Internet-Draft Wood, M., Erlinger, M.
Internet Engineering Task Force Internet Security Systems
Intrusion Detection Exchange Format Working Group February 20, 2001
Expires 20 August 2001

Intrusion Detection Message Exchange Requirements

<draft-ietf-idwg-requirements-05.txt>

Status of this Memo

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

Distribution of this memo is unlimited.

This Internet Draft expires August 20, 2001.

1. Abstract

The purpose of the Intrusion Detection Exchange Format Working Group (IDWG) is to define data formats and exchange procedures for sharing information of interest to intrusion detection and response systems, and to the management systems which may need to interact with them. This Internet-Draft describes the high-level requirements for such a communication mechanism, including the rationale for those requirements where clarification is needed. Scenarios are used to illustrate the requirements.

2. Conventions used in this document

This is not an IETF standards track document and thus the keywords MUST, MUST NOT, SHOULD, and MAY are NOT as in RFC 2119, but rather:

a. MUST: This word, or the terms "REQUIRED" or "SHALL", means that the described behavior or characteristic is an absolute requirement for a proposed IDWG specification.

b. MUST NOT: This phrase, or the phrase "SHALL NOT", means that the described behavior or characteristic is an absolute prohibition of a proposed IDWG specification.

c. SHOULD: This word, or the adjective "RECOMMENDED", means that there may exist valid reasons in particular circumstances for a proposed IDWG specification to ignore described behavior or characteristics.

d. MAY: This word, or the adjective "OPTIONAL", means that described behavior or characteristics are truly optional for a proposed IDWG specification. One proposed specification may choose to include the described behavior or characteristic while another proposed specification may omit the same behavior or characteristic.

3. Introduction

This document defines requirements for the Intrusion Detection Message Exchange Format (IDMEF), which is the intended work product of the Intrusion Detection Exchange Format Working Group (IDWG). IDMEF is planned to be a standard format which automated Intrusion Detection Systems (IDS) can use for reporting what they have deemed to be suspicious or of interest. This document also specifies requirements for a communication protocol for communicating IDMEF. As chartered IDWG, has the responsibility to first evaluate existing communication protocols before choosing to specify a new one. Thus the requirements in this document can be used to evaluate existing communication protocols. If IDWG determines that a new communication protocol is necessary, the requirements in this document can be used to evaluate proposed solutions.

3.1 Rationale for IDMEF

The reasons such a format should be useful are as follows:

1) A number of commercial and free Intrusion Detection Systems are available and more are becoming available all the time. Some products are aimed at detecting intrusions on the network, others are aimed at host operating systems, while still others are aimed at applications. Even within a given category, the products have very different strengths and weaknesses. Hence it is likely that users will deploy more than a single product, and users will want to observe the output of these products from one or more manager(s). A standard format for reporting will simplify this task greatly.

2) Intrusions frequently involve multiple organizations as victims, or multiple sites within the same organization. Typically, those sites will use different IDSs. It would be very helpful to correlate such distributed intrusions across multiple sites and administrative domains. Having reports from all sites in a common format would facilitate this task.

3) The existence of a common format should allow components from different IDSs to be integrated more readily. Thus, Intrusion Detection (ID) research should migrate into commercial products more easily.

4) In addition to enabling communication from an ID analyzer to an ID manager, the IDMEF notification system may also enable communication between a variety of IDS components. However, for the remainder of this document, we refer to the communication as going from an analyzer to a manager.

All of these reasons suggest that a common format for reporting anything deemed suspicious should help the IDS market to grow and innovate more successfully, and should result in IDS users obtaining better results from deployment of ID systems.

3.2 Intrusion Detection Terms

In order to make the rest of the requirements clearer, we define some terms about typical IDSs. These terms are presented in alphabetical order. The diagram at the end of this section illustrates the relationships of some of the terms defined herein.

3.2.1 Activity:

Elements of the data source or occurrences within the data source that are identified by the sensor or analyzer as being of interest to the operator. Examples of this include (but are not limited to) network session showing unexpected telnet activity, operating system log file entries showing a user attempting to access files to which he is not authorized to have access, application log files showing persistent login failures, etc.

