Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 12, 2016

C. Lonvick
April 10, 2016

A Taxonomy on Private Use Fields in Protocols
draft-lonvick-private-tax-10.txt

Abstract

This document attempts to provide some clarification for the way that private use fields have been used in protocols developed in the IETF. It is strictly a taxonomy of what has been published and offers a minimal amount of advice about how to design or use private use options.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on October 12, 2016.

Table of Contents

1. Introduction

Simply put, communications protocols are standardized ways for computing entities to convey information. Within each communications protocol, there must be standardized pieces of information that will be communicated, and there may be non-standardized pieces that can be communicated. Since one of the goals of standards is to provide interoperability, all parties participating in any communications protocol must be aware of how to deal with all fields in the protocol. Fields reserved for private use cannot provide interoperability unless their use is fully documented in openly available documents. This section uses examples of some well known protocols to demonstrate the differences between protocols that use private use options, and those that don't.

Existing standards permit private use options in different ways. The Time Protocol [RFC0868] is an example of a protocol that only conveys standardized information. There is no way to add anything other than what is specified in the document. On the other hand, DOD STANDARD TRANSMISSION CONTROL PROTOCOL [RFC0761] does have "options" but they must be registered through the IANA [IANAtcp] before use, which does not leave any room for optional information supplied by equipment vendors, network operators, or experimenters. Finally, Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4) [RFC3925] does allow for vendor specific options that do not need to be registered with anyone.

If a network operator wanted to add specific information to the Time Protocol [RFC0868], they could modify the code of all senders and receivers and run this within their own domain without any problems. However, if an equipment vendor wanted to include information specific to their equipment, they would have to ensure that all senders and receivers within all network domains would either accept the change in the protocol, or would not have problems with it. As a final case, if several equipment vendors desired to add equipment-specific information to this protocol, they would have to take great care that only their own receivers would accept information from their own transmitters. An extension to that would be that if one equipment vendor would like to transmit or receive the same information that another vendor is using.

For the case of TCP [RFC0761], standard options are expected; senders may use them and receivers may be configured to act upon that information, or to ignore it. If an experimenter wants to add an option, they will have to create a new IETF RFC with specific details, or obtain approval from the IESG to have the IANA add to the registry [IANAtcp]. Similarly, if equipment vendors Foo and Bar were to have a need for a similar option within TCP, they would each have

to go through the process to add to the registry. On the other hand, if a properly crafted multipurpose private use option were to be registered, such as in the case of multiple vendor instances within DHCPv4 [RFC3925], then vendors and experimenters would each be able to use it for their own purposes as long as all network participants could easily differentiate between the entities using the option.

This document explores the various ways that protocols have allowed optional information to be included using fields designated as "private use". It uses examples of some well known protocols. In well developed protocols, private use options may be useful in avoiding allocation conflicts, and in dynamically extending a feature. As with all good things, this will come with a cost. Adding any extra fields to a protocol will require additional processing for both the sender and the receiver. Also, larger packets will take up more bandwidth in transmission. In another aspect, a receiver will have to reserve buffers for an expected field in an inbound packet. Since one way of implementing private use options is to only enable the field if it is needed, then the allocation of buffers could be considered wasteful if it is actually not used.

2. Origins of the Private Use Namespace

Guidelines for Writing an IANA Considerations Section in RFCs [RFC2434] describes that values of specific namespaces may either be registered with the IANA, or not. In most cases, there are well defined values for namespaces. However, as the document explains, not all namespaces require centralized administration.

In that document, it seems to be assumed that private use namespaces will be domain specific and it will be up to the administrators of any domain to avoid conflicts. The first example given about private use namespaces refers to Dynamic Host Configuration Protocol [RFC2131] and presumably DHCP Options and BOOTP Vendor Extensions [RFC2132]. In this the example states that "site-specific options in DHCP have significance only within a single site". As noted below this became a problem that was rectified in a later revision of DHCP.

