OpenPacket.org
Draft Proposal
21 July 2006

0. Author: Richard Bejtlich
TaoSecurity LLC
9532 Liberia Ave Suite 141
Manassas, VA 20110
taosecurity [at] gmail dot com
www.taosecurity.com / taosecurity.blogspot.com / www.bejtlich.net

1. Purpose: The purpose of this document is to discuss the construction and operation of OpenPacket.org.

2. Description: OpenPacket.org is a Web site whose mission is to provide a centralized repository of network traffic traces for researchers, analysts, and other members of the digital security community.

3. Rationale: Knowledge of normal, suspicious, and malicious network traffic is a fundamental component of defending the modern enterprise. At present there is no single place where one might submit or download network traffic for research, operations, or educational purposes. The following sites represent the more well-known existing repositories of network traces.

- Wireshark Sample Captures (wiki.wireshark.org/SampleCaptures)
- Reliable Software Group (www.cs.ucsb.edu/~rsg/datasets/)
- Shmoo Group Def Con 8, 10 CTF (cctf.shmoo.com/data/)
- Planet Mirror Def Con 9 CTF (public.planetmirror.com/pub/cctf/defcon9/?fl=)
- DARPA (www.ll.mit.edu/IST/ideval/data/data_index.html)

Some sites provide ASCII representations of packet contents.

- Museum of Broken Packets (lcamtuf.coredump.cx/mobp/)
- SANS Internet Storm Center (isc.sans.org)

OpenPacket.org will provide a single, reliable, moderated, authoritative collection of network traces.

4. Required Features: OpenPacket.org will provide the following features.

- OpenPacket.org visitors will be able to browse and search a collection of network traffic traces. Traces will be free for download. No registration will be required to access traces.
- Traces will be stored in a wire capture format, such as that offered by Libpcap. This allows traces to be read by community tools like Wireshark or Snort, or replayed by Tcpreplay.
- OpenPacket.org users may submit traces for inclusion in the repository. Submitting traces requires registering with OpenPacket.org.
- An OpenPacket.org moderation team will inspect candidate traces. Only moderators can approve posting traces.
- Once posted, OpenPacket.org visitors can vote on the trace. Top trace submitters will be publicly recognized (e.g., see the Top Taggers at Splunk Base [www.splunk.com/base]).
- OpenPacket.org will host a forum for users to discuss traces and packet analysis. OpenPacket.org may host a mailing list for users who prefer discussions via email. The Web site should publish word of new traces via RSS and/or Atom feeds.
- OpenPacket.org will publish news of new traces via RSS and/or Atom feeds.
- OpenPacket.org may create an Internet Relay Chat channel on the Freenode network to host real-time trace discussions.
- OpenPacket.org will be built using open source software. The solution must run on FreeBSD. Heavy preference will be given to software in the FreeBSD ports tree.

5. Trace Restrictions: Traces published at OpenPacket.org will meet all of the following guidelines.

- Traffic submitted to OpenPacket.org is provided with the express consent of the enterprise from which the traffic was recorded. In many cases the “safest” traffic is that captured in a controlled laboratory setting.
- If necessary, traffic captured on production networks will be scrubbed to obscure any identifying characteristics, such as source and/or destination IP addresses. OpenPacket.org reserves the right to make these scrubbing decisions and actions.
- The traffic does not contain any proprietary or sensitive information that the submitting enterprise would not want published. Examples include (but are not limited to) usernames and/or passwords providing access to production systems or email accounts, Social Security Numbers, credit card numbers, medical records, sales reports, and commercial software binaries.
- The traffic does not contain lewd or inappropriate content that would offend reasonable parties who may reassemble or review the traffic.
- The traffic will be the reasonable minimum required to demonstrate the characteristics that make it interesting or valuable. For example, there is no need to publish a File Transfer Protocol (FTP) data channel that contains a ten megabyte file for the purpose of demonstrating the FTP protocol.

OpenPacket.org will consider large traces only if they add substantial value to the OpenPacket.org Database. For example, a large trace designed to provide "background traffic" for testing intrusion detection systems may be submitted for the "Testing" category.

