribute of the <tuple> element conveys the service
occurrence. Each <tuple> element with the same <contact> URI
represents a different occurrence of a particular service.

This specification introduces the <person> element, which can appear
as a child to <presence>. There can be zero or more occurrences of
this element per document. Each one has a mandatory "id" attribute,
which contains the occurrence identifier for the person. Each
<person> element contains any number of elements that indicate status
and characteristic information. This is followed by zero or more
optional <note> elements and an optional <timestamp>. Multiple
<note> elements would appear to convey the same note in multiple
languages.

RFC 3863 defines a <note> element, zero or more of which can be
present as a child to <presence>. As it relates to the model defined
here, these note elements, if present in a document, apply to all
person occurrences that do not have any of their own <note> elements.
In other words, if a <person> element has one or more <note>
elements, those are the <note> elements for that <person> element.
If a <person> element does not have any of its own <note> elements,
the <note> elements that are the direct children of <presence> are
the <note> elements for that <person>. If there are no <note>
elements underneath the <person> element, and there are no <note>
elements that are a direct child of <presence>, then that <person>
element has no <note> elements.

This specification also introduces the <device> element, which can
appear as a child to <presence>. There can be zero or more
occurrences of this element per document. The <device> element can
appear either before or after the <person> element; there are no
constraints on order. Each <device> element has a mandatory "id"
attribute, which contains the occurrence identifier for the device.
Like <person>, <device> contains any number of elements that indicate
status and characteristic information. This is followed by
<deviceID>, which contains the URN for the device ID for this device.
This is followed by zero or more optional <note> elements and an
optional <timestamp>. Multiple <note> elements would appear to
convey the same note in multiple languages.

A client that receives a PIDF document containing the <device> and
<person> elements, but does not understand them (because it doesn't
implement this specification), will ignore them. Furthermore, since
the semantics of service as defined here are aligned with the meaning
of a tuple as defined in RFC 2778 and RFC 3863, documents
incorporating the concepts defined in this model are compliant with
older implementations.

It's important to note that the mapping of the presence data model
into a PIDF document is merely an exercise in syntax.

Presence documents created according to this model MUST be valid,
with the following exception. A compositor is permitted to create a
presence document that it cannot fully validate but that otherwise
validates when processed according to the lax processing rules
allowed by the schema of the compositor. However, it is not expected
that entities receiving these documents would perform schema
validation; rather, they would merely access the information from the
document in the places they were expecting it to be. Implementations
SHOULD be prepared to receive documents that are not valid, and
extract whatever information from them that they can parse.

5.1. XML Schemas

The XML schemas are broken into a common schema, called common-
schema.xsd, which contains common type definitions, and the rest of
the data model, data-model.xsd.

5.1.1. Common Schema

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
 elementFormDefault="qualified" attributeFormDefault="unqualified">
 <xs:import namespace="http://www.w3.org/XML/1998/namespace"
  schemaLocation="http://www.w3.org/2001/xml.xsd"/>
 <xs:simpleType name="Timestamp_t">
  <xs:annotation>
   <xs:documentation>Timestamp type</xs:documentation>
  </xs:annotation>
  <xs:restriction base="xs:dateTime"/>
 </xs:simpleType>
 <xs:simpleType name="deviceID_t">
  <xs:annotation>
   <xs:documentation>Device ID, a URN</xs:documentation>
  </xs:annotation>
  <xs:restriction base="xs:anyURI"/>
 </xs:simpleType>
 <xs:complexType name="Note_t">
  <xs:annotation>
   <xs:documentation>Note type</xs:documentation>
  </xs:annotation>
  <xs:simpleContent>
   <xs:extension base="xs:string">
    <xs:attribute ref="xml:lang"/>
   </xs:extension>
  </xs:simpleContent>
```

```
      </xs:complexType>
     <xs:attributeGroup name="fromUntil">
      <xs:attribute name="from" type="xs:dateTime"/>
      <xs:attribute name="until" type="xs:dateTime"/>
     </xs:attributeGroup>
     <xs:complexType name="empty"/>
    </xs:schema>
```

5.1.2. Data Model