Later works identified a need to place a scope on private use namespaces. The second example of private use namespaces in the IANA guidelines [RFC2434] is from STANDARD FOR THE FORMAT OF ARPA INTERNET TEXT MESSAGES [RFC0822] which describes X- headers. Again, there is no effort made to control the namespace. It appears however that the users of X- headers have self-organized; most consistently use features that are universally useful and many have incorporated identifiers for useful features that may overlap.

3. Nomenclature

In this document, the following words are defined to prevent ambiguity. Some of these words have not been used in the referenced works but their meanings can be easily ascertained and applied.

- Communications protocol - a formal description of digital message formats and the rules for exchanging those messages in or between computing systems and in telecommunications [wpProt]

  Example: The File Transfer Protocol [RFC0959] is an example of a communications protocol. It has well defined fields and standard options. The Syslog Protocol [RFC5424] is another example of a communications protocol. It has well defined fields, standard options, and it also has standard and private use options. (See Section 5.1.5.)

- Protocol frame - a defined container of fields used to convey information in a communications protocol

  Example: An Internet Protocol packet [RFC0791] is considered to be a protocol frame. In the case of The File Transfer Protocol [RFC0959], an FTP message from the client to the server within the Internet Protocol [RFC0791] containing an FTP command is a protocol frame. In the case of The Syslog Protocol [RFC5424], a message from the client to the server within the Internet Protocol [RFC0791] containing a syslog message is also a protocol frame.

- Field - any defined container within a communications protocol frame

  Example: In the case of The File Transfer Protocol [RFC0959], a command will be contained within a field. In the case of The Syslog Protocol [RFC5424], the HOSTNAME is a field.

- Standard option - a field in a protocol frame that may only use values that are strictly defined within a specification

  Example: In the case of The File Transfer Protocol [RFC0959], an FTP command, such as CDUP or QUIT, is a standard option. The reason that a command is a standard option is that only the values listed by the IANA in the registry [IANAftp] may be used. The standard options are not limited to the values defined in the original RFC, but also include any additions to the registry. In the case of The Syslog Protocol [RFC5424], an SD-ID may be a standard option. The example given in Section 7.1.4 of [RFC5424] of

[timeQuality tzKnown="0" isSynced="0"]

is a standard option because all of the fields are listed in the document and in the IANA registry [IANAslg].

- Private use option - a field in a protocol frame that is reserved for private or local use only namespaces

  Example: In the case of The Syslog Protocol, an SD-ID may be a private use option. Example 3 given in Section 6.5 contains a private use option.

  <165>1 2003-10-11T22:14:15.003Z mymachine.example.com evntslog - ID47 [exampleSDID@32473 iut="3" eventSource= "Application" eventID="1011"] BOMAn application event log entry...

  Specifically, the SD-ID starting with "[exampleSDID@32473 ..." is not a specifically defined option in the RFC, nor is it registered in the IANA registry [IANAslg]. It is a way for an equipment vendor to insert their specific information without having to register anything. In this case if the receiver knows the format of that SD-ID then it can immediately interpret its meaning. However, if it does not know how to interpret that SD-ID, it can still log the message and an Operator or Administrator can look up its meaning at a later time.

- Namespace - the set of possible values a field may contain; its actual content may be a name, a number or another kind of value

  Example: In the same Example 3 from Section 6.5 of The Syslog Protocol [RFC5424], "exampleSDID@32473" provides the namespace so the context of the rest of the SD-ID may be interpreted. Specifically, the Private Enterprise Number [IANApen] (PEN) is used to associate the option with a private enterprise, and the text before the "@" identifies the option defined within that private enterprise.

Additionally, the terms "Source of Authority" and "Focus of the Namespace" are defined and further discussed below.

It should also be noted that some references use the term "name space" to refer to namespace. The IETF has been fairly consistent in using the term "namespace" in documents and this specification follows that precedence.

4. Characteristics of Useful Private Use Options

Private use options can be separated into discreet pieces of information. The interpretation of each piece of information places its context. The interpretation of the entirety of these pieces of information will uniquely describe the context of the information and the value associated with it. This must provide a single and unique interpretation of the information to each receiver.