Activity can range from extremely serious occurrences (such as an unequivocally malicious attack) to less serious occurrences (such as unusual user activity that’s worth a further look) to neutral activity (such as user login).

3.2.2 Administrator:

The human with overall responsibility for setting the security policy of the organization, and, thus, for decisions about deploying and configuring the IDS. This may or may not be the same person as the operator of the IDS. In some organizations, the administrator is associated with the network or systems administration groups. In other organizations, it's an independent position.

3.2.3 Alert:

A message from an analyzer to a manager that an event of interest has been detected.
An alert typically contains information about the unusual activity that was detected, as well as the specifics of the occurrence.

3.2.4 Analyzer:

The ID component or process that analyzes the data collected by the sensor for signs of unauthorized or undesired activity or for events that might be of interest to the security administrator. In many existing IDSs, the sensor and the analyzer are part of the same component. In this document, the term analyzer is used generically to refer to the sender of the IDMEF message.

3.2.5 Data Source:

The raw information that an intrusion detection system uses to detect unauthorized or undesired activity. Common data sources include (but are not limited to) raw network packets, operating system audit logs, application audit logs, and system-generated checksum data.

3.2.6 Event:

The occurrence in the data source that is detected by the sensor and
which may result in an IDMEF alert being transmitted. For example, 'N' failed logins in 'T' seconds might indicate a brute-force login attack.

3.2.7 IDS:

Intrusion detection system. Some combination of one or more of the following components: sensor, analyzer, manager.

3.2.8 Manager:

The ID component or process from which the operator manages the various components of the ID system. Management functions typically include (but are not limited to) sensor configuration, analyzer configuration, event notification management, data consolidation, and reporting.

3.2.9 Notification:

The method by which the IDS manager makes the operator aware of the alert occurrence
and thus the event.
In many IDSs, this is done via the display of a
colored icon on the IDS manager screen, the transmission of an e-mail or pager message, or the transmission of an SNMP trap, although other notification techniques are also used.

3.2.10 Operator:

The human that is the primary user of the IDS manager. The operator often monitors the output of the ID system and initiates or recommends further action.

3.2.11 Response:

The actions taken in response to an event. Responses may be undertaken automatically by some entity in the IDS architecture or may be initiated by a human. Sending a notification to the operator is a very common response. Other responses include (but are not limited to) logging the activity, recording the raw data (from the data source) that characterized the event, terminating a network, user, or application session, or altering network or system access controls.

3.2.12 Sensor:

The ID component that collects data from the data source. The frequency of data collection will vary across IDS offerings.
The sensor is setup to
forward events to the analyzer.

3.2.13 Signature:

A rule used by the analyzer to identify interesting activity to the security administrator. Signatures represent one of the mechanisms (though not necessarily the only mechanism) by which IDSs
detect
intrusions.

3.2.14 Security Policy:

The predefined, formally documented statement which defines what activities are allowed to take place on an organization's network or on particular hosts to support the organization's requirements. This includes, but is not limited to, which hosts are to be denied external network access.


```
 ________
|        |                --------
| Data   |________ ________|        |  ___________
| Source |    Activity     |Sensor  | |           |
|________|        |        |________| | Operator  |________
                  |            |      |___________|        |
                \ | /        Event         A               |
               ___V___         |          /|\              |
              |       |        |            \              |
              | Sensor|__      |        Notification       |
              |_______| Event  |              \          \ | /
                  A       |    V__________      \          V
                 /|\      |   |           |      \     Response
                  |       --->| Analyzer  |__     |        A
                  |           |           | Alert |       /|\
                  |           |___________|   |   |        |
                  |                 A         |   |        |
                  |                /|\      \ | / |        |
                  |_________________|      ___V___|        |
                        |                 |       |_       |
                        |                 |Manager|________|
                        |                 |_______|
                        |                     A
                   Security                  /|\
                        |  Policy_____________|
    ______________      |
   |              |     |
   | Administrator|_____|
   |______________|
```

The diagram above illustrates the terms above and their relationships.
Not every IDS will have all of these separate
components exactly as shown. Some IDSs will combine these
components into a single module; some will have multiple instances
of these modules.

3.3 Architectural Assumptions

In this document, as defined in the terms above, we assume that an analyzer determines somehow that a suspicious event has been seen by a sensor, and sends an alert to a manager. The format of that alert and the method of communicating it are what IDMEF proposes to standardize.