```
    <?xml version="1.0" encoding="UTF-8"?>
    <xs:schema targetNamespace="urn:ietf:params:xml:ns:pidf:data-model"
     xmlns:xs="http://www.w3.org/2001/XMLSchema"
     xmlns="urn:ietf:params:xml:ns:pidf:data-model"
     elementFormDefault="qualified" attributeFormDefault="unqualified">
     <xs:include schemaLocation="common-schema.xsd"/>
     <xs:element name="deviceID" type="deviceID_t">
      <xs:annotation>
       <xs:documentation>Device ID, a URN</xs:documentation>
      </xs:annotation>
     </xs:element>
     <xs:element name="device">
      <xs:annotation>
       <xs:documentation>Contains information about the
        device</xs:documentation>
      </xs:annotation>
      <xs:complexType>
       <xs:sequence>
        <xs:any namespace="##other" processContents="lax"
         minOccurs="0" maxOccurs="unbounded"/>
        <xs:element ref="deviceID"/>
        <xs:element name="note" type="Note_t" minOccurs="0"
         maxOccurs="unbounded"/>
        <xs:element name="timestamp" type="Timestamp_t" minOccurs="0"/>
       </xs:sequence>
       <xs:attribute name="id" type="xs:ID" use="required"/>
      </xs:complexType>
     </xs:element>
     <xs:element name="person">
      <xs:annotation>
       <xs:documentation>Contains information about the human
        user</xs:documentation>
      </xs:annotation>
      <xs:complexType>
       <xs:sequence>
        <xs:any namespace="##other" processContents="lax"
         minOccurs="0" maxOccurs="unbounded">
         <xs:annotation>
```

```
      <xs:documentation>Characteristic and status
        information</xs:documentation>
     </xs:annotation>
    </xs:any>
    <xs:element name="note" type="Note_t" minOccurs="0"
     maxOccurs="unbounded"/>
    <xs:element name="timestamp" type="Timestamp_t" minOccurs="0"/>
   </xs:sequence>
   <xs:attribute name="id" type="xs:ID" use="required"/>
  </xs:complexType>
 </xs:element>
</xs:schema>
```

6. Extending the Presence Model

   When new presence attributes are added, any such extension has to
   consider the following questions:

   1. Is the new attribute applicable to person, service, or device
       data components? If it is applicable to more than one, what is
       its meaning in each context? An extension should strive to have
       each attribute concisely defined for each area of applicability,
       so that a source can clearly determine to which type of data
       component it should be applied.

   2. Does it belong in a new namespace, or an existing one?
       Generally, new presence attributes defined within the same
       specification SHOULD belong to the same namespace. Presence
       attributes defined in separate specifications, but produced in a
       coordinated way by a centralized administration, MAY be placed in
       the same namespace. Doing so, however, requires the centralized
       administration to ensure that there are no collisions of element
       names across those specifications. Furthermore, if a new
       extension has elements meant to be placed as the children of
       another element at a point of extensibility defined by <any
       namespace="##other">, the new extension MUST use a different
       namespace than that of its parent elements.

   3. Does the extension itself require extensibility? If so, points
       of extension MUST be defined in the schema, and SHOULD be done
       using the <any namespace="##other"> construct.

7. Example Presence Document

   In this section, we give an example of a physical system, present the
   model of that system using the concepts described here, and then show
   the resulting presence document. The example makes use of presence
   attributes defined in [23] and [24].

7.1. Basic IM Client

   In this scenario, a provider is offering a service very similar to
   the instant messaging services offered today by the public providers
   like AOL, Yahoo!, and MSN. In this service, each user has a "screen
   name" that identifies the user in the service. A single client,
   generally a PC application, connects to the service at a time. When
   the client connects, this fact is made available to other watchers of
   that user in the system. The user has the ability to set a textual
   note that describes what they are doing, and this note is seen by the
   watchers in the system. The user can set one of several status
   messages (busy, in a meeting, etc.), which are pre-defined notes that
   the system understands. If a user does not type anything on their
   keyboard for some time, the user's status changes to idle on the
   screens of the various watchers of the system. The system also
   indicates the amount of time that the user has been idle.

   Whenever a user is connected to the system, they are capable of
   receiving instant messages. A user can set their status to
   "invisible", which means that they appear as offline to other users.
   However, if an IM is sent to them, it will still be delivered.

   This system is modeled by representing each presentity in the system
   with three data components: a person component, a service component,
   and a device component. The person component describes the state of
   the user, including the note and the pre-defined status messages.
   These represent information about the human user, so they are
   included in the person component. The service tuple represents the
   IM service. No characteristics are included. The service URI
   published by the client is set to the client's Address of Record
   (AOR). The device component is used to model the PC. The device
   component includes the <user-input> element [23], since the idleness
   refers to usage of the device, not the service.

   The document published by the client would look like this:

   <?xml version="1.0" encoding="UTF-8"?>
   <presence xmlns="urn:ietf:params:xml:ns:pidf"
    xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"
    xmlns:rp="urn:ietf:params:xml:ns:pidf:rpid"
    xmlns:caps="urn:ietf:params:xml:ns:pidf:caps"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
    <tuple id="sg89ae">
     <status>
      <basic>open</basic>
     </status>
     <dm:deviceID>mac:8asd7d7d70</dm:deviceID>
     <caps:servcaps>