6. Operation: Richard Bejtlich will create and lead the OpenPacket.org Research Team (OPRT). The purpose of the OPRT is to facilitate the provision of quality traces to the security community while operating within the trace restrictions. Mr. Bejtlich will personally assemble the OPRT by choosing analysts he trusts to make proper decisions concerning the traces to be published at OpenPacket.org. Mr. Bejtlich or a party he designates will have final approval over all traces to be published at OpenPacket.org.

The following describes the process by which OpenPacket.org will accept new traces for publication.

1. A new user registers with OpenPacket.org. The identity of the user will be verified by sending a confirmation email to the address provided by the new user. Registration is free. Unregistered users may download, but not upload, traces.
2. When the new user decides to contribute a trace, she is presented with an acceptance challenge. The challenge confirms that the traffic to be submitted meets the restrictions outlined earlier. By clicking "yes" the user agrees that she is authorized to submit a trace. The Web site records challenge acceptance.
3. The new user uploads the trace through a Web-based form into a non-public holding area accessible only to the OPRT.
4. One or more members of the OPRT review the newly submitted trace for conformance with traffic restriction guidelines.
5. If a simple majority of OPRT members who have reviewed the trace approve of the new trace, they give it an initial categorization and send it to Mr. Bejtlich or his designated deputy for review.
6. If Mr. Bejtlich or his designated deputy reviews and approves the trace, he categorizes and publishes it at OpenPacket.org.
7. The user who submitted the new trace is automatically given some sort of reward for providing the trace. A "karma" system as used at Slashdot.org, or a similar points system, is appropriate.
8. Other users who download and review the new trace will be allowed to vote on the quality or value of the trace. Good traces will be "moderated up," while poor traces will be moderated down.
9. Users who consistently provide quality traces are recognized as top contributors and displayed as such at OpenPacket.org.

7. Organization: We envision OpenPacket.org to be a Web site organized for the rapid retrieval of relevant network traces. Traces will be categorized according to the following methods. (Other methods of organization may be formulated in the future.)

The first method of organization involves classification by category. The categories are:

- *Normal*: traffic the OPRT labels as completely benign, as might be found in normal operation on an enterprise network. Examples include Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Secure Shell (SSH).
- *Suspicious*: traffic that may be unauthorized on an enterprise network, but which is most likely not associated with intrusive activity. Examples include Gnutella, Internet Control Message Protocol (ICMP) with excessively large payloads, and packets with odd Transmission Control Protocol (TCP) flags set.
- *Malicious*: traffic caused by an intruder seeking to perform reconnaissance, exploitation, and other intrusive activities. Examples include reconnaissance activity, covert channels, or traffic that exploits a vulnerable service.
- *Unknown*: traffic that has not yet been identified as one of the other three categories. This will not be a permanent category. Researchers looking for a challenge will spend time examining this sort of traffic.

These categories will not be used to make value judgments on the merits of various protocols. For example, one party may consider BitTorrent to be a suspicious protocol as it is sometimes used to distribute intellectual property in an unauthorized manner. Another party might categorize BitTorrent as completely normal, as that protocol is the means by which his organization distributes CD-ROM .iso images to customers.

The purpose of the categories is to give less sophisticated analysts an easy means to locate traffic of interest. DNS, for example, might have traces in all three categories. A normal trace might show benign requests and replies. A suspicious trace might contain an abnormally large DNS reply, or perhaps a request or reply with odd parameters. A malicious DNS trace might present an exploit against a vulnerable DNS resolver.

The second method of organization involves classification within the Open Systems Interconnection (OSI) model. These layers are generally recognized as the following:

1. Physical
2. Data link
3. Network
4. Transport
5. Session
6. Presentation
7. Application

For example, researchers may wish to compare examples of various transport layer protocols, like TCP or Stream Control Transmission Protocol (SCTP), or network layer protocols, like Internet Protocol version 4 (IP) and version 6 (IPv6).

