INCH Working Group
Internet-Draft
Category: Informational
Expires: December 24, 2006

Glenn M Keeni
Cyber Solutions Inc.
Roman Danyliw
CERT/CC
Yuri Demchenko
University of Amsterdam

June 25, 2006

Requirements for the Format for Incident Information Exchange (FINE)
<draft-ietf-inch-requirements-08.txt>

Status of this Memo

Abstract

This document describes the high-level functional requirements of an abstract format, the Format for Incident information Exchange (FINE), which will facilitate the exchange of incident information among computer security incident response teams (CSIRTs) and involved parties. A common and well-defined format will help in the exchange of incident related information across different administrative domains such as organizations, regions, and countries. Implementations of FINE will also be useful for reactionary analysis of current threats and support the proactive identification of trends that can lead to incident prevention.

Table of Contents

1. Introduction

Computer security incidents occur across administrative domains, often spanning different organizations and national borders. Hence, a response requires coordination and collaboration between the involved parties and the responsible computer security incident response teams (CSIRTs). The basis for this interaction is often data and statistics describing the nature of the incident. This information, referred to as an incident report in this document, will not only support response activity to the specific incident, but may also be used for historical analysis or proactive responses.

This document defines the high-level functional requirements for a format that can support the exchange of incident reports. The abstract format being discussed is referred to as the Format for INcident information Exchange (FINE). The implementation of the requirements, the format itself, is out of the scope of this document.

The intent of FINE is to enable rapid and effective response to incidents by improving the ability of CSIRTs to exchange and process incident reports. This will be achieved by ensuring that implementations of FINE require:

+ unambiguous semantics for the data;
+ a well-defined syntax for the data; and
+ support end-user processing (e.g., categorization and statistical analysis).

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [1].

2. Incident Handling Framework

2.1. Descriptive Terms

For the purpose of clarity, certain commonly used terms from the operational domain of CSIRTs are defined here. These are based on related documents [3, 4, 5, 6, 7]

2.1.1. Event

An event is an occurrence in a system or network that may be of interest and warrant attention. An event is not necessarily malicious or deliberate.

2.1.2. Attack

An attack is a series of events caused either directly or indirectly by a source that violates the security policy of the target. These violations may include a compromise of a user account, denial of service, information theft, etc.

2.1.3. Source

The origin of an attack as described by a host, user account, computer program, network address, person, or organization.

2.1.4. Target

The target of an attack as described by a host, user account, computer program, network address, person, or organization.

2.1.5. Computer security incident

A computer security incident, referred to as an incident, is a set of one or more related attacks.

2.1.6. Incident report

An incident report is the collection of information describing an incident. In this document the terms "incident report" and "incident information" are used interchangeably.

2.1.7. CSIRT

A computer security incident response team, CSIRT, is an individual or a group of individuals that has the responsibility to coordinate and support the response to incidents in a defined constituency [6]. A CSIRT creates, receives, processes, and maintains incident reports.

2.1.8. Impact

An impact describes the consequence of an incident on a target expressed in terms relevant to a user community.

2.2 The Operational Model

Incident reports are an important subset of information exchanged between a CSIRT and its constituency or other CSIRTs. These reports form the basis for resolving and understanding activity in a constituency. A CSIRT may create an incident report when an incident is reported, receive a report from another CSIRT, or send a report to a CSIRT. As investigation into the incident progresses, new information about an incident may be discovered. New information may trigger subsequent information exchange.

The creation and exchange of incident reports is often driven by a work-flow process that prioritizes and manages the information flow in a CSIRT. These systems often associate CSIRT personnel with

particular incidents or maintain status onto a given investigation. FINE does not provide a representation for these internal processes.

FINE is a representation for the data exchanged between different parties. In order to integrate FINE into the operational processes of CSIRTS, the parties will have to use an interface to convert to and from the internal data representation (of a propriety work-flow application or database) and FINE. Hence, the sender of an incident report must convert from the local format to FINE, while the recipient must translate FINE back into its own local format. The communicating CSIRTs need not have the same local format for storing incident reports. This information exchange is depicted in Figure 1.