```
        <caps:extensions>
         <caps:supported>
          <caps:pref/>
         </caps:supported>
        </caps:extensions>
        <caps:methods>
         <caps:supported>
          <caps:MESSAGE/>
          <caps:OPTIONS/>
         </caps:supported>
        </caps:methods>
       </caps:servcaps>
       <contact>sip:someone@example.com</contact>
      </tuple>
      <dm:person id="p1">
       <rp:activities>
        <rp:on-the-phone/>
       </rp:activities>
      </dm:person>
      <dm:device id="pc122">
       <rp:user-input>idle</rp:user-input>
       <dm:deviceID>mac:8asd7d7d70</dm:deviceID>
      </dm:device>
     </presence>
```

It is worth commenting further on the value of having a separate
device element just to convey the idle indicator. The idle
indication of interest is really an indicator that the device is
idle. By making that explicit, the idle indicator can be used by the
presence server to affect the state of other services running on the
same device. For example, let's say there is a VoIP application
running on the same device. This application reports its presence
state separately, but indicates that it runs on the same device.
Since it has indicated that it runs on the same device, the presence
server can use the status of the service to further refine the idle
indicator of the device. Specifically, if the user is using its VoIP
application, the presence server knows that the device is in use,
even if the IM application reports that the device is idle.
Typically, idleness is determined by lack of keyboard or mouse input,
neither of which might be used during a VoIP call.

In a more simplistic case, reporting the idle indicator as part of
the device status allows that indicator to be used for other services
on the same device. Taking, again, the example of the VoIP
application on the same device, if the VoIP application does not
report any device information, and a watcher is not provided
information on the IM service, the presence document sent to the
watcher can include the device status. Because of the usage of the

device IDs and the device information, the presence server can
correlate the device status as reported by the IM application with
the VoIP service, and use them together.

8. Security Considerations

The presence information described by the model defined here is very
sensitive. It is for this reason that privacy filtering plays a key
role in the processing of presence data. Privacy filtering is the
act of applying permissions to a presence document for the purposes
of removing information that a watcher is not authorized to see. In
more general terms, privacy filtering is a form of authorization.
Privacy filtering can also ensure that a watcher cannot see any
presence data for a presentity, and indeed, it can even ensure that
the presentity doesn't know that it is being blocked. The SIP
presence specifications (RFC 3856 [21]) require that such
authorization processing be performed before divulging presence
information. Specifications have also been defined for conveying
authorization policies to presence servers [26].

Integrity of presence information is also critical. Modification of
presence data by an attacker can lead to diverted communications, for
example. Protocols used to transport presence data, such as SIP for
presence, are used to provide necessary integrity functions.

9. Internationalization Considerations

This specification defines a data model that contains mostly tokens
that are meant for consumption by programs, not directly by humans.
Programs are expected to translate those tokens into language-
appropriate text strings according to the preferences of the watcher.

However, this specification defines a <note> element that can contain
free text. This element and other ones defined by extensions to PIDF
that can contain free text SHOULD be labeled with the 'xml:lang'
attribute to indicate their language and script. This specification
allows multiple occurrences of the <note> element so that the
presentity can convey the note in multiple scripts and languages. If
no 'xml:lang' attribute is provided, the default value is "i-default"
[8].

Since the presence model is represented in XML, it provides native
support for encoding information using the Unicode character set and
its more compact representations including UTF-8. Conformant XML
processors recognize both UTF-8 and UTF-16. Though XML includes
provisions to identify and use other character encodings through use
of an "encoding" attribute in an <?xml?> declaration, use of UTF-8 is

RECOMMENDED in environments where parser encoding support
incompatibility exists.

10. IANA Considerations

There are several IANA considerations associated with this
specification.

10.1. URN Sub-Namespace Registration

This section registers a new XML namespace, per the guidelines in [4]

    URI: The URI for this namespace is
    urn:ietf:params:xml:ns:pidf:data-model.

    Registrant Contact: IETF, SIMPLE working group, (simple@ietf.org),
    Jonathan Rosenberg (jdrosen@jdrosen.net).

    XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
    content="text/html;charset=iso-8859-1"/>
  <title>A Data Model for Presence</title>
</head>
<body>
  <h1>Namespace for Presence Data Model</h1>
  <h2>urn:ietf:params:xml:ns:pidf:data-model</h2>
  <p>See <a href="http://www.rfc-editor.org/rfc/rfc4479.txt">
    RFC4479</a>.</p>
</body>
</html>
END
```

