

Network Working Group
INTERNET-DRAFT
Category: Experimental
<[draft-aboba-roamops-adif-00.txt](#)>
[19](#) January 2001

Bernard Aboba
Microsoft Corporation
Dave Lidyad
Telco Research Corporation

The Accounting Data Interchange Format (ADIF)

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet- Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

[1](#). Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

[2](#). Abstract

This document describes an extensible human-readable accounting record format, the Accounting Data Interchange Format (ADIF). Based on MIME, ADIF is designed to compactly represent accounting data from any protocol using attribute/value pairs (AVPs) or variable bindings.

ADIF may be used within accounting systems in several ways. In some cases, Accounting Servers will produce ADIF records based on data obtained from accounting protocols. It is also possible for devices to store data in ADIF format and transfer ADIF records to the accounting server, using an accounting or file transfer protocol. The latter approach has the advantage of offloading the Accounting Server from the task of transcribing interim or session records, thus improving scalability. In either scenario, ADIF may be used to transfer a single accounting records, or a batch of accounting records.

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

[3.](#) Introduction

This document describes an extensible human-readable accounting record format, the Accounting Data Interchange Format (ADIF). Based on MIME, ADIF is designed to compactly represent accounting data from any protocol using attribute/value pairs (AVPs) or variable bindings.

ADIF may be used within accounting systems in several ways. In some cases, Accounting Servers will produce ADIF records based on data obtained from accounting protocols. It is also possible for devices to store data in ADIF format and transfer ADIF records to the accounting server, using an accounting or file transfer protocol. The latter approach has the advantage of offloading the Accounting Server from the task of transcribing interim or session records, thus improving scalability. In either scenario, ADIF may be used to transfer a single accounting records, or a batch of accounting records.

[3.1.](#) Terminology

This document uses the following terms:

Accounting

The collection of resource consumption data for the purposes of capacity and trend analysis, cost allocation, auditing, and billing. Accounting management requires that resource consumption be measured, rated, assigned, and communicated between appropriate parties.

Rating The act of determining the price to be charged for use of a resource.

Billing The act of preparing an invoice.

Archival accounting

In archival accounting, the goal is to collect all accounting data, to reconstruct missing entries as best as possible in the event of data loss, and to archive data for a mandated time period. It is "usual and customary" for these systems to be engineered to be very robust against accounting data loss. Legal or financial requirements frequently mandate archival accounting practices, and may often dictate that data be kept confidential, regardless of whether it is to be used for billing purposes or not.

Interim accounting

An interim accounting packet provides a snapshot of usage during a user's session. This may be useful in the event of a device reboot or other network problem that prevents the

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

reception or generation of a session summary packet or session record. Interim accounting packets can always be summarized without the loss of information.

Session record

A session record represents a summary of the resource consumption of a user over the entire session. Accounting gateways creating the session record may do so by processing interim accounting events or accounting events from several devices serving the same user.

Accounting Protocol

A protocol used to convey data for accounting purposes.

Accounting server

The accounting server receives accounting data from devices and translates it into session records. The accounting server may also take responsibility for the routing of session records to interested parties.

[4.](#) Accounting record format requirements

As detailed in [\[2\]](#), solution of the accounting problem in roaming requires a standardized accounting record format to enable exchange of accounting data between members of a roaming consortium. Since operational roaming services, described in [\[1\]](#), exhibit considerable diversity in their accounting implementations it is desirable that the chosen accounting record format be protocol-independent. Since accounting implementations are continually adding new attributes, extensibility is important.

For accounting session records, compactness of representation is a virtue. While compression can decrease the space taken up by a less compact accounting record representation, the computational load of compression and decompression will add to the overhead of accounting processing, and there are situations (such as embedded devices) in which

this will be problematic.

Session records can be stored within devices where memory or non-volatile storage is at a premium. Thus the compactness of the representation will determine how much data can be stored prior to experiencing data loss.

To satisfy legal and regulatory requirements it has become customary to archive session records. Thus, the compactness of the representation will determine how much warehouse space is required to store the tapes or CD-ROMs of the archived data. In addition, where session records are transferred over the wire the compactness of representation will

determine bandwidth consumption. For all these reasons, smaller is better.

It is also desirable that accounting session records be human readable. Since the processing of accounting data may ultimately result in the transfer of funds, it is important that the software producing and handling session records be correct. Human readable accounting formats are considerably easier to implement and debug than binary formats and thus software based on them is more likely to be free of defects.