This section summarizes the observed characteristics of private use options that are successful and deployed. Following sections will explain how these characteristics apply to specific protocols that are commonly used in the Internet.

There seem to be three characteristics of successful private use options:

A Source of Authority

A Focus of the Namespace

A Value of the Option

As an example, in SNMP the combination of the Source of Authority and the Focus of the Namespace (Focus) represent the OID. The combination of the Source of Authority, the Focus, and the Value of the Option (Value) constitute the VarBind.

4.1. Source of Authority

A private use option requires a path to an origin that has the authority to create and maintain the option. As shown above, this referent should be unique, and not be dependent upon local interpretation.

The name "Source of Authority" comes from the domain name system configuration file which enumerates a "SoA" as the person or entity who has ultimate control and decision making powers over the scope of the domain. Some liberties have been taken with using this name but the intent is to identify an authoritative source for the namespace.

The PEN (Section 5.1) is sourced by the Internet Assigned Numbers Authority (IANA). These may be viewed as being similar to domain names in that they are acquired by individuals, corporations, or other organizations. A notable difference is that when domain names fall into disuse they may be acquired and used by entirely different people or organizations - as per the conditions required by the Internet Corporation for Assigned Names and Numbers [ICANN], the

source of the domain names. The structure of the PEN registry does not place any limits on the time that a PEN will be active or associated with the requester. This is no different from many other registries maintained by the IANA; they are just a snapshot at the time of the reservation based on the information required by the IANA and provided by the applicant. This eternal association of the PEN, versus the ephemeral association of domain names, has not been shown to present any problems. This may, in fact, be a feature as this methodology ensures that these namespaces stay unique for the foreseeable future.

Domain names have similar problems as they can be more ephemeral than eternal. Unlike PENs that become unserviceable when their owning organization goes out of business, domain names that fall into disuse may be acquired and used by entirely different organizations. Similar to the use of PENs there have not been any problems reported from this.

It is vital to note that the usage of the option within the private space is the full responsibility of the private entity. In the example of the PEN, each entity registering a PEN must fully quantify the parameters of the use of the option within their purview.

4.2. Focus of the Namespace

Once the source of authority is established, an actual option, or multiple options, must be specified. This is usually an indicator of what value is expected. Within the domain established by the source of authority, the focus of each value must be unique. In a very simple example, a private use option may consist of "PEN"@"focus"="value". The PEN will be unique and will specify the source of authority. The focus will be unique as long as the source of authority maintains that uniqueness; e.g., it would be poor form for a private enterprise to define a focus, then to redefine it at a later time.

In some cases, multiple focuses and values need to be transmitted. When the PEN has been used, this has most often been achieved by nesting "type length value" (tlv's) within the field. Each type is then a focus for the private use option. More recently URIs have been used to point to a source of authority. This allows an organization to organize an abundance of information about their namespaces.

5. Examples of Successful Private Use Options

This section contains a review of RFCs that allow the use of private

use options. There seem to be three ways to address the namespace: via a global origin, via a truncated numerical origin, and via a namespace based upon a domain name.

5.1. Private Enterprise Number

Rather than using the entire SMI, protocol engineers started using just the Private Enterprise Number [IANApen]. This reduces the length of the identifier but continues to provide an identifier through a globally unique namespace. This section provides examples of how the PEN has been used to provide private use options.

5.1.1. SNMP

Likely, the first private use option was defined in the Structure and Identification of Management Information for TCP/IP-based Internets [RFC1155] which was first used in A Simple Network Management Protocol [RFC1067] (SNMP). The structure of management information (SMI) has been updated and is currently defined as the Structure of Management Information Version 2 (SMIv2) [RFC2578].

SMI is a well described tree of OBJECT IDENTIFIERs (OIDs). OIDs have an origin and a path for defined object identifiers which this document describes as standard options. It also allows for experimental and vendor specific object identifiers, which are described as private use options in this document. The IANA maintains a registry of these Network Management Parameters [IANAsmi].