10.2. XML Schema Registrations

   This section registers two XML schemas per the procedures in [4].

10.2.1. Common Schema

   URI: urn:ietf:params:xml:schema:pidf:common-schema.

   Registrant Contact: IETF, SIMPLE working group, (simple@ietf.org),
      Jonathan Rosenberg (jdrosen@jdrosen.net).

   The XML for this schema can be found as the sole content of
      Section 5.1.1.

10.2.2. Data Model

   URI: urn:ietf:params:xml:schema:pidf:data-model.

   Registrant Contact: IETF, SIMPLE working group, (simple@ietf.org),
      Jonathan Rosenberg (jdrosen@jdrosen.net).

   The XML for this schema can be found as the sole content of
      Section 5.1.2.

11. Acknowledgements

   This document is really a distillation of many ideas discussed over a
   long period of time. These ideas were contributed by many
   participants in the SIMPLE working group. Aki Niemi, Paul Kyzivat,
   Cullen Jennings, Ben Campbell, Robert Sparks, Dean Willis, Adam
   Roach, Hisham Khartabil, and Jon Peterson contributed many of the
   concepts that are described here. Example presence documents came
   from Robert Sparks' example presence documents specification, and
   ideas on defining services through characteristics, rather than
   enumeration, came from Adam Roach's service features document. A
   special thanks to Steve Donovan for discussions on the topics
   discussed here, and to Elwyn Davies for his final review of the
   document.

12. References

12.1. Normative References

   [1]   Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., and
         J. Peterson, "Presence Information Data Format (PIDF)", RFC
         3863, August 2004.

   [2]   Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Caller
         Preferences for the Session Initiation Protocol (SIP)", RFC
         3841, August 2004.

   [3]   Crocker, D. and P. Overell, "Augmented BNF for Syntax
         Specifications: ABNF", RFC 4234, October 2005.