The current state of the art in accounting data representation is reviewed in [4]. While the SNMP-oriented formats described in [14], [15] are appealing for representation of data collected via SNMP, they are not appealing for use with other protocols since processing the records would require ASN.1 encode and decoding operations that would typically not otherwise be necessary. Also, ASN.1 interchange formats are not human readable and require encoding and decoding of complex binary structures. This makes them complex to work with as computationally demanding as compared with plain-text files.

Binary formats such as those used in RADIUS [5]-[9] are relatively simple, but require development of specialized tools that would not be reusable for other protocols. Existing call detail records based on fixed record formats, while being human readable, do not offer the required extensibility.

XML-based accounting record formats such as TIPHON [29] are specified by an XML DTD. XML is extensible, protocol-indepdent, and human understandable. However, XML representations are not compact and

therefore may require compression prior to storage or transmission over the wire. XML formats are also limited by XML's inability to specify type or other validity checking for information within the tags. This will be improved by the XML Schema [\[31\]](#) efforts, but a stable reference implementation is not yet available. Thus, it appears that existing approaches cannot simultaneously satisfy the requirements for generality, extensibility, compactness and human readability.

One of the goals of ADIF is to provide a universal yet simple and easy to understand accounting record format. While an ADIF parser is required to make use of the new format, it is expected that the effort required to create processing tools would be leveragable across multiple protocols and services.

ADIF is based on MIME, described in [\[16\]](#). The use of MIME enables ADIF to represent both printable and non-printable characters as well as to allow representation of attributes of unlimited size. Through the use of attribute numbers and protocol defaults, it has been possible to produce an accounting record format that is simultaneously human readable,

general, extensible, and compact.

[4.1](#). Requirements language

In this document, the key words "MAY", "MUST", "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [\[10\]](#).

[5](#). Definition of the Accounting Data Interchange Format (ADIF)

ADIF may be used to represent either a single accounting record or a batch of accounting records of arbitrary size. A batch of ADIF records consists of a header providing basic information about the records in the file, followed by a series of records each separated by a separator. The header includes the ADIF version number (1 for this document), device name/description, and collection start date and time. A default protocol type may be optionally included in the header.

Note that it is possible to use ADIF in a real-time accounting environment where individual records are transmitted by the device to the Accounting Server in ADIF format. In such a case it is up to the mechanism to specify whether ADIF headers are included within each

record or are agreed to beforehand via some out-of-band mechanism.

Each record may consist of one or more lines, and as with MIME, described in [16], lines may be continued by putting a space or tab character on the succeeding line, allowing attributes to be of arbitrary length. Lines beginning with the "#" character are taken as comments and ignored.

Accounting records have traditionally been human-readable, so as to allow them to be more easily debugged. ADIF attributes can be expressed either in NVT ASCII (characters 32 through 126) or if non-printable characters are required, in base64.

5.1. Attribute and sub-attribute encoding

ADIF includes support for encoding of attributes for any protocol utilizing attribute/value pairs or variable bindings. The protocol type is indicated by prepending the protocol keyword as defined by IANA and a "/" to the attribute number, i.e. radius//46. To improve compactness, when a default protocol is indicated in the header, attributes of the default protocol do not need to include the protocol type. For example, if defaultProtocol: radius is indicated in the ADIF header, then 32 may be used instead of radius//32.

Protocols such as L2TP [13] include additional fields such as Vendor ID and flags in their Attribute Value Pairs (AVPs). To encode such AVPs,

ADIF includes support for attribute numbers expressed in OID form, as well as sub-attributes. Protocols such as COPS, defined in [22] include attribute numbers expressed as a combination of C-Num and C-Type as well as supporting complex objects.

Sub-attributes are included as additional fields, separated by a semi-colon, and are of the form <subattribute> = <value>. Sub-attributes used in complex objects are numbered starting from 1; letters are used for well-known sub-attributes. This document includes support for the following well-known sub-attributes: VendorId, VendorType, Mandatory and Hidden, which apply to all protocols. Other sub-attributes may be added as needed. Use of the VendorId and VendorType sub-attribute may be used for expression of vendor-specific attributes, such as those supported in RADIUS.

