

Title: IDWG: Progress towards an open IDS alert standard. Author: Stuart Staniford-Chen

Affiliation: Silicon Defense

Address: 513 2nd St, Eureka, CA 95501, USA. Topic Category: IDS interoperability Standards and Standardization

Abstract:

In this talk, I will give an overview of the work of the IETF's Intrusion Detection Working Group (IDWG). I'll cover the history of the group, it's progress to date, and sketch the technical aspects of the emerging standard.

IDWG began because of the attempt by the US Defense Advanced Research Project Agency (DARPA) to develop a common way for the intrusion detection research projects it was sponsoring to share data. This became known as the Common Intrusion Detection Framework (CIDF). The CIDF working group developed an S-expression based language for sharing data about intrusions, and a messaging layer to share them. As these developments started to mature, the working group approached the IETF with a view to possibly standardizing them. The IETF concluded that there was extensive interest within that organization for developing an intrusion detection standard, but that it was unclear that the CIDF work met the right requirements.

Therefore, a new IETF working group, IDWG, was created and tasked with initially studying the requirements of the problem, and then developing an appropriate data format and method for communicating the data between intrusion detection systems. The working group met for the first time in December of 1998. This working group has ended up not using the CIDF formats, but developing new technology of its own.

A great deal of thought went into the IDWG requirements document. The requirements fall into a number of areas, and a few key ones are highlighted here:

- requirements that govern the interaction of the standard with firewalls. These consider a number of deployment cases (sensor in the DMZ and console inside the organization, sensor inside the organization and console at a remote service provider, etc).
- requirements for privacy, authenticity, etc of alert transmission.
- requirements on the content of the alerts. These specify a variety of things that it must be possible to infer from the alert (such as it's time, the target of the event, identify of the sensor that detected the event, etc).

Next, the group developed a document which describes the data to be carried by an alert in an abstract way. This was done initially by studying the data carried by alerts in a variety of proprietary products today, abstracting the results, and then refined by a great deal of discussion. The outlines of this data model are now clear, although details continue to be refined and discussed.

A lot of discussion occurred as to what way to represent the data in alerts. It was finally settled to use an XML based representation for it's flexibility, ease of use with the web,

and property of being simultaneously machine and human readable. In the talk I will show several examples of alerts in the current representation.

Also, the group has developed a transport protocol, Intrusion Alert Protocol (IAP) for moving alerts around. This meets the various requirements defined in the requirements document, and is somewhat similar to HTTP, but with considerably more focus on what is required for security in both protocol and implementations. I will give an example of the operation of the protocol.