

## **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/1id-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 (2000). 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 or object identifiers.

While in many cases Accounting Servers will produce ADIF records based on data from accounting protocols, it is also possible for devices to store data in ADIF format and transfer ADIF records to the accounting server. The latter approach has the advantage of offloading the Accounting Server from the task of transcribing interim or session records, thus improving scalability. This approach also enables the transport of data from multiple sources within the same set of records. Where a record format is employed that supports batching, the transport of multiple records within the same accounting packet is enabled.

## **2.1. History**

-06 draft: added to terminology section. Fixed ABNF issues with respect to use of the space character. Added ability to support complex objects via numbered sub-attributes. Still needed: examples for protocols other than RADIUS, including SNMP, L2TP, LDAP and COPS.

## **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 or object identifiers.

While in many cases Accounting Servers will produce ADIF records based on data from accounting protocols, it is also possible for devices to store data in ADIF format and transfer ADIF records to the accounting server. The latter approach has the advantage of offloading the Accounting Server from the task of transcribing interim or session records, thus improving scalability. This approach also enables the transport of data from multiple sources within the same set of records. Where a record format is employed that supports batching, the transport of multiple records within the same accounting packet is enabled.

### **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 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 useful.

For accounting session records, compactness of representation is a virtue. 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 an 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.

In reviewing the current state of the art in accounting data representation, it was determined that existing approaches could not satisfy the requirements for generality, extensibility, compactness and human readability. For example, protocol-specific solutions such as the SNMP-oriented record format described in [10] do not offer sufficient flexibility in representing data from other protocols. They also have the disadvantage of not being human readable. Existing call detail records based on fixed record formats, while being human readable, do not offer the required extensibility.

While XML is human readable as well as being very extensible and flexible, it is not compact. As a result, XML-based session records are likely to consume many times more storage space than a MIME-based approach. While compression techniques can be used to reduce the size of XML-based session records, the resulting compressed records are still larger than comparable compressed MIME-based records. As a result, it is to be expected that XML-based session records will suffer in terms of compactness.

As a result of these considerations, it was decided to base ADIF on MIME, described in [11]. The use of MIME has enabled 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 [5].

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

ADIF is based on MIME, described in [11] and 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 [11], 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.

ADIF includes support for encoding of attributes for any protocol utilizing attribute/value pairs or object identifiers. This includes RADIUS, defined in [3] and [4], RTFM, defined in [14] and [15], L2TP, defined in [8] as well as SNMP, TACACS+ and COPS, defined in [16]. 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: rtfm is indicated in the ADIF header, then 32 may be used instead of rtfm//32.

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.

Protocols such as L2TP, defined in [8] include additional fields such as Vendor ID and flags in their Attribute Value Pairs (AVPs). Protocols such as COPS, defined in [16] include attribute numbers expressed as a combination of C-Num and C-Type as well as supporting complex objects. To encode such AVPs, ADIF includes support for attribute numbers expressed in OID form, as well as sub-attributes.





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"

### [5.1.](#) ADIF Examples

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
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
50: 73
51: 2
```

## **5.2. 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
for
                   US Eastern Standard Time. This may be supplemented by a
                   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                    = 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. 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] Rigney C., Rubens A., Simpson W., and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", [RFC 2138](#), April 1997.
- [4] Rigney C., "RADIUS Accounting", [RFC 2139](#), April 1997.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [6] Crocker, D., and P. Overrell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.
- [7] Reynolds, J., Postel, J., "Assigned Numbers," [RFC 1700](#), October 1994.
- [8] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and Palter, B., "Layer Two Tunneling Protocol L2TP", [RFC 2661](#), August 1999.
- [9] McCloghrie, K., Heinanen, J., Greene, W., Prasad, A., "Accounting Information for ATM Networks", [RFC 2512](#), February 1999.
- [10] 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.

- [11] 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.
- [12] Atkinson, R., Kent, S., "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [13] Mills, C., Hirsch, G. and Ruth, G., "Internet Accounting Background", [RFC 1272](#), November 1991.
- [14] Brownlee, N., Mills, C., and Ruth, G., "Traffic Flow Measurement: Architecture", [RFC 2063](#), January 1997.
- [15] Brownlee, N., "Traffic Flow Measurement: Meter MIB", Measurement: Architecture", [RFC 2064](#), January 1997.
- [16] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", [RFC 2748](#), January 2000.

## **7. 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 [\[12\]](#).

## **8. 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.

## **9. 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.

## **10. Authors' Addresses**

Bernard Aboba  
Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

Phone: 425-936-6605  
EMail: [bernarda@microsoft.com](mailto:bernarda@microsoft.com)

Dave Lidyard  
Telco Research Corporation  
**616 Marriott Drive**  
Nashville, TN 37214

Phone: 615-231-6110  
EMail: [dave@telcores.com](mailto:dave@telcores.com)

## **11. Full Copyright Statement**

Copyright (C) The Internet Society (2000). 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."

## **12.    Expiration Date**

This memo is filed as <[draft-ietf-roamops-actng-07.txt](#)>, and expires November 1, 2000.