As an example, the L2TP version number attribute (attribute 2) with the Mandatory bit set would be expressed as "l2tp//2: 1; M=1". A COPS In-Interface object (C-Num=3, C-Type=1, including an IPv4 address of [204.57.137.2](#) and an interface index of 1) would be expressed as "cops//3.1: 1=204.57.137.2; 2=1"

[5.2.](#) OID encoding

In order to allow for compact representation of object identifiers, the ADIF header allows the definition of oid-names that can be used in place of an object-identifier sub-tree. For example, the inclusion of a header statement:

```
oid-define: mib-2=1.3.6.1.2.1;
```

would permit the string "mib-2" to substitute for the sub-tree [1.3.6.1.2.1](#) wherever this occurred within the accounting record file. This would allow the OID 1.3.6.1.2.1.2 to be abbreviated as mib-2.2.

[5.3.](#) ADIF data types

In order for to enable unambiguous encoding of protocols within ADIF, and encourage interoperability, it is necessary to define exactly how attributes and varbinds of various data types are represented within ADIF. To achieve this, the ADIF BNF specifies the encoding of the following data types:

IPv4Address - A 32-bit IPv4 address.
IPv6Address - A 128-bit IPv6 address.
Boolean - A single bit, whose value is either TRUE (1) or FALSE (0).
Char - An unsigned 8-bit number.
Integer16 - A signed 16-bit number.
Unsigned16 - An unsigned 16-bit number.

Integer32 - A signed 32-bit integer.
Unsigned32 - An unsigned 32-bit integer.
Integer64 - A signed 64-bit integer.
Unsigned64 - An unsigned 64-bit integer.
Float32 - A 32-bit IEEE floating point number.
Float64 - A 64-bit IEEE floating point number.
Float128 - A 128-bit IEEE floating point number.
OctetString - 1+ Octets containing binary data (values 0 through 255 decimal).

Encoded in ADIF as Base-64.

- Text - 1+ Octets containing UTF-8 encoded 10646 [33] characters.
Time - 32 bit unsigned value, most significant octet first, representing
 seconds since 00:00:00 UTC, January 1, 1970.

5.4. Grammar

The following definition uses the ABNF specified in [6]:

```
adif-file           = header-spec *SEP 1*( SEP adif-record )
header-spec         = required-info SEP [optional-info SEP]
required-info       = device-spec SEP start-spec SEP
optional-info       = [version-spec SEP ] [description SEP] [def-protocol SEP]
                    [oid-def-spec SEP ]
device-spec         = "device:" *SP value
description         = "description:" *SP value
oid-def-spec        = "oid-define:" *SP 1*(oid-name "=" oid-num ";")
def-protocol        = "defaultProtocol:" *SP protocol
protocol            = <protocol keyword, defined by IANA>
version-spec        = "version:" *SP number
number              = 1*Digit ; number MUST be "1" for the
                    ; ADIF format described in this document

start-spec          = "date:" *SP datetime
datetime            = date SP time
date                = Dd SP Mon SP YYYY
time                = hh ":" mm ":" ss SP zone
Dd                  = <the one or two decimal integer day of the month in
                    the range 1 to 31.>
Mon                 = "JAN" / "FEB" / "MAR" / "APR" / "MAY" / "JUN" /
                    "JUL" / "AUG" / "SEP" / "OCT" / "NOV" / "DEC"
YYYY                = <the four decimal integer year in the range 0000 to
                    9999>
hh                  = <the two decimal integer hour of the day in the
                    range 00 to 24>
mm                  = <the two decimal integer minute of the hour in the
                    range 00 to 59>
ss                  = <the two decimal integer second of the minute in the
                    range 00 to 59>
zone                = <A four digit, signed time zone offset, such as -0600 fo
```