The Internet subtree of experimental OBJECT IDENTIFIERs starts with the prefix: 1.3.6.1.3., and the Internet subtree of private enterprise OBJECT IDENTIFIERs starts with the prefix: 1.3.6.1.4.1. This is followed by a Private Enterprise Number [IANApen] (PEN) and then the objects defined by that enterprise.

The globally unique origin in SNMP (Section 5.1.1) is the International Standards Organization [ISO] which is accredited by the United Nations to maintain this structure. However, the namespace resolves to the PEN (Section 5.1).

After the vendor identifier (the PEN) in the management information base (MIB), a vendor can create many different trees to identify objects. This may result in a very large number of OBJECT IDENTIFIERs; each of which is an identifier of the namespace described in this document. Each of these are uniquely identified by the vendor and do not require registration with any coordinating authority.

The last part of each OBJECT IDENTIFIER is the value corresponding to the focus, which is known as the varbind. In a GetRequest the server fills this field with a "0" and the client responds by replacing the "0" with the actual value. Since this field is defined by the vendor, it may actually be a concatenation of values. In a SetRequest transmitted to the receiver, this is the last field.

In this, each OBJECT IDENTIFIER contains a globally unique origin which is ISO, a focus which is the OBJECT IDENTIFIER down to the last field, and a value which is the last field in the SetRequest, and the last field in the response to a GetRequest.

Specific codes, known as error-indexes, are used to indicate when a request cannot be processed because a device does not understand a request.

While this is very practical for SNMP, fully qualified OIDs are not always well suited to be used as an indicator for private use options. In many other uses, the source of authority has been truncated to just the PEN (Section 5.1).

5.1.2. RADIUS

The Remote Authentication Dial In User Service (RADIUS) [RFC2058] specification documented how to use just the PEN (without the rest of the SMI path to the root) to allow "vendors" to articulate their own options. In that document, these are called Vendor-Specific Attributes (VSA).

The updated RADIUS document, [RFC2865], gives guidance for using the VSA.

- Servers not equipped to interpret the vendor-specific information sent by a client MUST ignore it (although it may be reported).
- Clients which do not receive desired vendor-specific information SHOULD make an attempt to operate without it, although they may do so (and report they are doing so) in a degraded mode.
- The Attribute-Specific field is dependent on the vendor's definition of that attribute.
- It SHOULD be encoded as a sequence of vendor type / vendor length / value fields.
- Multiple subattributes MAY be encoded within a single Vendor-Specific Attribute, although they do not have to be.

There are many attributes defined in RADIUS [RFC2058] which may be considered to be standard options. Each of these attributes is specified within a "type length value" (tlv) container. For this protocol, the attribute "type" is a specific numerical value which differentiates it from other attributes. As an example, the User-Name (type 1) and User-Password (type 2) may be considered to be standard options as they are well defined within the specification.

Type 26 denotes the Vendor Specified Attribute. It is "available to allow vendors to support their own extended Attributes not suitable for general usage". The PEN starts the "value" which should then include a subsequent nested tlv so the vendor may define and enumerate their own options within that field.

As noted above, the globally unique origin for RADIUS [RFC2865] is the PEN. The remainder of the Attribute field after the PEN is deliberately undefined in the specification. It is however suggested that the field contain embedded tlv's. This is again very practical. Each vendor may then have conflicting "types" (e.g. "1") which would be disambiguated by the origin. For example {PEN="N", type="1"} is different from {PEN="M", type="1"}. Since there is nothing to prevent vendors from registering multiple PENs, each vendor may have a plethora of {type="1"}. However, that is actually not needed since the focus may be extended by enumerating multiple types. For example, the vendor attribute may contain {PEN="M", type="1"(value), type="2"(value), type="3"(value)}.

The values for each type are bounded by the length of the attribute. Since the entire vendor attribute is defined by the vendor, the values may be human readable or not. Since the protocol tends to be machine-to-machine, it is likely that the values will not be human readable. In some cases, it is feasible that a value has no length. In that case, the transmission of the type alone, would be a signal of some sort to the receiver.