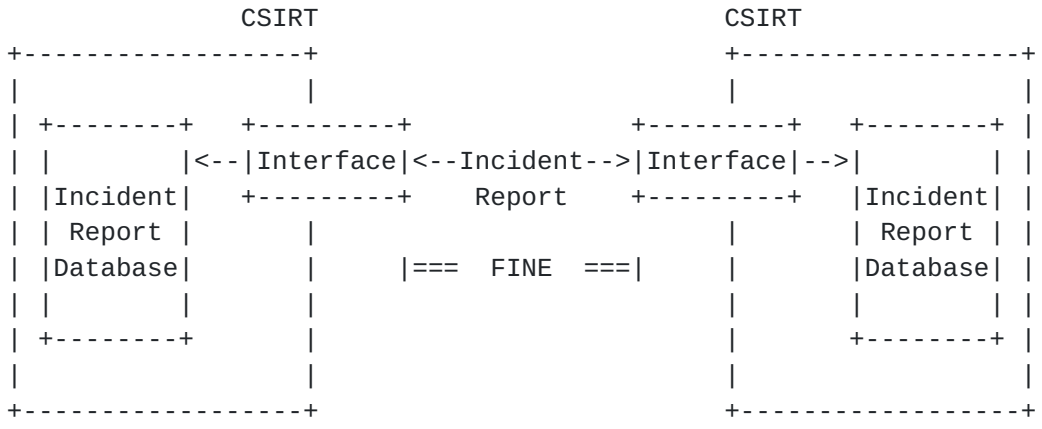

Fig. 1 Operational Model for FINE

3. General Requirements

3.1 FINE SHALL reference and use previously published RFCs where possible.

3.2 FINE MUST have well-defined semantics and specify a standard mechanism for extensibility.

The data elements of the various components of FINE should be typed, and the meaning should be well specified. Likewise, there should be a standardized method to address representing data not defined in the data model.

4. Format Requirements

4.1 FINE SHALL support full internationalization and localization.

A significant part of the incident report may consist of natural language text. Since some incidents may involve CSIRTs from different countries, FINE must have provisions for using local character sets and encodings.

In cases where local (non-standard) character sets and encodings are used, the data elements that carry encoding-sensitive information should be clearly indicated.

4.2 FINE MUST allow multilingual reports.

Different parts of the incident report may be written in a different natural language. FINE must support multiple translations of the same data element.

4.3 FINE MUST support aggregation and filtering of incident report data.

The structure of the FINE data elements and their associated semantics must lend themselves to aggregation and filtering by applications.

4.4 FINE MUST be able to document the evolution of an incident report.

As incidents are investigated new information may become available or old information may be invalidated. FINE must support the ability to convey this track record of an incident report.

4.5 FINE MUST support a granular access restriction policy on subsets of the incident report.

Different parts of an incident report may have information of varying degrees of sensitivity. It must be possible to label subsets of the incident report with their appropriate sensitivity. With this information, applications can then implement different levels of access restrictions for the different components of the incident report.

4.6 FINE SHOULD allow the application of external mechanisms to support authenticity, integrity, and non-repudiation checks of incident reports.

FINE itself need not guarantee authenticity, integrity, or non-

repudiation. However, the specification must detail a standardized mechanism to ensure these properties.

5. Communication Mechanism Requirements

5.1 The security properties of FINE reports SHOULD be independent of the communication mechanism.

The exchange of incident reports is typically conducted using standard communication protocols (e.g., SMTP, HTTP, FTP, XML Web Services). The security properties of FINE MUST NOT be tied to a particular communications protocol. Provisions for authenticity, integrity, and confidentiality should be made in FINE.

6. Content Requirements

6.1 FINE MUST be flexible enough to support various degrees of completeness, while still clearly defining the minimal information required for describing an incident.

6.2 FINE MUST support globally unique identifiers for each incident report.

It should be possible to reference an incident report unambiguously using a globally unique identifier. Furthermore, it should be possible to derive the constituency of the incident report from this identifier.

6.3 FINE MUST support the naming of the source and target.

6.4 FINE MUST support the description of various aspects of the source and target.

6.5 FINE MUST support the description of the methodology used by the attacker.

Well-known classifications or enumeration schemes should be used to describe the attack.

6.6 FINE SHOULD support the identification of the sender of the incident report.

FINE should indicate the source of each component of the incident report if it is different from the sender (e.g., the team handling the incident).

6.7 FINE SHOULD support the inclusion or referencing of information external to the incident report.

6.8 FINE MUST support natural language descriptions of the incident.

6.9 FINE SHOULD support references to the appropriate security advisories from coordination and analysis centers.

6.10 FINE SHOULD support a description of the impact of the incident.

6.11 FINE SHOULD support a description of the actions taken during the course of handling the incident.

6.12 FINE MUST use a standardized time specification.

Incident reports should represent time in such a way that it is possible to easily compare information reported from different time zones.

7. Security Considerations

There are no explicit security considerations for this document, since no protocol or information model is specified. However, a number of security relevant requirements are outlined for FINE implementers. By its nature, FINE will represent sensitive information. Hence, implementers should ensure support for access restriction (requirement 4.5), confidentiality, integrity, and non-repudiation (requirement 4.6) all through transport independent approaches (requirement 5.1).

8. IANA Considerations

This document requires no action from IANA.

9. References

9.1 Normative References

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

9.2 Informative References

[2] Arvidsson, J., Cormack, A., Demchenko, Y. and Meijer J., "TERENA's Incident Object Description and Exchange Format Requirements." RFC 3067. February 2001.