	time zone name in parentheses, e.g., "-0800 (PDT)">
adif-record	= 1*(attrval-series SEP)
attrval-series	= [rdate-spec SEP] 1*(attrval-spec)
rdate-spec	= "rdate:" *SP datetime ; date at which accounting data was received
attrval-spec	= attr ((std-encoding / base-64-encoding) [sub-attr-encoding]
comment	= ("#" *safe)
std-encoding	= (":" *SP value)
base-64-encoding	= ("::" *SP base64-value)
base64-value	= <base-64-encoded value, defined in [11] >
sub-attr-encoding	= *(";" sub-attr "=" <value>)
sub-attr	= "M" / "H" / "VID" / "VT" / Digit ; Mandatory, Hidden, Vendor ID, Vendor Type ; well known sub-attributes or digit
attr	= [protocol "/"] attribute-number
attribute-number	= number / oid
oid	= [oid-name "."] oid-num
oid-name	= Alpha *(ldh-str)
oid-num	= *(number ".") number
value	= IPv4Address / IPv6Address / Boolean / Char / Integer16 / Unsigned16 / Unsigned24 / Integer32 / Unsigned Integer64 / Unsigned64 / Float32 / Float64 / Float128 / Time / Text ; Octet-String values encoded in Base-64
value	= 1*safe-initval *safe
safe	= <ASCII values 040 - 0176 octal (32 - 126 decimal), excluding semi-colon (";", ASCII 59 decimal)
safe-initval	= <ASCII values 040 - 0176 octal (32 - 126 decimal), excluding colon (":", ASCII 58 decimal), SP, and semi-colon (";", ASCII 59 decimal)
SP	= %x20 ; Space character
SEP	= (CR LF) / LF
CR	= <ASCII CR, carriage return>
LF	= <ASCII LF, line feed>
ldh-str	= *(Alpha / Digit / "-") let-dig
let-dig	= Alpha / Digit
Alpha	= %x41-5A / %x61-7A ; A-Z / a-z
Digit	= %x30-39 ; 0-9

6. Profiles

In order to enable ADIF to be used for communicating accounting information within a protocol, it is necessary to define an ADIF profile for that protocol. The ADIF profile specifies additional well known sub-attributes to be used with the protocol, as well as the data types to be used for each variable.

This section defines the ADIF profile for RADIUS [5]-[9].

[6.1.](#) RADIUS profile

[6.1.1.](#) Sub-attribute support

[6.1.2.](#) Data types

Below is the mapping between RADIUS data types defined in [5]-[9] and the corresponding ADIF types:

RADIUS type	ADIF type
-----	-----
Address [5]	IPv4Address
IPv6Address [49]	IPv6Address
Tag [8]	Char
Prefix [49]	Char
Integer [5]	Unsigned32
Salt [8]	Unsigned16
String [5]	Octet-String
Text [5]	Text
Time [5]	Time
Integer24 [8]	Unsigned24

[6.1.3.](#) RADIUS example

Example 1: An ADIF file encoding RADIUS accounting data

```
version: 1
device: server3
description: Accounting Server 3
date: 02 Mar 1999 12:19:01 -0500
defaultProtocol: radius

rdate: 02 Mar 1999 12:20:17 -0500
#NAS-IP-Address
4: 204.45.34.12
#NAS-Port
5: 12
#NAS-Port-Type
61: 2
#User-Name
1: fred@bigco.com
#Acct-Status-Type
```

40: 2
#Acct-Delay-Time

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

41: 14
#Acct-Input-Octets
42: 234732
#Acct-Output-Octets
43: 15439
#Acct-Session-Id
44: 185
#Acct-Authentic
45: 1
#Acct-Session-Time
46: 1238
#Acct-Input-Packets
47: 153
#Acct-Output-Packets
48: 148
#Acct-Terminate-Cause
49: 11
#Acct-Multi-Session-Id
50: 73
#Acct-Link-Count
51: 2

Example 2: An ADIF file encoding RADIUS data with a vendor-specific attribute

version: 1
device: server3
description: Accounting Server 3
date: 02 Mar 1998 12:19:01 -0500
defaultProtocol: radius

rdate: 02 Mar 1998 12:25:23 -0500
4: 204.45.34.12
5: 12
61: 2
1: fred@bigco.com
#Vendor-Specific
26: 2; VID=301; VT=22
40: 2