   [4]   Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January
         2004.

   [5]   Yergeau, F., Paoli, J., Sperberg-McQueen, C., Bray, T., and E.
         Maler, "Extensible Markup Language (XML) 1.0 (Third Edition)",
         W3C REC REC-xml-20040204, February 2004.

   [6]   Maloney, M., Beech, D., Thompson, H., and N. Mendelsohn, "XML
         Schema Part 1: Structures Second Edition", W3C REC REC-
         xmlschema-1-20041028, October 2004.

   [7]   Malhotra, A. and P. Biron, "XML Schema Part 2: Datatypes Second
         Edition", W3C REC REC-xmlschema-2-20041028, October 2004.

   [8]   Alvestrand, H., "IETF Policy on Character Sets and Languages",
         BCP 18, RFC 2277, January 1998.

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

12.2. Informative References

   [10] Day, M., Rosenberg, J., and H. Sugano, "A Model for Presence
         and Instant Messaging", RFC 2778, February 2000.

   [11] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A.,
         Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP:
         Session Initiation Protocol", RFC 3261, June 2002.

   [12] Peterson, J., "Common Profile for Presence (CPP)", RFC 3859,
         August 2004.

   [13] Saint-Andre, P., "Internationalized Resource Identifiers (IRIs)
         and Uniform Resource Identifiers (URIs) for the Extensible
         Messaging and Presence Protocol (XMPP)", Work in Progress,
         December 2005.

   [14] Wilde, E. and A. Vaha-Sipila, "URI Scheme for GSM Short Message
         Service", Work in Progress, February 2006.

   [15] Schulzrinne, H., "The tel URI for Telephone Numbers", RFC 3966,
         December 2004.

   [16] Klyne, G., "A Syntax for Describing Media Feature Sets", RFC
         2533, March 1999.

   [17] Rosenberg, J., Schulzrinne, H., and P. Kyzivat, "Indicating
         User Agent Capabilities in the Session Initiation Protocol
         (SIP)", RFC 3840, August 2004.

   [18] Faltstrom, P. and M. Mealling, "The E.164 to Uniform Resource
         Identifiers (URI) Dynamic Delegation Discovery System (DDDS)
         Application (ENUM)", RFC 3761, April 2004.

   [19] Rosenberg, J., "The Session Initiation Protocol (SIP) and
         Spam", Work in Progress, March 2006.

   [20] Leach, P., Mealling, M., and R. Salz, "A Universally Unique
         IDentifier (UUID) URN Namespace", RFC 4122, July 2005.

   [21] Rosenberg, J., "A Presence Event Package for the Session
         Initiation Protocol (SIP)", RFC 3856, August 2004.

   [22] Rosenberg, J., "Obtaining and Using Globally Routable User
         Agent (UA) URIs (GRUU) in the Session Initiation Protocol
         (SIP)", Work in Progress, October 2005.

   [23] Schulzrinne, H., "RPID: Rich Presence Extensions to the
         Presence Information Data Format (PIDF)", RFC 4480, July 2006.

   [24] Lonnfors, M. and K. Kiss, "Session Initiation Protocol (SIP)
         User Agent Capability Extension to Presence Information Data
         Format (PIDF)", Work in Progress, January 2006.

   [25] Peterson, J., "A Presence-based GEOPRIV Location Object
         Format", RFC 4119, December 2005.

   [26] Rosenberg, J., "Presence Authorization Rules", Work in
         Progress, March 2006.

   [27] Jennings C. and R. Mahy, "Managing Client Initiated Connections
         in the Session Initiation Protocol (SIP)", Work in Progress,
         March 2006.

Author's Address

   Jonathan Rosenberg
   Cisco Systems
   600 Lanidex Plaza
   Parsippany, NJ 07054
   US

   Phone: +1 973 952-5000
   EMail: jdrosen@cisco.com
   URI: http://www.jdrosen.net

Full Copyright Statement

Intellectual Property

Acknowledgement