Again, these layers are often open for debate. For example, ICMP is sometimes classified as a network layer protocol, as it assists IP. Others argue ICMP is a transport layer protocol, because ICMP is assigned the IP protocol value 1. When such decisions need to be made, the protocol header itself will be the final judge. For example:

- Network-layer protocols are assigned EtherTypes, like 0x0806 for ARP, 0x0800 for IP version 4, and 0x86DD for IP version 6.
- Transport-layer protocols are assigned IP protocol values, like 1 for ICMP, 6 for TCP, 17 for User Datagram Protocol (UDP), 132 for SCTP, and so on.
- Application-layer protocols are assigned one or more SCTP, TCP, or UDP port numbers, like 22 for SSH, 23 for Telnet, and so on.

The IP Protocol Suite published at www.networksorcery.com/enp/topic/ipsuite.htm is helpful when categorizing protocols and will be referenced when disputes arise.

OpenPacket.org will consider accepting traces that contain a mix of traffic for the purposes of testing intrusion detection systems and other security equipment. In such cases, a new category, “Testing,” may be developed.

Traffic that is identified as being malicious will be identified by an element listing the Common Vulnerabilities and Exposures (CVE, cve.mitre.org) number, where possible.

8. Resources: The following are required to establish and operate OpenPacket.org.

- *Bandwidth*: OpenPacket.org will be sure to attract visitors as the quality and quantity of its database increases. We are looking for sponsors to provide the necessary bandwidth in exchange for public recognition of its sponsorship of the project.
- *Hardware*: OpenPacket.org will require a Web server and database server. It is possible for both components to operate on the same physical system, at least during the proof-of-concept and initial operation of the site. As the database accepts additional traces, a separate system will be needed. We are looking for sponsors to donate or pay for the necessary hardware. Alternatively, OpenPacket.org may turn to the community and request donations to pay for the necessary hardware.
- *Software*: OpenPacket.org requires a Web server such as Apache to present information to the visitors. This Web server should probably operate a Content Management System (CMS) to organize and make possible the trace acceptance, review, and publication system. The database can be an open source solution compatible with the CMS, such as MySQL or PostgreSQL.
- *Domain name*: Mr. Bejtlich has registered the OpenPacket.org domain name.

9. Caveats: The following are important aspects of OpenPacket.org that will be recognized by all parties.

- OpenPacket.org is not a repository for exploit code. It may be attractive to pair source code for an exploit with a trace that code generates. However, OpenPacket.org is not designed to provide exploits to visitors. The purpose is to provide quality traces for study and learning.

- OpenPacket.org, through the OPRT, will own the network traffic it accepts and publishes. Richard Bejtlich designed and owns the OpenPacket.org logo.
- When the OPRT decides a trace does not belong in the OpenPacket.org database, the OPRT will retain the right to immediately and reject and/or remove it.
- OpenPacket.org is vendor neutral and open to all visitors. Registration is not required to access traces. Submitting traces requires registration, so that the submitter accepts responsibility for the packets contained therein.
- OpenPacket.org is not a commercial venture. User registration will not be used as a way to market products or services to OpenPacket.org members. Discrete advertisements may be offered as a way to defray bandwidth costs, but these will not detract from the overall user experience.

10. Open Points: We are not sure if users should be allowed an easy means to download the entire contents, or large portions, of the OpenPacket.org database. The likely strain on the system makes this an undesirable feature. We also welcome ideas concerning mirrors, if thought necessary. Distribution of the entire repository via BitTorrent is another option.

11. News and Comments: News on OpenPacket.org will be posted at the OpenPacket.org blog (openpacket.blogspot.com). (When OpenPacket.org goes live, the openpacket.org URL will no longer redirect to openpacket.blogspot.com.)

The appropriate place to provide feedback on this document and the project in general is the OpenPacket.org development newsgroup (lists.sourceforge.net/lists/listinfo/openpacket-devel).

The list requires registration. Posts should be directed to openpacket-devel [at] lists dot sourceforge dot net.

Thank you for your interest in OpenPacket.org.