41: 14
42: 234732
43: 15439
44: 185
45: 1
46: 1238
47: 153
48: 148
49: 11

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

50: 73
51: 2

7. References

- [1] Aboba, B., Lu J., Alsop J., Ding J., and W. Wang, "Review of Roaming Implementations", [RFC 2194](#), September 1997.
- [2] Aboba, B., and G. Zorn, "Criteria for Evaluating Roaming Protocols", [RFC 2477](#), January 1999.
- [3] Aboba, B., Arkko, J., Harrington, D., "Introduction to Accounting Management", [RFC 2975](#), October 2000.
- [4] Brownlee, N., Blount, A., "Accounting Attributes and Record Formats", [RFC 2924](#), September 2000.
- [5] Rigney C., Rubens A., Simpson W., and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", [RFC 2865](#), June 2000.
- [6] Rigney C., "RADIUS Accounting", [RFC 2866](#), June 2000.
- [7] Zorn, G., Mitton, D. and B. Aboba, "RADIUS Accounting Modifications for Tunnel Protocol Support", [RFC 2867](#), June 2000.
- [8] Zorn, G., Leifer, D., Rubens, A., Shriver, J., Holdrege, M., Goyret, I., "RADIUS Attributes for Tunnel Protocol Support", [RFC 2868](#), June 2000.
- [9] Rigney, C., Willats, W. and Calhoun, P., "RADIUS Extensions", [RFC](#)

[2869](#), June 2000.

- [10] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [11] Crocker, D., and P. Overrell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.
- [12] Reynolds, J., Postel, J., "Assigned Numbers," [RFC 1700](#), October 1994.
- [13] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and Palter, B., "Layer Two Tunneling Protocol L2TP", [RFC 2661](#), August 1999.

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

- [14] McCloghrie, K., Heinanen, J., Greene, W., Prasad, A., "Accounting Information for ATM Networks", [RFC 2512](#), February 1999.
- [15] McCloghrie, K., Heinanen, J., Greene, W., Prasad, A., "Managed Objects for Controlling the Collection and Storage of Accounting Information for Connection-Oriented Networks", [RFC 2513](#), February 1999.
- [16] Borenstein, N., Freed, N. "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies", RFC 1521, Bellcore, Innosoft, December 1993.
- [17] Atkinson, R., Kent, S., "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [18] Mills, C., Hirsch, G. and Ruth, G., "Internet Accounting Background", [RFC 1272](#), November 1991.
- [19] Brownlee, N., Mills, C. and G. Ruth, "Traffic Flow Measurement: Architecture", [RFC 2722](#), October 1999.
- [20] Brownlee, N., "Traffic Flow Measurement: Meter MIB", Measurement: Architecture", [RFC 2720](#), October 1999.

- [21] Handelman, S., Brownlee, N., Ruth, G. and S. Stibler, "New Attributes for Traffic Flow Measurement", [RFC 2724](#), October 1999.
- [22] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", [RFC 2748](#), January 2000.
- [23] Information processing systems - Open Systems Interconnection - Specification of Abstract Syntax Notation One (ASN.1), International Organization for Standardization, International Standard 8824, December 1987.
- [24] Information processing systems - Open Systems Interconnection - Specification of Basic Encoding Rules for Abstract Notation One (ASN.1), International Organization for Standardization, International Standard 8825, December 1987.
- [25] "Data elements and interchange formats -- Information interchange -- Representation of dates and times", ISO 8601:1988.
- [26] Crocker, D., "STANDARD FOR THE FORMAT OF ARPA INTERNET TEXT MESSAGES", STD 11, [RFC 822](#), August 1982.

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

- [27] "Specification of TMN applications at the Q3 interface: Call detail recording", ITU-T Recommendation Q.825, 1998.
- [28] Handley, M., Schulzrinne, H., Schooler, E. and J. Rosenberg, "SIP: Session Initiation Protocol", [RFC 2543](#), March 1999.
- [29] "Telecommunications and Internet Protocol Harmonization Over Networks (TIPHON); Inter-domain pricing, authorization, and usage exchange", TS 101 321 V1.4.2, December 1998.
- [30] Bray, T., J. Paoli, and C. Sperberg-McQueen, "Extensible Markup Language (XML) 1.0", W3C Recommendation, February 1998.
- [31] "XML Schema Part 2: Datatypes", W3C Working Draft 07 April 2000, April 2000.
- [32] "XML Schema Part 1: Structures", W3C Working Draft 7 April 2000, April 2000.

- [33] Yergeau, F., "UTF-8, a transformation format of ISO 10646", [RFC 2279](#), January 1998.
- [34] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing SNMP Management Frameworks", [RFC 2271](#), Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998.
- [35] Rose, M., and K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based Internets", [RFC 1155](#), Performance Systems International, Hughes LAN Systems, May 1990.
- [36] Rose, M., and K. McCloghrie, "Concise MIB Definitions", [RFC 1212](#), Performance Systems International, Hughes LAN Systems, March 1991.
- [37] M. Rose, "A Convention for Defining Traps for use with the SNMP", [RFC 1215](#), Performance Systems International, March 1991.
- [38] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)", [RFC 1902](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.
- [39] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMPv2)", [RFC 1903](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.