5.1.3. Mobile IP

Mobile IP Vendor Specific Extensions [RFC3115] defines two extensions that can be used for making organization specific extensions by vendors/organizations for their own specific purposes for Mobile IP [RFC2002]. Mobile IP has been revised several times and is currently specified in IP Mobility Support for IPv4, Revised [RFC5944].

In that specification, two tlv's have been defined to contain private use options. These are collectively called Vendor/Organization Specific Extensions (VSE).

- When the Critical Vendor/Organization Specific Extension (CVSE) is encountered but not recognized, the message containing the extension MUST be silently discarded.

- When a Normal Vendor/Organization Specific Extension (NVSE) is encountered but not recognized, the extension SHOULD be ignored, but the rest of the Extensions and message data MUST still be processed.

Having two VSEs of this nature for private use options is consistent with the original Mobile IP specification [RFC2002] which states:

> When an Extension numbered in either of these sets within the range 0 through 127 is encountered but not recognized, the message containing that Extension MUST be silently discarded. When an Extension numbered in the range 128 through 255 is encountered which is not recognized, that particular Extension is ignored, but the rest of the Extensions and message data MUST still be processed.

The structure of the origin, type, and value of the CVSEs and NVSEs for Mobile IP [RFC3115] may be used in a manner very similar to that of RADIUS. The PEN is the origin and types and values may be stacked within the field following that.

It should be noted that this does not have to be the case. Specifying CVSEs and NVSEs in various packets can give a vendor another dimension in processing these private use fields. If a vendor placed all CVSEs in a single packet, and the receiver did not understand any one of them, the entire packet must be discarded. However, if the vendor places individual CVSEs in separate packets, only CVSEs that are not understood by the receiver will be discarded.

Similarly, a vendor may choose to not stack NVSEs so that a receiver won't discard the entire cluster of NVSEs if a single one is not understood.

The values are constrained by the length of the types or subtypes.

5.1.4. DHCP

The introduction to Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4) [RFC3925] states:

> The DHCP protocol for IPv4, [RFC2131], defines options that allow a client to indicate its vendor type (option 60), and the DHCP client and server to exchange vendor-specific information (option 43) [RFC2132]. Although there is no prohibition against passing

> multiple copies of these options in a single packet, doing so would introduce ambiguity of interpretation, particularly if conveying vendor-specific information for multiple vendors.

This meant that Dynamic Host Configuration Protocol [RFC2131] specified that there was one instance of the vendor type, and the receiver used that namespace to set the scope for the fields in the vendor-specific information option. This version of DHCP did not allow for multiple origins; only a single origin was permitted and the types were to be defined subsequent to that. Evidently this was found to be unworkable when different vendors needed to expand private use options in the protocol.

This situation was resolved with the publication of Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4) [RFC3925] which states:

> The Dynamic Host Configuration Protocol (DHCP) options for Vendor Class and Vendor-Specific Information can be limiting or ambiguous when a DHCP client represents multiple vendors.

That specification ([RFC3925]) then used the PEN [IANApen] to define a unique namespace for private use options in this protocol. Similar to other protocols of this era, tlv containers were used.

When this protocol was updated to conform to the requirements of IPv6, the PEN was again used as the way to identify the origin of the private use option.

5.1.5. Syslog

The Syslog Protocol [RFC5424] also uses the PEN to uniquely qualify the namespace for a private use option. Standard options do not contain the "@" character. Private use options must have the PEN following the "@" character. This allows a vendor or experimenter to have overlapping namespaces which the PEN will then uniquely identify. For example the standard option of tzKnown may only have associated values of "0" and "1". However tzKnown@32473 may have any value assigned to it by the owner of enterprise number 32473.

Syslog transport receivers are supposed to accept all correctly formatted Syslog messages. Unlike RADIUS, the receiving Syslog application does not have to have immediate knowledge of all variable options to continue operations. If a private use option is not immediately known to the receiving application, it may still store the message and an Operator or Administrator may look it up at a later time if they are really interested.

The Syslog protocol [RFC5424] uses the PEN as the origin and allows for the focus of the private use option to be fully defined by the vendor within the structured data. Specifically, a vendor may define a "type" of private use option by concatenating it with the PEN by using the @ character. Within the bounds of the structured data, multiple elements may be used that have identifiers and values.