For the purposes of this document, we assume that the analyzer and manager are separate components, and that they are communicating pairwise across a TCP/IP network. No other form of communication between these entities is contemplated in this document, and no other use of IDMEF alerts is considered. We refer to the communication protocol that communicates IDMEF as the IDMEF Communication Protocol (IDP).

We try to make no further architectural assumptions than those just stated. For example, the following points should not matter:

* Whether the sensor and the analyzer are integrated or separate.

* Whether the analyzer and manager are isolated, or embedded in some large hierarchy or distributed mesh of components.

* Whether the manager actually notifies a human, takes action automatically, or just analyzes incoming alerts and correlates them.

* Whether a component might act as an analyzer with respect to one component, while also acting as a manager with respect to another.

3.4 Organization of this document.

Besides this requirements document, the IDWG should
produce two other documents. The first should describe a data format or language for exchanging information about suspicious events. In this,
the requirements
document, we refer to that document as the "data-format specification". The
second document
to be produced should identify existing IETF protocols that are best used for conveying the data so formatted, and explain how to package this data in those existing formats or to specify a new protocol.
We refer to this as the
IDP (IDMEF Communication Protocol).

Accordingly, the requirements here are partitioned into five sections

* The first of these contains general requirements that apply to all aspects of the IDMEF specification.

* The second section describes requirements on the formatting of IDMEF messages.

* The third section outlines requirements on the communications mechanism, IDP, used to move IDMEF messages from the analyzer to the manager.

* The fourth section contains requirements on the content and semantics of the IDMEF messages.

* The final section places requirements on IDMEF event definitions and

the event definition process.

For each requirement, we attempt to state the requirement as clearly as possible without imposing an idea of what a design solution should be. Then we give the rationale for why this requirement is important, and state whether this should be an essential feature of the specification, or is beneficial but could be lacking if it is difficult to fulfill. Finally, where it seems necessary, we give an illustrative scenario. In some cases, we include possible design solutions in the scenario. These are purely illustrative.

3.5 Document Impact on IDMEF Designs

It is expected that proposed IDMEF designs will, at a minimum, satisfy the requirements expressed in this document. However, this document will be used only as one of many criteria in the evaluation of various IDMEF designs and proposed communication protocols.
It is recognized that the working group may use additional metrics to evaluate competing IDMEF designs and/or communication protocols.

4. General Requirements

4.1 The IDMEF SHALL reference and use previously published RFCs where possible.

4.1.1 Rationale: The IETF has already completed a great deal of research and work into the areas of networks and security. In the interest of time, it is smart to use already defined and accepted standards.

4.2 The IDMEF specification MUST take into account that IDMEF should be able to operate in environments that contain IPv4 and IPv6 implementations.

4.2.1 Rationale: Since pure IPv4, hybrid IPv6/IPv4, and pure IPv6 environments are expected to exist within the time frame of IDMEF implementations, the IDMEF specification MUST support IPv6 and IPv4 environments.

5. Message Format

The IDMEF message format is intended to be independent of the IDMEF communications protocol (IDP). It should be possible to use a completely different transport mechanism without changing the IDMEF format. The goal behind this requirement is to ensure a clean separation between semantics and communication mechanisms. Obviously the IDMEF communication protocol is recommended.

5.1 IDMEF message formats SHALL support full internationalization and localization.

5.1.1 Rationale: Since network security and intrusion detection are areas that cross geographic, political, and cultural boundaries, the IDMEF messages MUST be formatted such that they can be presented to an operator in a local language and adhering to local presentation customs.

5.1.2 Scenario: An IDMEF specification might include numeric event identifiers. An IDMEF implementation might translate these numeric event identifiers into local language descriptions. In cases where the messages contain strings, the information might be represented using the ISO/IEC IS 10646-1 character set and encoded using the UTF-8 transformation format to facilitate internationalization.

5.2 The format of IDMEF messages MUST support filtering and/or aggregation of data by the manager.

5.2.1 Rationale: Since it is anticipated that some managers might want to perform filtering and/or data aggregation functions on IDMEF messages, the IDMEF messages MUST be structured to facilitate these operations.