- [40] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMPv2)", [RFC 1904](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.
- [41] Case, J., Fedor, M., Schoffstall, M., and J. Davin, "Simple Network Management Protocol", [RFC 1157](#), SNMP Research, Performance Systems International, Performance Systems International, MIT Laboratory for Computer Science, May 1990.

- [42] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Introduction to Community-based SNMPv2", [RFC 1901](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.
- [43] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)", [RFC 1906](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.
- [44] Case, J., Harrington D., Presuhn R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", [RFC 2272](#), SNMP Research, Inc., Cabletron Systems, Inc., BMC Software, Inc., IBM T. J. Watson Research, January 1998.
- [45] Blumenthal, U., and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", [RFC 2274](#), IBM T. J. Watson Research, January 1998.
- [46] Case, J., McCloghrie, K., Rose, M., and S. Waldbusser, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", [RFC 1905](#), SNMP Research, Inc., Cisco Systems, Inc., Dover Beach Consulting, Inc., International Network Services, January 1996.
- [47] Levi, D., Meyer, P., and B. Stewart, "SNMPv3 Applications", [RFC 2273](#), SNMP Research, Inc., Secure Computing Corporation, Cisco Systems, January 1998
- [48] Wijnen, B., Presuhn, R., and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", [RFC 2275](#), IBM T. J. Watson Research, BMC Software, Inc., Cisco Systems, Inc., January 1998
- [49] Aboba, B., Zorn, G., Mitton, D., "RADIUS and IPv6", Internet draft (work in progress), [draft-aboba-radius-ipv6-03.txt](#), November 2000

[8.](#) Security Considerations

Since accounting data may include sensitive information, it may be desirable for this information to be kept confidential during

transmission. Several mechanisms may be used to accomplish this, including IPSEC, described in [\[12\]](#).

[9.](#) IANA Considerations

This draft creates two new name spaces that will need to be administered by IANA, namely the ADIF protocol name and attribute number spaces. In order to avoid creating any new administrative procedures, administration of the ADIF protocol name space will piggy-back on the allocation of IP protocol and UDP/TCP port numbers. Administration of the ADIF attribute number space will piggy-back on administration of the attribute numbers or object identifiers for the protocol in question.

ADIF protocol names are required to be unique, and are created coincident with allocation of an IP protocol number or UDP/TCP port number. In applying for a protocol number or UDP/TCP port, a unique keyword is assigned to the protocol, and this keyword is used as the ADIF protocol name.

Those wishing to use an ADIF protocol name should first acquire the rights to use the corresponding protocol or port number. Using an ADIF protocol name without first obtaining rights to a protocol or port number creates the possibility of conflict and therefore is to be discouraged.

Similarly, ADIF attribute numbers are allocated coincident with IANA allocation of attribute numbers or object identifiers for a given protocol. By default, the data types used with ADIF correspond to the data types of the attributes or variable bindings defined within the protocol specification.

Assuming that no new sub-attribute types need to be defined, and there is no ambiguity in the data types to be used for representation of the attributes or variable-bindings, then no specification is required for use of ADIF with a given protocol. However, if new sub-attributes are required, new data types need to be defined, or the mapping between the protocol data types and corresponding ADIF types is unclear, then a specification is required. In addition to addressing the outstanding issues, the specification will typically include one or more examples of ADIF use with that protocol.

10. Acknowledgments

Thanks to Glen Zorn of Cisco Systems, Jari Arkko of Ericsson, David Franscone and Pat Calhoun of Sun Microsystems, Thomas Narten of IBM, and Ryan Moats of AT&T for useful discussions of this problem space.

11. Authors' Addresses

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone: 425-936-6605
EMail: bernarda@microsoft.com

Dave Lidyard
Telco Research Corporation
[616](#) Marriott Drive
Nashville, TN 37214

Phone: 615-231-6110
EMail: dave@telcores.com

12. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.
This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English. The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns. This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

INTERNET-DRAFT The Accounting Data Interchange Format 19 January 2001

[13.](#) Expiration Date

This memo is filed as <[draft-aboba-roamops-adif-00.txt](#)>, and expires September 1, 2001.