5.2. Domain Name Strings

An alternative to using numerical indicators is to use textual strings. Again, the goal for using these strings is to disambiguate the identifiers and allow freedom of expression by the vendors and experimenters using them.

5.2.1. Secure Shell

The Secure Shell (SSH) Protocol Architecture [RFC4251] uses character strings rather than PENs. Similar to Syslog, but actually predating it, standard options must not have the "@" character in them. Private use options will have an origin identifier preceding an "@" character followed by a namespace field. For example, in The Secure Shell (SSH) Connection Protocol [RFC4254] SSH channels may be opened by specifying a channel type when sending the SSH_MSG_CHANNEL_OPEN message. Standard options for the channel type include "session" and "x11". A private use option for a channel type could be "example_session@example.com".

Obviously, these character strings are domain names [RFC1034] [RFC1035]. This is specified in The Secure Shell (SSH) Protocol Architecture [RFC4251]. Generally, the guidance given is that if a private use option of this nature is not understood it is to convey an error code to its peer.

In the SSH protocol [RFC4250], the origin is a domain name and the focus of the option is dependent upon context. For example, ourcipher-cbc@example.com can only be used when negotiating ciphers, while example_session@example.com can only be used when negotiating channel types, per the examples in [RFC4250].

5.3. URN-based Namespaces

Uniform Resource Names (URNs) have also been used to convey options. They are very flexible

(Need to add a lot here.) Uniform Resource Names (URN) Namespace Definition Mechanisms [RFC3406] An IETF URN Sub-namespace for Registered Protocol Parameters [RFC3555] The IETF XML Registry [RFC3688] Extensible Provisioning Protocol (EPP) [RFC5730] Extensible

Provisioning Protocol (EPP) Host Mapping [RFC5732] Namespaces in XML 1.0 (Third Edition) [W3C.REC-xml-names-20091208]

5.3.1. YANG and NETCONF

YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF) [RFC6020] and Network Configuration Protocol (NETCONF) [RFC6241] use URIs to indicate private use namespaces. The following is given as an example of a YANG and NETCONF configuration.

```
module my-config {
    namespace "http://example.com/schema/config";
    prefix "co";

    container system { ... }
    container routing { ... }
}
```

That example could be encoded in NETCONF as the following.

```
<rpc message-id="101"
     xmlns="urn:ietf:params:xml:ns:netconf:base:1.0"
     xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <running/>
    </target>
    <config>This eternal association
      <system xmlns="http://example.com/schema/config">
        <!-- system data here -->
      </system>
      <routing xmlns="http://example.com/schema/config">
        <!-- routing data here -->
      </routing>
    </config>
  </edit-config>
</rpc>
```

Section 8.3 of YANG [RFC6020] describes the parsing of the YANG payload. It contains a good deal of information about how to process elements or values that are not recognized.

Similarly, NETCONF [RFC6241] contains much information about processing requests that cannot be completed because elements or values are not recognized.

Both YANG [RFC6020] and NETCONF [RFC6241] use URIs to enumerate private use options of a device. The use of this comes from XPATH

[W3C.REC-xpath-19991116].

In both of these, the source of authority is the domain name in the URI and the origin is the full URI path. Many private use options may be described within YANG. From that, each private use option may be populated in NETCONF.

The following is used to demonstrate this. First the YANG module is shown, then a subset of the NETCONF is shown.

YANG module:

```
// Contents of "acme-system.yang"
module acme-system {
    namespace "http://acme.example.com/system";
    prefix "acme";

    organization "ACME Inc.";
    contact "joe@acme.example.com";
    description
        "The module for entities implementing the ACME system.";

    revision 2007-06-09 {
        description "Initial revision.";
    }

    container system {
        leaf host-name {
            type string;
            description "Hostname for this system";
        }

        leaf-list domain-search {
            type string;
            description "List of domain names to search";
        }

        container login {
            leaf message {
                type string;
                description
                    "Message given at start of login session";
            }

            list user {
                key "name";
                leaf name {
```

```
                        type string;
                    }
                    leaf full-name {
                        type string;
                    }
                    leaf class {
                        type string;
                    }
                }
            }
        }
    }
```

NETCONF exchange:

```
    <system>
      <login>
        <message>Good morning</message>
      </login>
    </system>
```

In this example, YANG describes the source of authority and focus for the login message, and the NETCONF exchange populates that specific value.