[3] Brownlee, N., Guttman, E., "Expectations for Computer Security Incident Response." BCP 21, RFC 2350. June 1998.

[4] Shirey, R., "Internet Security Glossary." FYI 36, RFC 2828. May 2000.

[5] "Establishing a Computer Security Incident Response Capability (CSIRC)." NIST Special Publication. 800-3. November 1991.

[6] West-Brown, M., Stikvoort, D., Kossakowski, K., Killcrece G., Ruefle, R., Zajicek, M., "Handbook for Computer Security Incident Response Teams (CSIRTs)." CMU/SEI-98-HB-002. Carnegie Mellon University, Pittsburgh, PA. April 2003.

[7] Howard, J. and Longstaff, A., "A Common Language for Computer Security Incidents." Sandia Report: SAND98-8667. Sandia National Laboratories. Albuquerque, NM. October 1998.

10. Acknowledgments

The precursor of this document is "RFC3067 TERENA's Incident Object Description Exchange Format Requirements" [2], which is based on the work done in the Incident Object Description Exchange Format Working Group at TERENA. Subsequent work and discussion have been carried out in the INCH-WG and in the WIDE-WG on Network Management and Security.

The following individuals, in alphabetic order, have made a substantial contribution to this document:

Hiroyuki Kido
Hiroyuki Ohno
Kathleen M. Moriarty
Jan Meijer

11. Authors' Addresses:

Glenn Mansfield Keeni
Cyber Solutions Inc.
Sendai, Japan
Email: glenn@cysols.com

Roman Danyliw
CERT Coordination Center
4500 Fifth Ave.
Pittsburgh, PA 15213
USA
Email: rdd@cert.org

Yuri Demchenko
University of Amsterdam, The Netherlands
Email: demch@chello.nl

Full Copyright Statement

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

Appendix - non-normative.

Major Changes (reverse count)
Information about changes to the document since publishing -00 version will be documented here.

Major changes in version-08
1) Editorial changes

Major changes in version-07
1) References [4], [5] (in -06) have been removed.
2) Editorial nits have been fixed
3) Authors' list and contributors' list have been updated.

Major changes in version-06
1) Reference [3] is deleted. The reference indices are renmbered.
2) Changed the wording in the abstract to bring it in line with the title
"INcident report" => "INcident information"
3) Added a sentence to the definition of Incident report
In this document the terms "incident report" and "incident information" are used interchangeably.
4) Modified 4.1 (clause about preserving the contents of encoding sensitive information when transferring is deleted).
5) Modified 4.11 (clause for supporting different time granularities is deleted).
6) Revised the requirement 5.1
7) Editorial nits

Major changes in version-05
1) In 2.1 the definitions have been rearranged. Incident Report (earlier 2.1.8 have been moved to 2.1.6)
2) Section 2.2, Operational model, revised
3) Editorial nits
4) IDnits
5) Added Roman Danyliw to the authors list.

Major changes in version -04
1) Operational model rewritten
2) Editorial nits
3) IPR notice updated

Major changes in version -03 (Second revision)
1) title changed to
Requirements for the Format for INcident information Exchange (FINE)
2) editorial nits
3) RFC2119 key words used

4) added description to 4.6
5) reformatted 4.7 and 5.1 to have single statement requirements followed by description of the requirements.
6) added an example to 4.2
7) moved 6.13 to Format requirements as 4.8
8) updated references #3, #5, #10
9) updated section 2.2

Major changes in version -03 (First revision)
1) editorial nits
2) in Security Considerations section an example is added to explain the impact of the contents of the IR on the security and privacy of individuals of organization.
3) Section 3 is deleted

Major changes in version -02

1) clarified definitions of some terms. Added a few definitions.

2) in 5.1, added requirement for handling non-standard/local encoding and/or character codes.

3) in 5.7, added requirement that multiple versions of the report should be consistent

4) in 7.5, added requirement that the source of each component of the Incident report must be identified (if different from the creator of the Incident report).

5) some editorial nits are fixed.

Major changes in version -01

1) clarified definition of some terms - still in the process, needs more discussion with concerned parties.

2) re-written section 2. Operational model

3) added text about multilingual support for non-utf-8 character sets

to item "5.1 FINE shall support full internationalization and localization" - results of discussion at IETF-56

4) included clear statement about unique identification of the Incident report to item "5.1 FINE shall support full internationalization and localization."

5) added item about the possibility of Incident description in

natural language:

7.7 The FINE may contain a description of the Incident or comprising security events in a natural language.

6) requirement about describing impact of the Incident extended (item 7.9) with recommendation to provide guidelines to describe the impact on the target to ensure a uniform interpretation of the description.

7) item 7.11 about time normalization extended with the possibility to describe time offset when normalization is not possible.