5.2.2 Scenario: An IDMEF specification proposal might recommend fixed format messages with strong numerical semantics. This would lend itself to high-performance filtering and aggregation by the receiving station.

6. IDMEF Communications Protocol (IDP) Requirements

6.1 The IDP MUST support reliable transmission of messages.

6.1.1 Rationale: IDS managers often rely on receipt of data from IDS analyzers to do their jobs effectively. Since IDS managers will rely on IDMEF messages for this purpose, it is important that IDP deliver IDMEF messages reliably.

6.2 The IDP MUST support transmission of messages between ID components across firewall boundaries without compromising security.

6.2.1 Rationale: Since it is expected that firewalls will often be deployed between IDMEF capable analyzers and their corresponding managers,
the ability to send IDMEF messages through firewalls is necessary. Setting up this communication MUST NOT require changes to the intervening firewall(s) that weaken the security of the protected network(s). Nor SHOULD this be achieved by conflating IDMEF messages with other kinds of traffic (e.g., by overloading the HTTP POST method) since that would make it difficult for an organization to apply separate policies to IDMEF traffic and other kinds of traffic.

6.2.2 Scenario: One possible design is the use of TCP to convey IDMEF messages. The general goal in this case is to avoid opening dangerous inbound "holes" in the firewall. When the manager is inside the firewall and the analyzers are outside the firewall, this is often achieved by having the manager initiate an outbound connection to each analyzer. However, it is also possible to place the manager outside the firewall and the analyzers on the inside; this can occur when a third-party vendor (such as an ISP) is providing monitoring services to a user. In this case, the outbound connections would be initiated by each analyzer to the manager. A mechanism that permits either the manager or the analyzer to initiate connections would provide maximum flexibility in manager and analyzer deployment.

6.3 The IDP MUST support mutual authentication of the analyzer and the manager to each other.

6.3.1 Rationale: Since the alert messages are used by a manager to direct responses or further investigation related to the security of an enterprise network, it is important that the receiver have confidence in the identity of the sender and that the sender have confidence in the identity of the receiver. This is peer-to-peer authentication of each party to the other. It MUST NOT be limited to authentication of the underlying communications mechanism, for example, because of the risk that this authentication process might be subverted or misconfigured.

6.4 The IDP MUST support confidentiality of the message content during message exchange. The selected design MUST be capable of supporting a variety of encryption algorithms and MUST be adaptable to a wide variety of environments.

6.4.1 Rationale: IDMEF messages potentially contain extremely sensitive information (such as passwords) and would be of great interest to an intruder. Since it is likely some of these messages will be transmitted across uncontrolled network segments, it is important that the content be shielded. Furthermore, since the legal environment for encryption technologies is extremely varied and changes often, it is important that the design selected be capable of supporting a number of different encryption options and be adaptable by the user to a variety of environments.

6.5 The IDP MUST ensure the integrity of the message content. The selected design MUST be capable of supporting a variety of integrity mechanisms and MUST be adaptable to a wide variety of environments.

6.5.1 Rationale: IDMEF messages are used by the manager to direct action related to the security of the protected enterprise network. It is vital for the manager to be certain that the content of the message has not been changed after transmission.

6.6 The IDP SHOULD be able to ensure non-repudiation of the origin of IDMEF messages.

6.6.1 Rationale: Given that sensitive security information is being exchanged via the IDMEF, it is important that the humans operating the system are able to associate messages with the originating IDMEF entity.

6.7 The IDP SHOULD resist protocol denial of service attacks.

6.7.1 Rationale: A common way to defeat secure communications systems is through resource exhaustion. While this does not corrupt valid messages, it can prevent any communication at all. It is desirable that IDP resist such denial of service attacks.

6.7.2 Scenario: An attacker penetrates a network being defended by an IDS. Although the attacker is not certain that an IDS is present, he is certain that application-level encrypted traffic (i.e., IDMEF traffic) is being exchanged between components on the network being attacked. He decides to mask his presence and disrupt the encrypted communications by initiating one or more flood events. If the IDP can resist such an attack, the probability that the attacker will be stopped increases.

6.8 The IDP SHOULD resist malicious duplication of messages.

6.8.1 Rationale: A common way to impair the performance of secure communications mechanisms is to duplicate the messages being sent, even though the attacker might not understand them, in an attempt to confuse the receiver. It is desirable that the IDP resist such message duplication.