As noted above, both of these specifications have good descriptions of actions to take if a namespace is not recognized.

6. Issues to Consider

This document is not an encouragement or recommendation to define private use fields in IETF protocols. Rather, since private use options are useful to the community and seem to be gaining popularity, this document is an attempt to document the ways in which they have been successful so others may benefit.

Private use options are a way to allow vendors, network operators, and experimenters to convey dynamic information without going through a rigorous process to register each variable. There is no "one size fits all" mechanism. The use of a very specific and fixed format works very well for RADIUS which requires speed in processing. On the other hand, the open nature of the private use options in Syslog are appropriate for that protocol where event messages need not be fully parsed at the time of reception.

There seem to be four essential features to using a private use

option.

- One requirement is to have a definable way for the community to ascertain the nature of all private use options. For example, several vendors have published their RADIUS VSAs on web pages which are easy to find. From that, anyone creating a new RADIUS server would have access to, and be able to incorporate the information available.

- Instructions are needed on how to deal with private use options that are not understood by a receiver. In some cases, a receiver may not need to understand the options immediately upon receipt as in the case of Syslog. In other cases, the options are immediately used and instructions must be clear on what to do if the receiver cannot process them. It appears that Mobile IP has the best thought-through instructions on this.

- Private use options must be extensible in a clearly designed way. RADIUS suggests that the string containing the option be another tlv. This allows a vendor to define multiple private use options within their own namespace field. These are becoming known as subattributes. This appears to be working in practice and it may be assumed that this has become a de facto rule for RADIUS.

- In most cases, a unique option (both standard and private use) will only be used once within the context of an exchange. RADIUS and DHCP either state or strongly imply this. However, while it is not explicitly discussed, there is nothing to prevent this within Syslog. Some guidance should be given about this in describing private use options in protocols.

Clear documentation in full and open standards is needed to achieve uniformity and interoperability in these features. Obviously implementers will need to adhere closely to these standards for complete interoperability.

Finally, the usage of any private use values on the wire before any namespace is properly reserved with the IANA is entirely inadvisable.

6.1. Value of the Option

The value of each private use option must be well defined and bounded. It is advisable that it be extensible to accomodate future requirements.

Generally speaking, values of private use options should follow the same guidance given for standard options.

6.2. Guidance on Incomplete Understanding

Within the protocol, an understanding needs to be established between the transmitter and receiver about what to do if the receiver does not understand a namespace. Some protocols have defined that a receiver will silently discard packets that contain private use options they do not understand. Other protocols have defined that they will only discard the private use option rather than the entire packet. While other protocols have no need for the receiver to have any understanding of any private use options when it receives them. Each of these behaviors is represented in the examples in this document.

Regardless of whether or not this understanding is established, the receiver of any protocol must have a defined path of action to follow when receiving anything that it may not understand.

7. Authors Notes

This section will be removed prior to publication.

This is version -10. A lot has gone on in my life during this year and I havn't been able to update this document as quickly as I would have liked.

8. Security Considerations

This document reviews ways that options are being used in various protocols. As such, there are no security considerations inherent in this document.

Readers and implementers should be aware of the context of implementing options in their own protocols.

9. IANA Considerations

This document does not propose a standard and does not require the IANA to do anything.

10. Acknowledgments

The idea for documenting this came from questions asked in the SIP-CLF Working Group and the author is grateful for the discussion around this topic.

The following people have contributed to this document. Listing their names here does not mean that they agree with or endorse the document, but that they have contributed to its substance.

David Harrington, Dan Romascanu, Bert Wijnen, Ralph Droms, Juergen Schoenwalder, Nevil Brownlee, Klaas Wierenga, and Brian Carpenter.