6.8.2 Scenario: An attacker penetrates a network being defended by an IDS. The attacker suspects that an IDS is present and quickly identifies the encrypted traffic flowing between system components as being a possible threat. Even though she cannot read this traffic, she copies the messages and directs multiple copies at the receiver in an attempt to confuse it. If the IDP resists such message duplication, the probability that the attacker will be stopped increases.

7. Message Content

7.1 There are many different types of IDSs, such as those based on: signatures, anomalies, correlation, network monitoring, host monitoring, or application monitoring. The IDMEF design MUST strive to accommodate these diverse approaches by concentrating on conveying *what* an IDS has detected, rather than *how* it detected it.

7.1.1 Rationale: There are many types of IDSs that analyze a variety of data sources. Some are profile based and operate on log files, attack signatures etc. Others are anomaly based and define normal behavior and detect deviations from the established baseline. Each of these IDSs report different data that, in part, depends on their intrusion detection methodology. All MUST be supported by this standard.

7.2 The content of IDMEF messages MUST contain the identified name of the event if it is known. This name MUST be drawn from a standardized list of events or will be an implementation-specific name if the event identity has not yet been standardized. It is not known how this list will be defined or updated, although requirements on the creation of this list are presented in the next section of this document.

7.2.1 Rationale: Given that this document presents requirements on standardizing ID message formats so that an ID manager is able to receive alerts from analyzers from multiple implementations, it is important that the manager understand the semantics of the reported events. There is, therefore, a need to identify known events and store information concerning their methods and possible fixes to these events. Some events are well known and this recognition can help the operator.

7.2.2 Scenario: Intruder launches an attack that is detected by two different analyzers from two distinct implementations. Both report the same event identity to the ID manager, even though the algorithms used to detect the attack by each analyzer might have been different.

7.3 The IDMEF message design MUST include information, which the sender should provide, that allows a receiver to locate background information on the kind of event that is being reported in the alert.

7.3.1 Rationale: This information is used by administrators to report and fix problems.

7.3.2 Scenario: Attacker performs a well-known attack. A reference to a URL to background information on the attack is included in the IDMEF message. The operator uses this information to initiate repairs on the vulnerable system.

7.4 The IDMEF message MUST be able to reference additional detailed data related to this specific underlying event. It is OPTIONAL for implementations to use this field. No requirements are placed on the format or content of this field. It is expected that this will be defined and described by the implementer.

7.4.1 Rationale: Operators might want more information on specifics of an event. This field, if filled in by the analyzer, MAY point to additional or more detailed information about the event.

7.5 The IDMEF message MUST contain the identity of the source of the event and target component identifier if it is known. In the case of a network-based event, this will be the source and destination IP address of the session used to launch the event. Note that the identity of source and target will vary for other types of events, such as those launched/detected at the operating system or application level.

7.5.1 Rationale: This will allow the operator to identify the source and target of the event.

7.6 The IDMEF message MUST support the representation of different types of device addresses.

7.6.1 Rationale: Devices are anything that might be involved in an intrusion (i.e., not limited to computers or networks) and might have addresses in various levels of the network protocol hierarchy (e.g., level 2 and level 3 addresses). Additionally, devices involved in an intrusion event might use addresses that are not IP-centric.

7.6.2 Scenario: The IDS recognizes an intrusion on a particular device and includes both the IP address and the MAC address of the device in the IDMEF message. In another situation, the IDS recognizes an intrusion on a device which has only a MAC address and includes only that address in the IDMEF message. Another situation involves analyzers in an ATM switch fabric which use E.164 address formats.

7.7 The IDMEF message MUST contain an indication of the possible impact of this event on the target. The value of this field MUST be drawn from a standardized list of values.

7.7.1 Rationale: Information concerning the possible impact of the event on the target system provides an indication of what the intruder is attempting to do and is critical data for the operator to perform damage assessment. Not all systems will be able to determine this, but it is important data to transmit for those systems that can. This requirement places no requirements on the list itself (e.g., properties of the list, maintenance, etc.), rather the requirement only specifies that the IDMEF must contain field for specifying impact.

7.8 The IDMEF message MUST provide information about the automatic actions taken by the analyzer in response to the event (if any).

7.8.1 Rationale: It is very important for the operator to know if there was an automated response and what that response was. This will help determine what further action to take, if any.

7.9 The IDMEF message MUST include information which would make it possible to later identify and locate the individual analyzer which reported the event.

7.9.1 Rationale: The identity of the detecting analyzer often proves to be a valuable piece of data to have in determining how to respond to a particular event.

7.9.2 Scenario: One interesting scenario involves the progress of an intrusion event throughout a network. If the same event is detected and reported by multiple analyzers, the identity of the analyzer (in the case of a network-based analyzer) might provide some indication of the network location of the target systems and might warrant a specific type of response. This might be implemented as an IP address.

7.10 The IDMEF message MUST be able to contain the identity of the implementer and the analyzer that detected the event.

7.10.1 Rationale: Users might run multiple IDSs to protect their enterprise. This data will help the systems administrator determine which implementer and analyzer detected the event.

7.10.2 Scenario: Analyzer X from implementer Y detects a potential intrusion. A message is sent reporting that it found a potential break-in with X and Y specified. The operator is therefore able to include the known capabilities or weaknesses of analyzer X in his decision regarding further action.

7.11 The IDMEF message MUST be able to state the degree of confidence of the report. The completion of this field by an analyzer is OPTIONAL, as this data might not be available at all analyzers.

7.11.1 Rationale: Many IDSs contain thresholds to determine whether or not to generate an alert. This might influence the degree of confidence one has in the report or perhaps would indicate the likelihood of the report being a false alarm.

7.11.2 Scenario: The alarm threshold monitor is set at a low level to indicate that an organization wants reports on any suspicious activity, regardless of the probability of a real attack. The degree of confidence measure is used to indicate if this is a low probability or high probability event.

7.12 The IDMEF message MUST be uniquely identifiable in that it can be distinguished from other IDMEF messages.

7.12.1 Rationale: An IDMEF message might be sent by multiple geographically-distributed analyzers at different times. A unique identifier will allow an IDMEF message to be identified efficiently for data reduction and correlation purposes.

7.12.2 Scenario: The unique identifier might consist of a unique originator identifier (e.g. IPv4 or IPv6 address) concatenated with a unique sequence number generated by the originator. In a typical IDS deployment, a low-level event analyzer will log the raw sensor information into, e.g., a database while analyzing and reporting results to higher levels. In this case, the unique raw message identifier can be included in the result message as supporting evidence. Higher level analyzers can later use this identifier to retrieve the raw message from the database if necessary.

7.13 The IDMEF MUST support reporting alert creation date and time in each event. The IDMEF MAY support reporting the date and time the event began in addition to the date and time the alert was created.

7.13.1 Rationale: Time is important from both a reporting and correlation point of view.
Event onset time might differ from the alert creation time because it might take some time for the sensor to accumulate information about a monitored activity before generating the event, and additional time for the analyzer to receive the event and create an alert. The event onset time is therefore more representative of the actual time that the reported activity began than is the alert creation time.

7.13.2 Scenario: If an event is reported in the quiet hours of the night, the operator might assign a higher priority to it than she would to the same event reported in the busy hours of the day. Furthermore, an event (like a lengthy port scan) may take place over a long period of time and it would be useful for the analyzer to report the time of the alert as well as the time the event began.

7.14 Time SHALL be reported such that events from multiple analyzers in different time zones can be received by the same manager and that the local time at the analyzer can be inferred.

7.14.1 Rationale: For event correlation purposes, it is important that the manager be able to normalize the time information reported in the IDMEF alerts.

7.14.2 Scenario: A distributed ID system has analyzers located in multiple timezones, all reporting to a single manager. An intrusion occurs that spans multiple timezones as well as multiple analyzers. The central manager requires sufficient information to normalize these alerts and determine that all were reported near the same "time" and that they are part of the same attack.

7.15 The format for reporting the date MUST be compliant with all current standards for Year 2000 rollover, and it MUST have sufficient capability to continue reporting date values past the year 2038.

7.15.1 Rationale: It is desirable that the IDMEF have a long lifetime and that implementations be suitable for use in a variety of environments. Therefore, characteristics that limit the lifespan of the IDMEF (such as 2038 date representation limitation) MUST be avoided.