11. References

[IANAtcp] Internet Assigned Numbers Authority, "IANA Transmission Control Protocol (TCP) Parameters, TCP Option Kind Numbers", 2011, <http://www.iana.org/assignments/tcp-parameters/tcp-parameters.txt>.

[IANAftp] Internet Assigned Numbers Authority, "IANA FTP Commands and Extensions", 2010, <http://www.iana.org/assignments/ftp-commands-extensions/ftp-commands-extensions.txt>.

[IANAslg] Internet Assigned Numbers Authority, "IANA syslog Parameter", 2010, <http://www.iana.org/assignments/syslog-parameters>.

[IANAsmi] Internet Assigned Numbers Authority, "Network Management Parameters", 2011, <http://www.iana.org/assignments/smi-numbers>.

[IANApen] Internet Assigned Numbers Authority, "IANA PRIVATE ENTERPRISE NUMBERS", 2011, <http://www.iana.org/assignments/enterprise-numbers>.

[wpProt] Wikipedia - the Free Dictionary, "Wikipedia entry for communication protocol", 2011, <http://en.wikipedia.org/wiki/Communications_protocol>.

[ISO] International Standards Organization, "International Standards Organization", 2011, <http://www.iso.org>.

[ICANN] Internet Corporation for Assigned Names and Numbers, "Internet Corporation for Assigned Names and Numbers", 2011, <http://www.icann.org>.

[RFC0761] Postel, J., "DoD standard Transmission Control Protocol", RFC 761, January 1980.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC0822] Crocker, D., "Standard for the format of ARPA Internet text messages", STD 11, RFC 822, August 1982.

[RFC0868] Postel, J. and K. Harrenstien, "Time Protocol", STD 26, RFC 868, May 1983.

[RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, October 1985.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC1067] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol", RFC 1067, August 1988.

[RFC1155] Rose, M. and K. McCloghrie, "Structure and identification of management information for TCP/IP-based internets", STD 16, RFC 1155, May 1990.

[RFC2002] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC2058] Rigney, C., Rubens, A., Simpson, W., and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2058, January 1997.

[RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3115] Dommety, G. and K. Leung, "Mobile IP Vendor/ Organization-Specific Extensions", RFC 3115, April 2001.

[RFC3406] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", BCP 66, RFC 3406, October 2002.

[RFC3555] Casner, S. and P. Hoschka, "MIME Type Registration of RTP Payload Formats", RFC 3555, July 2003.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.

[RFC3925] Littlefield, J., "Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4)", RFC 3925, October 2004.

[RFC4250] Lehtinen, S. and C. Lonvick, "The Secure Shell (SSH) Protocol Assigned Numbers", RFC 4250, January 2006.

[RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture", RFC 4251, January 2006.

[RFC4254] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH) Connection Protocol", RFC 4254, January 2006.

[RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.

[RFC5730] Hollenbeck, S., "Extensible Provisioning Protocol (EPP)", STD 69, RFC 5730, August 2009.

[RFC5732] Hollenbeck, S., "Extensible Provisioning Protocol (EPP) Host Mapping", STD 69, RFC 5732, August 2009.

[RFC5944] Perkins, C., "IP Mobility Support for IPv4, Revised", RFC 5944, November 2010.

[RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.

[RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

[W3C.REC-xpath-19991116] Clark, J. and S. DeRose, "XML Path Language (XPath) Version 1.0", World Wide Web Consortium

Recommendation REC-xpath-19991116, November 1999,
<http://www.w3.org/TR/1999/REC-xpath-19991116>.

[W3C.REC-xml-names-20091208]
Bray, T., Hollander, D., Layman, A., Tobin, R., and H.
Thompson, "Namespaces in XML 1.0 (Third Edition)", World
Wide Web Consortium Recommendation REC-xml-names-20091208,
December 2009,
<http://www.w3.org/TR/2009/REC-xml-names-20091208>.

Author's Address

Chris Lonvick
1307 Kent Oak Dr.
Houston, Texas 77077
US

Email: lonvick.ietf@gmail.com