7.16 Time granularity and time accuracy in event messages SHALL NOT be specified by the IDMEF.

7.16.1 Rationale: The IDMEF cannot assume a certain clock granularity on sensing elements, and so cannot impose any requirements on the granularity of the event timestamps. Nor can the IDMEF assume that the clocks being used to timestamp the events have a specified accuracy.

7.17 The IDMEF message MUST support an extension mechanism used by implementers to define implementation-specific data. The use of this mechanism by the implementer is OPTIONAL. This data contains implementation-specific information determined by each implementer. The implementer MUST indicate how to interpret these extensions, although there are no specific requirements place on how implementers describe their implementation-specific extensions.

7.17.1 Rationale: Implementers might wish to supply extra data such as the version number of their product or other data that they believe provides value added due to the specific nature of their product. Implementers may publish a document or web site describing their extensions; they might also use an in-band extension mechanism that is self-describing.

7.18 The semantics of the IDMEF message MUST be well defined.

7.18.1 Rationale: Good semantics are key to understanding what the message is trying to convey so there are no errors. Operators will decide what action to take based on these messages, so it is important that they can interpret them correctly.

7.18.2 Scenario: Without this requirement, the operator receives an IDMEF message and interprets it one way. The implementer who constructed the message intended it to have a different meaning from the operator's interpretation. The resulting corrective action is, therefore, incorrect.

8. Alert Identifiers and the Alert Identifier Definition Process

8.1 The standard list of IDMEF alerts MUST be extensible. As new events are defined by the community and as new methods of detecting them are available, the IDMEF MUST be able to grow with the technology.

8.1.1 Rationale: New intrusions are rapidly created; some are variations of existing intrusions and some are newly created intrusion techniques. If IDMEF is not extensible then the usefulness of the standard will quickly diminish.

8.2 The IDMEF itself MUST be extensible. As new ID technologies emerge and as new information about events becomes available, the IDMEF message format MUST be able to include this new information.

8.2.1 Rationale: As intrusion detection technology continues to evolve, it is likely that additional information relating to detected events will become available. The IDMEF message format MUST be able to be extended by a specific implementation to encompass this new information.

8.3 The standard list of alert identifiers MUST be extensible by implementers and administrators.

8.3.1 Rationale: The IDMEF will specify the basic information for each intrusion. Additionally, specific implementations might want to use the IDMEF for non-standard events.

8.4 The process by which new alert identifiers are defined and standardized MUST be implementation-independent.

8.4.1 Rationale: The process for new alert identifier definition MUST NOT favor one IDS implementation over another, otherwise a specific IDS implementation might determine that making event information available to the community has a negative effect on that implementation and might elect not to do so.

Acknowledgements:

The following individuals contributed substantially to this document and should be recognized for their efforts. This document would not exist without their help:

Mark Crosbie, Hewlett-Packard
David Curry, IBM Emergency Response Services
David Donahoo, Air Force Information Warfare Center
Mike Erlinger, Harvey Mudd College
Fengmin Gong, Microcomputing Center of North Carolina
Dipankar Gupta, Hewlett-Packard
Glenn Mansfield, Cyber Solutions, Inc.
Jed Pickel, CERT Coordination Center
Stuart Staniford-Chen, Silicon Defense
Maureen Stillman, Nokia IP Telephony

Editor's Address:

Mark Wood
Internet Security Systems, Inc.
6303 Barfield Road
Atlanta, GA 30328
Phone: +1 (404) 236-2600
E-mail: mark1@iss.net

Mike Erlinger
The Aerospace Corporation
PO Box 92927 M1/102
Los Angeles, CA, 90009
Phone: +1 (909) 621-8912
E-mail: mike@cs.hmc.edu

Intrusion Detection Exchange Format Working Group (IDWG):

The Intrusion Detection Exchange Format Working Group (IDWG) can be contacted via the working group's mailing list (idwg-public@zurich.ibm.com) or through its chairs:

Stuart Staniford
stuart@SiliconDefense.com
Silicon Defense

Mike Erlinger
mike@cs.hmc.edu
Harvey Mudd College

Full Copyright Statement